
Autoregressive Convolutional Neural Networks for Asynchronous Time Series

Mikołaj Bińkowski^{1,2} Gautier Marti^{2,3} Philippe Donnat²

Abstract

We propose *Significance-Offset Convolutional Neural Network*, a deep convolutional network architecture for regression of multivariate asynchronous time series. The model is inspired by standard autoregressive (AR) models and gating mechanisms used in recurrent neural networks. It involves an AR-like weighting system, where the final predictor is obtained as a weighted sum of adjusted regressors, while the weights are data-dependent functions learnt through a convolutional network. The architecture was designed for applications on asynchronous time series and is evaluated on such datasets: a hedge fund proprietary dataset of over 2 million quotes for a credit derivative index, an artificially generated noisy autoregressive series and UCI household electricity consumption dataset. The proposed architecture achieves promising results as compared to convolutional and recurrent neural networks.

1. Introduction

Time series forecasting is focused on modeling the predictors of future values of time series given their past. As in many cases the relationship between past and future observations is not deterministic, this amounts to expressing the conditional probability distribution as a function of the past observations:

$$p(X_{t+d}|X_t, X_{t-1}, \dots) = f(X_t, X_{t-1}, \dots). \quad (1)$$

This forecasting problem has been approached almost independently by econometrics and machine learning communities.

In this paper we examine the capabilities of convolutional

neural networks (CNNs), (Lecun et al., 1998) in modeling the conditional mean of the distribution of future observations; in other words, the problem of *autoregression*. We focus on time series with multivariate and noisy signals. In particular, we work with financial data which has received limited *public* attention from the deep learning community and for which nonparametric methods are not commonly applied. Financial time series are particularly challenging to predict due to their low signal-to-noise ratio (cf. applications of Random Matrix Theory in econophysics (Laloux et al., 2000; Bun et al., 2017)) and heavy-tailed distributions (Cont, 2001). Moreover, the predictability of financial market returns remains an open problem and is discussed in many publications (cf. efficient market hypothesis of Fama (1970)).

A common situation with financial data is that the same signal (e.g. value of an asset) is observed from different *sources* (e.g. financial news, analysts, portfolio managers in hedge funds, market-makers in investment banks) in asynchronous moments of time. Each of these sources may have a different bias and noise with respect to the original signal that needs to be recovered (cf. time series in Figure 1). Moreover, these sources are usually strongly correlated and lead-lag relationships are possible (e.g. a market-maker with more clients can update its view more frequently and precisely than one with fewer clients). Therefore, the significance of each of the available past observations might be dependent on some other factors that can change in time. Hence, the traditional econometric models such as AR, VAR, VARMA (Hamilton, 1994) might not be sufficient. Yet their relatively good performance motivates coupling such linear models with deep neural networks that are capable of learning highly nonlinear relationships.

For these reasons, we propose *Significance-Offset Convolutional Neural Network*, a Convolutional Network extension of standard autoregressive models (Sims, 1972; 1980) equipped with a nonlinear weighting mechanism, and provide empirical evidence on its competitiveness with multi-layer CNNs, recurrent Long-Short Term Memory network (*LSTM*, Hochreiter & Schmidhuber (1997)) and *Phased LSTM* (Neil et al., 2016). The mechanism is inspired by the gating systems that proved successful in recurrent neural

¹Department of Mathematics, Imperial College London, London, UK ²Hellebore Capital Limited, London, UK ³Laboratoire d'informatique, Ecole Polytechnique, Palaiseau, France. Correspondence to: Mikołaj Bińkowski <mikbinkowski at gmail dot com>.




Figure 1. Quotes from four different market participants (sources) for the same CDS² throughout one day. Each trader displays from time to time the prices for which he offers to buy (*bid*) and sell (*ask*) the underlying CDS. The filled area marks the difference between the best sell and buy offers (*spread*) at each time.

networks (Hochreiter & Schmidhuber, 1997; Chung et al., 2014) and highway networks (Srivastava et al., 2015).

2. Related work

2.1. Time series forecasting

Literature in time series forecasting is rich and has a long history in the field of econometrics which makes extensive use of stochastic models such as AR, ARIMA and GARCH to mention a few. Unlike in machine learning, research in econometrics is more focused on explaining variables rather than improving out-of-sample prediction power. In practice, one can notice that these models ‘over-fit’ on financial time series: their parameters are unstable and out-of-sample performance is poor.

Reading through recent proceedings of the main machine learning venues (e.g. ICML, NIPS, AISTATS, UAI), one can notice that time series are often forecast using Gaussian processes (Petelin et al., 2011; Tobar et al., 2015; Hwang et al., 2016), especially when time series are irregularly sampled (Cunningham et al., 2012; Li & Marlin, 2016). Though still largely independent, researchers have started to “bring together the machine learning and econometrics communities” by building on top of their respective fundamental models yielding to, for example, the Gaussian Copula Process Volatility model (Wilson & Ghahramani, 2010). Our paper is in line with this emerging trend by coupling AR models and neural networks.

