Symphony in the Latent Space: Provably Integrating High-Dimensional Techniques with Non-linear Machine Learning Models

Qiong Wu[1] [*], Jian Li[2], Zhenming Liu[1], Yanhua Li[3], Mihai Cucuringu[4,5]

[1]College of William & Mary
[2] Tsinghua University
[3] Worcester Polytechnic Institute
[4] University of Oxford
[5] The Alan Turing Institute

Abstract

This paper revisits building machine learning algorithms that involve interactions between entities, such as those between financial assets in an actively managed portfolio, or interactions between users in a social network. Our goal is to forecast the future evolution of ensembles of multivariate time series in such applications (e.g., the future return of a financial asset or the future popularity of a Twitter account). Designing ML algorithms for such systems requires addressing the challenges of high-dimensional interactions and non-linearity. Existing approaches usually adopt an **ad-hoc** approach to integrating high-dimensional techniques into non-linear models and recent studies have shown these approaches have questionable efficacy in time-evolving interacting systems.

To this end, we propose a novel framework, which we dub as the ***additive influence model***. Under our modeling assumption, we show that it is possible to decouple the learning of high-dimensional interactions from the learning of non-linear feature interactions. To learn the high-dimensional interactions, we leverage kernel-based techniques, with provable guarantees, to embed the entities in a low-dimensional latent space. To learn the non-linear feature-response interactions, we generalize prominent machine learning techniques, including designing a new statistically sound non-parametric method and an ensemble learning algorithm optimized for vector regressions. Extensive experiments on two common applications demonstrate that our new algorithms deliver significantly stronger forecasting power compared to standard and recently proposed methods.

Introduction

We revisit the problem of building machine learning algorithms that involve interactions between entities, such as those between users and items in a recommendation system, or between financial assets in an actively managed portfolio, or between populations in different counties in a disease-spreading process. Our proposed forecasting model uses information available up to time t to predict $\mathbf{y}_{t+1,i}$, the future behavior of entity i at time $t + 1$ (e.g., the future price of stock i at time $t + 1$), for a total number of d entities (Laptev et al. 2017; Farhangi et al. 2022). Designing such models has

___
[*]Currently working at AT&T Labs.

proven remarkably difficult, as one needs to circumvent two main challenges that require often incompatible solutions.

1. Cross-entity interaction: high-dimensionality. In many ensembles of multivariate time series systems, it is often the case that the current state of one entity could potentially impact the future state of another. When considering the equity market as an example, Amazon's disclosure of its revenue change in cloud services could indicate that the revenues of other cloud providers (e.g., competitors) could also change.

The interaction is high-dimensional because the total possible number of interactions is usually much larger than the number of available observations. For example, in a portfolio of 3,000 stocks, the total number of potential links between pairs of stocks is $3,000 \times 3,000 \approx 10^7$, but we often have only 2,500 data points (e.g., 10 years of daily data), and thus capturing the cross-entity interactions becomes a very challenging problem.

2. Feature-response interactions: non-linearity. Linear models are usually insufficient to characterize the relationship between the response/label and the available information (features), thus techniques beyond simple linear regressions are heavily needed. For example, in a financial context, economic productivity is non-linear in temperature for most countries; similarly, electricity consumption is a nonlinear function of temperature, and modeling this relationship is crucial for pricing electricity derivative contracts. As shown in Fig. 1(a), the existing relevant learning models can be categorized into the following two groups.

1. Provable cross-entity models (CEM) for high-dimensionality. Cross-entity models solve a vector regression problem $\mathbf{y}_{t+1} = f(\mathbf{x}_t) + \xi_t$, to forecast the future behavior of all entities, where $\mathbf{y}_{t+1} \triangleq (y_{t+1,1}, \ldots, y_{t+1,d})$, and $\mathbf{x}_t$ denotes the features of all entities, constructed from their historical data. Since the features of one entity can be used to predict the future behavior of another, CEMs have stronger expressive and predictive power. CEMs are both computationally and statistically challenging because we need to solve the "high-dimensional" (overparametrized) problem and mathematically understand the root cause of the overfitting. Extensive research has been undertaken to design regularization techniques (Chen, Dong, and Chan 2013; Wu et al. 2021) to address the issue, and most algorithms in this category are linear and have theoretical guarantees.
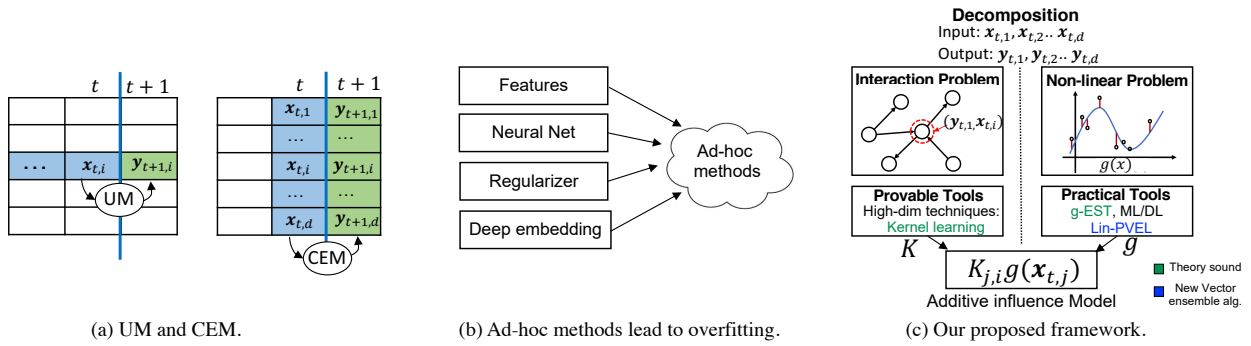
Figure 1: (a) UM for non-linearity and CEM for high-dimensionality. (b) Exsiting ad-hoc methods have questionable efficacy. (c) Our framework decouples the high-dimensional learning of entity interactions and non-linear learning of feature interactions.

2. Practical univariate models (UM) for non-linearity. Univariate models fit a function $\mathbf{y}_{t+1,i} = f(\mathbf{x}_{t,i}) + \xi_{t,i}$ to forecast one entity's feature behavior by using features constructed from that entity's historical data. Univariate models primarily learn the feature-response interaction by using off-the-shelf ML techniques such as Deep learning (DL) (Hochreiter and Schmidhuber 1997; Wu et al. 2019) or gradient boosted algorithms (Chen and Guestrin 2016; Ke et al. 2017; Dorogush, Ershov, and Gulin 2018). These practical models are effective in extracting non-linear signals but they often do not come with theoretical guarantees.

Existing integration techniques: ad-hoc methods It remains unclear how to integrate two seemingly incompatible modeling processes (i.e., UM and CEM) with different design philosophies. In Fig. 1 (b), we show that existing integration solutions predominately follow an **ad-hoc** approach, in part due to the belief that deep learning is the "holly-grail" for practical problems (Sejnowski 2018). For example, one often adds an ℓ_1- or ℓ_2-regularizer to a neural net's cost function, hoping such regularizers will also magically work in neural nets (Paszke et al. 2017). However, the mathematical properties of a provable technique often break when combined into a neural net. Furthermore, latent embedding models have also been recently introduced (Wang et al. 2019; Feng et al. 2019; Chen et al. 2019). The central idea is to project the entities into points in a low-dimensional space so that similar entities (i.e., stocks in the above works) are closer to each other in this embedding. Because point interactions are more restrictive in the latent space, they have the potential to address the overfitting issues (Wang et al. 2019). However, these lines of work do not offer any theoretical guarantees and are often not robust in practice. Recent studies have demonstrated that the efficacy of such ad-hoc approaches is questionable in many interacting systems (Dacrema, Cremonesi, and Jannach 2019; Rendle, Zhang, and Koren 2019; Qiong et al. 2021).

