

Learnability of Competitive Threshold Models

Yifan Wang and Guangmo Tong

Department of Computer and Information Sciences, University of Delaware, USA

{yifanw, amotong}@udel.edu

Abstract

Modeling the spread of social contagions is central to various applications in social computing. In this paper, we study the learnability of the competitive threshold model from a theoretical perspective. We demonstrate how competitive threshold models can be seamlessly simulated by artificial neural networks with finite VC dimensions, which enables analytical sample complexity and generalization bounds. Based on the proposed hypothesis space, we design efficient algorithms under the empirical risk minimization scheme. The theoretical insights are finally translated into practical and explainable modeling methods, the effectiveness of which is verified through a sanity check over a few synthetic and real datasets. The experimental results promisingly show that our method enjoys a decent performance without using excessive data points, outperforming off-the-shelf methods.

1 Introduction

Social contagion phenomena, such as the propagation of behavior patterns or the adoption of new technologies, have attracted huge research interests in various fields (e.g., sociology, psychology, epidemiology, and network science [Cozzo *et al.*, 2013; Goldstone and Janssen, 2005; Camacho *et al.*, 2020]). Formal methods for modeling social contagion are crucial to many applications, for example, the recommendation and advertising in viral marketing [Kempe *et al.*, 2003; Tong *et al.*, 2018] and misinformation detection [Budak *et al.*, 2011; Tong, 2020]. Given the operational diffusion models, there is a fundamental question regarding the learnability of the models: to which extent can we learn such models from data?. In this paper, we study the classic threshold models in which the core hypothesis is that the total influence from active friends determines the contagion per a threshold. We seek to derive its PAC learnability and sample complexity by which efficient learning algorithms are possible.

Motivation and Related Works. Learning contagion models from data has been a central topic in social computing. One research branch focuses on estimating the total influence resulting from an initial status [Li *et al.*, 2017;

Gomez-Rodriguez *et al.*, 2016; Tang *et al.*, 2009]; in contrast, our learning task attempts to predict the diffusion status for each node. For node status prediction, it has been proved that deep learning methods are promising when rich features are available (e.g., [Qiu *et al.*, 2018; Leung *et al.*, 2019]), but their generalization performance in theory is often unclear. Recent works have demonstrated the sample complexity of influence learning for single-cascade models [He *et al.*, 2016; Du *et al.*, 2014], where certain technical conditions are required and the approach therein is different from ours in a substantial way. Our research is inspired by the recent works that show the possibility of simulating single-cascade contagion models through neural networks [Narasimhan *et al.*, 2015; Adiga *et al.*, 2019]. Departing from the common surrogate modeling methods where the learned mappings are less transparent, the pursued idea seeks to design neural networks with their layer structures mimicking real networks. In this paper, we take a step towards the case of general threshold models, without limiting the number of cascades. Comparing to the single-cascade case where threshold layers are sufficient, multi-cascade diffusion requires more sophisticated designs in order to properly propagate influence summation between layers without losing the information of cascade identity.

Contribution. Our work can be summarized as follows.

- **Model Design.** For a finite-precision system, we present elementary designs showing that the general threshold model can be simulated explicitly by neural networks with piecewise polynomial activation functions.
- **Theoretical Analysis.** Based on our hypothesis design, we prove that the general threshold model is PAC learnable, and design efficient empirical risk minimization schemes with a provable sample complexity.
- **Empirical Studies.** We experimentally examine the proposed methods over synthetic and real datasets. The results suggest that methods based on explicit simulation outperform off-the-shelf learning methods by an evidence margin. They also verify that our methods work in the way it supposed to.

The proofs, full experimental results, and source code are provided in the supplementary material¹.

¹<https://github.com/cdslabamotong/LTInfLearning>

2 Preliminaries

2.1 Competitive Linear Threshold Model

We follow the standard competitive linear threshold (CLT) model [Borodin *et al.*, 2010], which has been widely adopted in the formal study of social contagions [He *et al.*, 2012; Tzoumas *et al.*, 2012; Bozorgi *et al.*, 2017]. A social network is represented as a directed graph $G = (V, E)$, where $V = \{v_1, \dots, v_N\}$ is the set of nodes and E is the set of edges, with $N = |V|$ and $M = |E|$ being the numbers of nodes and edges, respectively. For each node $v \in V$, we denote by $N(v) \subseteq V$ the set of its in-neighbors.

