



Quantile convolutional neural networks for Value at Risk forecasting

Gábor Petneházi *

Doctoral School of Mathematical and Computational Sciences, University of Debrecen, Hungary

ARTICLE INFO

Keywords:

Quantile regression
Convolutional neural network
Value at Risk
Forecasting

ABSTRACT

This article presents a new method for forecasting Value at Risk. Convolutional neural networks can do time series forecasting, since they can learn local patterns in time. A simple modification enables them to forecast not the mean, but arbitrary quantiles of the distribution, and thus allows them to be applied to *VaR*-forecasting. The proposed model can learn from the price history of different assets, and it seems to produce fairly accurate forecasts.

1. Introduction

Convolutional neural networks have shown great results in time series forecasting. However, the applications so far, as time series forecasting in general, focused mainly on predicting the mean. This article presents a convolutional neural network for forecasting quantiles.

The QCNN model is applied to Value at Risk (*VaR*) forecasting. *VaR* is a quantile of the loss distribution, thus it may be forecasted using quantile regression—either simple or deep quantile regression.

Deep neural networks are powerful machine learning algorithms, but they require large amounts of data. It means that the price history of a single asset may not be enough to properly train the network. Here we show that it might be beneficial to use data on multiple assets to forecast the Value at Risk of any single one.

Contributions. As far as I know, this study is among the first to apply a dilated causal convolutional neural network to quantile regression, and to apply it to Value at Risk forecasting, particularly. Though the use of the pinball loss function is not a novelty, and not a major contribution, CNN-based time series forecasting has just recently become a popular research topic, and this area of application has not yet been explored. Another contribution of this study is to show that joint model training may work well for *VaR*-forecasting. This latter finding has key importance regarding the application of convolutional neural networks (and deep learning, in general) to *VaR*-forecasting.

2. QCNN

This section describes the proposed Quantile Convolutional Neural Network (QCNN) model.

2.1. CNNs for time series forecasting

Convolutional neural networks use slided local receptive fields to find local features in the input data. It enables them to model certain data types particularly well, for example, images (with strong 2D structure) or time series (with strong 1D structure). Variables spatially or temporally nearby are often correlated, so we should take advantage of the topology of the inputs (LeCun et al., 1995). This is what convolutional neural networks can do very neatly.

CNNs are most often associated with images, yet they might be applied to one-dimensional data as well. Even the earliest convolutional architectures were applied to the time domain (Waibel et al., 1995). Yet, several notable time series forecasting convolutional architectures were just proposed recently.

Recurrent neural networks (LSTMs and GRUs) are often considered the best or even the optimal neural network architectures for sequential data (including time series). However, Bai et al. (2018) compared generic recurrent and convolutional architectures on various sequence modeling tasks, and found that the latter might be a better choice.

Learning long temporal dependencies is a challenging task for CNNs, but dilations (Yu & Koltun, 2015) can help. A dilated convolution means a convolution with holes, that is, the convolutional filter is enlarged by skipping a few points. The l th layer's dilation rate can be set to 2^{l-1} in order to allow an exponential growth in the effective receptive field: by increasing the number of layers, we can exponentially increase the time horizon that the network can see. The convolution should also be causal, meaning that outputs can never depend on future inputs.

The WaveNet model uses dilated causal convolutions for generating audio waveforms (Oord et al., 2016). Borovykh et al. (2017) used an adaptation of the WaveNet for time series forecasting, and found it an easy to implement and time-efficient alternative to recurrent networks.

QCNN is a one-dimensional dilated causal convolutional network with an appropriately chosen quantile loss function.

* Correspondence to: H-4032 Debrecen, Egyetem tér 1, Hungary.

E-mail address: gabor.petnehazi@science.unideb.hu.

2.2. Quantile regression

While simple least squares regressions estimate the conditional mean of a given variable, quantile regressions estimate conditional quantiles (Koenker & Hallock, 2001). This requires a new objective (1). Mean squared error targets the mean, this targets arbitrary θ -quantiles. Thus, we should just change loss function.

$$\theta \sum_{y_i \geq \hat{q}_i} (y_i - \hat{q}_i) + (\theta - 1) \sum_{y_i < \hat{q}_i} (y_i - \hat{q}_i) \quad (1)$$

2.3. Training

The loss function can be minimized using any proven optimizer algorithm (e.g., stochastic gradient descent or its variants).

