

Federated Learning in Breast Cancer Diagnosis

Aprendizado Federado no Diagnóstico do Câncer de Mama

Gleidson V. G. Barbosa^{1*}, Leonardo G. Ferreira Rodrigues^{1*}, Larissa F. Rodrigues Moreira^{1,2*}, André Ricardo Backes^{2*}

Abstract: Breast cancer, which accounts for 29.7% of cancer cases among women and has a mortality rate of approximately 16.5%, it is the most common cancer in Brazil. To enhance early detection and minimize the occurrence of advanced cases, pathologists are increasingly utilizing deep learning techniques to analyze histopathological images. This study investigates and compares the performance of two convolutional neural networks, AlexNet and ResNet50, in federated and centralized training settings. By employing the BreakHis dataset, the research explores the effectiveness of federated learning in securely distributing model training across multiple clients and contrasts it with traditional centralized training approaches. This study investigates whether federated learning can achieve performance close to that of centralized approaches while ensuring privacy for breast cancer histopathology image classification. It is hypothesized that federated learning can provide comparable results to centralized approaches while ensuring the confidentiality of patient data. Our results confirm that federated learning performs similarly to centralized learning with the added benefits of maintaining privacy and supporting collaboration among multiple clients. The findings aim to advance the state of the art in breast cancer detection, addressing both the challenges of privacy and data distribution and the impact of different training methodologies on model performance.

Keywords: Breast Cancer — Federated Learning — Histopathology — Distributed Learning

Resumo: O câncer de mama, que representa 29,7% dos casos de câncer entre mulheres e tem uma taxa de mortalidade de aproximadamente 16,5%, é o câncer mais comum no Brasil. Para melhorar a detecção precoce e minimizar a ocorrência de casos avançados, os patologistas estão cada vez mais utilizando técnicas de aprendizado profundo para analisar imagens histopatológicas. Este estudo investiga e compara o desempenho de duas redes neurais convolucionais, AlexNet e ResNet50, em configurações de treinamento federado e centralizado. Ao empregar o conjunto de dados BreakHis, a pesquisa explora a eficácia do aprendizado federado na distribuição segura do treinamento do modelo entre vários clientes e o contrasta com as abordagens tradicionais de treinamento centralizado. Este estudo investiga se o aprendizado federado pode alcançar uma performance próxima à das abordagens centralizadas, garantindo a privacidade na classificação de imagens histopatológicas de câncer de mama. A hipótese é que o aprendizado federado pode fornecer resultados comparáveis às abordagens centralizadas, enquanto assegura a confidencialidade dos dados dos pacientes. Nossos resultados confirmam que o aprendizado federado apresenta desempenho semelhante ao do aprendizado centralizado, com os benefícios adicionais de manter a privacidade e suportar a colaboração entre vários clientes. As descobertas visam avançar o estado da arte na detecção do câncer de mama, abordando os desafios de privacidade e distribuição de dados e o impacto de diferentes metodologias de treinamento no desempenho do modelo.

Palavras-Chave: Câncer de Mama — Aprendizado federado — Histopatologia — Aprendizado Distribuído

¹ School of Computer Science (FACOM), Federal University of Uberlândia (UFU), Brazil

² Department of Computing (DC), Federal University of São Carlos (UFSCar), Brazil

*Corresponding authors: {gleidson.barbosa, leonardo.g.rodrigues, larissarodrigues}@ufu.br, arbackes@yahoo.com.br

DOI: <http://dx.doi.org/10.22456/2175-2745.143480> • Received: 23/10/2024 • Accepted: 23/12/2024

CC BY-NC-ND 4.0 - This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

1. Introduction

Breast cancer, characterized by the uncontrolled proliferation of abnormal cells within the breast tissue, can develop into

invasive tumors. Although it is relatively rare among men, accounting for approximately 1% of cases, breast cancer remains a critical health issue. In 2020, it constituted 29.7% of

all new cancer diagnoses and was associated with a high mortality rate of 16.5% [1]. Despite technological advancements and ongoing efforts in early detection, challenges persist, particularly within Brazil's public health system (SUS), where conventional methods often lead to prolonged waiting periods. This delay allows for the continuous proliferation of abnormal cells, complicating early diagnosis and treatment [2].

