

Marketing Material
For Professional Investors Only

Go Beyond Sentiment: Stock Prediction Enhanced with Financial News

RAM AI Systematic Equity

Emmanuel Hauptmann, Valentin Betrix, Nicolas Jamet, Tian Guo,
Louis-Alexandre Piquet

Leave the **herd** behind.

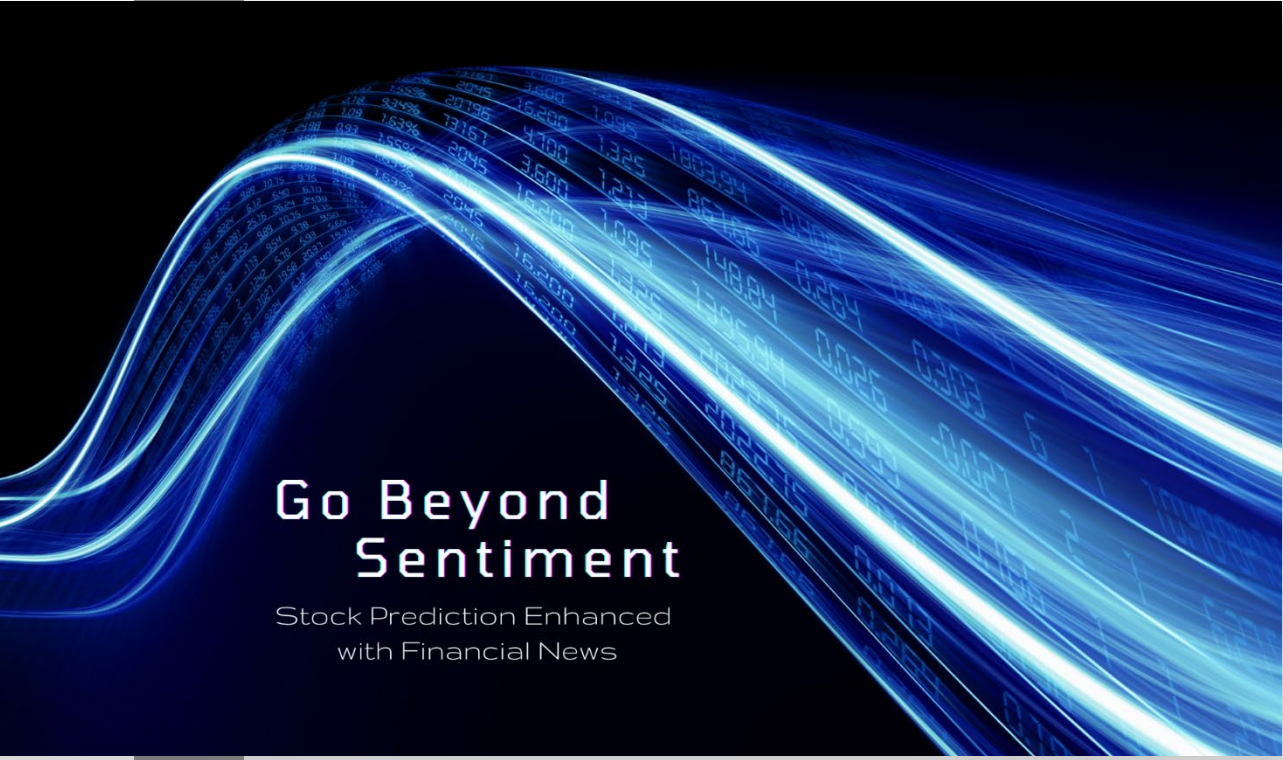
ram^{ai}



Introduction

Extensive research has been focused on extracting quantitative features from financial text (e.g., news, transcripts, etc.) to seek alpha signals such as sentiments, polarity and language complexity, among others.^{5, 6, 11} Though this type of feature, which is usually in the form of scalar values, reflects specific text data characteristics, it neglects the semantic and contextual information needed to evaluate text meanings' relevance or similarity.^{1, 3, 7}

Nowadays, large language models (e.g., BERT, GPT-3, etc.) are among the most impactful natural language processing techniques^{1, 3, 8, 10} and have gained popularity in various applications, such as machine translation, language understanding, etc. These models use specialised neural network architectures (e.g., Transformers, Mixture of Experts, etc.)^{1, 2, 10} and training paradigms to learn the transformation of text to high-dimensional numerical vectors (such as text representations or embeddings). They are trained on a large text corpus (e.g., Wikipedia, Common Crawl, etc.) to capture texts' semantic similarity in the numerical space. Moreover, the language models, which are pre-trained on a large generic corpus, can be further fine-tuned on domain-specific (small) datasets to fit into the corresponding applications.^{4, 8, 9}



