
Learning Diffusion using Hyperparameters

Dimitris Kalimeris¹ Yaron Singer¹ Karthik Subbian² Udi Weinsberg²

Abstract

In this paper we advocate for a hyperparametric approach to learn diffusion in the independent cascade (IC) model. The sample complexity of this model is a function of the number of edges in the network and consequently learning becomes infeasible when the network is large. We study a natural restriction of the hypothesis class using additional information available in order to dramatically reduce the sample complexity of the learning process. In particular we assume that diffusion probabilities can be described as a function of a global hyperparameter and features of the individuals in the network. One of the main challenges with this approach is that training a model reduces to optimizing a non-convex objective. Despite this obstacle, we can shrink the best-known sample complexity bound for learning IC by a factor of $|E|/d$ where $|E|$ is the number of edges in the graph and d is the dimension of the hyperparameter. We show that under mild assumptions about the distribution generating the samples one can provably train a model with low generalization error. Finally, we use large-scale diffusion data from Facebook to show that a hyperparametric model using approximately 20 features per node achieves remarkably high accuracy.

1. Introduction

For well over a decade there has been extensive work on learning social network influence models (Liben-Nowell & Kleinberg, 2003; Netrapalli & Sanghavi, 2012; Abrahao et al., 2013; Friggeri et al., 2014; Anderson et al., 2015; Subbian et al., 2017), and the independent cascade model in particular (Saito et al., 2008; Gomez Rodriguez et al., 2010; Goyal et al., 2010; Gomez Rodriguez et al., 2011; Du et al., 2014; Lemonnier et al., 2014; Bourigault et al.,

^{*}Equal contribution ¹Department of Computer Science, Harvard University ²Facebook, Menlo Park. Correspondence to: Dimitris Kalimeris <kalimeris@g.harvard.edu>.

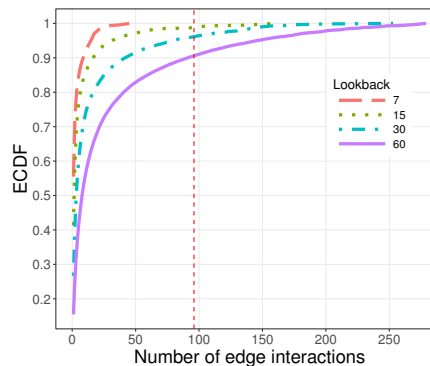


Figure 1. The plot shows the empirical CDF of the number of per edge interactions in Facebook (Events dataset, see Section 5.3).

2014; Narasimhan et al., 2015). Independent cascade (IC) was popularized by the seminal work of Kempe, Kleinberg, and Tardos (Kempe et al., 2003) and is a stochastic model that predicts the likelihood of information diffusing from one individual to another in a social network. In this model for every pair of individuals connected in the network u, v there is a probability $p_{u,v}$ that v adopts the behavior of u (i.e. information is *diffused* from u to v).

The main challenge with learning the IC model is that the sample complexity is often overwhelmingly large or simply infeasible. To illustrate this point, Figure 1 shows the cumulative distribution of edge interactions for millions of public events on Facebook over varying periods of time, ranging from one week to two months (see detailed description of the dataset in Section 5). The vertical line marks the minimal number of observations per edge required to infer the likelihood of influence with error 0.1 and confidence 95%. In this data set, more than 90% of the edges do not have enough observations to learn the respective diffusion probabilities accurately, even over a period of two months. Furthermore, in a single week (the timeframe in which inference about an event is often most relevant), *none of the edges in the data set have sufficiently many observations*.

Given that even with the data available on Facebook there are not enough observations to learn the model, one needs to impose additional assumptions. A natural approach is to assume that the diffusion probabilities are a function of network and individuals' characteristics and some underlying

global hyperparameter θ . In the case of events for example, it seems reasonable that influence could be estimated as a function of some global unknown multidimensional parameter θ and individuals' characteristics such as location, gender, and age, and topological features like the ratio between the intersection and the union of the individuals' neighborhoods. Using $\mathbf{x}_{u,v}$ to denote the characteristics of u and v , the hyperparametric approach assumes that the probability of u to influence v denoted $p_{u,v}$ is not arbitrary and can be faithfully estimated via some function p that maps θ and $\mathbf{x}_{u,v}$ to $[0, 1]$, i.e. $p_{u,v} = p(\theta, \mathbf{x}_{u,v})$. Given a set of characteristics, learning the IC model then reduces to recovering the underlying hyperparameter θ .

Intuitively, learning a hyperparametric model necessitates far fewer samples than a general diffusion model for two main reasons. First, since the diffusion probabilities are correlated, each observation provides information about all edges in the network. Second, it seems reasonable that the sample complexity of learning the hyperparameter should largely depend on the dimension of the hyperparameter rather than the number of edges the network.

A simple example. To solidify our intuition, consider a simple bipartite network $G = (U, V, E)$ where nodes in U attempt to activate nodes from V as depicted in Figure 1.1 and each activation attempt together with its outcome (label) constitutes one sample. Our goal is to find a $\hat{p}_{u,v}$ for every edge $(u, v) \in E$, s.t. with prob. at least $1 - \delta$ for all edges:

$$|p_{u,v} - \hat{p}_{u,v}| \leq \epsilon$$

Hoeffding's inequality and a union bound imply that $\Theta\left(\frac{|E|}{\epsilon^2} \log \frac{|E|}{\delta}\right)$ samples are necessary and sufficient to learn the diffusion probabilities on all the edges in the graph. In comparison, suppose that the diffusion probability of each edge is a function of a hyperparameter $\theta \in [0, 1]^d$ and some known features of the edge $\mathbf{x}_{u,v} \in [0, 1]^d$ as follows:

$$p_{u,v} = \frac{1}{1 + e^{-\langle \theta, \mathbf{x}_{u,v} \rangle}}$$

Then, learning the diffusion probabilities becomes a logistic regression problem and thus only $\mathcal{O}\left(\frac{d}{\epsilon^2} \log \frac{d}{\delta}\right)$ samples are required, independent of the number of edges. This reduces the sample complexity by a factor of $\frac{|E|}{d}$ which is quite dramatic when the number of edges in the network $|E|$ is large and the dimension of the hyperparameter d is small.

