Breast Cancer Detection from Thermal Images using Machine Learning

Sijche Pechkova
Faculty of Technology and Metallurgy, Skopje, North Macedonia
Lyudmyla Venger
IPectus Project, Berlin, Germany
Dragana Andonovski
North Kansas City Hospital, Missouri, United States
Beti Andonovic
Faculty of Technology and Metallurgy, Skopje, North Macedonia

Abstract

In this study, the authors propose an advanced strategy to analyze thermal images for breast cancer detection employing machine learning techniques. By focusing on critical features that capture geometric and structural information in thermal images, the aim is to elevate the precision and uniformity of breast cancer diagnostics. The dataset comprises thermal images from patients with breast cancer; these vital features are extracted and integrated into proposed decision tree model, resulting in a classification accuracy of 92%. This highlights the utility of combining specialized features with machine learning algorithms in medical image analysis. Consequently, the findings suggest that this approach can substantially enhance traditional imaging methods, establishing a robust basis for early and accurate breast cancer detection.

Keywords: breast cancer, thermal images, machine learning
JEL classification: Y80

Paper type: Research article
Received: 12 February 2024
Accepted: 29 Jun 2024

DOI: 10.54820/entrenova-2024-0042

Introduction

There are many studies that support thermography as a promising supplementary tool rather than a standalone solution for breast cancer screening. They suggest it may benefit patients at high risk or with specific breast characteristics, but that standardization and further research are needed to optimize its role. The consensus is that thermography's non-invasive nature and potential to detect heat irregularities make it a valuable addition to mammography, especially with advances in AI and machine learning improving its diagnostic capacity.

Studies that evaluate the effectiveness of thermography as a diagnostic tool for breast cancer, focusing on its potential to complement mammography. Kim et al. (2015) and Mok et al. (2013) conducted meta-analyses and found that thermography can detect abnormal heat patterns linked to cancerous tissue, but its sensitivity and specificity are inconsistent. Their results indicate that thermography alone cannot reliably replace mammography but may have value when used as an adjunct tool. Ahmad et al. (2014) further explored thermography's ability to assist mammography, particularly for patients with dense breast tissue, where traditional screening may miss early signs of cancer.

Zhang et al. (2017) reviewed infrared thermography's role in screening and found that while it can detect cancer-related thermal anomalies, its diagnostic precision falls short of mammography. Czerwinski et al. (2015) echoed these findings, pointing out that while thermography is non-invasive and radiation-free, its accuracy varies, likely due to differences in technique and technological limitations.

Rastghalam & Pourghassem (2013) and Selamat et al. (2018) focus on refining the process of feature extraction in thermal imaging. Rastghalam & Pourghassem developed a spectral feature extraction approach that successfully identifies asymmetries in thermograms, which often indicate abnormalities. Selamat et al. provide a comprehensive review of infrared imaging techniques, discussing how increased thermal activity in cancerous tissue can be effectively distinguished from healthy tissue through digital analysis, thus enhancing the diagnostic precision of thermographic imaging.

Pramanik et al. (2019) conducted detailed studies on segmentation methods for thermographic breast images. One study by Pramanik et al. uses level-set segmentation to isolate suspicious regions by applying statistical and texture-based analysis, while another examines advanced statistical segmentation techniques to differentiate malignant areas, offering greater precision in pinpointing abnormal regions.

Mohamed et al. (2022) and Goncalves et al. (2022) explore the application of convolutional neural networks (CNNs) to classify and predict breast cancer in thermographic data. Mohamed et al. achieved high diagnostic accuracy by using a CNN-based classification system, highlighting its potential for automated detection in thermography. Goncalves et al. further refined CNN performance by applying bio-inspired algorithms to optimize the architecture, which significantly improved diagnostic results.

Chatterjee et al. (2022) developed deep feature selection methods for texture analysis, which allow for the enhancement of classification outcomes in thermal imaging. In one of their studies, they use Grunwald-Letnikov-based feature selection with a deep learning model, demonstrating improved classification accuracy for detecting early cancer indicators.

Studies on breast cancer detection using thermal imaging increasingly integrate machine learning (ML) and artificial intelligence (AI) to improve diagnostic accuracy. Thermal imaging captures surface temperature distributions, highlighting heat anomalies linked to cancerous tissue. Models like convolutional neural networks (CNNs) have achieved high accuracy in detecting breast cancer on thermograms, especially with image preprocessing techniques like segmentation. Approaches involving texture analysis, feature extraction and color morphology further refine detection.

