

No. 680

Rafael Reisenhofer, Xandro Bayer,
and Nikolaus Hautsch

HARNet: A Convolutional Neural Network for Realized Volatility Forecasting

The CFS Working Paper Series

presents ongoing research on selected topics in the fields of money, banking and finance. The papers are circulated to encourage discussion and comment. Any opinions expressed in CFS Working Papers are those of the author(s) and not of the CFS.

The Center for Financial Studies, located in Goethe University Frankfurt's House of Finance, conducts independent and internationally oriented research in important areas of Finance. It serves as a forum for dialogue between academia, policy-making institutions and the financial industry. It offers a platform for top-level fundamental research as well as applied research relevant for the financial sector in Europe. CFS is funded by the non-profit-organization Gesellschaft für Kapitalmarktforschung e.V. (GfK). Established in 1967 and closely affiliated with the University of Frankfurt, it provides a strong link between the financial community and academia. GfK members comprise major players in Germany's financial industry. The funding institutions do not give prior review to CFS publications, nor do they necessarily share the views expressed therein.

HARNet: A Convolutional Neural Network for Realized Volatility Forecasting

Rafael Reisenhofer^{†,‡,*}, Xandro Bayer^{‡,§,*}, and Nikolaus Hautsch^{‡,§}

[†]*Faculty of Mathematics, University of Vienna, Oskar-Morgenstern-Platz 1, A-1090 Vienna, Austria*

[‡]*Research Network Data Science @ Uni Vienna, Kolingasse 14-16, A-1090 Vienna, Austria*

[§]*Department of Statistics and Operations Research, University of Vienna, Oskar-Morgenstern-Platz 1, A-1090 Vienna, Austria*

**These authors contributed equally*

Abstract

Despite the impressive success of deep neural networks in many application areas, neural network models have so far not been widely adopted in the context of volatility forecasting. In this work, we aim to bridge the conceptual gap between established time series approaches, such as the Heterogeneous Autoregressive (HAR) model (Corsi, 2009), and state-of-the-art deep neural network models. The newly introduced HARNet is based on a hierarchy of dilated convolutional layers, which facilitates an exponential growth of the receptive field of the model in the number of model parameters. HARNets allow for an explicit initialization scheme such that before optimization, a HARNet yields identical predictions as the respective baseline HAR model. Particularly when considering the QLIKE error as a loss function, we find that this approach significantly stabilizes the optimization of HARNets. We evaluate the performance of HARNets with respect to three different stock market indexes. Based on this evaluation, we formulate clear guidelines for the optimization of HARNets and show that HARNets can substantially improve upon the forecasting accuracy of their respective HAR baseline models. In a qualitative analysis of the filter weights learnt by a HARNet, we report clear patterns regarding the predictive power of past information. Among information from the previous week, yesterday and the day before, yesterday's volatility makes by far the most contribution to today's realized volatility forecast. Moreover, within the previous month, the importance of single weeks diminishes almost linearly when moving further into the past.

1 Introduction

Volatility is of critical importance in risk management, asset pricing and portfolio construction. Over the past years, extensive research efforts have been devoted to the development and evaluation of volatility forecasting models. A substantial body of research focuses on the construction of estimators for daily volatility exploiting high-frequency information. Seminal contributions (see, e.g., Andersen et al., 2001; Barndorff-Nielsen and Shephard, 2002) provided the foundation for using the realized variance, i.e., the sum of squared intraday returns, as a consistent estimator for the latent integrated variance of an underlying diffusion process. These approaches have been extended and robustified, and built the starting point for an active line of research on the development of high-frequency-based volatility estimators (see, e.g., Ait-Sahalia and Jacod, 2014). At the same time, these developments triggered the need to develop appropriate time series models to predict realized volatility. One of the most popular and empirically successful models is the heterogeneous autoregressive (HAR) model proposed by Corsi (2009), which captures the typically strong persistency of daily realized variances in a simple and efficient way.

In the past decade, deep neural networks (DNNs) have revolutionized many areas of machine learning such as image classification, machine translation or speech recognition. The most widely

¹rafael.reisenhofer@uni-bremen.de

used neural network (NN)-based approaches for forecasting sequential data are recurrent NNs, such as the so-called long short-term memory (LSTM) architecture (Hochreiter and Schmidhuber, 1997). Originally developed for image data sets, NNs with convolutional layers – so-called convolutional neural networks (CNNs) – have more recently been applied with impressive success for time series prediction in the realm of natural language processing (Oord et al., 2016; Kalchbrenner et al., 2016; Dauphin et al., 2017; Gehring et al., 2016; Gehring et al., 2017). In the context of financial time series, Miura, Pichl, and Kaizoji (2019) investigated numerous machine learning approaches, including CNNs, for predicting the realized volatility of Bitcoin returns, while Borovykh, Bohte, and Oosterlee (2017) employed a dilated convolutional network for forecasting stock index returns. A main advantage of dilated convolutional layers is that a linear increase in depth allows for an exponential increase of the considered time horizon. Such architectures should thus be particularly powerful for tasks, where large time horizons are potentially useful, but machine learning models often suffer from overfitting with a growing number of model parameters.

The applicability of feed-forward and recurrent NN models in the context of realized volatility forecasting has recently been explored by Bucci (2020), Christensen, Siggaard, Veliyev, et al. (2021) and Rahimikia and Poon (2020). The fact that so far, NN-based approaches have not been widely adopted in the area of financial econometrics could be explained by the fact that NN models are often used as off-the-shelf black boxes whose results are typically difficult to interpret, in particular in comparison to a simple and well-understood approach like the HAR model. In this work, we aim to bridge the conceptual gap between established methods such as the HAR model and state-of-the-art deep learning-based approaches by introducing so-called 'HARNets' as a novel NN model for forecasting realized volatility.

A HARNet is based on hierarchies of dilated convolutional layers. It is strongly inspired by the well-known HAR model in the sense that each layer of a HARNet computes features with respect to different time horizons, which are comparable to weekly and monthly aggregates as in the original HAR model. In particular, HARNets allow for an explicit initialization scheme for the respective model weights such that before optimization, a HARNet yields identical predictions as the respective baseline HAR model. This approach not only significantly stabilizes the optimization process but also facilitates an in-depth analysis of the performance of HARNets relative to different fits of their respective HAR baseline models.

We briefly summarize important concepts from realized volatility forecasting and the application of CNNs for time series forecasting in Sections 2 and 3, respectively.

The HARNet model as well as our approaches to initializing and optimizing the respective model parameters are presented in Section 4.

In Section 5, we conduct a detailed evaluation of the performance of HARNets relative to their respective baseline HAR models. In particular, we consider three different loss functions, the mean absolute error (MAE), the mean squared error (MSE) and the QLIKE error. Moreover, we evaluate three different estimation approaches for fitting the HAR baseline model, namely ordinary least squares (OLS), weighted least squares (WLS) and log-OLS, where the input time series is logarithmized before being processed by the model. Based on our initial findings, we further address questions regarding the stability of the optimization process and the generalization properties of optimized HARNets, and eventually formulate clear guidelines for the optimization of HARNets. In particular, we also conduct a qualitative analysis of the filter weights learnt by HARNets on the different convolutional layers, which is presented in Section 5.4.

In Section 6, we consider extended HARNets for higher-dimensional input time series, which also contain information regarding the realized semivariance as well as jump variation and again evaluate their performance relative to their extended HAR-like baseline models.

2 Forecasting realized volatility with the HAR model

We assume that the dynamics of the log price of a financial asset evolve according to the stochastic differential equation

$$dp(t) = \mu(t)dt + \sigma(t)dW(t), \quad (1)$$

where the càdlàg finite variation process $\mu(t)$ denotes the drift, $W(t)$ is a standard Brownian motion, and $\sigma(t)$ is the spot volatility, which is independent of $W(t)$. The daily integrated variance

$$IV_t = \int_{t-1}^t \sigma^2(\omega) d\omega \quad (2)$$

at a day t is unobservable in practice, even ex-post. However, it has been shown that the daily sum of N intraday squared returns

$$RV_t = \sum_{n=0}^{N-1} r_{t,n}^2, \quad (3)$$

where $r_{t,n} = p(t - n\Delta) - p(t - (n + 1)\Delta)$ with $\Delta = 1/N$, is a consistent estimator for IV_t when $N \rightarrow \infty$ (Andersen and Bollerslev, 1998).

The most popular model for forecasting RV_t is the so-called heterogeneous autoregressive model introduced by Corsi (2009). The model has a simple autoregressive structure and includes averages of daily realized variances over different time horizons. For an aggregation period $j \in \mathbb{N}$, we denote the respective average of daily realized variances by

$$RV_t^{(j)} = \frac{1}{j} \sum_{n=0}^{j-1} RV_{t-n}. \quad (4)$$

Corsi (2009) proposes using aggregation periods of 1, 5, and 22, corresponding to daily, weekly, and monthly RVs, respectively. The HAR model is then defined as

$$RV_t = \beta_0 + \beta_1 RV_{t-1} + \beta_2 RV_{t-1}^{(5)} + \beta_3 RV_{t-1}^{(22)} + \epsilon_t, \quad (5)$$

where the parameters β_0, \dots, β_3 , are usually estimated by ordinary least squares (OLS). However, since it is well-known that time series of realized variances exhibit properties such as conditional heteroscedasticity and non-Gaussianity, OLS is not ideal. To overcome this, we follow Patton and Sheppard (2015) and alternatively estimate the HAR model using weighted least squares (WLS) with weights chosen as the inverse of the (OLS) fitted value. Another popular approach to deal with heteroscedasticity and non-Gaussianity is to transform the data logarithmically. For the remainder of this paper, this method will be denoted as log-OLS.

3 Dilated convolutional NNs for time series forecasting

In this section, we briefly review the basic concepts and definitions underlying dilated convolutional neural networks (NNs) in the context of time series forecasting.

Like the majority of popular NN models, dilated convolutional NNs belong to the general class of feedforward NNs. Such networks transform a high-dimensional input vector, also called the *input layer*, by consecutively applying affine linear mappings and non-linear functions. An affine linear mapping followed by a non-linear function, which acts component-wise on the respective input, is called a *hidden layer*. The layer that yields the final output of the NN is typically called the *output layer*. The single entries of the vectors obtained at the different layers of the NN are called *input neurons*, *hidden neurons*, and *output neurons*, respectively. The number of consecutively applied layers defines the *depth* of the NN, while the number of neurons in a single layer defines the *width* of the NN. The non-linearity usually remains the same across all hidden layers of the network. The parameters that define the affine mappings can be different from layer to layer and are usually referred to as the *weights* of the NN. Note that the terms 'parameters' and 'weights' both refer to the values that define what a given NN model actually computes and are used interchangeably in the machine learning literature. The weights of a NN are *learnt* during *training* in the sense that they are iteratively updated with the goal of minimizing a predefined loss function (cf. Section 4.3). Before training, NN weights are usually initialized randomly. Sometimes it can be helpful to keep certain parameters of a NN model fixed throughout training. In this case, one differentiates between *trainable* and *non-trainable* weights. Informally, the extent to which a NN can compute different functions when varying the values of its trainable weights is often called

its *capacity*, or *expressivity*. Increasing the depth or the width of a NN model usually also increases its capacity.