Over the past 5 years, deep neural networks have surpassed results from most of the existing literature in many fields (Schmidhuber, 2015): computer vision (Krizhevsky et al., 2012), audio signal processing and speech recognition (Sak et al., 2014), natural language processing (NLP) (Bengio

²iTraxx Europe Main Index, a tradable Credit Default Swap index of 125 investment grade rated European entities.

et al., 2003; Collobert & Weston, 2008; Grave et al., 2016; Jozefowicz et al., 2016). Although sequence modeling in NLP, i.e. prediction of the next character or word, is related to our forecasting problem (1), the nature of the sequences is too dissimilar to allow using the same cost functions and architectures. Same applies to the adversarial training proposed by Mathieu et al. (2016) for video frame prediction, as such approach favors *most plausible* scenarios rather than outputs *close* to all possible outputs, while the latter is usually required in financial time series due to stochasticity of the considered processes.

Literature on deep learning for time series forecasting is still scarce (cf. Gamboa (2017) for a recent review). Literature on deep learning for *financial* time series forecasting is even scarcer though interest in using neural networks for financial predictions is not new (Mozer, 1993; McNelis, 2005). More recent papers include Sirignano (2016) that used 4-layer perceptrons in modeling price change distributions in Limit Order Books, and Borovykh et al. (2017) who applied more recent WaveNet architecture (van den Oord et al., 2016a) to several short univariate and bivariate time-series (including financial ones). Despite claim of applying deep learning, Heaton et al. (2016) use autoencoders with a single hidden layer to compress multivariate financial data. Neil et al. (2016) present augmentation of LSTM architecture suitable for asynchronous series, which stimulates learning dependencies of different frequencies through *time gate*. In this paper, we investigate the capabilities of several architectures (CNN, Residual Network, multi-layer LSTM and Phased LSTM) on AR-like artificial asynchronous and noisy time series, household electricity consumption dataset, and on real financial data from the credit default swap market where some inefficiencies may exist, i.e. time series may not be totally random.

2.2. Gating and weighting mechanisms

Gating mechanisms for neural networks were first proposed by Hochreiter & Schmidhuber (1997) and proved essential in training recurrent architectures (Jozefowicz et al., 2016) due to their ability to overcome the problem of vanishing gradient. In general, they can be expressed as

$$f(x) = c(x) \otimes \sigma(x), \quad (2)$$

where f is the output function, c is a ‘candidate output’ (usually a nonlinear function of x), \otimes is an element-wise matrix product and $\sigma : \mathbb{R} \rightarrow [0, 1]$ is a sigmoid nonlinearity that controls the amount of the output passed to the next layer (or to further operations within a layer). Appropriate compositions of functions of type (2) lead to the popular recurrent architectures such as LSTM (Hochreiter & Schmidhuber, 1997) and GRU (Chung et al., 2014).

A similar idea was recently used in construction of highway

networks (Srivastava et al., 2015) which enabled successful training of deeper architectures. van den Oord et al. (2016b) and Dauphin et al. (2016) proposed gating mechanisms (respectively with hyperbolic tangent and linear ‘candidate outputs’) for training deep convolutional neural networks.

The gating system that we propose is aimed at weighting a number of different ‘candidate predictors’ and therefore is most closely related to the *softmax* gating used in MuFuRU (Multi-Function Recurrent Unit, Weissenborn & Rocktäschel (2016)), i.e.

$$f(x) = \sum_{l=1}^L p^l(x) \otimes f^l(x), \quad p(x) = \text{softmax}(\hat{p}(x)), \quad (3)$$

where $(f^l)_{l=1}^L$ are candidate outputs (*composition operators* in MuFuRu) and $(\hat{p}^l)_{l=1}^L$ are linear functions of the inputs.

The idea of weighting outputs of the intermediate layers within a neural networks is also used in attention networks (Cho et al., 2015) that proved successful in such tasks as image captioning and machine translation. Our approach is similar as the separate inputs (time series steps) are weighted in accordance with learned functions of these inputs, yet different since we model these functions as multi-layer CNNs (instead of projections on learned directions) and since we do not use recurrent layers. The latter is important in the above mentioned tasks as it enables the network to remember the parts of the sentence/image already translated/described.

3. Motivation

Time series observed in irregular moments of time make up significant challenges for learning algorithms. Gaussian processes provide useful theoretical framework capable of handling asynchronous data; however, due to assumed Gaussianity they are inappropriate for financial datasets, which often follow fat-tailed distributions (Cont, 2001). On the other hand, prediction of even a simple autoregressive time series such as AR(2) given by $\hat{X}(t) = \alpha X(t-1) + \beta X(t-2) + \varepsilon(t)$ ³ may involve highly non-linear functions when sampled irregularly. Precisely, it can be shown that the conditional expectation

$$\mathbb{E}[X(t)|X(t-1), X(t-k), k] = a_k X(t-1) + b_k X(t-k), \quad (4)$$

where a_k and b_k are rational functions of α and β .⁴ It would not be a problem if k was fixed, as then one would be interested in estimating a_k and b_k directly; however, this is not the case with asynchronous sampling. When X is an autoregressive series of higher order and more

³Where $\varepsilon(t)$ is an error term independent of $\{X(s) : s < t\}$.

