
Multi-Graph Tensor Networks

Yao Lei Xu, Kriton Konstantinidis, Danilo P. Mandic

Department of Electrical and Electronic Engineering

Imperial College London

{yao.xu15,k.konstantinidis19,d.mandic}@imperial.ac.uk

Abstract

The irregular and multi-modal nature of numerous modern data sources poses serious challenges for traditional deep learning algorithms. To this end, recent efforts have generalized existing algorithms to irregular domains through graphs, with the aim to gain additional insights from data through the underlying graph topology. At the same time, tensor-based methods have demonstrated promising results in bypassing the bottlenecks imposed by the Curse of Dimensionality. In this paper, we introduce a novel *Multi-Graph Tensor Network* (MGTN) framework, which exploits both the ability of graphs to handle irregular data sources and the compression properties of tensor networks in a deep learning setting. The potential of the proposed framework is demonstrated through an MGTN based deep Q agent for Foreign Exchange (FOREX) algorithmic trading. By virtue of the MGTN, a FOREX currency graph is leveraged to impose an economically meaningful structure on this demanding task, resulting in a highly superior performance against three competing models and at a drastically lower complexity.

1 Introduction

Deep learning techniques have been at the core of machine learning research for over a decade and have proved successful in a number of areas, including computer vision and speech processing [1]. However, as we enter the era of Big Data, the associated multi-modal and irregular nature of data is posing stern challenges to traditional learning systems, also owing to the sheer volume, variety, veracity and velocity of modern data sources [2]. To this end, it is necessary to generalize deep learning approaches to handle such irregular and multi-modal data.

Some of the most successful approaches for data analytics on irregular domains resort to graph signal processing techniques, because of their ability to provide insights into both the data acquisition process and their generative mechanisms [3]. By virtue of their ability to account for the underlying data structure, graph-based learning algorithms have proved advantageous in applications where the graph is known a-priori [4]. When it comes to exceedingly large multi-modal data, tensor-based methods have demonstrated their potential in areas including multi-modal learning [5], compression of large-dimensional data [6], and interpretability of neural networks [7]. In particular, tensor decomposition (TD) and tensor networks (TN) leverage the multi-modality of many Big Data applications to compress large-dimensional data while preserving their structure and interpretability, thus effectively bypassing the bottlenecks imposed by the Curse of Dimensionality [2].

Despite progress in these individual fields, the full potential arising from the combination of tensors, graphs, and neural networks has only begun to be explored. One such recent approach is the *Recurrent Graph Tensor Network* (RGTN) [8], which provides a framework for modelling multi-modal sequential data, through a unifying account of the expressive power of graphs and tensor networks. The initial RGTN model has been introduced primarily for sequential data and is therefore only defined on a single graph domain, often impractical for Big Data applications.

To provide a general framework that fully exploits the advantages of both graphs and tensors in a deep learning setting, we here generalize the RGTN concept in [8] to introduce the novel *Multi-Graph Tensor Network* (MGTN). In this way, the proposed MGTN is capable of handling irregular data residing on multiple graph domains, while simultaneously leveraging the compression properties of tensor networks to enhance modelling power and reduce parameter complexity.

The proposed model is verified on the task of Foreign Exchange (FOREX) algorithmic trading, a notoriously challenging paradigm characterized by highly irregular and noisy data [9]. More specifically, we consider a deep Q trading agent [10], where an MGTN is used to approximate the action values. By combining the advantages of graphs, tensors, and neural networks, the proposed MGTN agent is shown to yield highly superior performance against three competing agents. This setting is general enough to suggest the use of MGTN in a range of other application domains, including social networks, communication networks, and cognitive neuroscience.

The rest of the paper is organized as follows. We first briefly present the theoretical background necessary to follow this work in Section 2. The MGTN framework is introduced in Section 3, followed by an in-depth analysis of how FOREX algorithmic trading can be naturally modelled with an MGTN in Section 4. Numerical experiments are provided in Section 5 and Conclusion along with promising future research directions and potential MGTN application domains are given in Section 6.

2 Preliminaries

2.1 Tensors and Tensor Networks

A real-valued tensor is a multidimensional array, denoted by a calligraphic font, e.g., $\mathcal{X} \in \mathbb{R}^{I_1 \times \dots \times I_N}$, where N is the order of the tensor and I_n ($1 \leq n \leq N$) the size of its n^{th} mode. Matrices (denoted by bold capital letters, e.g., $\mathbf{X} \in \mathbb{R}^{I_1 \times I_2}$) can be seen as order-2 tensors ($N = 2$), vectors (denoted by bold lower-case letters, e.g., $\mathbf{x} \in \mathbb{R}^I$) can be seen as order-1 tensors ($N = 1$), and scalars (denoted by lower-case letters, e.g., $x \in \mathbb{R}$) are tensors of order $N = 0$. A specific entry of a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ is given by $x_{i_1, \dots, i_N} \in \mathbb{R}$. The tensor indices in this paper are grouped according to the Little-Endian convention [11].