Beyond potential improvements in sample complexity, a hyperparametric model is convenient due to the structure it imposes. A recent line of work on influence maximization in bandit models (Wen et al., 2015; Vaswani et al., 2017), assumes that the diffusion probabilities are a linear function of edge features, i.e. $p_{u,v} = \langle \theta, x_{u,v} \rangle$ and this structure is leveraged in order to develop faster algorithms.




Figure 2. Learning the diffusion probabilities on simple bipartite graphs. Green symbolizes active nodes while red inactive ones.

1.1. A Hyperparametric Approach

Our goal in this paper is to explore a hyperparametric approach for learning the independent cascade diffusion model. Doing so requires addressing three open questions:

- *Does restriction to a low-dimensional hyperparameter substantially decrease the sample complexity?* As discussed above, the motivation for a hyperparametric approach is that *intuitively* its sample complexity should depend on the dimension of the hyperparameter rather than the number of edges. While intuitive, when the indegree of nodes is greater than 1, minimizing empirical risk becomes a non-convex optimization problem and analyzing sample complexity is not trivial;
- *Can a hyperparametric model be learned efficiently?* As learning a hyperparametric model requires solving a non-convex optimization problem, it is not clear it can be learned efficiently, in theory or in practice;¹
- *Are low-dimensional hyperparametric models predictive?* Assuming that sample complexity heavily depends on its dimension, our approach is relevant only if reasonable estimates of the IC model are achievable with a low-dimensional hypothesis class.

In this paper we address the above questions. We first show that the sample complexity can indeed be dramatically reduced when restricting the hypothesis class to a low-dimensional hyperparameter. Specifically, when comparing with the state-of-the-art bound for learning the independent cascade model (without the hyperparametric assumption) we show that the sample complexity can be reduced by a factor of $|E|/d$, as foreshadowed by the example above.

Despite being a non-concave optimization problem we show that the problem has a great deal of structure. Under mild assumptions about the distribution generating the samples, we show how this structure can be leveraged to efficiently train a model with arbitrarily small generalization error.

Lastly, we show that the hyperparametric approach does work in practice. To do so, we ran experiments on large scale cascades recorded on the Facebook social network. We show that with a hyperparameter of dimension 40 one

¹We note that even without the hyperparametric assumption PAC learning the IC model is a non-convex optimization problem (Narasimhan et al., 2015).

can estimate the diffusion probabilities with remarkably high accuracy. Naturally, there are data sets that do not contain individuals' characteristics. Nonetheless, most social networking services include useful information about its members that help predict diffusion, and the topology of the network alone often may serve as a good proxy.

2. The Hyperparametric Model

A social network is a finite directed graph $G = (V, E)$, where the set of nodes V represents the individuals in the network and the set of edges E represents their social links.

Independent Cascade (IC) model. The IC model assumes that every node $v \in V$ can either be *active* or *inactive*. All nodes begin as inactive. At time step $t = 0$ a subset of the nodes X called *the seed* becomes active, and activations in the network continue according to the following stochastic process: Each node u that became active at time step $t = \tau$ attempts to influence every one of its neighbors v *only once* at time step $t = \tau + 1$ independently, and succeeds with some probability $p_{u,v}$. A node v that became active during the process will never go back to being inactive. Our work generalizes the standard IC model by assuming that the probabilities $p_{u,v}$ are not arbitrary but correlated and specifically consequences of nodes' features. This is the hyperparametric assumption, as described below.

Hyperparametrization. Every node $u \in V$ is associated with a vector of features containing information about it. Every edge $(u, v) \in E$, is also associated with a feature vector, the concatenation of the feature vectors of its endpoints, denoted by $x_{u,v}$. The diffusion probability of each edge is a function of a global hyperparameter θ and its feature vector. Formally, we assume that there exists a function $p : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$ s.t. $p_{u,v} = p(\theta, x_{u,v})$ for any edge $(u, v) \in E$. In this work we define p as the sigmoid function:

$$p_{u,v} = \sigma(\theta, x_{u,v}) = \frac{1}{1 + e^{-\langle \theta, x_{u,v} \rangle}}.$$

We restrict the hyperparameter θ to lie in a hypothesis class $\mathcal{H} = [-B, B]^d$ for some constant $B > 0$, and w.l.o.g. we assume that the feature vector of every edge lies in $[0, 1]^d$. Additionally, we assume that $p_{u,v}$ is bounded away from 0 and 1 for all edges, i.e. $p_{u,v} \in [\lambda, 1 - \lambda]$, for some $\lambda > 0$.

Further discussion about the hyperparametric model and the choice of the sigmoid function can be found in Appendix A.

Samples. The input to a learning algorithm is a collection of labeled samples. We assume that there is some unknown distribution \mathcal{D}_0 over subset of nodes, that activates the initial seed of the cascade, V_0 . Subsequently, as we discussed before, we can partition $V \setminus V_0$ into subsets of nodes V_1, V_2, \dots, V_{n-1} that become activated at steps

$\tau = 1, 2, \dots, n - 1$, respectively². Notice that the cascade can be further decomposed into a sequence of simpler samples as follows: for every $\tau \in \{0, 1, \dots, n - 1\}$ consider all the nodes $v \notin \cup_{t=0}^{\tau-1} V_t$ that are within distance of 1 from V_τ . For every v that became activated by V_τ (i.e. $v \in V_{\tau+1}$) create the sample $((V_\tau, v), 1)$, and for every v that remained inactive create the sample $((V_\tau, v), 0)$. Throughout this paper we assume that the input to our learning algorithm is of the form $\{(X_i, v_i), y_i\}_{i=1}^m$ where $X_i \subseteq V$ is a subset of active nodes, v_i is a node in distance 1 from X_i and $y_i \in \{0, 1\}$ is its label. In Appendix C we map every seed-generating distribution \mathcal{D}_0 to a sample-generating distribution \mathcal{D} .