⁴See Appendix A for the proof. Appendices are available in the online version at <https://arxiv.org/abs/1703.04122>.

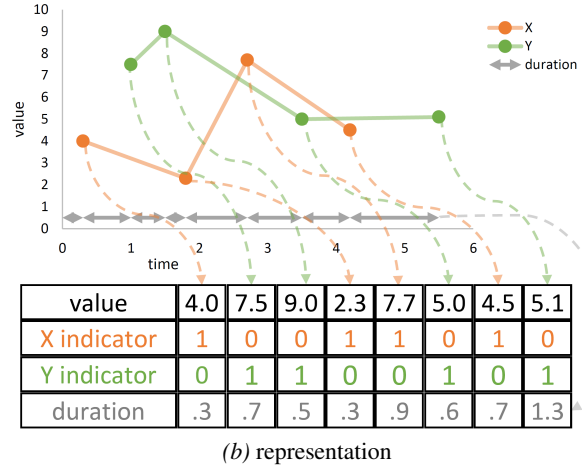



Figure 2. (a) Fixed sampling frequency and its drawbacks; keeping all available information leads to much more datapoints. (b) Proposed data representation for the asynchronous series. Consecutive observations are stored together as a single *value* series, regardless of which series they belong to; this information, however, is stored in *indicator* features, alongside durations between observations.

past observations are available, the analogous expectation $\mathbb{E}[X(t)|\{X(t-m), m = 1, \dots, M\}]$ would involve more complicated functions that in general may not allow closed-form expression.

In real-world applications we often deal with multivariate time series whose dimensions are observed separately and asynchronously. This adds even more difficulty to assigning appropriate weights to the past values, even if the underlying data generating process is linear. Furthermore, appropriate representation of such series might be not obvious as aligning such series at fixed frequency may lead to loss of information (if too low frequency is chosen) or prohibitive enlargement of the dataset (especially when durations vary by levels of magnitude), see Figure 2a. As an alternative, we might consider representing separate dimensions as a single one with dimension and duration indicators as additional features. Figure 2b presents this approach, which is going to be at the core of the proposed architecture. Note that this representation is also natural for Phased LSTM, where time dimension (a cumulative sum of durations) is used to open/close the time gate.

However, even regular LSTMs have been successfully applied with time/durations as explanatory variables (Żoła & Romański, 2017; Kumar et al., 2018). Let $x_n = (X(t_n), t_n - t_{n-1})$ be a series of pairs of consecutive input

values and respective durations. One may expect that, given such input series, LSTM in each step would memorize the values and weight them at the output according to the durations. Yet, there is a drawback of such approach, which perhaps lies in imbalance between the needs for memory and for nonlinearity: the weights that such network needs to assign to the memorized observations potentially require several layers of nonlinearity to be computed properly, while past observations might just need to be memorized as they are.

For these reasons we shall consider a model that combines simple autoregressive approach with a neural network in order to allow learning meaningful data-dependent weights

$$\begin{aligned} \mathbb{E}[x_n | \{x_{n-m}, m = 1, \dots, M\}] &= \\ &= \sum_{m=1}^M \alpha_m(x_{n-m}) \cdot x_{n-m} \end{aligned} \quad (5)$$

where $(\alpha_m)_{m=1}^M$ satisfying $\alpha_1 + \dots + \alpha_M \leq 1$ are modeled using a neural network. To allow more flexibility and cover situations when e.g. observed values of x are biased, we should consider the summation over terms $\alpha_m(x_{n-m}) \cdot f(x_{n-m})$, where f is also a neural network. We formalize this idea in Section 4.

4. Model Architecture

Suppose that we are given a multivariate time series $(x_n)_{n=0}^{\infty} \subset \mathbb{R}^d$ and we aim to predict the conditional future values of a subset of elements of x_n

$$y_n = \mathbb{E}[x_n^I | \{x_{n-m}, m = 1, 2, \dots\}], \quad (6)$$

where $I = \{i_1, i_2, \dots, i_{d_I}\} \subset \{1, 2, \dots, d\}$ is a subset of features of x_n . Let $\mathbf{x}_n^{-M} = (x_{n-m})_{m=1}^M$. We consider the following estimator of y_n

$$\hat{y}_n = \sum_{m=1}^M [F(\mathbf{x}_n^{-M}) \otimes \sigma(S(\mathbf{x}_n^{-M}))]_{\cdot, m}, \quad (7)$$

where

- $F, S : \mathbb{R}^{d \times M} \rightarrow \mathbb{R}^{d_I \times M}$ are neural networks described below,
- σ is a normalized activation function independent at each row, i.e.

$$\sigma((a_1^T, \dots, a_{d_I}^T)^T) = (\sigma(a_1)^T, \dots, \sigma(a_{d_I})^T)^T \quad (8)$$

for any $a_1, \dots, a_{d_I} \in \mathbb{R}^M$ and σ such that $\sigma(a)^T \mathbf{1}_M = 1$ for any $a \in \mathbb{R}^M$.

- \otimes is Hadamard (element-wise) matrix multiplication.

