Optimizing targeted vaccination across cyber-physical networks: an empirically-based mathematical simulation study
Supplementary Information

Enys Mones, Arkadiusz Stopczynski, Alex 'Sandy' Pentland, Nathaniel Hupert & Sune Lehmann

October 1, 2017

Contents

S1   Dataset

We consider data collected in the Copenhagen Networks Study, spanning years between 2013 and 2015. The data has been collected from a densely-connected freshman population of approximately 1 000 students at a large European university (Technical University of Denmark) and contains high-resolution traces including close-proximity interactions, telecommunication, online social networks, and geographical location. Here, densely-connected refers to the high frequency of physical proximity contacts as well as online communication between the individuals. Majority of the data has been collected with custom-built application installed on smartphones provided to the participants (Google Nexus 4). Full details of the study can be found in Ref. [1].

Network	V_{active}	E_{dynamic}	E_{static}	$\langle k \rangle_{\text{dynamic}}$	$\langle k \rangle_{\text{static}}$
Full-range proximity	532	2 670 547	69 055	1.245	259.6
Short-range proximity	518	428 481	20 690	0.200	77.78
Facebook	410	3 321	1 261	0.007	4.741
Call	345	2 134	354	0.004	1.331

Table 1: Basic statistics of the physical proximity and social networks: active individuals (V_{active}), those having at least one link; total number of temporal links (E_{dynamic}); number of diads or static links (E_{static}); dynamic degree ($\langle k \rangle_{\text{dynamic}}$), that is the average degree for single time bins, averaged over all time bins; static degree ($\langle k \rangle_{\text{static}}$). Average degrees are calculated using the total population size, i.e., $V_{\text{total}} = 532$.

All metrics are calculated over the month February 2014.

In this manuscript we select individuals with a data quality of at least 60% during period we focus on (February 2014), defined by the time coverage of the Bluetooth scans. Data quality is calculated by binning all Bluetooth scans (that are used for inferring physical proximity interactions) and removing participants that have appeared in scans in less than 60% of the total bins (i.e., 4838 out of 8064 time bins). This results in a population of 532 individuals. Using Bluetooth, call detail records (CDRs), and Facebook feed data, we construct four distinct networks of interactions between the selected individuals. We have explored the effects of subsampling the dataset elsewhere[2] and find that the dataset is robust to subsampling (within reasonable limits)

S1.1 Network creation

Physical proximity networks are built using Bluetooth scans between devices. Each time the device ID of participant i is listed in a Bluetooth scan of the device of participant j at time t, there is a link $\gamma_{ijt} = s$ between the corresponding participants, where s denotes the received signal strength indicator (RSSI). From the raw counts, we build temporally-binned contacts with bin size of 5 minutes, such that if there is any observation in the n-th time bin, we add a link. As we do not expect false positives in a Bluetooth scan, all edges are undirected even if they are observed in only one direction. We also extend local star-like graphs to cliques, to avoid missing links in case of a single (or very few) devices reporting data in particular time bin. Such pre-processing results in physical proximity network capturing interactions at the distance of $10 - 15$ meters, which we denote *full-range* network. The network is described by a binary-valued adjacency matrix $A_{i \times j \times t}$ with $a_{ijt} = 1$ when interaction is present and $a_{ijt} = 0$ otherwise. We also create *short-range* network by thresholding the links based on their RSSI values, i.e., we restrict the links only to those for which $\gamma_{ijt} > s_0$ with $s_0 = -75\,\text{dB}$. It has been shown that this value corresponds to a distance up to approximately 1m [3, 4]. Thus, we create two types of proximity networks, one with full-range interactions and one with only very close proximity events. This constraint on the signal strength results in a more sparse and tree-like network (18% of the original interactions are present after thresholding).

Digital social networks, corresponding to social behavior and social ties, are based on calls and Facebook activity. Each call between two individuals results in a corresponding link in the call network, and every interaction on Facebook creates a link in the *functional* Facebook network (interactions include comments, tags, messages, mentions, posts, etc). Here we do not use the temporal information in the social networks, i.e., we use networks as aggregated over the month of interest. We also do not consider directionality of calls nor directionality or type of Facebook interactions. It is also possible to build other types of networks by including text messages in addition to calls metadata or using the static friendship links between users (for *structural* Facebook graph). However, as we show in the sections Monitoring and Vaccination, these approaches do not result in significant differences in social monitoring and vaccination. Finally, we summarize the basic properties for the four networks in the study calculated over February 2014 in Table 1.
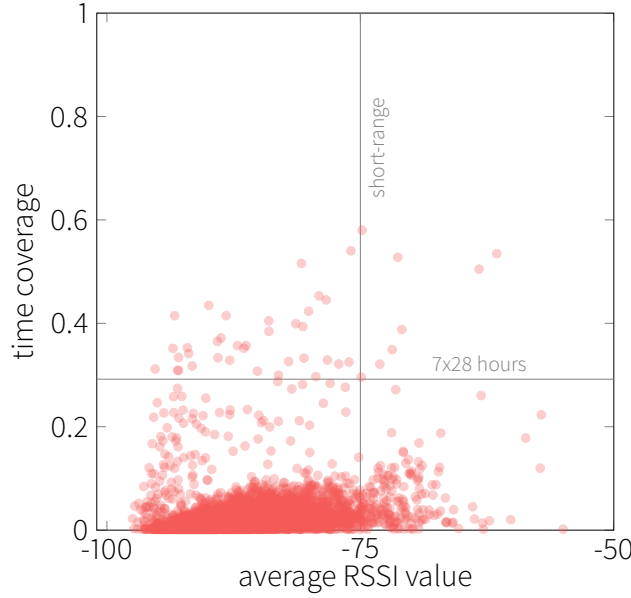
Figure 1: **Distribution of Bluetooth based proximity contacts** Total temporal coverage of proximity links versus their average RSSI, that is, the signal strength. Above the horizontal line are the links that describe contacts of frequent periods with a total length of more than 7×28 hours. To the left from the vertical lines lie the links that correspond to the short-range interactions ($RSSI > -75\,\text{dB}$).

S1.2   Limitations of physical proximity contacts

As Bluetooth signals can pass through walls, some fraction of the contact networks built from the Bluetooth scans are expected to be false contacts, as they are not able to transmit real physical diseases. To address this problem, we consider the signal strength (RSSI) of the links: it's feasible to assume that links passing through walls have a vanishing strength irrespective of the actual distance. Therefore, in the short-range network we don't expect fake contacts as it includes only links with a certain high signal strength. In case of full-range network, weak links are more common. However, from the simulation perspective those weak links have a relatively high impact on the spread that are also frequently re-appearing, i.e., signals corresponding to dorm walls. To investigate these links, in Fig. 1 we plot the total time coverage of the links versus their average signal strength. Signals crossing dormitory walls are located in the upper-left corner, i.e., the frequent weak links, however, as the figure shows, these links are rare and they cover a negligible fraction of the total temporal links.

Furthermore, the majority of infections take place during lunch breaks, classes and social events where individuals are located in the same closed area.

S1.3   Target group selection

For the epidemic monitoring and vaccination, we construct target groups based on different types of ranking: random, colocation, and social.

In some cases, ranking n individuals does not result in a well-defined order, for example, when the centrality of two or more individuals is the same. In case of such ties we randomly select individuals from the ranked list of individuals. For example, let us assume we can select a group of 10 individuals where the last two persons in the group have the same centrality as with 5 others among the rest of the population. Then after selecting the 8 individuals with distinct centralities, we randomly pick the remaining 2 people from the $5 + 2 = 7$ with tied centrality-scores. In each of these cases, we averaged the results over 10 groups with randomly sampled tied centralities as

described above.

For colocation groups, we consider the aggregated weight of individuals in the proximity networks. For each participant we calculate $w_i = \sum_{j,t} a_{ijt}$, that is, the total time (expressed in 5-minute bins) spent in the proximity of others. Note that in this case, each temporal link is considered separately, thus someone spending a given amount of time in a large group has higher weight than that of spending the same time in a small group. After ranking individuals by their weight, we select the top n individuals from the population as our target group. As shown in Ref. [5], in case of epidemic monitoring the colocation-based target group performs close to the optimal strategy achieved with a greedy algorithm. Vaccination of the colocation-based group is not strictly optimal, since it does not use the full temporal information. This group, however, approximates vaccination performance in the "full information" scenario very well [6]. To keep the terminology as simple as possible, we refer to the colocation strategy as optimal.