Our approach & contributions We propose a general latent position model dubbed as the *additive influence model* to enable us to seamlessly orchestrate mathematically rigorous high-dimensional techniques with practically effective machine learning algorithms. In Fig. 1(c), we show that it is possible to decouple the learning of high-dim interactions between entities from the learning of the non-linear signals. We assume each entity is associated with an embedded

position $\mathbf{z}_i$ and at timestamp t, entity i is also associated with an unobserved signal $\mathbf{s}_{i,t} \in \mathbf{R}$ that is a function of $\mathbf{x}_{i,t}$. We assume the generative model $\mathbf{y}_{i,t} = \sum_{i \leq j} \kappa(\mathbf{z}_i, \mathbf{z}_j)\mathbf{s}_{j,t} + \epsilon_{i,t}$, where $\kappa(\mathbf{z}_i, \mathbf{z}_j)$ is a function that measures the interaction strength between $\mathbf{z}_i$ and $\mathbf{z}_j$, and can be any kernel function, such as a Gaussian kernel or simply an inner product, and $\epsilon_{i,t}$ denotes noise. Each entity could potentially influence $\mathbf{y}_{i,t}$. The influence of j on i depends on the "distance" or "similarity" between $\mathbf{z}_i$ and $\mathbf{z}_j$. On the other hand, we assume $\mathbf{s}_{j,t} = g(\mathbf{x}_{j,t})$ for some $g(\cdot)$, so that the model captures high-dimensional interactions via $\mathbf{z}_i$ and non-linearity via $g(\cdot)$.

Our proposed model allows for feature interactions through $g(\cdot)$, and addresses the overfitting problem arising from entity interactions because the distances (interaction strength) between entities are constrained by the latent Euclidean space: when both $(\mathbf{z}_i - \mathbf{z}_j)$ and $(\mathbf{z}_j - \mathbf{z}_k)$ are small, then $(\mathbf{z}_i - \mathbf{z}_k)$ is also small, and thus the degree of freedom for entity interactions becomes substantially smaller than $O(d^2)$.

Our goal is to learn both the $\mathbf{z}_i$'s and $g(\cdot)$. We note that these two learning tasks can be *decoupled*: high-dimensional methods can be developed to *provably* estimate the $\mathbf{z}_i$'s *without the knowledge of* $g(\cdot)$, and when estimates of $\mathbf{z}_i$'s are given, an experiment-driven process can be used to learn $g(\cdot)$ by examining prominent machine learning methods such as neural nets and boosting. In other words, when we learn entity interactions, we do not need to be troubled by the overfitting problem escalated by fine-tuning $g(\cdot)$, and when we learn feature interactions, the generalization error will not be jeopardized by the curse of dimensionality from entity interactions.

• To learn the $\mathbf{z}_i$'s, we design a simple algorithm that uses low-rank approximation of $\mathbf{y}_t$'s covariance matrix to infer the closeness of the entities and develop a novel theoretical analysis based on recent techniques from high dimensionality and kernel learning (Belkin 2018; Tang et al. 2013; Wu et al. 2020a).

• To learn $g(\cdot)$, we generalize major machine learning techniques, including neural nets, non-parametric, and boosting methods, to the additive influence model when estimates of $\mathbf{z}_i$'s are known. We specifically develop a moment-based algorithm for non-parametric learning of $g(\cdot)$, and a computationally efficient boosting algorithm.

• Finally, we perform extensive experiments on a major equity market and social network datasets to confirm the efficacy of our modeling approaches and analysis.

Related Work and Comparison

Univariate machine learning models handle feature-response interactions and mostly rely on deep learning and GBRT (Goodfellow, Bengio, and Courville 2016; Wu et al. 2020b; Goodfellow, Bengio, and Courville 2016; Wüthrich, Permunetilleke et al. 1998; Chen and Guestrin 2016; Ke et al. 2017; Dorogush, Ershov, and Gulin 2018; Gong et al. 2017; Yang and Ding 2020; Ding et al. 2015; Zhang, Aggarwal, and Qi 2017; Feng, He, and Polson 2018; Han et al. 2018; Wu et al. 2020b; Chen, Pelger, and Zhu 2019; Kelly, Pruitt et al. 2019; Ke et al. 2019; Chen et al. 2019; Li et al. 2019; Wu et al. 2015). These models aim to optimize their empirical performance and limit theoretical investigations. Recent cross-entity models consider the high-dimensional interactions, where overfitting easily happens and theoretical justifications are essential to avoid spurious result in practice. Cross-entity models are mostly linear models (Bunea, She, and Wegkamp 2011; Koltchinskii, Lounici et al. 2011; Negahban and Wainwright 2011; Huang, Li, and Zhou 2019) that have theoretical guarantees, but they cannot effective for non-linear feature-response interactions. Efforts for building CEMs include (Tibshirani 1996; Candès and Wakin 2008; Tao and Series 2009; Hoerl and Kennard 1970; Tsigler and Bartlett 2020; Liu et al. 2019).

Ad-hoc approach for integration. Recent integrating solutions for high-dimensionality and nonlinearity challenges has been a frustrating endeavor, which we can call the ad-hoc approaches and many were shown to have questionable efficacy in interacting systems. *1. Deep learning + Lasso/Ridge* For example, one (Paszke et al. 2017) often adds an l_1- or l_2-regularizer to a neural net's cost function, hoping these regularizers can also magically work in neural nets. *2. Deep embedding.* Recent studies have addressed high-dimensional entity interactions by using deep embedding, based on the idea that when entities are embedded in low-dim Euclidean space, they can interact in a quite restricted way, therefore preventing overfitting (Zhao et al. 2020; Shen et al. 2022; Zhang, Aggarwal, and Qi 2017; Hu, Liu et al. 2018; Li et al. 2019; Wang et al. 2019). While this idea is effective for linear models (Abraham et al. 2015; Li et al. 2017), deep embedding-based solutions may have very high false positive rates, for instance, when forecasting the returns of financial assets (Qiong et al. 2021; Wang et al. 2019).

Remark: *(i) Modeling framework.* Our framework proposes a key algorithmic insight that the latent position estimation should be decoupled from the learning link function $g(\cdot)$. We develop the first algorithm that can provably estimate the entity's latent positions and provide theoretical guarantees. Our novel analysis leverages a diverse set of tools from kernel learning, non-parametric methods, and random walks. *(ii) Comparison to deep embedding.* While embedding can be learned by deep learning, it usually does not provide any theoretical guarantee, whereas our framework makes stricter assumptions (e.g., how embedding and features should interact) and delivers a quality guarantee. Deep embedding also requires every component including the function $g(\cdot)$ in the architecture to be represented by a neural net to run SGD, whereas we allow $g(\cdot)$ to be learned by a wide range of algorithms such as boosting or non-parametric techniques.

Problem Definition

Notations. For a matrix A, $\mathcal{P}_r(A)$ denotes its rank-r approximation obtained by keeping the top r singular values and the corresponding singular vectors. $\sigma_i(A)$ (resp. $\lambda_i(A)$) is the i-th singular value (resp. eigenvalue) of A. We use Python/MATLAB notation when we refer to a specific row or column. For example, $A_{1,:}$ is the first row of A, and $A_{:,1}$ is the first column. $\|A\|_F$ and $\|A\|_2$ denote the Frobenius and spectral norms, respectively, of A. In general, we use boldface upper case (e.g., $\mathbf{X}$) to denote data matrices and boldface lower case (e.g., $\mathbf{x}$) to denote one sample. $\mathbf{x}_{t,i}$, which refers to the features associated with stock i at time t, can be one or multi-dimensional. Let $(\mathbf{x}_{t,i})_j$ be the j-th coordinate (feature) of $\mathbf{x}_{t,i}$. An event occurring with high probability (whp) means that it happens with probability $\geq 1 - n^{-10}$, where 10 is an arbitrarily chosen large constant and is not optimized. A bivariate function is a Gaussian kernel if $\kappa(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|^2/\sigma^2)$, an inverse multi-quadratic (IMQ) kernel if $\kappa(\mathbf{x}, \mathbf{x}') = (c^2 + \|\mathbf{x} - \mathbf{x}'\|^2)^{-\alpha}$ ($\alpha > 0$), and an inner product kernel if $\kappa(\mathbf{x}, \mathbf{x}') = \langle \mathbf{x}, \mathbf{x}' \rangle$. A function $g(\cdot)$ is Lipschitz-continuous if $|g(\mathbf{x}_1) - g(\mathbf{x}_2)| \leq c\|\mathbf{x}_1 - \mathbf{x}_2\|$ for a constant c. A distribution $\mathcal{D}$ with bounded domain and probability density function $f_\mathcal{D}$ is near-uniform if $\frac{\sup f_\mathcal{D}(\mathbf{x})}{\inf f_\mathcal{D}(\mathbf{x})} = O(1)$.