Neural networks provide greater flexibility than simple linear models. However, a single time series might not be enough to exploit this flexibility. Thus, we may expect QCNN to work better when trained jointly on a set of similar time series.

3. Empirical study

This section presents the application of QCNN to Value at Risk forecasting.

3.1. Value at risk forecasting

Value at Risk (VaR) is an important risk measure in finance. It aims to find a realistically worst-case outcome in the sense that anything worse happens with a given (small) probability (Shin, 2010). VaR is the worst loss over a horizon that will not be exceeded with a given level of confidence (Jorion, 2000). It is reported as a positive number, by convention.

Mathematically, the VaR of X with a confidence level $1 - \theta$ can be defined as the quantile function of $-X$ at $1 - \theta$ (2).

$$VaR_{\theta}(X) = F_{-X}^{-1}(1 - \theta) = -\inf \{x \in \mathbb{R} : \theta \leq F_X(x)\} \quad (2)$$

VaR can be estimated in many ways. A usual procedure estimates the return variance, and assumes a normal (or t) distribution to compute the quantile estimates. ARCH/GARCH models are often used to analyze and forecast volatility (Engle, 2001). However, these distributional assumptions are often invalid. Efforts have been made to relax them, see, for example, Hull and White (1998) or Glasserman et al. (2002). One of the great advantages of the quantile regression approach is that it does not require any such assumptions.

Quantile regression methods are often applied to Value at Risk forecasting. Engle and Manganelli (1999) proposed the CAViaR (Conditional Value at Risk By Quantile Regression) model. Taylor (2000) applied a quantile regression neural network approach to VaR estimation. Chen and Chen (2002) found that the quantile regression approach outperforms the variance-covariance approach. Taylor (2007) developed an exponentially weighted quantile regression. White et al. (2015) proposed a vector autoregressive extension to quantile models. Xu et al. (2016) applied a quantile autoregression neural network (QARNN) as a flexible generalization of quantile autoregression (Koenker & Xiao, 2006). Yan et al. (2018) used a long short-term memory neural network in a quantile regression framework to learn the tail behavior of financial asset returns. Just to mention a few notable studies.

VaR deals with rare events, and rare events produce small data. The higher the confidence level we choose, the fewer loss events we have to learn from. Data volume is thus a crucial issue in Value at Risk forecasting. It is, therefore, beneficial to use several stocks' data to learn to forecast VaR . More stocks have experienced more extreme events, and we may expect them to have similar sources. If they do so, then a joint VaR -forecasting model might outperform individual models.

3.2. Data

Our stock market dataset was obtained from Kaggle.¹ We have randomly chosen 100 stocks listed on NASDAQ, NYSE, and AMEX in the 10-year period under study (2009-01-01 to 2018-12-31). Daily logarithmic returns were computed and fed to the algorithms. The last 30% of each time series was used as a test set, that is, about the last 3 years.

3.3. Models and experiments

VaRs are forecasted for 3 confidence levels (95%, 99%, and 99.9%), using 5 different methods:

- a constant quantile estimate,
- a GARCH(1,1) with normal distribution,
- a linear quantile regression,
- a QCNN,
- a joint QCNN trained on all available training data.

The QCNN network contains 6 causal convolutional layers (each with 8 filters of kernel size 2, with *relu* activation, and exponentially increasing dilation rates), and a convolutional layer with a single filter of kernel size 1 and a linear activation. The Python implementation of the QCNN network is available on GitHub.²

The inputs are overlapping 128-step sequences extracted from the time series of stock returns. The output sequences are the inputs shifted by 1 step, so that at any time we predict VaR one day ahead.

The dataset was scaled by subtracting the mean and dividing by the standard deviation, and it was fed to the algorithms in 128-batches. The QCNN models were trained for 128 epochs, using the adadelta (Zeiler, 2012) optimizer.

The linear model is an autoregression using 4 lags of the target variable.