Kronecker Product A (left) Kronecker product between two tensors, $\mathcal{A} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ and $\mathcal{B} \in \mathbb{R}^{J_1 \times \dots \times J_N}$, denoted by \otimes , yields a tensor $\mathcal{C} \in \mathbb{R}^{I_1 J_1 \times \dots \times I_N J_N}$, of the same order, with entries $c_{\overline{i_1 j_1}, \dots, \overline{i_N j_N}} = a_{i_1, \dots, i_N} b_{j_1, \dots, j_N}$, where $\overline{i_n j_n} = j_n + (i_n - 1)J_n$ [2].

Matricization and Tensorization The mode- n matricization of a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ reshapes the multidimensional array into a matrix $\mathbf{X}_{(n)} \in \mathbb{R}^{I_n \times I_1 I_2 \dots I_{n-1} I_{n+1} \dots I_N}$ with $(x_{(n)})_{i_n, \overline{i_1 \dots i_{n-1} i_{n+1} \dots i_N}} = x_{i_1, \dots, i_N}$. The inverse process, *Tensorization*, is denoted by $\text{ten}(\cdot)$.

Tensor Contraction An (m, n) -contraction [2] denoted by \times_n^m , between an order- N tensor $\mathcal{A} \in \mathbb{R}^{I_1 \times \dots \times I_n \times \dots \times I_N}$ and an order- M tensor $\mathcal{B} \in \mathbb{R}^{J_1 \times \dots \times J_m \times \dots \times J_M}$, where $I_n = J_m$, yields a third order- $(N + M - 2)$ tensor, $\mathcal{C} \in \mathbb{R}^{I_1 \times \dots \times I_{n-1} \times I_{n+1} \times \dots \times I_N \times J_1 \times \dots \times J_{m-1} \times J_{m+1} \times \dots \times J_M}$, where $c_{i_1, \dots, i_{n-1}, i_{n+1}, \dots, i_N, j_1, \dots, j_{m-1}, j_{m+1}, \dots, j_M} = \sum_{i_n=1}^{I_n} a_{i_1, \dots, i_{n-1}, i_n, i_{n+1}, \dots, i_N} b_{j_1, \dots, j_{m-1}, i_n, j_{m+1}, \dots, j_M}$.

Tensor Network A Tensor Network (TN) is a tensor architecture comprised of smaller-order core tensors which are connected by tensor contractions, where each tensor is represented as a node, while the number of edges that extends from that node corresponds to tensor order [6]. If two nodes are connected through an edge, it represents a linear contraction between two tensors over modes of equal dimensions.

Tensor Decomposition Special instances of tensor networks include those based on Tensor Decomposition (TD) methods, which approximate high-order, large-dimension tensors via contractions of smaller core tensors, therefore drastically reducing the computational complexity in tensor manipulation while preserving the data structure [6]. We here consider the Tensor-Train decomposition (TTD) [12], a highly efficient TD method that can decompose a large order- N tensor, $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$, into interconnected smaller core tensors, $\mathcal{G}^{(n)} \in \mathbb{R}^{R_{n-1} \times I_n \times R_n}$, as $\mathcal{X} = \mathcal{G}^{(1)} \times_2 \mathcal{G}^{(2)} \times_3 \mathcal{G}^{(3)} \times_3 \dots \times_3 \mathcal{G}^{(N)}$, where the set of R_n for $n = 0, \dots, N$ and $R_0 = R_N = 1$ is referred to as the *TT-rank*. The compression properties of TTD can be applied to significantly compress neural networks while maintaining comparable performance [13].

2.2 Graph Signal Processing

A graph, $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, is defined by a set of N vertices (or nodes) $v_n \subset \mathcal{V}$ for $n = 1, \dots, N$, and a set of edges connecting the n^{th} and m^{th} vertex $e_{n,m} = (v_n, v_m) \in \mathcal{E}$, for $n = 1, \dots, N$ and $m = 1, \dots, N$. A signal on a given graph is defined by a vector $\mathbf{f} \in \mathbb{R}^N$ such that $\mathbf{f} : \mathcal{V} \rightarrow \mathbb{R}$, which associates a signal value to every node on the graph [14]. A graph can be fully described in terms of its weighted adjacency matrix, $\mathbf{A} \in \mathbb{R}^{N \times N}$, such that $a_{n,m} > 0$ if $e_{n,m} \in \mathcal{E}$, and $a_{n,m} = 0$ if $e_{n,m} \notin \mathcal{E}$. The adjacency matrix can also be represented in its normalized form as $\tilde{\mathbf{A}} = \mathbf{D}^{\frac{1}{2}} \mathbf{A} \mathbf{D}^{\frac{1}{2}}$, where $\mathbf{D} \in \mathbb{R}^{N \times N}$ is the diagonal degree matrix such that $d_{n,n} = \sum_m a_{n,m}$.

