

Deep learning of Value at Risk through generative neural network models : the case of the Variational Auto Encoder

Pierre Brugière* and Gabriel Turinici†
CEREMADE, University Paris Dauphine-PSL
Paris, FRANCE

November 2022

Abstract

We present in this paper a method to compute, using generative neural networks, an estimator of the "Value at Risk" for a financial asset. The method uses a Variational Auto Encoder with a 'energy' (a.k.a. Radon-Sobolev) kernel. The result behaves according to intuition and is in line with more classical methods.

1 Motivation

1.1 The law of the returns: sample size, Gaussian assumption and beyond

The goal of this paper is to briefly present a method for computing the "Value at Risk" for a financial asset using generative neural network approaches.

The "Value at Risk", introduced by J.P. Morgan and Reuters [1], measures the maximum loss incurred by a holder of a financial asset, at the end of a given time horizon T (for instance $T = 10$ days) when excluding the worst scenarios (5% worst for the $q = 95\%$ Var, 1% for the $q = 99\%$ VaR and so on)¹; we refer to [1] and to the huge financial and prudential literature on it for more detailed considerations on the VaR.

Of course, as VaR is about the future returns, there is a large uncertainty on the statistical properties of the future returns at the time T . Usual implementations consider the returns as normal variables (thus the prices are log-normal) and fit the Gaussian to historical data. Several problems can appear:

* brugiere@ceremade.dauphine.fr

† gabriel.turinici@dauphine.fr

¹It is thus related to the $1 - q$ quantile of the loss/profit on the asset at the end of the time horizon.

- a lack of robustness because the returns parameters may change with respect to the past or because the available data is too scarce to estimate them accurately

- a fragility due to model risk when some assumptions are made on the law of the distribution of the returns.

Thus it is relevant, to mitigate this risk, to investigate some model free approaches.

1.2 Generative neural networks: Variational Auto Encoders

To address these concerns, we use generative neural networks; these are methods that can sample new objects (similar to ones already in the training set) without parametric assumptions on the probability law on the set of objects.

There are many flavors of generative neural networks, e.g. the Generative Adversarial Networks [2, 4], the Variational Auto Encoders (hereafter abbreviated VAE) [3], etc. We use the latter ones and refer to [5] for details concerning the technical considerations and more precisely on the special kind of VAE named the 'energy' or 'Radon-Sobolev' VAE, that is employed.

2 Deep neural network for VaR computation

A VAE (Variational Auto-Encoder) deep generative network was trained to sample sequences of $252 + 10 = 262$ consecutive trading days. The training set is composed of all blocks of 262 consecutive trading days for the S&P500 index available from Jan 1st 1928 to Nov. 1st 2022, totaling 23561 (overlapping) sequences. In order to avoid scaling issues, a multiplicative factor is applied to ensure that any sequence starts with value 1.

The goal is to produce the distribution probability of the return over the next 10 days, conditional to the last 252 historical data available. This distribution is used to compute the VaR as of Nov 1st 2022.

The VAE is built within the class of 'energy'/'Radon-Sobolev' kernels as described in [5], which, in this case, turns out to exhibit better convergence properties than the GAN, (cf. [4]). After convergence, one obtains a procedure to sample an arbitrary number of scenarios, see figure 1 for an example.

Once possible scenarios are sampled at random, a similarity metric (a weight) is computed for each of them : the closer the first 252 days of the scenario are to the Nov. 1st 2021 - Nov. 1st 2022 historical data (252 days), the more important its weight is; see figure 2 for an illustration. In practice this metric is a softmax function applied to the set of the squares of the euclidian norms of differences between the reference scenario and the members of the set of 10'000 sampled scenarios².

²That is, supposing v is the reference vector and u a generated one, the weight w_u of




Figure 1: Ten random price scenarios as constructed by the deep NN VAE algorithm. For the practical computation we use 10'000 of them. None of them corresponds to a precise historical sequence but all are considered "plausible" by the VAE constructing them. In mathematical terms, we sample from a continuous probability law on \mathbb{R}^{262} close to the discrete probability law consisting of the 23561 (overlapping) sequences in the training set.

