



The Strength of Friendship Ties in Proximity Sensor Data

Vedran Sekara^{1*}, Sune Lehmann^{1,2}

1 Cognitive Systems, Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kgs. Lyngby, Denmark, **2** Niels Bohr Institute, University of Copenhagen, Østerbro, Denmark

Abstract

Understanding how people interact and socialize is important in many contexts from disease control to urban planning. Datasets that capture this specific aspect of human life have increased in size and availability over the last few years. We have yet to understand, however, to what extent such electronic datasets may serve as a valid proxy for real life social interactions. For an observational dataset, gathered using mobile phones, we analyze the problem of identifying transient and non-important links, as well as how to highlight important social interactions. Applying the Bluetooth signal strength parameter to distinguish between observations, we demonstrate that weak links, compared to strong links, have a lower probability of being observed at later times, while such links—on average—also have lower link-weights and probability of sharing an online friendship. Further, the role of link-strength is investigated in relation to social network properties.

Citation: Sekara V, Lehmann S (2014) The Strength of Friendship Ties in Proximity Sensor Data. PLoS ONE 9(7): e100915. doi:10.1371/journal.pone.0100915

Editor: Christopher M. Danforth, University of Vermont, United States of America

Received: January 17, 2014; **Accepted:** May 31, 2014; **Published:** July 7, 2014

Copyright: © 2014 Sekara, Lehmann. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: The authors were funded by a Young Investigator Grant from the Villum Foundation (High Resolution Networks, awarded to SL). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors confirm that one of the authors, Sune Lehmann, is a PLOS ONE Editorial Board member. This does not alter the authors' adherence to PLOS ONE Editorial policies and criteria.

* Email: vese@dtu.dk

Introduction

Recognizing genuine social connections is a central issue within multiple disciplines. When do connections happen? Where do they take place? And with whom is an individual connected? These questions are important when working to understand and design urban areas [1,2], studying close-contact spreading of infectious diseases [3–5], or organizing teams of knowledge workers [6–9]. In spite of their importance, measuring social ties in the real world can be difficult.

In classical social science the standard approach is to use self-reported data. This method, however, is only practical for relatively small groups and suffers from cognitive biases, errors of perception, and ambiguities [10]. Further, it has been shown that the ability to capture behavioral patterns via self-reported data is limited in many contexts [11]. A different approach for uncovering social behavior is to use digital records from emails and cell phone communication [12–19]. Although such analyses have improved our understanding of social ties, they have left many important questions unanswered—are electronic traces a valid proxy for real social connections? Eagle et al. [20] began to answer this question by including a spatial component as part of their data, using the short range ($\sim 10m$) Bluetooth sensor embedded in study participants' smartphones to measure physical proximity. Their results show that proximity data closely reflects social interactions in many cases. But since it is easy to think of examples where reciprocal Bluetooth detection does not correspond to social interaction (e.g. transient co-location in dining hall) the question remains, which observations correspond to actual social connections and which are just noise?

Multiple alternatives have been proposed to Bluetooth for sensor-driven measurement of social interactions, each with particular strengths and weaknesses [21–31]. For example, Radio

Frequency Identification (RFID) badges have short interaction ranges ($1–4m$) and measure only face-to-face interactions, thus solving many of the resolution problems posed by Bluetooth [30,31]. This approach, however, confines interactions to occur within specific areas covered by special radio receivers and requires participants to wear custom radio tags on their chests at all times—unlike Bluetooth which is ubiquitous across many types of modern electronic devices.

Our investigation digs into the role of Bluetooth signal strength, using a dataset obtained from applications running on the cell phones of 134 students at a large academic institution. Each phone records and sends data to researchers about call and text logs, Bluetooth devices in nearby proximity, WiFi hotspots in proximity, cell towers, GPS location, and battery usage [32]. In addition, we combine the data collected via the phones with online data, such as social graphs from Facebook for a majority of the participants. The study continuously gathers data, but in this paper we focus on Bluetooth proximity data gathered for 119 days during the academic year of 2012–2013. Specifically, we focus on the received signal strength parameter and propose a methodology that applies signal strength to distinguish between social and non-social interactions. We concentrate on the signal parameter because it is present in a majority of digitally recorded proximity datasets [30,32,33] and in addition, it also suggests a rough estimate for the distance between two devices. Applying the method on our data, we compare the findings to a null model and demonstrate how removing links with low signal strength influences network structure. Moreover, we use estimated link-weights and an online dataset to validate the friendship-quality of removed links.