Graph Shift Filters The weighted adjacency matrix can be used as a shift operator to filter signals on graphs. Such a graph filter represents a linear combination of vertex-shifted graph signals, which captures graph information at a local level [15]. For example, the operation $\mathbf{g} = (\mathbf{I} + \mathbf{A})\mathbf{f}$ produces a filtered signal, $\mathbf{g} \in \mathbb{R}^N$, such that $g_n = f_n + \sum_{m \in \Omega_n} a_{n,m} f_m$, where Ω_n denotes the 1-hop neighbours that are directly connected to the n -th node. For M graph signals stacked in a matrix form as $\mathbf{F} \in \mathbb{R}^{N \times M}$, the resulting graph filter can be compactly written as $\mathbf{G} = (\mathbf{I} + \mathbf{A})\mathbf{F}$ [15].

2.3 Recurrent Graph Tensor Networks

A Recurrent Graph Tensor Network (RGTN) [8] models sequential data through a time-based, multi-linear graph filter in tensor-network format.

gRGTN A general RGTN (gRGTN) extracts a feature map $\mathbf{Y} \in \mathbb{R}^{J_1 \times I_1}$ from sequential data, $\mathbf{X} \in \mathbb{R}^{J_0 \times I_1}$, via the forward pass given by $\mathbf{Y} = \sigma(\mathcal{R} \times_{3,4}^{1,2} \mathbf{W}^{(x)} \times_2^1 \mathbf{X})$, where $\mathbf{W}^{(x)} \in \mathbb{R}^{J_1 \times J_0}$ is the input weight matrix and $\mathcal{R} \in \mathbb{R}^{J_1 \times I_1 \times J_1 \times I_1}$ is the multi-linear time-graph filter. Specifically, \mathcal{R} is defined as $\mathcal{R} = \text{ten}(\mathbf{I} + (\mathbf{A} \otimes \mathbf{W}^{(r)}))$, where $\mathbf{I} \in \mathbb{R}^{J_1 I_1 \times J_1 I_1}$ is the identity matrix, $\mathbf{A} \in \mathbb{R}^{I_1 \times I_1}$ is the time-vertex based graph adjacency matrix with I_1 time-steps represented as graph nodes, and the weight matrix $\mathbf{W}^{(r)} \in \mathbb{R}^{J_1 \times J_1}$ models information propagation between successive time-steps over J_1 features.

fRGTN The fast RGTN (fRGTN) is defined by approximating $\mathbf{W}^{(r)} \approx \mathbf{I}$ in gRGTN, which leads to a reduced forward pass, $\mathbf{Y} = \sigma(\mathbf{R} \times_2^2 \mathbf{W}^{(x)} \times_2^1 \mathbf{X})$, where $\mathbf{R} \in \mathbb{R}^{I_1 \times I_1}$ is a standard graph shift filter defined as $\mathbf{R} = (\mathbf{I} + \mathbf{A})$, as discussed in Section 2.2. The lower parameter complexity over gRGTN is due to the identity approximation of $\mathbf{W}^{(r)}$.

fRGTN-TT If the problem is inherently multi-modal, then the large dense layer matrices of the fRGTN can be tensorized and represented in the Tensor-Train format, as discussed in [13]. This leads to the highly efficient fRGTN-TT model, which preserves the inherent multi-modality and has drastically lower parameter complexity.

2.4 Reinforcement Learning

In reinforcement learning (RL), at each time step t , a RL agent observes the state s_t of the environment, takes an action a_t and receives a reward r_t from the environment. The RL agent then learns a decision policy π^* from the pairs (a_t, r_t) , that is optimal in maximizing the long-term reward $r = \sum_t r_t$.

Q-Learning In Q-Learning, the goal is to learn a function $Q(s, a)$ that maps from the state s and action a to the expected accumulated reward. In other words, it estimates the reward that an agent will receive when it performs action a under the environment state s . More precisely, the function Q that we aim to learn is given by: $Q_\pi(s, a) = E[R_t | s_t = s, a_t = a, \pi]$, with π as the decision policy and R_t the accumulated reward given by $R_t = \sum_t \gamma^{t-t'} r_{t'}$, where $\gamma < 1$ is the discount factor used to give immediate rewards a higher value. The optimal Q-function that results in the best policy π^* which maximizes the expected reward is given by $Q^*(s, a) = \max_\pi Q_\pi(s, a)$. Within Q-Learning, this function is estimated iteratively using the Bellman equation, $Q^*(s, a) = r_t + \gamma \max_{a'} Q^*(s', a')$, where s' is the new state after taking action a under state s .