The forecasting problem. We operate in a time-dependent setting, where each timestamp t can be construed as the t^{th} round. An interacting system consisting of d entities (e.g., denoting stocks in the equity market or user accounts in a network), that are updated at each round, for a total number of T rounds. Let $\mathbf{y}_{t,i} \in \mathbf{R}$ denote the next-period forecast of entity i at the t-th round, and $\mathbf{y}_t = (\mathbf{y}_{t,1}, \ldots, \mathbf{y}_{t,d}) \in \mathbf{R}^d$. Our goal is to forecast $\mathbf{y}_t$ based on all information available up to (but excluding) round t.

Model Assumptions. Under the additive influence model, a generic model takes the form

$$\mathbf{y}_{t,i} = \sum_{j \leq d} \kappa(\mathbf{z}_i, \mathbf{z}_j) g(\mathbf{x}_{t,j}) + \xi_{t,i}, \tag{1}$$

and our goal is to learn $g(\cdot)$ and $\mathbf{z}_i$'s with a total number of n observations. Let $K \in \mathbf{R}^{d \times d}$ such that $K_{i,j} = \kappa(\mathbf{z}_i, \mathbf{z}_j)$. Here, we assume that • *(A.1)* the vector representations $\mathbf{z}_i$'s of the stocks and features $\mathbf{x}_{t,i}$ are i.i.d. samples from (two different) near-uniform distributions on bounded supports, • *(A.2)* $\mathbf{x}_{t,i} \in [-1, 1]$ and $\mathbb{E}[g(\mathbf{x}_{t,i})] = 0$, • *(A.3)* $g(\cdot)$ is Lipschitz-continuous, and • *(A.4)* $\xi_{t,i}$'s are zero-mean i.i.d. Gaussian random variables with standard deviation σ_ξ.

We remark that (A.1) is standard in the literature (Abraham et al. 2015; Sussman, Tang, and Priebe 2013; Tang et al. 2013; Li et al. 2017; Rastelli, Friel, and Raftery 2016). Assuming (A.2) simplifies the calculation and is without loss of generality, and (A.4) can also be relaxed to settings in which the $\xi_{t,i}$ variables are sub-Gaussian. See App. A for a more detailed discussion of the assumptions.
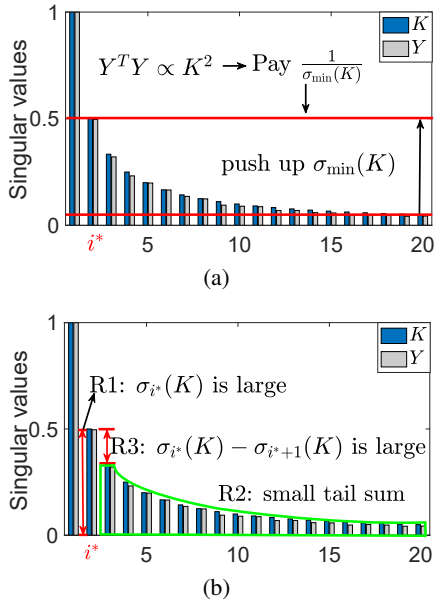
Figure 2: (a) We use the square root of $\mathcal{P}_{i^*}(\mathbf{Y}^T\mathbf{Y})$ to approximate K so that we pay a factor of $1/\sigma_{i^*}(K)$, instead of $1/\sigma_{\min}(K)$. (b) Three key requirements for i^*: ● (R1) $\sigma_{i^*}(K)$ is large, ● (R2) $\mathcal{P}_{i^*}(K^2)$ is close to K^2, and ● (R3) $\sigma_{i^*}(K) - \sigma_{i^*+1}(K)$ is large.

Our Algorithms

This section introduces our algorithmic pipeline in full detail. Sec. 4 describes an algorithm for learning the embedding without knowing $g(\cdot)$. Sec. 4 explains the estimation of $g(\cdot)$ using machine learning techniques. Due to the space limit, detailed proofs of all the Props are deferred to App. B.

Learning Vector Representation Provably

This section presents a provable algorithm to estimate the kernel matrix K and the embedding $\mathbf{z}_i$'s. Our algorithm does not require knowledge of $g(\cdot)$, thus providing a conceptually new approach to construct CEMs: high-dimensional learning of entity interactions can be decoupled from using ML techniques to fit the features. Because learning entity interactions could be a major source of causing overfitting, disentangling it from the downstream task of learning $g(\cdot)$ enables us to leverage the function-fitting power of ML techniques without the cost of amplifying generalization errors.

We next walk through our design intuition and start by introducing additional notation. Let $\mathbf{Y} \in \mathbf{R}^{n \times d}$ be such that $\mathbf{Y}_{t,i} = \mathbf{y}_{t,i}$ ($\mathbf{Y}$ is a matrix and $\mathbf{y}$ a random variable), $\mathbf{S} \in \mathbf{R}^{n \times d}$ with $\mathbf{S}_{t,i} = \mathbf{s}_{t,i} \triangleq g(\mathbf{x}_{t,i})$, and $E \in \mathbf{R}^{n \times d}$ with $E_{t,i} = \xi_{t,i}$. Recall that $K \in \mathbf{R}^{d \times d}$ s.t. $K_{i,j} = \kappa(\mathbf{z}_i, \mathbf{z}_j)$, and $\mathcal{P}_r(A)$ denotes A's rank-r approximation obtained by keeping the top r singular values and vectors. Finally, for any PSD matrix A with SVD $A = U\Sigma U^T$, let $\sqrt{A} \triangleq U\Sigma^{\frac{1}{2}}U^T$.

Eq. (1) can be re-written as $\mathbf{Y} = \mathbf{S}K + E$, in which we need to infer K using only $\mathbf{Y}$. We first observe that while none of the entries in $\mathbf{S}$ are known, the $\mathbf{S}_{t,i}$'s are i.i.d. random variables (because the $\mathbf{x}_{t,i}$'s are i.i.d.); therefore, our

problem resembles a *dictionary learning* problem, in which K can be viewed as the dictionary to be learned, and $\mathbf{S}$ is the measurement matrix (see e.g., (Arora et al. 2014)). However, in our case, K is neither low-rank nor sparse, and we cannot use standard dictionary learning techniques.

First, we observe that, if infinitely many samples were available, then $\mathbf{Y}^T\mathbf{Y}/n$ approaches to K^2. Hence, intuitively we could use $\sqrt{\mathbf{Y}^T\mathbf{Y}/n}$ to approximate $\sqrt{K^2} = K$. However, the existing standard matrix square root result has the notorious "$1/\sigma_{\min}$-blowup" problem, i.e., it gives us only $\|\sqrt{\mathbf{Y}^T\mathbf{Y}/n} - KW\|_F \propto 1/\sigma_{\min}(K)$ (W a unitary matrix), where typically $\sigma_{\min}(K)$ is extremely small, thus rendering the bound too loose to be useful (Bhojanapalli, Kyrillidis, and Sanghavi 2016).