The constant quantile estimate was computed using linear interpolation (3).

$$x_{[i]} + (i - [i])(x_{[i]+1} - x_{[i]}), \quad i = (N - 1)\theta + 1 \quad (3)$$

3.4. Evaluation

We apply the Dynamic Quantile test of Engle and Manganelli (1999). A *Hit* variable (4) is constructed, which takes the value $1 - \theta$ at VaR exceedances, and takes $-\theta$ else, and so has an expected value of zero. This variable is regressed on $X(5)$, which may contain *Hit*'s past lags, the Value at Risk (and its lags), and possibly further variables. This regression should have no explanatory power, so we test the hypothesis $H_0 : \delta = 0$. Applying the central limit theorem, we can construct an asymptotically chi-square distributed test statistic (6). q denotes the number of input variables, t denotes the time steps.

$$Hit_t = I(y_t < -VaR_t) - \theta \quad (4)$$

$$Hit = X\delta + u \quad (5)$$

$$\frac{Hit' X [X' X]^{-1} X' Hit}{\theta(1 - \theta)} \stackrel{a}{\sim} \chi^2(q) \quad (6)$$

This test is simple and easy to implement. Here we use VaR , and 3 lags of *Hit* as input variables.

The average estimated VaR values and VaR exceedance rates are also reported for a more detailed assessment of model performance.

3.5. Results

The means, medians, and standard deviations of VaR exceedances, the average rejection rates of the DQ test at two different significance levels, and the average VaR values are reported in Tables 1, 2, and 3. Zero-exceedance forecasts are assigned a 0 p -value for the DQ test. The DQ test results are not reported for the 99.9% VaR level.

¹ <https://www.kaggle.com/qks1lver/amex-nyse-nasdaq-stock-histories>

² <https://gist.github.com/bandeeka/10b8cd050a912cacdb0f7b41c4fa1422>

Table 1
95% VaR forecasts.

	Exceedances			DQ Test Rejections		VaR
	Mean	Median	SD	0.01	0.05	Mean
Constant	0.0399	0.0358	0.0291	0.5300	0.6600	0.0727
GARCH	0.0347	0.0338	0.0134	0.2100	0.4100	0.0739
QR	0.0409	0.0344	0.0293	0.6200	0.7200	0.0689
QCNN	0.1269	0.1245	0.0362	0.7400	0.8000	0.0423
Joint QCNN	0.0433	0.0444	0.0101	0.0500	0.1600	0.0583

Table 2
99% VaR forecasts.

	Exceedances			DQ Test Rejections		VaR
	Mean	Median	SD	0.01	0.05	Mean
Constant	0.0084	0.0061	0.0117	0.3800	0.4100	0.1935
GARCH	0.0133	0.0119	0.0066	0.1400	0.2000	0.1039
QR	0.0097	0.0066	0.0125	0.5400	0.5900	0.1852
QCNN	0.0599	0.0576	0.0280	0.9500	0.9600	0.0676
Joint QCNN	0.0115	0.0119	0.0056	0.0900	0.1600	0.1023

Table 3
99.9% VaR forecasts.

	Exceedances			VaR
	Mean	Median	SD	Mean
Constant	0.0023	0.0007	0.0052	0.3764
GARCH	0.0060	0.0053	0.0041	0.1375
QR	0.0041	0.0013	0.0075	0.3630
QCNN	0.0249	0.0212	0.0186	0.1079
Joint QCNN	0.0023	0.0013	0.0026	0.2062

The single-stock QCNN overshot the targeted *VaR*-exceedance rates, in some cases by orders of magnitude. All other methods produced more conservative, and seemingly more reliable forecasts. Even the constant *VaR* estimate produced mean exceedance rates close to the desired levels. However, the joint QCNN achieved similar exceedance accuracy with lower average *VaR* estimates. It means that while a simple historical quantile estimate might be enough to set a budget against a realistically worst-case outcome, the QCNN gives a cheaper solution. Also, the joint QCNN produced *VaR* exceedance rates with consistently lower standard deviation than the benchmark methods. The Dynamic Quantile test is rejected for much fewer stocks in case of the joint QCNN than any other method, which also justifies that this one produces the highest quality *VaR* estimates.

The experiments were repeated for the previous 10 years' data (1999–2008) with a different set of randomly chosen stocks, and the results were quite similar.

4. Summary

This article proposed a one-dimensional convolutional neural network architecture for Value at Risk forecasting. The model takes as input a series of observations (returns), and directly makes one-step ahead forecasts for arbitrary quantiles. When trained jointly on multiple time series (that is, using data on multiple assets), the proposed method seems to perform well.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

There is no funding to report.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.mlwa.2021.100096>.

