



Rome, February 18th, 2025

A NATURAL LANGUAGE PROCESSING TOOLBOX FOR THE NATIONAL BANK OF ROMANIA

Claudia Voicilă, Senior Economist

Note 1: The opinions are those of the author and do not necessarily reflect the views of the National Bank of Romania

Note 2: Part of my dissertation thesis for the Artificial Intelligence Master's program at the Faculty of Mathematics and Informatics, University of Bucharest.

Agenda

Introduction

Chapter 1. English Monetary Policy Statements
Database

Chapter 2. Romanian Financial News Articles
Database

Conclusions

Abstract

- This paper introduces a specially designed toolbox aimed at leveraging Natural Language Processing (NLP) to unlock insights for the National Bank of Romania (NBR), enhancing economists' capacity to process text data such as financial news and press releases, exploring areas of Financial Stability and Central Bank Communication.
- We propose implementing scalable Natural Language Processing methods within the National Bank of Romania's analysis toolkit, focusing on lexicon-based sentiment analysis of daily news and monetary policy decisions. This study seeks to align the institution with the best practices of other central banks and highlights the untapped potential of textual data as a valuable resource in central banking.


Introduction

- The paper is structured in two chapters: the first one addressed monetary policy statements in English from the National Bank of Romania, the European Central Bank, the United States's Federal Reserves and 25 other emerging market countries, with the scope of building a list of positive and negative words specific for central bank text, the starting point of the Romanian Financial Stability Dictionary.
- The second chapter focus on the Romanian news archive database, and we finalize the Financial Stability Dictionary that contains specialized positive and negative words. We augment the method with GenAI scores to measure text-based sentiments in financial texts.

Chapter 1. The English Monetary Policy Statements Database

- The first part focuses on English monetary policy statements starting from the moment NBR adopted inflation targeting. We extended the dataset by including ECB, FED, and 25 other emerging market countries.
- We fine-tuned a Transformer model, FinBERT, on the binary sentiment labels we created for the monetary policy statement based on movements in the monetary policy interest rate of each country (positive sentiment if the interest rate is lowered, negative sentiment if the interest rate is raised).
- After trying unspecialized lexicons, we demonstrated that using specialized Financial Stability Dictionary is crucial for constructing a sentiment index for monetary policy, and we measure the polarity of words using GenAI.


Confusion Matrix of
FinBERT on test sample
from the extended
English database




Classification
report on test
sample (10
epochs)

	precision	recall	F1-score	support
positive	86.52%	89.19%	87.82%	259
negative	92.53%	90.60%	91.56%	383
accuracy			90.03%	642
macro accuracy	89.53%	89.89%	89.69%	642
Weighted average	90.11%	90.03%	90.05%	642

Sentiment Indicator of Monetary Policy Decisions using RO dictionary




Sentiment Indicator of Monetary Policy Decisions using RO dictionary with Scores



- The version of the sentiment index with scores outperforms the simple sentiment index by 2% (78% versus 76%), and the final index that we obtained by mixing the two methods matches 92.5% of the labels. This performance is considered excellent, given that the neutral decisions were built using the last decision rule (negative sentiment persists if last decision was to raise the interest rate, positive sentiment persists if the last decision was to lower interest rate) and not manually evaluated by a human.
- To assess the performance of our NLP analysis, we examined the correlation between the sentiment index of monetary policy decisions and the most relevant economic variable: the inflation rate, since controlling inflation is the primary goal of monetary policy. We found an expected negative correlation between our computed sentiment and the all-items HICP (Harmonized Index of Consumer Prices) of over -40%, while the correlation between the monetary policy interest rate and HICP was around -60%.

Sentiment Index of Monetary Policy compared to Binary Labels



Chapter 2. The Romanian Financial News Articles Database

- The second part addresses the NBR's archive of daily financial news, for which we conduct topic modeling using LDA to track key trends in the news, and propose a tailor-made, Romanian Financial Stability Dictionary to measure sentiments.
- We computed polarity intensity scores by prompting as a financial stability expert a Gen AI agent.
- In absence of labels, we assess the performance by comparing our method with K-Means++ clustering applied to embeddings obtained with multilingual BERT Transformer and reduced with PCA, and by demonstrating correlations with macroeconomic and market variables as well as interactions with monetary policy decisions.