To tackle the problem, our algorithm uses $\sqrt{\mathcal{P}_{i^*}(\mathbf{Y}^T\mathbf{Y})/n}$ to approximate K for a carefully chosen i^* so that we pay a factor of $\sigma_{i^*}(K)$, instead of $\sigma_{\min}(K)$, to substantially tighten the error. See Alg. 2 in App. B and Fig. 2(a). To implement this idea, we need to show that there always exists an i^* such that ● *(R1):* $\sigma_{i^*}(K)$ is sufficiently large, ● *(R2):* $\mathcal{P}_{i^*}(K^2)$ is close to K^2, and ● *(R3):* the spectral gap $\sigma_{i^*}(K) - \sigma_{i^*+1}(K)$ is sufficiently large so that we can use the Davis-Kahan theorem to prove that $\mathcal{P}_{i^*}(K^2) \propto \mathcal{P}_{i^*}(Y^T Y)$ (Stewart and Sun 1990). See also Fig. 2(b).

These three requirements may not always be met simultaneously. For example, when $\sigma_i(K^2) \propto \frac{1}{i}$, the gap is insufficient and the tail diverges (R2 and R3 are violated). Therefore, we integrate the following two results. ● *(i)* The eigenvalues decay fast. This stems from two classical results from the *kernel learning* literature. First, when $\kappa(\cdot, \cdot)$ is sufficiently smooth (such as the Gaussian, IMQ, or inner product kernels), the eigenvalues of the kernel operator $\mathcal{K}$ associated with $\kappa(\cdot, \cdot)$ decay exponentially (e.g., $\lambda_i(\mathcal{K}) \le \exp(-Ci^{\frac{1}{r}})$ for Gaussian kernels (Belkin 2018)). Second, it holds true that $\sum_{i \ge 1} |\lambda_i(\mathcal{K}) - \lambda_i(K/d)|_F^2 \propto \frac{1}{n}$, a convergence result under the PAC setting (Tang et al. 2013). Therefore, $\lambda_i(K)$ also approximately decays exponentially. ● *(ii)* Combinatorial analysis between gaps and tails. We then leverage a recent analysis (Wu et al. 2020a) showing that when $\lambda_i(K)$ decays fast, it is always possible to find an i^* such that $\lambda_{i^*}(\mathcal{K}) - \lambda_{i^*+1}(\mathcal{K})$ is sufficiently large (R1 & R3 are satisfied) and $\sum_{j \ge i^*} \lambda_j^2(\mathcal{K}) = o(1)$ (R2 is satisfied). Putting all these together leads to the following statement.

Proposition 4.1. *Consider the additive influence model. Let $\kappa(\mathbf{z}_i, \mathbf{z}_j)$ be a Gaussian, inverse multi-quadratic (IMQ) or inner product kernel. Let $n \ge d$ be the number of observations and $\epsilon = \frac{c_0 \log^3 d}{\sqrt{d}}$. Assume that the noise level $\sigma_\xi = O(\sqrt{d})$. Let δ be a tunable parameter (also appeared in Alg. 2 in App. B) such that $\delta^3 = \omega(\epsilon^2)$. There exists an efficient algorithm that outputs $\hat{K}$ such that $\frac{1}{d^2}\|\hat{K} - K\|_F^2 = O(\frac{\epsilon^2}{\delta^3} + \delta^{\frac{4}{5}})(= \tilde{O}(d^{-\Theta(1)}))$.*

We remark that *(i)* the algorithm does not need to know the exact form of κ, so long as it is one of Gaussian, IMQ, or inner product kernels, *(ii)* once K is estimated, an Isomap-flavored algorithm may be used to estimate $\mathbf{z}_i$'s (Li et al. 2017), and *(iii)* knowing $\hat{K}$ (without reconstructing $\mathbf{z}_i$'s) is

Algorithm 1: nparam-gEST:

Input $\mathbf{X}$, $\mathbf{Y}$, $\hat{K}$;
Output μ_1 (estimating other μ_i's is similar)

1: **procedure** NPARAM-GEST($\hat{K}$, $\mathbf{X}$, $\mathbf{Y}$)
2:     **for all** $t \leftarrow 1$ **to** n **do**
3:        $q_t = \text{Rand}(d)$
4:        $L_{(t,q_t),j} = \text{MAP-REGRESS}(q_t, \hat{K}, \mathbf{X}_{t,:})$
5:     **return** $\mu_1 \leftarrow \text{FLIPSIGN}\left(q_t, \{\mathbf{y}_t, L_{(t,q_t),j}\}_{t \leq n}\right)$
6: **procedure** MAP-REGRESS(q_t, $\hat{K}$, $\mathbf{x}_t$)
7:     Let $L_{(t,q_t),j} = 0$
8:     **for all** $k \leftarrow 1$ **to** d **do**
9:        $L_{(t,q_t),j} += \hat{K}_{q_t,k}$ with j s.t. $\mathbf{x}_{t,k} \in \Omega_j$.
10:     **return** $L_{(t,q_t),j}$
11: **procedure** FLIPSIGN(q_t, $\{\mathbf{y}_t, L_{(t,q_t),j}\}_{t \leq n}$)
12:     **for all** $t \leftarrow 1$ **to** n **do**
13:        $\hat{\Pi}_1^{(q_t)}(t) \triangleq L_{(t,q_t),1} - \frac{1}{\ell-1}\left(\sum_{j \neq 1} L_{(t,q_t),j}\right)$
14:        $\tilde{b}_{t,q_t} = \begin{cases} 1 & \text{if } \hat{\Pi}_1^{(q_t)}(t) \geq \frac{c}{\log d}\sqrt{\frac{d}{\ell}} \\ -1 & \text{if } \hat{\Pi}_1^{(q_t)}(t) < -\frac{c}{\log d}\sqrt{\frac{d}{\ell}} \\ 0 & \text{otherwise} \end{cases}$
15:     **return** $\mu_1 = \frac{\sum_{t \leq n} \tilde{b}_{t,q_t}\mathbf{y}_{t,q_t}}{\sum_{t \leq n} \tilde{b}_{t,q_t}\hat{\Pi}_1^{(q_t)}(t)}$

$$y_{t,1} = g(x_{t,1}) + g(x_{t,2}) + ... + g(x_{t,d}) + \xi_{t,1} \ (K_{ij} = 1)$$


Figure 3: A toy example of nparam-gEST when $K_{i,j} = 1$ for all i and j and $\Omega = [-1,1]$ and is uniformly partitioned into 10 pieces. Sampling a $g(\mathbf{x}_{t,i})$ corresponds to randomly placing a ball into a total number of 10 bins. For example, $\mathbf{x}_{t,2}$ falls into the 8-th interval so μ_8 is used to approximate $g(\mathbf{x}_{t,2})$, which may be viewed as a new ball of type μ_8 (or in 8-th bin) is created. The mean load for each bin is $d/\ell = d/10$. We calculate $\sum_{i \leq d} g(\mathbf{x}_{t,i})$ by counting the balls in each bin: $\mathbf{y}_{t,1} = 5 \times \mu_1 + 1 \times \mu_2 + ... + 6 \times \mu_8 + 3 \times \mu_9 + 1 \times \mu_{10} + \xi_{t,1}$.

sufficient for the downstream $g(\cdot)$-learners.

Learning $g(\cdot)$

Here, we explain how prominent machine learning techniques, including neural nets (deep learning), non-parametric methods, and boosting, can be used to learn $g(\cdot)$. These techniques make different functional form assumptions of $g(\cdot)$, and possess different "iconic" properties: deep learning assumes that $g(\cdot)$ can be represented by a possibly sophisticated neural net and uses stochastic gradient descent to train the model; non-parametric methods learn a Lipschitz-continuous $g(\cdot)$ with statistical guarantees; boosting consolidates forecasts produced from computationally efficient weak learners.