Although during vaccination the colocation-based group is not strictly optimal considering that no temporal information is utilized, it provides a reasonable upper bound of the vaccination performance with access to full information about close-proximity interactions. Keeping the above in mind, we refer to the colocation strategy as *optimal*.

Social target groups are based purely on information extracted from the social networks, with the assumption that those networks describe, at a fundamental level, social structure of the population. In order to assess the social role and importance of an individual, we calculate their centrality in the aggregated social networks, rank the population by the centrality, and pick the n most central individuals. The results in the manuscript are reported with closeness centrality, however, the exact centrality used does not affect the findings qualitatively. Similarly, even if link weights (i.e., the number of contacts between individuals) are considered, the corresponding results do not change and we do not observe significant gain in performance, suggesting that the extent of the ego-network is more relevant than the strength of the specific links.

In some cases, ranking n individuals does not result in a well-defined order, for example, when the centrality of two or more individuals is the same. In case of such ties we randomly select individuals from the ranked list of individuals. For example, let us assume we can select a group of 10 individuals where the last two persons in the group have the same centrality as with 5 others among the rest of the population. Then after selecting the 8 individuals with distinct centralities, we randomly pick the remaining 2 people from the $5 + 2 = 7$ with tied centrality-scores In each of these cases, we averaged the results over 10 groups with randomly sampled tied centralities as described above.

S1.4   Testing for quality bias

We perform a quality check on the data to ensure that the reported results are not simply a consequence of quality bias. We test if individuals with high-quality data (those having a high temporal coverage in Bluetooth scans) are also the ones that have many calls or increased activity on Facebook. As Bluetooth scanning happens passively on the phones and we consider data quality based on the number of scans (regardless whether they discovered other participants of the study or not), we expect no correlation between how socially active the participants are and their data quality. We plot call and Facebook event counts for each user versus their corresponding coverage of Bluetooth data, as shown in Fig. 2. For the comparison, we plot data calculated over a complete year and show both total and internal communications of participants. (for better visibility, we omitted the top < 10 largely deviating data points). Correlations are low and in many cases are not significant: $r(w, c_{\text{call}}) = -0.0343(0.4286)$, $r(w, c_{\text{call total}}) = 0.2012(0.0000)$, $r(w, c_{\text{Facebook}}) = 0.0009(0.9833)$, $r(w, c_{\text{Facebook total}}) = -0.0282(0.5149)$. We report Spearman correlations, not assuming linear relation between the intensity of the two activities (social and proximity), numbers in the parentheses denote the corresponding p-values. The results indicate that the image of social activity we get about the population is not driven by data quality variation. Note that the manuscript, we base our results only on internal phone calls.

As for the social channels, the likelihood of missing data in the Facebook channel is minimal as the feed of each user has been crawled every 24 hours, providing a high level of redundancy for data quality. We consider the quality of the call data instead. To investigate whether we perceive some students that are highly connected in the call network merely due to the quality of recording
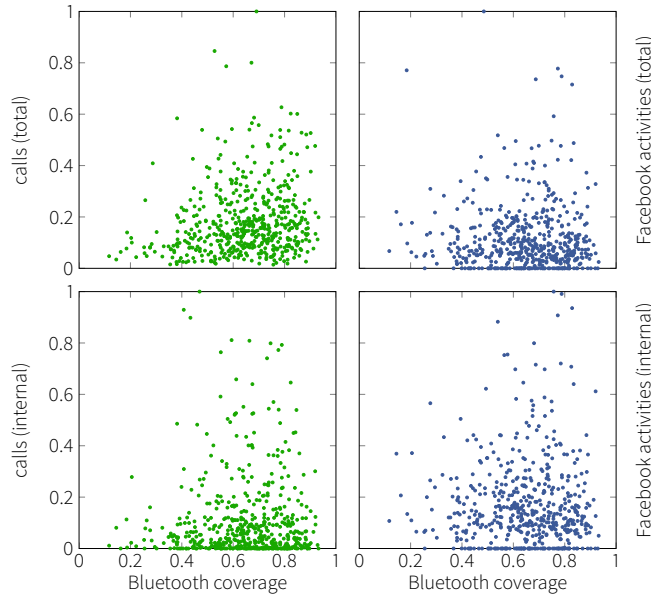
Figure 2: **Social activity versus data quality in the proximity layer** Call and Facebook activity intensity compared to the coverage seen by the Bluetooth scans for each participant in the current study: (top) total number of calls and Facebook events involving the given participant, (bottom) internal communication events, i.e., only between individuals considered in the simulations. Social network counts are normalized by the largest data point shown.

their events better than others', we compare the calling behavior pattern of all participants in the month of interest (February 2014) to their yearly behavior.

Figure 3 shows the individual deviations from their yearly average in February compared to the call intensity in the index month. We first sorted participants according to the total duration of their calls in the index month of February and assigned the number of calls in this month (n_i^0) to each participant i. As a second step, for every individual i we calculated their standard score over the year, that is, $(n_i^0 - \overline{n}_i)/\sigma(n_i)$, where $\overline{n}_i$ and $\sigma(n_i)$ are the mean and standard deviation calculated over the other months. If there was systemic missing data, we would observe a shift in the vertical distribution of the dots towards negative values, however, as we can see the dots are centered around a standard score of zero. This also means that February is not particularly different from the other months with respect to call behavior. Furthermore, as centrality values (color of the dots) indicate, high social centrality is not driven by count number nor deviation from the yearly behavior. This indicates that although we may expect certain amount of noise in the data, relevant measures on the population are not biased.

S2 Model

S2.1 SIR model

For simulation of epidemic outbreaks, we use a susceptible-infected-recovered (SIR) model taking place on the network of physical interactions of the individuals. The SIR model assumes that the entire population—except for the index cases—is initially susceptible. Whenever a susceptible individual has contact to an infected one, the susceptible individual becomes infected with probability β. Infected individuals recover after an average infectious period T_{inf} which is modeled in this study by a probability γ to recover at any time. Recovered individuals can not be re-infected. We use the temporal proximity networks obtained from the Bluetooth scans as the contact network. In all simulations, the size of the index (initially infected) group is 5 (approximately 1% of the population).
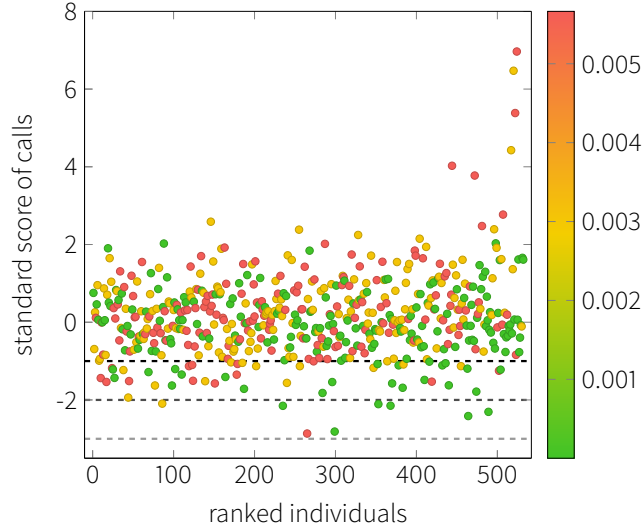
Figure 3: **Call activity deviations versus call counts** Standard score of monthly number of calls for each participant considered in the current study compared to their rank based on the actual number of calls in the index month (February, 2014). Participants are ranked according to the total duration of their calls in the index month. Color denotes the closeness centrality of the individuals in the aggregated network of calls.

S2.2 Transmission types

We model two different types of disease transmission: full-range and short-range infections, by modifying the structural properties of the underlying networks as introduced in the previous section. In the full-range transmission (up to 10-15 meters) we loosely approximate spreading of airborne diseases (e.g. measles) [4]. Short-range transmission (approximately up to 1 meter) corresponds roughly to droplet diseases (e.g. influenza). While the proximity networks do not correspond to disease transmission exactly, they have been considered in the literature a telling approximation of the potential spreading paths [5, 7, 8, 4].

S2.3 Parameter selection

The applied epidemiological model has two parameters: the probability of infection β and the probability of recovery γ. Given the temporal resolution ($\Delta t = 1/288$day) and temporal average degree ($\langle k \rangle = \frac{1}{\mathcal{N}(t)} \int_{t_i}^{t_i + \Delta t} \langle k \rangle(t') \mathrm{d}t'$, $\mathcal{N}(t)$ denoting the number of time bins), the simulation probabilities correspond to real physical rates in the following way. The rate of infection is given by

$$\beta_{\mathrm{physical}} = \frac{\langle k \rangle \beta}{\Delta t}, \tag{1}$$

whereas the expected infectious period is

$$T_{\mathrm{inf}} = \frac{\Delta t}{\gamma}. \tag{2}$$

Based on the physical rates the basic reproduction number is $R_0 = \beta_{\mathrm{physical}} T_{\mathrm{inf}}$.