Deep Q-Learning A deep neural network in this context is trained by minimizing the difference between the two sides of the Bellman equation, using the loss function $L(\theta_i) = E_{(s,a) \sim \rho(\cdot)} [(y - Q(s, a; \theta_i))^2]$, where i denotes the training iteration and θ the weights of the deep Q-network. The training examples are in the form (s, a, r, s') , where $\rho(s, a)$ denotes their distribution, while y is

the prediction of $Q(s, a)$ given by the Bellman equation $y = r + \gamma \max_{a'} Q^*(s', a'; \theta_{i-1})|(s, a)$. To alleviate the issue of non-stationary targets, experience replay is used, whereby past experiences are stored in a buffer, from which a batch is sampled at every time instant to train the network with back-propagation and stochastic gradient descent.

Double Deep Q-Learning A notorious issue with the deep Q-Learning approach is its tendency towards instability and overestimation of action values. This is a result of the fact that both the current Q value, $Q(s, a)$, and estimated Q value, $y = r + \gamma \max_{a'} Q^*(s', a')$, are computed by a single Q-network. To mitigate this issue, we employ Double Deep Q-Learning [16], where a separate target network, \tilde{Q} , is used to compute the estimated Q value. The weights of \tilde{Q} are updated at the end of every training episode by hard copying the weights of Q .

3 Multi-Graph Tensor Networks

The RGTN model introduced in [8] was developed to model time-series problems related to sequential data, and is only defined for a single graph domain. To make the concept suitable for applications beyond time-series and in a Big Data setting, we next generalize the results from [8] to develop a *Multi-Graph Tensor Network*, which is capable of handling multi-modal data defined on multiple graph domains that is not limited to time-series modelling.

3.1 General Multi-Linear Graph Filter

The time-based multi-linear graph filter, \mathcal{R} , was developed in [8] to model time-series problems through a time-graph adjacency matrix that reflects the temporal flow of information, as discussed in Section 2.3. For this filter to be extended to other domains, the underlying graph topology needs to be modified. More generally, given a weighted graph adjacency matrix, $\mathbf{A} \in \mathbb{R}^{I_1 \times I_1}$, we can construct a multi-linear graph filter in the tensor domain, $\mathcal{F} \in \mathbb{R}^{J_1 \times I_1 \times J_1 \times I_1}$, as

$$\mathcal{F} = \text{ten}(\mathbf{I} + (\mathbf{A} \otimes \mathbf{P})) \quad (1)$$

where the propagation matrix $\mathbf{P} \in \mathbb{R}^{J_1 \times J_1}$ models the flow of information between neighbouring vertices (as opposed to successive time-steps in the RGTN case). This allows us to generalize the multi-linear graph filter \mathcal{F} to any given graph domain of any data modality.

3.2 General Multi-Graph Tensor Network

Consider a multi-graph learning problem where the input is an order- $(M + 1)$ tensor $\mathcal{X} \in \mathbb{R}^{J_0 \times I_1 \times I_2 \times \dots \times I_M}$ with J_0 features indexed along M physical modes $\{I_1, I_2, \dots, I_M\}$, such that a graph $\mathcal{G}^{(m)}$ is associated with each of the I_m modes, $m = 1, \dots, M$. For this problem, we define:

1. $\mathcal{A} = \{\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \dots, \mathbf{A}^{(M)}\}$, a set of adjacency matrices $\mathbf{A}^{(m)} \in \mathbb{R}^{I_m \times I_m}$ constructed from the corresponding graphs $\mathcal{G}^{(m)}$.
2. $\mathcal{W} = \{\mathbf{W}^{(1)}, \mathbf{W}^{(2)}, \dots, \mathbf{W}^{(M)}\}$, a set of weight matrices $\mathbf{W}^{(m)} \in \mathbb{R}^{J_m \times J_{m-1}}$ used for feature transforms, where J_m , for $m = 1, \dots, M$ controls the number of feature maps at every m .
3. $\mathcal{P} = \{\mathbf{P}^{(1)}, \mathbf{P}^{(2)}, \dots, \mathbf{P}^{(M)}\}$, a set of propagation matrices $\mathbf{P}^{(m)} \in \mathbb{R}^{J_m \times J_m}$, modelling the propagation of information over the neighbours of the graph $\mathcal{G}^{(m)}$.