We consider the scenario where there are $S \in \mathbb{Z}^+$ information cascades, each of which starts to spread from their seed set $\psi_0^s \subseteq V$ for $s \in [S]$. Associated with each node v_i and each cascade s , there is a threshold $\theta_i^s \in [0, 1]$; associated with each edge $(v_i, v_j) \in E$ and each cascade s , there is a weight $w_{i,j}^s \in [0, 1]$. Without loss of generality, we assume that the seed sets are disjoint and the weights are normalized such that $\sum_{v_j \in N(v_i)} w_{j,i}^s \leq 1$ for each node $v_i \in V$ and cascade s . The nodes are initially inactive and can be s -active if activated by cascade s . In particular, the diffusion process unfolds step by step as follows:

- **Step 0:** For each cascade $s \in [S]$, nodes in ψ_0^s become s -active.
- **Step $t > 0$:** Let $\psi_{t-1}^s \subseteq V$ be the set of s -active nodes after $t - 1$ step. There are two phases in step t .
 - **Phase 1:** For an inactive node $v_i \in V$, let ζ_i^s be the summation of the weights from its s -active in-neighbors:

$$\zeta_i^s = \sum_{v_j \in N(v_i) \cap \psi_{t-1}^s} w_{j,i}^s. \quad (1)$$

The node v_i is activated by cascade s if and only if $\zeta_i^s \geq \theta_i^s$. After phase 1, it is possible that a node can be activated by more than one cascades.

- **Phase 2:** For each node v_i , it will be finally activated by the cascade s^* with the largest weight summation:

$$s^* = \arg \max_{s: \zeta_i^s \geq \theta_i^s} \zeta_i^s. \quad (2)$$

For tie-breaking, cascades with small indices are preferred.

Clearly, there are at most n diffusion steps, and without loss of generality, we may assume that there are always n diffusion steps. An illustration is given in Figure 1. The following notations are useful for formal discussions.

Definition 1. We use a binary matrix $\mathbf{I}_0 \in \{0, 1\}^{N \times S}$ to denote the initial status, where $\mathbf{I}_{0,i,s} = 1$ if and only if node v_i is in the seed set ψ_0^s of cascade s . For a time step $t > 0$, the diffusion status is denoted by a binary tensor $\mathbf{I}_t \in \{0, 1\}^{N \times S \times 2}$, where, for a cascade s , $\mathbf{I}_{t,i,s,p} = 1$ if and only if node v_i is s -active after phase $p \in \{1, 2\}$ in time step t .

We are interested in the learnability of the influence function of CLT models.




Figure 1: Diffusion process example. In this example, we have two cascades with seed sets $\{v_2\}$ and $\{v_1\}$. The weights and thresholds are given as graph shows. According to the CLT model, v_3 will become 1-active after one time step. v_4 will become 1-active at the end of diffusion process.

Definition 2. Given a CLT model, the influence function $F : \{0, 1\}^{N \times S} \mapsto \{0, 1\}^{N \times S}$ maps from the initial status \mathbf{I}_0 to the final status $\mathbf{I}_{N,::,2}$.

2.2 Learning Settings

Assuming that the social graph G is unknown to us, we seek to learn the influence function using a collection D of $K \in \mathbb{Z}^+$ pairs of initial status and corresponding diffusion statuses:

$$D = \left\{ (\mathbf{I}_0^1, \mathbf{I}_{1:N,::,2}^1), \dots, (\mathbf{I}_0^K, \mathbf{I}_{1:N,::,2}^K) \right\} \quad (3)$$

where $\mathbf{I}_{1:N,::,1:2}^k$ denotes the diffusion status after each phase in each diffusion step:

$$\mathbf{I}_{1:N,::,2}^k = \left[\mathbf{I}_{1,::,1}^k, \mathbf{I}_{1,::,2}^k, \mathbf{I}_{2,::,1}^k, \mathbf{I}_{2,::,2}^k, \dots, \mathbf{I}_{N,::,1}^k, \mathbf{I}_{N,::,2}^k \right].$$

We use the zero-one loss ℓ_{0-1} to measure the difference between the prediction $\hat{\mathbf{I}}_{N,::,2}$ and the ground truth $\mathbf{I}_{N,::,2}$, which is defined as

$$\ell_{0-1}(\mathbf{I}_{N,::,2}, \hat{\mathbf{I}}_{N,::,2}) := \frac{1}{2NS} \sum_{i \in [N]} \sum_{s \in [S]} \mathbb{1}[\mathbf{I}_{N,i,s,2} \neq \hat{\mathbf{I}}_{N,i,s,2}]$$

For a distribution \mathcal{D} over the input \mathbf{I}_0 and a CLT model generating the training data D , we wish to leverage D to learn a mapping $\hat{F} : \{0, 1\}^{N \times S} \mapsto \{0, 1\}^{N \times S}$ such that the generalization loss can be minimized:

$$L(\hat{F}) := \mathbb{E}_{\mathbf{I}_0 \sim \mathcal{D}} \left[\ell_{0-1}(F(\mathbf{I}_0), \hat{F}(\mathbf{I}_0)) \right]. \quad (4)$$

3 Learnability Analysis

In this section, we present a PAC learnability analysis for the influence function. The following assumptions will be made throughout the analysis.

Assumption 1. Following the conventions regarding finite precision system [Holi and Hwang, 1993; Weiss *et al.*, 2018; Dembo *et al.*, 1989], we assume that the diffusion model is a decimal system of a finite precision: for the involved parameters (i.e., $w_{i,j}^s$ and θ_i^s), number of their decimal places is upper bounded by a constant $Q \in \mathbb{Z}^+$.