Before investigating epidemic surveillance and immunization, we have to set infection probability and infectious time by bearing two aspects in mind: first, we intend to have an infectious time in the order of magnitude of days and second, to measure changes in outbreak sizes accurately we aim for an outbreak that lies near the boundary of the epidemic threshold: a disease that does not extinct but it does not result in the full infection of the population either. Figure 4 shows the
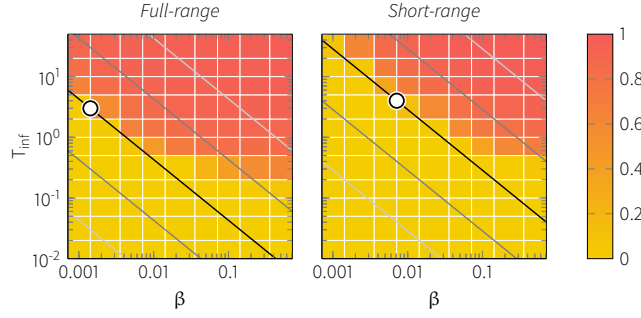
Figure 4: **Outbreak size as the function of the various parameters.** Absolute size of the outbreak (total number of infections normalized by the population size) at different values of the contact rate (β) and average infectious period (T_{inf}). White circles show the position of the actual values that were used for the two transmission types. Straight lines correspond to the derived phyical R_0 values. Each tile corresponds to the median outbreak size of 100 SIR simulations.

Network	$\langle k \rangle$	β	γ	β_{physical}	T_{inf}	R_0
Full-range	1.245	0.002	$11.57 \cdot 10^{-4}$	$0.717 \, \text{day}^-1$	$3 \, \text{day}$	2.151
Short-range	0.200	0.01	$8.68 \cdot 10^{-4}$	$0.591 \, \text{day}^-1$	$4 \, \text{day}$	2.364

Table 2: Simulation parameters and the corresponding physical quantities: average temporal degree ($\langle k \rangle$), probability of infection for a single infected-susceptible contact (β), physical rate of infection (β_{physical}), probability of recovery for a single time step (γ), expected infectious period (T_{inf}) and the basic reproduction number (R_0).

relative outbreak size (i_{rel}, the total number of infected divided by the population size) at different values of the parameters during airborne and droplet diseases.

Note that parameter values were probed on a logarithmic scale (we tested parameter values of 10^n, $2 \cdot 10^n$ and $5 \cdot 10^n$ for various exponents n) and therefore the size of the tiles in the figure is not uniform. Based on the outbreak size landscape, the following parameters were chosen for the airborne and droplet networks (marked by the white circles): $\beta^{\text{full}} = 0.002$, $T_{\text{inf}}^{\text{full}} = 3 \, \text{days}$ and $\beta^{\text{short}} = 0.01$, $T_{\text{inf}}^{\text{short}} = 4 \, \text{days}$, resulting in typical outbreak sizes of $i_{\text{rel}}^{\text{full}} = 0.727$ and $i_{\text{rel}}^{\text{short}} = 0.634$. Table 2 summarizes the simulation parameters and related physical rates for the two proximity networks. The average degree is calculated for a single time step of δ_t, similarly, the rate parameters β and γ are defined for a single step of the simulation. However, as physical parameters of the epidemic are independent of the temporal resolution, they are corrected for the time window and therefore are independent of δ_t. In other words, the physical parameters β_{physical} and T_{inf} correspond to the rate parameters used in the differential equations in the compartment model. Finally, it is also noteworthy that setting R_0 to the same value in the two experiments does not necessarily result in the same outbreak sizes, contrary to the compartment (differential equation) model. The contour lines of constant R_0 do not follow the constant outbreak size contours (by comparing the lines to the color). This is due to the inherent difference between compartment models and network simulations. We conclude that having simulations with the same R_0 values would not provide comparable statistics and the comparison of outbreaks in the two networks cannot simplified significantly by standardizing the physical epidemic parameters.

As we are interested only in the actual outbreak size (instead of the probability of outbreak or extinction), in all cases we accept only those simulations with a relative outbreak size of at least 5%, which results in a corresponding saturation value of the relative outbreak size in our vaccination curves. As for the monitoring, to obtain a larger ensemble for the statistics, the threshold of relative outbreak size is 20%, however, the results are robust against the choice of threshold.

S2.4 Extinction

While we acknowledge the fundamental interest in investigating the peculiar circumstances surrounding extinctions, we did not include the simulations that went extinct intentionally. This is because these rare (and not systematic) events would result in a bias of the calculated relative outbreak size that we cannot correct for when comparing relative outbreak sizes. Furthermore, our preliminary results with the proximity networks revealed that a significant rate of extinction only appears at very high levels of vaccinations (for the specific R_0 values), due to the high network density (many and frequent contacts). Therefore in practice we are only removing a very small fraction of runs.

S2.5 Epidemic dynamics

Here we provide a quick overview of the actual dynamics of the epidemics and how it is realized in the various monitoring and vaccination scenarios. For the two cases of preventive strategies, we plot the fraction of infected in Fig. 5 in time. During epidemic surveillance, we consider a non-vaccinated scenario and measure the corresponding statistics inside the target groups of 30 individuals, which is shown in Fig. 5a. Here we show the mean over 1 000 realizations along with the standard deviation as the median of the normalized curves by a low number (30 individuals) would result in a step-function shape. As the curves show, for both transmission types, the mean of the fraction of infected peaks at an earlier stage in social target groups compared to the population average. Furthermore, in case of the short-range network, the group dynamics is almost identical to that of the optimal group. The cumulative curves also indicate these differences (Fig. 5a inset). The significantly large error is the result of averaging over realizations with different starting times and is a fingerprint of the non-Gaussian samples at each time. Due to the impact of the circadian rhythm on the momentary fraction of infected, each data point is averaged over a sample with large deviations. Therefore, the infection dynamics shown in the figure are rather illustrative and in the analyses provided in the manuscript and the rest of the Supplementary Information are based on other more accurate statistics.

In the presence of vaccination (Fig. 5b), we immunize those in the target group and plot the fraction of infected in the total population. Curves represent the median over 1 000 simulations in each case, error bands show the lower and upper quartiles. For the sake of comparison, we show the dynamics of a target group of 50 individuals for the short-range interactions and a group of 120 students for the full-range network, corresponding to 0.96% and 22.6% of the population). Although social vaccination displays pronounced effects on the median fraction of infected compared to the unvaccinated and the random vaccination cases, the slow-down in the dynamics and the the shrink of the peak is more emphasized during short-range interactions, also approaching the results of immunizing the optimal target group more. Again, the above effect is more visible in the integrated number of infected as shown in the inset of Fig. 5b.

S3 Monitoring

In epidemic surveillance candidates in effective target groups are preferred to be in the high-risk subset of the population and preferably are infected early in the outbreak. This way, status of target individuals can be used to notice epidemic outbreak before it infects a high fraction of the population and may help in forecasting outbreak statistics. Therefore, the good target group is small compared to the population (so that monitoring is feasible and requires a lower level of resource allocation) and has high infection probability and low infection time. Here we consider various target group sizes and, in order to obtain a large number of samples for the statistics, we only consider outbreaks above 20% of relative outbreak size. In each case, we measure two quantities for target individuals related to epidemic surveillance. First we calculate the probability of infection—the fraction of simulations in which an individual becomes infected, provided that they are not part of the index group:

