

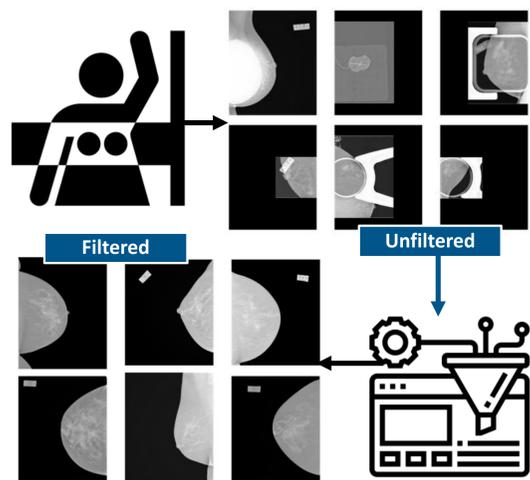
Simple 2-Step Process for Removing Unwanted Images in Large Mammography Dataset

#67

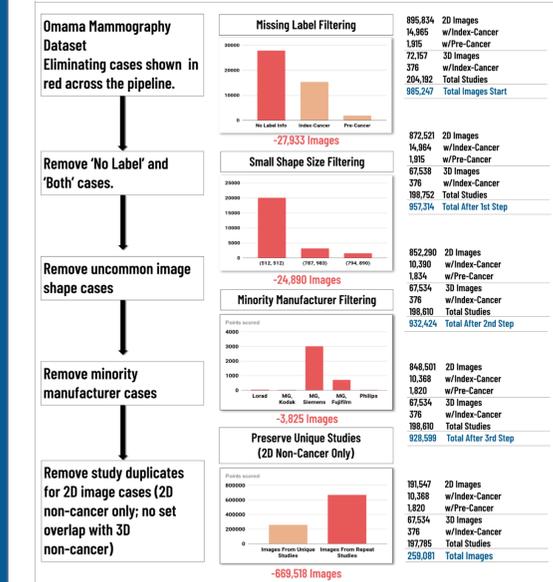
Ryan Zurrin, Neha Goyal, Pablo Bendiksen, Muskaan Manocha, Daniel Haehn

Abstract

We propose an automated pipeline designed to filter low-quality mammogram images..