Log-likelihood of a sample. For every node v , the event “ v becomes influenced by X , when the hyperparameter has value θ ” is a Bernoulli random variable with probability of success $f_v^\theta(X) = 1 - \prod_{u \in X \cap N(v)} (1 - p_{u,v}(\theta))$ where $N(v)$ is the set of *in-neighbors* of node v . Hence, the likelihood of a sample $s = ((X, v), y)$, where $v \notin X$ is $f_v^\theta(X)^y \cdot (1 - f_v^\theta(X))^{1-y}$, and the respective log-likelihood of s is:

$$\mathcal{L}(s, \theta) = y \ln(f_v^\theta(X)) + (1 - y) \ln(1 - f_v^\theta(X)) \quad (1)$$

We want to recover a hyperparameter θ that yields accurate estimates. To do so, given a training set $S = \{s_i\}_{i=1}^m$, we seek the most probable hyperparameter generating S by maximizing the cumulative log-likelihood function:

$$\hat{\theta} = \arg \max_{\theta \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \mathcal{L}(s_i, \theta) \quad (2)$$

Learning a diffusion model. Our goal is to bound the sample complexity, i.e. the number of i.i.d. samples generated by a distribution \mathcal{D} that we need to observe to PAC learn \mathcal{H} . That is, guarantee that $\sup_{\theta \in \mathcal{H}} \mathbb{E}_{s \sim \mathcal{D}} [\mathcal{L}(s, \theta)] - \mathbb{E}_{s \sim \mathcal{D}} [\mathcal{L}(s, \hat{\theta})] \leq \epsilon$, with probability at least $1 - \delta$ (see definition of PAC learnability in Appendix B).

Notice that while there are $|E|$ edges in the network, which translates to $|E|$ diffusion probabilities, in the optimization problem (2) there are only d parameters to be learned.

We would like to note that PAC learning guarantees are required to hold for any distribution \mathcal{D} that generates the data. Hence, it is easy to see that the diffusion probabilities or the hyperparameter θ are not learnable without extra assumptions on \mathcal{D} . For details refer to Appendix D.

3. Learning a Hyperparametric Model

In this section we prove Theorem 2 which is the main technical result of the paper. The main takeaway is that a hyperparameteric approach makes learning an influence model feasible. Informally, the theorem states that the number of samples required to (ϵ, δ) -PAC learn the model is

²The influence process terminates after at most $|V| - 1$ steps.

$\tilde{\mathcal{O}}\left(\Delta^2\left(\frac{\Delta \cdot d + \log \frac{1}{\delta}}{\epsilon^2}\right)\right)$, where Δ is the maximum degree in the network and d is the dimension of the hyperparameter. As we later show in the experiments section, very small *constant* values of d suffice to learn an influence model with almost no error on real data. This is in sharp contrast to the best sample complexity guarantees due to (Narasimhan et al., 2015) for learning the model without the hyperparametric assumption which is $\tilde{\mathcal{O}}\left(\Delta^2\left(\frac{\Delta \cdot |E| + \log \frac{1}{\delta}}{\epsilon^2}\right)\right)$.

Furthermore, imposing assumptions on the distribution \mathcal{D} , can reduce the dependence on Δ making the sample complexity (almost) independent of the size of the network.

3.1. Sample Complexity via Radamacher Complexity

The main technical challenge is due to the fact that the MLE objective in (2) is non-concave and we cannot immediately derive sample complexity bounds from convergence guarantees of Stochastic Gradient Descent for example. Instead, to analyze the sample complexity we will argue about the Rademacher complexity of our hypothesis class by using covering numbers. Informally, the Rademacher complexity measures the expressive power of a hypothesis class \mathcal{H} with respect to a probability distribution and the covering number of a set is the number of balls of a certain radius whose union contains the set (see Definitions 2 and 3 in Appendix B). Recall that the sample complexity of a hypothesis class can be derived from its Rademacher complexity.

Theorem 1 ((Shalev-Shwartz & Ben-David, 2014)). *Assume that for every sample $s \sim \mathcal{D}$ and every $\theta \in \mathcal{H}$ we have that: $|\mathcal{L}(s, \theta)| \leq C$. Let $S \sim \mathcal{D}^m$ and $\hat{\theta} = \arg \max_{\theta \in \mathcal{H}} \left\{ \frac{1}{m} \sum_{i=1}^m \mathcal{L}(s_i, \theta) \right\}$. Then with probability at least $1 - \delta$ over the choice of S we have that:*

$$\mathbb{E}_{s \sim \mathcal{D}}[\mathcal{L}(s, \hat{\theta})] \geq \sup_{\theta \in \mathcal{H}} \mathbb{E}_{s \sim \mathcal{D}}[\mathcal{L}(s, \theta)] - \mathcal{R}(S, \mathcal{H}) - \mathcal{O}\left(C \sqrt{\frac{\log \frac{1}{\delta}}{m}}\right)$$

where $\mathcal{R}(S, \mathcal{H})$ is the Rademacher complexity of the class \mathcal{H} with respect to S .

Hence, our goal reduces to bounding $\mathcal{R}(S, \mathcal{H})$ for a training set S of size m . We do so by discretizing \mathcal{H} by ϵ , and prove that if the discretization is dense enough, then we do not sacrifice a lot by searching for the most likely hyperparameter in the discrete space \mathcal{H}_ϵ instead of the continuous \mathcal{H} .

To this end, we construct an ϵ -cover of the hypothesis class $\mathcal{H} = [-B, B]^d$. Proving that the log-likelihood of any fixed sample s , is bounded and Lipschitz³ in θ with respect to

³intuitively for a Lipschitz function a small change in the argument cannot lead to a large change in the value of the function, see Definition 4, Appendix B.

the ℓ_1 -norm, where the Lipschitz parameter depends on λ (Lemma 3 in Appendix E), allows us to translate the cover of the space of the hyperparameter, into a cover of the space of the log-likelihood functions, by slightly increasing the number of points we include in it, as stated in Lemma 1. The proof is deferred to Appendix E.