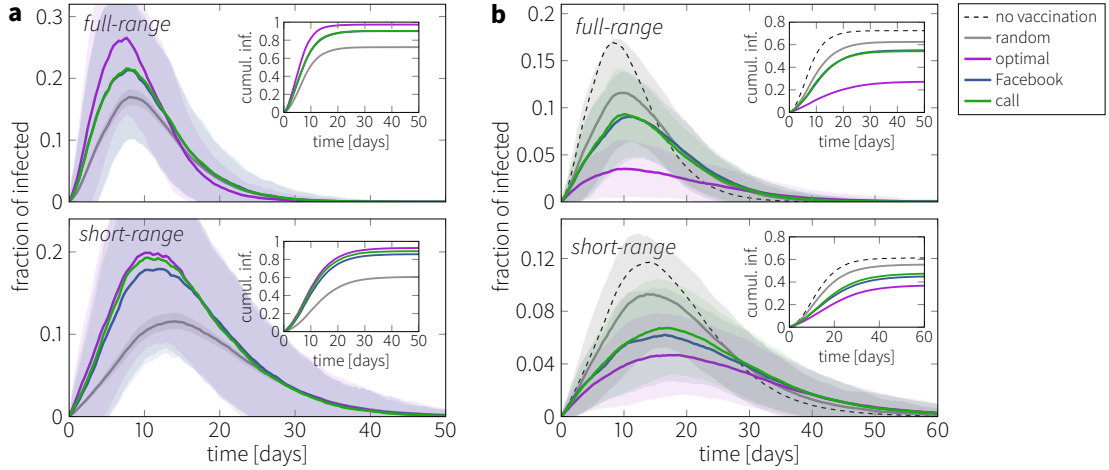$$\mathcal{P}_{\text{inf}}(i) = \frac{N_{\text{inf}}(i)}{N(i)},\tag{3}$$

Figure 5: **Typical dynamics of the outbreak. a)** Fraction of infected of the disease inside the monitor group compared to random target groups (that shows the same dynamics as the population average). For each curve, monitor groups include the 30 individuals with the highest rank according to the given strategy. Curves show the mean over 1 000 realizations, error bands mark the standard deviation. **b)** Fraction of infected measured in the total population during vaccination. Each curve represents the dynamics of the infection when the corresponding target groups are immunized, compared to the unvaccinated case and also to randomly selected target groups. Size of the target groups is 120 for full-range and 50 for short-range transmission. In both panels, insets show the cumulative fraction of infected. Curves indicate the median over 1 000 simulations, error bands display the standard deviation.

where $N(i)$ is the number of simulations in which individual i was not an index case and $N_{\mathrm{inf}}(i)$ is the number of simulations in which individual i was infected during the outbreak. We also measure a corrected infection time that accounts for the level of exposure of an individual to the disease—the average time of their infection, divided by the probability of infection:

$$\tau_{\mathrm{inf}}(i) = \frac{t_{\mathrm{inf}}(i)}{\mathcal{P}_{\mathrm{inf}}(i)}. \tag{4}$$

where $t_{\mathrm{inf}}(i)$ is the average time of getting the infection for individual i. In case of individuals who are never infected, the total duration of the outbreak is considered.

S3.1 Social monitoring

Figure 6 shows the infection probability and infection time for the four types of target groups during full- and short-range infection outbreaks. In both transmission types, social monitor groups outperform random groups in the probability of infection, with performance close to optimal (colocation) groups, indicating that socially central individuals are indeed members of the high-risk subpopulation, irrespective of the infection range. Infection times in these groups are significantly lower compared to the population average, displaying temporal gain of 1.5 - 4 days, translating to leading times of 15-25% compared to the population mean, depending on the transmission type and group size.

Social networks contain sufficient information about the population to effectively monitor the population in case of an epidemic, and various layers of social interaction (calls, text messages, Facebook events) show similar performance in inferring good candidates for monitoring. Furthermore, as long as the selection of target individuals is based on a centrality that grasps the global structure of the corresponding networks, the details of the method is less relevant. These observations are summarized in Fig. 7, where we compare various centralities as well as social network channels. Comparison of different centralities considered in the call-based network results in similar
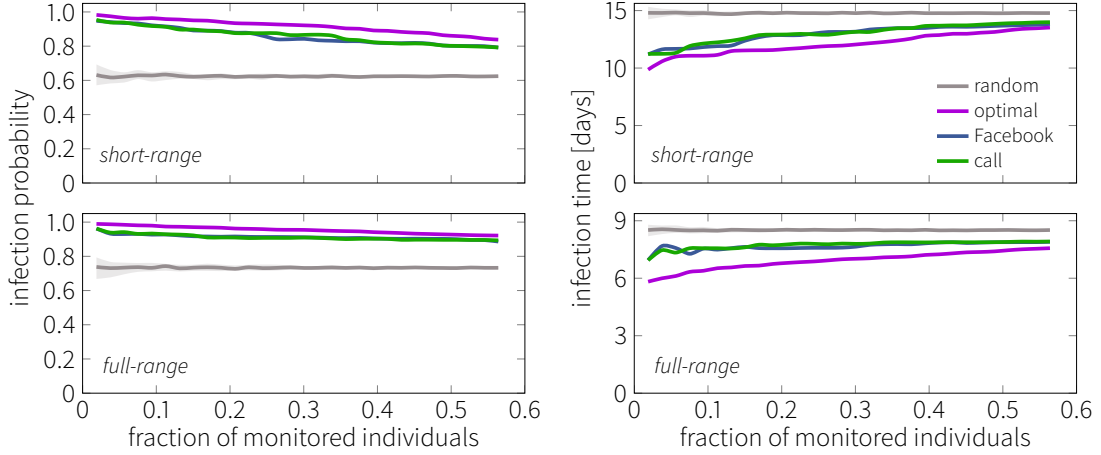
Figure 6: **Monitoring performance of social target groups** Median infection probability (left) and median infection time (right) of social target groups of different size, compared to the population average (random, grey) and optimal (purple) groups. Each curve is the result of 10^4 simulations with a minimum relative outbreak size of 20% in an unvaccinated population. Curves corresponding to the random groups show the median over 100 random realizations, error bars denote lower and upper quartiles of the samples.

performance (with the exception of the raw count-based ranking), and we also notice that closeness centrality shows the best performance, although only marginally better than other strategies.

In addition to various centrality measures, we also consider social sensors strategy [9, 10, 11]. We simulate friendship based target groups from the phone network (calls and text messages) in the following way. First, we consider the three most frequent contacts for each individual and denote them as *friends*. Based on these top-contact lists, we rank the participants according to the number of times they appear on the friends list. As Fig. 7b shows, the resulting target groups have similar performance as the centrality-based ones. Finally, if additional or other channels as text messages or structural Facebook connections are considered, the monitoring performance does not change significantly (Fig. 7c-d).

S3.2   Effect of network structure

Since monitoring does not alter the structure of the network but probes the outbreaks dynamics using small number of individuals, we can simplify the question of finding good candidates to possible correlations between network properties and monitoring efficiency, e.g., infection probability. We remove various correlations from the proximity networks, and plot infection probability versus the weight of individuals in the network. Results are shown in Fig. 8 for both full- and short-range proximity networks. We apply three types of correlation removal:

Temporal We remove all circadian rhythm effects and re-distribute links homogeneously in time, keeping both aggregated degree-distribution and degree-correlations. We simply re-assign time stamps so that their distribution became uniform in time. Individual weights are kept fixed.

Spatial All degree-correlations are removed, but both circadian rhythm and the static degree-distributions are kept fixed. This corresponds to a redistribution of temporal links according to the degree-distributions and temporal edge density. Individual weights remain the same.