Assumption 2. The cascade number $S = 2^Z$ is a power of 2 for some constant $Z \in \mathbb{Z}^+$.

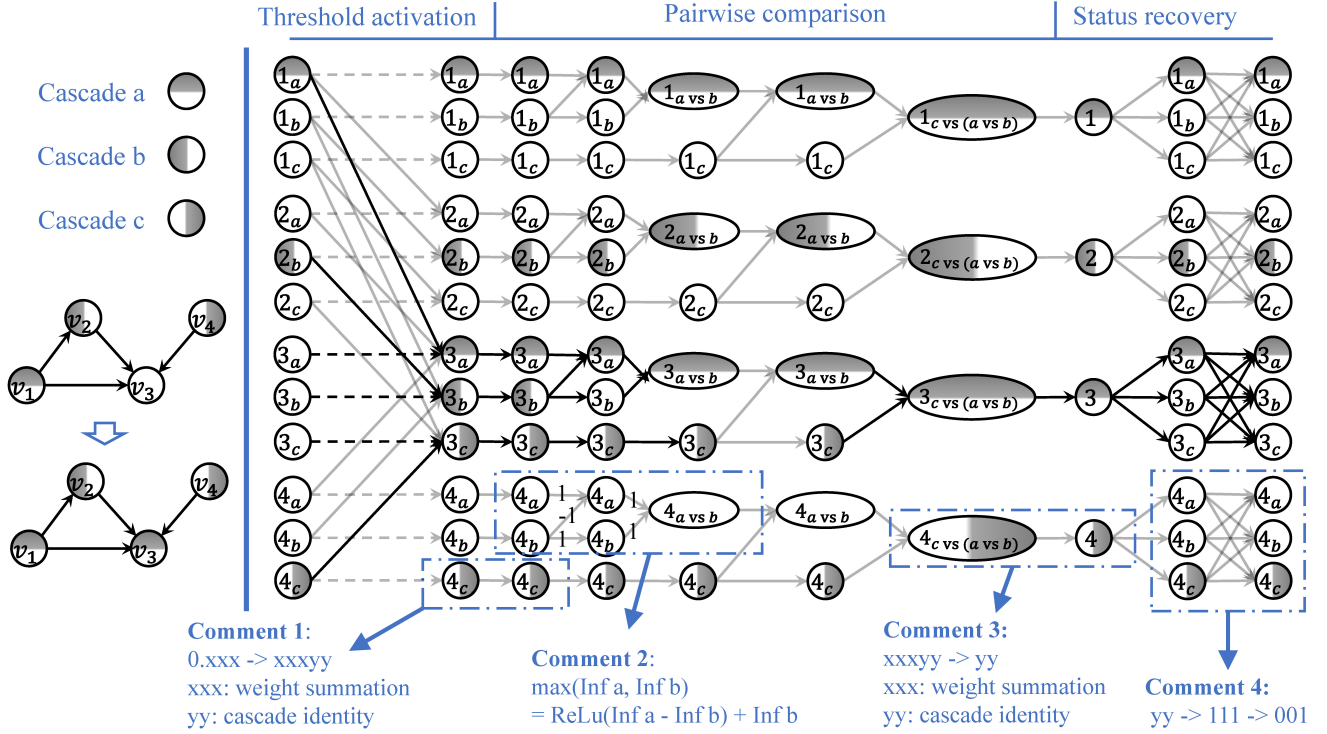


Figure 2: **One-step Simulation.** The left part depicts the scenario where the node v_3 is activated by three cascades among which cascade a finally wins the competition. The right parts demonstrates how to simulate such a scenario through three our design; the computation path for node v_3 is highlighted.

The overall idea is to design a realizable hypothesis with a finite VC dimension by which the PAC learnability can be established. In order to obtain a realizable hypothesis space, we seek to explicitly simulate the diffusion function through neural networks. While the universal approximation theory ensures that such neural networks exist, we wish for elementary designs of a polynomial size so as to admit efficient learning algorithms. To this end, we first demonstrate how to simulate the one-step diffusion under the CLT model, which can be used repeatedly to simulate the entire diffusion.

3.1 Realizable Hypothesis Space

For the one-step diffusion, its simulation is done through two groups of layers that are used in turn to simulate phase 1 and phase 2. As illustrated in Figure 2, the simulation starts by receiving the diffusion status $\mathbf{H}^1 \in \{0, 1\}^{N \cdot S}$ after the last step, where for $i \in [N]$ and $s \in [S]$, $\mathbf{H}^1_{(i-1) \cdot S + s} = 1$ if and only if node v_i is s -active. After $2Z + 4$ layers of transformations, it ends by outputting a vector $\mathbf{H}^{2Z+6} \in \{0, 1\}^{N \cdot S}$, which corresponds exactly to the diffusion status after one diffusion step under the CLT model. The layer transformations all follow the canonical form:

$$\mathbf{H}^{l+1} = \sigma^{l+1}(\mathbf{W}^l \mathbf{H}^l) \quad (5)$$

with \mathbf{W}^l being the weight matrix and σ^l being a collection of activation functions over the elements of $\mathbf{W}^l \mathbf{H}^l$. In the rest of this subsection, we will present our designs of the weight matrices and the activation functions that can fulfill the goal of simulation.

