



HAL
open science

Predicting Stock Returns with Batched AROW

Rachid Guennouni Hassani, Alexis Gilles, Emmanuel Lassalle, Arthur
Dénouveaux

► **To cite this version:**

Rachid Guennouni Hassani, Alexis Gilles, Emmanuel Lassalle, Arthur Dénouveaux. Predicting Stock Returns with Batched AROW. [Research Report] Machina Capital. 2020. hal-02496048v1

HAL Id: hal-02496048

<https://hal.science/hal-02496048v1>

Submitted on 5 Mar 2020 (v1), last revised 9 Mar 2020 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Predicting Stock Returns with Batched AROW

Rachid Guennouni Hassani^{♣¹}, Alexis Gilles^{♦²}, Emmanuel Lassalle^{♥²}
and Arthur Dénouveaux^{♣²}

¹École Polytechnique

²Machina Capital

March 5, 2020

Abstract

We extend the AROW regression algorithm developed by Vaits and Crammer in [VC11] to handle synchronous mini-batch updates and apply it to stock return prediction. By design, the model should be more robust to noise and adapt better to non-stationarity compared to a simple rolling regression. We empirically show that the new model outperforms more classical approaches by backtesting a strategy on S&P500 stocks.

♣ rachid.guennouni-hassani@polytechnique.edu
♦ agilles@machinacap.com
♥ elassalle@machinacap.com
♣ adenouveaux@machinacap.com

1 Introduction

Financial markets exhibit highly non-stationary behaviors, making it difficult to build predictive signals that do not decay too rapidly (see [SCSG13, Con01] for empirical studies of return time series). A standard method for capturing these changes in time series data consists in using a rolling regression, that is, a linear regression model trained on a rolling window and kept as static model during a prediction period. However, the size of historical training data as well as the duration of the prediction period have a direct impact on the performance of the resulting model: using too many training data would result in a model that does not react quickly enough to sudden changes while short training and prediction windows would make the model unstable (see for instance [IJR17]).

Online learning algorithms are suited to situations where data arrives sequentially. New information is taken into account by updating the model parameters in a supervised fashion. More precisely, an online learning algorithm repeats the following steps indefinitely: receive a new instance x_t , make a prediction \hat{y}_t , receive the correct label y_t for the instance and update the model accordingly.

In the particular case of regression, online models are also good candidates to handle the non-stationarity inherent in financial time series while keeping a certain memory of what has been learnt from the beginning. The recursive least squares (RLS) algorithm is a well known approach to online linear regression problems (e.g. [SD91]), yet it updates the model parameters using one sample at a time. However, building predictive models on stock markets, one should take into account the very low signal over noise ratio, and one way to do so is to fit a single model for predicting all the stock returns of a trading universe (hence fitting on much more data). It allows us to have more training data covering shorter periods, but it comes with the difficulty of updating the model parameters synchronously with all the data available at a given time.

In this paper we extend AROW for regression ([VC11]), an algorithm similar to RLS, in order to take into account a batch of instances for the online update instead of doing one update per sample. Indeed the latter approach would introduce a spurious order in information that actually occur synchronously and should be captured as such. Like RLS (see [Hay96]), AROW suffers logarithmic regret in the stationary case, but also comes with a bound on regret in the general case. We test it on the universe of S&P500 stocks on a daily strategy and show it outperforms the rolling regression on the same set of features in backtest.

2 Adaptive Weight Regularization

When training online models, we try to find the right balance between reactivity to new information and accumulation of predictive power over time. For instance, the family of models introduced in [CDK⁺06] aggressively update their parameters by guaranteeing a certain performance on new data (expressed in terms of geometric margin). But even regularized versions of these algorithms do not take into account the fact that some features might be less noisy than others. In other words, they are not suited for cases where we would like updates to be more aggressive on particular parameters than on the rest of the weight vector.