Spatio-temporal Both temporal and spatial correlations are removed, together with the aggregated degree-distribution. Only the individual weights are conserved from the original structure of the networks.
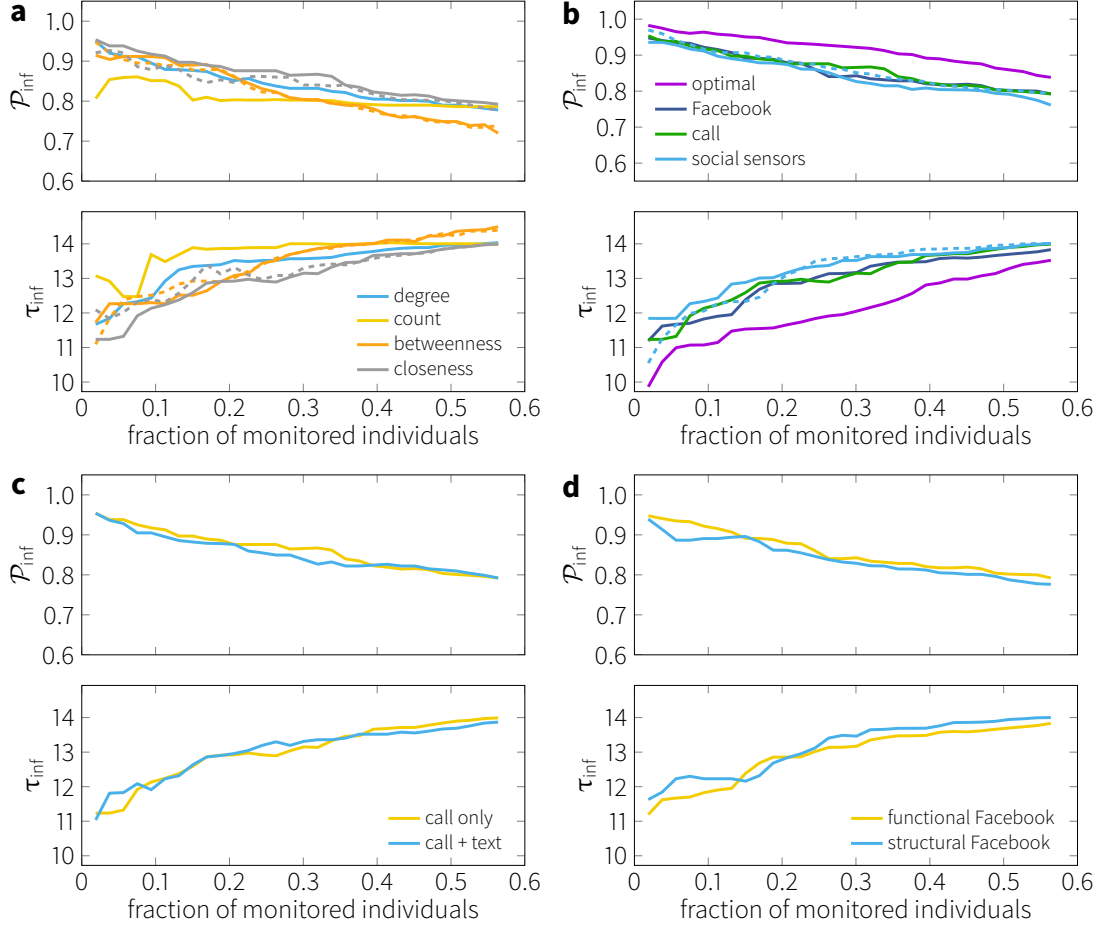
Figure 7: **Different monitoring strategies.** Monitoring power of target groups selected by various methods and using different channels: (a) testing centrality measures based on the call network, dashed lines denote weighted measures; (b) social sensors estimated from the call + text messages network as well as from the short-range proximity network, dashed lines correspond to social sensors estimated from the short-range proximity network; (c) networks constructed from calls only compared to calles + text messages; (d) functional Facebook network (based on user activity) compared to structural Facebook network (that is, the sole structure of friendships). All lines are the result of 1000 simulations with a relative outbreak size larger than 20%.
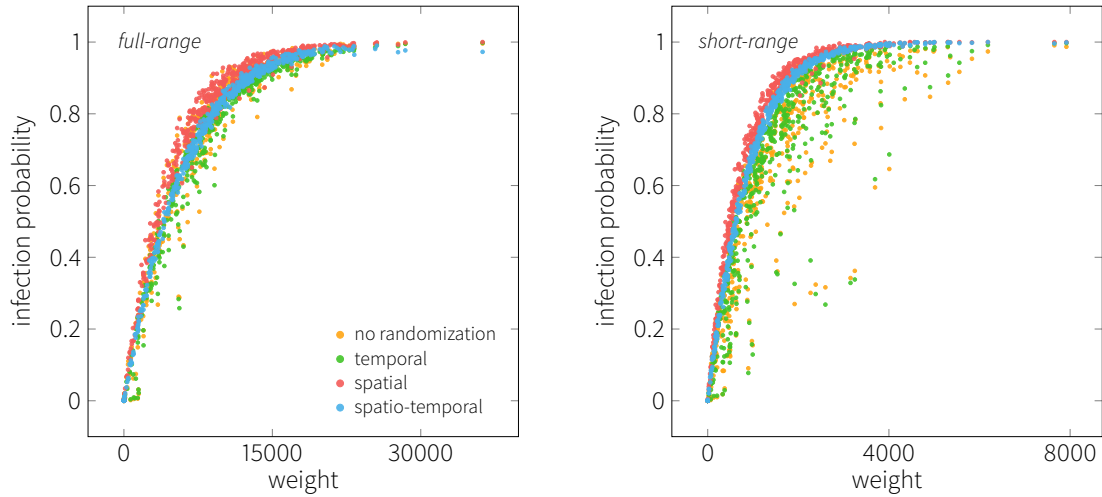
Figure 8: **Structural effects on monitoring performance** Infection probability of individuals versus their corresponding weight in the proximity networks after several correlations are removed: the original networks (yellow), after temporal correlations are removed (green), after structural correlations are removed (red) and when all correlations are eliminated (blue). In all cases, the personal weights are kept fixed as it forms the basis of the curves shown.

In both short- and full-range networks, removing spatio-temporal correlations also removes the vast majority of noise from the structure, and results in a clear trend of infection probability as a function of weight, saturating at some point (Figure 8). The higher level of initial noise in the original networks seen in the short-range case indicates that structure plays a major role in the strength of weight-based target groups. Spatial randomization also removes more noise from the trends, suggesting that the temporal nature of the networks has a smaller effect during epidemic surveillance. These observations point our that weight is the most relevant property when the susceptibility of individuals is considered. This is supported by the fact that highly central individuals in social networks are more likely to have high weight, we find the correlations between social network centralities and physical proximity networks to be approximately $r_{\mathrm{Spearman}} = 0.4$. However, as we have seen, the structural correlation between social networks and proximity networks does not prove to be sufficient for efficient vaccination, as vaccination corresponds to a fundamentally different challenge, by effectively removing individuals from the network, thus changing the underlying network structure.

S4   Vaccination

S4.1   Criteria for assessing vaccination efficiency

The main finding in this paper is that vaccination strategies based on communication networks (which are fundamentally different from the person-to-person networks on which the disease is transmitted) have comparable performance in case of short-range transmission, whereas vaccination based on communication networks proves to be less efficient in the full-range transmission case.

Because the inherent efficiency of vaccination differs in the two networks (short/full-range), our focus is not on the performance of vaccination, but how different strategies compare to each other. For this reason, our argument is based on comparing the typical relative outbreak size as a function of the vaccinated population (as shown in Fig 2B of the main text), rather than looking at absolute numbers.

S4.2 Characteristic differences

Vaccination efficacy was shown to be qualitatively different in the two types of transmission modeled by the short-range and full-range contact networks. In the manuscript we compared the median along with the SEM of the measured relative outbreak sizes, and Fig. 9 depicts the statistical comparison of the curves at the level of individual simulations. In the figure, we compare the distribution of relative outbreak sizes to the optimal relative outbreak sizes at each vaccination level using the Mann–Whitney statistics. On the vertical axis, we show the derived z-index calculated as

$$z = \frac{U - \frac{n_1 n_2}{2}}{\sqrt{\frac{n_1 n_2 (n_1 + n_2 + 1)}{12}}}, \tag{5}$$

where n_1 and n_2 denote the size of the two populations (number of simulations for the two strategies at the same vaccination level), and U is the Mann–Whitney statistics. We chose this statistics based on the observation that the distributions (shown in Fig. 2 in the main text) do not follow Gaussian distribution. Clearly, cyber-based vaccination shows a significantly lower statistics compared to random in the short-range network, suggesting that the separation observed in the median is an overall behavior of the outbreak sizes over a wide range of vaccination levels. However, in case of full-range transmission, we observe no statistical difference between the cyber-based vaccination and random strategy.
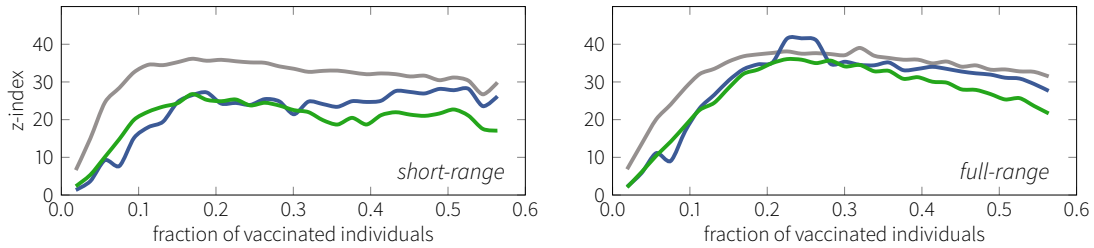

Figure 9: **Mann–Whitney test z-indices.** Curves show the calculated z-index from Mann–Whitney tests with a null hypothesis that the relative outbreak size has the same distribution for the optimal and another strategy.