Phase 1

The single-cascade activation can be simulated by

$$\mathbf{H}^2 = \sigma^1(\mathbf{W}^1 \mathbf{H}^1), \quad (6)$$

where for each $i, j \in [N]$ and $s_1, s_2 \in [S]$, we have

$$\mathbf{W}^1_{(i-1) \cdot S + s_1, (j-1) \cdot S + s_2} := \begin{cases} w_{i,j}^{s_1} & \text{if } (v_i, v_j) \in E \text{ and } s_1 = s_2 \\ 1 & \text{if } i = j \text{ and } s_1 = s_2 \\ 0 & \text{otherwise} \end{cases},$$

and for each $i \in [N]$ and $s \in [S]$, we have

$$\sigma^1_{(i-1) \cdot S + s}(x) := \begin{cases} x & \text{if } x \geq \theta_i^s \\ 0 & \text{Otherwise} \end{cases}. \quad (7)$$

For each $i \in [N]$ and $s \in [S]$, the resulting $\mathbf{H}^2_{(i-1) \cdot S + s}$ is equal to the influence summation of cascade s at node v_i if v_i is activated by cascade s , and otherwise zero.

Phase 2

In simulating phase 2, the general goal is to figure out among the candidate cascades which one wins the competition. Intuitively, this can be implemented through a hierarchy of pairwise comparisons. One technical challenge is that comparison made directly by linear functions tells only the difference between the summations, while we need to keep the summation for future comparisons; we observe that this problem can be solved by using two consecutive layers of piecewise linear functions. A more serious problem is that simply passing

the largest influence summation to subsequent layers incurs the loss of the information about the cascade identity (i.e., index), making it impossible to identify the node status – which is required as the input to simulating the next diffusion step. We propose to address such a challenge by appending the cascade identity to the weight summation as the lowest bits. The total shifted bits are determined by the system precision (Assumption 1). Doing so does not affect the comparison result while allowing us to retrieve the cascade identity by the scale and modulo operations. In particular, two groups of transformations are employed to achieve the pairwise comparison and status recovery.

Pairwise Comparison. With the input \mathbf{H}^2 from phase 1, we encode the cascade identity through

$$\mathbf{H}^3 = \sigma^2(\mathbf{W}^2 \mathbf{H}^2), \quad (8)$$

where for each $i \in [N \cdot S]$ and $j \in [N \cdot S]$, we have

$$\mathbf{W}_{i,j}^2 := \begin{cases} 10^{(Q + \lceil \log_{10} S \rceil + 1)} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}, \quad (9)$$

and for each $i \in [N]$ and $s \in [S]$, we have

$$\sigma_{(i-1) \cdot S + s}^2(x) = x + s. \quad (10)$$

The design of \mathbf{W}^2 ensures that sufficient bits are available at the low side, and the activation function σ^2 then writes the cascade identity into the empty bits (Comment 1 in Figure 2).

Using the information in \mathbf{H}^3 , for a node $v_i \in V$, the largest influence summation can be determined by the sub-array

$$[\mathbf{H}_{(i-1) \cdot S + 1}^3, \mathbf{H}_{(i-1) \cdot S + 2}^3, \dots, \mathbf{H}_{(i-1) \cdot S + S}^3]$$

through pairwise comparisons. For the comparison between two cascades s_1 and s_2 at node v_i ($s_1 < s_2$), as illustrated in Comment 2 in Figure 2, the difference

$$\mathbf{H}_{(i-1) \cdot S + s_1}^3 - \mathbf{H}_{(i-1) \cdot S + s_2}^3$$

is first fed into a ReLU function of which the result is then added by $\mathbf{H}_{(i-1) \cdot S + s_2}^3$. We can verify that this returns exactly

$$\max(\mathbf{H}_{(i-1) \cdot S + s_1}^3, \mathbf{H}_{(i-1) \cdot S + s_2}^3).$$

Therefore, two layers are used to eliminate a half of the candidate cascades, thereby resulting in $2Z$ layers in total. The output of this part is $\mathbf{H}^{2Z+3} \in \mathbb{R}^N$ in which we have

$$\mathbf{H}_i^{2Z+3} = \max(\mathbf{H}_{(i-1) \cdot S + 1}^3, \mathbf{H}_{(i-1) \cdot S + 2}^3, \dots, \mathbf{H}_{(i-1) \cdot S + S}^3)$$

for each $i \in [N]$.

Notably, the lowest Q bits in \mathbf{H}^{2Z+3} stores the cascade index, which can be recovered by

$$\mathbf{H}^{2Z+4} = \sigma^{2Z+3}(\mathbf{W}^{2Z+3} \mathbf{H}^{2Z+3}), \quad (11)$$

where \mathbf{W}^{2Z+3} is the identity matrix and we have

$$\sigma^{2Z+3}(x) := x \bmod 10^{\lceil \log_{10} S \rceil + 1}, \quad (12)$$

which is shown in Comment 3 in Figure 2.