The integration of computational technologies into breast cancer diagnosis has gained prominence, with innovations in image analysis and Artificial Intelligence (AI) playing a crucial role in early detection and accurate lesion assessment. Techniques such as digital mammography, histopathology, and breast tomosynthesis, when combined with machine learning algorithms, have markedly enhanced diagnostic accuracy. These advancements reduce variability among different specialists and improve the detection rates of early-stage cancers. Such innovations are vital for the progression of personalized medicine, offering more precise and less invasive treatment options [3, 4].

This study proposes the evaluation of distributed learning methodologies using the BreakHis dataset to advance computational techniques for breast cancer prediction and diagnosis. Specifically, it explores the Federated Learning approach where local devices train separate models on local data, sharing updates with a central server to improve global accuracy without sharing personal data, which facilitates the integration of processing power from multiple devices while preserving patient data confidentiality [5, 6]. By leveraging Federated Learning, this approach aims to develop generic models from aggregated training weights, enabling collaborative efforts between healthcare institutions without compromising data privacy. This strategy is anticipated to enhance both internal predictive accuracy and diagnostic capabilities.

The remaining of this paper is organized as follows: Section 2 provides an overview of related works in the literature that are similar to the proposed approach. Section 3 presents the proposed method for evaluating CNNs. In Section 4, we present and discuss the results achieved in this work. Finally, Section 5 concludes the discussion and offers some final remarks.

2. Related Work

Several studies explore AI methods to enhance breast cancer diagnosis, with an emphasis on Convolutional Neural Networks (CNNs) for analyzing histopathological images. In 2019, Spanhol [7] aimed to distinguish between benign and malignant tumors in histopathological images with high accuracy using Deep Learning and Multiple Instance Learning. They achieved an accuracy of 92.1% with centralized training, although they did not aim to surpass current works in precision.

Vo, Nguyen, and Lee [8] used CNNs in combination with gradient-boosting tree algorithms to extract useful information and improve the performance of the CNNs. Moreover, a majority voting strategy was employed to combine the boosting

trees, resulting in the best classification performance among them. Using the BreakHis dataset, the results demonstrate the efficiency of the combined model, which achieved the best accuracy for all cases tested.

In 2022, there was growing interest in the use of FL to predict breast cancer due to its ability to preserve data privacy. In [9], the authors applied Federated Learning (FL) in the detection and classification of Invasive Carcinoma of No Special Type (IC-NST), one of the most prevalent types of breast cancer. The proposed approach combines features extracted by residual neural networks and Gabor filters, resulting in a feature fusion that was used with a custom classifier. The experiments showed promising results, with accuracy exceeding 80%, with the best result being 86.57% accuracy, indicating the viability of FL for applications requiring privacy and collaboration, especially in medical contexts.

In [10] the potential of FL in applications with medical images and distributed data is highlighted, using FL with MLP and CNN models, combined with pre-trained models (MobileNet, Densenet121, Xception, ResNet50) to improve breast cancer classification while preserving patient privacy, in contrast to centralized learning (CL). Using datasets such as the Digital Database for Screening Mammography (DDSM) and the Curated Breast Imaging Subset of DDSM (CBIS-DDSM), the results indicate that the FedAvg-CNN + MobileNet model not only protects privacy but also improves classification performance, surpassing previous models with accuracy close to 98% and a recall capacity of 100%.

The authors of [11] present a series of experiments to define the best parameters for training models using histopathology image datasets. They employed various CNNs, such as AlexNet, EfficientNet, SqueezeNet, ResNet50-50, VGG-16, and ShuffleNet, generating a guide where one can observe the best combination of parameters to start training using one of these networks on the BreakHis and Biglycan datasets.

This paper presents an opportunity for advancement through the incorporation of Federated Learning. Table 1 summarizes the related works and contrasts with the proposed work.