Here we explore further epidemiologically relevant measures that describe the efficacy of vaccination. Vaccination has a decreasing effect on the outbreak duration (number of days until the disease dies out) in case of both transmission types (Fig. 10), however, in the full-range network, social vaccination does not display observable difference compared to the random strategy. Notably, as shown by the error bands, differences are blurred compared to the actual size of the outbreak, even between optimal and random vaccination.

Metrics related to the peak of the outbreak are depicted in Figs. 11-12. The peak of the outbreak is defined by the highest momentary fraction of infected throughout the course of the infection. The same trends can be observed as in the relative outbreak size: as the level of vaccination increases, curves corresponding to the cyber-based vaccination are separated from the random strategy in the short-range network, contrary to the case of full-range interactions. Similarly, peak times are more noisy and differences are not significant in this respect.

Finally, we calculated the estiated attack-rate of the disease for various vaccination levels. The attack rate is defined by the fraction of new infections relative to the susceptible population, and we estimate this rate by calculating the average rate during the first 5 days of the outbreak. Figure 13 shows the results. Differences are less pronounced as this metric is prone to noise in our simulations. However, note that results are robust against the number of days attack rate is approximated for. Using Facebook activity as proxy, the initial attack rate in the short-range network shows a clear separation from the random strategy, but it saturates at a level similar to the random strategy. Calls are capable of significantly decreasing the attack rate in both transmission types.

The above metrics further support the main message that communication networks prove to

be a good candidates for the basis of efficient vaccinations provided that the disease spreads by short-range interactions.
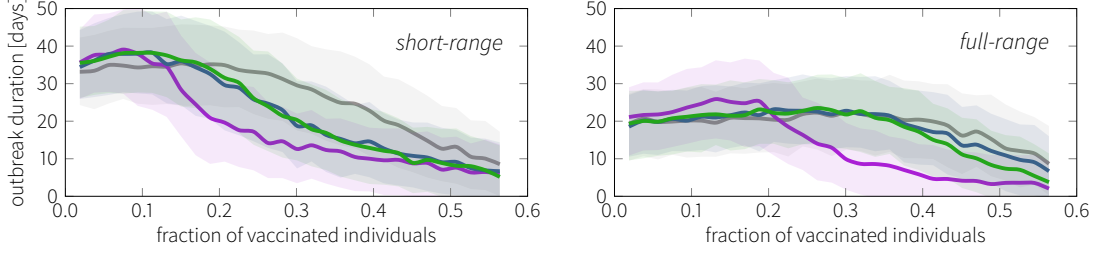

Figure 10: **Outbreak duration in days.** Curves show the median duration of the outbreaks at different vaccination levels for the two networks: short-range (left) and full-range (right).
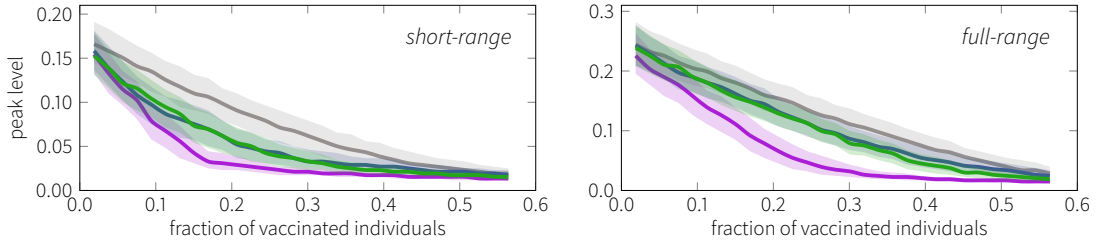

Figure 11: **Peak levels during the outbreaks.** Curves show the median peak level (highest fraction of infected during the spread) at different vaccination levels for the two networks: short-range (left) and full-range (right).

S4.3   Social vaccination

In the manuscript we have shown that social vaccination is effective in short-range transmission, here we explore robustness of the observed performance against various definition of 'central individuals'. Several strategies exploiting the structure of social networks have been proposed for efficient immunization [12, 13, 11]. Figure 14a shows the performance of social vaccination (in short-range transmission) based on different strategies for determining target individuals. The network measures in the population are highly correlated and lead to qualitatively similar results, independent of whether degree, k-cores, or betweenness centrality are considered. We also examine different configurations of the social network channels used for social vaccination, adding text message metadata to call metadata and using the Facebook friendship graph. We find no significant difference in the efficacy by using these different views, indicating that the epidemiologically relevant individuals can be robustly identified based on the digital social networks.

To overcome the difficulty of collecting a complete network of interactions for targeted interventions, a scheme of social sensors (acquaintance immunization) has been proposed [9, 10, 11]. By vaccinating friends of randomly selected individuals—as such friends tend to be more central in the social network—strictly *local* information can be used (majority of the above measures rely on global network information). Here we evaluate the performance of social sensors scheme directly, choosing three strongest contacts from the call and SMS lists of each individual as a proxy for naming someone a friend[1]. Individuals nominated as friends most often were then included in the target groups. We show that social sensors are a useful strategy for selecting individuals using strictly local information in the network, with the performance close to using full network structure (Fig. 14a lower right panel).
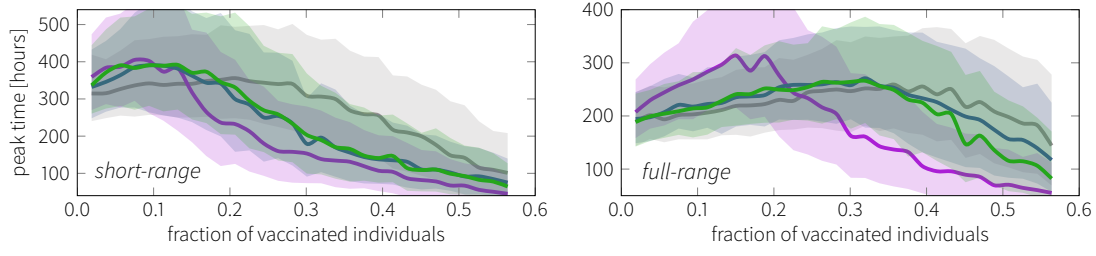
Figure 12: **Peak times during the outbreaks.** Curves show the median peak time (time to arrive at the highest level of infection during the outbreak) in hours at different vaccination levels for the two networks: short-range (left) and full-range (right).
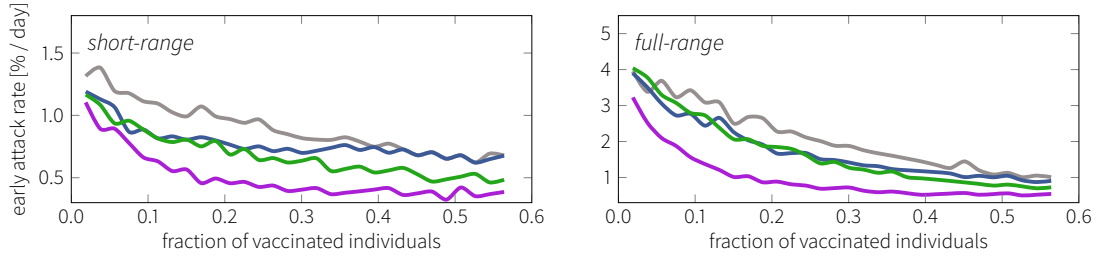

Figure 13: **Attack rate during the first 5 days.** Curves show the median attack rate (fraction of new infections divided by the number of susceptible) in percentages of the whole population at different vaccination levels for the two networks: short-range (left) and full-range (right).

S4.4 Models of social target groups

To understand the performance of social vaccination, here we provide a qualitative explanation of the observed vaccination curves. The performance of the social vaccination in the short- and full-range transmission networks can we be decomposed into the result of two basic mechanisms. For a given vaccination size, social networks manage to identify certain fraction of individuals from the optimal (colocation) group, the *the core*. Those target individuals not included in the colocation group (the *the periphery*) also contribute to the performance and in the case of an effective target group these individuals also display increased relevance during immunization. To understand the impact of these two components, we consider three modified models of target groups:

Model A We keep the core fixed and replace the periphery with individuals sampled from outside the optimal group randomly.

Model B We keep the periphery fixed and replace the core by randomly chosen individuals from the optimal group.

Model C We keep the periphery fixed and replace the core with the same number of most connected participants from the optimal group.