Finally, the cascade indices in \mathbf{H}^{2Z+4} are converted to binary indicators through two layers. The first layer transforms \mathbf{H}^{2Z+4} into a binary vector in $\{0, 1\}^{N \cdot S}$ by

$$\mathbf{H}^{2Z+5} = \sigma^{2Z+4}(\mathbf{W}^{2Z+4} \mathbf{H}^{2Z+4}), \quad (13)$$

where for each $i, j \in [N]$ and $s \in [S]$, we have

$$\mathbf{W}_{i, (j-1) \cdot S + s}^{2Z+4} := \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}, \quad (14)$$

and

$$\sigma_{(i-1) \cdot S + s}^{2Z+4}(x) := \begin{cases} 1 & \text{if } x \geq s \\ 0 & \text{otherwise} \end{cases}. \quad (15)$$

This ensures that in each group

$$[\mathbf{H}_{(i-1) \cdot S + 1}^{2Z+5}, \mathbf{H}_{(i-1) \cdot S + 2}^{2Z+5}, \dots, \mathbf{H}_{(i-1) \cdot S + S}^{2Z+5}],$$

$\mathbf{H}_{(i-1) \cdot S + s}^{2Z+5}$ is 1 if and only if s is no larger than the index of the cascade that activates v_i . Finally, \mathbf{H}^{2Z+5} is transformed into the form that matches exactly the diffusion status after ones-step diffusion, which can be achieved through

$$\mathbf{H}^{2Z+6} = \sigma^{2Z+5}(\mathbf{W}^{2Z+5} \mathbf{H}^{2Z+5}), \quad (16)$$

where for each $i, j \in [N]$ and $s_1, s_2 \in [S]$, we have

$$\mathbf{W}_{(i-1) \cdot S + s_1, (j-1) \cdot S + s_2}^{2Z+5} := \begin{cases} 1 & \text{if } i = j \text{ and } s_1 \leq s_2 \\ -1 & \text{if } i = j \text{ and } s_1 = s_2 \\ 0 & \text{otherwise} \end{cases}, \quad (17)$$

and

$$\sigma_{(i-1) \cdot S + s}^{2Z+5}(x) := \begin{cases} 1 & \text{if } x \geq s \\ 0 & \text{otherwise} \end{cases}. \quad (18)$$

The last two layers are illustrated in Comment 4 in Figure 2.

Lemma 1. *One-step diffusion in a CLT model can be simulated by a feed-forward neural network composed of $O(Z)$ layers, with $O(S(N + M))$ adjustable weights and $O(N \cdot S)$ piecewise linear computation units each of which with $O(2^Q)$ pieces.*

Definition 3 (Hypothesis Space). Repeating such one-step simulations for N steps, the CLT model can be explicitly simulated. Taking the weights $w_{i,j}^s$ and the thresholds θ_i^s as parameters, we denote by \mathcal{F} the hypothesis space formed by such neural networks.

One important property is that the entire neural network is composed of piecewise polynomial activation functions, which will be used in deriving its sample complexity.

3.2 Efficient ERM

Theorem 1 suggests that for any sample set D , there always exists a perfect hypothesis in \mathcal{F} . In what follows, we show that such an empirical risk minimization solution can be efficiently computed. It is sufficient to find the parameters that can ensure the output of each activation function coincides with the diffusion result. Since the activation functions are all pairwise linear functions, the searching process can be done by linear programming with a polynomial number of constraints. Formally, we have the following statement.

Lemma 2. For a CLT model and each sample set D , there exists a hypothesis in \mathcal{F} with zero training error, and it can be computed in polynomial time in terms of N, M, S , and $|D|$.

3.3 Generalization Performance

To establish the learnability of class \mathcal{F} , we first analyze the VC dimension of the binary-valued function class $\mathcal{F}_{i,j}$ that is obtained by restricting the output of to the diffusion status of node v_i regarding cascade s in the last layer of the neural network. That is,

$$\mathcal{F}_{i,s} = \left\{ f_{i,s} : i \in [N], s \in [S], f_{i,s}(\mathbf{H}^1) = \mathbf{H}_{i,S+s}^{1+N(2Z+5)} \right\}.$$

Lemma 3. For each $i \in [N]$ and $s \in [S]$, the VC dimension of class $\mathcal{F}_{i,s}$ is $\text{VC}(\mathcal{F}_{i,s}) = \tilde{O}(S \cdot Q \cdot (N + M))$.

proof sketch. The proof is obtained by first showing that the functions in $\mathcal{F}_{i,s}$ have a manageable VC dimension when restricted to the one-step simulation. Given the fact that repeating such one-step simulations does not increase the model capacity in terms of shattering points, the VC dimension of the entire class $\mathcal{F}_{i,s}$ can be derived immediately. \square

Remark 1. The above suggests that the VC dimension is linear in the network size, the number of cascades, and the coding size of the numerical values. When there is only one cascade, pairwise comparisons and status recovery are no longer needed, and therefore, the size of the neural network does not depend on the system precision (i.e., Q). This gives a VC dimension of $\tilde{O}(N + M)$, which recovers the results in [Narasimhan *et al.*, 2015]. Interestingly, similar results apply to the case of $S = 2$ but not to larger S . We provide more details in Appendix B.