- $A_{\cdot, m}$ denotes the m -th column of a matrix A . Note that $\sum_{m=1}^M A_{\cdot, m} = A \cdot \underbrace{(1, 1, \dots, 1)^T}_{\times M}$.

The summation in Eq. (7) goes over the columns of the matrix in bracket; hence each i -th element of the output vector \hat{y}_n is a linear combination of the respective i -th row of the matrix $F(\mathbf{x}_n^{-M})$. We are going to consider S to be a fully convolutional network (composed solely of convolutional layers) and F of the form

$$F(\mathbf{x}_n^{-M}) = W \otimes [\text{off}(x_{n-m}) + x_{n-m}^I]_{m=1}^M \quad (9)$$

where $W \in \mathbb{R}^{d_I \times M}$ and $\text{off} : \mathbb{R}^d \rightarrow \mathbb{R}^{d_I}$ is a multilayer perceptron. In that case F can be seen as a sum of projection ($\mathbf{x} \mapsto \mathbf{x}^I$) and a convolutional network with all kernels of length 1. Equation (7) can be rewritten as

$$\hat{y}_n = \sum_{m=1}^M W_{\cdot, m} \otimes (\text{off}(x_{n-m}) + x_{n-m}^I) \otimes \sigma(S_{\cdot, m}(\mathbf{x}_n^{-M})). \quad (10)$$

We will call the proposed network a *Significance-Offset Convolutional Neural Network* (SOCNN), while off and S respectively the *offset* and *significance* (sub)networks. The network scheme is shown in Figure 3. Note that when $\text{off} \equiv 0$ and $\sigma \equiv 1$ the model simplifies to the collection of d_I separate $AR(M)$ models for each dimension.

Interpretation of the components

Note that the form of Equation (10) enforces the separation of temporal dependence (obtained in weights W_m), the local significance of observations S_m (S as a convolutional network is determined by its filters which capture local dependencies and are independent of the relative position in time) and the predictors $\text{off}(x_{n-m})$ that are completely independent of position in time. This provides some amount of interpretability of the fitted functions and weights. For instance, each of the past observations provides an adjusted single regressor for the target variable through the offset network. Note that due to asynchronous sampling procedure, consecutive values of x might be heterogenous (e.g. if they represent different signals); therefore, adjustment from the offset network is important in such cases. On the other hand, significance network provides data-dependent weights for all regressors and sums them up in an autoregressive manner. Sample significance and offset activations of the trained network are presented in Appendix E.2.

Relation to asynchronous data

As mentioned before, one of the common problems with time series are the varying durations between consecutive observations. A simple approach at data-preprocessing level is aligning the observations at some chosen frequency by e.g. duplicating or interpolating observations. This, however, might extend the size of an input and, therefore, model complexity.




Figure 3. A scheme of the proposed SOCNN architecture. The network preserves the time-dimension up to the top layer, while the number of features per timestep (filters) in the hidden layers is custom. The last convolutional layer, however, has the number of filters equal to dimension of the output. The *Weighting* frame shows how outputs from offset and significance networks are combined in accordance with Eq. (10).

The other idea is to treat the duration and/or time of the observation as another feature, as presented in Figure 2b. This approach is at the core of the SOCNN architecture: the significance network is aimed at learning the high-level features that indicate the relative importance of past observations, which, as shown in Section 3, could be predominantly dependent on time and duration between observations.

Loss function

L^2 error is a natural loss function for the estimators of expected value

$$L^2(y, y') = \|y - y'\|^2. \quad (11)$$

As mentioned above, the output of the offset network can be seen as a collection of separate predictors of the changes between corresponding observations x_{n-m}^I and the target variable y_n

$$\text{off}(x_{n-m}) \simeq y_n - x_{n-m}^I. \quad (12)$$

For that reason, we consider the *auxiliary loss* function equal to mean squared error of such intermediate predictions

$$L^{aux}(\mathbf{x}_n^{-M}, y_n) = \frac{1}{M} \sum_{m=1}^M \|\text{off}(x_{n-m}) + x_{n-m}^I - y_n\|^2. \quad (13)$$

The total loss for the sample (\mathbf{x}_n^{-M}, y_n) is therefore given by

$$L^{tot}(\mathbf{x}_n^{-M}, y_n) = L^2(\hat{y}_n, y_n) + \alpha L^{aux}(\mathbf{x}_n^{-M}, y_n), \quad (14)$$

where \hat{y}_n is given by Eq. (10) and $\alpha \geq 0$ is a constant. In Section 5.3 we discuss the empirical findings on the impact of positive values of α on the model training and performance, as compared to $\alpha = 0$ (lack of auxiliary loss).

5. Experiments

We evaluate the proposed model on artificially generated datasets, household electric power consumption dataset available from UCI repository (Lichman, 2013), and the financial dataset of bid/ask quotes sent by several market participants active in the credit derivatives market, comparing its performance with simple CNN, single- and multi-layer LSTM (Hochreiter & Schmidhuber, 1997), Phased LSTM (Neil et al., 2016) and 25-layer ResNet (He et al., 2015).

Apart from performance evaluation of SOCNNs, we discuss the impact of the network components, such as auxiliary loss and the depth of the offset sub-network.