References

- Bai, Shaojie, Kolter, J. Zico, & Koltun, Vladlen (2018). An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv preprint [arXiv:1803.01271](https://arxiv.org/abs/1803.01271).
- Borovykh, Anastasia, Bohte, Sander, & Oosterlee, Cornelis W. (2017). Conditional time series forecasting with convolutional neural networks. arXiv preprint [arXiv:1703.04691](https://arxiv.org/abs/1703.04691).
- Chen, Mei-Yuan, & Chen, Jau-Er (2002). Application of quantile regression to estimation of value at risk. *Review of Financial Risk Management*, 1(2), 15.
- Engle, Robert (2001). GARCH 101: The use of ARCH/GARCH models in applied econometrics. *Journal of Economic Perspectives*, 15(4), 157–168.
- Engle, Robert F., & Manganelli, Simone (1999). *CAViaR: conditional value at risk by quantile regression: Technical report*, National Bureau of Economic Research.
- Glasserman, Paul, Heidelberger, Philip, & Shahabuddin, Perwez (2002). Portfolio value-at-risk with heavy-tailed risk factors. *Mathematical Finance*, 12(3), 239–269.
- Hull, John, & White, Alan (1998). Value at risk when daily changes in market variables are not normally distributed. *Journal of Derivatives*, 5, 9–19.
- Jorion, Philippe (2000). Value at risk.
- Koenker, Roger, & Hallock, Kevin F. (2001). Quantile regression. *Journal of Economic Perspectives*, 15(4), 143–156.
- Koenker, Roger, & Xiao, Zhijie (2006). Quantile autoregression. *Journal of the American Statistical Association*, 101(475), 980–990.
- LeCun, Yann, & Bengio, Yoshua (1995). Convolutional networks for images, speech, and time series. *The Handbook of Brain Theory and Neural Networks*, 3361(10), 1995.
- Oord, Aaron van den, Dieleman, Sander, Zen, Heiga, Simonyan, Karen, Vinyals, Oriol, Graves, Alex, Kalchbrenner, Nal, Senior, Andrew, & Kavukcuoglu, Koray (2016). Wavenet: A generative model for raw audio. arXiv preprint [arXiv:1609.03499](https://arxiv.org/abs/1609.03499).
- Shin, Hyun Song (2010). *Risk and liquidity*. Oxford University Press.
- Taylor, James W. (2000). A quantile regression neural network approach to estimating the conditional density of multiperiod returns. *Journal of Forecasting*, 19(4), 299–311.
- Taylor, James W. (2007). Using exponentially weighted quantile regression to estimate value at risk and expected shortfall. *Journal of financial Econometrics*, 6(3), 382–406.
- Waibel, Alexander, Hanazawa, Toshiyuki, Hinton, Geoffrey, Shikano, Kiyohiro, & Lang, Kevin J (1995). Phoneme recognition using time-delay neural networks. In *Backpropagation: theory, architectures and applications* (pp. 35–61). Mahwah, NJ: Lawrence Erlbaum Assoc..
- White, Halbert, Kim, Tae-Hwan, & Manganelli, Simone (2015). VAR for VaR: Measuring tail dependence using multivariate regression quantiles. *Journal of Econometrics*, 187(1), 169–188.
- Xu, Qifa, Liu, Xi, Jiang, Cuixia, & Yu, Keming (2016). Quantile autoregression neural network model with applications to evaluating value at risk. *Applied Soft Computing*, 49, 1–12.
- Yan, Xing, Zhang, Weizhong, Ma, Lin, Liu, Wei, & Wu, Qi (2018). Parsimonious quantile regression of financial asset tail dynamics via sequential learning. In *Advances in neural information processing systems* (pp. 1575–1585).
- Yu, Fisher, & Koltun, Vladlen (2015). Multi-scale context aggregation by dilated convolutions. arXiv preprint [arXiv:1511.07122](https://arxiv.org/abs/1511.07122).
- Zeiler, Matthew D. (2012). ADADELTA: an adaptive learning rate method. arXiv preprint [arXiv:1212.5701](https://arxiv.org/abs/1212.5701).