

Private algorithms for the protected in social network search

Michael Kearns^{a,1}, Aaron Roth^a, Zhiwei Steven Wu^a, and Grigory Yaroslavtsev^a

^aComputer and Information Science, University of Pennsylvania, Philadelphia, PA 19104

Edited by Salil Vadhan, Harvard University, Cambridge, MA, and accepted by the Editorial Board November 23, 2015 (received for review May 30, 2015)

Motivated by tensions between data privacy for individual citizens and societal priorities such as counterterrorism and the containment of infectious disease, we introduce a computational model that distinguishes between parties for whom privacy is explicitly protected, and those for whom it is not (the targeted subpopulation). The goal is the development of algorithms that can effectively identify and take action upon members of the targeted subpopulation in a way that minimally compromises the privacy of the protected, while simultaneously limiting the expense of distinguishing members of the two groups via costly mechanisms such as surveillance, background checks, or medical testing. Within this framework, we provide provably privacy-preserving algorithms for targeted search in social networks. These algorithms are natural variants of common graph search methods, and ensure privacy for the protected by the careful injection of noise in the prioritization of potential targets. We validate the utility of our algorithms with extensive computational experiments on two large-scale social network datasets.

data privacy | social networks | counterterrorism

The tension between the useful or essential gathering and analysis of data about citizens and the privacy rights of those citizens is at an historical peak. Perhaps the most striking and controversial recent example is the revelation that US intelligence agencies systemically engage in “bulk collection” of civilian “metadata” detailing telephonic and other types of communication and activities, with the alleged purpose of monitoring and thwarting terrorist activity (1). Other compelling examples abound, including in medicine (patient privacy vs. preventing epidemics), marketing (consumer privacy vs. targeted advertising), and many other domains.

Debates about (and models for) data privacy often have an “all or nothing” flavor: privacy guarantees are either provided to every member of a population, or else privacy is deemed to be a failure. This dichotomy is only appropriate if all members of the population have an equal right to, or demand for, privacy. Few would argue that actual terrorists should have such rights, which leads to difficult questions about the balance between protecting the rights of ordinary citizens, and using all available means to prevent terrorism. [A recent National Academies study (2) reached the conclusion that there are not (yet) technological alternatives to bulk collection and analysis of civilian metadata, in the sense that such data are essential in current counterterrorism practices.] A major question is whether and when individual privacy should be sacrificed in service of such societal priorities. Similarly, in the medical domain, epidemics (such as the recent international outbreak of Ebola; ref. 3) have raised serious debate about the clear public interest in controlling contagion versus the privacy rights of the infected and those that care for them.

The model and results in this paper represent a step toward explicit acknowledgments of such trade-offs, and algorithmic methods for their management. The scenarios sketched above can be broadly modeled by a population divided into two types. There is a protected subpopulation that enjoys (either by law, policy, or choice) certain privacy guarantees. For instance, in the

examples above, these protected individuals might be nonterrorists, or uninfected citizens (and perhaps informants and health care professionals). They are to be contrasted with the “unprotected” or targeted subpopulation, which does not share those privacy assurances. A key assumption of the model we will introduce is that the protected or targeted status of individual subjects is not known, but can be discovered by (possibly costly) measures, such as surveillance or background investigations (in the case of terrorism) or medical tests (in the case of disease). Our overarching goal is to allow parties such as intelligence or medical agencies to identify and take appropriate actions on the targeted subpopulation, while also providing privacy assurances for the protected individuals who are not the specific targets of such efforts, all while limiting the cost and extent of the background investigations needed.

As a concrete example of the issues we are concerned with, consider the problem of using social network data (for example, telephone calls, emails, and text messages between individuals) to search for candidate terrorists. One natural and broad approach would be to use common graph search methods: beginning from known terrorist “seed” vertices in the network, neighboring vertices are investigated, in an attempt to grow the known subnetwork of targets. This general practice is sometimes referred to as contact chaining (section 3.1 of ref. 2). A major concern is that such search methods will inevitably encounter protected citizens, and that even taking action against only discovered targeted individuals may compromise the privacy of the protected.

To rigorously study the trade-offs between privacy and societal interests discussed above, our work introduces a formal model for privacy of network data that provides provable assurances only to the protected subpopulation, and gives algorithms that allow effective investigation of the targeted population. These algorithms are deliberately “noisy” and are privacy-preserving versions of the widely used graph search methods mentioned

Significance

Motivated by tensions between data privacy for individual citizens, and societal priorities such as counterterrorism, we introduce a computational model that distinguishes between parties for whom privacy is explicitly protected, and those for whom it is not (the “targeted” subpopulation). Within this framework, we provide provably privacy-preserving algorithms for targeted search in social networks. We validate the utility of our algorithms with extensive computational experiments on two large-scale social network datasets.

Author contributions: M.K., A.R., Z.S.W., and G.Y. designed research, performed research, analyzed data, and wrote the paper. Ordering of authors is alphabetical.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. S.V. is a guest editor invited by the Editorial Board.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. Email: mkearns@cis.upenn.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1510612113/-DCSupplemental.

above, and as such represent only mild, but important, departures from commonly used approaches. At the highest level, one can think of our algorithms as outputting a list of confirmed targeted individuals discovered in the network, for whom any subsequent action (e.g., publication in a most-wanted list, further surveillance, or arrest in the case of terrorism; medical treatment or quarantine in the case of epidemics) will not compromise the privacy of the protected.

The key elements of our model include the following:

- i) Network data collected over a population of individuals and consisting of pairwise contacts (physical, social, electronic, financial, etc.). The contacts or links of each individual comprise the private data they desire to protect. We assume a third party (such as an intelligence agency or medical organization) has direct access to this network data, and would like to discover and act upon targeted individuals.
- ii) For each individual, an immutable “status bit” that determines their membership status in the targeted subpopulation (such as terrorism or infection). These status bits can be discovered by the third party, but only at some nontrivial cost (such as further surveillance or medical testing), and thus there is a budget limiting the number of status bits that an algorithm can reveal. One might assume or hope that in practice, this budget is sufficient to investigate a number of individuals that is of the order of the targeted subpopulation size (so they can all be discovered), but considerably less than that needed to investigate every member of the general population.
- iii) A mathematically rigorous notion of individual data privacy (based on the widely studied differential privacy; ref. 4) that provides guarantees of privacy for the network data of only the protected individuals, while allowing the discovery of targeted individuals. Informally, this notion guarantees that compared with a counterfactual world in which any protected individual arbitrarily changed any part of their data, or even removed themselves entirely from the computation, their risk (measured with respect to the probability of arbitrary events) has not substantially increased.