Subsequent online learning algorithms introduced in [DCP08] and [CKD13] maintain a Gaussian distribution on weight vectors representing the confidence the model has in its parameters. From this point of view, the mean (resp. the variance) of the distribution represents the knowledge (resp. the confidence in the parameters). In that framework, when receiving new feature vector and target (x_t, y_t) , we would like the Gaussian parameters to be updated by solving the following optimization problem:

$$\begin{aligned} (\mu_t, \Sigma_t) = \arg \min_{\mu, \Sigma} D_{KL}(\mathcal{N}(\mu, \Sigma) \parallel \mathcal{N}(\mu_{t-1}, \Sigma_{t-1})) \\ \text{s.t. } \mathbb{P}_{w \sim \mathcal{N}(\mu, \Sigma)}(\ell(y_t, w^\top x_t) \leq \epsilon) \geq \eta \end{aligned}$$

where D_{KL} is the Kullback-Leibler divergence and ℓ is the classical mean square loss. The idea behind this optimization is to find the minimal changes in knowledge and confidence such that minimum regression performance is achieved with a given probability threshold η .

Although that formulation is tractable (in particular there is an explicit formula for the KL divergence between two Gaussian distributions), it is not convex in μ and Σ . With AROW, Vaits and Crammer use a simpler but convex objective ([VC11]):

$$\mathcal{C}(\mu, \Sigma) = D_{KL}(\mathcal{N}(\mu, \Sigma) \parallel \mathcal{N}(\mu_{t-1}, \Sigma_{t-1})) + \lambda_1 \ell(y_t, \mu^\top x_t) + \lambda_2 x_t^\top \Sigma x_t \quad (1)$$

where λ_1 and λ_2 are hyperparameters to adjust the tradeoff between the three components of the objectives. These three components should be understood as follows:

1. the parameters should not change much per update;
2. the new model should perform well on the current instance;
3. the uncertainty about the parameters should reduce as we get additional data.

As such, the model is well suited for non-stationary regression. However the updates only take into account a single instance at the time. Because we want a single model for all the stocks here, we need to extend this approach to synchronously update the parameters on all cross-sectional information available at a given time, which is the purpose of the next section.

3 Synchronous Batch Updates

We now assume that between the updates at times t and $t-1$, we have K synchronous instances (x_t^k, y_t^k) , $k = 1, \dots, K$ of the features and targets. They correspond to the observation of a complete universe of stocks at a given time. Applying an AROW update would introduce a fake order between the instances and potentially hurt the performance (see Section 5 for more details). In order to take into account all the new information at once in a single batch, we extend the cost function from equation 1 as follows:

$$\begin{aligned} \mathcal{C}(\mu, \Sigma) &= D_{KL}(\mathcal{N}(\mu, \Sigma) \parallel \mathcal{N}(\mu_{t-1}, \Sigma_{t-1})) \\ &\quad + \frac{\lambda_1}{K} \sum_{k=1}^K \ell(y_t^k, \mu^\top x_t^k) + \frac{\lambda_2}{K} x_t^{k\top} \Sigma x_t^k. \end{aligned}$$

For simplicity of presentation, we set $\lambda_1 = \lambda_2 = \frac{1}{2r}$ where $r > 0$ and $R = rK$. Using

$$\begin{aligned} D_{KL}(\mathcal{N}(\mu_1, \Sigma_1) \parallel \mathcal{N}(\mu_2, \Sigma_2)) &= \frac{1}{2} \log \left(\frac{\det \Sigma_2}{\det \Sigma_1} \right) + \text{Tr}(\Sigma_2^{-1} \Sigma_1) \\ &\quad + \frac{1}{2} (\mu_2 - \mu_1)^\top \Sigma_2^{-1} (\mu_2 - \mu_1) - \frac{d}{2}, \end{aligned}$$

and setting

$$X_t = [x_t^1 \cdots x_t^K]^\top, \quad Y_t = [y_t^1 \cdots y_t^K]^\top.$$

we get, remembering that ℓ is the mean square loss,

$$\begin{aligned} \mathcal{C}(\mu, \Sigma) = & \frac{1}{2} \log \left(\frac{\det \Sigma_{t-1}}{\det \Sigma} \right) + \frac{1}{2} \text{Tr}(\Sigma_{t-1}^{-1} \Sigma) + \frac{1}{2} (\mu_{t-1} - \mu)^\top \Sigma_{t-1}^{-1} (\mu_{t-1} - \mu) \\ & - \frac{d}{2} + \frac{1}{2R} \left(\|Y_t - \mu^\top X_t\|^2 + \text{Tr}(X_t \Sigma X_t^\top) \right). \end{aligned}$$