Using the above objects and an optional activation function, $\sigma(\cdot)$, we can now define the general *Multi-Graph Tensor Network* (gMGTN) layer characterized by the following forward pass:

$$\mathcal{Y} = \sigma \left(\mathcal{F}^{(M)} \times_{3,4}^{1,M+1} \mathbf{W}^{(M)} \times_2^1 \dots \times_2^1 \mathcal{F}^{(2)} \times_{3,4}^{1,3} \mathbf{W}^{(2)} \times_2^1 \mathcal{F}^{(1)} \times_{3,4}^{1,2} \mathbf{W}^{(1)} \times_2^1 \mathcal{X} \right) \quad (2)$$

where $\mathcal{F}^{(m)} = \text{ten}(\mathbf{I} + (\mathbf{A}^{(m)} \otimes \mathbf{P}^{(m)}))$. The so defined forward pass generates a feature map, $\mathcal{Y} \in \mathbb{R}^{J_M \times I_1 \times \dots \times I_M}$, from the input tensor, \mathcal{X} , through a series of multi-linear graph filter and weight matrix contractions, which essentially iterates the graph filtering operation across all M graph domains.

3.3 Fast Multi-Graph Tensor Network

The gMGTN introduced in the previous section learns a propagation matrix, $\mathbf{P}^{(m)}$, and a weight matrix, $\mathbf{W}^{(m)}$, for each of the M graphs. For simplicity, let $J_m = J$ for $m = 1, \dots, M$; this results in a parameter complexity of $\mathcal{O}(MJ^2)$, which is linear in the number of graphs, M , and quadratic in the size of feature maps, J . This hinders the performance of gMGTN, since computation can become intractable for high dimensional multi-graph problems. To that end, we develop the fast *Multi-Graph Tensor Network* (fMGTN) as a low-complexity variant of the gMGTN.

Similar to [8], we can reduce the parameter complexity of gMGTN by: (i) approximating $\mathbf{P}^{(m)} \approx \mathbf{I}$ for $m = 1, \dots, M$; and (ii) using one single weight matrix, $\mathbf{W}^{(x)} \in \mathbb{R}^{J_1 \times J_0}$, for all of the graph domains, where J_1 controls the number of hidden units (feature maps). This allows us to define the fMGTN with the following reduced forward pass

$$\mathbf{y} = \sigma \left(\mathbf{F}^{(M)} \times_2^{M+1} \dots \times_2^4 \mathbf{F}^{(2)} \times_2^3 \mathbf{F}^{(1)} \times_2^2 \mathbf{W}^{(1)} \times_2^1 \mathcal{X} \right) \quad (3)$$

where $\mathbf{F}^{(m)} = (\mathbf{I} + \mathbf{A}^{(m)})$ is a standard graph shift filter as discussed in Section 2.2. In contrast to the gMGTN model, the proposed fMGTN does not have to learn $\mathbf{P}^{(m)}$ or $\mathbf{W}^{(m)}$, which reduces the parameter complexity of the forward pass to $\mathcal{O}(J^2)$ but at the cost of lower expressive power.

Finally, after extracting the feature map, $\mathbf{y} \in \mathbb{R}^{J_1 \times I_1 \times \dots \times I_M}$, it is customary to flatten the extracted features and pass them through dense layers to generate the final output. To further reduce the complexity, the weight matrices of the dense layers can be tensorized and represented in TT format, as discussed in [13]. This further reduces the number of parameters, while maintaining compatibility with the inherent multi-modal nature of the problem. For clarity, an example of a fMGTN model which implements this series of contractions is shown in Figure 1, using tensor network notation.

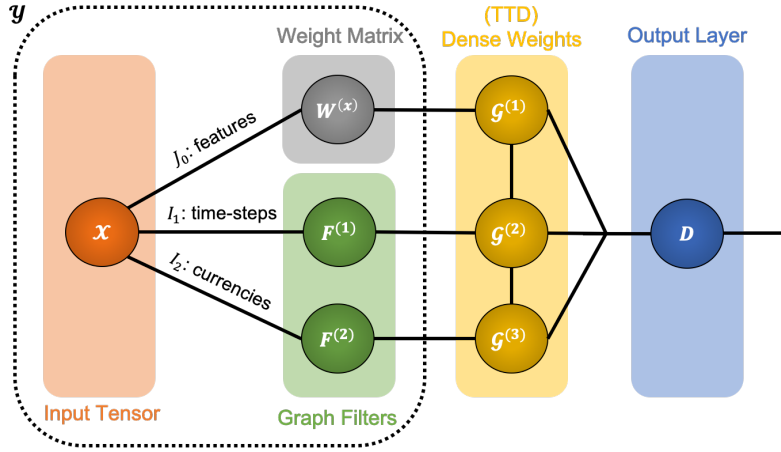


Figure 1: Tensor network representation of the fast Multi-Graph Tensor Network (fMGTN) used for our proposed experiment. The section encircled in dotted line denotes the multi-graph filtering operation for $M = 2$ as per equation (3). The yellow region denotes a tensorized dense layer weight matrix, represented in the Tensor-Train format. The input data used for our experiment is an order-3 tensor with $J_0 = 4$ pricing features, $I_1 = 30$ past time-steps, and $I_2 = 9$ currencies, as discussed in Section 5. Note that we define a time-domain graph filter and a currency-domain graph filter for input data modes of dimensions I_1 and I_2 , respectively.