We emphasize two important points about our model. First, we assume that the process of investigating an individual to determine their status bit is unobservable, and leaks no information itself. This assumption is justified in some settings—for example, when the investigation involves secretly intercepting digital communications, like emails and phone calls, or when it involves performing tests on materials (like blood samples) or information already obtained. However, our model does not fit situations in which the investigations themselves are observable—for example, if the investigation requires interviewing an individual’s family, friends, and colleagues—because the very fact that an individual was chosen for an investigation (regardless of its outcome) might disclose their private data. The second point is that “privacy” is a word that has many meanings, and it is important to distinguish between the types of privacy that we aim to protect (see, for example, Solove’s taxonomy of privacy; ref. 5). Our goal is to quantify informational privacy—that is, how much information about a protected individual can be deduced from the output of an analysis. However, it is important to note that the status bit investigations our algorithms make, even if unobservable, are a privacy loss that Solove calls “intrusion.” Our results can be viewed as providing a quantitative trade-off between informational privacy and intrusion: using our algorithms, it is possible to guarantee more informational privacy at the cost of a higher degree of intrusion, and vice versa.

Our main results are:

- i) The introduction of a broad class of graph search algorithms designed to find and identify targeted individuals. This class of algorithms is based on a general notion of a statistic of proximity—a network-based measure of how “close” a given individual v is to a certain set of individuals S . For instance, one such closeness measure is the number of short paths in the network from v to members of S . Our (necessarily randomized) algorithms add noise to such statistics to prioritize which status bits to query (and thus how to spend the budget).
- ii) A theoretical result proving a quantitative privacy guarantee for this class of algorithms, where the level of privacy depends on a measure of the sensitivity of the statistic of proximity to small changes in the network.
- iii) Extensive computational experiments in which we demonstrate the effectiveness of our privacy-preserving algorithms on real social network data. These experiments demonstrate that in addition to the privacy guarantees, our algorithms are also useful, in the sense that they find almost as many members of the targeted subpopulation as their nonprivate counterparts. The experiments allow us to quantify the loss in effectiveness incurred by the gain in privacy.

Our formal framework is the first to our knowledge to introduce explicit protected and targeted subpopulations with qualitatively differing privacy rights. This is in contrast to the quantitative distinction proposed by Dwork and McSherry (6), which still does not allow for the explicit discovery of targeted individuals. We also note that our definition of privacy can be expressed in the Blowfish privacy framework (7) (although this had not previously been done). Our algorithms are the first to provide mathematically rigorous privacy guarantees for the protected while still allowing effective discovery of the targeted. More generally, we believe our work makes one of the first steps toward richer privacy models that acknowledge and manage the tensions between different levels of privacy guarantees to different subgroups.

Preliminaries

Consider a social network in which the individuals have an immutable status bit which specifies whether they are members of a targeted subpopulation or not (in which case we say they are part of the protected subpopulation). These status bits induce a partition on the population—we write \mathcal{T} to denote the targeted subpopulation and \mathcal{P} to denote protected subpopulation. Individuals correspond to the vertices V in the network, and the private data of each individual v is the set of edges incident to v . We assume that the value of an individual’s status bit is not easily observed, but can be discovered through (possibly costly) investigation. Our goal is to develop search algorithms to identify members of the targeted subpopulation, while preserving the privacy of the edge set of the protected population.

Any practical algorithm must operate under an investigation budget, which limits the number of status bits that are examined. Our goal is a total number of status bit examinations that is on the order of the size of the targeted subpopulation \mathcal{T} , which may be much smaller than the size of the protected population \mathcal{P} . This is the source of the tension we study—because the budget is limited, it is necessary to exploit the private edge set to guide our search (i.e., we cannot simply investigate the entire population), but we wish to do so in a way that does not reveal much about the edges incident to any specific protected individual.

The privacy guarantee we provide is a variant of differential privacy, an algorithmic definition of data privacy. It formalizes the requirement that arbitrary changes to a single individual’s private data should not significantly affect the output distribution of the data analysis procedure, and so guarantees that the analysis leaks little information about the private data of any single individual. We introduce the definition of differential privacy




Fig. 1. Visual comparison of the nonprivate algorithm Target (Left) and the private algorithm PTarget (Right) on a small portion of the IMDB network (see *Experimental Evaluation* for more details). For each algorithm, blue indicates protected vertices that have been examined, red indicates targets that have been examined, and gray vertices have not been examined yet. Both algorithms begin with the same seed target vertex, and by directed statistic-first search discover a subnetwork of targeted individuals (central red edges). As a consequence, many protected vertices are discovered and examined as well. Due to the added noise, PTarget explores the network in a more diffuse fashion, which in this case permits it to find an additional subnetwork of targets toward the right side of the network. The primary purpose of the noise, however, is for the privacy of protected vertices.

We next describe the nonprivate version of our search algorithm $\text{Target}(k, f)$. Our motivation in choosing this particular algorithm is simplicity: it is the most straightforward type of contact chaining algorithm that ignores privacy entirely, and simply uses the given SoP to prioritize investigations.

For any fixed SoP f , Target proceeds in k rounds, each corresponding to the identification of a new connected component in the subgraph induced by T . The algorithm must be started with a seed vertex—a preidentified member of the targeted population. Each round of the algorithm consists of two steps:

i) **Statistic-first search:** Given a seed targeted vertex, the algorithm iteratively grows a discovered component of targeted vertices, by examining, in order of their SoP values (computed with respect to the set S of individuals already identified as being members of the targeted population), the vertices that neighbor the previously discovered targeted vertices.

This continues until every neighbor of the discovered members of the targeted population has been examined, and all of them have been found to be members of the protected population. We note that this procedure discovers every member of the targeted population that is part of the same connected component as the seed vertex, in the subgraph induced by only the members of the targeted population.

ii) **Search for a new component:** Following the completion of statistic-first search, the algorithm must find a new vertex in the targeted population to serve as an initial vertex to begin a new round of statistic-first search. To do this, the algorithm computes the value of the SoP for all vertices whose status bit has not already been examined, using as the input set S the set of already discovered members of the targeted population. It then sorts all of the vertices in decreasing order of their SoP value, and begins examining their status bits in this order. The first vertex that is found to be a member of the