Go Beyond Sentiment

Stock Prediction Enhanced
with Financial News

Source: AI generated with the Stable Diffusion Model on <https://beta.dreamstudio.ai/dream>



● Our Methodology

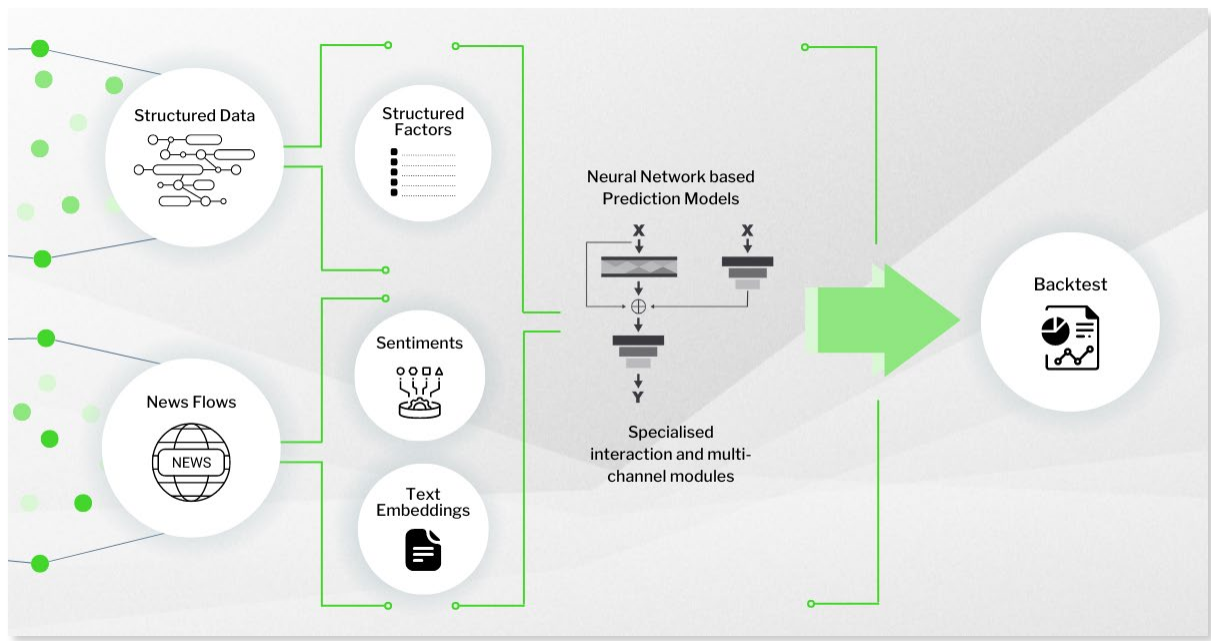


Figure 1: Illustration of the evaluation workflow.

In this paper, we will study the effect of financial news articles’ sentiments and text embeddings on predicting stock returns, as shown in Figure 1. Our in-house developed deep learning framework is able to model interactions between different input features and facilitates the comparison of feature combinations.




Figure 2: News categories and the respective amounts. Source: RAM AI, Factset.



○ The Predictive Power of Stand-Alone News

As shown in Figure 2, our newsflow dataset contains more than 2 million financial news articles from different categories. Therefore, we applied FinBERT, the BERT model³ trained on a financial corpus, and a labelled financial sentiment dataset⁴ to our news data to extract sentiments in multiple lookback time horizons and news' text embeddings. As compared to BERT, FinBERT's output sentiment values and text embeddings have shown to be more accurate and effective at capturing the semantic similarity of financial text.⁴

What is BERT?

BERT (Bidirectional Encoder Representations from Transformers) is a deep learning model that reads entire sequences at once, unlike directional models, which read inputs from left to right or vice-versa. FinBERT was developed by training BERT with financial corpora and has become a cutting-edge tool to analyse financial text sentiment.

In the first group of backtests, we trained two predictive models using sentiments only and sentiments combined with news embeddings, respectively, which allowed us to analyse the predictive power of extracted news features on future stock returns.

○ Combining Structured Data and Unstructured Financial News Data

Apart from unstructured data (i.e., financial news in this paper), many factors built from structured data ("structured factors") such as analyst sentiments, fundamental analysis, price, liquidity and position-based factors have shown predictive power on returns and were widely used in building the predictive models. Figure 3 shows different factor styles and the corresponding number of factors in our structured dataset.