Results are shown in Figure 15. The decrease of performance in model A indicates that social networks are able to locate more appropriate candidates (compared to random) outside of the optimal group in addition to the optimal ones (Fig. 15a). This observation is more pronounced in the short-range network (Fig. 15b). In case of the core however, social networks highlight optimal individuals with the same efficiency as one would obtain by sampling the optimal group randomly, as indicated by negligible change in performance in model B. Finally, the performance change of model C compared to the original strategy is higher in the full-range network than that in the
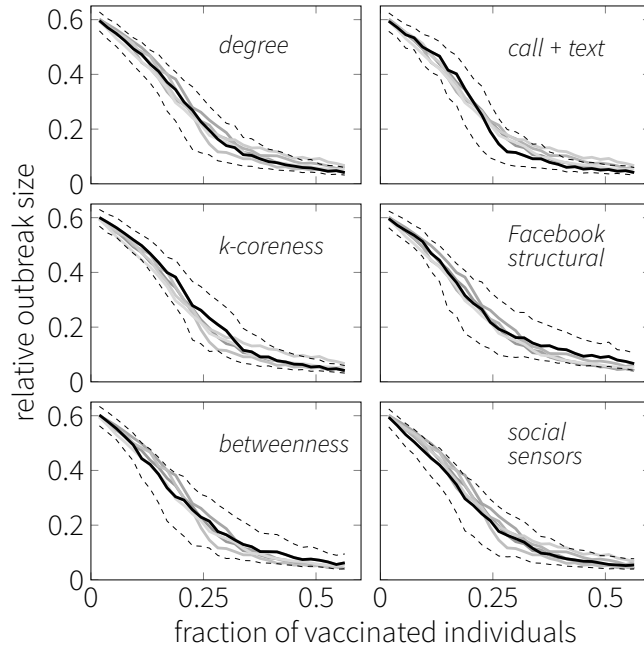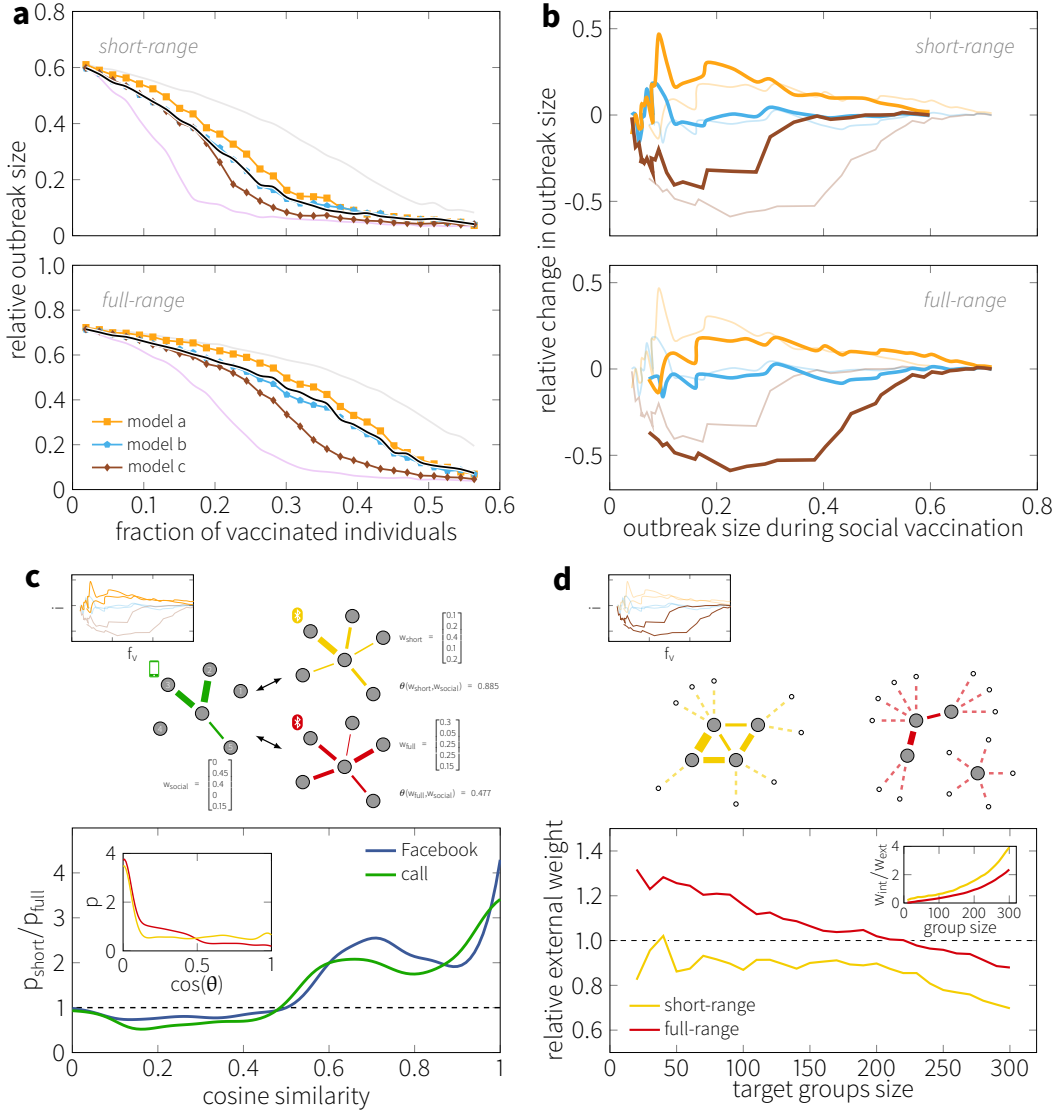
Figure 14: **Effect of various centrality measures and communication channels considered.** Panels show the vaccination performance of target groups based on different centralities and layers, with a single centrality or channel being highlighted in each panel. In case of different centralities (left panels), the call network was used as the basis for the analysis. For the different channels (right panels, with the exception of the bottom panel), we selected individuals based on the closeness centrality. Bottom right panel shows the results with the social sensors, the simulation of Christakis et al. social sensor selection method based on the call + text message data among students.

Figure 15: **Role of target individuals in social target groups. a)** Median relative outbreak size in three models (see text for their definitions) compared to random (light gray), optimal (light purple) and the call network (light green). **b)** Relative differences in the outbreak sizes of the original and the modified social groups as the function of the outbreak size in the optimal vaccination, the curves of the corresponding proximity network are highlighted. **c)** Top: illustration of the cosine similarity. For each individual, we define their weighted degree as the total fraction of interactions with all other participants (weighted degrees are normalized to have unit component sum). For a given proximity network, the cosine similarity between the weighted degrees of the same individual is calculated, and the distribution of similarities is constructed. Bottom: point-wise ratio of the probabilities of similarity values in short-range and full-range networks, when compared to the call (green) and Facebook (blue) graphs. Inset shows the original similarity distributions for the short-range (yellow) and full-range (red) graphs with the call network. **d)** Weight removed from the remainer of the graph after the removal of the vaccinated individuals. Top: illustration of the weights distributed among the removed (gray) and remaining (white) individuals. Solid lines correspond to internal, dashed denote external weights being removed from the network. Bottom: total external weight removed from the graphs when the overlap of the social and optimal target groups are replaced by the highest-weight (solid) or randomly chosen (dashed) individuals from the optimal group. Inset shows the ratio of internal / external weights in the two networks for different target groups.

17

short-range, indicating that this effect is a major driver (together with difference uncovered by model A) behind the superior performance of the social vaccination in the short-range network.

The difference in the performance due to the contribution of the periphery can be explained by the structural differences between the proximity networks. Calculating cosine similarity of the participants' weighted degree v (the fraction of interactions they have with the rest of the population) between digital social and physical proximity networks for all participants $\theta(v_{\text{digital}}, v_{\text{physical}})$, we show that short-range network displays higher frequency of high similarity values ($\theta > 0.5$) (Fig. 15c). Digital social networks capture the local structure of the short-range network to a higher extent than that of the full-range network, which includes a high fraction of incidental interactions[4]. As a consequence, social networks are able to locate epidemiologically relevant target individuals outside of the optimal group, as central participants in the social networks are also exhibit high centrality in the short-range network.

On the other hand, performance change with respect to the core is related to whether the contacts of the vaccinated nodes are *internal* or *external*. Vaccinating individuals effectively removes them from the network along with all of their contacts. Internal links are connecting target individuals, and so their interactions do not contribute to the improvement in the vaccination. Thus, even though the total weight of the removed links may be equal in short- and full-range networks, the actual impact of the removed links can be different if the distribution of internal/external links differ. Interactions in the short-range network tend to take place between participants socially more related[4], and we observe that target groups in that network also tend to be more interconnected (Fig. 15d). Thus choosing the individuals with the highest level of interactions (weight) in the short-range network does not lead to pronounced increment in performance, as we remove individuals from a more interconnected subgraph (compared to the full-range network).