Table 1. State-of-the-Art Approaches

Approach	Environment	Dataset	Accuracy
Spanhol et al. (2018) [7]	Centralized	BreakHis	92.1%
Vo et al. (2019) [8]	Centralized	BreakHis	96.9%
Agbley et al. (2022) [9]	Distributed	BHI and BreakHis	86.57%
Tan et al. (2023) [10]	Distributed	CBIS-DDSM	98.00%
Barbosa et al. (2024) [11]	Centralized	BreakHis and Biglycan	98.86%
Our approach	Distributed	BreakHis	97.91%

3. Proposed approach

This study proposes a method utilizing the BreakHis dataset [12] and applying FL technique with CNNs¹. The primary goal is to explore the potential of FL by adapting models traditionally

¹Source code available in: <https://github.com/Plufty/Estudos-Federated-Learning>

used for centralized training to a federated environment using the Flower framework [13] as illustrated in Figure 1. We conducted the experiments on an Acer Predator Helios 300 machine, equipped with an Nvidia GeForce RTX 2070 with Max-Q Design, Intel Core i7-10750H CPU @2.60GHz, and 32GB of RAM.




Figure 1. Proposed Method: Federated learning with AlexNet and ResNet50 on the BreakHis dataset, to compare federated and centralized training results.

To implement the method we used two models, AlexNet [14] and ResNet50 [15], which were trained on three clients. The Flower framework’s simulation function was employed to simulate the clients and the server, enabling the application of Federated Learning.

Step One: This step involved acquiring and preprocessing the BreakHis dataset [12] this dataset was chosen for its significance in breast cancer diagnosis, comprising a substantial collection of histopathological images. This dataset enabled an evaluation of the models’ robustness within a practical diagnostic context. The preprocessing step included dividing the dataset into two primary classes: *Benign* and *Malignant*. Subsequently, the dataset was further subdivided for the three clients, where the images in each class were randomized and split, with 80% allocated for training and 20% for validation. Additionally, we prepared a centralized dataset using the same split to conduct centralized training on a single machine. The BreakHis dataset contains 7,909 microscopic images of breast tumor tissue, divided between benign tumors (2,480 images) and malignant tumors (5,429 images), with dimensions of 700×460 pixels.

Step Two: Two Convolutional Neural Networks (CNNs), AlexNet [14] and ResNet50 [15], were trained using the SGD optimizer. The optimizer’s settings included a learning rate of 0.001, a batch size of 16, a momentum of 0.9, and a total of

50 epochs. These parameter choices were based on previously obtained results [11], identified as the most effective configurations for the dataset in question. AlexNet and ResNet50 were selected due to their prominence as effective architectures in medical image classification, providing a favorable balance between simplicity and performance, particularly in situations with limited data

Step Three: We independently trained the locally trained models on each client using their respective subsets of the dataset. The training process for each client involved the application of the chosen CNN architectures and the previously provided parameters.

Step Four: After completing the local training, the trained model weights from each client were sent to the central server. This step involved the secure transmission of model parameters without sharing the actual data, ensuring privacy and compliance with Federated Learning principles.

Step Five: Upon receiving the weights from all clients, the server calculated the new global model weights. This calculation was performed by averaging the locally trained weights, effectively combining the knowledge gained from each client into a single updated model.

Step Six: In this step, the server-distributed global model weights were sent back to each client. Each client then updated their local model with these new weights, preparing for the next round of local training.

Step Seven: This step involved centralized training for both models using the same configuration as the federated training. The key difference was that we calculated the epochs to simulate the total training sum across all clients for comparison purposes, considering both time and the impact of extensive training in a centralized environment.

Step Eight: In the final step, the results from both federated and centralized training were collected. Graphs were generated for each round of federated learning, as well as for centralized learning for each model. These results will be analyzed in the following Section 4.

4. Results and Discussions

Experiments were conducted with three different clients. We trained the network models for 50 epochs in each round. We performed three rounds of federated training, as this number observed sufficient to achieve noticeable improvements, beyond which no significant gains were observed[16]. At the end, we conducted a centralized training with the same total number of epochs ($\text{rounds} \times \text{clients} \times \text{epochs per client}$) for comparison purposes.

In the subsequent figures, the X-axis represents the epochs, while the Y-axis denotes the loss (left graphs) and accuracy (right graphs) unless otherwise specified, these figures illustrate the loss and accuracy results of the training over the epochs for the validation and training phases.

4.1 AlexNet

4.1.1 Federated Learning

Round 1: The loss and accuracy trends for this round demonstrate clear convergence. For Client 1, the loss steadily decreases over the epochs, reaching a relatively low value, while accuracy progressively increases. This indicates that the model is learning effectively and adjusting its parameters to improve classification accuracy.