Our setting has a different cost structure: in univariate models, $g(\mathbf{x}_{t,j})$ controls only one response $\hat{\mathbf{y}}_{t,j}$, but here, $g(\mathbf{x}_{t,j})$ impacts all responses $\hat{\mathbf{y}}_{t,i}$, $i \in [d]$, as $\hat{\mathbf{y}}_{t,i} = \sum_j K_{i,j}g(\mathbf{x}_{t,j})$. We generalize ML techniques under the new cost functions, while retaining the iconic properties of each technique.

Technique 1. Learn $g(\cdot)$ using neural nets. When an estimate $\hat{K}$ is given, the training cost is $\sum_{t,i}(\mathbf{y}_{t,i} - \sum_{j \in [d]} \hat{K}_{i,j}g(\mathbf{x}_{t,j}))^2$, in which case one can employ stochastic gradient descent when $g(\cdot)$ is a neural net.

Technique 2. Learn $g(\cdot)$ using non-parametric methods. When the response is univariate, e.g., $\mathbf{y}_{t,i} = g(\mathbf{x}_{t,i}) + \xi_{t,i}$, we can use a neighbor-based approach to estimate $g(\mathbf{x})$ for a new $\mathbf{x}$: we identify one (or multiple) $\mathbf{x}_{t,i}$'s in the training set that are close to the new $\mathbf{x}$, and output $\mathbf{y}_{t,i}$ (or their averages, when multiple $\mathbf{x}_{t,i}$ are chosen), using $g(\mathbf{x}) \approx g(\mathbf{x}_{t,i})$, whenever $\mathbf{x}$ is close to $\mathbf{x}_{t,i}$.

Here, we do not directly observe the values of individual $g(\mathbf{x}_{t,i})$'s. Instead, each response is a linear combina-

tion of multiple $g(\cdot)$'s evaluated at different points, e.g., $\mathbf{y}_{t,1} = K_{i,1} \cdot g(\mathbf{x}_{t,1}) + \cdots + K_{i,d} \cdot g(\mathbf{x}_{t,d}) + \xi_{t,i}$. We show that finding neighbors reduces to solving a linear system. Furthermore, we design a moment-based algorithm, namely "nparam-gEST", which estimates $g(\cdot)$ with provable guarantees, as summarized in the following result.

Proposition 4.2. *Consider the problem of learning an additive influence model with the same setup/parameters as in Prop. 4.1. Assume that $\mathbf{x}_{t,i} \in \mathbf{R}^{O(1)}$. Let ℓ be a tunable parameter. There exists an efficient algorithm to compute $\hat{g}(\cdot)$, based on $\hat{K}$ such that $\sup_{\mathbf{x}} |\hat{g}(\mathbf{x}) - g(\mathbf{x})| \leq (\log^6 n)(\sqrt{\gamma} + \sqrt{\frac{\ell}{n}} + \frac{1}{\ell}) = \tilde{O}(d^{-c})$ for suitable parameters, where $\gamma \triangleq \frac{\epsilon^2}{\delta^3} + \delta^{\frac{4}{5}}$.*

Our algorithm (Alg. 1) consists of the following 3 steps:

Step 1. Approximation of $g(\cdot)$. Partition $\Omega = [-1,1]^k$ into subsets $\{\Omega_j\}_{j \leq \ell}$, and use piece-wise constant function to approximate $g(\cdot)$, i.e., $\tilde{g}(\mathbf{x}_{t,i})$ takes the same value for all $\mathbf{x}_{t,i}$ in the same Ω_j. We partition $\{\Omega_j\}_{j \leq \ell}$ in a way such that $\Pr[\mathbf{x}_{t,i} \in \Omega_j]$ are the same for all j.

Step 2. Reduction to linear regression. Each observation can be construed as a linear combination of μ_j's ($j \in [\ell]$), where $\mu_j = \mathbb{E}[g(\mathbf{x}_{t,i}) \mid \mathbf{x}_{t,i} \in \Omega_j]$. For example, $\mathbf{y}_{t,1} = \sum_{i \leq d} K_{1,i}\mu_{j_i} + \xi_{t,1} + o(1)$, where $\mathbf{x}_{t,i} \in \Omega_{j_i}$, and in general, we have

$$\mathbf{y}_{t,i} = \sum_{j \leq \ell} L_{(t,i),j}\mu_j + \xi_{t,i} + o(1), \qquad (2)$$

where $L_{(t,i),j} = \sum_{m \in \mathcal{L}_{t,j}} K_{i,m}$ and $\mathcal{L}_{t,j} = \{m : \mathbf{x}_{t,m} \in \Omega_j\}$.

Therefore, our learning problem reduces to a linear regres-

sion problem, in which the $L_{(t,i),j}$'s are features and the $\{\mu_j\}_{j\le\ell}$ are coefficients to be learned.

Step 3. Moment-based estimation. An MSE-based estimator is consistent but finding its confidence interval (error bound) requires knowing the spectrum of the features' covariance matrix, which is remarkably difficult in our setting. Therefore, we propose a moment-based algorithm with provable performance (FLIPSIGN in Alg. 1).

We illustrate each steps above through a toy example, in which we assume $K_{i,j} = 1$ for all i and j so the model simplifies to $\mathbf{y}_{t,1} = \sum_{j\le d} g(\mathbf{x}_{t,j}) + \xi_{t,1}$. See Fig. 3 for additional details.

Steps 1 & 2. First, we view the generation of samples as a balls-and-bins process so that the $g(\cdot)$-estimation problem reduces to a regression problem (Steps 1 & 2). Specifically, we generate $(\mathbf{y}_{t,1}, \{\mathbf{x}_{t,i}\}_{i\le d})$ as first sequentially sampling $\{\mathbf{x}_{t,i}\}_{i\le d}$ and computing the corresponding $g(\mathbf{x}_{t,i})$, then summing each term up together with $\xi_{t,1}$ to produce $\mathbf{y}_{t,1}$. When an $\mathbf{x}_{t,i}$ is sampled, it falls into one of Ω_i's with uniform probability. Let j_i be the bin that $\mathbf{x}_{t,i}$ falls into. Then $g(\mathbf{x}_{t,i})$ is approximated by μ_{j_i} according to Step 1. Thus, we may view a ball of "type μ_{j_i}" (or in j_i-th bin) is created. For example, in Fig. 3, $\mathbf{x}_{t,2}$ falls into the 8-th interval so a ball is added in the 8-th bin. After all $\mathbf{x}_{t,i}$'s are sampled, compute $\mathbf{y}_{t,1}$ by counting the numbers of balls in different bins. Recalling that the load of j-th bin is $L_{(t,1),j}$, we have $\mathbf{y}_{t,1} \approx \sum_{j\le d} L_{(t,1),j} \cdot \mu_j + \xi_{t,1}$. Let $\Delta_{t,j} = L_{(t,1),j} - d/\ell$ and using that $\mathbb{E}[L_{(t,1),j}] = d/\ell$ and $\sum_{j\le d}\mu_j = 0$, we have

$$\mathbf{y}_{t,1} = \Delta_{t,1}\mu_1 + \cdots + \Delta_{t,\ell}\mu_\ell + \xi_{t,1}. \qquad (3)$$

Eq. (3) is a standard (univariate) regression: for each t, we know $\mathbf{y}_{t,1}$, and know all $\Delta_{t,j}$'s because all $\mathbf{x}_{t,j}$'s are observed so the number of balls in each bin can be calculated. We need to estimate the unknown μ_j's. Note that $\mathbb{E}[\Delta_{t,j}] = 0$.