Materials and Methods

Dataset

We distributed phones among students from four study lines (majors), where each major was chosen based on the fraction of students interested in participating in the project. This selection method yielded a coverage of >93% of students per study line, enabling us to capture a dense sample of the social interactions between subjects. Such high coverage of internal connections within a social group, with respect to the density of social interactions combined with the duration of observation, has not been achieved in earlier studies [20,30].

The data collector application installed on each phone follows a predefined scanning time table, which specifies the activation and duration of each probe. Proximity data is obtained by using the Bluetooth probe. Every 300 seconds each phone performs a Bluetooth scan that lasts 30 seconds. During the scan it registers all discoverable devices within its vicinity (5–10m) along with the associated received signal strength indicator (RSSI) [34]. Recorded proximity data is of the form (i, j, t, s) , denoting that person i has observed j at time t with signal strength s . Only links between experiment participants are considered, comprising a dataset of 2 183 434 time ordered edges between 134 nodes, see Table 1 for more information. Data collection, anonymization, and storage was approved by the Danish Data Protection Agency, and complies with both local and EU regulations. Written informed consent was obtained via electronic means, where all invited participants digitally signed the form with their university credentials. Along with the mobile phone study we also collected Facebook graphs of the participants. Not all users donated their data since this was voluntary, however we obtained a user participation of ~88% (119 users and 1018 Facebook friendships). For the missing 12% of users, we assume they do not share any online friendships with the bulk of participants.

Identifying links

Independent of starting conditions, the scanning framework on one phone will drift out of sync with the framework on other phones after a certain amount of time, thus the phones will inevitably scan in a desynchronized manner. This desynchronization can mainly be attributed to: internal drift in the time-protocol of each phone, depletion of the battery,

and users manually turning phones off. To account for irregular scans, we divide time into windows (bins) of fixed width and aggregate the Bluetooth observations within each time-window into a weighted adjacency matrix. The complete adjacency matrix is then given by: $W = (W^{(\Delta t_1)}, W^{(\Delta t_2)}, \dots, W^{(\Delta t_n)})$, where each link is weighted by its signal strength and where Δt_i indicates window number i . These matrices generally assume a non-symmetric form, i.e. person A might observe B with signal strength s while person B observes A with strength s' , or not at all. The scanning frequency of the application sets a natural lower limit of the network resolution to 5 minutes. If we are interested in the social dynamics at a different temporal resolution we can aggregate the adjacency matrices and retain entries according to some heuristic (e.g. with the strongest signal). Depending on the level of description (monthly, weekly, daily, hourly, or every 5 minutes) the researcher must think carefully about the definition of a network connection. Frameworks for finding the best temporal resolution, so called *natural timescales* have for specific problems been investigated by Clauset and Eagle [35], and Sulo et al. [36]. In this paper, however, we are interested in the identification and removal of non-social proximity links, so aggregating multiple time-windows is not a concern here. Henceforth we solely work with 5 minutes time-bins.

The Bluetooth probe logs all discoverable devices within a sphere with a radius of 5–10 meters—walls and floor divisions reduce the radius, but the reduction in signal depends on the construction materials [37]. Blindly taking proximity observations as a ground truth for social interactions will introduce both false negative and false positive links in the social network. False negative links are typically induced by hardware errors beyond our control, thus we focus on identifying false positive links. We therefore propose to identify non-social or noisy proximity links via the signal strength parameter. The parameter can be thought of as a proxy for the relative distance between devices, since most people carry their phones on them, it in principle also suggests the separation distance between individuals.

Previous work has applied Bluetooth signals to estimate the position of individuals [38–41] but studies by Hay [42], and Hossein et al. [43] have revealed signal strength as an unsuitable candidate for accurately estimating location. However, the complexity of the problem can greatly be reduced by focusing on the relative distance between individuals rather than position. In theory, the transmitted power between two antennae is inversely proportional to the distance squared between them [44]. Reality is more complicated, due to noise and reflection caused by obstacles.

We use the ideal result as a reference while we perform empirical measurements to determine how signal strength depends on distance. Two devices are placed on the ground in a simulated classroom setting, where we are able to control the relative distance between them. The resulting measurements are plotted in Fig. 1A. As is evident from the figure, there is a large variance in the measured signal strength values for each fixed distance. However, as both phones exhibit the same variance we can exclude faulty hardware; further, environmental noise such as interference from other devices, or solar radiation can also be dismissed since there appear no daily patterns in the data. But we observe multiple bands or so-called modes onto which measurements collapse. Ladd et al. [33] noted a similar behavior for the received signal strength of WiFi connections, both are phenomena caused by non-Gaussian distributed noise. The empirical measurements form a foundation for understanding signal variance as

Table 1. Data overview.