The model demonstrates solid performance with accuracies ranging from 93.93% to 94.88%. Precision and recall metrics for the benign class are slightly lower, suggesting that the model may have difficulties in accurately identifying benign cases. F1-scores also reflect this trend, with relatively lower values for the benign class.

Round 2: The training process for Round 2 across Clients 1, 2, and 3, revealed good trends in both loss and accuracy over the 50 epochs. Each client demonstrated an effective behavior pattern, indicating that the model was able to learn and adapt well during the training phase.

A notable improvement in performance metrics is observed, with accuracies ranging from 94.50% to 96.58%. Precision and recall for both classes improve, especially for the malignant class, indicating that the model is adjusting its predictions more effectively with federated training. The F1-score also increases, showing a better balance between precision and recall.

Round 3: The results for Round 3 of Clients 1, 2, and 3 are illustrated in Figures 2, 3, and 4 and show an improvement trend over the training, as evidenced by the observed patterns in loss and accuracy. The data indicates that the model's performance improves consistently across these clients during the training process.

- Client 1:** This client demonstrated consistency in training accuracy, maintaining levels above 99% (see Figure 2). Training loss exhibited a steady decrease, from 0.0799 in the first epoch to 0.0001 by the end. However, validation loss showed some fluctuation, ranging from 0.1333 to 0.1674, which suggests occasional challenges in generalizing to the validation set. Despite these fluctuations, validation accuracy remained high, varying between 92.79% and 96.58%, reflecting overall strong model performance. This indicates that while the model is improving internally, it faces some difficulties in consistently adapting to the validation set.
- Client 2:** This client demonstrated shows improvement in training accuracy, reaching 100% by the end of the 50 epochs (see Figure 3). Training loss showed a consistent decline from 0.0829 in the first epoch to 0.0004 by the final epoch. Validation loss initially increased but stabilized around 0.15 in the later epochs, with a peak at 0.3070 in the 5th epoch. Validation accuracy varied between 95.07% and 96.83%, reflecting strong overall performance despite minor fluctuations. This suggests that the model successfully improved its performance




Figure 2. AlexNet Loss and Accuracy C1 R3

during training, while also handling validation data with high accuracy.




Figure 3. AlexNet Loss and Accuracy C2 R3

- Client 3:** This client showed a good performance, with training accuracy reaching 100% by epoch 11 (see Figure 4). Training loss decreased steadily from 0.0985 in the first epoch to 0.0003 by the final epoch. Validation loss showed some fluctuation, peaking at 0.3180 in the 9th epoch, but ultimately stabilized around 0.2339. Validation accuracy varied between 92.22% and 95.26%, indicating robust performance with occasional variability. This reflects that the model achieved high accuracy on both training and validation data, despite some challenges in generalization.




Figure 4. AlexNet Loss and Accuracy C3 R3

The model achieves its best performance metrics, with accuracies ranging from 94.88% to 96.58%. The F1-score

for both classes reaches good levels, reflecting the generalization ability after several rounds of training. The model shows good precision and recall for the malignant class while also improving precision for the benign class. The evolution of the AlexNet model's performance across different training rounds is illustrated in Figure 5, where the X-axis represents the rounds, and the Y-axis represents the loss (left graph) and accuracy (right graph). The figure displays the model's accuracy and loss trends throughout the training process, illustrating key changes and overall performance improvements.




Figure 5. AlexNet Loss and Accuracy across the Rounds

4.1.2 Centralized Learning

Centralized training demonstrates even stronger performance as shown in Figure 6, with an accuracy of 96.96%. This performance highlights the centralized model's ability to process and classify data effectively. However, federated learning offers advantages by enabling training on distributed data without centralization, which can be beneficial in scenarios involving sensitive or distributed data. The combination of federated and centralized results provides a comprehensive view of the AlexNet model's potential in different training configurations.