Steps 3. We solve the regression (Step 3). Our algorithm "tweaks" the observations so that the features associated with μ_1 are always positive: let $b_{t,1} = 1$ if $\Delta_{t,1} > 0$ and -1 otherwise. Multiply $b_{t,1}$ to both sides of Eq. (3) for each t,

$$b_{t,1}\mathbf{y}_{t,1} = |\Delta_{t,1}|\mu_1 + \cdots + b_{t,1}\cdot\Delta_{t,\ell}\cdot\mu_\ell + b_{t,1}\xi_{t,1}. \quad (4)$$

We sum up the LHS and RHS of (4) and obtain

$$\sum_{t\le n} b_{t,1}\mathbf{y}_{t,1} = \big(\sum_{t\le n}|\Delta_{t,1}|\big)\mu_1 + \cdots + \qquad (5)$$

$$\big(\sum_{t\le n}b_{t,1}\cdot\Delta_{t,\ell}\big)\mu_\ell + \sum_{t\le n}b_{t,1}\xi_{t,1}.$$

Next, we have $\sum_{t\le n}|\Delta_{t,1}| = \Theta(n)$ whp. Also, we can see that $b_{t,1}$ and $\Delta_{t,j}$ are "roughly" independent for $j\neq 1$ (careful analysis will make it rigorous). Therefore, for any $j\neq 1$, $\mathbb{E}[b_{t,1}\cdot\Delta_{t,j}] = 0$, and thus $\sum_{t\le n}b_{t,1}\cdot\Delta_{t,j} = O(\sqrt{n})$ whp. Now (5) becomes $\sum_{t\le n}b_{t,1}\cdot\mathbf{y}_{t,1} = \big(\sum_t|\Delta_{t,1}|\big)\mu_1 + O(\ell\cdot\sqrt{n})$. Thus our estimator is $\hat{\mu}_1 \triangleq \frac{\sum_t b_{t,1}\cdot\mathbf{y}_{t,1}}{\big(\sum_t|\Delta_{t,1}|\big)} = \mu_1 + \frac{O(\ell\cdot\sqrt{n})}{\Theta(n)} = \mu_1 + O\big(\frac{\ell}{\sqrt{n}}\big)$. Here, the covariance analysis for the $\Delta_{t,j}$'s is circumvented because $\Delta_{t,j}$'s interactions are compressed into the term $O\big(\frac{\ell}{\sqrt{n}}\big)$. We remark that the above

analysis contains some crude steps and can be tightened up, as we have done in App. C.

Technique 3. Learn $g(\cdot)$ using boosting. In the univariate setting, we have $\mathbf{y}_{t,i} = \sum_{m\le b} g_m(\mathbf{x}_{t,i}) + \xi_{t,i}$, in which each $g_m(\mathbf{x}_{t,i})$ is a weak learner. Standard boosting algorithms, such as (Quinlan 1986; Chen and Guestrin 2016), assume that each $g_m(\cdot)$ is represented by a regression tree and constructed sequentially. A greedy strategy is used to build a new tree, e.g., iteratively splitting a node in a tree by choosing a variable that optimizes prediction improvement. In our setting, $\mathbf{y}_{t,i}$ depends on evaluating $g_m(\cdot)$ at d different locations $\mathbf{x}_{t,1}, \ldots, \mathbf{x}_{t,d}$, so the splitting procedure either is d (e.g. 3000 for equity market) times slower in a standard implementation, or requires excessive engineering tweak of existing systems.

Here, we propose a simple and effective weak learner based on the intuition of the tree structure. Let

$$(\mathbf{x}_t)_i = ((\mathbf{x}_{t,1})_i, (\mathbf{x}_{t,2})_i, \ldots, (\mathbf{x}_{t,d})_i) \in \mathbf{R}^d,$$

$$(\mathbf{x}_t)_{i,j} = ((\mathbf{x}_{t,1})_i\cdot(\mathbf{x}_{t,1})_j, \ldots, (\mathbf{x}_{t,d})_i\cdot(\mathbf{x}_{t,d})_j) \in \mathbf{R}^d,$$

and $(\mathbf{x}_t)_{i,j,k}$ can be defined in a similar manner. We observe that regression trees used in GBRT models for equity return are usually shallow and can be *linearized*: we may unfold a tree into disjunctive normal form (DNF) (Abasi, Bshouty, and Mazzawi 2014), and approximate the DNF by a sum of multiple interaction terms, e.g., $I((\mathbf{x}_{t,i})_1 > 0) \cdot I((\mathbf{x}_{t,i})_2 > 0)$ can be approximated by $(\mathbf{x}_{t,i})_1 \cdot (\mathbf{x}_{t,i})_2$.

Our algorithm, namely LIN-PVEL (linear projected vector ensemble learner), consists of weak learners in linear forms. Each linear learner consists of a subset of features and their interactions. The number of features included and the depth of their interactions are hyper-parameters corresponding to the depth of the decision tree. For example, if the first three features are included in the learner, we need to fit $\mathbf{y}_{t,i}$ against

$$\sum_{j\in[d]} \underbrace{\hat{K}_{i,j}}_{\text{given}} \cdot \Big[\underbrace{\beta_1(\mathbf{x}_{t,j})_1 + \ldots}_{\text{linear terms}} + \underbrace{\beta_4(\mathbf{x}_{t,j})_{1,2} + \cdots + \beta_7(\mathbf{x}_{t,j})_{1,2,3}}_{\text{interaction terms}}\Big]$$
$$(6)$$

by MSE. Conceptually, although we use linearized models to approximate the trees, the "target" trees are unavailable (for the computational efficiency reasons above). We need a new procedure to select features for each learner. Our intuition is that, if an interaction term could have predictive power, each feature involved in the interaction should also have predictive power. Our procedure is simply to select a fixed number of i's with the largest $\mathrm{corr}((\mathbf{y}_{\mathrm{Res}})_t, \hat{K}(\mathbf{x}_t)_i)$, where $(\mathbf{y}_{\mathrm{Res}})_t$ is the residual error.

Using feature interactions to approximate DNF $(I((\mathbf{x}_{t,i})_1 > 0) \cdot I((\mathbf{x}_{t,i})_2 > 0) \approx (\mathbf{x}_{t,i})_1 \cdot (\mathbf{x}_{t,i})_2)$ may not always be accurate, however, in our setting, linear interaction models often outperform decision trees or DNFs. We believe this occurs because interaction terms are continuous (whereas DNFs are discrete functions), and thus they are more suitable to model smooth changes.

Evaluation

We evaluate our algorithms on two real-world data sets: an equity market to predict stock returns, and a social network data

	Universe 800				Full universe				Backtesting	
Models	corr	w_corr	t-stat	w_t-stat	corr	w_corr	t-stat	w_t-stat	PnL	Sharpe
Ours: LIN-PVEL	**0.0764**	**0.0936**	**6.7939**	**6.3362**	**0.0944**	**0.1009**	**8.2607**	**6.4435**	**0.5261**	**10.97**
Ours: nparam-gEST	**0.0446**	**0.0320**	**3.2961**	**1.5753**	**0.0618**	**0.0553**	**5.7327**	**3.5212**	**0.3386**	**7.59**
Ours: MLP	**0.0550**	**0.0567**	**6.4782**	**5.0172**	**0.0738**	**0.0692**	**9.2034**	**6.4151**	**0.4202**	**9.43**
Ours: LSTM	**0.0286**	**0.0347**	**3.4517**	**3.0261**	**0.0473**	**0.0491**	**6.3615**	**4.2385**	**0.2487**	**7.10**
UM: poor man Lin-PVEL	0.0674	0.0866	6.0947	5.7312	0.0827	0.0884	7.4297	5.6659	0.4565	9.76
UM: poor man nparam-gEST	0.0432	0.0309	3.1505	1.4912	0.0584	0.0509	5.0098	3.0844	0.3070	6.59
UM: MLP	0.0507	0.5050	6.0234	4.4966	0.0606	0.0467	8.2857	4.4555	0.2782	6.38
UM: LSTM	0.0178	0.0200	2.2136	1.8077	0.0352	0.0297	4.0602	2.3619	0.175	4.33
UM: Linear models	0.0106	0.0192	1.6471	2.3030	0.0290	0.0251	4.4711	2.6010	0.1888	4.79
UM: GBRT	0.0516	0.0591	7.5739	5.6310	0.0673	0.0747	9.3379	7.8931	0.3858	4.45
UM: SFM	0.0027	0.0032	0.4688	0.4050	0.0147	0.0051	1.2683	0.3892	0.0169	0.54
Existing CEM: VR	0.0156	0.0159	2.4997	1.7046	0.0041	-0.0025	0.8847	-0.3021	0.0430	1.20
Existing CEM: ARRR	0.0314	0.0382	2.5336	2.4213	0.0222	0.0273	1.8557	1.8968	0.1674	3.24
Ad-hoc integration: AlphaStock	0.0085	0.0063	2.1045	1.2516	0.0027	0.0032	0.4688	0.4050	0.0045	0.10
Ad-hoc integration: HAN	0.0105	0.0081	1.7992	1.0017	0.0080	0.0050	1.5716	0.7340	0.0570	2.02
Consolidated: All Ours	**0.0775**	**0.0950**	**6.8687**	**6.4108**	**0.0958**	**0.1025**	**8.5703**	**6.6487**	**0.5346**	**11.30**