Lemma 1. *Let $S = \{(X_i, v_i), y_i\}_{i=1}^m$ be a non-empty set of samples and let $\Delta_S = \max_{s \in S} |X \cap N(v)|$ (maximum in-degree of a node that was activated, across all samples). The covering number of the class of all log-likelihood functions for S is $\mathcal{O}\left(\left(\frac{B\rho d}{\lambda \Delta_S \epsilon}\right)^d\right)$, i.e. we can choose a discrete cover $\mathcal{H}_\epsilon \subseteq \mathcal{H}$ of size $\mathcal{O}\left(\left(\frac{B\rho d}{\lambda \Delta_S \epsilon}\right)^d\right)$, such that for all $\theta \in \mathcal{H}$, there exists a $\theta_\epsilon \in \mathcal{H}_\epsilon$ with*

$$\sup_{s \in S} |\mathcal{L}(s, \theta) - \mathcal{L}(s, \theta_\epsilon)| \leq \epsilon.$$

Given the above lemma, we can invoke Massart's lemma (Lemma 5 in Appendix E) on \mathcal{H}_ϵ which upper bounds the Rademacher complexity of *finite* hypothesis classes. Subsequently, we use Lemma 4 (Appendix E) to upper bound the Rademacher complexity of \mathcal{H} from that of \mathcal{H}_ϵ . We are now ready to prove the main theorem of the section.

Theorem 2 (Sample Complexity of MLE). *Let $G = (V, E)$ be a directed graph and \mathcal{D} be a distribution that generates samples of the form $s = ((X, v), y)$. Let $\Delta = \max_{s \sim \mathcal{D}} |X \cap N(v)|$. Then, for any $\epsilon, \delta \in (0, 1)$, if we use Maximum Likelihood Estimation on a training set of size $m \geq m(\epsilon, \delta) = \mathcal{O}\left(\Delta^2 \log^2(1/\lambda) \frac{d \log(B\rho d / \lambda \Delta \epsilon) + \log(1/\delta)}{\epsilon^2}\right)$ samples drawn i.i.d. from \mathcal{D} , with probability at least $1 - \delta$ (over the draw of the training set) it holds:*

$$\sup_{\theta \in \mathcal{H}} \mathbb{E}_{s \sim \mathcal{D}}[\mathcal{L}(s, \theta)] - \mathbb{E}_{s \sim \mathcal{D}}[\mathcal{L}(s, \hat{\theta})] \leq \epsilon.$$

Proof. Define $\Delta := \max_{s \sim \mathcal{D}} |X \cap N(v)|$, i.e. the maximum active in-degree that any sample generated by \mathcal{D} can have. Then, applying Lemma 1 we can create a discrete ϵ -cover of the space of the log-likelihoods, $\mathcal{H}_\epsilon \subseteq \mathcal{H}$, of size $|\mathcal{H}_\epsilon| = \mathcal{O}\left(\left(\frac{B\rho d}{\lambda \Delta \epsilon}\right)^d\right)$ for any training set S of any size. Invoking Lemma 4 (Appendix E), we can associate the Rademacher complexity of \mathcal{H} with that of its cover \mathcal{H}_ϵ , for any S and any $\epsilon > 0$, as follows:

$$\mathcal{R}(S, \mathcal{H}) \leq \mathcal{R}(S, \mathcal{H}_\epsilon) + 2\epsilon.$$

Hence, we can focus on bounding the Rademacher complexity of \mathcal{H}_ϵ instead of that of \mathcal{H} . Since \mathcal{H}_ϵ is finite, the well-known Massart's lemma apply yielding:

$$\begin{aligned} \mathcal{R}(S, \mathcal{H}) &\leq \mathcal{R}(S, \mathcal{H}_\epsilon) + 2\epsilon \\ &\leq 2 \max_{\theta \in \mathcal{H}_\epsilon} \left\| \left(\mathcal{L}(s_i, \theta) \right)_{i=1}^m \right\|_2 \frac{\sqrt{2 \log(|\mathcal{H}_\epsilon|)}}{m} + 2\epsilon \\ &\leq 2\sqrt{m} \Delta \ln \frac{1}{\lambda} \cdot \frac{\sqrt{2 \log(|\mathcal{H}_\epsilon|)}}{m} + 2\epsilon \end{aligned}$$

$$\begin{aligned}
 &= 2\Delta \ln \frac{1}{\lambda} \sqrt{\frac{2 \log(|\mathcal{H}_\epsilon|)}{m}} + 2\epsilon \\
 &= \mathcal{O}\left(\Delta \log \frac{1}{\lambda} \sqrt{\frac{d \log(B\rho d/\lambda^\Delta \epsilon)}{m}}\right) + 2\epsilon
 \end{aligned}$$

where the first inequality holds because of Lemma 4 (discretization), the second because of Massart’s lemma (finite hypothesis class), the third because of Lemma 3 (bounded \mathcal{L}), and the last one because of Lemma 1 (covering number).

Setting $\epsilon = 1/m$ yields: $\mathcal{R}(S, \mathcal{H}) = \mathcal{O}\left(\Delta \log(1/\lambda) \sqrt{\frac{d \log(B\rho d m/\lambda^\Delta)}{m}}\right)$. Now using the generalization bound of Theorem 1, one can see that in order to achieve $\mathbb{E}_{s \sim D}[\mathcal{L}(s, \hat{\theta})] \geq \sup_{\theta \in \mathcal{H}} \mathbb{E}_{s \sim D}[\mathcal{L}(s, \theta)] - \epsilon$, with probability at least $1 - \delta$, we need S to be of size $m = \mathcal{O}\left(\Delta^2 \log^2(1/\lambda) \frac{d \log(B\rho d/\lambda^\Delta \epsilon) + \log(1/\delta)}{\epsilon^2}\right)$, which concludes the proof. \square

Given that B and λ are constants the above bound simplifies to $\tilde{\mathcal{O}}\left(\Delta^2 \left(\frac{\Delta \cdot d + \log \frac{1}{\delta}}{\epsilon^2}\right)\right)$, which allows immediate comparison with the bounds derived in (Narasimhan et al., 2015). Additionally, when the degree of every node is constant (which is the case for real social networks like Facebook) or when the distribution \mathcal{D} activates only seeds of constant size, Δ is a constant and the sample complexity becomes $\tilde{\mathcal{O}}\left(\frac{d + \log \frac{1}{\delta}}{\epsilon^2}\right)$, independent of the size of the network.