With Lemmas 2 and 3, the generalization bound of the ERM scheme for predicting the status of one node regarding one cascade can be derived through the class result of statistical learning theory, and the sample complexity of \mathcal{F} follows from the union bound over all the nodes and cascades.

Theorem 1. The influence function class \mathcal{F} is PAC learnable and the sample complexity for achieving a generalization error no larger than ϵ with probability at least $1 - \delta$ is $O\left(\frac{\text{VC}(\mathcal{F}_{i,s}) \log(1/\epsilon) + \log(NS/\delta)}{\epsilon}\right)$.

4 Experiment

In this section, we present empirical studies for experimentally examining the practical utility of the proposed method.

4.1 Experimental Settings

Graphs. We adopt two classic graph structures (Kronecker [Leskovec *et al.*, 2010] and power-law) and one real graph Higgs, which was built after monitoring the spreading process of the messages posted between 1st and 7th July 2012 regarding the discovery of a new particle with the features of the elusive Higgs boson [De Domenico *et al.*, 2013]. The graph statistics are provided in the supplementary material.

Diffusion Model. We simulate the case with the number of cascades S selected from $\{2, 3, 4\}$. The thresholds θ_i^s are

generated uniformly at random from $[0, 1]$. We consider two settings for weight generation, following the weighted cascade setting [Borodin *et al.*, 2010], the weights on Kronecker and Power-law are set as $w_{i,j}^s = \frac{1}{|N(v_j)|+s/7}$ for each pair $(v_i, v_j) \in E$ and each $s \in [S]$. On Twitter, the weights are sampled uniformly at random from $[0, 1]$ and then normalized such that $\sum_{v_i \in N(v_j)} w_{i,j}^s = 1$ for each $v_j \in V$. The parameters are set to be numbers with three decimal places.

Samples. To generate one sample of $(\mathbf{I}_0^1, \mathbf{I}_{1:N,::,2}^1)$, the total number of seed nodes $|\bigcup_{s=1}^S \psi_0^s|$ is generated from $[0.1N, 0.5N]$ uniformly at random, and then a subset of $|\bigcup_{s=1}^S \psi_0^s|$ nodes are randomly selected from V . For each selected node, it is assigned to one cascade uniformly at random. Given the initial status, the ground truth is generated following the diffusion model. For each graph, we generate a pool of 2,000 samples among which the average diffusion step is four and the average number of total active nodes is 254.

Methods. Given the samples, the parameters of our model are computed by the linear programming solver Cplex [Cplex, 2009]. Alternatively, one can use off-the-shelf supervised learning methods to train a one-step diffusion model on each node, and make predictions by repeatedly simulating the one-step diffusion process. Given the nature of binary classification, we implement three methods, logistic regression (LR), support vector machine (SVM), and multilayer perceptron (MLP), using scikit-learn [Pedregosa *et al.*, 2011]. For such baselines, the number of repetitions for testing is selected from $\{1, 2, 3, 4, 5\}$. We also implement the random method that makes a random prediction.

Training and Testing. The size of the training set is selected from $\{50, 100, 500\}$ and the testing size is 500. In each run, the training and testing samples are randomly selected from the sample pool, and we evaluate the performance of each method in terms of F1 score, precision, and accuracy. The entire process is repeated for five times, and we report the average performance together with the standard deviation.

4.2 Results

Major Observations

The main results are given in Table 1. We see from the table that all the prediction results are statistically robust per the standard deviation. The main observation is that our method outperforms other methods in a significant manner in terms of any of the measurements. MLP is better than LR and SVM, but still produces low-quality predictions on the power-law graph. In addition, for classic learning methods, the Kronecker graph is relatively easy to handle, while the predictions on the power-law graph are consistently less satisfactory, which suggests that they are sensitive to graph structures.

Minor Observations

Hidden Layers Meeting Diffusion Status. From the standpoint of surrogate modeling, it may be interesting to notice that the hidden features of the proposed neural networks meet exactly the node status during the diffusion process, which somehow makes our model explainable and more transparent.

