

PADME-SoSci: A Platform for Analytics and Distributed Machine Learning for the Social Sciences

Zeyd Boukhers^{1,2} Arnim Bleier³ Yeliz Ucer Yediel^{1,5} Mio Hienstorfer-Heitmänn³
 Mehrshad Jaberansary^{2,4} Adamantios Koumpis^{2,4} Oya Beyan^{1,2,4}

¹Fraunhofer Institute for Applied Information Technology -FIT-, Sankt Augustin, Germany

²University Hospital Cologne, Cologne, Germany

³GESIS – Leibniz Institute for the Social Sciences, Cologne, Germany

⁴University of Cologne, Cologne, Germany

⁵RWTH Aachen University, Aachen, Germany
 Germany

{zeyd.boukhers,yeliz.ucer.yediel}@fit.fraunhofer.de

{arnim.bleier,mio.hienstorfer-heimann}@gesis.org

{mehrshad.jaberansary,adamantios.koumpis,oya.beyan}@uk-koeln.de

ABSTRACT

Data privacy and ownership are significant in social data science, raising legal and ethical concerns. Sharing and analyzing data is difficult when different parties own different parts of it. An approach to this challenge is to apply de-identification or anonymization techniques to the data before collecting it for analysis. However, this can reduce data utility and increase the risk of re-identification. To address these limitations, we present **PADME**, a distributed analytics tool that federates model implementation and training. **PADME** uses a federated approach where the model is implemented and deployed by all parties and visits each data location incrementally for training. This enables the analysis of data across locations while still allowing the model to be trained as if all data were in a single location. Training the model on data in its original location preserves data ownership. Furthermore, the results are not provided until the analysis is completed on all data locations to ensure privacy and avoid bias in the results.

CCS CONCEPTS

• **Applied computing** → *Sociology*; Digital libraries and archives; • **Computing methodologies** → **Distributed computing methodologies**; **Learning paradigms**.

KEYWORDS

distributed analytics, data privacy, social sciences, data science

Reference Format:

Boukhers et al. 2023. PADME-SoSci: A Platform for Analytics and Distributed Machine Learning for the Social Sciences

1 INTRODUCTION

Data privacy and ownership are crucial in social data science as the data often includes sensitive personal information. For example, it is expected that political survey data will typically be gathered by various parties across different groups, but sharing may not

occur due to privacy considerations. As a result, each party independently analyzes their own data and only shares the outcome. This approach can be limited as the aggregated outcome may not accurately reflect the entire data. An alternative solution is to mask personal information, but this can also have limitations as masked attributes may be important for analysis.

To address this issue, we present **PADME**¹ [4], a distributed analytics tool that enables the training of models on large-scale datasets without the need to centralize the data. **PADME** trains the analytic model incrementally by visiting each data station one by one. This ensures that the data remains decentralized and protected from unauthorized access. Additionally, **PADME** only produces output after visiting all data stations, thereby avoiding the sharing of derived knowledge that could potentially reveal information about the data within a specific station. This provides an added layer of security and ensures that the data remains confidential. To further guarantee the security and privacy of the data, it is recommended that the model is implemented in a federated manner. This ensures that the data is protected by multiple parties and that any potential vulnerabilities are identified and addressed.

2 PADME-SOSCI

The distributed analytics tool depicted in Figure 1 has two primary components: the *Service Center* and the *Data Stations*. The *Service Center* acts as the central hub for the entire analytics process. It oversees the development and approval of analytic models, manages the authentication and authorization of scientists and researchers, and ensures the proper execution of analytics tasks. The process of developing and approving the analytic models involves the collaboration and agreement of both the researcher and the data owners. This consensus mechanism ensures that the parameters and architecture of the global model are agreed upon by all parties. The *Service Center* also ensures the security of the analytics process by securely packaging the analytics model into containers and managing the execution routes through authentication. This minimizes the risk of unauthorized access or misuse of the analytic models. Additionally, the *Service Center* employs encryption techniques to protect the privacy and ownership of data throughout the entire

¹<https://padme-analytics.de/>

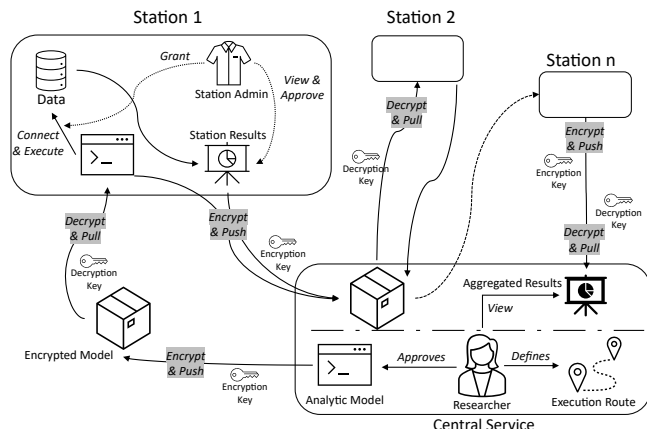


Figure 1: An overview of the distributed analytic tool with n data stations

process. The encrypted containers carrying the analytics models are then transferred to the next target station.