4. Algorithms

As we mentioned in the previous section, the maximization problem (2) is non-concave and it cannot be solved efficiently. However, the cumulative log-likelihood function we aim to optimize has a great deal of structure we can utilize.

To understand this structure, note that there are only three distinct cases for a sample $s = ((X, v), y)$ in the training set S : (i) node v was not influenced, (ii) node v was influenced and there is only one neighbor of v in X and (iii) node v was influenced and there is more than one neighbor of v in X . The only case that makes the respective log-likelihood non-concave is (iii) since we are unable to identify which of the parents of v actually influenced it and how to update the hyperparameter (equation (1) yields that formally). We refer to such samples as *obfuscated*.

One can partition S into S_o and $S \setminus S_o$ where S_o contains the obfuscated samples. We can then write $\tilde{f}(\theta) := \frac{1}{m} \sum_{s \in S} \mathcal{L}(s, \theta) = \frac{1}{m} \sum_{s \in S \setminus S_o} \mathcal{L}(s, \theta) + \frac{1}{m} \sum_{s \in S_o} \mathcal{L}(s, \theta) =: f(\theta) + \xi(\theta)$. Optimizing \tilde{f} can be perceived as optimizing a concave function f under noise ξ . The magnitude of the noise depends on the probability of seeing an obfuscated sample, which characterizes the difficulty of the optimization problem and can be computed in simple cases (see e.g. Lemma 8 in Appendix F).

There are three distinct approaches that we can follow:

1. Ignore the obfuscated samples and optimize f instead of \tilde{f} , using standard methods like Gradient Descent. The fact that the likelihood of each sample is bounded (Lemma 3 in Appendix E) will assure that the recovered solution will approximately optimize \tilde{f} as well.
2. Optimize \tilde{f} directly by applying techniques from (Belloni et al., 2015) for concave optimization under noise.
3. Attempt to optimize \tilde{f} using standard concave optimization techniques (for example Stochastic Gradient Descent (SGD) which is widely used in the training of deep networks, a non-convex optimization problem).

The first two methods provide theoretical guarantees for noise of small magnitude, if the noise is large however, they can lead to large error. See Appendix F for a detailed description of these two approaches. The third heuristic approach works remarkably well in practice, even when the noise is large, as the experiments of Section 5.1 demonstrate. Additionally, in Section 5.2 we include experiments indicating that, even if the noise is small, it is still in our best interest to utilize all the available samples since the shortage of the training set hurts us more than non-concavity.

5. Experiments

We conduct two sets of experiments. First, using synthetic datasets we show that if the hyperparametric assumption holds in a network, we can accurately learn the edge probabilities despite the non-concavity of (2), and significantly outperform methods that do not include information about the node features, for small training sets⁴. We also investigate which properties of the network and the model affect the convergence rate. Secondly, we validate our approach using real Facebook data, by showing that low-dimensional hyperparametric models are predictive in practice.

5.1. Learning the Diffusion Probabilities

Real Graphs: We also use the “ego-facebook”, “wiki-Vote”, “bitcoin-otc” and “bitcoin-alpha” datasets from (Leskovec & Krevl, 2014), which are publicly available real-world social networks, enabling the reproducibility of our experiments.

Graphs. We synthetically generate the social graph and the hyperparameter that determines the diffusion probabilities.

Synthetic Graphs: We simulate a social network using standard graph models. Since different models yield graphs with different topology, we selected four of the widely used ones: Barabási-Albert, Kronecker, Erdős-Rényi and the con-

⁴Notice that learning the diffusion probabilities allows us to compute other quantities of interest as well, like the probability of a node becoming influenced, the final size of a cascade initiated from a given set, or the influence function.

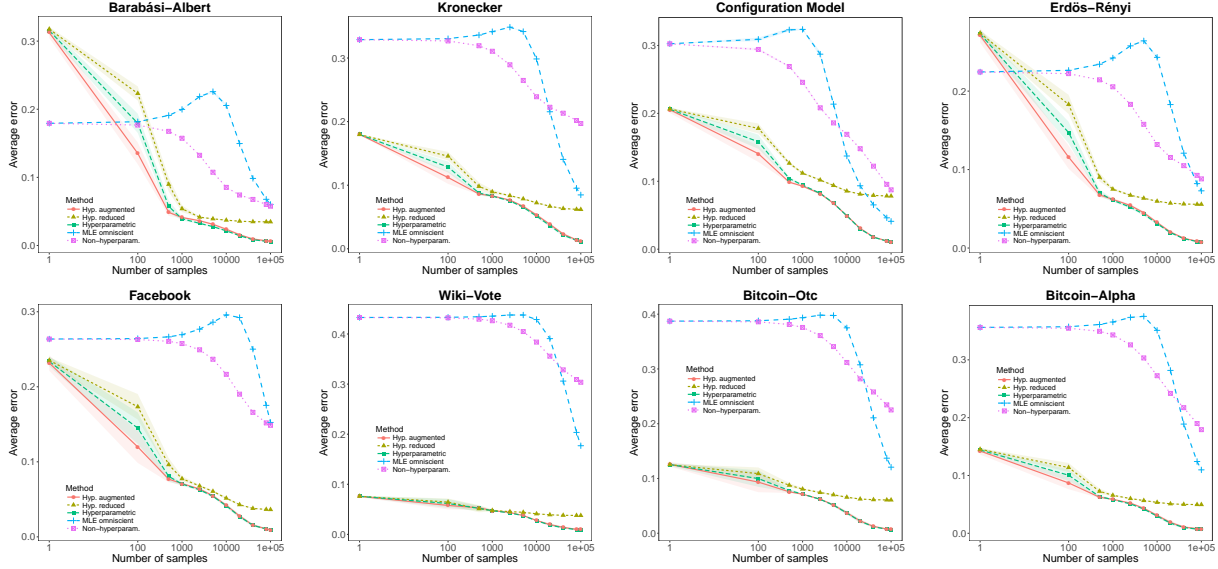


Figure 3. Learning the diffusion probabilities with vs without the hyperparametric assumption.

figuration model. For a more detailed description of these models and the construction process refer to Appendix G.