Let d denote the dimension of the input vector x and L the number of layers in a feedforward NN where the l -th layer consists of N_l neurons. Each layer is then defined by an affine linear mapping

$$A_l : \mathbb{R}^{N_{l-1}} \rightarrow \mathbb{R}^{N_l} : x \mapsto W_l x + b_l, \quad 1 \leq l \leq L, \quad (6)$$

with $N_0 = d$, weight matrices $W_l \in \mathbb{R}^{N_{l-1} \times N_l}$, and bias vectors $b_l \in \mathbb{R}^{N_l}$. Let σ denote a real-valued non-linear function that acts component-wise on the input. The respective feedforward NN defines a mapping from $\mathbb{R}^d \rightarrow \mathbb{R}^{N_L}$, namely

$$\Phi(x) = A_L \sigma(A_{L-1} \sigma(\dots \sigma(A_1(x)))) , \quad x \in \mathbb{R}^d. \quad (7)$$

A convolutional neural network (CNN) is a special case of a feedforward NN where the affine mappings A_l are chosen to define convolution operators. The one-dimensional convolution of a sequence x with a filter h is defined as

$$(x * h)_t = \sum_{n=-\infty}^{\infty} x_{t-n} h_n, \quad x, h \in \mathbb{R}^{\mathbb{Z}}. \quad (8)$$

CNNs first rose to prominence when deep CNNs revolutionized the field of image classification by outperforming all previously existing approaches by significant margins (Krizhevsky, Sutskever, and Hinton, 2012). The basic idea behind deep CNNs is that sliding a filter over a given input is a simple means of detecting local features. In the context of image classification, these features are usually certain types of edges or corners in early layers but can become increasingly complex in deeper layers where filters are sensitive to certain geometrical shapes or even objects. A main advantage of CNNs is that only considering convolution operators imposes a strong structural constraint on the respective linear mapping A_l , which drastically reduces the number of model parameters when working with high-dimensional input data such as images.

Deep CNNs have also been widely applied for time series prediction (Oord et al., 2016; Borovykh, Bohte, and Oosterlee, 2017). When working with time series, it is usually required that a filter can only aggregate information from the past and present, but not from the future. This type of convolution is known as *causal* convolution and can be written as

$$(x *^c h)_t = \sum_{n=0}^{N-1} x_{t-n} h_n, \quad x \in \mathbb{R}^{\mathbb{Z}}, h \in \mathbb{R}^N, \quad (9)$$

where h denotes a finite filter of length N .

The *receptive field* of a model describes the part of an input time series which is actually used by the model to compute a prediction for future values of the time series. That is, the receptive field determines how far a model can see into the past. When only considering convolutions of the form (9), the length of the receptive field can only grow linearly with the number of filter coefficients, that is, the number of trainable parameters in the model. In situations with long-range dependencies, this can lead to inefficient training and overfitting due to the high number of trainable parameters. A possible remedy is to use compositions of *dilated* convolutions. In a dilated convolution, the inner product with the convolution filter is not based on the consecutive entries of the time series but on entries that are a fixed number of steps apart from each other. Formally, a dilated causal convolution with dilation factor k is defined as

$$(x *^c_k h)_t = \sum_{n=0}^{N-1} x_{t-kn} h_n, \quad x \in \mathbb{R}^{\mathbb{Z}}, h \in \mathbb{R}^N, \quad (10)$$

where h denotes a finite filter of length N . By composing dilated convolutional layers that have exponentially increasing dilation factors, it is possible to cover receptive fields that grow exponentially with the number of trainable model parameters. All three types of convolutions presented in this section are illustrated in Figure 1. Examples for causal convolutions of an RV time series with different average filters are shown in Figure 2.

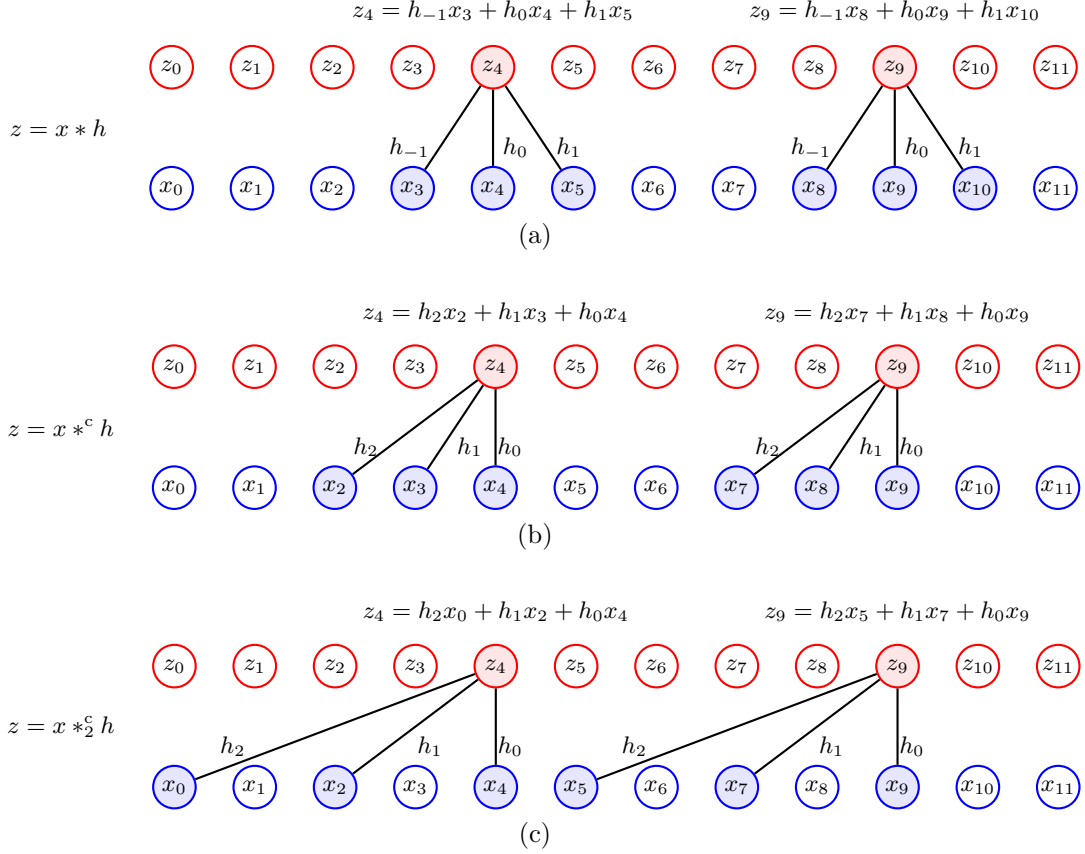


Figure 1: Three different types of one-dimensional convolutions. (a) One-dimensional convolution of a time series x with a filter h . (b) Causal convolution of a time series x with a filter h . (c) Dilated causal convolution of a time series x with a filter h .

3.1 Optimization of NNs

Let $\Phi_\theta: \mathbb{R}^d \rightarrow \mathbb{R}^{N_L}$ be the function defined by a NN model as in equation (7), where θ denotes the set of all trainable weights. That is, θ contains all parameters of the matrices W_l and bias vectors b_l that are learnt during training. Here, we consider the case of *supervised learning*. This means that the model is optimized such that its output best matches input-output pairs from a carefully annotated and selected *training data set*. The outputs within the training data set are usually referred to as *labels*. When training a NN architecture to yield one-day-ahead forecasts for RV, the input x in a single input-label pair from the training set typically consists of the complete training RV time series up to a day t , while the corresponding label y is chosen as the RV on day $t + 1$, that is, $x = (\dots, RV_{t-1}, RV_t)^T$ and $y = RV_{t+1}$.

For a sequence $Z = (x^{(n)}, y^{(n)})_{n=1}^N$ of N input-label pairs, we define the *empirical risk* via

$$J(\theta, Z) = \frac{1}{N} \sum_{n=1}^N \ell(\Phi_\theta(x^{(n)}), y^{(n)}), \quad (11)$$

where ℓ is a real-valued loss function. In the context of this work, the labels $y^{(n)}$ are always the consecutive values of the RV time series $x^{(n)}$, which are both taken from a predefined training time series. $\Phi_\theta(x^{(n)})$ represents the forecast of the model with parameters θ .

The sequence Z , which can contain the complete training data set, or only selected samples from different parts of the full training time series, is usually referred to as a *batch*. The length N of the sequence Z is called the *batch size*.

The weights θ are updated iteratively via gradient descent. In its most basic form, the gradient

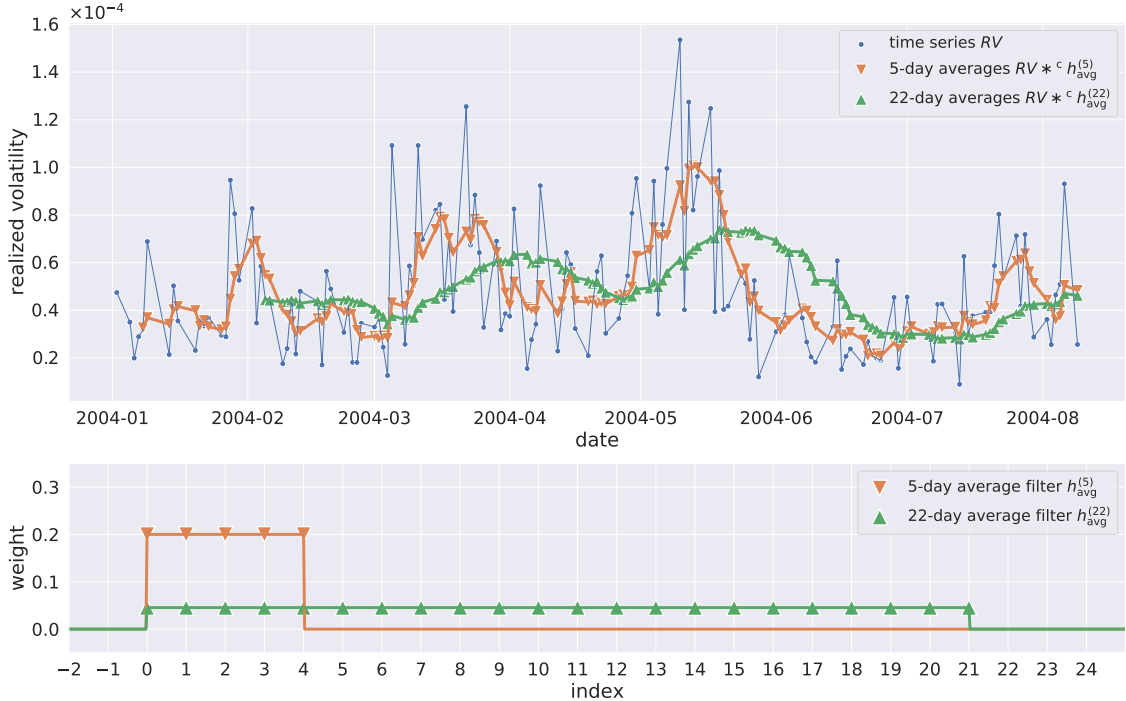


Figure 2: Causal convolution of a realized volatility time series with 5-day and 22-day average filters. The filters $h_{\text{avg}}^{(5)}$ and $h_{\text{avg}}^{(22)}$ are defined in equation (14).

descent update rule can be written as

$$\theta^{(i+1)} = \theta^{(i)} - \alpha \nabla_{\theta} J(\theta^{(i)}, Z^{(i)}), \quad (12)$$

where $\alpha > 0$ denotes the learning rate, $Z^{(i)}$ is the training data batch used for the i -th update of the parameters θ , and the gradient vector is defined as

$$\nabla_{\theta} J(\theta^{(i)}, Z^{(i)}) = \begin{pmatrix} \frac{\partial J(\theta^{(i)}, Z^{(i)})}{\partial \theta_1} \\ \frac{\partial J(\theta^{(i)}, Z^{(i)})}{\partial \theta_2} \\ \vdots \end{pmatrix}. \quad (13)$$

Two of the most popular first-order gradient descent algorithms are ADAM (Kingma and Ba, 2017) and RMSProp (Hinton, Srivastava, and Swersky, 2012), which both implement adaptive schemes for the learning rate α . The gradient vector $\nabla_{\theta} J(\theta^{(i)}, Z^{(i)})$ is usually obtained via back-propagation (Rumelhart, Hinton, and Williams, 1986).