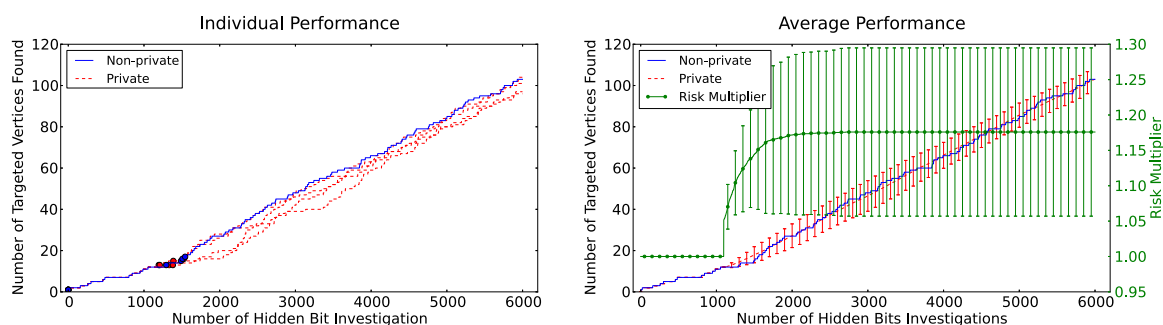


Fig. 2. Performance for the case in which there is a dominant component in the targeted subpopulation. In *Left*, we show the number of targeted vertices found as a function of the budget used for both the (deterministic) nonprivate algorithm Target (blue), and for several representative runs of the randomized private algorithm PTarget (red). Colored circles indicate points at which the corresponding algorithm has first discovered a new targeted component. In *Right*, we show average performance over 200 trials for the private algorithm with 1-SD error bars. We also show the private algorithm risk multiplier with error bars. In this regime, after a brief initial flurry of small component discovery, both algorithms find the dominant component, so the private performance closely tracks nonprivate, and the private algorithm's risk multiplier quickly levels off at around only 1.17.

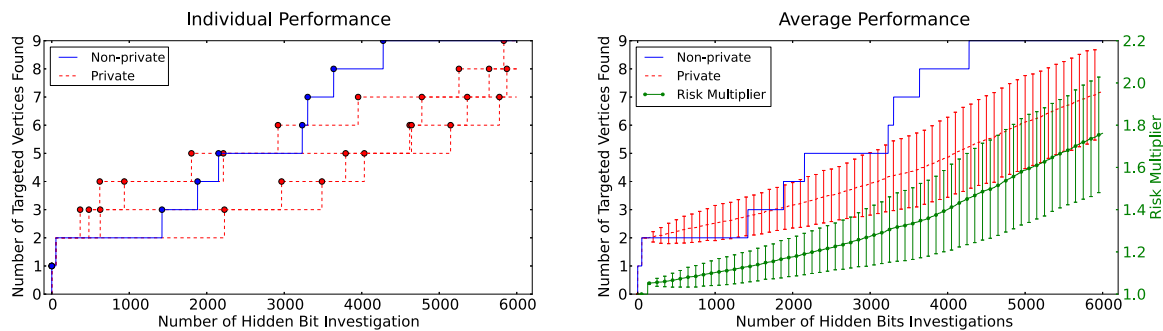


Fig. 4. A case with a highly fragmented targeted subpopulation. Both algorithms now spend most of their budget investigating nontargeted vertices and suffer accordingly.

sensitivity 1, and so can be used in our algorithm while adding only a small amount of noise. In particular, the private algorithm PTarget adds noise sampled from the Laplace distribution $\text{Lap}(20)$ to the SoP when performing new component search. By Theorem 1, such an instantiation of PTarget guarantees $((k-1)/20)$ -protected differential privacy if it finds k targeted components.

The main trade-off we explore is the number of members of the targeted population that are discovered by the algorithms (the y axis in the ensuing plots), as a function of the budget, or number of status bits that have been investigated so far (the x axis in the ensuing plots). In each plot, the parameters of the diffusion model described above were fixed and used to stochastically generate targeted subpopulations of the fixed networks given by our social network data. By varying these parameters, we can investigate performance as a function of the underlying component structure of the targeted subnetwork. As we shall see, in terms of relative performance, there are effectively three different regimes of the diffusion model (i.e., targeted subpopulation) parameter space. In all of them PTarget compares favorably with Target, but to different extents and for different reasons that we now discuss. We also plot the growth of the risk multiplier for PTarget, which remains less than 2 in all three regimes.

On each plot, there is a single blue curve showing the performance of the (deterministic) algorithm Target, and multiple red curves showing the performance across 200 runs of our (randomized) algorithm PTarget.

The first regime (Fig. 2) occurs when the largest connected component of the targeted subnetwork is much larger than all of the other components. In this regime, if both algorithms begin at a seed vertex inside the largest component, there is effectively no difference in performance, as both algorithms remain inside this component for the duration of their budget and find identical sets of targeted individuals. More generally, if the algorithms begin at a seed outside the largest component, relative performance is a race to find this component; the private algorithm lags slightly due to the added noise, but is generally quite competitive; see Fig. 2 for details.

The second regime (Fig. 3) occurs when the component sizes are more evenly distributed, but there remain a few significantly larger components. In this setting both algorithms spend more of their budget outside the targeted subpopulation “searching” for these components. Here the performance of the private algorithm lags more significantly—because both algorithms behave the same when inside of a component, the smaller the components are, the more detrimental the noise is to the private algorithm (though again we see particular runs in which the randomness of the private algorithm permits it to actually outperform the nonprivate).

The third regime (Fig. 4) occurs when all of the targeted components are small, and thus both algorithms suffer accordingly, discovering only a few targeted individuals; but again the private algorithm compares favorably with the nonprivate, finding only a few less targeted vertices.

Conclusion

We view the work presented here as a proof of concept: despite the fact that using network analysis to identify members of a targeted population is intrinsically contrary to the privacy of the targeted individuals, we have shown that there is no inherent reason why informational privacy guarantees cannot be given to individuals who are not members of the targeted population, and that these privacy guarantees need not severely harm our ability to find targeted individuals. Our work is of course not a complete solution to the practical problem, which can differ from our simple model in many ways. Here we highlight just one interesting modeling question for future work: Is it possible to give rigorous privacy guarantees to members of the protected population when membership in the targeted population is defined as a function of the individuals’ private data? In our model, we avoid this question by endowing the algorithm with a costly “investigation” operation, which we assume can infallibly determine an individual’s targeted status—but it would be interesting to extend our style of analysis to situations in which this kind of investigation is not available.