	Total	Average pr. time-bin
Nodes (Users)	134	17.32
Edges (Dyads)	2 183 434	62.50
Time-bins	34 272	-
Average clustering	0.85	0.26
Average degree	103.51	2.41

Statistics showing the number of total (aggregated) and average values of network properties. Time-bins span five minutes and cover the entire 119 day period, including weekends and holidays. For the average values we only take active nodes into account, i.e. people that have observed another person or been observed themselves in that specific time-bin. Network properties are calculated for the full aggregated network and as averages over each temporal network slice.

doi:10.1371/journal.pone.0100915.t001

a function of distance, but they were performed in a controlled environment. In reality, there are a multitude of ways to carry a smartphone: some carry it around in a pocket, others in a bag. Liu and Striegel [45] investigated how these various scenarios influence the received signal strength—their results indicate only minor variations, hence we conclude that the general behavior is similar to the measurements shown in the figure. Further, social interactions are not only limited to office environments, so we have re-produced the experiment outdoors and in basement-like settings; the results are similar.

Bi-directional observations yield at most two observations per dyad per 5-minute time-bin, we can average over the measurements (Fig 1B), or take the maximal value (Fig 1C). Fig. 2 shows the distributions of signal strength for each respective distance. For raw data, Fig. 2A, we observe a localized zero-distance distribution while the 1, 2, and 3-m distributions overlap considerably. Averaging over values per time-bin smoothes out and compresses the distributions, but the bulk of the distributions still overlap (Fig. 2B). Taking only the maximal signal value into account separates the distributions more effectively (Fig. 2C). The reasoning behind choosing the maximal signal value is that phones are physically at different locations and we expect the distance to be maximally reflected in the distributions.

Thus, by thresholding observations on signal strength, we can filter out proximity links that are likely to be further away than a certain distance. By doing so we are able to emphasize links that are more probable of being genuine social interactions, while minimizing noise and filtering away non-social proximity links. From the behavioral data we count the number of appearances per dyad and assign the values as weights for each link. Link weights follow a heavy-tailed distribution, with a majority of pairs only observed a few times (low weights), a social behavior that has previously been observed by Onnela et al. [15]. Based on their weight we divide links into two categories: weak and strong. A link is defined as ‘weak’ if it has been observed (on average) less than once per day during the data collection period, remaining links are characterized as ‘strong’. An effective threshold should maximize the number of removed weak links, while minimizing the loss of strong links. Fig. 3 depicts the number of weak and strong links as a function of threshold value. We observe that, as we increase the threshold, the number of weak links decreases linearly, while the number of strong links remains roughly constant and then drops off suddenly. Taking into account both the maximum-value distance distributions (Fig. 2C) and link weights (Fig. 3), we choose the value (-80 dBm) that optimizes the ratio between strong and weak links. In a large majority of cases, this corresponds to




Figure 1. Bluetooth signal strength (RSSI) as a function of distance. **A:** Scans between two phones. Measurements are per distance performed every five minutes over the course of 7 days. Mean value and standard deviation per distance are respectively $\mu_0 = -45.13 \pm 1.56$ dBm, $\mu_1 = -77.48 \pm 4.15$ dBm, $\mu_2 = -82.03 \pm 4.57$ dBm, and $\mu_3 = -85.49 \pm 2.75$ dBm. **B:** Average of the values in respective time-bins. Summary statistics are: $\mu_0^{\text{avg}} = -45.13 \pm 1.20$ dBm, $\mu_1^{\text{avg}} = -77.46 \pm 2.90$ dBm, $\mu_2^{\text{avg}} = -81.99 \pm 3.17$ dBm, and $\mu_3^{\text{avg}} = -85.45 \pm 1.88$ dBm. **C:** Maximal value per time-bin. The mean value and standard deviation per distance are: $\mu_0^{\text{max}} = -44.41 \pm 1.11$ dBm, $\mu_1^{\text{max}} = -75.09 \pm 3.24$ dBm, $\mu_2^{\text{max}} = -79.25 \pm 3.47$ dBm, and $\mu_3^{\text{max}} = -83.88 \pm 2.00$ dBm. The measurements cover hypothetical situations where individuals are far from each other and on either side of a wall. doi:10.1371/journal.pone.0100915.g001