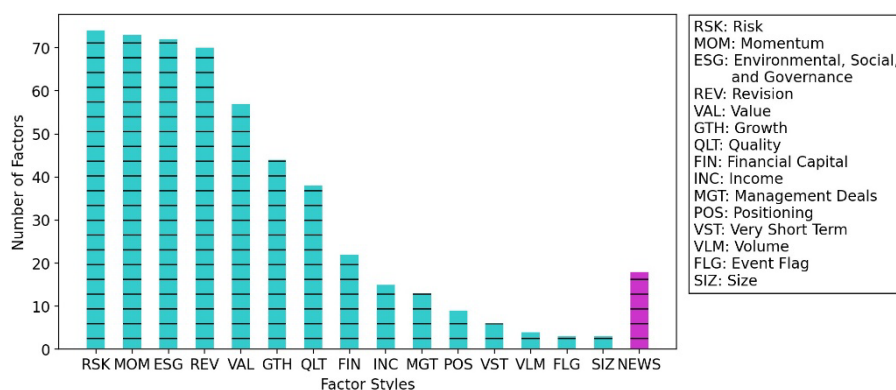


Figure 3: Structured factor styles and the corresponding number of factors. Source: RAM AI, Factset.

We then evaluated the benefit of adding financial news sentiment and embedding features to predictive models on top of the structured factors that already cover most market inefficiencies. For this, in the second group of backtests, we first trained the predictive model solely using our structured factors. Then, we trained two more models on two groups of features; structured factors and sentiments; and structured factors, sentiments, and embeddings. This allowed us to undertake a comparative analysis and determine the advantages of combining unstructured financial news data and large structured datasets in predicting stock returns.



● Backtest Results

In this section, we report backtest results on an investment universe that included around 1000 names from the European market. The results include two parts; the decile returns and the performance of three investing strategies based on return predictions (i.e., equal-weighted long, short, and long/short). The decile returns were derived by collecting the stocks with the predictions falling in a certain decile and aggregating their ex-post monthly returns. The investing strategy was designed to select the stocks based on return predictions. All the results presented below are from the out-of-sample testing period.

○ Stock Prediction with News

Figures 4 and 5 show the results of the predictive models trained separately on news sentiments and the combination of sentiments and text embeddings.

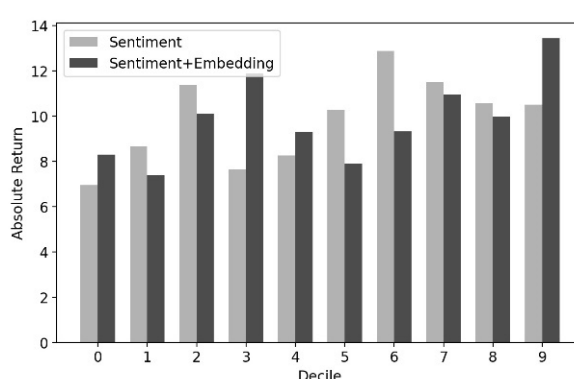


Figure 4: Decile returns of purely news-based predictions. On the x-axis, No. 0 or the bottom decile represents the set of the names with the lowest return predictions. Source: RAM AI, Factset.

Figure 4 shows that adding embeddings noticeably improves the top decile return which consists of the names with the highest return predictions at each rebalancing date. The decile returns present a slightly better monotonic pattern than for the model without embedding.




Figure 5: Performance charts. (a) the long strategy consists of the equal-weighted names in the top decile vs. the equal-weighted universe (b) the short strategy consists of the names in the equal-weighted universe benchmark vs. the bottom decile (c) long/short strategy based on the names in the top and bottom deciles. Source: RAM AI, Factset.

Past performance is not a reliable indicator of future returns.



In Figure 5, the long strategy notably benefits from the use of both sentiments and embeddings, while the short strategy underperforms. A potential explanation of this underperformance is that negative news sentiment is predictive enough and that the relation to negative future returns is quite explicit, so adding embeddings might overly complexify the predictive relation. However, the long/short strategy using both feature sets outperforms.

○ Structured Factors plus News Features

Figures 6 and 7 show the results of the predictive models trained on different combinations of structured factors and news features.