In practice, each batch $Z^{(i)}$ usually contains only a small subset of the complete training data, which is selected randomly for each gradient descent update of the form (12). This approach is known as mini-batch stochastic gradient descent (SGD). When only using a single input-label pair in each gradient descent step, the learning algorithm often has difficulties identifying patterns in the data, in particular when the training data is very noisy. Using the entire training data as a single batch, on the other hand, often leads to overfitting. Mini-batch SGD can be seen as a compromise between these two extremes which has been found to be highly applicable in practice.

Parameters such as the (initial) learning rate α , the batch size, the stopping criteria for the learning algorithm, but also the width and depth and the NN are called *hyperparameters*. Hyperparameters define important characteristics of the architecture and the applied learning algorithm which are not changed during training.

For a more comprehensive overview on deep learning, see for example Goodfellow, Bengio, and Courville (2016).

4 Method

We aim to exploit the close relationship between the original HAR model (cf. equation (5)) and NN models that are based on layers of dilated causal convolutions of the form (10). In particular, we will consider a NN network model whose weights can be initialized such that, before optimization, it exactly replicates the predictions of a fitted HAR model. This initialization approach significantly stabilizes the optimization process and allows us to perform a rigorous evaluation of the nonlinear NN model relative to the HAR model with respect to different types of data as well as different optimization settings.

We first note that the HAR model can easily be written in terms of convolutions with average filters. Let

$$h_{\text{avg}}^{(j)} = \left(\frac{1}{j}, \dots, \frac{1}{j} \right) \in \mathbb{R}^j \quad (14)$$

denote an average filter of length $j \in \mathbb{N}$. Then, the averages of daily realized volatility are defined by $RV_t^{(j)} = (RV *^c h_{\text{avg}}^{(j)})_t$, and thus, the HAR model (5) can be written as

$$RV_t = \left(\beta_0 + \beta_1 RV *^c h_{\text{avg}}^{(1)} + \beta_2 RV *^c h_{\text{avg}}^{(5)} + \beta_3 RV *^c h_{\text{avg}}^{(22)} \right)_{t-1} + \epsilon. \quad (15)$$

For the model defined in equation (15), the weights of the filters $h_{\text{avg}}^{(1)}$, $h_{\text{avg}}^{(5)}$, and $h_{\text{avg}}^{(22)}$ are fixed to obtain the respective 1-, 5-, and 22-day averages (cf. Figure 2). A straightforward but impractical approach to define a CNN with learnable filters based on this model would be to consider a network with a single convolutional layer, which consists of three different filters of length 1, 5, and 22, respectively, with trainable weights, and a linear output layer, which aggregates all filter outputs with trainable weights β_0, \dots, β_3 . With the additional non-linearity, which is applied to all filter outputs, and a total number of 32 trainable parameters (β_0, \dots, β_3 and the weights of the three filters in the convolutional layer), this model would be significantly more expressive than the HAR model, which has only four parameters. Moreover, such a shallow model with only one hidden layer does not exploit the hierarchical structure of deep CNNs while having the same receptive field as the original HAR model with eight times as many model parameters. As a consequence, this model would be extremely prone to overfitting.

An approach that keeps the number of model parameters low while also increasing the expressiveness of the HAR model such that meaningful filters for different time horizons can be learnt is to consider dilated convolutions and multiplicatively nested time horizons. The latter means that each aggregation period in a sequence of time horizons is an integer multiple of its predecessor (cf. equation 16). In particular, this approach yields receptive fields that can grow exponentially in the number of model parameters.

Let (j_1, \dots, j_L) be an increasing sequence of L aggregation periods, where each period is an integer multiple of its predecessor, that is,

$$\exists c_l \in \mathbb{N}: c_l j_l = j_{l+1} \quad \forall l \in \{1, \dots, L-1\}. \quad (16)$$

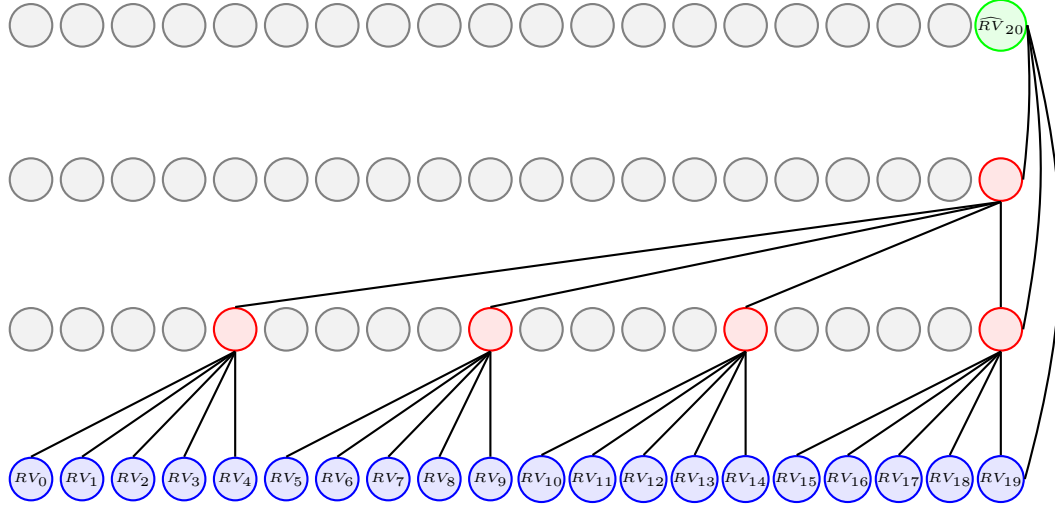
Then, the respective averages of daily RVs can be defined recursively via

$$RV^{(j_l)} = \begin{cases} RV *^c h_{\text{avg}}^{(j_1)} & \text{if } l = 1, \\ RV^{(j_{l-1})} *_{j_{l-1}}^c h_{\text{avg}}^{(j_l/j_{l-1})} & \text{else.} \end{cases} \quad (17)$$

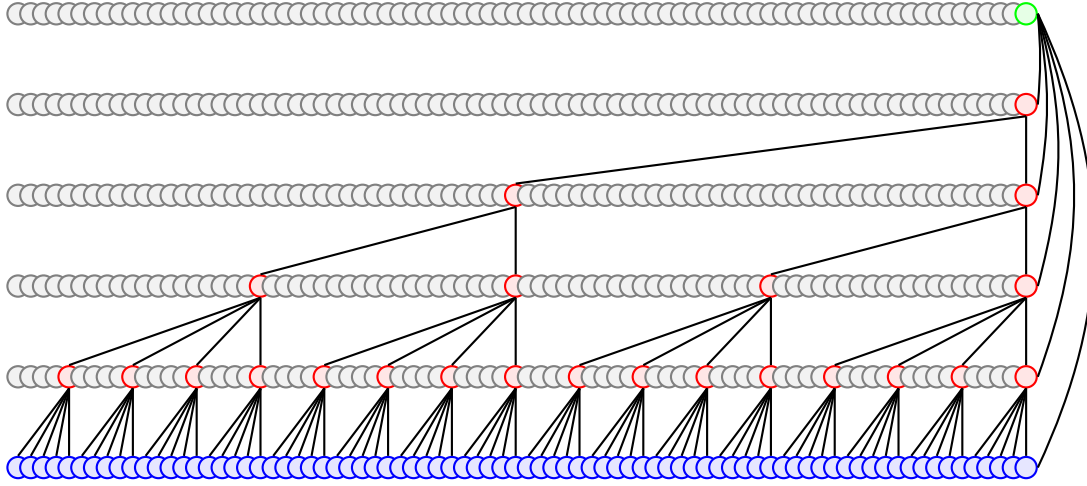
For the multiplicatively nested sequence of time horizons (1, 5, 20), which closely resembles the sequence of time horizons considered in the original HAR model, this approach would yield a hierarchical model of dilated convolutions with three filters of lengths 1, 5, and 4, respectively, where the second filter computes weekly weighted averages and the third filter monthly weighted averages. The respective model, however, would only contain 14 trainable parameters.

4.1 HARNet model

The HARNet is based on a hierarchy of dilated causal convolutions and utilizes the recursive relationship defined in equation (17). For a multiplicatively nested sequence (cf. (16)) of L aggregation



(a) Network topology of $\Phi_\theta^{\text{HAR}^{20}}$ with aggregation periods (1, 5, 20).



(b) Network topology of $\Phi_\theta^{\text{HAR}^{80}}$ with aggregation periods (1, 5, 20, 40, 80).

Figure 3: Topologies of the NN models used in our experiments. The blue nodes in the bottom row denote the input time series while the green nodes in the top row denote the respective one-day-ahead predictions.

periods (j_1, \dots, j_L) , the HARNet consists of L convolutional layers, where the l -th layer is defined by a finite filter $h^{(l)}$ and computes

$$f^{(l)}(x) = \begin{cases} \sigma(x *^c h^{(l)}), & h^{(l)} \in \mathbb{R}^{j_1} & \text{if } l = 1, \\ \sigma(f^{(l-1)}(x) *_{j_{l-1}}^c h^{(l)}), & h^{(l)} \in \mathbb{R}^{j_l/j_{l-1}} & \text{else,} \end{cases} \quad (18)$$

where $x \in \mathbb{R}^Z$ is an input time series and σ denotes the non-linear activation function. Here, we always use the rectified linear unit (ReLU) as an activation function, which is defined as

$$\sigma_{\text{relu}}(x) = \begin{cases} x & \text{if } x > 0, \\ 0 & \text{else.} \end{cases} \quad (19)$$

The complete function implemented by a HARNet is then given by

$$\Phi_\theta^{\text{HAR}}(x) = \beta_0 + \sum_{l=1}^L \beta_l f^{(l)}(x), \quad (20)$$

where the set θ of all trainable parameters consists of the weights β_0, \dots, β_L as well as the filters $h^{(1)}, \dots, h^{(L)}$. Note, however, that in the case of $j_1 = 1$, which is true for all models considered in this work, we will omit the first convolutional layer and simply compute the identity instead. Then, the total number of model parameters can be computed as

$$|\theta| = 1 + L + \sum_{l=2}^L \frac{j_l}{j_{l-1}}. \quad (21)$$

We will write $\Phi_{\theta}^{\text{HAR20}}$ and $\Phi_{\theta}^{\text{HAR80}}$ to denote the HARNet models that are based on the aggregation periods (1, 5, 20) and (1, 5, 20, 40, 80), respectively. The corresponding network topologies are shown in Figure 3. The models $\Phi_{\theta}^{\text{HAR20}}$ and $\Phi_{\theta}^{\text{HAR80}}$ thus contain 13 and 19 trainable parameters, respectively.