Experimental setup. We generate 15 random features in $[0, 1]$ for every node (we found consistent results across a large range of features). We define the first 10 of them to be the ones in which the hyperparametric assumption is based, and the rest is redundant information (30 features for each edge, where only 20 of them are important). Additionally, we generate a random hyperparameter in $[-1, 1]^d$ ($d = 20$). We use the sigmoid function over the important features of each edge e and the hyperparameter to compute p_e , as described in Section 2, imposing correlation between the probabilities of different edges by construction.

Subsequently, we generate 100,000 samples, and attempt to solve the optimization problem (2) using SGD, initializing the hyperparameter to 0 and using a learning rate of $1/\sqrt{T}$, where T is the size of the training set. Details on how we create the training set can be found in Appendix G.

Benchmarks. We tested the hyperparametric model against the following benchmarks:

- *Omniscient MLE*: The true diffusion probability of an edge is approximated by $\hat{p}_e = n_e^+ / n_e$, where n_e^+ is the number of activations of edge e , while n_e is the total number of exposures of e (activation attempts). Here, we assume that for every sample $((X, v), y)$ we observe the activation or not of all the edges $e = (u, v)$, where u is an active neighbor of v . This is a strong benchmark since in practice, we can observe whether v became active but not which node activated it.
- *Non-hyperparametric MLE*: We implemented the algorithm of (Narasimhan et al., 2015), that allows one

to learn the diffusion probabilities only by observing whether a node v was influenced by the seed X or not.

- *Hyperparametric MLE, reduced information*: We compare ourselves against a hyperparametric model that is unaware of the exact features that are important, and selects only a subset of them. Here we select only 5 out of 10 important features of every node.
- *Hyperparametric MLE, augmented information*: Similarly, we compare ourselves against a hyperparametric model that is unaware of the important features, thus it selects all the available ones (15 features per node).

Results. We repeat each experiment 10 times, and provide the mean and the standard deviation in Figure 3. The y -axis corresponds to $\frac{1}{|E|} \sum_{e \in E} |p_e - \hat{p}_e|$, the average absolute error between the real probability (known by construction) and the empirical one across the network. In all networks, the hyperparametric approach greatly outperforms the non-hyperparametric benchmarks, *even assuming omniscience*.

Note that in the non-hyperparametric benchmarks, since samples do not carry global information, there exist edges that have no exposures given the samples that we have seen so far. In that case, we define $\hat{p}_e = 0$. This explains the initial increase in the error in the omniscient MLE since, if p_e is small and the first exposure of edge e is an activation, the error on e increases from p_e to $1 - p_e$. Once we see enough samples, \hat{p}_e converges to p_e . One can also notice, the effect of omniscience since it leads to faster convergence than actual implementable non-hyperparametric methods.

Regarding the two benchmarks that involve the hyperparametric assumption it is worth noting that reduced information does not allow convergence to 0 error, while augmented

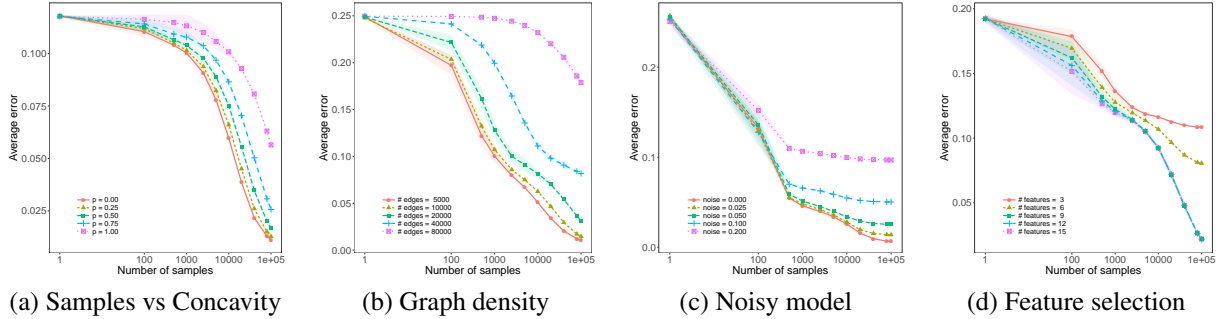


Figure 4. Characteristics of the graph or the model that affect convergence.

information does (see also Figure 4d for a more detailed investigation). Finally, the initial difference in the errors of different networks is related to how good predictor the initialization of SGD is (i.e. $\theta = 0$), as well as how large the average diffusion probability in the network is. The learning effect is evident and universal though, since the error converges to 0 independently of the underlying network.

5.2. Convergence Rate Investigation

In these experiments we use Erdős-Rényi graphs with 1000 nodes and 20000 edges (unless otherwise stated).

Samples vs Concavity. In Section 4 we classified the samples into categories based on whether they yield concave log-likelihood or not. Recall that if we ignore the obfuscated samples, the optimization problem becomes concave. A natural question is whether sacrificing samples for concavity leads to faster convergence. To this end, we generate samples and if a sample is obfuscated, we discard it with probability $p \in \{0.0, 0.25, 0.5, 0.75, 1.0\}$. The results can be found in Figure 4a. It is evident that even though our optimization problem becomes concave and hence theoretically easier to solve, the price due to data shortage is huge.