The main result of this article is the following: minimizing the cost function \mathcal{C} has an explicit solution given by

$$\Sigma_t = \Sigma_{t-1} - \Sigma_{t-1} X_t^\top \left(R \text{Id}_d + X_t \Sigma_{t-1} X_t^\top \right)^{-1} X_t \Sigma_{t-1} \quad (2a)$$

$$\mu_t = \mu_{t-1} - \Sigma_{t-1} X_t^\top \left(R \text{Id}_d + X_t \Sigma_{t-1} X_t^\top \right)^{-1} (X_t \mu_{t-1} - Y_t). \quad (2b)$$

From now on we refer to this batch version of AROW as BAROW, and detail the steps to the solution in the next section.

4 Deriving the Update Formulas

As for AROW, BAROW's cost function \mathcal{C} is convex. To prove 2a and 2b it is thus enough to look at the critical points of \mathcal{C} . We start by computing $\partial \mathcal{C} / \partial \Sigma$ and, using formulas from [PP12] chapter 2, we find

$$\frac{\partial \mathcal{C}}{\partial \Sigma}(\mu, \Sigma) = -\frac{1}{2} \left(\Sigma^{-1} + \Sigma_{t-1}^{-1} + \frac{1}{R} X_t^\top X_t \right).$$

The Kailath variant of the Woodbury identity (see [PP12] chapter 3) yields

$$\left(\Sigma_{t-1}^{-1} + \frac{1}{R} X_t^\top X_t \right)^{-1} = \Sigma_{t-1} - \Sigma_{t-1} X_t^\top \left(R \text{Id}_d + X_t \Sigma_{t-1} X_t^\top \right)^{-1} X_t \Sigma_{t-1}$$

and allows us to deduce that $\partial \mathcal{C} / \partial \Sigma$ vanishes for

$$\Sigma = \Sigma_{t-1} - \Sigma_{t-1} X_t^\top \left(R \text{Id}_d + X_t \Sigma_{t-1} X_t^\top \right)^{-1} X_t \Sigma_{t-1},$$

which is formula 2a.

Similarly, one finds

$$\frac{\partial \mathcal{C}}{\partial \mu}(\mu, \Sigma) = \Sigma_{t-1}^{-1}(\mu - \mu_{t-1}) + \frac{1}{R}(X_t^\top X_t \mu - X_t^\top Y_t),$$

so that $\partial \mathcal{C} / \partial \mu = 0$ if and only if

$$\begin{aligned} \mu &= \left(\Sigma_{t-1}^{-1} + \frac{1}{R} X_t^\top X_t \right)^{-1} \left(\Sigma_{t-1}^{-1} \mu_{t-1} + \frac{1}{R} X_t^\top Y_t \right) \\ &= \left(\Sigma_{t-1}^{-1} + \frac{1}{R} X_t^\top X_t \right)^{-1} \Sigma_{t-1}^{-1} \mu_{t-1} + \frac{1}{R} \left(\Sigma_{t-1}^{-1} + \frac{1}{R} X_t^\top X_t \right)^{-1} X_t^\top Y_t. \end{aligned}$$

We deal with the two terms separately. For the first one we use again the Kailath variant of the Woodbury equality and get

$$\begin{aligned} \left(\Sigma_{t-1}^{-1} + \frac{1}{R} X_t^\top X_t \right)^{-1} \Sigma_{t-1}^{-1} \mu_{t-1} &= \left(\text{Id} - \Sigma_{t-1} X_t^\top (R \text{Id} + X_t \Sigma_{t-1} X_t)^{-1} X_t \right) \mu_{t-1} \\ &= \mu_{t-1} - \Sigma_{t-1} X_t^\top (R \text{Id} + X_t \Sigma_{t-1} X_t)^{-1} X_t \mu_{t-1}. \end{aligned}$$