Figure 6. AlexNet Loss and Accuracy (Centralized Learning)

4.2 ResNet50

4.2.1 Federated Learning

Round 1: The observed metrics from this round indicate all clients demonstrated a consistent reduction in loss across the 50 epochs. This trend underscores the effectiveness of the model's training process in achieving steady improvement

throughout the evaluation. The model demonstrates good initial performance, with accuracy ranging from 94.88% to 95.83%. Precision and recall metrics for the benign class are lower, suggesting that the model may be having difficulties in correctly identifying benign cases.

Round 2: The results for this round demonstrate good stabilization, as evidenced by the loss metrics reaching lower values more swiftly compared to Round 1. Additionally, the accuracy remained high throughout the epochs, indicating better performance stability. A notable increase in accuracy is observed, with results ranging from 96.58% to 97.53%. Precision and recall improve, especially for the malignant class, suggesting that federated learning is tuning the model for better predictions.

Round 3: The analysis of the model's performance during Round 3 and can be observed in Figures 7, 8, and 9 reveals a continuous improvement in its effectiveness over time. The observed decrease in loss and the sustained high levels of accuracy suggest that the model's learning process was progressing as expected.




Figure 7. ResNet50 Loss and Accuracy C1 R3

- **Client 1:** This client showed consistency in training accuracy, with values exceeding 99% (see Figure 7). However, validation loss fluctuated, ranging between 0.0951 and 0.2000, suggesting some variation in the model's ability to generalize to the validation set. Validation accuracy remained robust, staying predominantly between 96% and 98%, indicating that despite the fluctuations in loss, the model still performs well in predictions. This may suggest that while the model continues to improve internally, there are challenges in generalization that need to be addressed.
- **Client 2:** This client demonstrated consistency in training accuracy, which remained above 99% (see Figure 8). However, validation loss ranged between 0.0942 and 0.2106, indicating difficulties in generalizing the validation set. Despite these fluctuations in loss, validation accuracy remained robust, varying from 96.39% to 98.29%, reflecting good overall model performance in predictions. This pattern suggests that although the model continues to improve its internal performance,

there may be specific challenges in adapting to the validation set that needs to be addressed.




Figure 8. ResNet50 Loss and Accuracy C2 R3

- Client 3:** This client consistently achieved good training accuracies, staying above 99% (see Figure 9). Despite this, the validation loss fluctuated between 0.0848 and 0.1827, hinting at some trouble the model had in generalizing to the validation set. Even with these ups and downs in validation loss, the model's accuracy remained solid, mostly between 95.58% and 97.77%. This suggests that, although the model is improving its performance internally, it still faces challenges when adapting to new data that might need further attention.




Figure 9. ResNet50 Loss and Accuracy C3 R3

Accuracy reaches values ranging from 95.64% to 97.91%. The model demonstrates excellent performance across both classes, with a high F1-score, indicating good generalization ability after several rounds of training. The evolution of the ResNet-50 model's performance across different training rounds is illustrated in Figure 10, where the X-axis represents the rounds, and the Y-axis represents the loss (left graph) and accuracy (right graph). The figure displays the model's accuracy and loss trends during the training process, illustrating key changes and overall performance improvements.

4.2.2 Centralized Learning

Centralized training shows solid performance as shown in Figure 11 with an accuracy of 97.53%, which is comparable to or better than the model's performance after federated learning rounds.




Figure 10. ResNet50 Loss and Accuracy across the Rounds




Figure 11. ResNet50 Loss and Accuracy (Centralized Learning)

However, federated learning offers advantages by enabling training on distributed data without the need for centralization, which can be beneficial in scenarios involving sensitive or distributed data [16, 17]. This approach aligns with ethical principles by prioritizing confidentiality and minimizing the need for data sharing across institutions [5, 17].

5. Conclusion

This study compared the performance of ResNet-50 and AlexNet models in federated and centralized learning scenarios, demonstrating that both models are effective for classification tasks in a federated environment. ResNet-50 exhibited a slight advantage in terms of accuracy and performance capability, attributed to its deeper residual architecture, with accuracy ranging from 95.64% to 97.91%. AlexNet also showed robust performance, although it faced some limitations due to its shallower architecture, with accuracy ranging from 94.88% to 96.96%.

The results indicate that while both models are suitable for federated learning, the choice between them may depend on specific problem requirements, such as the need for higher accuracy or model complexity. The comparison of loss and accuracy results, along with the analysis of classification reports, confirms the effectiveness of the models in both federated and centralized environments, providing meaningful conclusions for the practical application of these models.