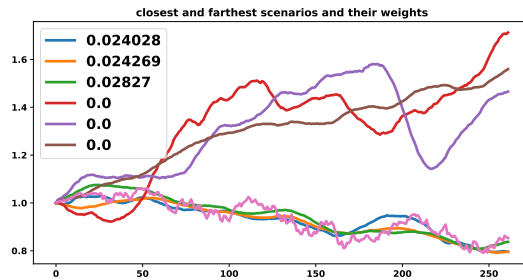


Figure 2: Illustration of the 3 closest and 3 furthest scenarios from the 10'000 synthetic sampled scenarios. The reference is the S&P500 evolution from Nov. 1st 2021 to Nov. 1st 2022 (magenta color). The weights are given in the legend (no weight is printed for the S&P500 historical data in magenta); those appearing as zero are in fact smaller than 10^{-6} .




Figure 3: For each of the 10'000 sampled scenarios we consider the return corresponding to the last 10 days and plot the histogram of the generated data in blue. For comparison the histogram of a Gaussian with same mean and variance is given in orange. The VAE histogram is much wider than the normal equivalent which shows that in particular the distribution is heavy tailed, which is a positive behavior as financial data is known to display such deviations from the (log-) normal model.

The last 10 days of each scenario are now used as samples of 10-days returns as illustrated in figure 3 (without weights); the distribution is renormalised to match the last 252 days 10-day returns mean and variance. The result is more heavy tailed than the Gaussian counterpart, in agreement with econometric information on the financial data.

Finally, from this distribution of 10 days returns, a 99% *VaR* is now inferred; we find a value of 12.8% which is close to the 10 day bootstrap value without shuffling (that is 12.4%). As a comparison, the Gaussian *VaR* on the same data (daily return variance calibrated on the last 252 days i.e. Nov. 1st 2021 to Nov. 1st 2022) is 12.1%³.

3 Conclusion

We present a procedure that leverages on historical data to build, through a generative neural network (VAE), an estimation of the 10-days Value at Risk.

More analysis and fine tuning of the model will be pursued but so far we find the results encouraging especially as we are able to simulate much more scenarios "model free" than with the bootstrap methods while finding results coherent with existing techniques.

the scenario u is $w_u = c_n \cdot e^{-\|v-u\|^2/2\sigma^2}$. Here c_n is a normalization parameter to ensure that the weights sum up to 1 and the parameter σ^2 is the mean error (computed during training) between the initial sequences in the database u and the reconstructed ones ($\sigma^2 = \mathbb{E}_X \|D_{\theta_d} \circ E_{\theta_e}(X) - X\|^2$ with notations in [5, p 299]), see VAE literature for details.

³Recall that the Gaussian 99% 10-days *VaR* is computed as $\sqrt{10}$ times the 0.01 quantile of the empirical distribution from the historical data.

References

- [1] *1996 RiskMetrics Technical Document*. en-US. URL: <https://www.msci.com/www/research-report/1996-riskmetrics-technical/018482266>.
- [2] Ian Goodfellow et al. “Generative Adversarial Nets”. In: *Advances in Neural Information Processing Systems*. Ed. by Z. Ghahramani et al. Vol. 27. Curran Associates, Inc., 2014. URL: <https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf>.
- [3] Diederik P. Kingma and Welling Max. *An Introduction to Variational Autoencoders*. English. Now Publishers Inc, Nov. 2019. ISBN: 978-1-68083-622-6.
- [4] Gabriel Turinici. “Convergence Dynamics of Generative Adversarial Networks: The Dual Metric Flows”. en. In: *Pattern Recognition. ICPR International Workshops and Challenges*. Ed. by Alberto Del Bimbo et al. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2021, pp. 619–634. ISBN: 978-3-030-68763-2. DOI: 10.1007/978-3-030-68763-2_47.
- [5] Gabriel Turinici. “Radon–Sobolev Variational Auto-Encoders”. In: *Neural Networks* 141 (Sept. 2021), pp. 294–305. ISSN: 0893-6080. DOI: 10.1016/j.neunet.2021.04.018. URL: <https://www.sciencedirect.com/science/article/pii/S0893608021001556>.