4 Financial Environment

4.1 The Foreign Exchange Market and the Carry Graph

The FOREX market allows participants to trade pairs of currencies at a given *spot rate*, which measures the value of a currency with respect to another currency at a given instant (e.g. EUR/USD spot rate of 1.2 implies that 1 Euro can be exchanged for 1.2 US Dollars). Alternatively, the

participants can engage in forward contracts that allows them to exchange pairs of currencies on an agreed future date and at a specified *forward rate*. If the forward rate of a given currency pair is higher than the current spot rate, then the numerator currency is expected to increase in value against the denominator currency and vice-versa.

There are many factors that can affect the movements of spot rates, although the most important factor is arguably the *carry* factor: a tendency for high interest rate currencies to generate higher returns than the low interest rate ones. According to the *interest-rate-parity* theory [17], the expectation of currency pairs moving in different directions, depending on the interest rate differential, is reflected in the difference between the spot rate and the forward rate.

Therefore, for a pair of currencies, i and j , we can construct a pairwise carry signal by computing $c_{i,j} = 1 - \frac{r_f}{r_s}$, where r_f and r_s denote respectively the forward rate and the spot rate of the currency pair. Finally, we can construct a carry graph adjacency matrix \mathbf{A} such that its entries $a_{i,j}$ depend on the magnitude of the carry signal $c_{i,j}$. Figure 2 shows an example of the so constructed carry graph.




Figure 2: An example of FOREX carry graph. Thicker edges indicate stronger carry signal.

4.2 Characteristics of FOREX Data

The FOREX data are characterized by a number of properties that make classical machine learning techniques inadequate for their modelling; these include:

- FOREX data are multi-modal in nature, since they contain multiple pricing information indexed over time and across several related assets, which results in large dimensional tensors whose computation suffers from the Curse of Dimensionality.
- Financial data is known to have low signal-to-noise ratio due to the arbitrage forces in the market [9], which makes training particularly susceptible to overfitting.
- Various market factors can influence the pricing at different degrees depending on the time-horizon; this constitutes a multi-resolution problem that not many machine learning algorithms can handle.

The proposed MGTN is particularly suited to address the above challenges, as:

- The multi-modal nature of FOREX data naturally leads to a tensor representation, which can be readily handled by the tensor network structure of our proposed model.
- The model can leverage the powerful low-rank compression and regularization properties of tensor networks, which are inherently immune to the Curse of Dimensionality and provide a regularization framework via TD that does not degrade the underlying data structure.
- Long-term market factors such as *carry* can be encapsulated in graph filters that naturally appeal to the pair-wise formulation of the FOREX data; this allows to process high frequency pricing data through an economically meaningful low-frequency graph topology.

5 Experiments

To demonstrate the applicability and superiority of the proposed framework for the task of algorithmic trading, we use the fMGTN model as a feature extraction part of the deep Q network of a trading agent, and evaluate its performance against three commonly used agents based on: (i) Gated Recurrent Unit (GRU) Neural Network [18], (ii) Tensor-Train Neural Network (TTNN) [13], and (iii) Graph Convolutional Network (GCN) [19].

Data Description Minute-wise spot-rate pricing data were used for the period between October 1st 2019 and October 9th 2019 for a total of 9 currencies. Training took place over the first 7 days, while out-of-sample performance evaluation was performed in the last 2 days. Our features include Open/High/Low/Close (OHLC) spot-rates of the 9 currencies presented in Figure 2.

Data Pre-Processing The pricing data were pre-processed by computing the log-returns defined as $r_t = \ln(p_t) - \ln(p_{t-1})$, which measure the log-price difference between successive time-steps. The log-returns were then aggregated into multi-modal input samples, such that each sample, $\mathcal{X} \in \mathbb{R}^{J_0 \times I_1 \times I_2}$, contains log-returns indexed along $J_0 = 4$ features (OHLC), $I_1 = 30$ past time-steps, and $I_2 = 9$ currencies. Such multi-modal input data samples can be readily processed by the TTNN and the proposed fMGTN in their natural tensor form, as shown in Figure 1. However, the input samples were matricized as $\mathbf{X} \in \mathbb{R}^{I_1 \times J_0 I_2}$ and $\mathbf{X} \in \mathbb{R}^{I_2 \times J_0 I_1}$ for compatibility with the GRU and the GCN agents, respectively.