As a future work, we aim to explore the application of additional regularization and optimization techniques to enhance model performance in federated scenarios. Furthermore, in-

investigating even deeper or task-specific neural network architectures could offer a more detailed understanding of how different models perform in complex federated environments. Additional studies could also evaluate the impact of various model aggregation strategies and privacy considerations on the overall performance of federated learning.

Acknowledgment

André R. Backes gratefully acknowledges the financial support of CNPq (National Council for Scientific and Technological Development, Brazil) (Grant #307100/2021-9). This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brazil (CAPES) - Finance Code 001.

Author contributions

GVGB: Conceptualization, Methodology, Software, Investigation, Validation, Visualization, Writing - original draft. LGFR: Conceptualization, Methodology, Validation, Writing - original draft. LFRM: Methodology, Validation, Writing - Review & Editing. ARB: Supervision, Funding acquisition, Resources, Writing - Review & Editing, Project administration.

References

- [1] INSTITUTO NACIONAL DE CâNCER - INCA. *Câncer*. 2022. Disponível em: <https://www.gov.br/inca/pt-br/assuntos/cancer>. Acesso em: 25 oct. 2022.
- [2] INCA, I. N. de Câncer José Alencar Gomes da S. *Deteção Precoce do Câncer*. [S.l.]: Ministério da Saúde, 2021. v. 1.
- [3] SHEN, D.; WU, G.; SUK, H.-I. Deep learning in medical image analysis. *Annual review of biomedical engineering*, Annual Reviews, v. 19, n. 1, p. 221–248, 2017.
- [4] LITJENS, G. et al. A survey on deep learning in medical image analysis. *Medical image analysis*, Elsevier, v. 42, p. 60–88, 2017.
- [5] LI, T. et al. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, IEEE, v. 37, n. 3, p. 50–60, 2020.
- [6] SABAH, F. et al. Model optimization techniques in personalized federated learning: A survey. *Expert Systems with Applications*, Elsevier, p. 122874, 2023.
- [7] SPANHOL, F. A. *Automatic breast cancer classification from histopathological images: a hybrid approach*. Tese (Ph.D. Thesis) — Federal University of Paraná (UFPR), Curitiba, Brazil, 2018.
- [8] VO, D. M.; NGUYEN, N.-Q.; LEE, S.-W. Classification of breast cancer histology images using incremental boosting convolution networks. *Information Sciences*, v. 482, p. 123–138, 2019. ISSN 0020-0255.
- [9] AGBLEY, B. L. Y. et al. Federated learning-based detection of invasive carcinoma of no special type with histopathological images. *Diagnostics*, MDPI, v. 12, n. 7, p. 1669, 2022.
- [10] TAN, Y. N. et al. A transfer learning approach to breast cancer classification in a federated learning framework. *IEEE Access*, v. 11, p. 27462–27476, 2023.
- [11] BARBOSA, G. et al. Optimization and Learning Rate Influence on Breast Cancer Image Classification. In: INSTICC. *Proceedings of the 19th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 3: VISAPP*. [S.l.]: SciTePress, 2024. p. 792–799. ISBN 978-989-758-679-8.
- [12] SPANHOL, F. A. et al. A Dataset for Breast Cancer Histopathological Image Classification. *IEEE Transactions on Biomedical Engineering*, v. 63, n. 7, p. 1455–1462, 2016.
- [13] FLOWER LABS. *A Friendly Federated Learning Framework*. 2023. Disponível em: <https://flower.dev>. Acesso em: 07 nov. 2023.
- [14] KRIZHEVSKY, A. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997*, 2014.
- [15] HE, K. et al. Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. [S.l.: s.n.], 2016. p. 770–778.
- [16] MOREIRA, R. et al. An intelligent native network slicing security architecture empowered by federated learning. *Future Generation Computer Systems*, v. 163, p. 107537, 2025. ISSN 0167-739X. Disponível em: <https://www.sciencedirect.com/science/article/pii/S0167739X24005016>.
- [17] ADNAN, M. et al. Federated learning and differential privacy for medical image analysis. *Scientific reports*, Nature Publishing Group UK London, v. 12, n. 1, p. 1953, 2022.