Code for the experiments is available online at <https://github.com/mbinkowski/nntimeseries>.

5.1. Datasets

Artificial data

We test our network architecture on the artificially generated datasets of multivariate time series. We consider two types of series:

1. *Synchronous series*. The series of K noisy copies ('sources') of the same univariate autoregressive series ('base series'), observed together at random times. The noise of each copy is of different type.
2. *Asynchronous series*. The series of observations of one of the sources in the above dataset. At each time, the source is selected randomly and its value at this time is added to form a new univariate series. The final series is composed of this series, the durations between random times and the indicators of the 'available source' at each time.

The details of the simulation process are presented in Appendix C⁵.




Figure 4. Simulated synchronous (left) and asynchronous (right) artificial series. Note the different durations between the observations from different sources in the latter plot. For clarity, we present only 6 out of 16 total dimensions.

We consider synchronous and asynchronous series $X_{K \times N}$ where $K \in \{16, 64\}$ is the number of sources and $N = 10,000$, which gives 4 artificial series in total.⁶ Figure 4 presents a sample from two of the generated datasets.

Electricity data

The UCI household electricity dataset⁷ contains measurements of 7 different quantities related to electricity consumption in a single household, recorded every minute for 47 months, yielding over 2 million observations. Since we aim to focus on asynchronous time-series, we alter it so that a single observation contains only value of one of the seven features, while durations between consecutive observations range from 1 to 7 minutes.⁸ The regression aim is to predict all of the features at the next time step.

Non-anonymous quotes

The proposed model was designed primarily for forecasting incoming non-anonymous quotes received from the credit default swap market. The dataset contains 2.1 million quotes from 28 different *sources*, i.e. market participants. Each quote is characterized by 31 features: the offered price, 28 indicators of the quoting source, the *direction* indicator (the quote refers to either a buy or a sell offer) and duration from the previous quote. For each source and direction we aim at predicting the next quoted price from this given source and direction considering the last 60 quotes. We formed 15

⁶Note that a series with K sources is $K + 1$ -dimensional in synchronous case and $K + 2$ -dimensional in asynchronous case. The base series in all processes was a stationary AR(10) series. Although that series has the true order of 10, in the experimental setting the input data included past 60 observations. The rationale behind that is twofold: not only is the data observed in irregular random times but also in real-life problems the order of the model is unknown.

⁷Dheeru & Karra Taniskidou (2017), available at UCI Machine Learning Repository website <https://archive.ics.uci.edu/ml/datasets/individual+household+electric+power+consumption>

⁸The exact details of preprocessing can be found in Appendix D.

separate prediction tasks; in each task the model was trained to predict the next quote by one of the fifteen most active market participants.⁹

This dataset, which is proprietary, motivated the construction of the aforementioned artificial *asynchronous* time series datasets which mimic some of its statistical features.

5.2. Training details

To evaluate the model and the significance of its components, we perform a grid search over some of the hyperparameters, more extensively on the artificial and electricity datasets. These include the offset sub-network’s depth (we consider depths of 1 and 10 for artificial and electricity datasets; 1 for Quotes data) and the auxiliary weight α (compared values: $\{0, 0.1, 0.01\}$). For all networks we have chosen LeakyReLU activation function (15)

$$\sigma^{\text{LeakyReLU}}(x) = x \text{ if } x \geq 0, \quad ax \text{ otherwise.} \quad (15)$$

with leak rate $a = .1$ as an activation function.

Benchmark networks

We compare the performance of the proposed model with CNN, ResNet, multi-layer LSTM and Phased LSTM networks, and linear (VAR) model, choosing the best hyperparameter setting through grid search. The benchmark networks were designed so that they have a comparable number of parameters as the proposed model. Consequently, LeakyReLU activation function (15) with leak rate .1 was used in all layers except the top ones where linear activation was applied. For CNN we provided the same number of layers, same stride (1) and similar kernel size structure. In each trained CNN, we applied max pooling with the pool

⁹This separation is related to data normalization purposes and different magnitudes of the levels of predictability for different market participants. The quotes from the remaining 13 participants were not selected for prediction as their market presence was too short or too irregular to form reliable training, validation and test samples.

Table 1. Configurations of the trained models. f - number of convolutional filters/memory cell size in LSTM, ks - kernel size, p - dropout rate, $clip$ - gradient clipping threshold, $conv$ - ($k \times 1$) convolution with kernel size k indicated in the ks column, $conv1$ - (1×1) convolution. Apart from the listed layers, each network has a single fully connected layer on the top. Kernel sizes (3, 1) ((1, 3, 1)) denote alternating kernel sizes 3 and 1 (1, 3 and 1) in successive convolutional layers. We also optimized a parameter specific to Phased LSTM, the *initialized period*, considering two settings: [1, 1000] and [.01, 10].