Graph Filter Formulation For the given FOREX input data, we formulate the problem as a multi-graph learning problem, where each sample contains J_0 features indexed along I_1 time-steps and I_2 currencies, whereby a time-graph and a currency-graph are associated with the time mode and the currency mode respectively. More precisely, we formulate the time-graph in the same way as in [8] for a total of $I_1 = 30$ time-steps, while the currency-graph was based on the carry-graph as discussed in Section 4.1. Finally, the respective graph filter for our fMGTN model was computed as discussed in Section 3.3. This results in graph filters $\mathbf{F}^{(1)}$ and $\mathbf{F}^{(2)}$, as illustrated in Figure 1.

Agent Specification For comparable results, the action value approximation network had the same specifications across all agents, with the sole difference being the feature extraction method. More specifically, each agent was based on a 3-layer architecture: (i) a feature extraction layer with 16 units and ReLU activation, (ii) a dense layer with 27 units and ReLU activation, and (iii) a linear output layer with 2 units corresponding to buy and sell action values. Each agent uses a different feature extraction layer, which can be based on fMGTN, GRU, TTNN, or GCN. In addition, due to their inherent tensor representation, the dense layer following the fMGTN and TTNN layers was tensorized and represented in TT format. All agents were trained using ADAM with a learning rate of $2 \cdot 10^{-4}$ and a mini-batch size of 64 for 15 episodes. Finally, the reward of the agents consisted of minute-wise log-returns. Our models were implemented using Tensorflow 2.3^{1 2}.

Performance Metrics Five different financial metrics were used to assess the performance of the agents: *Total return* measures the total percentage returns generated by the agent in the episode; *Sharpe ratio* measures the risk-adjusted return generated by the agent, which is computed as $\frac{\mu_r}{\sigma_r}$, where μ_r is the average log-return and σ_r the standard deviation of log-returns; *Sortino ratio* computes the risk-adjusted return as $\frac{\mu_r}{\sigma_r^d}$, where σ_r^d measures the standard deviation of negative log-returns; *Max Drawdown* measures the maximum percentage loss incurred by the agent during a consecutive period; *Hit Rate* measures the percentage of profitable trades to total trades.

Experimental Results Highly superior performance for the fMGTN based agent was obtained across a basket of European currencies, both in terms of generated profits and other common financial metrics. As shown in Figure 3, the fMGTN agent generated substantial profit (0.8%) during the out-of-sample testing period. Table 5 summarizes the performance of the considered agents, with fMGTN outperforming the other considered agents across a multitude of the most commonly used financial performance metrics. Finally, fMGTN achieved the best performance at a drastically lower parameter complexity, using up to 90% less trainable parameters compared to the GCN agent, and up to 80% less compared to the GRU agent, as shown in Table 2. This confirms the potential of the proposed framework for integrating graphs, tensors, and neural networks, which proved superior to any of the three components acting individually.

¹github.com/gylx/GTNR-Trading

²The GCN implementation was adapted from github.com/vermaMachineLearning/keras-deep-graph-learning




Figure 3: *Out-of-sample trading performance of the considered agents, averaged over five European currencies. The vertical axis represents the investment growth of an initial value of 1000\$, while the horizontal axis represents time in minutes.*

Agent	Total Return (%)	Sharpe Ratio	Sortino Ratio	Max Drawdown (%)	Hit Rate (%)
fMGTN	0.8018	0.0445	0.0604	0.2893	52.8056
GRU	0.0260	0.0012	0.0015	0.3477	50.4008
TTNN	0.1628	0.0064	0.0083	0.3493	50.6346
GCN	-0.0538	-0.0032	-0.0040	0.4180	50.2338

Table 1: *Performance comparison for the considered agents*

fMGTN	GRU	TTNN	GCN
531	3107	451	5891

Table 2: *Number of trainable parameters of the considered agents*

6 Conclusion

We have introduced a novel framework that embarks upon the advantages of both graphs and tensors, to provide an efficient and meaningful modelling strategy in a deep learning setting. The so introduced *Multi-Graph-Tensor-Network* (MGTN) has been shown to be capable of handling irregular data residing on multiple graph domains, while simultaneously leveraging on the compression properties of tensor networks to enhance the modelling power and drastically reduce parameter complexity. The effectiveness of the proposed model has been demonstrated on FOREX algorithmic trading, a challenging task owing to multiple sources of uncertainty and multi-modality. Experimental results have demonstrated the superiority of the proposed MGTN framework, which has generalized better in the highly irregular and noisy financial environment than the Recurrent Neural Network, Tensor-Train Neural Network, and Graph Convolutional Network based agents.