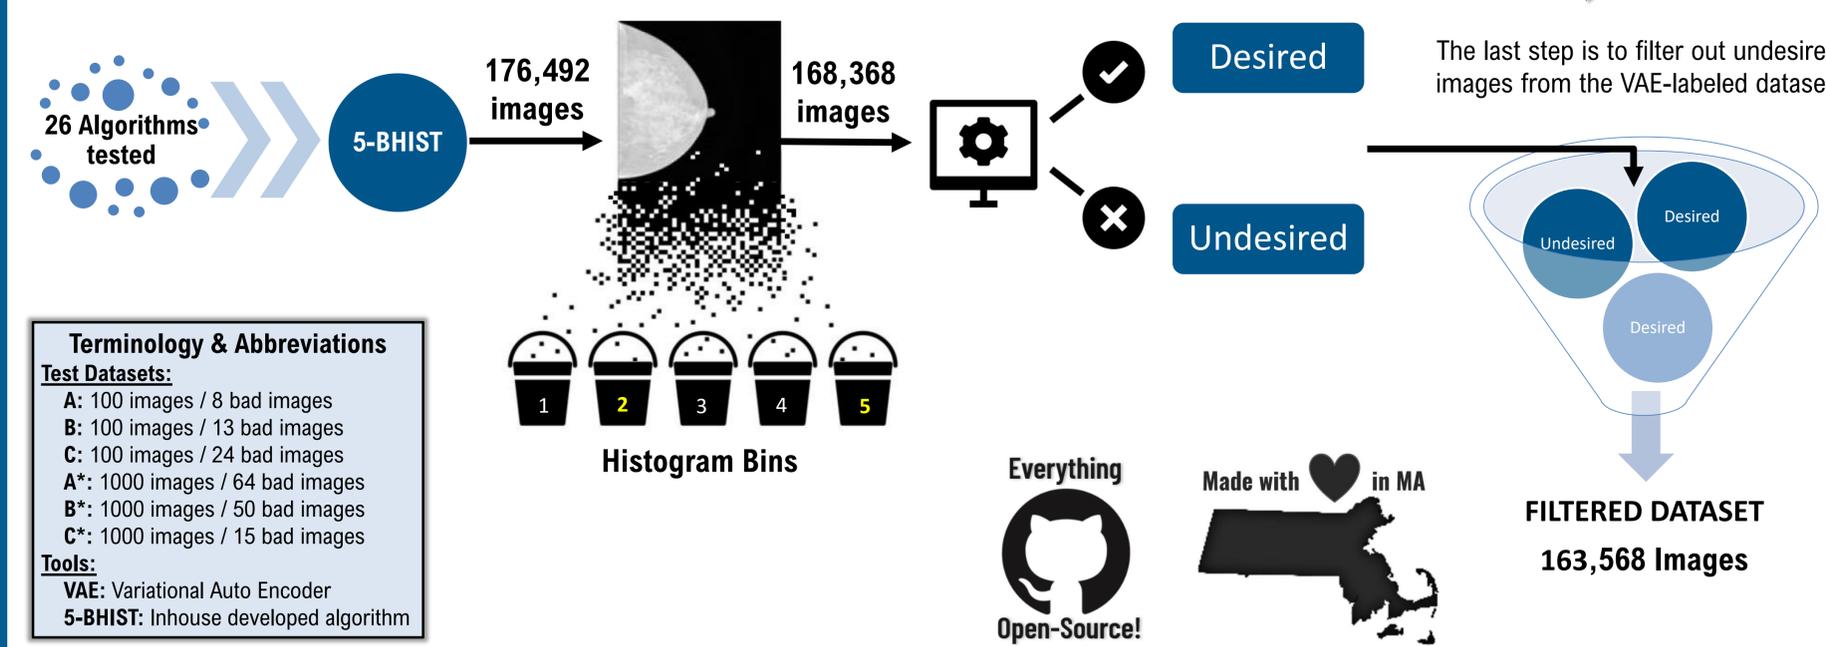
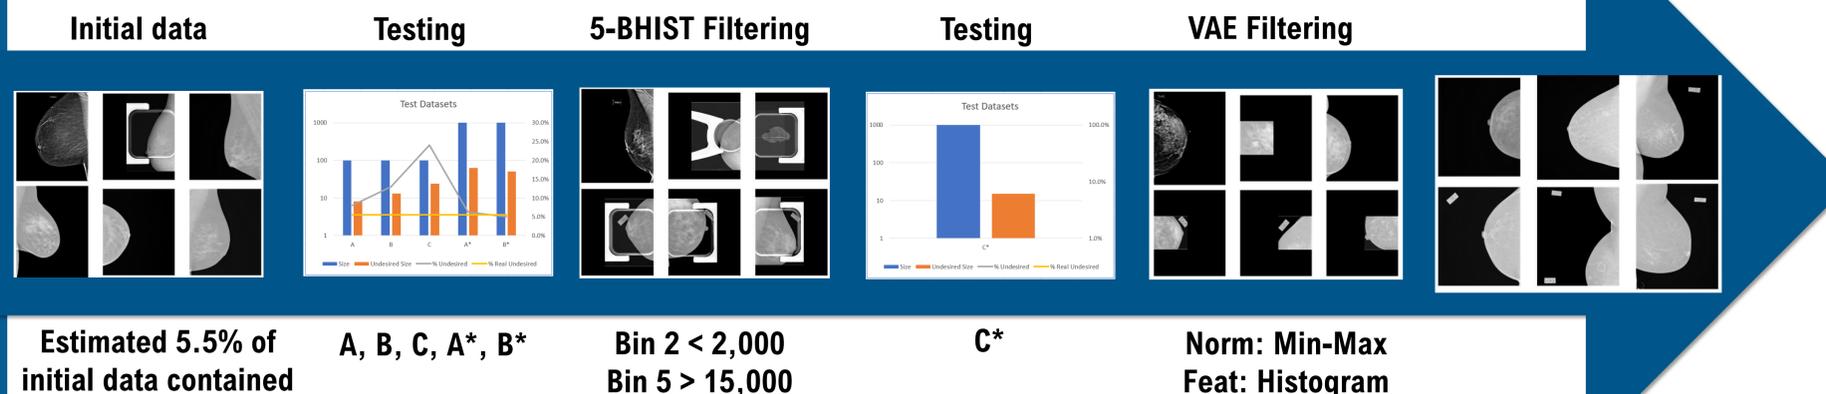


Introduction



Methodology

This poster showcases a two-stage approach for removing undesirable images (outliers) from a large collection of mammograms. Stage 1 uses *threshold-based 5-bin histogram filtering (5-BHIST)* to eliminate unwanted images, and then for stage 2 it employs a *variational autoencoder (VAE)* to further refine and remove any remaining subpar scans. This innovative method was devised after rigorously testing 26 unsupervised outlier detection algorithms.



Results

F1 Score: A measure of an algorithm's accuracy, combining precision (true positives) and recall (true positives among all relevant elements). Ranging from 0 (worst) to 1 (best), it helps evaluate the effectiveness of outlier detection methods

Cleaned Dataset

- 163,568 images
- From 5.5 % to 1% unwanted images
- Total reduction by factor 5.93x
- Average F1 score of .8862
- < 6 hours to process over 176,000 images

Conclusion

Utilizing min-max normalized histogram binning with a variational autoencoder, we've reduced low-quality scans in our dataset by nearly 6x, achieving an average F1 Score of 0.8862. Consequently, only 1% of images in our final dataset are unwanted, verified by manual inspection.

Acknowledgments



Zurrin, Ryan, et al. "Outlier Detection for Mammograms". Medical Imaging with Deep Learning, Short Paper Track, 2023, <https://openreview.net/forum?id=4E93Xdg98u>.