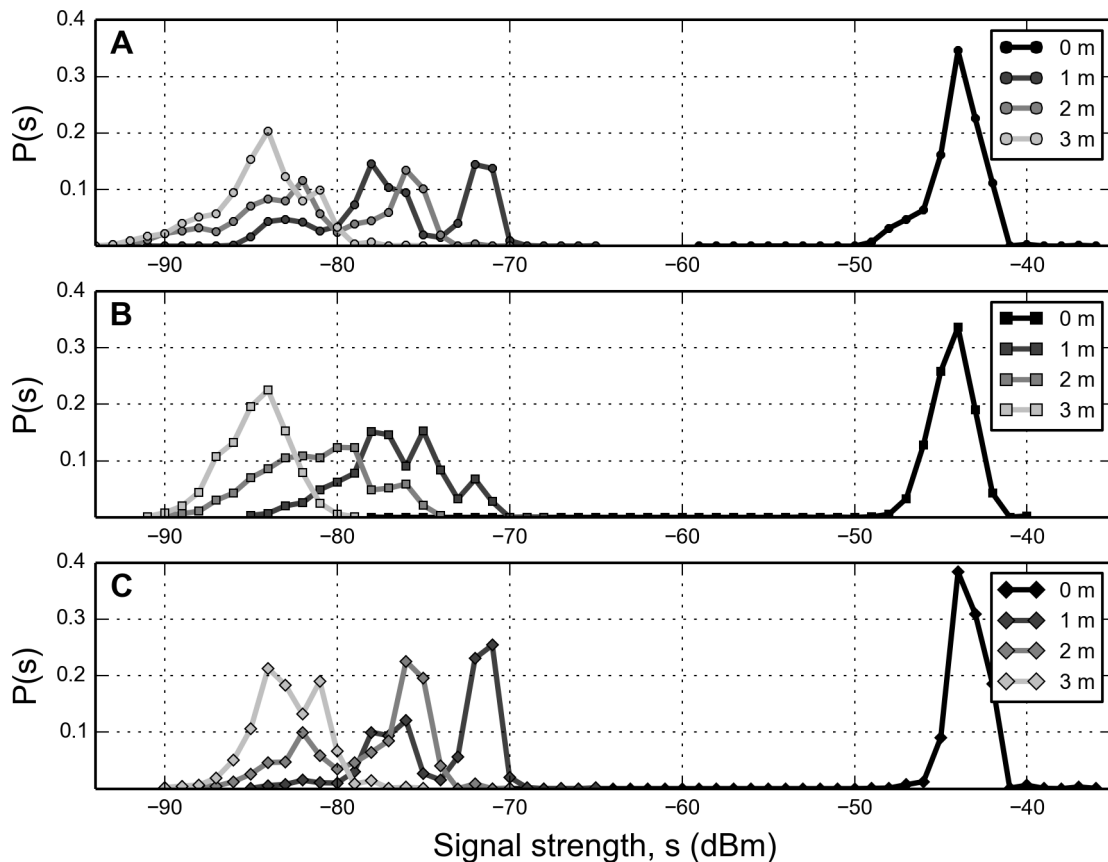


Figure 2. Distributions of signal strength for the respective distances. **A:** Raw data. Measurements from both phones are statistically indistinguishable and are collapsed into single distributions, i.e. there is no difference between whether A observes B or vice versa. **B:** Average of signal strength per time-bin. **C:** Maximal value of signal strength per time-bin.
doi:10.1371/journal.pone.0100915.g002

interactions that occur within a radius of 0–2 meters—a distance which Hall [46] notes as a typical social distance for interactions among close acquaintances.

Removing links

This section outlines various strategies for removing non-social links from the network. Fig. 4A shows an illustration of the raw proximity data for a single time-bin, a link is drawn if either $i \rightarrow j$ or $j \rightarrow i$. Thickness of a link represents the strength of the received signal. For the thresholded network (Fig. 4B) we remove links according to the strength of the signal (where we assume the weaker the signal the greater the relative distance between two persons). To estimate the effect of the threshold we compare it to a null model, where we remove the same number of links, but where the links are chosen at random, illustrated in Fig. 4C. To minimize any noise the random removal might cause, we repeat the procedure $n=100$ times, each time choosing a new set of random links, with statistics averaged over the 100 repetitions. As a reference, to check whether thresholding actually emphasizes social proximity links, we additionally compare it to a control network, where we remove the same amount of links, but where the links have signal strengths *above* or *equal* to the threshold, Fig. 4D. This procedure is also repeated n times. In a situation where there are more links below the threshold than above, we will remove fewer links for the latter compared to the other networks.