Approximate models. Here we investigate how properties of the graph or the model affect the convergence rate.

- *Graph Density:* We create an Erdős-Rényi graph with 1000 nodes and varying number of edges, exploring how does graph density affect the convergence of the error. As the density increases, so does the average degree in the network and, as a result, the number of obfuscated samples. Hence, the information obtained becomes “noisier” and convergence is slower.
- *Noisy model:* Until now we assumed that the hyperparametric model is the ground truth. Here we relax this assumption. We generate each edge probability as before, and subsequently add noise uniform in $[-N, N]$, for increasing values of N . Now, the average error does not converge to 0 and increases with N .
- *Features Effect:* In many cases we might not know the exact features that support the hyperparametric model.

We explore the effect of this lack of information by including varying number of significant features in our model. Our results show that in terms of convergence more information does not hurt, despite being more costly computationally. However, if we fail to include all the significant features, we do not converge to 0 error, and the error grows with the removal of features.

The results can be found in Figure 4b-d. In all the cases the way that we enforced the hyperparametric assumption, created the samples and ran SGD is the same as in Section 5.1.

5.3. Are Low-dimensional Models Predictive?

Importantly, we evaluate the validity of the hyperparametric assumption on real cascade data. To this end, we use the following aggregated and public Facebook data sets containing only de-identified data (i.e. they don’t include personally identifying information about individuals in the dataset).

Events. In Facebook a user can invite a set of other users for an event, who can then forward the invite to their friends to join. Moreover, when friends join an event a user may be notified in their Facebook feed and join in turn. The *cascade* in this scenario is an event and an *exposure* is either a direct invite or a feed notification. A user is influenced if she marked herself as “going” to the event. This dataset is a random sample of events that happened over a two-month period in late 2017. We have included only *public events* that are visible to everyone and also excluded users who created the event or joined without being invited. The number of cascades in our dataset is roughly 3 million with 90 million users participating and 130 million exposures.

Video and Photo Reshares. A *cascade* in this dataset is a video or photo content. Whenever a user watches a video (photo) posted from a friend we consider it as an *exposure* and when the user shares that video (photo) after watching we consider it as an *adoption* (we consider only videos that were explicitly seen and not auto-played). We collected a random sample of photo/video data on a random day in January 2018, and included only public photo and video

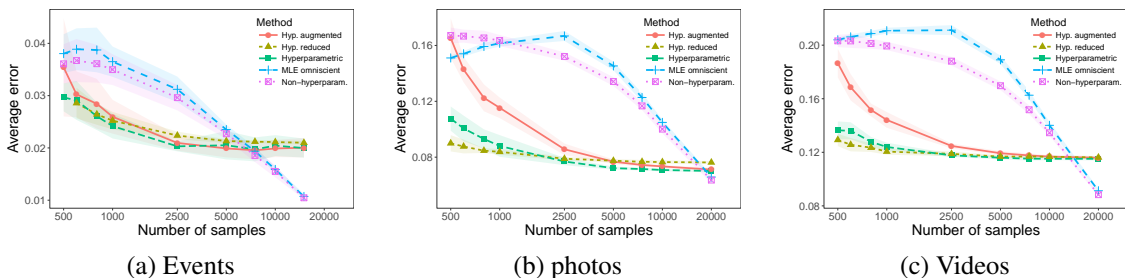


Figure 5. Learning with the hyperparametric assumption in Facebook cascade data sets.

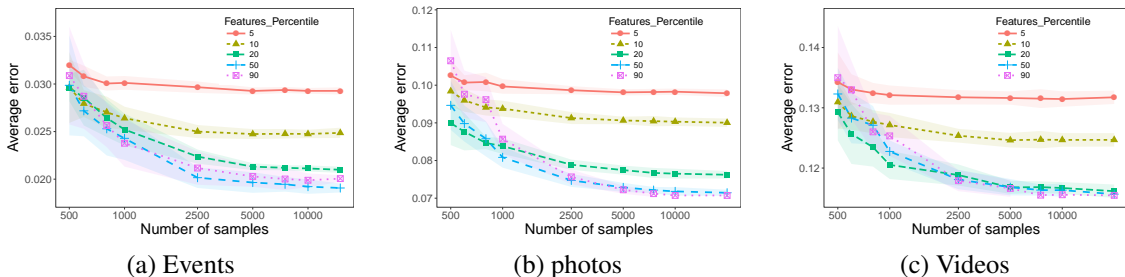


Figure 6. Feature sensitivity analysis of hyperparametric model for Facebook data sets.

posts. Our data sets contain roughly 10 million cascades with more than 100 million users and 500 million exposures.

Experimental setup. The features that we include in our model in both cases are user attributes such as Facebook age in days, friend count, number of initiated friendship requests, subscriber count and subscription count, city, country, and language, number of days active in last 7 days, and 28 days. The categorical features were binarized in the model. An important difference with the experiments of Sections 5.1 and 5.2 is that here we don’t know the true diffusion probability of every edge by construction. Instead, we estimate it from samples as $\hat{p}_e = \frac{n_e^+}{n_e}$. In order to estimate \hat{p}_e accurately, we need enough samples for edge e . Hence, we restrict our evaluation set (the set of edges where we measure the error) only to edges that have at least 67 interactions, meaning that $|p_e - \hat{p}_e| \leq 0.15$ with probability at least 90%. Each experiment is repeated 50 times and the averages together with the standard deviations are reported in Figures 5 and 6.

Results. Our first set of experiments is to validate the hyperparametric assumption in real data. We observe that using the optimization problem (2), with very few samples, the hyperparametric model achieves significant reduction in average error (up to 60%) over methods that don’t utilize node features. The results, reported in Fig. 5, are consistent with our synthetic experiments, where the hyperparametric assumption holds by construction. Note that the non-hyperparametric methods will eventually converge to zero error as they correspond to the ground truth while the hyperparametric model is only a good approximation of it.