- Greenwald G (June 6, 2013) NSA collecting phone records of millions of Verizon customers daily. *The Guardian*. Available at www.theguardian.com/world/2013/jun/06/nsa-phone-records-verizon-court-order.
- National Research Council (2015) *Bulk Collection of Signals Intelligence: Technical Options* (The National Academies Press, Washington, DC).
- Allen J (October 25, 2014) U.S. nurse quarantined over Ebola calls treatment “frenzy of disorganization.” *Reuters*. Available at www.reuters.com/article/health-ebola-usa-obama-idUSL6N0SK0IN20141026.
- Dwork C, McSherry F, Nissim K, Smith A (2006) Calibrating noise to sensitivity in private data analysis. *Proceedings of Third Theory of Cryptography Conference* (Springer, New York), pp. 265–284.
- Daniel J (2006) Solove. A taxonomy of privacy. *Univ Pa Law Rev* 154(3):477–564.
- Dwork C, McSherry FD (2010) Selective privacy guarantees. US Patent 7,818,335.
- He X, Machanavajjhala A, Ding B (2014) Blowfish privacy: Tuning privacy-utility trade-offs using policies. *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data* (ACM, New York), pp 1447–1458.
- Hay M, Li C, Miklau G, Jensen D (2009) Accurate estimation of the degree distribution of private networks. *Ninth IEEE International Conference on Data Mining* (IEEE, New York), pp 169–178.
- Kasiviswanathan SP, Nissim K, Raskhodnikova S, Smith A (2013) Analyzing graphs with node differential privacy. *Theory of Cryptography Conference* (Springer, New York), pp 457–476.
- Blocki J, Blum A, Datta A, Sheffet O (2013) Differentially private data analysis of social networks via restricted sensitivity. *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science* (ACM, New York), pp 87–96.
- Dwork C, Roth A (2014) The algorithmic foundations of differential privacy. *Found Trends Theor Comput Sci* 9(3–4):211–407.
- Bassily R, Groce A, Katz J, Smith A (2013) Coupled-worlds privacy: Exploiting adversarial uncertainty in statistical data privacy. *IEEE 54th Annual Symposium on Foundations of Computer Science* (IEEE, New York), pp 439–448.
- Kearns M, Pai M M, Roth A, Ullman, J (2014) Mechanism design in large games: Incentives and privacy. *Am Econ Rev* 104(5):431–435.
- Alaggar M, Gams S, Keramarrec A-M (2015) Heterogeneous differential privacy. arXiv: 1504.06998.

Supporting Information for “Privacy for the Protected (Only)”

Michael Kearns Aaron Roth Zhiwei Steven Wu Grigory Yaroslavtsev

October 13, 2015

1 Model & Preliminaries

We study graph search algorithms which operate on graphs $G = (V, E)$ defined over a vertex set V and edge set $E \subseteq V \times V$. The vertex set V is partitioned into two fixed subsets $V = \mathcal{T} \cup \mathcal{P}$, where \mathcal{T} represents the *targeted* subpopulation, and \mathcal{P} represents the *protected* subpopulation. The algorithms we consider are initially given a single *seed vertex* $s \in \mathcal{T}$ (or several such vertices), and the goal of the algorithm will be to find as many other members of the targeted subpopulation \mathcal{T} as possible.

The algorithm cannot directly observe which subpopulation a particular vertex $v \in V$ belongs to since otherwise the problem is trivial, but it has the ability to make a query on a vertex $v \in V$ to determine its subpopulation membership. We model this ability formally by giving the algorithm access to an *identity oracle* $\mathcal{I}: V \rightarrow \{0, 1\}$, defined such that $\mathcal{I}(v) = 1$ if and only if $v \in \mathcal{T}$. A call to this oracle is the abstraction we use to represent the possibly costly operation (instantiated in our example applications by e.g. surveillance, or medical tests) which determines whether a particular member of the population is protected or not. Because we view these operations as expensive, we want our algorithm to operate by making as few calls to this oracle as possible. Hence, the algorithm must use the network data represented by the graph G to *guide its search* for which vertices to query. This creates a source of privacy tension since the edges in this network are what we view as private information.

Thus our goal is to give algorithms which discover members of the targeted population using the edges in the network to guide their search. We wish to protect the privacy of the protected individuals: we do not want the outcome of the search to reveal too much about the edge set incident to any protected individual. However, we want to exploit the edges incident to targeted individuals in ways that will not necessarily be privacy preserving. The notion of privacy that we employ is a variant of *differential privacy*. To formally define differential privacy, we consider databases which are multisets of elements from an abstract domain \mathcal{X} , representing the set of all possible data records. (In our case, the data domain \mathcal{X} can be identified with subsets of the vertex set V – it represents the set of all possible *neighbors* that a vertex might be adjacent to in the network).

Definition 1 (Differential Privacy [2]). *Two databases $D, D' \subset \mathcal{X}$ are neighbors if they differ in at most one data record: that is, if there exists an index i such that for all indices $j \neq i$, $D_j = D'_j$. An algorithm $\mathcal{A}: \mathcal{X}^n \rightarrow \mathcal{R}$ satisfies (ϵ, δ) -differential privacy if for every set of outcomes $S \subseteq \mathcal{R}$ and for all neighboring databases $D, D' \in \mathcal{X}^n$, the following holds:*

$$\Pr[\mathcal{A}(D) \in S] \leq \exp(\epsilon) \Pr[\mathcal{A}(D') \in S] + \delta.$$

If $\delta = 0$, we say \mathcal{A} satisfies ε -differential privacy.

This notion of privacy is very strong — indeed, it is *too* strong for our purposes. It provides a symmetric guarantee that does not allow the output of the algorithm to change substantially as a function of any person’s data changing. However, in our case, in order to achieve good utility guarantees we want our algorithm to be allowed to be highly sensitive in the data of the targeted individuals. We will modify this definition in the following way: first, we will view the partition of the vertices into protected and targeted individuals as a fixed, immutable characteristic, separate from the private data of the individuals, and view the private data of individual v as being the edges incident on v ,¹

$$D_v(G) = \{(u, v) \in E \mid u \in V\}. \quad (1)$$

We then redefine the neighboring relation in the setting of networks: for any protected and targeted subpopulations \mathcal{P} and \mathcal{T} , two networks G and G' are *neighboring* if G' can be obtained by only changing a single *protected* node’s edges in G . Specifically, $G = (V, E)$ and $G' = (V, E')$ are *neighbors* with respect to a partition $V = \mathcal{P} \cup \mathcal{T}$ if there exists a $v \in \mathcal{P}$ such that for all $v' \neq v$: $D_{v'}(G) \cup \{(v, v')\} = D_{v'}(G') \cup \{(v, v')\}$. Note that for neighboring graphs G and G' , the edge sets in the subgraph induced on the vertices \mathcal{T} must also be the same.

In the following, we denote the set of all possible networks over the vertices V by \mathcal{G} , and denote the set of all possible outcomes of an algorithm by \mathcal{O} .

Definition 2 (Protected Differential Privacy). *An algorithm $\mathcal{A}: \mathcal{G} \rightarrow \mathcal{O}$ satisfies (ε, δ) -protected differential privacy if for every partition of n vertices V into sets \mathcal{P} and \mathcal{T} , for every pair of graphs G, G' that are neighbors with respect to the partition $(\mathcal{P}, \mathcal{T})$, and for any set of outcomes $S \subseteq \mathcal{O}$*

$$\Pr[\mathcal{A}(G) \in S] \leq \exp(\varepsilon) \Pr[\mathcal{A}(G') \in S] + \delta.$$

If $\delta = 0$, we say \mathcal{A} satisfies ε -protected differential privacy.