Artificial & Electricity Datasets					
Model	layers	f	ks	p	clip
SOCNN	10conv + {1, 10}conv1	{8, 16}	{(3, 1), 3}	0	{1, .001}
CNN	7conv + 3pool	{16, 32}	{(3, 1), 3}	{0, .5}	{1, .001}
LSTM	{1, 2, 3, 4}	{16, 32}	-	{0, .5}	{1, .001}
Phased LSTM	1	{16, 32}	-	0	{1, .001}
ResNet	22conv + 3pool	16	(1, 3, 1)	{0, .5}	{1, .001}
Quotes Dataset					
Model	layers	f	ks	p	clip
SOCNN	7conv + {1, 7}conv1	8	{(3, 1), 3}	.5	.01
CNN	7conv + 3pool	{16, 32}	{(3, 1), 3}	.5	.01
LSTM	{1, 2, 3}	{32}	-	.5	.0001 ¹⁰
Phased LSTM	1	{16, 32}	-	0	.01
ResNet	22conv + 3pool	16	(1, 3, 1)	.5	.01

size of 2 every two convolutional layers.¹¹ Table 1 presents the configurations of the network hyperparameters used in comparison.

Network Training

The training and validation sets were sampled randomly from the first 80% of timesteps in each series, with ratio 3 to 1. The remaining 20% of data was used as a test set. All models were trained using Adam optimizer (Kingma & Ba, 2015) which we found much faster than standard Stochastic Gradient Descent in early tests. We used a batch size of 128 for artificial and electricity data, and 256 for quotes dataset. We also applied batch normalization (Ioffe & Szegedy, 2015) in between each convolution and the following activation. At the beginning of each epoch, the training samples were shuffled. To prevent overfitting we applied *dropout* and *early stopping*.¹² Weights were initialized following the *normalized uniform* procedure proposed by Glorot & Bengio (2010). Experiments were carried out using implementation relying on Tensorflow (Abadi et al., 2016) and Keras (Chollet, 2015). For artificial and electricity data we optimized the models using one K20s NVIDIA GPU while for quotes dataset we used 8-core Intel Core

¹⁰We found LSTMs very unstable on quotes dataset without gradient clipping or with higher clipping boundary.

¹¹Hence layers 3, 6 and 9 were pooling layers, while layers 1, 2, 4, 5, . . . were convolutional layers.

¹²Whenever 10 consecutive epochs did not bring improvement in the validation error, the learning rate was reduced by a factor of 10 and the best weights obtained till then were restored. After the second reduction and another 10 consecutive epochs without improvement, the training was stopped. The initial learning rate was set to .001.

i7-6700 CPU machine only.

5.3. Results

Table 2 presents the detailed results from all datasets. The proposed networks outperform significantly the benchmark networks on the asynchronous, electricity and quotes datasets. For the synchronous datasets, on the other hand, SOCNN almost matches the results of the benchmarks. This similar performance could have been anticipated - the correct weights of the past values in synchronous artificial datasets are far less nonlinear than in case when separate dimensions are observed asynchronously. For this reason, the significance network’s potential is not fully utilized.

We can also observe that the depth of the offset network has negligible or negative impact on the results achieved by the SOCNN network. This means that the significance network has crucial impact on the performance, which is in-line with the potential drawbacks of the LSTM network discussed in Section 3: obtaining proper weights for the past observations is much more challenging than getting good predictors from the single past values.

It is worth noticing that Phased LSTM did not perform well with asynchronous artificial and electricity datasets.

For the Quotes dataset, the proposed model was the best one for 13 out of 15 tasks and the only one to always beat VAR model. Surprisingly, for each of the other networks it was difficult to exceed the benchmark set by simple linear model. We also found benchmark networks to have very unstable test loss during training in some cases, despite convergence of the training error. Especially, for two of the tasks, Phased LSTM and ResNet gave very high test errors.¹³ The auxiliary loss was usually found to have minor importance, though in many cases it led to best results. Phased LSTM performed very well in most of the tasks, being the best model in two of them and runner-up in another eight. However, in some cases it achieved quite bad error rates, which causes rather average overall mean error.

The small positive auxiliary weight helped achieve more stable test error throughout training in many cases. The higher weights of auxiliary loss considerably improved the test error on asynchronous datasets (See Table 3); however for other datasets its impact was negligible. In general, the proposed SOCNN had significantly lower variance of the test and validation errors, especially in the early stage of the training process and for quotes dataset. This effect can be seen in the learning curves for *Asynchronous 64* artificial dataset presented in Figure 5.

In Appendix E.2 we visualize significance and offset activa-

¹³The exact results for all tasks for Quotes dataset can be found in Appendix E.

Table 2. Detailed results for all datasets. For each model, we present the average and standard deviation (in parentheses) mean squared error obtained on the out-of-sample test set. The best results for each dataset are marked by bold font. *SOCNN1* (*SOCNN1+*) denote proposed models with one (10 or 7) offset sub-network layers. For quotes dataset the presented values are averaged mean-squared errors from 6 separate prediction tasks, normalized according to the error obtained by VAR model.