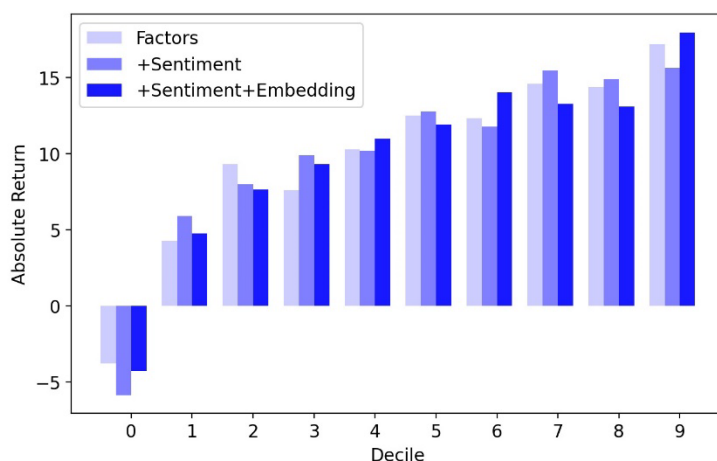


Figure 6: Decile return of news- and factor-based predictions. Source: RAM AI, Factset.

Figure 6 shows that adding sentiments and embeddings enhances the top decile return compared to using only structured factors, though it leads to a slightly lagged performance on the bottom decile. This is in line with the pattern observed in Figure 4.




Figure 7: Performance charts. (a) the long strategy consisted of the equal-weighted names in the top decile vs. the equal-weighted universe benchmark. (b) the short strategy consisted of the names in the bottom decile vs. equal-weighted universe benchmark. (c) the long/short strategy based on the names in the top and bottom deciles. Source: RAM AI, Factset.

Figure 7 (a) demonstrates that adding sentiments and embeddings helps enhance the long strategy's performance. Figure 7 (b) shows that news sentiments provide more complementary predictive power to structured factors and lead to better performance than only using factors. As per Figure 7 (c), the long/short strategy is also improved by the use of both sentiments and embeddings.

Past performance is not a reliable indicator of future returns.



● Interpreting the Model

To give an idea of the news that contribute to high and low return predictions, in Figure 8 we extracted some news examples associated with return predictions falling in the top and bottom decile.

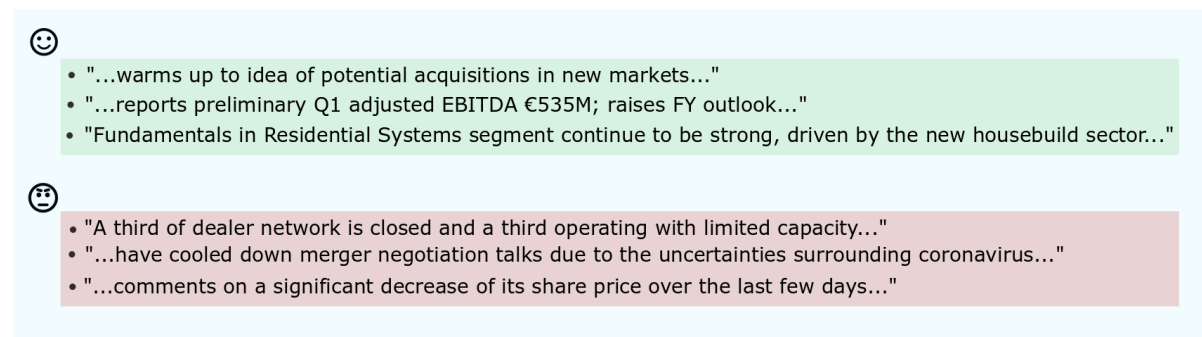


Figure 8: News snippets associated with the stocks falling in the top (green) and bottom (red) decile of the return prediction. Source: RAM AI, Factset.

Figure 9 illustrates the strength of the interaction between factors and news features. The values in this heatmap are aggregated from the model parameters learned during the training of the predictive model. Thus, it is different from conventional correlation metrics calculated on features. Meanwhile, due to the non-linear nature of neural networks, these interaction values reflect the non-linear contribution of feature pairs to the prediction.

For instance, at the bottom line of the heatmap and highlighted in blue, the negative value between NEWS and MOM implies that positive news would lower the contribution of MOM-related factors to the prediction. In contrast, the positive interaction value between NEWS and ESG suggests a reinforced positive effect of news and ESG-related factors on the return prediction.




Figure 9: Visualisation of the factor styles interaction learned during the training of the predictive model using both factors and news. Source: RAM AI, Factset.



● Conclusion

This case study illustrates that text information obtained from financial newsflow has predictive power on stock returns. Sentiment and embedding features extracted from newsflow also seem to be complementary to structured factors that capture price, fundamental, analyst sentiment, liquidity, or positioning-driven inefficiencies in predicting future stock returns.