Results

Network properties

Now that we have determined a threshold for filtering out non-social proximity links, let us study the effects on the network properties. Thresholding weak links does not significantly influence the number of nodes present (N) in the network (Fig. 5A), while the number of links (M) is substantially reduced (Fig. 5B). On average we remove 2.38 nodes and 32.18 links per time-bin. Social networks differ topologically from other kinds of networks by having a larger than expected number of triangles [47], thus clustering is a key component in determining the effects of thresholding. Fig. 6 suggests that we are, in fact, keeping real social interactions: random removal disentangles the network and dramatically decreases the clustering coefficient, while thresholding conserves most of the average clustering. Calculating the average ratio ($\langle\langle c_T \rangle\rangle / \langle\langle c_N \rangle\rangle$) between clustering in the thresholded ($\langle\langle c_T \rangle\rangle$) and the null networks ($\langle\langle c_N \rangle\rangle$) reveals that c_T on average is 2.38 larger. These findings emphasize that a selection process based on signal strength greatly differs from a random one.

Link evaluation

Sorting links by signal strength and disregarding weak ones greatly reduces the number of links, but do we remove the correct links, i.e. do we get rid of noisy, non-social links? The fact that clustering remains high in spite of removing a large fraction of




Figure 3. Number of links per type as a function of threshold value. Links are classified as weak if they are observed less than 120 times in the data, i.e. links that on average are observed less than once per day—otherwise they are classified as strong. Grouping students into study lines, reveals that links within each study line have an almost uniform distribution of weights while links across study lines are distributed according to a heavy-tailed distribution. A threshold of -80 dBm (gray area) removes 1159 weak and 387 strong links and classifies 97.6% of inter-study line links as weak and 86.7% of intra-study line links as strong.
doi:10.1371/journal.pone.0100915.g003

links is a good sign, but we want to investigate this question more directly. To do so, we divide the problem into two timescales; a short one where we consider the probability that a removed link might reappear a few time-steps later, and a long where we evaluate the quality of a removed link according to certain network properties. Let's first consider the short time-scale. We assume that human interactions take place on a time-scale that is mostly longer than the 5-minute time-bins we analyze here. Thus, if a noisy link is removed, the probability that it will re-appear in one of the immediately following time-steps should be low, since no interaction is assumed to take place. However, we expect the probability to be significantly greater than zero, since even weak (non-social) links imply physical proximity. Similarly, if we (accidentally) remove a social link, the probability that it will appear again should be high, since the social activity is expected to continue to take place.

Let us formalize this notion. Consider a link e that is removed at time t , the probability that the link will appear in the next time-

step is $p(t+1|e,t)$. Generalizing this we can write the probability that any removed link will appear in all the following n time-steps as:

$$p(t+1, \dots, t+n|t) = \frac{\text{no. links removed at } t \text{ present at } t+1 \cap \dots \cap t+n}{\text{no. links removed at } t} \quad (1)$$

Fig. 7A illustrates that thresholded links in subsequent time-steps are observed less frequently than both null and control links. To compare with the worst possible condition, we compare data from each thresholded time-bin with the *raw data* from the next bin (where the raw data contains many weak links). In spite of this, we observe a clear advantage of distinguishing between links with weak and strong signal strengths. If we look at values for $t+1$, the first subsequent time-step, the probability of re-occurrence in the thresholded network is about 12% lower than for the null model, and as we look to later time-steps, the gap widens.

A different set of social dynamics unfolds on longer timescales where the class schedule imposes certain links to appear periodically, e.g every week. Here we determine impact of removing links in two ways. First, we use total link weights and second, we use online friendship status. Friends meet frequently; we capture this behavior by using the total number of observations of a certain dyad to estimate the weight of a friendship (again, counted in the raw network). Thus, we evaluate the quality of a removed links by considering its total weight compared to the weight of other links present in the same time-bin. However, since multiple links are removed per time-bin we are more interested in the average,

$$q_t = \frac{\text{Avg. weight of removed links at } t}{\text{Avg. weight of all links present at } t} \quad (2)$$

This estimates, per time-bin, whether removed links on average have weights below, close to, or above the mean. Note that the measure is intended to estimate the quality of removed links and is therefore not defined for bins where zero links are removed. Fig. 7B indicates difference in link selection processes. Choosing links at random (null network) removes both strong and weak links with equal probability, thus on average this corresponds to the mean weight of links present. Compared to null, the thresholded network removes links with weights below average, indicating that removed links are less frequently observed and therefore also less likely to be real friendships. The control case displays an diametrical behavior, on average, it removes links with higher weights.