4.2 Initialization of model parameters

For a HARNet defined by a sequence of L aggregation periods (j_1, \dots, j_L) , we initialize the parameters β_0, \dots, β_L with the fitted parameters of the HAR model (5), which corresponds to the same sequence of aggregation periods. We consider three different techniques for estimating the baseline HAR model, namely OLS, WLS, and log-OLS (cf. Section 2).

The convolutional filter in the l -th layer is initialized with the average filter of the corresponding length, that is,

$$h^{(l)} = \begin{cases} h_{\text{avg}}^{(j_1)} & \text{if } l = 1, \\ h_{\text{avg}}^{(j_l/j_{l-1})} & \text{else.} \end{cases} \quad (22)$$

With this initialization and the ReLU as an activation function, the function $\Phi_{\theta}^{\text{HAR}}$ exactly replicates the predictions of the respective fitted HAR model, given that the input time series does not contain any negative values.

4.3 Loss functions and optimization

Fitting the HAR model by OLS or WLS can be done via a closed formula and numerically only requires a single matrix inversion. However, even for linear problems such as the HAR model, this is only possible when considering the (weighted) mean squared error as the objective function. While gradient descent-based optimization is usually much more tedious and in most relevant situations not guaranteed to converge to a global minimum, it comes with the great advantage that it can be performed with arbitrary loss functions.

Here, we consider three different loss functions during optimization and for evaluating the forecast accuracy of trained models. These are the mean absolute error (MAE), the mean squared error (MSE) and the QLIKE error as proposed by Patton (2011). For an observation y and a

Table 1: Optimization hyperparameters. Choosing a batch size of 4, and 5 consecutive labels per training sample implies that a single batch contains labels for 20 different one-day-ahead forecasts that are distributed across 4 randomly sampled segments of the complete training time series. For a training time series that spans four years – which roughly corresponds to 1,000 labels – these settings correspond to an approximate number of 200 epochs.

Optimizer	ADAM
Learning rate	1×10^{-4}
Batch size	4
Consecutive labels per training sample	5
Iterations	10,000

prediction \hat{y} produced by the model, these functions are defined by

$$\ell_{\text{MAE}}(\hat{y}, y) = |\hat{y} - y|, \quad (23)$$

$$\ell_{\text{MSE}}(\hat{y}, y) = (\hat{y} - y)^2, \quad (24)$$

$$\ell_{\text{Q}}(\hat{y}, y) = \frac{y}{\hat{y}} - \log\left(\frac{y}{\hat{y}}\right) - 1. \quad (25)$$

According to Patton (2011), ℓ_{MAE} and ℓ_{Q} belong to a parametric family of robust and homogeneous loss functions for volatility forecasting, i.e., they are robust to noise in the volatility proxy and invariant to the choice of units of measurement.

For a given time series of daily RVs, HARNets are trained via mini-batch stochastic gradient descent. For each iteration of the gradient descent algorithm, a batch is constructed by randomly sampling a fixed number of training samples that are taken from different regions of the training time series. Each of these training samples is a segment of consecutive observations from the training time series whose length exceeds the length of the receptive field of the model by a fixed number of entries. Each of these additional entries defines an input-observation pair in the sense that the sample contains sufficient data for the model to compute the respective one-day-ahead prediction. This format is chosen because predictions for consecutive days can be computed efficiently during training due to the convolutional structure of the architecture. On the other hand, by considering batch sizes that are greater than one, it is also possible to combine samples from completely different parts of the training time series in a single batch.

In our numerical experiments, we apply 10,000 iterations of the ADAM Kingma and Ba (2017) optimizer with a fixed learning rate of 10^{-4} . We use a batch size of 4 and consider 5 consecutive labels per training sample. That is, each input sample consists of a time series which exceeds the length of the receptive field of the model by 5 entries, and a single batch thus contains a total of 20 labels that come from 4 different regions of the complete training time series.

In the subsequently presented results, we usually consider training time series that span four consecutive years, which corresponds to roughly 1,000 one-day-ahead predictions that can be used as labels. This means that with 10,000 iterations of stochastic gradient descent and 20 labels per batch, our optimization algorithm runs for roughly 200 epochs. That is, each label in the complete training time series is used approximately 200 times. All hyperparameters are summarized in Table 1.

5 Results

To evaluate the forecast accuracy of the proposed models and optimization techniques, we consider time series of daily RV for three different indexes, namely the S&P500, the FTSE, and the DJI. The data is obtained from Oxford-Man Institutes realised library (Heber et al., 2009), from which we use the 5-min sub-sampled realized variance measure. Each of the three daily RV time series spans more than 19 years, ranging from 2001-08-29 to 2020-12-31. We split each time series into 15 pairs of training and test sets, where each training set spans four consecutive years and the succeeding year is used as a test set. Throughout this section, we will compare different methods with respect to their performance across all 15 pairs of training and test sets for each of the three indexes.

We evaluate the HARNet models $\Phi_{\theta}^{\text{HAR20}}$ and $\Phi_{\theta}^{\text{HAR80}}$ by considering three different choices for the HAR baseline fit that is used for initializing the respective model parameters and three different loss functions. The considered loss functions are ℓ_{MAE} , ℓ_{MSE} , and ℓ_{Q} (cf. Section 4.3). The three choices for the HAR baseline fit are OLS, WLS, and log-OLS (cf. Section 2). To ensure that predictions are strictly larger than zero, we clip the predicted value for both models at half of the smallest realized variance in the training set. All input RV time series are scaled linearly before being passed to the model such that 0 is mapped to 0 and 10^{-3} is mapped to 1. When working with logarithmized time series, -13 is mapped to 0 and -2.5 is mapped to 1. Scaling input features to roughly cover the interval $[0, 1]$ is common practice in many machine learning applications and often significantly improves gradient descent optimization.

Note that when using the log-OLS fit of the HAR model as a baseline, the NN model is also optimized for the logarithmized time series. In this case, the model does not predict the realized

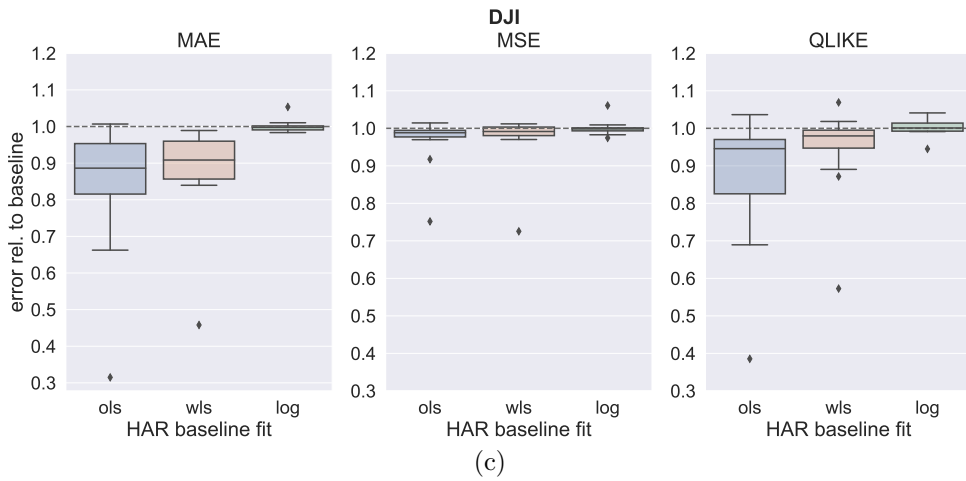
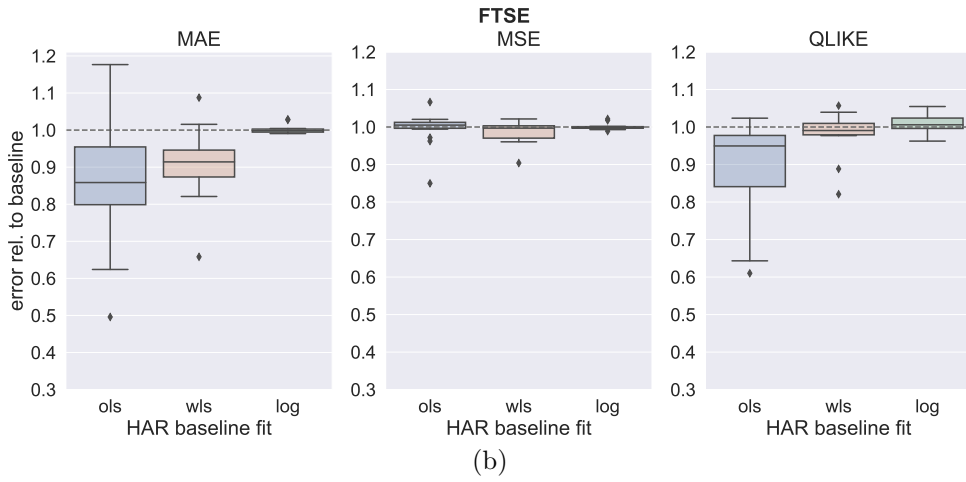
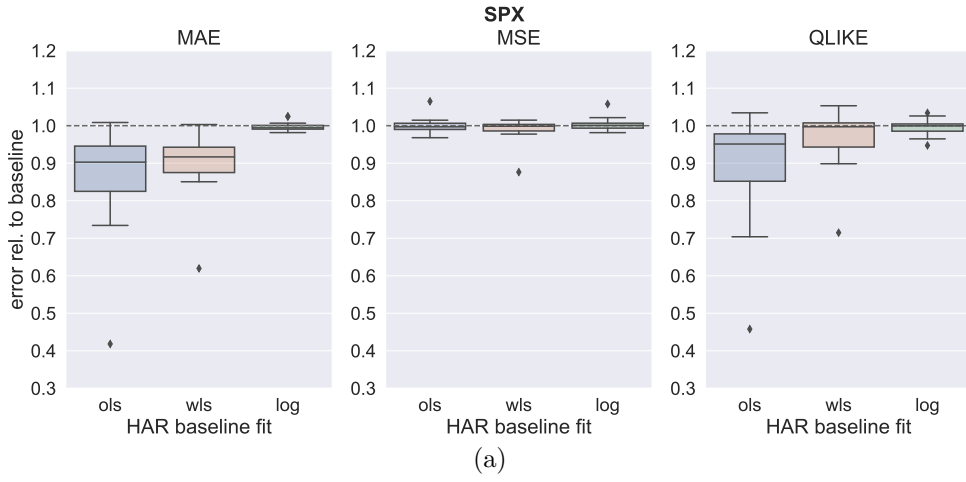
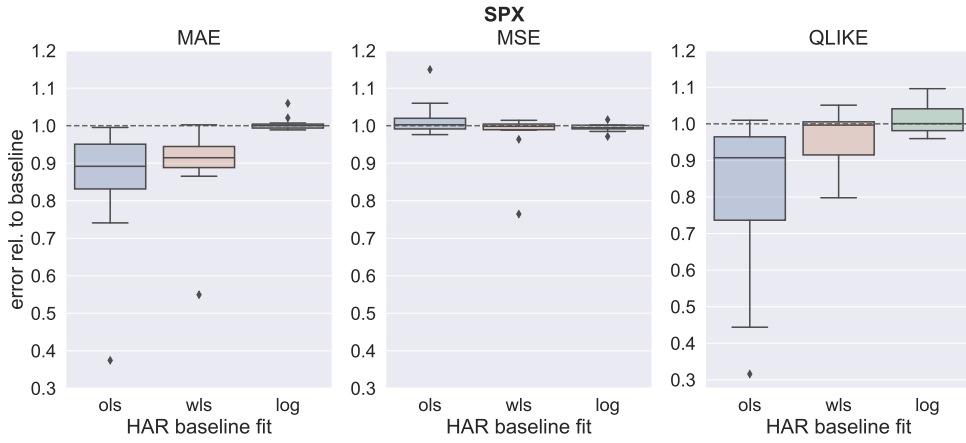
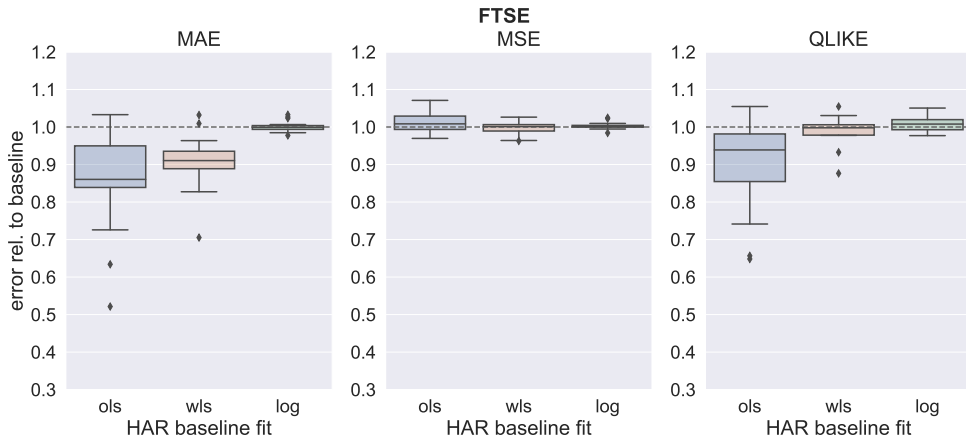