Table 1: Summary of results for equity raw return forecasts. LIN-PVEL is the gradient boosting method with the linear learner. Boldface denotes the best performance in each group. Backtesting results pertain to the *Full universe*.

set to predict user popularity, respectively. We demonstrate the main results in the equity market and additional details and experiments for the equity market and Twitter data sets are in APP. G. We remark that this is a *theoretical* paper; examining the performance on more data sets and baselines is a promising direction for future work.

Models under our framework. We estimate K and $g(\cdot)$ separately. To estimate K, we use both the algorithm discussed in Sec. 4 and other refinements discussed in App. B. To estimate $g(\cdot)$, we use SGD-based algorithms (MLP and LSTM), nparam-gEST, and LIN-PVEL.

Baselines. Our baselines include the commonly used models and domain specific models. *(i) The UMs* include linear, MLP, LSTM, GBRT, and SFM (Zhang, Aggarwal, and Qi 2017). We also implement a "poor man's version" of both LIN-PVEL and nparam-gEST for UM, which assumes that influences from other entities are 0; *(ii) The CEMs* include a standard linear VAR (Negahban and Wainwright 2011), ARRR (Wu et al. 2020a). *(iii) Ad-hoc integration* AlphaStock (Wang et al. 2019), and HAN (Hu, Liu et al. 2018) for the equity data set;

Predicting equity returns. We use 10 years of equity data from an emerging market to evaluate our algorithms and focus on predicting the next 5-day returns, for which the last three years are out-of-sample. The test period is substantially longer than those employed in recent works (Zhang, Aggarwal, and Qi 2017; Hu, Liu et al. 2018; Li et al. 2019), adding to the robustness of our results. We constructed 337 standard technical factors to serve as a feature database for all models. We consider two universes: (i) *Universe 800* can be construed as an equivalence to the S&P 500 in the US, and consists of 800 stocks, and (ii) *Full universe* consists of all stocks except for the very illiquid ones.

We next describe our evaluation metrics and argue why they are more suitable and different from those employed in standard ML problems (see App. G) ● *(i) Correlation vs MSE.* While the MSE is a standard metric for regression problems, correlations are better-suited metrics for equity

data sets. ● *(ii) Significance testing.* The use of t-statistics estimators (Newey and West 1986) can account for the serial and cross-sectional correlations (App. G) ● *(iii) Stock capacity/liquidity considerations.* Predicting illiquid stocks is less valuable compared to predicting liquid ones because they cannot be used to build large portfolios. We use a standard approach to weight correlations (w_corr) and t-statistics by a function of historical *notional (dollar) traded volume* to reflect the capacity of the signals.

Results. See Table 1 for the results and the simulated Profit & Loss (PnL). The experiments confirm that ● *(i)* Models under our framework consistently outperform prior works. In addition, our LIN-PVEL model has the best performance; ● *(ii)* By using a simple consolidation algorithm, the aggregated signal outperforms all individual ones. Our new models pick up signals that are orthogonal to existing ones because we rely on a new mechanism to use stock and feature interactions.

Conclusion

This paper revisits the problem of building machine learning algorithms that involve interactions between entities. We propose an *additive influence framework* that enables us to decouple the learning of the entity-interactions from the learning of feature-response interactions. Our upstream entity interaction learner has provable performance guarantees, whereas our downstream $g(\cdot)$-learners can leverage a wide set of effective ML techniques. All these methods under our framework are proven to be superior to the existing baselines.

Acknowledgments

References

Abasi, H.; Bshouty, N. H.; and Mazzawi, H. 2014. On exact learning monotone DNF from membership queries. In *ALT*.

Abraham, I.; Chechik, S.; Kempe, D.; and Slivkins, A. 2015. Low-distortion inference of latent similarities from a multiplex social network. *SICOMP*.

Amihud, Y. 2002. Illiquidity and stock returns: cross-section and time-series effects. *Journal of financial markets*.

Arora, S.; Bhaskara, A.; Ge, R.; and Ma, T. 2014. More algorithms for provable dictionary learning. *arXiv preprint arXiv:1401.0579*.

Azevedo, D.; and Menegatto, V. A. 2015. Eigenvalues of dot-product kernels on the sphere. *Proceeding Series of the Brazilian Society of Computational and Applied Mathematics*.

Belkin, M. 2018. Approximation beats concentration? An approximation view on inference with smooth radial kernels. In *COLT*.

Bhojanapalli, S.; Kyrillidis, A.; and Sanghavi, S. 2016. Dropping convexity for faster semi-definite optimization. In *COLT*.

Bunea, F.; She, Y.; and Wegkamp, M. H. 2011. Optimal selection of reduced rank estimators of high-dimensional matrices. *ANN STAT*.

Candès, E. J.; and Wakin, M. B. 2008. An introduction to compressive sampling. *IEEE signal processing magazine*, 25(2): 21–30.

Chen, C.; Zhao, L.; Bian, J.; et al. 2019. Investment behaviors can tell what inside: Exploring stock intrinsic properties for stock trend prediction. In *KDD*.

Chen, K.; Dong, H.; and Chan, K.-S. 2013. Reduced rank regression via adaptive nuclear norm penalization. *Biometrika*.

Chen, L.; Pelger, M.; and Zhu, J. 2019. Deep learning in asset pricing. *Available at SSRN 3350138*.

Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *KDD*.

Chung, F.; and Lu, L. 2006. Concentration inequalities and martingale inequalities: a survey. *Internet Mathematics*.

Colby, R. W.; and Meyers, T. A. 1988. *The encyclopedia of technical market indicators*. Dow Jones-Irwin Homewood, IL.

Dacrema, M. F.; Cremonesi, P.; and Jannach, D. 2019. Are We Really Making Much Progress? A Worrying Analysis of Recent Neural Recommendation Approaches. In *Proceedings of the 13th ACM Conference on Recommender Systems*, RecSys '19.

Ding, X.; Zhang, Y.; Liu, T.; and Duan, J. 2015. Deep learning for event-driven stock prediction. In *IJCAI*.

Dorogush, A. V.; Ershov, V.; and Gulin, A. 2018. CatBoost: gradient boosting with categorical features suppfort. *arXiv preprint arXiv:1810.11363*.

Fama, E. F.; and French, K. R. 1993. Common risk factors in the returns on stocks and bonds. *JFE*.

Farhangi, A.; Bian, J.; Huang, A.; Xiong, H.; Wang, J.; and Guo, Z. 2022. AA-Forecast: Anomaly-Aware Forecast for Extreme Events. *arXiv preprint arXiv:2208.09933*.