Figure 4. Networks. **A:** Raw network; shows all observed links for a specific time-bin. Thickness of a link symbolizes the maximum of the received signal strengths. **B:** Thresholded network, we remove links with received signal strengths below a certain threshold, where dotted lines indicate the removed links. **C:** Null model; with respect to the previous network we remove the same amount of links, but where the links are chosen at random. **D:** Control network, a similar amount of links with signal strength above or equal to the threshold are removed.
doi:10.1371/journal.pone.0100915.g004




Figure 5. Network statistics. Properties are highly dynamic but on average we observe 17.32 nodes and 62.50 links per time-bin. **A:** Number of nodes N as a function of time. Only active nodes are counted, i.e. people that have observed another person or been observed themselves. Dynamics are shown for two weeks during the 2013 spring semester, clearly depicting both daily and weekly patterns. Data markers are omitted to avoid visual clutter. On average thresholding removes 3.06 nodes during weekends and holidays, and 2.38 during regular weekdays. **B:** Number of links M as a function of time. 10.60 links are on average removed during weekends/holidays, and 32.21 are removed during weekdays. doi:10.1371/journal.pone.0100915.g005

The second method to evaluate the link-selection processes compares the set of removed links with the structure of an online social network, i.e. if a removed proximity link has an equivalent online counterpart. We estimate the quality by measuring the fraction of removed links with respect to those present at time t .

$$q_t^{FB} = \frac{\text{no. of FB links removed at } t}{\text{no. of FB links present at } t} \quad (3)$$

The quality measure is essentially a ratio, i.e. it can assume values




Figure 6. Average clustering. Only active nodes, i.e. nodes that are part of at least one dyad contribute to the average, the rest are disregarded. Average clustering is calculated according to the definition in [48]. Since social activity in groups larger than two individuals results in network triangles, the fact that clustering is not significantly reduced by thresholding (compared to the null model) provides evidence that we are preserving social structure in spite of link removal. doi:10.1371/journal.pone.0100915.g006

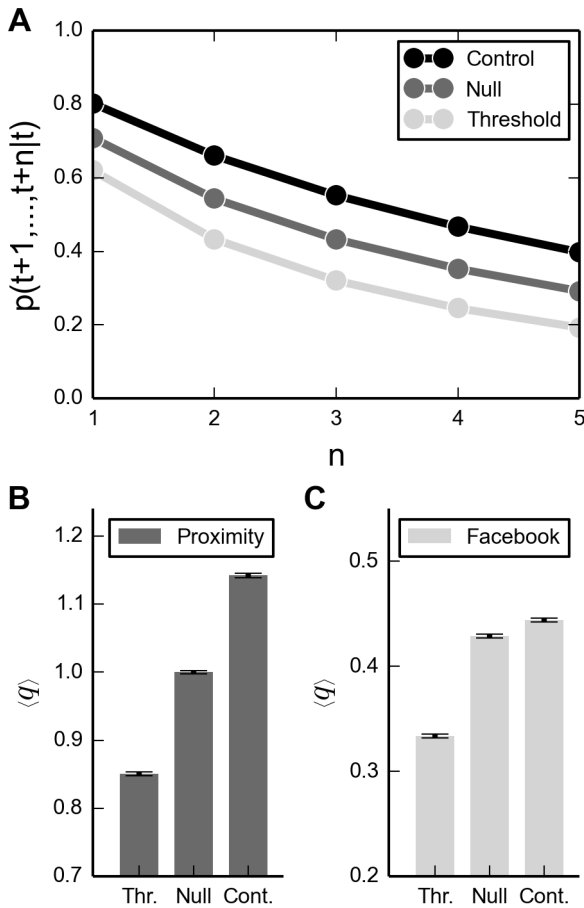


Figure 7. Link evaluation. **A:** Probability of link reappearance. For each selection process we remove a specific set of links. In the thresholded network, we remove links with weak signal strength. For the null network, we remove links at random. Lastly, in the control network case we remove strong links. The probability for links to reappear within all the next n time-steps is calculated using Eq. 1 and averaging over all time-bins. Boundary conditions are not applied and the reappearance probability for the last $n=5$ bins is not taken into account. **B:** Quality measure for proximity data. **C:** Quality measure for the online data. For each time-bin we calculate q_t as defined in Eq. 2 and 3. Brackets indicate a temporal average across all time-bins and value are shown for all three network types. doi:10.1371/journal.pone.0100915.g007

$0 \leq q_t^{FB} \leq 1$ depending on the fraction of links that are removed. Bins with zero Facebook friendships are disregarded since they contain no information regarding the online social network. Fig. 7C shows that random removal (null network), on average, removes $\sim 43\%$ of online friendships, while the thresholded network removes $\sim 33\%$, a 10 percent point difference. For

References

- Sun J, Yuan J, Wang Y, Si H, Shan X (2011) Exploring space–time structure of human mobility in urban space. *Physica A: Statistical Mechanics and its Applications* 390: 929–942.
- Sevtsuk A, Ratti C (2010) Does urban mobility have a daily routine? learning from the aggregate data of mobile networks. *Journal of Urban Technology* 17: 41–60.
- Liljeros F, Edling CR, Amaral LAN, Stanley HE, Åberg Y (2001) The web of human sexual contacts. *Nature* 411: 907–908.
- Mossong J, Hens N, Jit M, Beutels P, Auranen K, et al. (2008) Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS Medicine* 5: e74.

comparison, the control network removes $\sim 44\%$ of the online links. Further, redoing the analysis for a dataset comprised only of users for which we have both proximity and online data for, does not significantly alter the results.