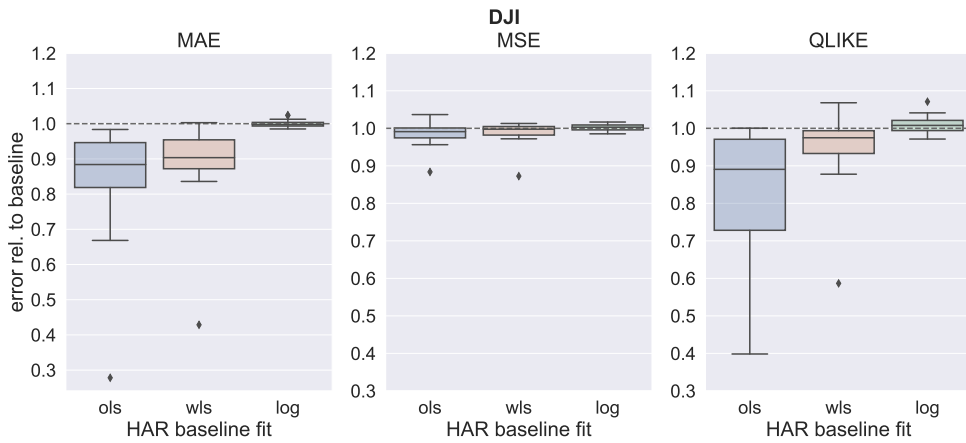
Figure 4: Performance of the $\Phi_\theta^{\text{HAR}20}$ model with aggregation periods (1, 5, 20) relative to the respective HAR baseline fit. Each boxplot depicts the distribution of the relative test error for 15 training set/test set pairs taken from a time series of the respective index that ranges from 2002 to 2020.



(a)



(b)



(c)

Figure 5: Performance of the $\Phi_{\theta}^{\text{HAR}80}$ model with aggregation periods (1, 5, 20, 40, 80) relative to the respective HAR baseline fit. Each boxplot depicts the distribution of the relative test error for 15 training set/test set pairs taken from a time series of the respective index that ranges from 2002 to 2020.

variance, but the log-realized variance, which needs to be reconverted to the original scale of measurement. Here, we apply the bias correction proposed by Proietti and Lütkepohl (2013).

Figure 4 depicts test errors relative to the test error of the respective HAR baseline fit for the $\Phi_{\theta}^{\text{HAR20}}$ model after optimization with three different loss functions. In each case, the error on the test set is computed with the same loss function, which was already used during optimization. Note that before optimization, $\Phi_{\theta}^{\text{HAR20}}$ and the HAR baseline model yield identical predictions. Our results therefore provide a detailed quantitative analysis as to how the performance of the model changes from the HAR baseline during optimization. Analogous results for the $\Phi_{\theta}^{\text{HAR80}}$ model are shown in Figure 5.

For both models $\Phi_{\theta}^{\text{HAR20}}$ and $\Phi_{\theta}^{\text{HAR80}}$, we observe that the largest relative improvement can be achieved when using ℓ_{MAE} as a loss function and the standard OLS approach for fitting the baseline HAR model. In the case of $\Phi_{\theta}^{\text{HAR20}}$, across all indexes, we achieve an average reduction of the median test MAE of about 11.74%. Substantial improvements in terms of the MAE can also be observed when choosing the WLS fit as a starting point. However, when considering logarithmized RV time series, the relative improvement in accuracy is greatly reduced. Using ℓ_{Q} as a loss functions yields slightly reduced but still considerable improvements relative to the HAR baseline fit, in particular when considering the standard OLS approach. In the case of ℓ_{MSE} , however, the test error of the NN model relative to the baseline HAR fit remains almost the same after optimization.

Based on the results presented in Figures 4 and 5, we believe that the following observations are of significant importance and merit further analysis:

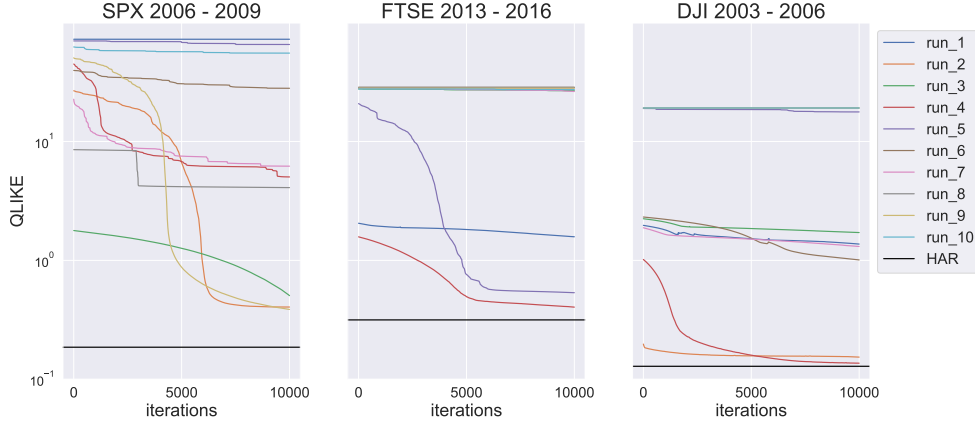
1. While Rahimikia and Poon (2020) do not recommend using the QLIKE error as an objective function during optimization, we found that this approach can yield consistent and substantial improvements. Can these diverging experimental results be explained by our approach of initializing the parameters of the HARNet model?
2. The QLIKE error has been described as a superior metric for ranking different volatility forecasting models (Patton and Sheppard, 2009). Does applying ℓ_{Q} as a loss function also yield improvements with respect to other metrics such as the MAE or the MSE and can thus be generally recommended in this setting?
3. Our method yields substantial improvements relative to the baseline HAR model when considering the standard OLS fit but often little to no improvements when using WLS or log-OLS. This raises the question whether HARNets that used an OLS fit as a starting point also outperform HAR models that were fitted with an WLS or log-OLS approach.
4. On average, the optimized $\Phi_{\theta}^{\text{HAR20}}$ model matches the performance on the test data of the respective baseline HAR fit and often clearly outperforms it. That is, it often generalizes better despite having significantly more model parameters. In particular, $\Phi_{\theta}^{\text{HAR20}}$ learns weighted filters for the weekly and monthly aggregation periods, whereas the original HAR model only considers fixed average filters. Can an analysis of the learnt filter weights provide further insights as to which aggregates of past data are of significant importance when forecasting RV?

In the following subsections, we will perform a detailed analysis of the questions raised above.

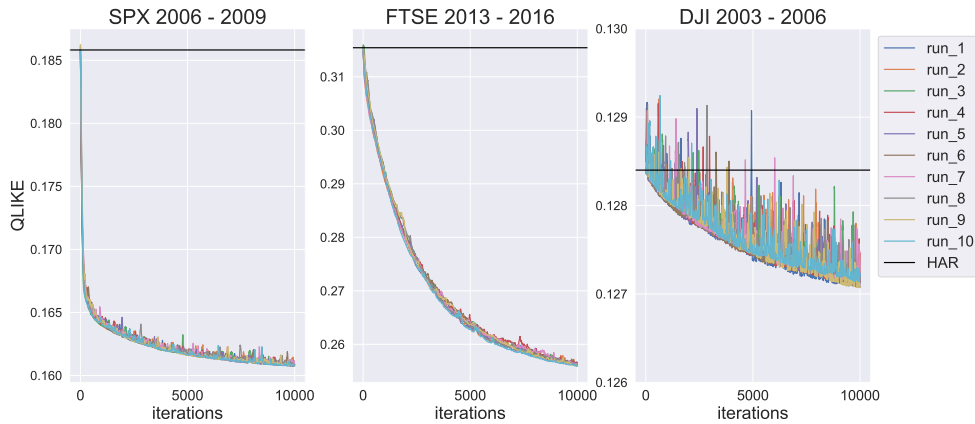
5.1 Initialization with a HAR baseline model stabilizes optimization with the QLIKE loss

Rahimikia and Poon (2020) report an inferior performance of machine learning models for RV forecasting that were optimized with a QLIKE loss as compared to an MSE loss. They argue that this could be due to the insensitivity of the QLIKE error in cases where the true RV is very high. In practice, such a flat loss function would make it very difficult for the optimizer to overcome large training errors on volatility jump days during optimization.

We suspect that this issue can at least partially be overcome in our setting due to the proposed initialization approach. After initialization, the HARNet yields identical predictions as the baseline HAR model. Thus, it can produce extreme prediction errors only to the same extent as produced



(a) Random initialization of weights.



(b) Proposed initialization with the HAR baseline model.

QLIKE	SPX 2006 - 2009			FTSE 2013 - 2016			DJI 2003 - 2006		
	median	mean	std	median	mean	std	median	mean	std
Random init	5.62038	23.84032	28.00004	27.08056	19.66301	12.34607	1.54270	8.09141	8.76879
HAR baseline init	0.16084	0.16083	0.00005	0.25620	0.25621	0.00021	0.12720	0.12722	0.00013

(c) Summary of QLIKE training errors after 10,000 iterations.

Figure 6: Initialization with a baseline HAR model stabilizes the optimization of a HARNet when using the QLIKE error as a loss function. (a) Optimization curves for ten different optimization runs of the $\Phi_{\theta}^{\text{HAR}20}$ model with randomly initialized weights. (b) Optimization curves for ten different optimization runs of the $\Phi_{\theta}^{\text{HAR}20}$ model with the proposed initialization with a HAR baseline model. (c) QLIKE training errors after 10,000 iterations of the ADAM algorithm.