Our results show that augmenting a large structured input set with news-based sentiment and embedding features can improve the accuracy of a return-prediction model and significantly enhance the returns of a deep-learning-based L/S strategy.




Figure 10: AI generated with the Stable Diffusion Model on <https://beta.dreamstudio.ai/dream>



References

1. Vaswani, Ashish, et al. "Attention is all you need." Advances in neural information processing systems 30. 2017.
2. Guo, Tian, Tao Lin, and Nino Antulov-Fantulin. "Exploring interpretable LSTM neural networks over multi-variable data." International conference on machine learning. PMLR, 2019.
3. Kenton, Jacob Devlin Ming-Wei Chang, and Lee Kristina Toutanova. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." Proceedings of NAACL-HLT. 2019.
4. Huang, Allen, Hui Wang, and Yi Yang. "FinBERT—A Deep Learning Approach to Extracting Textual Information." Available at SSRN 3910214. 2020.
5. Guo, Tian, et al. "ESG2Risk: A deep learning framework from ESG news to stock volatility prediction." arXiv preprint arXiv:2005.02527. 2020.
6. Sawhney, Ramit, et al. "FAST: Financial news and tweet based time aware network for stock trading." Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume. 2021.
7. Sevastjanova, Rita, et al. "Explaining Contextualization in Language Models using Visual Analytics." ACL-IJCNLP 2021-59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, Proceedings of the Conference. 2021.
8. Wei, Jason, et al. "Finetuned Language Models are Zero-Shot Learners." International Conference on Learning Representations. 2021.
9. Chen, Qinkai. "Stock movement prediction with financial news using contextualized embedding from bert." arXiv preprint arXiv:2107.08721. 2021.
10. Du, Nan, et al. "Glam: Efficient scaling of language models with mixture-of-experts." International Conference on Machine Learning. PMLR, 2022.
11. Zhai, Shuang, and Zhu Zhang. "Read the News, Not the Books: Forecasting Firms' Long-term Financial Performance via Deep Text Mining." ACM Transactions on Management Information Systems. 2022.



Disclaimer

Important Information:

The information and analyses contained in this document are based on sources deemed to be reliable. However, RAM Active Investments S.A. cannot guarantee that said information and analyses are up-to-date, accurate or exhaustive, and accepts no liability for any loss or damage that may result from their use. All information and assessments are subject to change without notice.

This document has been drawn up for information purposes only. It is neither an offer nor an invitation to buy or sell the investment products mentioned herein and may not be interpreted as an investment advisory service. It is not intended to be distributed, published or used in a jurisdiction where such distribution, publication or use is forbidden, and is not intended for any person or entity to whom or to which it would be illegal to address such a document. In particular, the investment products are not offered for sale in the United States or its territories and possessions, nor to any US person (citizens or residents of the United States of America). The opinions expressed herein do not take into account each customer's individual situation, objectives or needs. Customers should form their own opinion about any security or financial instrument mentioned in this document. Prior to any transaction, customers should check whether it is suited to their personal situation, and analyse the specific risks incurred, especially financial, legal and tax risks, and consult professional advisers if necessary.

This document is strictly confidential and addressed solely to its intended recipient; its reproduction and distribution are prohibited. It has not been approved by any financial Authority. Issued in Switzerland by RAM Active Investments S.A. ((Rue du Rhône 8 CH-1204 Geneva)) which is authorised and regulated in Switzerland by the Swiss Financial Market Supervisory Authority (FINMA). Issued in the European Union and the EEA by the Management Company RAM Active Investments (Europe) S.A., 51 av. John F. Kennedy L-1855 Luxembourg, Grand Duchy of Luxembourg. Distributed in the United Kingdom to professional investors by RAM AI Advisory, a trading name of Sturgeon Ventures LLP utilised under exclusive license. Sturgeon Ventures LLP (FRN: 452811) is authorised and regulated by the Financial Conduct Authority (FCA).

No part of this document may be copied, stored electronically or transferred in any way, whether manually or electronically, without the prior agreement of RAM Active Investments S.A.

RAM Active Investments SA
Geneva
Rue du Rhône 8
1204 Geneva – SWITZERLAND
Tel : +41 58 726 87 00

RAM Active Investments (Europe) SA
Luxembourg
51, avenue John F. Kennedy
1855 Luxembourg – LUXEMBOURG
Tel : +352 28 56 14 1