S4.5 Epidemic spread on the social networks

The main aim of the work presented in the manuscript is to find individuals based on their social network (which is approximated from their digital fingerprint) while probing their impact in case of an epidemic outbreak that takes place on the physical proximity networks. Our results show that social and proximity networks are structured in a way that allows for a decreased performance for vaccination, however, the efficacy of state-of-the-art strategies proposed in the literature is related to the fact that both finding target individuals and evaluating their impact was performed on the same network. Figure 16b shows how we would understand the performance of targeted vaccination if both the selection and the spread of took place on the social networks. As we can see, if the disease spreads in the same network that is the basis for the target group selection, degree immunization (that is, one of the most efficient strategy) performs well in both dynamic and static cases, and significant performance change is observed compared to the, e.g., closeness centrality based strategy.

S5 Robustness

So far we have investigated the performance of the target group in the same period that formed the basis for their selection, however, in case of real epidemic scenarios it is preferred to collect data from a preceding periods so that we can infer to upcoming months. In this spirit, here we use the target groups of February (in the index month) for monitoring and vaccination in the consecutive months March, April and May. As an evaluation, we compare social target groups to the colocation based groups from both the index and later months.

Figure 17 shows the infection probability and infection time observed in the social target groups for different months, using the short-range proximity network. In the top panel, we show the infection probability relative to the probability measures in the optimal target groups based on the colocation data in the actual month. In all months, social target groups outperform random monitoring, and also show similar probability values to the optimal that can be obtained using proximity data of the same period. However, as the period in interest becomes further in time from the index month, efficacy becomes less stable and it decreases as well. Regarding infection time, we show the raw results in units of days compared to the optimal of the month. Again, social groups display significantly earlier times of infection than that of random groups but they are also
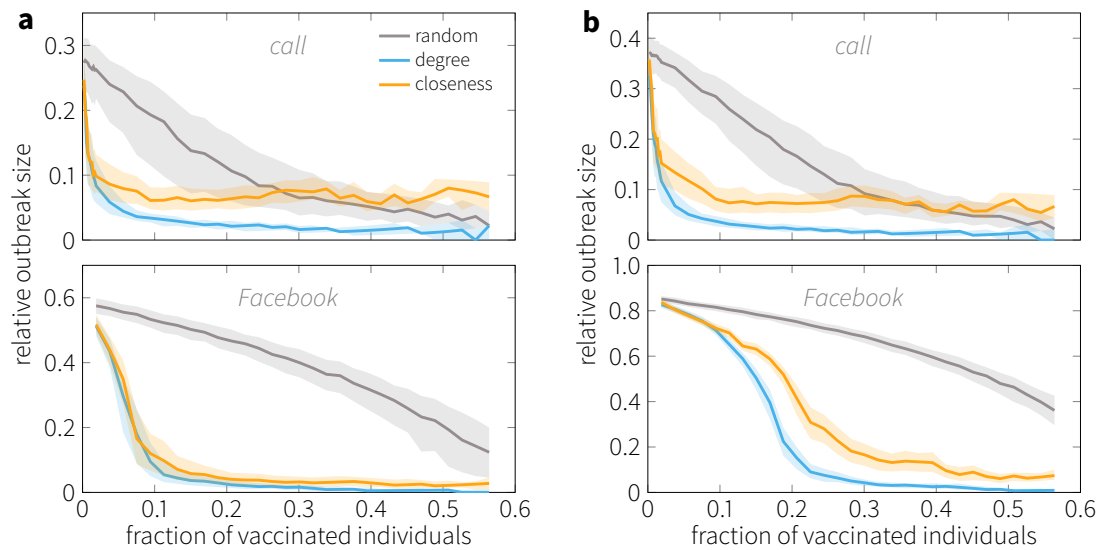
Figure 16: **Vaccination performance of outbreaks taking place on the social networks.** Efficient immunization strategies against epidemic that spreads via the ties in the digital social networks. In this case, both the target individuals are selected by and the epidemic takes place on the same social network. **a)** The network dynamics is also considered as they evolve in time, **b)** the static aggregated networks are used and the epidemic spreads rapidly on the static links.

less effective as the true optimal is. Interestingly, the decreasing performance trend observed in case of infection probability is not present here and social target groups perform equally well in all months.

References

[1] Stopczynski, A. *et al.* Measuring large-scale social networks with high resolution. *PLoS ONE* **9**, e95978 (2014).

[2] Stopczynski, A., Sapiezynski, P. & Lehmann, S. Temporal fidelity in dynamic social networks. *The European Physical Journal B* **88** (2015).

[3] Sekara, V. & Lehmann, S. The strength of friendship ties in proximity sensor data. *PLoS ONE* **9**, e100915 (2014).

[4] Stopczynski, A., Pentland, A. S. & Lehmann, S. Physical proximity and spreading in dynamic social networks. *arXiv preprint arXiv:1509.06530* (2015).

[5] Salathé, M. *et al.* A high-resolution human contact network for infectious disease transmission. *Proceedings of the National Academy of Sciences* **107**, 22020–22025 (2010).

[6] Smieszek, T. & Salathé, M. A low-cost method to asses the epidemilogical importance of individuals in controlling infecitous disease outbreaks. *BMC Medicine* **11**, 1–8 (2013).

[7] Stehlé, J. *et al.* High-resolution measurements of face-to-face contact patterns in a primary school. *PLoS ONE* **6**, e23176 (2011).

[8] Isella, L. *et al.* Close encounters in a pediatric ward: measuring face-to-face proximity and mixing patterns with wearable sensors. *PloS one* **6**, e17144 (2011).

[9] Christakis, N. A. & Fowler, J. H. Social network sensors for early detection of contagious outbreaks. *PLoS ONE* **5**, e12948 (2010).
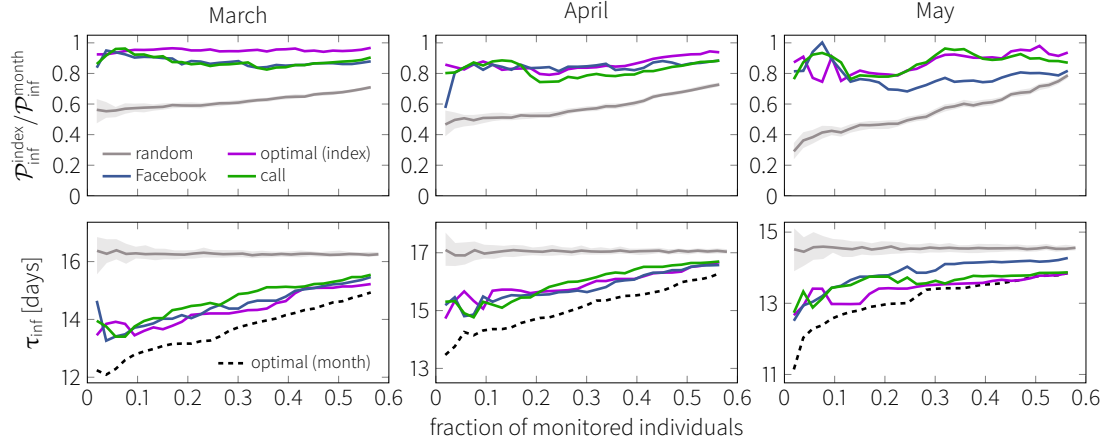
Figure 17: **Robustness of social target groups in external periods of time.** Top: infection probability relative to the current optimal. Bottom: infection time in target groups. Each curve shows the median over 10^4 simulations with a minimum relative outbreak size of 20%. Error bars on the random group results denote lower and upper quartiles of 100 realizations.

[10] Garcia-Herranz, M., Moro, E., Cebrian, M., Christakis, N. A. & Fowler, J. H. Using friends as sensors to detect global-scale contagiuos outbreaks. *PLoS ONE* **9**, e92413 (2014).

[11] Cohen, R., Havlin, S. & ben Avraham, D. Efficient immunization strategies for computer networks and populations. *Physical Review Letters* **91**, 247901 (2003).

[12] Kitsak, M. *et al.* Identification of influential spreaders in complex networks. *Nature* **6**, 888–893 (2010).

[13] Chen, Y., Paul, G., Havlin, S., Liljeros, F. & Stanley, H. E. Finding a better immunization strategy. *Physical Review Letters* **101**, 058701 (2008).