For the second term, first notice that:

$$\begin{aligned} \left(\Sigma_{t-1}^{-1} + \frac{1}{R} X_t^\top X_t \right) \Sigma_{t-1} X_t^\top &= \left(X_t^\top + \frac{1}{R} X_t^\top X_t \Sigma_{t-1} X_t^\top \right) \\ &= \frac{1}{R} X_t^\top (R \text{Id} + X_t \Sigma_{t-1} X_t^\top) \end{aligned}$$

so that

$$\left(\Sigma_{t-1}^{-1} + \frac{1}{R} X_t^\top X_t \right) \Sigma_{t-1} X_t^\top (R \text{Id} + X_t \Sigma_{t-1} X_t^\top)^{-1} Y_t$$

is equal to

$$\frac{1}{R} X_t^\top (R \text{Id} + X_t \Sigma_{t-1} X_t^\top) (R \text{Id} + X_t \Sigma_{t-1} X_t^\top)^{-1} Y_t.$$

Simplifying $(R \text{Id} + X_t \Sigma_{t-1} X_t^\top)$ we get

$$\frac{1}{R} \left(\Sigma_{t-1}^{-1} + \frac{1}{R} X_t^\top X_t \right)^{-1} X_t^\top Y_t = \Sigma_{t-1} X_t^\top (R \text{Id} + X_t \Sigma_{t-1} X_t^\top)^{-1} Y_t.$$

Finally, combining the two terms, we get

$$\begin{aligned} \mu &= \mu_{t-1} - \Sigma_{t-1} X_t^\top (R \text{Id} + X_t \Sigma_{t-1} X_t^\top)^{-1} X_t \mu_{t-1} + \Sigma_{t-1} X_t^\top (R \text{Id} + X_t \Sigma_{t-1} X_t^\top)^{-1} Y_t \\ &= \mu_{t-1} - \Sigma_{t-1} X_t^\top (R \text{Id} + X_t \Sigma_{t-1} X_t^\top)^{-1} (X_t \mu_{t-1} - Y_t), \end{aligned}$$

which was the claimed formula.

5 Backtesting a Strategy using BAROW

BAROW combines the adaptability of AROW to non-stationary data and the advantage of taking into account all new information synchronously. We tested the model against two baselines as return predictor for a trading strategy on the S&P500 universe.

We backtested a long-short strategy taking daily positions proportional to the prediction generated by a regression model on stock returns and showed it outperforms the following baselines:

1. A rolling regression updated daily and using the past 12 months of data for training.
2. AROW regression with single instance updates (500 updates per day).




Figure 1 – Estimated Return of each model

We ran the backtest on 2,250 days from 2011 to 2019, allowing a burn-in period of 12 months for AROW and BAROW before starting to use them as return predictors. We also tweaked the R hyperparameter on 2010 for AROW and BAROW. To avoid the strategy having too high sensitivity to market moves, we neutralized daily returns using a multi-factor model (including beta, volatility and a variety of other indicators, see [KL15] for a detailed discussion about risk models). These neutralized returns are the regression targets y_t . We used features based on MACD indicators such as those described in [CNL14].

We estimated daily returns of a strategy taking positions proportional to the predictions as the cross-sectional correlation between the predictions and the realized returns, multiplied by the cross-sectional standard deviation of the returns.

As expected, we can see that the baseline using AROW updates stock per stock is not suited to learn a generic predictor for the universe. Its cumulative return appears random compared to the two other models. That empirically justifies the batch update computed in BAROW.