When the partition $(\mathcal{P}, \mathcal{T})$ is clear from the context we will omit it to simplify the presentation. In the context of graph search algorithms that we consider here the algorithm \mathcal{A} is given an oracle \mathcal{I} which encodes the partition V into \mathcal{P} and \mathcal{T} . We denote such algorithms as $\mathcal{A}_{\mathcal{I}}$. The output \mathcal{O} in the above definition of *protected differential privacy* in the context of graph search algorithms is an ordered list of targeted individuals.

Remark 1. *A careful reader may already have noticed that there is a trivial graph search algorithm that achieves 0-protected differential privacy while outputting the entire set of targeted individuals \mathcal{T} — it simply queries $\mathcal{I}(v)$ for every $v \in V$, and outputs every v such that $\mathcal{I}(v) = 1$. This algorithm satisfies perfect (i.e. with $\varepsilon = 0$) protected differential-privacy because it operates independently of the private network G . The problem with this approach is that it requires querying the status of every vertex $v \in V$, which can be impractical both because of cost (the query might itself require a substantial investment of resources) and because of societal norms (it may not be defensible to subject every individual in a population to background checks). Hence, here we aim to design algorithms that use the graph data G to effectively guide the search for which vertices v to query. This is what leads to the tension with privacy, and our goal is to effectively trade off the privacy parameter ε with the number of queries to \mathcal{I} that the algorithm must make.*

¹We do not provide any privacy guarantees for what we can reveal about v ’s membership in \mathcal{T} . This is an inherent characteristic of the problem that we consider since the goal of our algorithms is to identify the members of \mathcal{T} .

One way to interpret protected differential privacy is differential privacy applied to an appropriately defined input. Let the algorithm have two inputs: the set of edges incident to the protected vertices in \mathcal{P} , and the edges in $E(\mathcal{T}) = \{(u, v) \mid u, v \in \mathcal{T}\}$ (i.e. all of the other edges)². In this view, protected differential privacy only requires the algorithm to be differentially private in its first argument, and not in its second. This view, formalized in the following lemma, will allow us to apply some of the basic tools of differential privacy in order to achieve protected differential privacy.

Lemma 1. *Given a graph $G = (V, E)$ and a partition its edges into $E_2 = E(\mathcal{T})$ and $E_1 = E \setminus E_2$ an algorithm $\mathcal{A}_{\mathcal{I}}(G) := \mathcal{A}_{\mathcal{I}}(E_1, E_2)$ satisfies (ε, δ) -protected differential privacy if it is (ε, δ) -differentially private in its first argument.*

1.1 Basic Privacy Tools

We include some basic privacy tools here to facilitate the discussion of our algorithm for the rest of the paper. For simplicity, we will state these tools in the generic setting, in which we view algorithms to be arbitrary randomized mappings from \mathcal{X}^n to \mathcal{R} .

A basic, but extremely useful result is that differential privacy is robust to arbitrary *post-processing*:

Lemma 2 (Post-Processing [2]). *For any algorithm $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{R}$ and any (possibly randomized) function $p : \mathcal{R} \rightarrow \mathcal{R}'$, if $\mathcal{A}(\cdot)$ is (ε, δ) -differentially private then $p(\mathcal{A}(\cdot))$ is (ε, δ) -differentially private.*

Another extremely useful property of differential privacy is that it is compositional — given two differentially private algorithms, their combination remains differentially private, with parameters that degrade gracefully. In fact, there are two such composition theorems. The first, simpler one lets us simply add the privacy parameters when we compose mechanisms:

Lemma 3 (Basic Composition [2]). *If $M_1 : \mathcal{X}^n \rightarrow \mathcal{R}_1$ is $(\varepsilon_1, \delta_1)$ -differentially private, and $M_2 : \mathcal{X}^n \times \mathcal{R}_1 \rightarrow \mathcal{R}_2$ is $(\varepsilon_2, \delta_2)$ -differentially private in its first argument, then $M : \mathcal{X}^n \rightarrow \mathcal{R}_2$ is $(\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2)$ differentially private where*

$$M(D) = M_2(D, M_1(D)).$$

We can of course apply the composition theorem repeatedly, and so the composition of m mechanisms, each of which is ε -differentially private is $m\varepsilon$ -differentially private. The second composition theorem, due to [3], allows us to compose m mechanisms while letting the ε parameter degrade *sublinearly* in m (at a rate of only $O(\sqrt{m})$), at the cost of a small increase in the δ parameter.

Lemma 4 (Advanced Composition [3]). *Fix $\delta > 0$. The class of (ε', δ') -differentially private mechanisms satisfies $(\varepsilon, m\delta' + \delta)$ -differential privacy under m -fold adaptive composition³ for*

$$\varepsilon' = \frac{\varepsilon}{\sqrt{8m \log(1/\delta)}}.$$

²It is crucial here that such a simplification can only be made for the purposes of the analysis only. Since all our algorithms are only given access to a membership oracle \mathcal{I} there is no way for them to explicitly construct these two inputs without incurring a cost associated with oracle queries.

³See [3] for a formal exposition of adaptive composition.

When designing private algorithms, we will work extensively with function *sensitivity* for functions defined on data sets — which informally, measures how much the function value can change when a single data entry in the input data set is altered.

Definition 3 (Sensitivity). *The sensitivity Δf of a function $f : \mathcal{X}^n \rightarrow \mathbb{R}$ is defined as*

$$\Delta(f) = \max_{D \sim D'} |f(D) - f(D')|,$$

where $D \sim D'$ indicates that D and D' are neighboring databases.

We will give different notions of sensitivity in the next section, which are more appropriate for some tasks in our setting. Finally, we introduce two simple algorithms that provide differential privacy by adding noise proportional to the sensitivity of a function.

For any function $f : \mathcal{X}^n \rightarrow \mathbb{R}$, the *Laplace mechanism* applied to function f is the algorithm which on input D releases $\hat{f}(D) := f(D) + \nu$, where $\nu \sim \text{Lap}(\Delta(f)/\varepsilon)$ and $\text{Lap}(b)$ denotes the Laplace distribution centered at 0 with probability density function

$$\frac{1}{2b} \exp\left(-\frac{|x|}{b}\right).$$

Lemma 5 ([2]). *The Laplace mechanism is ε -differentially private.*

Another simple algorithm, useful for answering non-numeric queries, is the *Report Noisy Max* mechanism: given a database $D \in \mathcal{X}^n$ and a collection of k functions f_1, f_2, \dots, f_k each with sensitivity at most γ , Report Noisy Max performs the following computation:

- Compute the noisy estimate of each function evaluated on D : $\hat{f}_i := f_i(D) + \nu$ where $\nu \sim \text{Lap}(\gamma/\varepsilon)$;
- Output the index $i^* = \arg \max_i \hat{f}_i$, and also the noisy value \hat{f}_{i^*} .