Kronecker		F1 Score	Precision	Accuracy
Our Method		0.994(1E-3)	0.995(1E-4)	0.993(1E-4)
LR	1	0.499(1E-3)	0.676(2E-3)	0.709(2E-3)
	2	0.474(1E-3)	0.752(1E-2)	0.705(4E-3)
	3	0.460(3E-3)	0.802(1E-2)	0.705(6E-3)
	4	0.452(3E-3)	0.827(5E-3)	0.701(4E-3)
	5	0.451(3E-3)	0.832(6E-3)	0.701(4E-3)
SVM	1	0.556(2E-3)	0.716(2E-3)	0.727(3E-3)
	2	0.496(1E-3)	0.731(2E-3)	0.705(4E-3)
	3	0.480(3E-3)	0.770(9E-3)	0.702(5E-3)
	4	0.470(4E-3)	0.798(9E-3)	0.697(3E-3)
	5	0.465(3E-3)	0.818(6E-3)	0.698(6E-3)
MLP	1	0.909(1E-3)	0.975(1E-4)	0.935(1E-4)
	2	0.909(1E-3)	0.975(1E-4)	0.934(1E-3)
	3	0.909(1E-3)	0.975(1E-4)	0.935(1E-3)
	4	0.909(1E-3)	0.975(1E-3)	0.935(1E-3)
	5	0.909(1E-3)	0.975(1E-4)	0.934(1E-4)
Random		0.209(1E-2)	0.250(1E-2)	0.250(1E-2)

Table 1: Main results on Kronecker graph with 100 training samples. Each cell presents the average performance plus the standard deviation. Testing results for LR, SVM and MLP are provided for different numbers of iterations.

While the original goal is to predict the final diffusion status, we seek to examine to what extent the hidden features are semantically meaningful. To this end, we compare the hidden features with the node status in the corresponding steps, and the results are given in Figure 3. For our method, we in general observe that the one-step prediction is very accurate, while the match becomes less perfect with more steps. Therefore, given the fact that most of the samples contain no more than six diffusion steps, our method is overall effective.

Imbalance Class Distribution. Imbalance class distribution is an inherent issue in predicting the node status in that the majority of the nodes tend to be inactive, and therefore high accuracy does not necessarily imply good prediction performance. According to the results in Table 1, the standard methods often simply predict all the nodes as inactive, which gives a decent accuracy but low F1 score and precision. In such a sense, our method is more robust and effective.

Impact of Testing Iterations. For LR and MLP, one can see that repeating the single-step prediction does not help increase prediction quality. One plausible reason is that the single-step prediction is not sufficiently good and its error is amplified by the repetitions. For MLP, the results suggest that the prediction almost converges after the first single-step prediction.

Impact of the Number of Training Samples. Tables 5 and 6 in Appendix C presents the full results on Kronecker and power-law using different training sizes. We see that the performance increases very slightly when more samples are

Higgs	F1 Score	Precision	Accuracy
Our Method	0.999(1E-4)	0.999(1E-4)	0.998(2E-4)
LR	0.427(4E-3)	0.322(4E-3)	0.663(2E-3)
SVM	0.492(6E-3)	0.674(6E-3)	0.648(4E-3)
MLP	0.939(3E-3)	0.962(3E-3)	0.960(2E-3)
Random	0.451(4E-4)	0.450(3E-4)	0.450(4E-4)

Table 2: Main results on Higgs with the cascades using 100 training samples. The testing of LR, SVM, and MLP are performed with one iteration.




Figure 3: Results of step-step matching.

provided, suggesting that training size is not the main bottleneck.

Impact of the Number of Cascades. As our work focuses on the CLT model, we are interested in the performance under different numbers of cascades, and the results of this part can be found in Tables 3 and 4 in Appendix C. Our method exhibits a stable performance, while for other methods, the task becomes more challenging in the presence of more cascades. The above tables also present the running time of each method, which confirms that our proposed linear programming can be easily handled by standard software packages.

5 Further Discussions

In this paper, we demonstrate the possibility of design methods for simulating social contagions by using neural networks that directly mimic the diffusion pattern. We present a learnability study for the competitive threshold models as well as an efficient learning algorithm with a provable sample complexity. While our method outperforms classic supervised learning methods, it does not rule out the possibility that advanced representation techniques can achieve better learning performance with sufficient samples, which is one direction of our future work. In addition, there exist other classic operational models, such as the independent cascade model and voter model, and it remains unknown how to explicitly simulate them through neural networks, or more generally, graphical models. Given the encouraging experimental results we acquire in this paper, we believe it is worth investigating more models beyond the single-cascade case.