Now if we compare BAROW to the rolling linear regression, we directly see that performances drift away from each other. BAROW significantly outperforms the baselines, illustrating the benefits of online learning models for non-stationary data.

model	Return	Sharpe	MaxDD	Calmar
AROW	0.5%	0.1	-3.02%	0.18
Linear	31.3%	4.0	-0.58%	53.6
BAROW	38.1%	5.1	-0.61%	62.2

Table 1 – Performance statistics of each model

We provide some usual performance statistics in table 1. Specifically, we report the total return over the period, the Sharpe ratio of expected return (*Sharpe*), the maximum drawdown (*MaxDD*), and the Calmar ratio (*Calmar*, defined as return over the period divided by *MaxDD*). One should be aware that using such short targets to predict induces a very high turnover in the portfolio, resulting in lower net performance after transaction costs are applied.

6 Discussion

Our BAROW algorithm outperformed its baseline in a backtest, yet the question of how Σ impacts the model updates remains. The model forces Σ to converge, causing the model to be less and less able to adapt quickly. That trade-off between adaptability and robustness might not be the right one in volatility periods when one could require more dynamic updates. One way to address it is to schedule reset of the covariance matrix as experimented in [VC11]. In our case it might be beneficial to reset the covariance matrix conditionally on any market event we consider as a trigger for more dynamic updates.

The outperformance of BAROW vs its baseline and the adaptability/robustness trade-off mentioned above could both be further analyzed by increasing the backtest frequency (e.g. 1-min bars). A higher frequency would likely enable shorter training and prediction windows, hereby enhancing the potential overperformance of BAROW vs the linear baseline.

Also, the backtest results described above are a mere sum of single-stock returns and do not include any form of risk management at the overall portfolio level, which could further enhance the risk-adjusted performance of each approach.

Finally, it is worth pointing out that BAROW results are significantly impacted by the method used to neutralize returns that serve as the regression target. Usual standardization methods have exhibited periods of higher correlation of residuals in recent years. These likely stem from stakeholders relying on similar risk models and an ever-spreading use of machine learning type of strategies.

References

- [CDK⁺06] Koby Crammer, Ofer Dekel, Joseph Keshet, Shai Shalev-Shwartz, and Yoram Singer. Online passive-aggressive algorithms. *J. Mach. Learn. Res.*, 7:551–585, 2006.
- [CKD13] Koby Crammer, Alex Kulesza, and Mark Dredze. Adaptive regularization of weight vectors. *Machine Learning*, 91(2):155–187, May 2013.
- [CNL14] Terence Chong, Wing-Kam Ng, and Venus Liew. Revisiting the performance of macd and rsi oscillators. *Journal of Risk and Financial Management*, 7(1):1–12, Feb 2014.

- [Con01] Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1:223–236, 2001.
- [DCP08] Mark Dredze, Koby Crammer, and Fernando Pereira. Confidence-weighted linear classification. In *Proceedings of the 25th international conference on Machine learning - ICML '08*, pages 264–271, Helsinki, Finland, 2008. ACM Press.
- [Hay96] Monson H. Hayes. *Statistical Digital Signal Processing and Modeling*. John Wiley & Sons, Inc., USA, 1st edition, 1996.
- [IJR17] Atsushi Inoue, Lu Jin, and Barbara Rossi. Rolling window selection for out-of-sample forecasting with time-varying parameters. *Journal of Econometrics*, 196(1):55 – 67, 2017.
- [KL15] Zura Kakushadze and Jim Liew. Custom v. standardized risk models. *Risks*, 3(2):112–138, May 2015.
- [PP12] K. B. Petersen and M. S. Pedersen. The matrix cookbook, nov 2012. Version 20121115.
- [SCSG13] Thilo A. Schmitt, Desislava Chetalova, Rudi Schäfer, and Thomas Guhr. Non-stationarity in financial time series: Generic features and tail behavior. *EPL (Europhysics Letters)*, 103(5):58003, sep 2013.
- [SD91] L.L. Scharf and C. Demeure. *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Addison-Wesley series in electrical and computer engineering. Addison-Wesley Publishing Company, 1991.
- [VC11] Nina Vaits and Koby Crammer. Re-adapting the Regularization of Weights for Non-stationary Regression. In *Algorithmic Learning Theory*, pages 114–128, 2011.