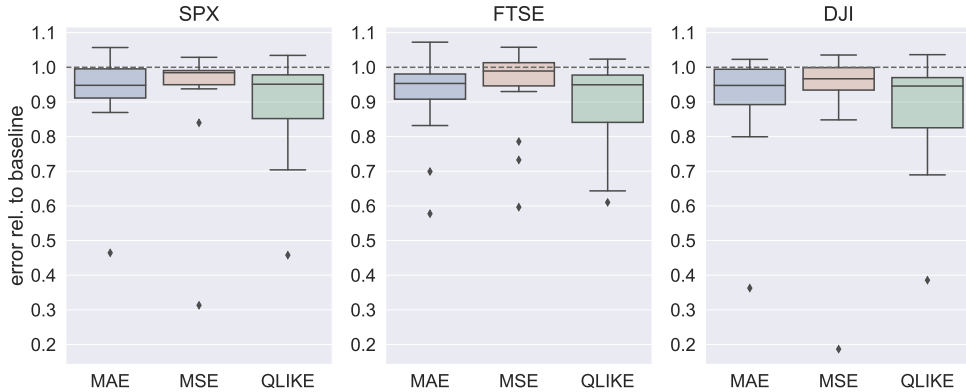
by the baseline model. However, when considering a random initialization of the model weights, this is not the case and extreme prediction errors are very common during early optimization.

To test this hypothesis experimentally, we compare optimization curves for the model $\Phi_{\theta}^{\text{HAR}20}$ when using random initialization techniques and when using an OLS fit of the HAR model as a baseline. For random initialization, we consider the Glorot uniform initializer (Glorot and Bengio, 2010), which is one of the most common initialization methods for NN weights. The Glorot uniform initializer draws samples from a uniform distribution over an interval whose limits depend on the number of input and output units in the weight tensor of a layer.

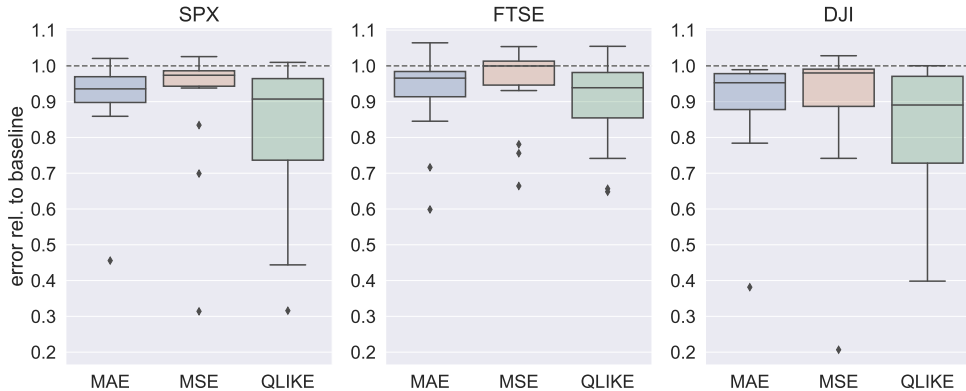
Figure 6 contains the results of an experiment where we first randomly selected a time span of four years as a training set for each of the three indexes, and then performed 10 independent optimizations with both initialization methods, respectively. Our results clearly show that the QLIKE training error highly depends on the initialization when considering random initialization. As shown in Figure 6a, random initialization causes a high variation in the final training error

after 10,000 iterations of the ADAM algorithm, while all of the 10 optimization runs yield training errors that are greater than the training error of the fitted HAR model with the same aggregation periods. When considering initialization with a baseline HAR model, on the other hand, all of the 10 optimization curves follow basically the same path and improve upon the QLIKE training error of the baseline model (cf. Figure 6b).

5.2 QLIKE optimization also generalizes with respect to the MAE and MSE test metrics



(a) HARNet with aggregation periods (1, 5, 20).



(b) HARNet with aggregation periods (1, 5, 20, 40, 80).

Figure 7: Performance of the $\Phi_{\theta}^{\text{HAR}20}$ and $\Phi_{\theta}^{\text{HAR}80}$ models after optimization with the QLIKE loss function in terms of the MAE, MSE, and the QLIKE test error relative to the HAR baseline fit. The baseline fit was obtained by OLS. Each boxplot depicts the distribution of the relative test error for 15 training set/test set pairs taken from a time series of the respective index that ranges from 2002 to 2020.

The QLIKE error is often considered to be the preferred evaluation metric in the context of RV forecasting. In our initial results, we found that when using the QLIKE error as a loss function, a HARNet can substantially improve upon the QLIKE test error of a HAR baseline model which was fitted by OLS (cf. Figures 4 and 5). This raises the question whether optimizing with a QLIKE loss function can also improve other test metrics such as the MAE or the MSE.

Our results compiled in Figure 7 indeed show that HARNets that were optimized with a QLIKE loss not only yield better test errors with respect to the QLIKE metric than the baseline HAR model but also with respect to the MAE and the MSE. While the improvements in terms of the MAE are smaller than when directly optimizing with an MAE loss, we found that this approach leads to consistent albeit small improvements in terms of the MSE test error. This is particularly surprising as in our initial results, we found that using the MSE as a loss function usually fails at

improving the MSE test error of the baseline model.

5.3 HARNets that are initialized with an OLS fit of the HAR baseline model outperform WLS and log-OLS fits of the baseline model

Our main results presented in Figures 4 and 5 show that HARNets can substantially improve upon the forecasting accuracy when using an OLS fit of the HAR baseline model. However, this is in general not true when fitting the baseline model by WLS, while using the log-OLS fit of the baseline as a starting point for the HARNet usually yields no significant changes in the accuracy on the test set whatsoever.

The results presented in Figure 8 show that HARNets that were initialized with the OLS fit of the baseline model also consistently outperform the respective WLS and log-OLS fits of the baseline model in terms of the QLIKE error. We even found that HARNets that were initialized with the OLS fit of the baseline model consistently match the performance of the HAR baseline model which was obtained by the *best* method for fitting its parameters, i.e., for which the decision whether to use OLS, WLS, or log-OLS was made after knowing the respective test errors.

In the log-OLS approach, the training time series is logarithmized before solving the linear system. This transformation introduces non-linearity to the forecasting model and implicitly reduces the relative importance of days with a very high RV in the OLS fit. When using WLS, on the other hand, one explicitly defines weights, which usually also reduce the importance of days of RV jumps. Both approaches are explicitly defined schemes that aim to overcome the limitations of simple linear models, such as the HAR model, without adding any complexity to the model itself. We believe that the results presented in this section suggest that, due to the significantly increased expressiveness, the HARNet model is capable of learning similarly effective relationships directly from the data. Explicitly incorporating the WLS or log-OLS fit of the HAR baseline model, however, seems to either cause the HARNet model to get stuck in unfavorable local minima or increase the risk of overfitting.

In summary, our findings suggest that when used for forecasting of daily RV, HARNets are best optimized with a QLIKE loss function and an initialization scheme which uses the OLS fit of the HAR baseline model as a starting point.

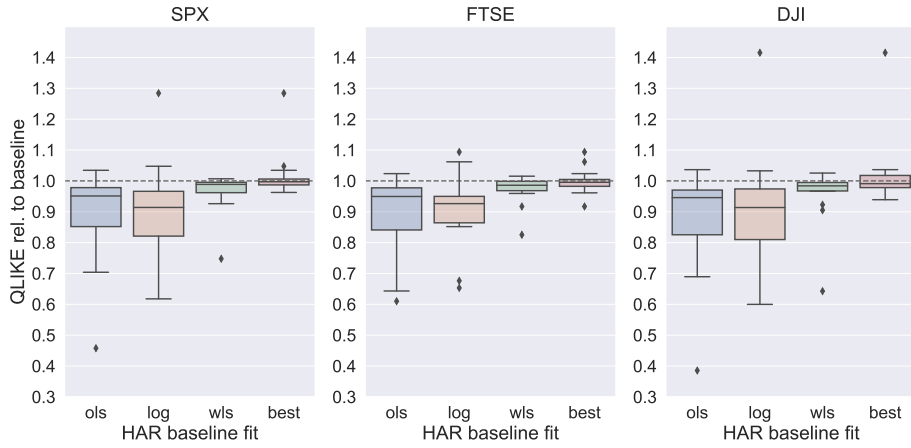
5.4 Interpretability of HARNet filter weights

Due to the close relationship to the original HAR model, HARNets allow for a high degree of interpretability when forecasting RV, setting them apart from generic machine learning models which are often treated as black boxes. A particularly interesting property of a HARNet is that the time series corresponding to a certain aggregation period is computed via weighted averages rather than normal averages, as it is the case in the original HAR model. The filter kernels which define the respective weights are learnt during optimization and can thus inform us as to which of the previous days or weeks are in fact significant for RV forecasting within a given training time series.

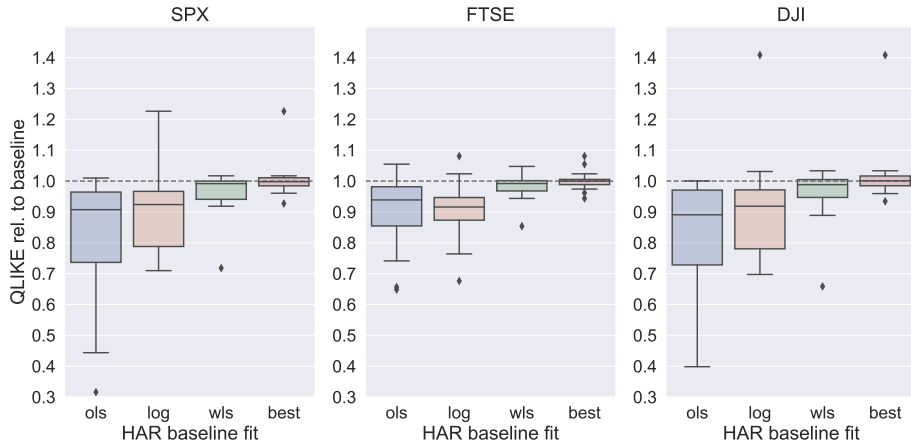
Figure 9 depicts the distribution of the weights in the two convolutional layers of a $\Phi_{\theta}^{\text{HAR20}}$ model which was initialized with the OLS fit of the HAR baseline model and optimized with a QLIKE loss function for the 15 training sets obtained from the SPX time series. With respect to the first layer, it can be observed that the HARNet strongly focuses on the RV at time points $t - 1$ and $t - 2$, while the importance of the other days within the past week gradually diminishes. In particular, the RV at the day before yesterday played a significant role in all fits of the model. A possible explanation for this is that single RV jumps can only be distinguished from periods of generally high RV when at least considering the values at $t - 1$ and $t - 2$ in the time series. In the second layer, we can observe an almost linearly diminishing importance of the weeks in the past month.