We also included reduced and augmented hyperparametric models for comparison as in the synthetic experiments. In the case of the reduced model we used only 20% of the most important features of each edge (measured using Mutual Information). For the augmented version on the other hand, we augment the feature vector of each node with redundant information (increase its dimension by 50% and fill the extra coordinates with random noise) and investigate whether convergence still occurs. As in the experiments of Section 5.1, the reduced model converges to higher average error than the models that use more information, while the augmented model successfully ignores all the redundant features.

We also evaluated the sensitivity of the hyperparametric model when we include all versus few selected features. The picture that we see matches the synthetic experiments (Figure 4d), i.e. the hyperparametric model is supported on several features and if we fail to include all of them our error won’t converge to 0. However, an important difference with the synthetic experiments is that here not all the features are equally important, hence by applying feature-selection algorithms we can collect a small subset that performs almost as well as using the entire feature vector (see e.g. the difference in the error between 20% and 90% of the features).

6. Acknowledgements

This research was supported by NSF grant CAREER CCF-1452961, BSF grant 2014389, NSF USICCS proposal 1540428, and a Facebook research award.

References

- Abrahao, Bruno D., Chierichetti, Flavio, Kleinberg, Robert, and Panconesi, Alessandro. Trace complexity of network inference. In *The 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2013, Chicago, IL, USA, August 11-14, 2013*, pp. 491–499, 2013.
- Anderson, Ashton, Huttenlocher, Daniel P., Kleinberg, Jon M., Leskovec, Jure, and Tiwari, Mitul. Global diffusion via cascading invitations: Structure, growth, and homophily. In *Proceedings of the 24th International Conference on World Wide Web*, pp. 66–76, 2015.
- Belloni, Alexandre, Liang, Tengyuan, Narayanan, Hariharan, and Rakhlin, Alexander. Escaping the local minima via simulated annealing: Optimization of approximately convex functions. In *Proceedings of The 28th Conference on Learning Theory*, pp. 240–265, 2015.
- Bourigault, Simon, Lagnier, Cedric, Lamprier, Sylvain, Denoyer, Ludovic, and Gallinari, Patrick. Learning social network embeddings for predicting information diffusion. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pp. 393–402, 2014.
- Du, Nan, Liang, Yingyu, Balcan, Maria, and Song, Le. Influence function learning in information diffusion networks. In *International Conference on Machine Learning*, pp. 2016–2024, 2014.
- Friggeri, Adrien, Adamic, Lada, Eckles, Dean, and Cheng, Justin. Rumor cascades. In *International AAAI Conference on Web and Social Media*, 2014.
- Gomez Rodriguez, M, Balduzzi, D, Schölkopf, B, Scheffer, Getoor T, et al. Uncovering the temporal dynamics of diffusion networks. In *28th International Conference on Machine Learning (ICML 2011)*, pp. 561–568. International Machine Learning Society, 2011.
- Gomez Rodriguez, Manuel, Leskovec, Jure, and Krause, Andreas. Inferring networks of diffusion and influence. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1019–1028. ACM, 2010.
- Goyal, Amit, Bonchi, Francesco, and Lakshmanan, Laks VS. Learning influence probabilities in social networks. In *Proceedings of the third ACM international conference on Web search and data mining*, pp. 241–250. ACM, 2010.
- Kempe, David, Kleinberg, Jon, and Tardos, Éva. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146. ACM, 2003.
- Lazarsfeld, Paul and Merton, Robert K. Friendship as a social process: A substantive and methodological analysis. Morroe Berger, Theodore Abel, and Charles H. Page, editors, *Freedom and Control in Modern Society*.
- Lemonnier, Rémi, Scaman, Kevin, and Vayatis, Nicolas. Tight bounds for influence in diffusion networks and application to bond percolation and epidemiology. In *Advances in Neural Information Processing Systems*, pp. 846–854, 2014.
- Leskovec, Jure and Krevl, A. Snap datasets: Stanford large network dataset collection. 2014. URL <http://snap.stanford.edu/data>.
- Leskovec, Jure, Chakrabarti, D, Kleinberg, Jon, and Faloutsos, Christos. Realistic, mathematically tractable graph generation and evolution, using kronecker multiplication. In *PKDD 2005, volume 3721 of Lecture Notes in Computer Science*, pp. 133–145. Springer, 2005.
- Liben-Nowell, David and Kleinberg, Jon M. The link prediction problem for social networks. In *Proceedings of the 2003 ACM CIKM International Conference on Information and Knowledge Management, New Orleans, Louisiana, USA, November 2-8, 2003*, pp. 556–559, 2003.
- McPherson, Miller, Smith-Lovin, Lynn, and Cook, James M. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 2001.
- Narasimhan, Harikrishna, Parkes, David C, and Singer, Yaron. Learnability of influence in networks. In *Advances in Neural Information Processing Systems*, pp. 3186–3194, 2015.
- Netrapalli, Praneeth and Sanghavi, Sujay. Learning the graph of epidemic cascades. In *SIGMETRICS*, 2012.
- Saito, Kazumi, Nakano, Ryohei, and Kimura, Masahiro. Prediction of information diffusion probabilities for independent cascade model. In *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*, pp. 67–75. Springer, 2008.
- Shalev-Shwartz, Shai and Ben-David, Shai. *Understanding machine learning: From theory to algorithms*. 2014.
- Subbian, Karthik, Prakash, B Aditya, and Adamic, Lada. Detecting large reshare cascades in social networks. In *Proceedings of the 26th International Conference on World Wide Web*, pp. 597–605, 2017.
- Vaswani, Sharan, Kveton, Branislav, Wen, Zheng, Ghavamzadeh, Mohammad, Lakshmanan, Laks V. S., and Schmidt, Mark. Diffusion independent semi-bandit influence maximization. In *arXiv preprint*, pp. arXiv:1703.00557, 2017.
- Wen, Zheng, Kveton, Branislav, and Ashkan, Azin. Efficient learning in large-scale combinatorial semi-bandits. In *International Conference on Machine Learning*, 2015.