Feng, F.; He, X.; Wang, X.; Luo, C.; Liu, Y.; and Chua, T.-S. 2019. Temporal relational ranking for stock prediction. *TOIS*.

Feng, G.; He, J.; and Polson, N. G. 2018. Deep learning for predicting asset returns. *arXiv preprint arXiv:1804.09314*.

Ge, R.; Jin, C.; et al. 2017. No spurious local minima in nonconvex low rank problems: A unified geometric analysis. In *ICML*.

Gong, R.; Fonseca, E.; Bogdanov, D.; Slizovskaia, O.; Gomez, E.; and Serra, X. 2017. Acoustic scene classification by fusing LightGBM and VGG-net multichannel predictions. In *Proc. IEEE AASP Challenge Detection Classification Acoust. Scenes Events*.

Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep learning*. MIT press.

Grover, A.; and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *KDD*.

Gu, S.; Kelly, B.; and Xiu, D. 2020. Empirical asset pricing via machine learning. *The Review of Financial Studies*.

Ha, C.-W. 1986. Eigenvalues of differentiable positive definite kernels. *SIAM Journal on Mathematical Analysis*.

Hamilton, A.; and Tegmark, M. 2000. Decorrelating the power spectrum of galaxies. *Monthly Notices of the Royal Astronomical Society*.

Han, Y.; He, A.; Rapach, D.; and Zhou, G. 2018. What Firm Characteristics Drive US Stock Returns? *Available at SSRN 3185335*.

Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*.

Hoerl, A. E.; and Kennard, R. W. 1970. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1): 55–67.

Hu, Z.; Liu, W.; et al. 2018. Listening to chaotic whispers: A deep learning framework for news-oriented stock trend prediction. In *WSDM*.

Huang, D.; Li, J.; and Zhou, G. 2019. Shrinking factor dimension: A reduced-rank approach. *Available at SSRN 3205697*.

Hurst, H.; Black, R.; and Simaika, Y. 1965. Long-term storage: an experimental study Constable. *London UK*.

Kakushadze, Z. 2016. 101 formulaic alphas. *Wilmott*.

Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; and Liu, T.-Y. 2017. Lightgbm: A highly efficient gradient boosting decision tree. In *NeurIPS*.

Ke, G.; Xu, Z.; Zhang, J.; Bian, J.; and Liu, T.-Y. 2019. DeepGBM: A deep learning framework distilled by GBDT for online prediction tasks. In *KDD*.

Kelly, B. T.; Pruitt, S.; et al. 2019. Characteristics are covariances: A unified model of risk and return. *JFE*.

Koltchinskii, V.; Lounici, K.; et al. 2011. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *ANN STAT*.

Krishnapur, M. 2016. Anti-concentration inequalities. *Lecture notes*.

Laptev, N.; Yosinski, J.; Li, L. E.; and Smyl, S. 2017. Time-series extreme event forecasting with neural networks at uber. In *ICML*.

Li, C.; Wong, F.; Liu, Z.; and Kanade, V. 2017. From which world is your graph. In *NeurIPS*.

Li, Z.; Yang, D.; Zhao, L.; Bian, J.; Qin, T.; and Liu, T.-Y. 2019. Individualized indicator for all: Stock-wise technical indicator optimization with stock embedding. In *KDD*.

Liu, A.; Wu, Q.; Liu, Z.; and Xia, L. 2019. Near-neighbor methods in random preference completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 4336–4343.

Ming, F.; Wong, F.; Liu, Z.; and Chiang, M. 2014. Stock market prediction from WSJ: text mining via sparse matrix factorization. In *ICDM*.

Negahban, S.; and Wainwright, M. J. 2011. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *ANN STAT*.

Newey, W. K.; and West, K. D. 1986. A simple, positive semi-definite, heteroskedasticity and autocorrelationconsistent covariance matrix. *National Bureau of Economic Research Cambridge, Mass., USA*.

Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch. In *NIPS 2017 Workshop on Autodiff*.

Posner, R. A. 2014. *Economic analysis of law*. Aspen Publishing.

Qiong, W.; Brinton, C. G.; Zhang, Z.; Pizzoferrato, A.; Liu, Z.; and Cucuringu, M. 2021. Equity2Vec: End-to-end Deep Learning Framework for Cross-sectional Asset Pricing. *International Conference on AI in Finance*.

Quinlan, J. R. 1986. Induction of decision trees. *Machine learning*.

Rastelli, R.; Friel, N.; and Raftery, A. E. 2016. Properties of latent variable network models. *Network Science*.

Rather, A. M.; Agarwal, A.; and Sastry, V. 2015. Recurrent neural network and a hybrid model for prediction of stock returns. *EXPERT SYST APPL*.

Rendle, S.; Zhang, L.; and Koren, Y. 2019. On the difficulty of evaluating baselines: A study on recommender systems. *arXiv preprint arXiv:1905.01395*.

Rudelson, M.; and Vershynin, R. 2010. Non-asymptotic theory of random matrices: extreme singular values. In *ICM*.

Sak, H.; Senior, A. W.; and Beaufays, F. 2014. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. *arXiv preprint: arXiv:1402.1128*.

Sejnowski, T. J. 2018. *The deep learning revolution*. MIT press.

Shen, H.; Oh, J.; Zhao, S.; Wang, G.; Taghavi, T.; and Lee, S. 2022. Learning Personalized Representations using Graph Convolutional Network. In *17th International Workshop on Mining and Learning with Graphs, co-located with KDD 2022*.

Stewart, G. W.; and Sun, J.-g. 1990. *Matrix perturbation theory*. Academic press.

Sussman, D. L.; Tang, M.; and Priebe, C. E. 2013. Consistent latent position estimation and vertex classification for random dot product graphs. *TPAMI*.

Tang, M.; Sussman, D. L.; Priebe, C. E.; et al. 2013. Universally consistent vertex classification for latent positions graphs. *ANN STAT*.

Tao, T.; and Series, M. L. 2009. Compressed Sensing. *University of California*.

Tibshirani, R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*.

Tsigler, A.; and Bartlett, P. L. 2020. Benign overfitting in ridge regression. *arXiv preprint arXiv:2009.14286*.

Wang, J.; Zhang, Y.; Tang, K.; Wu, J.; and Xiong, Z. 2019. AlphaStock: A Buying-Winners-and-Selling-Losers Investment Strategy using Interpretable Deep Reinforcement Attention Networks. In *KDD*.

Wu, Q.; Hare, A.; Wang, S.; Tu, Y.; Liu, Z.; Brinton, C. G.; and Li, Y. 2021. Bats: a spectral biclustering approach to single document topic modeling and segmentation. *TIST*.

Wu, Q.; Hsu, W.-L.; Xu, T.; Liu, Z.; Ma, G.; Jacobson, G.; and Zhao, S. 2019. Speaking with actions-learning customer journey behavior. In *ICSC*.

Wu, Q.; Hui, L. C.; Yeung, C. Y.; and Chim, T. W. 2015. Early car collision prediction in VANET. In *ICCVE*. IEEE.

Wu, Q.; Wong, F.; Liu, Z.; Li, Y.; and Kanade, V. 2020a. Adaptive Reduced Rank Regression. In *NeurIPS*.

Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; and Philip, S. Y. 2020b. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*.

Wüthrich, B.; Permunetilleke, D.; et al. 1998. Daily prediction of major stock indices from textual www data. *Hkie transactions*.

Yang, X.; and Ding, J. 2020. A computational framework for iceberg and ship discrimination: Case study on Kaggle competition. *IEEE Access*, 8: 82320–82327.

Zhang, L.; Aggarwal, C.; and Qi, G.-J. 2017. Stock price prediction via discovering multi-frequency trading patterns. In *KDD*.

Zhao, S.; Hsu, W.-L.; Ma, G.; Xu, T.; Jacobson, G.; and Rustamov, R. 2020. Characterizing and Learning Representation on Customer Contact Journeys in Cellular Services. In *KDD*.