Facebook links are not necessary good indicators for strong friendships, but are more likely to correspond to real social interactions. In spite of this, both Fig. 7B and C support that distinguishing between strong and weak proximity links tends to emphasize real social interactions: on average thresholded links have lower edge weights and remove fewer Facebook friendships compared to both the null-model and the control.

Discussion

The availability of electronic datasets is increasing, so the question of how well can we use these electronic *clicks* to infer actual social interactions is important for effectively understanding processes such as relational dynamics, and contagion. Sorting links based on their signal strength allows us to distinguish between strong and weak ties, and we have argued that thresholding the network emphasizes social proximity links while eliminating some noise.

Simply thresholding links based on signal strength is not a perfect solution. In certain settings we remove real social connections while noisy links are retained. Our results indicate that the proposed framework is better at identifying strong links than removing them. A trend which the link-reappearance probability, link-weights, and online friendship analysis support. Compared to the baseline we achieve better results than just assuming all proximity observations as real social interactions. But determining whether a close proximity link corresponds to an actual friendship interaction is much more difficult. Multiple scenarios exist where people are in close contact but are not friends, one obvious example is queuing. Each human interaction has a specific social context, so an understanding of the underlying social fabric is required to fully discern when a close proximity link is an actual social meeting. This brings us back to the question of how to determine a real friendship from digital observations (cf. [10]). Close proximity may not be the best indicator of friendship; call logs, text logs, and geographical positions are all factors which coupled with information from the Bluetooth probe could give us a better insight into social dynamics and interactions.

Acknowledgments

We thank L. K. Hansen, A. Stopczynski, and P. Sapie ́ z yński for many useful discussions and A. Cuttone for proofreading the manuscript.

Author Contributions

Conceived and designed the experiments: VS SL. Performed the experiments: VS SL. Analyzed the data: VS SL. Wrote the paper: VS SL.