Future research directions include work leveraging on the versatility of the proposed framework to investigate its potential in numerous applications that share the modelling setup considered here. For example, the graph filters within the MGTN allow for the modelling of irregular data defined on one or multiple graph domains, a typical setting in social networks, recommender systems, and traffic forecasting. In addition, the tensor network structure of the MGTN allows for the modelling of high-dimensional data at a low complexity, which appeals to problems including multi-sensor processing, video classification, and natural language processing. Future research on spectral MGTN models can also potentially improve the modelling power of MGTN through spectral graph filtering techniques.

Broader Impact

The proposed multi-graph tensor network framework exploits the convergence of two successful but independently considered areas of graphs and tensor networks. This equips the MGTN framework with the ability to have considerable impact in applications where data sources come from irregular domains and are intrinsically high dimensional. This is a typical scenario in modern applications in finance, social networks, and physical sciences, to name but a few.

An immediate impact of our work, as illustrated through the considered experiment, may be in the creation of successful FOREX algorithmic trading strategies that efficiently model the underlying variable coupling through the proposed multi-graph domain. Another potential outreach of our work is on the modelling of molecule structures through large graphs, that could in turn lead to more efficient drug discovery.

The impact arising from model failure and data bias is identical to those of supervised learning models in general. Finally, we have not identified anyone that can be put at disadvantage from this work.

Acknowledgments and Disclosure of Funding

Y.L.X. is supported by an EPSRC Doctoral Scholarship. K.K. is supported by an EPSRC International Doctoral Scholarship.

References

- [1] Q. Zhang, L. T Yang, Z. Chen, and P. Li. A survey on deep learning for big data. *Information Fusion*, 42: 146–157, 2018.
- [2] A. Cichocki. Era of big data processing: A new approach via tensor networks and tensor decompositions. In *Proceedings of the International Workshop on Smart Info-Media Systems in Asia*, March 2014.
- [3] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98, 2013.
- [4] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–21, 2020.
- [5] A. Cichocki, D. Mandic, L. De Lathauwer, G. Zhou, Q. Zhao, C. Caiafa, and H. A. PHAN. Tensor decompositions for signal processing applications: From two-way to multiway component analysis. *IEEE Signal Processing Magazine*, 32(2):145–163, March 2015.
- [6] A. Cichocki, N. Lee, I. Oseledets, A. Phan, Q. Zhao, D. P. Mandic, et al. Tensor networks for dimensionality reduction and large-scale optimization: Part 1 low-rank tensor decompositions. *Foundations and Trends® in Machine Learning*, 9(4-5):249–429, 2016.
- [7] N. Cohen, O. Sharir, and A. Shashua. On the expressive power of deep learning: A tensor analysis. In *Proceedings of the Conference on Learning Theory*, pages 698–728, 2016.
- [8] Y. L. Xu and D. P. Mandic. Recurrent graph tensor networks. *arXiv preprint arXiv:2009.08727*, September 2020.
- [9] M. L. de Prado. *Machine Learning for Asset Managers*. Cambridge University Press, 2020.
- [10] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing Atari with deep reinforcement learning. *ArXiv e-prints*, December 2013.
- [11] S.V. Dolgov and D.V. Savostyanov. Alternating minimal energy methods for linear systems in higher dimensions. *SIAM Journal on Scientific Computing*, 36(5):A2248–A2271, 2014.
- [12] I. V. Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.
- [13] A. Novikov, D. Podoprikin, A. Osokin, and D. P. Vetrov. Tensorizing neural networks. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, pages 442–450, 2015.

- [14] L. Stankovic, D. Mandic, M. Dakovic, M. Brajovic, B. Scalzo, and T. Constantinides. Graph signal processing—Part I: Graphs, graph spectra, and spectral clustering. *arXiv preprint arXiv:1907.03467*, 2019.
- [15] L. Stankovic, D. Mandic, M. Dakovic, M. Brajovic, B. Scalzo, and A. G. Constantinides. Graph signal processing—Part II: Processing and analyzing signals on graphs. *arXiv preprint arXiv:1909.10325*, 2019.
- [16] H. V. Hasselt, A. Guez, and D. Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, September 2015.
- [17] R. Z. Aliber. The interest rate parity theorem: A reinterpretation. *Journal of political economy*, 81(6): 1451–1459, 1973.
- [18] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *Proceedings of NIPS 2014 Workshop on Deep Learning*, December 2014.
- [19] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.