The *Data Station* operates independently of the *Service Center*, but it remains closely connected to it through a secure communication channel. The station pulls the encrypted container from the *Service Center* after authenticating itself. The station admin then provides the necessary connection information to allow the decrypted container to execute the analytics model. The station’s unique key is a crucial security feature that ensures that only the designated station can execute the analytic model within the decrypted container. After the analytics task is completed, the Station Admin reviews the results and, upon approval, packages them into a secure container using the station’s specific key. This process is repeated for each station in the route until the final results are available in an encrypted container. The authenticated scientist can then retrieve, decrypt and view the aggregated results.

2.1 Prerequisites

To effectively utilize the distributed analytic tool, compliance with the following prerequisites is crucial:

- **Data standardization:** All used data must be in a standardized format that the model is designed to work with.
- **Distributed model:** The analytic model must be capable of being distributed. For example, the Latent Dirichlet Allocation (LDA) model requires the entire vocabulary of the corpus before running, making it unsuitable for direct use with PADME.
- **Computational resources:** Every data station needs to have adequate computational resources to execute the analytic model.

3 USE CASES

3.1 Sentiment Analysis

Online reproducibility services, such as mybinder.org, have become popular in the Computational Social Science community [1]. These services allow researchers to share complex computational analysis pipelines in the form of notebooks that can be easily executed in a browser without the need for additional software installation. However, the pipelines are currently limited to public data. In this

paper, we demonstrate the use of PADME for a sentiment analysis test within the following prototypical workflow²:

- (1) Schema data is created and publicly shared for a sensitive dataset that can not be shared.
- (2) An interested researcher that would like to access the sensitive dataset develops an analysis using the publicly available schema data.
- (3) The researcher submits the analysis to PADME to execute the analysis of the private data.
- (4) An exit control on the analysis results is performed, and the results of the analysis are sent back to the researcher if the exit control is passed.

The publicly available part of the German Federal Election 2017 Twitter dataset[3] (DBK: ZA6926) is used for this use case.

3.2 Supervised Author Name Disambiguation

In this demonstration, we are showcasing the capability of PADME by distributing the Author Name Disambiguation (AND) task to be carried out at two separate stations, each of which holds a portion of the data. The data used for AND is typically open, but the purpose of this demonstration is to highlight the versatility of PADME in handling various types of data and analytical models. In particular, we are utilizing a supervised neural network model [2] that has been trained on a preprocessed DBLP dataset³. The DBLP dataset has been split and distributed between the two stations for this demonstration.

ACKNOWLEDGEMENT

This work was partially supported by the project NFDI4DS⁴ with grant number 460234259.

REFERENCES

- [1] Arnim Bleier, Kenan Erdogan, Christian Kahmann, and Lisa Posch. 2022. GESIS Notebooks: Online Reproducible Computational Analysis for the Social Sciences. <https://doi.org/10.5281/zenodo.7362188>
- [2] Zeyd Boukhers and Nagaraj Bahubali Asundi. 2022. Whois? Deep Author Name Disambiguation Using Bibliographic Data. In *26th International Conference on Theory and Practice of Digital Libraries, TPDL 2022*. Springer, 201–215.
- [3] Sebastian Stier, Arnim Bleier, Malte Bonart, Fabian Mörshheim, Mahdi Bohlouli, Margarita Nizhegorodov, Lisa Posch, Jürgen Maier, Tobias Rothmund, and Steffen Staab. 2018. Systematically monitoring social media: The case of the German federal election 2017. *arXiv preprint arXiv:1804.02888* (2018).
- [4] Sascha Welten, Yongli Mou, Laurenz Neumann, Mehrshad Jaberansary, Yeliz Yediel Ucer, Toralf Kirsten, Stefan Decker, and Oya Beyan. 2022. A privacy-preserving distributed analytics platform for health care data. *Methods of information in medicine* 61, S 01 (2022), e1–e11.

²Originally discussed on GitHub https://github.com/gesiscss/btw17_sample_scripts/issues/4

³<https://doi.org/10.5281/zenodo.7506562>

⁴<https://gepris.dfg.de/gepris/projekt/460234259>