6 Additional regressors

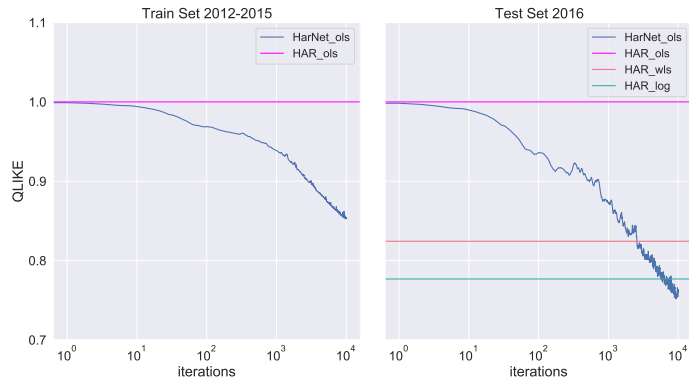
Incorporating additional information contained in the time series of high-frequency returns can significantly improve the accuracy of RV forecasting models. Motivated by Patton and Sheppard



(a) HARNet with aggregation periods (1, 5, 20).



(b) HARNet with aggregation periods (1, 5, 20, 40, 80).



(c) A HARNet model initialized with the OLS fit of the baseline model outperforms the OLS, WLS, and log-OLS fit of the baseline model on the test set after optimization

Figure 8: Performance of the $\Phi_{\theta}^{\text{HAR}20}$ and $\Phi_{\theta}^{\text{HAR}80}$ models relative to different fits of the HAR baseline model. All HARNet models were initialized with the OLS fit of the baseline model and optimized with the QLIKE loss function. The method *best* means that the decision whether to use the OLS, WLS, or log-OLS approach for fitting the baseline model used in the comparison with the respective HARNet was made *after* knowing the test errors. Each boxplot depicts the distribution of the relative test error for 15 training set/test set pairs taken from a time series of the respective index that ranges from the year 2002 until 2020.

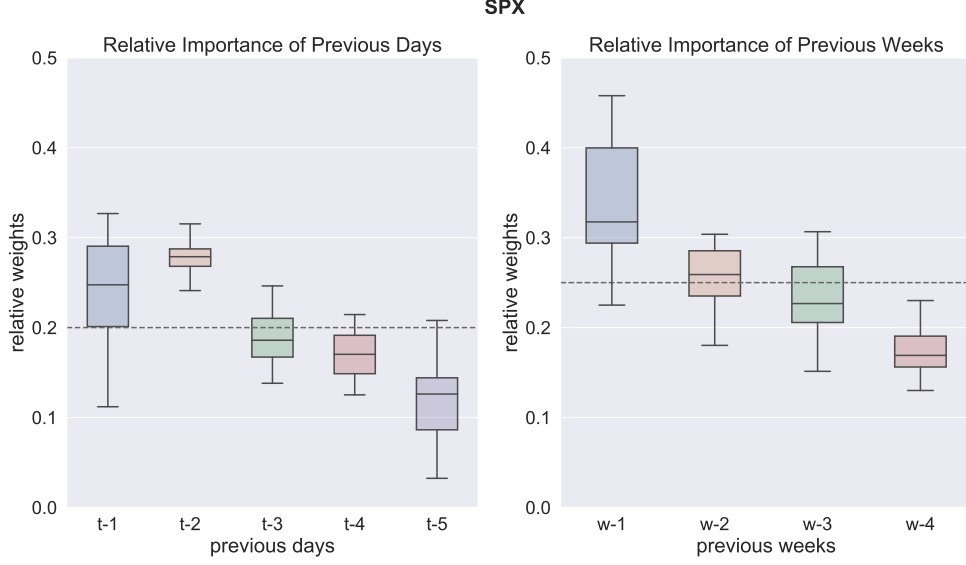


Figure 9: Distribution of normalized filter weights in the layers of a $\Phi_{\theta}^{\text{HAR20}}$ model which was initialized with the OLS fit of the HAR baseline model and optimized with a QLIKE loss function for the 15 training sets obtained from the SPX time series.

(2015) we consider the negative realized semivariance and the signed jump variation as additional inputs to our models. The ability of neural networks to identify nonlinear relations in the data might prove beneficial in a setting where this type of information is available. In line with Barndorff-Nielsen, Kinnebrock, and Shephard (2008), the realized semivariance decomposes the realized variance into components associated with positive and negative returns, respectively. Analogously to RV_t (cf. equation (3)), these components are defined as

$$RS_t^+ = \sum_{n=0}^{N-1} r_{t,n}^2 \mathbb{1}_{r_{t,n} > 0}, \quad (26)$$

$$RS_t^- = \sum_{n=0}^{N-1} r_{t,n}^2 \mathbb{1}_{r_{t,n} < 0}, \quad (27)$$

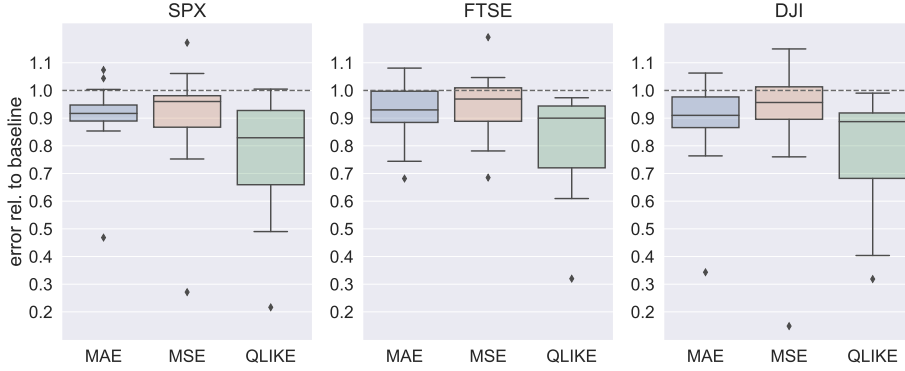
where $r_{t,n}$ are the intraday returns and $\mathbb{1}$ denotes the indicator function. We denote $RS_t^{+(j)}$ and $RS_t^{-(j)}$ as the respective j -day averages. The signed jump variation is defined as the difference between RS_t^+ and RS_t^- , that is,

$$\Delta J_t^2 = RS_t^+ - RS_t^-. \quad (28)$$

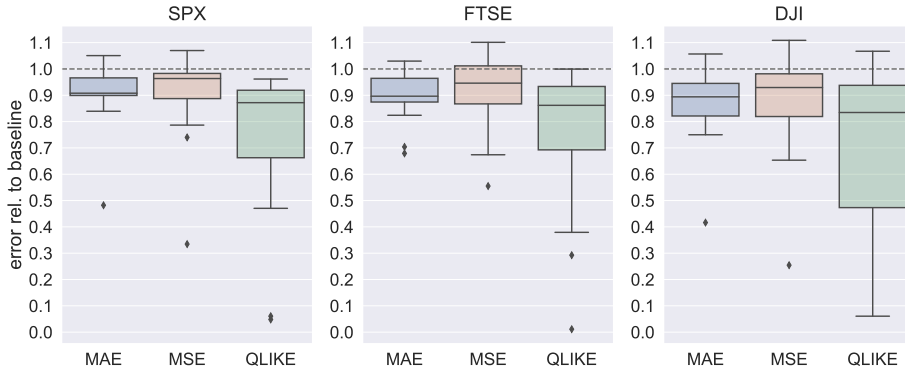
Accordingly, we consider the following linear HAR-type baseline model:

$$\begin{aligned} RV_t = & \beta_0 + \beta_1 RV_{t-1} + \beta_2 RV_{t-1}^{(5)} + \beta_3 RV_{t-1}^{(20)} \\ & + \beta_4 RS_{t-1}^- + \beta_5 RS_{t-1}^{-(5)} + \beta_6 RS_{t-1}^{-(20)} \\ & + \beta_7 \Delta J_{t-1}^2, \end{aligned} \quad (29)$$

with parameters $\beta_0, \dots, \beta_7 \in \mathbb{R}$. The corresponding HARNets with aggregation periods (1, 5, 20) and (1, 5, 20, 40, 80) are defined by extending the model such that the input to the first convolutional layer is a two-dimensional time series, consisting of the entries for RV_t and RS_t^- , and each layer has two two-dimensional filters with the same length as in equation (18). Using a similar scheme as introduced in the case of a one-dimensional input time series, the model parameters of these extended HARNets can again be initialized such that they exactly reproduce the predictions of the linear baseline model (29). Notice that, when using the log-OLS approach, we cannot directly use signed jump variations as inputs since these are not always positive. In this case, as suggested



(a) Extended HARNet with aggregation periods (1, 5, 20).



(b) Extended HARNet with aggregation periods (1, 5, 20, 40, 80).

Figure 10: Performance of the extended $\Phi_{\theta}^{\text{HAR}20}$ and $\Phi_{\theta}^{\text{HAR}80}$ models after optimization with the QLIKE loss function in terms of the MAE, MSE, and the QLIKE test error relative to the HAR baseline fit. The baseline fit was obtained by OLS. The input time series additionally contained values for the daily realized semivariance and the signed jump variation. Each boxplot depicts the distribution of the relative test error for 15 training set/test set pairs taken from a time series of the respective index that ranges from the year 2002 until 2020.

in the appendix to Patton and Sheppard (2015), we replace ΔJ_t^2 by a percentage jump variation measure defined as $1 + \Delta J_t^2 / RV_t$, prior to logarithmizing.

Following the recommendations obtained in Section 5, we restrict our numerical experiments in this setting to the case where the QLIKE error is used as a loss function and the NN weights are initialized by the OLS fit of the baseline extended HAR model. Using the same pairs of test and training sets as in Section 5, Figure 10 compares the respectively optimized HARNets to the OLS fit of the baseline HAR model in terms of the MAE, MSE, and QLIKE test error. Again, we find that the extended HARNets can consistently and often substantially improve upon the accuracy of the baseline model with respect to all three test metrics and with larger margins than in our previous results from Section 5.2.

Figure 11 shows an analysis of the performance of QLIKE optimized HARNets with an OLS baseline fit relative to the OLS, WLS, and log-OLS fits of the respective extended baseline models. Similar to our results in Section 5.3, the extended $\Phi_{\theta}^{\text{HAR}20}$ model still outperforms all of the three approaches for fitting the baseline model and matches the performances of the method where the optimal fitting method was selected a-posteriori. However, we also found a slight decrease in the relative performance of the extended $\Phi_{\theta}^{\text{HAR}80}$ model, which is slightly outperformed by the log-OLS fit of the HAR baseline model with the same aggregation periods.