model	VAR	CNN	ResNet	LSTM	Phased LSTM	SOCNN1	SOCNN1+
Synchronous 16	0.841 (0.000)	0.154 (0.003)	0.152 (0.001)	0.151 (0.001)	0.166 (0.026)	0.152 (0.001)	0.172 (0.001)
Synchronous 64	0.364 (0.000)	0.029 (0.001)	0.029 (0.001)	0.028 (0.000)	0.038 (0.004)	0.030 (0.001)	0.032 (0.001)
Asynchronous 16	0.577 (0.000)	0.080 (0.032)	0.059 (0.017)	0.038 (0.008)	1.021 (0.090)	0.019 (0.003)	0.026 (0.004)
Asynchronous 64	0.318 (0.000)	0.078 (0.029)	0.087 (0.014)	0.065 (0.020)	0.924 (0.119)	0.035 (0.006)	0.044 (0.118)
Electricity	0.729 (0.005)	0.371 (0.005)	0.394 (0.013)	0.461 (0.011)	0.744 (0.015)	0.163 (0.010)	0.165 (0.012)
Quotes	1.000 (0.019)	0.897 (0.019)	2.245 (0.179)	0.827 (0.024)	0.945 (0.034)	0.387 (0.016)	–




Figure 5. Learning curves with different auxiliary weights for SOCNN model trained on *Asynchronous 64* dataset. The solid lines indicate the test error while the dashed lines indicate the training error.




Figure 6. Experiment comparing robustness of the considered networks for *Asynchronous 16* dataset. The plots show how the error would change if an additional noise term was added to the input series. The dotted curves show the total significance and average absolute offset (not to scale) outputs for the noisy observations. Interestingly, significance of the noisy observations increases with the magnitude of noise; i.e. noisy observations are far from being discarded by SOCNN.

Table 3. MSE for different values of α for two artificial datasets.

α	Async 16	async 64
0.0	0.0284	0.0624
0.01	0.0253	0.0434
0.1	0.0172	0.0323

tions of the network trained for *electricity* dataset.

Model robustness

To understand better why SOCNN obtained better results than the other networks, we check how these networks react to the presence of additional noise in the input terms.¹⁴ Figure 6 presents changes in the mean squared error and significance and offset network outputs with respect to the level of noise. SOCNN is the most robust out of the compared networks and, together with single-layer LSTM, least prone to overfitting. Despite the use of dropout and cross-validation, the other models tend to overfit the training set and have non-symmetric error curves on test dataset.

¹⁴The details of the added noise terms are presented in the Appendix B.

6. Conclusion and discussion

In this article, we proposed a weighting mechanism that, coupled with convolutional networks, forms a new neural network architecture for time series prediction. The proposed architecture is designed for regression tasks on asynchronous signals in the presence of high amount of noise. This approach has proved to be successful in forecasting several asynchronous time series outperforming popular convolutional and recurrent networks.

The proposed model can be extended further by adding intermediate weighting layers of the same type in the network structure. Another possible generalization that requires further empirical studies can be obtained by leaving the assumption of independent offset values for each past observation, i.e. considering not only 1x1 convolutional kernels in the offset sub-network.

Finally, we aim at testing the performance of the proposed architecture on other real-life datasets with relevant characteristics. We observe that there exists a strong need for common ‘econometric’ datasets benchmark and, more generally, for time series (stochastic processes) regression.

ACKNOWLEDGEMENTS

Authors would like to thank Hellebore Capital Ltd. for providing data for the experiments. M.B. thanks Engineering and Physical Sciences Research Council (EPSRC) for financial support for this research.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: A system for large-scale machine learning, May 2016. URL <http://arxiv.org/abs/1605.08695>.
- Bengio, Y., Ducharme, R., Vincent, P., and Janvin, C. A neural probabilistic language model. *J. Mach. Learn. Res.*, 3:1137–1155, March 2003. ISSN 1532-4435. URL <http://portal.acm.org/citation.cfm?id=944966>.
- Borovykh, A., Bohte, S., and Oosterlee, C. W. Conditional time series forecasting with convolutional neural networks, March 2017. URL <http://arxiv.org/abs/1703.04691v1.pdf>.
- Bun, J., Bouchaud, J.-P., and Potters, M. Cleaning large correlation matrices: tools from random matrix theory. *Physics Reports*, 666:1–109, 2017.
- Cho, K., Courville, A., and Bengio, Y. Describing multimedia content using attention-based Encoder–Decoder networks. *IEEE Transactions on Multimedia*, 17(11):1875–1886, July 2015. ISSN 1520-9210. doi: 10.1109/tmm.2015.2477044. URL <http://dx.doi.org/10.1109/tmm.2015.2477044>.
- Chollet, F. Keras. <https://github.com/fchollet/keras>, 2015.
- Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. Empirical evaluation of gated recurrent neural networks on sequence modeling, December 2014. URL <http://arxiv.org/abs/1412.3555>.
- Collobert, R. and Weston, J. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning*, pp. 160–167. ACM, 2008.
- Cont, R. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.
- Cunningham, J. P., Ghahramani, Z., Rasmussen, C. E., Lawrence, N., and Girolami, M. Gaussian processes for time-marked time-series data. In *AISTATS*, volume 22, pp. 255–263, 2012.
- Dauphin, Y. N., Fan, A., Auli, M., and Grangier, D. Language modeling with gated convolutional networks, December 2016. URL <http://arxiv.org/abs/1612.08083.pdf>.
- Dheeru, D. and Karra Taniskidou, E. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Fama, E. F. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970.
- Gamboa, J. C. B. Deep learning for time-series analysis. *arXiv preprint arXiv:1701.01887*, 2017.
- Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In *In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS10)*. Society for Artificial Intelligence and Statistics, 2010. URL <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.207.2059>.
- Grave, E., Joulin, A., Cissé, M., Grangier, D., and Jégou, H. Efficient softmax approximation for GPUs, December 2016. URL <http://arxiv.org/abs/1609.04309.pdf>.
- Hamilton, J. D. *Time series analysis*, volume 2. Princeton university press Princeton, 1994.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition, December 2015. URL <http://arxiv.org/abs/1512.03385.pdf>.
- Heaton, J. B., Polson, N. G., and Witte, J. H. Deep learning in finance, February 2016. URL <http://arxiv.org/abs/1602.06561.pdf>.
- Hochreiter, S. and Schmidhuber, J. Long Short-Term memory. *Neural Computation*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL <http://dx.doi.org/10.1162/neco.1997.9.8.1735>.
- Hwang, Y., Tong, A., and Choi, J. Automatic construction of nonparametric relational regression models for multiple time series. In *Proceedings of the 33rd International Conference on Machine Learning*, 2016.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift, March 2015. URL <http://arxiv.org/abs/1502.03167v2.pdf>.