Lemma 6 ([1]). *The Report Noisy Max mechanism is 2ε -differentially private.*

2 Statistics of Proximity (SoP)

Our family of graph search algorithms will rely on various network-centric statistics of proximity (SoP) that ascribe a numerical measure of proximity of an individual vertex v based on its position in the network relative to a set S of vertices from the targeted population (which will in our usage always be the set of targeted individuals discovered so far by the search algorithm). Specifically, a statistic of proximity is a function f that maps a network G , a node v , and a set of nodes $S \subseteq \mathcal{T}$ to a real number. Since the value $f(G, v, S)$ can reveal information about the links in the network, we will often need to perturb the values of these statistics with noise, calibrated with scale proportional to the *targeted sensitivity* — the maximum change in any *targeted* node's SoP relative to any set S when a *protected node*'s adjacency list is changed.

Definition 4 (Targeted Sensitivity). *Let $f : \mathcal{G} \times V \times 2^{\mathcal{T}} \rightarrow \mathbb{R}$ be a statistic of proximity. The targeted sensitivity of f is*

$$\Delta(f) = \max_{G \sim G', t \in \mathcal{T}, S \subseteq \mathcal{T}} |f(G, t, S) - f(G', t, S)|,$$

where $G \sim G'$ indicates that G and G' are neighboring graphs in \mathcal{G} relative to a fixed partition of V into \mathcal{P} and \mathcal{T} .

Note that when computing the *targeted sensitivity* we are not concerned with the effect that a change in the edges incident on vertices in \mathcal{T} has on the statistic, nor on the effect of any change on the statistic computed on vertices $v \in \mathcal{P}$.

Another quantity of interest is *impact cardinality* — the maximum *number* of nodes whose SoP's can change as the result of a change to the adjacency list of a single node $v \in \mathcal{P}$:

Definition 5 (Impact Cardinality). *Let $f: \mathcal{G} \times V \times 2^{\mathcal{T}} \rightarrow \mathbb{R}$ be a statistic of proximity. The impact cardinality of f is*

$$\text{IC}(f) = \max_{G \sim G', S \subseteq \mathcal{T}} |\{v \in V \mid f(G, v, S) \neq f(G', v, S)\}|.$$

We include some examples of candidate SoPs and their sensitivities. A desirable property for good statistics is that they should have low sensitivity (relative to the scale of the statistic) and small impact cardinality (relative to the target number of queries to the identity oracle), which will allow us to achieve protected differential privacy by adding only small amounts of noise to the various parts of our computations.

- $\text{Flow}_k(G, v, S)$: the value of the maximum flow that can be routed between node v and the nodes in S using only paths of length at most k ;
- $\text{Path}_k(G, v, S)$: the number of paths from v to nodes in S with length at most k ;
- $\text{Triangle}(G, v, S) = |\{\{a, b\} \subseteq S \mid a, b, v \text{ forms a triangle in } G\}|$, the number of triangles formed by the vertex v in G ;
- $\text{CN}(G, v, S) = |\{u \mid (v, u) \in E \text{ and } (u, v') \in E \text{ for some } v' \in S\}|$, the number of common neighbors v has with vertices in S .

In graphs with maximum degree d , the sensitivity of these SoPs are as follows:

- $\Delta(\text{Flow}_k) \leq d$ since a vertex of degree d can only affect the size of the flow by at most d .
- $\Delta(\text{Path}_k) \leq (k-1)d^{k-1}$ since the total number of paths from v to S on which a vertex $u \in \mathcal{P}$ might lie is at most $\sum_{j=1}^{k-1} d^{j-1}d^{k-j} = (k-1)d^{k-1}$. Here we used the index j to denote the index of u along the path starting from v together with the fact that the total number of different paths of length ℓ from u is at most d^ℓ .
- $\Delta(\text{Triangle}) \leq d$ since each triangle is associated with an edge and the total number of edges affected is at most d .
- $\Delta(\text{CN}) \leq 1$ since a single vertex can change the count of common neighbors by at most 1.

Note that $\text{Path}_1(G, v, S)$, which simply counts the number of edges between v and $S \subseteq \mathcal{T}$ actually has targeted sensitivity *zero*. This is because, since $S \subseteq \mathcal{T}$, if $v \in \mathcal{T}$ is also a member of the targeted population, then the statistic is a function only of $E(\mathcal{T})$, the edge set of the subgraph defined over the targeted sub-population \mathcal{T} . Since $E(\mathcal{T})$ is identical on all neighboring graphs, and because targeted sensitivity only measures the sensitivity of the SoP evaluated on *targeted* nodes to changes in *protected* nodes, we get zero sensitivity. This will be important to our analysis.

3 SoP Based Targeting Algorithms

Before we present the full algorithm, we will first present some useful subroutines together with analysis of their privacy properties.

3.1 Statistic-First Search

First, we introduce statistic-first search (SFS), a search algorithm that explores the entire targeted connected component given a seed targeted node. It is a search strategy that only inspects the neighbors of verified targeted nodes. The formal description is presented in Algorithm 1.

Algorithm 1 SFS(G, t)

Input: known targeted t in a network G

Initialize:

$$\tilde{T} = \{t\} \quad I = \{t\} \quad \mathcal{N} = \text{neighbors of } t$$

while $\mathcal{N} \setminus I \neq \emptyset$

Let

$$v' = \operatorname{argmax}_{x \in \mathcal{N} \setminus I} \text{Path}_1(G, x, \tilde{T})$$

Query $\mathcal{I}(v')$ to determine v' 's targeted status.

$I = I \cup \{v'\}$

if $\mathcal{I}(v') = 1$ **then** $\tilde{T} = \tilde{T} \cup \{v'\}$ and $\mathcal{N} = \text{neighbors of } \tilde{T}$

Output: list \tilde{T}

We now establish the simple but remarkable privacy guarantee of SFS — the algorithm can often identify a targeted connected component free of privacy cost.

Lemma 7. *The graph search algorithm SFS satisfies 0-protected differential privacy.*

Proof. Let G and G' be two neighboring networks in \mathcal{G} with respect to the same partition $(\mathcal{P}, \mathcal{T})$. We know that both networks have the same set of targeted nodes \mathcal{T} and targeted links $E(\mathcal{T})$. Since we know that $\Delta(\text{Path}_1) = 0$, and SFS only branches on the evaluations of f on nodes $v \in \mathcal{T}$, the behavior of SFS depends only on \mathcal{T} and $E(\mathcal{T})$, and hence $\text{SFS}(G, v, f)$ and $\text{SFS}(G', v, f)$ always produce the same output. \square

3.2 Private Search for Targeted Component

With SFS, we can start with a seed node $v \in \mathcal{T}$, and at no additional privacy cost, find the entire connected component $T \subseteq \mathcal{T}$ in the subgraph defined on vertices in \mathcal{T} that v belongs to. This is a useful subroutine in a graph search algorithm: however, once we have exhausted our seed node's connected component T , we need a way to search for a new seed node v' that is part of a new connected component. This is what our subroutine **SearchCom** does:

1. Given a list of already identified members of the targeted subpopulation $\tilde{T} \subseteq \mathcal{T}$ and a SoP, we compute a noisy SoP value for each node $\hat{f}(v)$;

2. we sort the nodes in decreasing order of their noisy SoP value, and query each vertex v in this order to determine whether $v \in \mathcal{T}$ or $v \in \mathcal{P}$ until we find a node such that $v \in \mathcal{T}$.
3. If we query K nodes without having found any members of the targeted sub-population, we halt the search. The stopping condition needs to be checked privately, so K is in fact a randomly perturbed value.