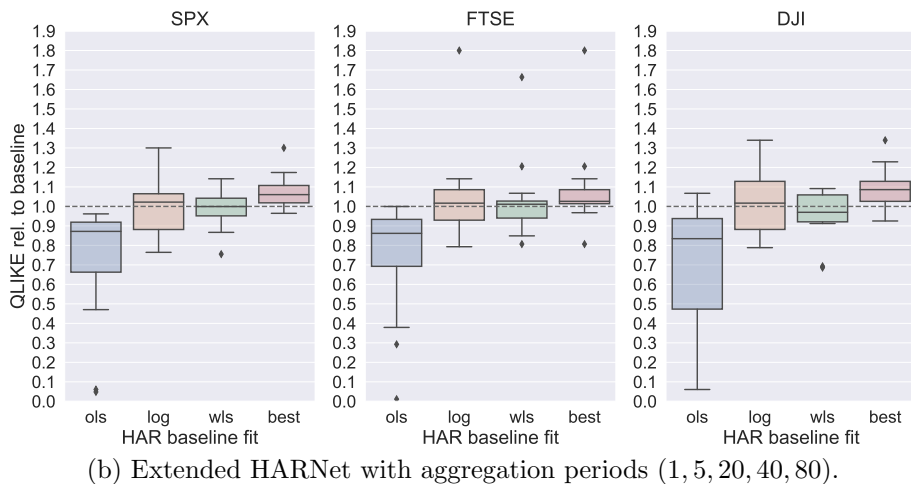
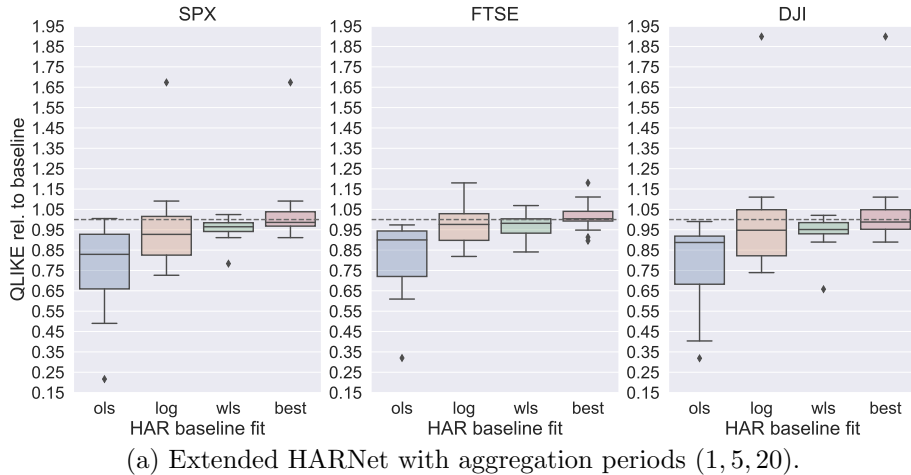


Figure 11: Performance of the extended $\Phi_{\theta}^{\text{HAR}20}$ and $\Phi_{\theta}^{\text{HAR}80}$ models relative to different fits of the HAR baseline model. All HARNets were initialized by OLS estimation of the baseline model and optimized with the QLIKE loss function. The method *best* means that the decision whether to use the OLS, WLS, or log-OLS approach for fitting the baseline model used in the comparison with the respective HARNet was made *after* knowing the test errors. The input time series additionally contained values for the daily realized semivariance and the signed jump variation. Each boxplot depicts the distribution of the relative test error for 15 training set/test set pairs taken from a time series of the respective index that ranges from the year 2002 until 2020..

7 Conclusions

We proposed a novel NN model, dubbed HARNet, for the forecasting of RV time series. HARNets are based on hierarchies of dilated causal convolutions. Their close relationship to the well-known linear HAR model allows for an explicit initialization scheme after which the initialized HARNet is computationally equivalent to the fitted HAR baseline model. This approach significantly stabilizes the optimization of HARNets and facilitates a comprehensive evaluation of their performance relative to the respective baseline models.

We found that HARNets can substantially improve upon the forecasting accuracy of the respective HAR baseline models. In particular, HARNets that are initialized by OLS fit of the HAR baseline model and optimized by a QLIKE loss function outperform OLS, WLS, and log-OLS fits of the HAR baseline model when considering the daily RV time series as the sole input.

While in practice, NN models are often treated as black boxes, the model parameters of a HARNet can be clearly interpreted. In this respect, we found that the filter weights, which define the features at weekly and monthly time horizons of a HARNet follow a distinct pattern after

optimization. In particular, we found that within the previous week, yesterday and the day before yesterday make by far the most significant contributions to today’s RV forecast. We suspect that this reflects the fact that single jumps can only be discriminated from extended periods of high RV when looking at least two days in the past. Regarding the monthly time horizon, we found that within the previous month, the importance of single weeks for the RV forecast seems to diminish almost linearly when moving further into the past.

Hierarchies of dilated convolutional layers allow for an exponential growth of the receptive field with the number of model parameters. This suggests that models such as the HARNet should be particularly well-suited for resolving possible long-term dependencies in the data without the necessity of using an overparameterized model. While a HARNet with a receptive field that covers the past four months outperforms the respective HAR baseline model, this advantage disappears when also considering realized semivariances and signed jump variation as additional inputs. In general, our current findings do not suggest that HARNets are consistently better at understanding long-term relationships in a time series than the respective linear baseline models. We think that this issue is a clear candidate for future research and suspect that a well-considered integration of significantly more training data and a diverse set of exogenous input variables could be crucial to succeed in this matter.

Contrary to previous investigations (Rahimikia and Poon, 2020), the QLIKE loss function turned out to be highly applicable for the optimization of HARNets. We collected strong experimental evidence that this can be explained by a stabilizing effect of our proposed initialization scheme on the optimization process. In general, the idea of initializing a complex non-linear model such that it is computationally equivalent to a linear baseline model before optimization is not restricted to RV forecasting and the proposed scheme could easily be adapted for any autoregression task where linear models already yield competitive results.

A TensorFlow implementation of the HARNet architecture is available on github at <https://github.com/mdsunivie/HARNet>.

8 Acknowledgements

R.R. gratefully acknowledges support from the Austrian Science Fund (FWF M 2528). The computational results presented have been achieved using the Vienna Scientific Cluster (VSC).

References

- Aït-Sahalia, Y. and J. Jacod (2014). *High-frequency financial econometrics*. Princeton University Press.
- Andersen, T. G. and T. Bollerslev (1998). “Answering the skeptics: Yes, standard volatility models do provide accurate forecasts”. In: *International economic review*, pp. 885–905.
- Andersen, T. G. et al. (2001). “The Distribution of Realized Exchange Rate Volatility”. In: *Journal of the American Statistical Association* 96.453, pp. 42–55. DOI: 10.1198/016214501750332965. eprint: <https://doi.org/10.1198/016214501750332965>. URL: <https://doi.org/10.1198/016214501750332965>.
- Barndorff-Nielsen, O. E., S. Kinnebrock, and N. Shephard (2008). “Measuring downside risk-realised semivariance”. In: *CREATES Research Paper* 2008-42.
- Barndorff-Nielsen, O. E. and N. Shephard (2002). “Econometric analysis of realized volatility and its use in estimating stochastic volatility models”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64.2, pp. 253–280.
- Borovykh, A., S. Bohte, and C. W. Oosterlee (2017). “Conditional time series forecasting with convolutional neural networks”. In: *arXiv preprint arXiv:1703.04691*.
- Bucci, A. (June 1, 2020). “Realized Volatility Forecasting with Neural Networks”. In: *Journal of Financial Econometrics* 18.3, pp. 502–531.
- Christensen, K., M. Siggaard, B. Veliyev, et al. (2021). *A machine learning approach to volatility forecasting*. Vol. 3. Department of Economics and Business Economics, Aarhus University.
- Corsi, F. (2009). “A simple approximate long-memory model of realized volatility”. In: *Journal of Financial Econometrics* 7.2, pp. 174–196.

- Dauphin, Y. N. et al. (2017). “Language modeling with gated convolutional networks”. In: *International conference on machine learning*. PMLR, pp. 933–941.
- Gehring, J. et al. (2016). “A convolutional encoder model for neural machine translation”. In: *arXiv preprint arXiv:1611.02344*.
- Gehring, J. et al. (2017). “Convolutional sequence to sequence learning”. In: *International Conference on Machine Learning*. PMLR, pp. 1243–1252.
- Glorot, X. and Y. Bengio (2010). “Understanding the difficulty of training deep feedforward neural networks”. In: *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, pp. 249–256.
- Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep Learning*. <http://www.deeplearningbook.org>. MIT Press.
- Heber, G. et al. (2009). *Oxford-Man Institute’s realized library, version 0.3*. Oxford-Man Institute, University of Oxford.
- Hinton, G., N. Srivastava, and K. Swersky (2012). “Neural networks for machine learning lecture 6a overview of mini-batch gradient descent”. In: *Cited on 14.8*, p. 2.
- Hochreiter, S. and J. Schmidhuber (1997). “Long short-term memory”. In: *Neural computation* 9.8, pp. 1735–1780.
- Kalchbrenner, N. et al. (2016). “Neural machine translation in linear time”. In: *arXiv preprint arXiv:1610.10099*.
- Kingma, D. P. and J. Ba (2017). *Adam: A Method for Stochastic Optimization*. arXiv: 1412.6980 [cs.LG].
- Krizhevsky, A., I. Sutskever, and G. E. Hinton (2012). “Imagenet classification with deep convolutional neural networks”. In: *Advances in neural information processing systems* 25, pp. 1097–1105.
- Miura, R., L. Pichl, and T. Kaizoji (2019). “Artificial Neural Networks for Realized Volatility Prediction in Cryptocurrency Time Series”. In: *Advances in Neural Networks – ISNN 2019*. Ed. by H. Lu, H. Tang, and Z. Wang. Vol. 11554. Series Title: Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 165–172.
- Oord, A. v. d. et al. (2016). “Wavenet: A generative model for raw audio”. In: *arXiv preprint arXiv:1609.03499*.
- Patton, A. J. (2011). “Volatility forecast comparison using imperfect volatility proxies”. In: *Journal of Econometrics* 160.1, pp. 246–256.
- Patton, A. J. and K. Sheppard (July 2015). “Good Volatility, Bad Volatility: Signed Jumps and The Persistence of Volatility”. In: *Review of Economics and Statistics* 97.3, pp. 683–697. ISSN: 0034-6535, 1530-9142.
- Patton, A. J. and K. Sheppard (2009). “Optimal combinations of realised volatility estimators”. In: *International Journal of Forecasting* 25.2, pp. 218–238.
- Proietti, T. and H. Lütkepohl (2013). “Does the Box–Cox transformation help in forecasting macroeconomic time series?” In: *International Journal of Forecasting* 29.1, pp. 88–99.
- Rahimikia, E. and S.-H. Poon (2020). “Machine learning for realised volatility forecasting”. In: *Available at SSRN 3707796*.
- Rumelhart, D. E., G. E. Hinton, and R. J. Williams (1986). “Learning representations by back-propagating errors”. In: *Nature* 323.6088, pp. 533–536.

Recent Issues

All CFS Working Papers are available at www.ifk-cfs.de.

No.	Authors	Title
679	Refet S. Gürkaynak, Burçin Kısacıkoğlu, and Sang Seok Lee	<i>Exchange Rate and Inflation under Weak Monetary Policy: Turkey Verifies Theory</i>
678	Hans Gersbach and Sebastian Zelzner	<i>Why Bank Money Creation?</i>
677	Volker Brühl	<i>Green Financial Products in the EU – a critical review of the status quo</i>
676	Brandon Gipper, Luzi Hail, and Christian Leuz	<i>Determinants and Career Consequences of Early Audit Partner Rotations</i>
675	Matthias Breuer, Christian Leuz, and Steven Vanhaverbeke	<i>Reporting Regulation and Corporate Innovation</i>
674	Andrew Ellul, Marco Pagano, and Annalisa Scognamiglio	<i>Careers in Finance</i>
673	Marco Pagano, Christian Wagner, and Josef Zechner	<i>Disaster Resilience and Asset Prices</i>
672	Carlo Altavilla, Andrew Ellul, Marco Pagano, Andrea Polo, and Thomas Vlassopoulos	<i>Loan Guarantees, Bank Lending and Credit Risk Reallocation</i>
671	Stefano Mengoli, Marco Pagano, and Pierpaolo Pattitoni	<i>The Geography of Investor Attention</i>
670	Lutz Kilian and Xiaoqing Zhou	<i>The Impact of Rising Oil Prices on U.S. Inflation and Inflation Expectations in 2020-23</i>