9. de Montjoye YA, Stopczynski A, Shmueli E, Pentland A, Lehmann S (2014) The strength of the strongest ties in collaborative problem solving. *Scientific reports* 4.
10. Wuchty S (2009) What is a social tie? *Proceedings of the National Academy of Sciences* 106: 15099–15100.
11. Watts DJ (2007) A twenty-first century science. *Nature* 445: 489–489.
12. Eckmann JP, Moses E, Sergi D (2004) Entropy of dialogues creates coherent structures in e-mail traffic. *Proceedings of the National Academy of Sciences of the United States of America* 101: 14333–14337.
13. Barabasi AL (2005) The origin of bursts and heavy tails in human dynamics. *Nature* 435: 207–211.
14. Kossinets G, Watts DJ (2006) Empirical analysis of an evolving social network. *Science* 311: 88–90.
15. Onnela JP, Saramäki J, Hyvönen J, Szabó G, Lazer D, et al. (2007) Structure and tie strengths in mobile communication networks. *Proceedings of the National Academy of Sciences* 104: 7332–7336.
16. Gonzalez MC, Hidalgo CA, Barabasi AL (2008) Understanding individual human mobility patterns. *Nature* 453: 779–782.
17. Lazer D, Pentland A, Adamic L, Aral S, Barabási AL, et al. (2009) Computational social science. *Science* 323: 721–723.
18. Song C, Qu Z, Blumm N, Barabási AL (2010) Limits of predictability in human mobility. *Science* 327: 1018–1021.
19. Bagrow JP, Wang D, Barabási AL (2011) Collective response of human populations to large-scale emergencies. *PLoS ONE* 6: e17680.
20. Eagle N, Pentland A, Lazer D (2009) Inferring friendship network structure by using mobile phone data. *Proceedings of the National Academy of Sciences* 106: 15274–15278.
21. Haritaoglu I, Harwood D, Davis LS (2000) W4: Real-time surveillance of people and their activities. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 22: 809–830.
22. Polastre J, Szewczyk R, Culler D (2005) Telos: enabling ultra-low power wireless research. In: *Information Processing in Sensor Networks, 2005. IPSN 2005. Fourth International Symposium on. IEEE*, pp. 364–369.
23. Salathé M, Kazandjieva M, Lee JW, Lewis P, Feldman MW, et al. (2010) A high-resolution human contact network for infectious disease transmission. *Proceedings of the National Academy of Sciences* 107: 22020–22025.
24. Rosenstein B (2008) Video use in social science research and program evaluation. *International Journal of Qualitative Methods* 1: 22–43.
25. Olguín DO, Waber BN, Kim T, Mohan A, Ara K, et al. (2009) Sensible organizations: Technology and methodology for automatically measuring organizational behavior. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on* 39: 43–55.
26. Kjærsgaard MB, Nurmi P (2012) Challenges for social sensing using wifi signals. In: *Proceedings of the 1st ACM workshop on Mobile systems for computational social science. ACM*, pp. 17–21.
27. Wyatt D, Choudhury T, Kautz H (2007) Capturing spontaneous conversation and social dynamics: A privacy-sensitive data collection effort. In: *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on. IEEE*, volume 4, pp. IV–213.
28. Carreras I, Matic A, Saar P, Osmani V (2012) Comm2sense: Detecting proximity through smartphones. In: *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2012 IEEE International Conference on. IEEE*, pp. 253–258.
29. Wang D, Pedreschi D, Song C, Giannotti F, Barabasi AL (2011) Human mobility, social ties, and link prediction. In: *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM*, pp. 1100–1108.
30. Cattuto C, Van den Broeck W, Barrat A, Colizza V, Pinton JF, et al. (2010) Dynamics of person-to-person interactions from distributed RFID sensor networks. *PLoS ONE* 5: e11596.
31. Barrat A, Cattuto C (2013) *Temporal networks of face-to-face human interactions*. In: *Temporal Networks*, Springer. pp. 191–216.
32. Stopczynski A, Sekara V, Sapiezynski P, Cuttone A, Madsen MM, et al. (2014) Measuring largescale social networks with high resolution. *PLoS ONE* 9: e95978.
33. Ladd AM, Bekris KE, Rudys A, Kavradi LE, Wallach DS (2005) Robotics-based location sensing using wireless ethernet. *Wireless Networks* 11: 189–204.
34. Shorey R, Miller BA (2000) The bluetooth technology: merits and limitations. In: *Personal Wireless Communications, 2000 IEEE International Conference on. IEEE*, pp. 80–84.
35. Clauset A, Eagle N (2012) Persistence and periodicity in a dynamic proximity network. *arXiv preprint arXiv:12117343*.
36. Sulo R, Berger-Wolf T, Grossman R (2010) Meaningful selection of temporal resolution for dynamic networks. In: *Proceedings of the Eighth Workshop on Mining and Learning with Graphs. ACM*, pp. 127–136.
37. Cheung KC, Intille SS, Larson K (2006) An inexpensive bluetooth-based indoor positioning hack. *Proc UbiComp06 Extended Abstracts*.
38. Anastasi G, Bandelloni R, Conti M, Delmastro F, Gregori E, et al. (2003) Experimenting an indoor bluetooth-based positioning service. In: *Distributed Computing Systems Workshops, 2003. Proceedings. 23rd International Conference on. IEEE*, pp. 480–483.
39. Bruno R, Delmastro F (2003) Design and analysis of a bluetooth-based indoor localization system. In: *Personal wireless communications. Springer*, pp. 711–725.
40. Madhavapeddy A, Tse A (2005) A study of bluetooth propagation using accurate indoor location mapping. In: *UbiComp 2005: Ubiquitous Computing, Springer*. pp. 105–122.
41. Zhou S, Pollard JK (2006) Position measurement using bluetooth. *Consumer Electronics, IEEE Transactions on* 52: 555–558.
42. Hay S, Harle R (2009) Bluetooth tracking without discoverability. In: *Location and context awareness, Springer*. pp. 120–137.
43. Hossain A, Soh WS (2007) A comprehensive study of bluetooth signal parameters for localization. In: *Personal, Indoor and Mobile Radio Communications, 2007. PIMRC 2007. IEEE 18th International Symposium on. IEEE*, pp. 1–5.
44. Friis HT (1946) A note on a simple transmission formula. *proc IRE* 34: 254–256.
45. Liu S, Striegel A (2011) Accurate extraction of face-to-face proximity using smartphones and bluetooth. In: *Computer Communications and Networks (ICCCN), 2011 Proceedings of 20th International Conference on. IEEE*, pp. 1–5.
46. Hall ET (1990) *The hidden dimension*. Anchor Books New York.
47. Newman ME, Park J (2003) Why social networks are different from other types of networks. *Physical Review E* 68: 036122.
48. Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393: 440–442.