References

- [Adiga *et al.*, 2019] Abhijin Adiga, Chris J Kuhlman, Madhav Marathe, S Ravi, and Anil Vullikanti. Pac learnability of node functions in networked dynamical systems. In *Proc. of ICML*, 2019.
- [Anthony *et al.*, 1999] Martin Anthony, Peter L Bartlett, and Peter L Bartlett. *Neural network learning: Theoretical foundations*. Cambridge university press, 1999.
- [Borodin *et al.*, 2010] Allan Borodin, Yuval Filmus, and Joel Oren. Threshold models for competitive influence in social networks. In *International workshop on internet and network economics*. Springer, 2010.
- [Bozorgi *et al.*, 2017] Arastoo Bozorgi, Saeed Samet, Johan Kwisthout, and Todd Wareham. Community-based influence maximization in social networks under a competitive linear threshold model. *Knowledge-Based Systems*, 2017.
- [Budak *et al.*, 2011] Ceren Budak, Divyakant Agrawal, and Amr El Abbadi. Limiting the spread of misinformation in social networks. In *Proc. of WWW*, 2011.
- [Camacho *et al.*, 2020] David Camacho, Ángel Panizol-Lledot, Gema Bello-Orgaz, Antonio Gonzalez-Pardo, and Erik Cambria. The four dimensions of social network analysis: An overview of research methods, applications, and software tools. *Information Fusion*, 2020.
- [Cozzo *et al.*, 2013] Emanuele Cozzo, Raquel A Banos, Sandro Meloni, and Yamir Moreno. Contact-based social contagion in multiplex networks. *Physical Review E*, 2013.
- [Cplex, 2009] IBM ILOG Cplex. V12. 1: User’s manual for cplex. *International Business Machines Corporation*, 2009.
- [De Domenico *et al.*, 2013] Manlio De Domenico, Antonio Lima, Paul Mougél, and Mirco Musolesi. The anatomy of a scientific rumor. *Scientific reports*, 2013.
- [Dembo *et al.*, 1989] Amir Dembo, Kai-Yeung Siu, and Thomas Kailath. Complexity of finite precision neural network classifier. In *Proc. of NIPS*, 1989.
- [Du *et al.*, 2014] Nan Du, Yingyu Liang, Maria Balcan, and Le Song. Influence function learning in information diffusion networks. In *Proc. of ICML*. PMLR, 2014.
- [Goldstone and Janssen, 2005] Robert L. Goldstone and Marco A. Janssen. Computational models of collective behavior. *Trends in Cognitive Sciences*, 9:424–430, 2005.
- [Gomez-Rodriguez *et al.*, 2016] Manuel Gomez-Rodriguez, Le Song, Nan Du, Hongyuan Zha, and Bernhard Schölkopf. Influence estimation and maximization in continuous-time diffusion networks. *ACM Transactions on Information Systems (TOIS)*, 2016.
- [He *et al.*, 2012] Xinran He, Guojie Song, Wei Chen, and Qingye Jiang. Influence blocking maximization in social networks under the competitive linear threshold model. In *Proc. of ICDM*. SIAM, 2012.
- [He *et al.*, 2016] Xinran He, Ke Xu, David Kempe, and Yan Liu. Learning influence functions from incomplete observations. In *Proc. of NIPS*, 2016.
- [Holi and Hwang, 1993] Jordan L Holi and J-N Hwang. Finite precision error analysis of neural network hardware implementations. *IEEE Transactions on Computers*, 1993.
- [Kempe *et al.*, 2003] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proc. of SIGKDD*, 2003.
- [Leskovec *et al.*, 2010] Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos, and Zoubin Ghahramani. Kronecker graphs: an approach to modeling networks. *Journal of Machine Learning Research*, 2010.
- [Leung *et al.*, 2019] Carson K Leung, Alfredo Cuzzocrea, Jiaxing Jason Mai, Deyu Deng, and Fan Jiang. Personalized deepinf: enhanced social influence prediction with deep learning and transfer learning. In *Proc. of Big Data*. IEEE, 2019.
- [Li *et al.*, 2017] Cheng Li, Jiaqi Ma, Xiaoxiao Guo, and Qiaozhu Mei. Deepcas: An end-to-end predictor of information cascades. In *Proc. of WWW*, 2017.
- [Narasimhan *et al.*, 2015] Harikrishna Narasimhan, David C Parkes, and Yaron Singer. Learnability of influence in networks. In *Proc. of NIPS*, 2015.
- [Pedregosa *et al.*, 2011] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 2011.
- [Qiu *et al.*, 2018] Jiezhong Qiu, Jian Tang, Hao Ma, Yuxiao Dong, Kuansan Wang, and Jie Tang. Deepinf: Social influence prediction with deep learning. In *Proc. of SIGKDD*, 2018.
- [Shalev-Shwartz and Ben-David, 2014] Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [Tang *et al.*, 2009] Jie Tang, Jimeng Sun, Chi Wang, and Zi Yang. Social influence analysis in large-scale networks. In *Proc. of SIGKDD*, 2009.
- [Tong *et al.*, 2018] Guangmo Tong, Weili Wu, and Ding-Zhu Du. Coupon advertising in online social systems: Algorithms and sampling techniques. *arXiv preprint arXiv:1802.06946*, 2018.
- [Tong, 2020] Guangmo Tong. Stratlearner: Learning a strategy for misinformation prevention in social networks. In *Proc. of NeurIPS*, 2020.
- [Tzoumas *et al.*, 2012] Vasileios Tzoumas, Christos Amanatidis, and Evangelos Markakis. A game-theoretic analysis of a competitive diffusion process over social networks. In *International Workshop on Internet and Network Economics*. Springer, 2012.
- [Weiss *et al.*, 2018] Gail Weiss, Yoav Goldberg, and Eran Yahav. On the practical computational power of finite precision rnns for language recognition. *arXiv preprint arXiv:1805.04908*, 2018.