We include a formal description of SearchCom in Algorithm 2.

Algorithm 2 SearchCom($G, \tilde{T}, I, f, \varepsilon, K$)

Input: identified members of the targeted population $\tilde{T} \subseteq \mathcal{T}$ in a network G , the set of investigated nodes I , SoP f , privacy parameter ε , and stopping threshold K

Initialize:

Set noisy stopping threshold $\hat{K} = K + \nu$ where $\nu \sim \text{Lap}(2\text{IC}(f)/\varepsilon)$ and count = 0

for each $v \in V \setminus I$:

let $\hat{f}(v) = f(G, v, \tilde{T}) + \zeta_v$ where $\zeta_v \sim \text{Lap}(4\Delta(f)/\varepsilon)$

while $(V \setminus I) \neq \emptyset$ and count $\leq \hat{K}$

Let

$$v' = \underset{x \in V \setminus I}{\text{argmax}} \hat{f}(x)$$

Query $\mathcal{I}(v')$ to determine if $v' \in \mathcal{T}$.

Let $I = I \cup \{v'\}$ and count = count + 1

if $\mathcal{I}(v') = 1$ **then return** $\{v'\}$

return \emptyset

Lemma 8. *The targeting algorithm SearchCom instantiated with privacy parameter ε satisfies ε -protected differential privacy.*

Proof. Let G and G' be neighboring networks in \mathcal{G} . First suppose that $\tilde{T} = \mathcal{T}$. In this case, SearchCom will output \emptyset with probability 1 on both inputs G and G' , and hence satisfy 0-protected differential privacy. Hence, for the remainder of the argument, we can assume that there exists a vertex $v \in \mathcal{T} \setminus \tilde{T}$.

In this case, the algorithm can equivalently be viewed as the following 2-step procedure:

1. Use the Report Noisy Max algorithm to output the index of the targeted node t which maximizes $\hat{f}(t)$ together with the perturbed value $\hat{f}(t)$;
2. Let b be the number of nodes $i \in \mathcal{P}$ such that $\hat{f}(i) \geq \hat{f}(t)$. If $b - \nu \leq K$, output node t . Otherwise output \emptyset .

When the input to the algorithm $G = (V, E)$ is viewed as the pair of edge sets $(E \setminus E(\mathcal{T}), E(\mathcal{T}))$, we show below that each of these two steps satisfies $\varepsilon/2$ -differential privacy with respect to its first argument. By the basic composition theorem Lemma 3 and Lemma 1 the algorithm SearchCom satisfies ε -protected differential privacy.

The first step is an instantiation of the Report Noisy Max algorithm with privacy parameter $\varepsilon/2$, so it is $\varepsilon/2$ -differentially private by Lemma 6.

The second step is post-processing of $\hat{b} = b - \nu$. We just need to show that releasing \hat{b} is $\varepsilon/2$ differentially private, and the result will follow by the post processing lemma: Lemma 2. For the following analysis, we will fix the (arbitrary) values of $\{\zeta_v\}$ and compute probabilities as a function of the randomness of ν .

Since the SoP f has impact cardinality $\text{IC}(f)$, we know that there are at most $\text{IC}(f)$ many nodes v such that $f(G, v, \tilde{T}) \neq f(G', v, \tilde{T})$. It follows that

$$\left| \{v \mid f(G, v, \tilde{T}) + \zeta_v \neq f(G', v, \tilde{T}) + \zeta_v\} \right| \leq \text{IC}(f)$$

Let $b(G)$ and $b(G')$ denote the number of nodes $i \in \mathcal{P}$ with $\hat{f}(v) \geq \hat{f}(t)$ in G and G' respectively. Then we know that,

$$|b(G) - b(G')| \leq \text{IC}(f).$$

Since we are releasing \hat{b} by adding noise ν sampled from the Laplace distribution $\text{Lap}(2\text{IC}(f)/\varepsilon)$, \hat{b} is $\varepsilon/2$ -differentially private from the property of Laplace mechanism Lemma 5. \square

3.3 The Full Algorithm: Putting the Pieces Together

Our graph search algorithm alternates between two phases. In the first phase, the algorithm starts with a seed node $v \in \mathcal{T}$, and uses SFS to find every other vertex $v' \in \mathcal{T}$ that is part of the same connected component as v in the subgraph defined on \mathcal{T} . After this targeted component has been fully identified, the second phase begins. In the second phase, the algorithm uses SearchCom to search for a new vertex $v \in \mathcal{T}$ that will serve as a seed node for the next iteration of SFS. Once such a seed node has been found, the algorithm reverts to phase 1, and this continues for a specified number of iterations. The formal description of the algorithm is presented in Algorithm 3.

Algorithm 3 Private Search Algorithm: PTarget($G, \mathbf{s}, f, k, N, \varepsilon$)

Input: A network G , a seed node $\mathbf{s} \in \mathcal{T}$, a SoP f for SearchCom, a target number of components k , a stopping threshold N for SearchCom, and privacy parameter ε

Initialize: Use set I to keep track of the set of investigated nodes. Initially $I = \{\mathbf{s}\}$.

Let list $\tilde{T} = \text{SFS}(G, \mathbf{s})$

For $k - 1$ rounds:

let $a = \text{SearchCom}(G, \tilde{T}, I, f, \varepsilon, N)$

if $a = \emptyset$

Output \tilde{T}

else

let $\tilde{T} = \tilde{T} \cup \text{SFS}(G, a)$

Output \tilde{T}

We now establish the following privacy guarantee of Algorithm 3. Recall that the parameter k represents the maximum number of disjoint components of the subgraph defined on \mathcal{T} that the algorithm will identify.