Interactive LDA analysis, Example for year 2023


Selected Topic: **1** Previous Topic Next Topic Clear Topic

Slide to adjust relevance metric:⁽²⁾
 $\lambda = 1$ 0.0 0.2 0.4 0.6 0.8 1.0

Intertopic Distance Map (via multidimensional scaling)




Top-30 Most Relevant Terms for Topic 1 (38.6% of tokens)



Overall term frequency
 Estimated term frequency within the selected topic

1. saliency(term w) = frequency(w) * [sum_t p(t | w) * log(p(t | w)/p(t))]] for topics t; see Chuang et. al (2012)
2. relevance(term w | topic t) = $\lambda * p(w | t) + (1 - \lambda) * p(w | t)/p(w)$; see Sievert & Shirley (2014)


Heatmap of the Importance of the Main Five Topics Over Time



Word Cloud analysis, Example for year 2023




Daily Sentiment of News using Romanian Financial Stability Dictionary



$$\text{Sentiment Index} = \frac{P(t) - N(t)}{P(t) + N(t)}$$

where P number of positive lemmas and N number of negative lemmas

Daily Sentiment of News using Romanian Financial Stability Dictionary augmented with scores




$$\text{Sentiment Score Index} = \frac{\sum_{i=1}^{P(t)} S_p(i) + \sum_{j=1}^{N(t)} S_n(j)}{\sum_{i=1}^{P(t)} S_p(i) + \sum_{j=1}^{N(t)} |S_n(j)|}$$


where S_p scores of positive lemmas and S_n scores of negative lemmas

- To assess the performance of our daily sentiment indicator, we developed an unsupervised K-Means++ analysis according to Kanungo (2002) on embeddings using the BERT multilingual model by Devlin (2018), after reducing their size using PCA. The results show a 93% overlap between the clusters and the binary labels from the simple index, while a slightly higher 96% overlap for the index using scores.
- We compared the daily sentiment of financial news with the moments of monetary policy decisions and showed that there is strong evidence that the press respond to monetary policy decisions and that the communication of the central bank is efficient in influencing the markets.
- We also found a clear correlation with macroeconomic variables: BET, the stock market index, yields of short-term bonds, as well as interbank interest rates ROBOR3M.

Daily Sentiment of Financial News and Monetary Policy Decisions




Daily Sentiment of News and Stock Market Index



Quarterly Sentiment of News and Percentage Change in Stock Market Index



Daily Sentiment of News, Interbank Interest rates and Sovereign Yield Rates



Conclusions

- This paper proposes stepping aside from classical numerical, tabular data and investing more time and effort in the use of nontraditional data, not just time series, thus implementing natural language processing methods at a large scale for some types of text that exist in a central bank. The primary goal is to provide the NBR with a guideline for pre-processing Romanian text, building NLP tools, and benchmarking sentiment analysis.
- We can only hope that by following this NLP Toolbox for the National Bank of Romania, the list of potential use cases for NLP will continue to be extended at the NBR and that economists will commit to a long-term adoption of text data in their research, as there are many other text sources available.

Thank you!
claudia.voicila@bnro.ro



References

- Abodayeh, A., Hejazi, R., Najjar, W., Shihadeh, L., & Latif, R. (2023, March). Web Scraping for Data Analytics: A BeautifulSoup Implementation. In *2023 Sixth International Conference of Women in Data Science at Prince Sultan University (WiDS PSU)* (pp. 65-69). IEEE.
- Aprigliano, V., Emiliozzi, S., Guaitoli, G., Luciani, A., Marcucci, J., & Monteforte, L. (2023). The power of text-based indicators in forecasting Italian economic activity. *International Journal of Forecasting*, 39(2), 791-808.
- Araci, D. (2019). FinBERT: A Pretrained Language Model for Financial Communications. arXiv preprint arXiv:1908.10063.
- Araujo, D., Bruno, G., Marcucci, J., Schmidt, R., & Tissot, B. (2023). Data science in central banking: applications and tools. *IFC Bulletin*, (59).
- Araujo, D., Bruno, G., Marcucci, J., Schmidt, R., & Tissot, B. (2023). Machine learning applications in central banking. *Journal of AI, Robotics & Workplace Automation*, 2(3), 271-293.
- Araujo, Douglas KG. 2023. “Gingado: A Machine Learning Library Focused on Economics and Finance.” Working Paper 1122. BIS Working Paper. Bank for International Settlements.
- Astuti, V., Bruno, G., Marchetti, S., & Marcucci, J. (2023). *News and banks' equities: do words have predictive power?* (Vol. 59). Bank for International Settlements.
- Baccianella, S., Esuli, A., & Sebastiani, F. (2010, May). Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining. In *Lrec* (Vol. 10, No. 2010, pp. 2200-2204).
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan), 993-1022.
- Chong, E., & Ho, S. J. (2023). *Measuring text-based sentiments from monetary policy statements—a Malaysian case study using natural language processing* (Vol. 59). Bank for International Settlements.
- .