- Jozefowicz, R., Vinyals, O., Schuster, M., Shazeer, N., and Wu, Y. Exploring the limits of language modeling, February 2016. URL <http://arxiv.org/abs/1602.02410.pdf>.
- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. January 2015. URL <http://arxiv.org/abs/1412.6980>.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Kumar, A., Biswas, A., and Sanyal, S. ecommercegan: A generative adversarial network for e-commerce. *arXiv preprint arXiv:1801.03244*, 2018.
- Laloux, L., Cizeau, P., Potters, M., and Bouchaud, J.-P. Random matrix theory and financial correlations. *International Journal of Theoretical and Applied Finance*, 3(03):391–397, 2000.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998. ISSN 00189219. doi: 10.1109/5.726791. URL <http://dx.doi.org/10.1109/5.726791>.
- Li, S. C.-X. and Marlin, B. M. A scalable end-to-end gaussian process adapter for irregularly sampled time series classification. In *Advances in Neural Information Processing Systems*, pp. 1804–1812, 2016.
- Lichman, M. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Mathieu, M., Couprie, C., and LeCun, Y. Deep multi-scale video prediction beyond mean square error, February 2016. URL <http://arxiv.org/abs/1511.05440.pdf>.
- McNelis, P. D. *Neural networks in finance: gaining predictive edge in the market*. Academic Press, 2005.
- Mozer, M. C. Neural net architectures for temporal sequence processing. In *Santa Fe Institute Studies in the Sciences of Complexity*, volume 15, pp. 243–243, 1993.
- Neil, D., Pfeiffer, M., and Liu, S.-C. Phased lstm: Accelerating recurrent network training for long or event-based sequences. In *Advances In Neural Information Processing Systems*, pp. 3882–3890, 2016.
- Petelin, D., Šindelář, J., Přikryl, J., and Kocijan, J. Financial modeling using gaussian process models. In *Intelligent Data Acquisition and Advanced Computing Systems (IDAACS), 2011 IEEE 6th International Conference on*, volume 2, pp. 672–677. IEEE, 2011.
- Sak, H., Senior, A. W., and Beaufays, F. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In *Interspeech*, pp. 338–342, 2014.
- Schmidhuber, J. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- Sims, C. A. Money, income, and causality. *The American economic review*, 62(4):540–552, 1972.
- Sims, C. A. Macroeconomics and reality. *Econometrica: Journal of the Econometric Society*, pp. 1–48, 1980.
- Sirignano, J. Extended abstract: Neural networks for limit order books, February 2016. URL <http://arxiv.org/abs/1601.01987>.
- Srivastava, R. K., Greff, K., and Schmidhuber, J. Highway networks, November 2015. URL <http://arxiv.org/abs/1505.00387>.
- Tobar, F., Bui, T. D., and Turner, R. E. Learning stationary time series using gaussian processes with nonparametric kernels. In *Advances in Neural Information Processing Systems*, pp. 3501–3509, 2015.
- van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., and Kavukcuoglu, K. WaveNet: A generative model for raw audio, September 2016a. URL <http://arxiv.org/abs/1609.03499>.
- van den Oord, A., Kalchbrenner, N., Vinyals, O., Espeholt, L., Graves, A., and Kavukcuoglu, K. Conditional image generation with PixelCNN decoders, June 2016b. URL <http://arxiv.org/abs/1606.05328>.
- Weissenborn, D. and Rocktäschel, T. MuFuRU: The Multi-Function recurrent unit, June 2016. URL <http://arxiv.org/abs/1606.03002v1.pdf>.
- Wilson, A. and Ghahramani, Z. Copula processes. In *Advances in Neural Information Processing Systems*, pp. 2460–2468, 2010.
- Żołna, K. and Romański, B. User modeling using lstm networks. In *AAAI*, pp. 5025–5027, 2017.