Theorem 1. Fix any $0 < \delta < 1$. $\text{PTarget}(\cdot, \cdot, \cdot, \cdot, k, \cdot, \cdot, \varepsilon)$ satisfies ε_1 -protected differential privacy for

$$\varepsilon_1 = (k - 1)\varepsilon,$$

and satisfies (ε_2, δ) -protected differential privacy for

$$\varepsilon_2 = 2\sqrt{2(k - 1)\ln(1/\delta)}\varepsilon.$$

Proof. The algorithm is a composition of at most k instantiations of **SFS** and $(k - 1)$ instantiations of **SearchCom** with privacy parameter ε . Recall that each call to **SFS** is 0-differentially private, and each call to **SearchCom** is ε -differentially private with respect to the edges incident on vertices in \mathcal{P} . By the composition theorem, we know that the algorithm is $(k - 1)\varepsilon$ -differentially private by Lemma 3, and at the same time $(\sqrt{8k\ln(1/\delta)}, \delta)$ -differentially private for any $\delta \in (0, 1)$ by Lemma 4. Our result then easily follows from Lemma 1. □

4 Experiments

4.1 Subpopulation Construction

Our experiments are conducted on two real social networks:

- the scientific collaboration network in DBLP (“Digital Bibliography and Library Project”), where nodes represent authors and edges connect authors that have coauthored a paper;
- the movie costarring network in IMDB (“Internet Movie Database”), where nodes represent actors and edges connect actors that have appeared in a movie together.

Pre-processing step on the networks: We sparsify the IMDB and DBLP networks by removing a subset of the edges. This will allow us to generate multiple targeted components more easily. In both networks, there is a natural notion of weights for the edges. In the case of DBLP, the edge weights correspond to the number of papers the individuals have co-written. In the case of IMDB, the edge weights correspond to the number of movies two actors have co-starred in. In our experiments, we only remove edges with weights less than 2.

However, the networks we use do not have an identified partition of the vertices into a targeted and protected subpopulation. Instead, we generate the targeted subpopulation synthetically using the following diffusion process. We use the language of “infection”, which is natural, but we emphasize that this process is not specific to our motivating example of the targeted population representing people infected with a dangerous disease. The goal of the infection process is to generate a targeted subpopulation \mathcal{T} such that:

1. The subnetwork restricted to \mathcal{T} has multiple distinct connected components (so that the search problem is algorithmically challenging, and isn’t solved by a single run of statistic-first search), and
2. The connected components of \mathcal{T} are close to one another in the underlying network G , so that the network data is useful in identifying new members of \mathcal{T} .

The process $\text{infect}(G, s, p, q, k)$ takes as input a seed infected node s , two values $p, q \in (0, 1)$, and a number of rounds k , and proceeds with two phases:

1. **Infection phase:** Initially, only the node s is in the infected set $\tilde{\mathcal{I}}$. Then in each of the k rounds, each neighbor v of the infected nodes $\tilde{\mathcal{I}}$ becomes infected independently with probability p .
2. **Immune phase:** After the infection process above, we will set some of the infected nodes as immune. For each node i in the infected node set $\tilde{\mathcal{I}}$, let i become “immune” (non-infected) with probability q .

The infection phase above is common in the literature on contagion in networks; e.g. relatively recently it was considered as the “independent cascades model” of [4]. The immune phase is introduced in order to permit control of the component structure while still keeping targeted vertices proximate to each other in the resulting network. We include a formal description of the algorithm in Algorithm 4.

Algorithm 4 $\text{infect}(G, s, p, q, k)$

Input: a network G , a seed node s in G , infection probability p , and immune probability q
Initially the infected population contains only the seed node:

$$\tilde{I} = \{s\}$$

for $t = 1, \dots, k$:
 for each node v that is neighbor to \tilde{I} :
 Let ν be a uniformly random number from $[0, 1]$
 if $\nu \leq p$ **then** $\tilde{I} = \tilde{I} \cup \{v\}$
let $\mathcal{T} = \emptyset$
for each node $v' \in \tilde{I}$:
 let ν be a uniformly random number from $[0, 1]$
 if $\nu > q$ **then** $\mathcal{T} = \mathcal{T} \cup \{v\}$
Output: \mathcal{T} as the targeted subpopulation

4.2 Non-Private Benchmark Target

We experimentally evaluate the performance of our algorithm Algorithm 3 on two social network data-sets with a partition of vertices into \mathcal{P} and \mathcal{T} using the infection process described in the previous section. We compare the performance of the private version of our algorithm with the non-private version Algorithm 5 which uses the SoP directly, without adding noise. The metric we are interested in is how many queries to the identify oracle \mathcal{I} are needed by each algorithm to find a given number of members of the protected sub-population \mathcal{T} . We here give a formal description of the non-private version of our graph search algorithm in Algorithm 5.

Algorithm 5 Non-Private Targeting Algorithm: $\text{Target}(G, \mathbf{s}, f, k, N)$

Input: A network G , a seed node $\mathbf{s} \in \mathcal{T}$, a SoP f for SearchCom, a target number of components to find k , and a stopping threshold N for SearchCom

Initialize: Use I to keep track of the set of investigated nodes. Initially $I = \{\mathbf{s}\}$.

Let list $\tilde{T} = \text{SFS}(G, \mathbf{s})$

For $k - 1$ rounds:

let count = 0 and $a = \emptyset$

while $(V \setminus I) \neq \emptyset$ and count $\leq N$

Let

$$v' = \underset{x \in V \setminus I}{\operatorname{argmax}} f_2(G, v, \tilde{T})$$

 Query $\mathcal{I}(v')$ to determine if $v' \in \mathcal{T}$.

Let $I = I \cup \{v'\}$ and count = count + 1

if $\mathcal{I}(v') = 1$ **then let** $a = v'$ and **break**

if $a = \emptyset$ **then Output** \tilde{T}

else let $\tilde{T} = \tilde{T} \cup \text{SFS}(G, a, f_1)$

Output \tilde{T}

4.3 SoP Instantiations

In our experiments, we will use the SoP CN for the SearchCom subroutine, which is the number of common neighbors between the node v and the subset of nodes S representing the already discovered members of the targeted population. The targeted sensitivity of CN is bounded by 1.

Lemma 9. *The SoP CN has targeted sensitivity $\Delta(\text{CN})$ bounded by 1.*

Proof. Let G and G' be two neighboring networks over the same protected and targeted populations \mathcal{P} and \mathcal{T} . Let $t \in \mathcal{T}$ be a targeted node and $S \subseteq V$ be a subset of nodes. Since G and G' only differ by the edges associated with a protected node i , we know that the neighbor sets of both t and S can differ by at most one node between G and G' . Note that the $\text{CN}(G, t, S)$ computes the cardinality of the intersection between these two sets, and the intersection sets of these two networks can differ by at most one node. It follows that $\Delta(\text{CN}) \leq 1$. \square

References

- [1] Cynthia Dwork and Aaron Roth. [The algorithmic foundations of differential privacy](#). *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- [2] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. [Calibrating noise to sensitivity in private data analysis](#). In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, pages 265–284, 2006.
- [3] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. [Boosting and differential privacy](#). In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 51–60, 2010.

- [4] David Kempe, Jon Kleinberg, and Eva Tardos. [Maximizing the spread of influence through a social network](#). In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003.