- Chong, E., Ho, C. C., Ong, Z. F., & Ong, H. H. (2022). *Using news sentiment for economic forecasting: a Malaysian case study* (Vol. 57). Bank for International Settlements
- Chuang, J., Manning, C. D., & Heer, J. (2012, May). Termite: Visualization techniques for assessing textual topic models. In *Proceedings of the international working conference on advanced visual interfaces* (pp. 74-77).
- Correa, R., Garud, K., Londono-Yarce, J. M., & Mislant, N. (2017). Constructing a dictionary for financial stability.
- De Bandt, O., Jean-Charles, B., Julien, D., Alexandre, D., De Gaye, A., & Pierre-Antoine, R. (2023). *Utiliser la presse pour construire un nouvel indicateur de perception d'inflation en France*. Banque de France.
- Denes, J., Lestrade, A., & Richardet, L. (2022). *Using twitter data to measure inflation perception* (Vol. 57). Bank for International Settlements.
- Devlin, J. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Doerr, S., Gambacorta, L., & Garralda, J. M. S. (2021). Big data and machine learning in central banking. *BIS Working Papers*, (930).
- Dumitrescu, S. D., Avram, A. M., Morogan, L., & Toma, S. A. (2018, June). Rowordnet—a python api for the romanian wordnet. In *2018 10th International Conference on Electronics, Computers and Artificial Intelligence (ECAI)* (pp. 1-6). IEEE.
- Edison, H., & Carcel, H. (2021). Text data analysis using Latent Dirichlet Allocation: an application to FOMC transcripts. *Applied Economics Letters*, 28(1), 38-42.
- Evdokimova, T., Nagy-Mohacsi, P., Ponomarenko, O., & Ribakova, E. (2023). Central banks and policy communication: How emerging markets have outperformed the Fed and ECB. *Peterson Institute for International Economics Working Paper*, (23-10).
- Hu, M., & Liu, B. (2004, August). Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 168-177).
- Hutto, C., & Gilbert, E. (2014, May). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media* (Vol. 8, No. 1, pp. 216-225).

Istrefi, K., Odendahl, F., & Sestieri, G. (2022). ECB communication and its impact on financial markets.

Istrefi, K., Odendahl, F., & Sestieri, G. (2023). Fed communication on financial stability concerns and monetary policy decisions: Revelations from speeches. *Journal of Banking & Finance*, 151, 106820.

Jabbar, M. A., Wibisono, O., Widjanarti, A., & Zulen, A. A. (2023). Machine learning for measuring central bank policy credibility and communication from news. *IFC Bulletins chapters*, 59.

Jolliffe, I. T. (1986). Principal component analysis. *Springer Series in Statistics*. Springer-Verlag.

Kanungo, T., Mount, D. M., Netanyahu, N. S., Piatko, C. D., Silverman, R., & Wu, A. Y. (2002). An efficient k-means clustering algorithm: Analysis and implementation. *IEEE transactions on pattern analysis and machine intelligence*, 24(7), 881-892.

Lee, Y., & Seo, B. (2023). Extracting economic sentiment from news articles: the case of Korea. *IFC Bulletins chapters*, 59.

Loper, E., & Bird, S. (2002). Nltk: The natural language toolkit. *arXiv preprint cs/0205028*.

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of finance*, 66(1), 35-65.

Mabey, B. (2018). pyLDAvis documentation.

Qi, P., Zhang, Y., Zhang, Y., Bolton, J., & Manning, C. D. (2020). Stanza: A Python natural language processing toolkit for many human languages. *arXiv preprint arXiv:2003.07082*.

Richardson, L. (2007). Beautiful soup documentation.

Schmitt, X., Kubler, S., Robert, J., Papadakis, M., & LeTraon, Y. (2019, October). A replicable comparison study of NER software: StanfordNLP, NLTK, OpenNLP, SpaCy, Gate. In *2019 sixth international conference on social networks analysis, management and security (SNAMS)* (pp. 338-343). IEEE.

Vasilev, Y. (2020). *Natural language processing with Python and spaCy: A practical introduction*. No Starch Press.