In this study, we present a method of utilizing advanced machine learning techniques to analyze thermal images for breast cancer detection. Our focus lies in identifying critical features that capture crucial geometric and structural information from these images, thus improving the accuracy and consistency of breast cancer diagnosis. Employing a dataset comprised of thermal images from patients with confirmed breast cancer, we extract and incorporate these vital features into our decision tree model, achieving a classification accuracy of 92%. This demonstrates the effectiveness of combining specialized features with machine learning algorithms in medical image analysis. Our findings suggest that this approach can significantly enhance traditional imaging methods, laying a strong foundation for early and precise breast cancer detection.

The task involves processing and analysing thermal images of approximately 400 patients, each consisting of a series of thermal images. The patient cohort includes individuals both with and without confirmed cancer diagnoses. The goal is to determine if there are discernible patterns or features in the thermal images that can help identify potential cancer cases among the given dataset.

The Intuition behind our Approach

Asymmetry is a common characteristic of cancer growth as it spreads unevenly and irregularly in the body, making it difficult to identify consistent patterns or symmetries that can be reliably used for analysis.

To address this challenge, various techniques are explored in order to extract features from medical images in a symmetrical way. Symmetrical feature extraction refers to the process of identifying and analyzing patterns or features that have reflective or mirror-like properties across the left and right sides of an image. This approach is based on the assumption that if cancer exhibits asymmetry, then its reflection or mirror image on the other side of the body may provide valuable information for diagnosis and prognosis.

However, extracting symmetrical features from medical images can be a complex task due to several challenges. First, as mentioned in the text, it is often unclear which features are relevant for analysis since asymmetry may not be easily described or quantified. Second, medical images can contain significant variability and noise, making it challenging to reliably identify symmetrical patterns. Third, the presence of artifacts or other structural differences between the left and right sides of an image can further complicate the process of identifying symmetrical features.

Despite these challenges, researchers are exploring various approaches for symmetric feature extraction in medical imaging, including the use of deep learning models, registration techniques, and statistical analysis methods. These approaches aim to automatically identify and extract symmetrical features from medical images while accounting for variability and noise, as well as the presence of structural differences or artifacts. Ultimately, the goal is to develop robust and accurate methods for analyzing asymmetrical medical data that can help improve diagnosis, treatment planning, and patient outcomes.

Methodology

The work involves collecting thermal images from about 400 patients, labeled by their health status, to create a robust dataset for training a machine learning model. After preprocessing the images—resizing, normalizing, and augmenting them to ensure consistency and improve robustness—meaningful features will be extracted using different methods. Then selected machine learning algorithms, such as Random forests of Decision Trees will be trained on the dataset, and their performance evaluated through metrics to predict the likelihood of cancer in new patients based on their thermal images. This are the main points:

1. Data Collection: Gather thermal images from approximately 400 patients, ensuring to document their health status (either healthy or diagnosed with cancer) to create a comprehensive labeled dataset. This dataset will serve as the foundation for training and testing our machine learning model.
2. Feature Extraction: To derive significant features from the thermal images, we will employ techniques like Histogram of Oriented Gradients (HOG), Local Binary Patterns (LBP), or Convolutional Neural Networks (CNN). These extracted features will be utilized as inputs for our machine learning model.
3. Model Training: Select an appropriate machine learning algorithm, such as Support Vector Machines (SVM), Random Forest Classifier, or deep learning architectures. The dataset will be divided into training and testing subsets, employing methods like k-fold cross-validation to assess model performance.
4. Model Evaluation: Evaluate the performance of the trained machine learning model by analyzing various metrics. The model will then be used to predict whether a new patient represented by an unlabeled thermal image is likely to have cancer. It will produce a probability score indicating the likelihood of cancer presence.

Data Collection

We collected thermal images from approximately 400 patients, carefully recording each individual's health status (whether healthy or diagnosed with cancer) to construct a well-labeled and detailed dataset. This dataset is crucial for training and testing our machine learning model, offering a wide array of thermal patterns associated with both cancerous and non-cancerous conditions.

Figure 1
Example of patient data.

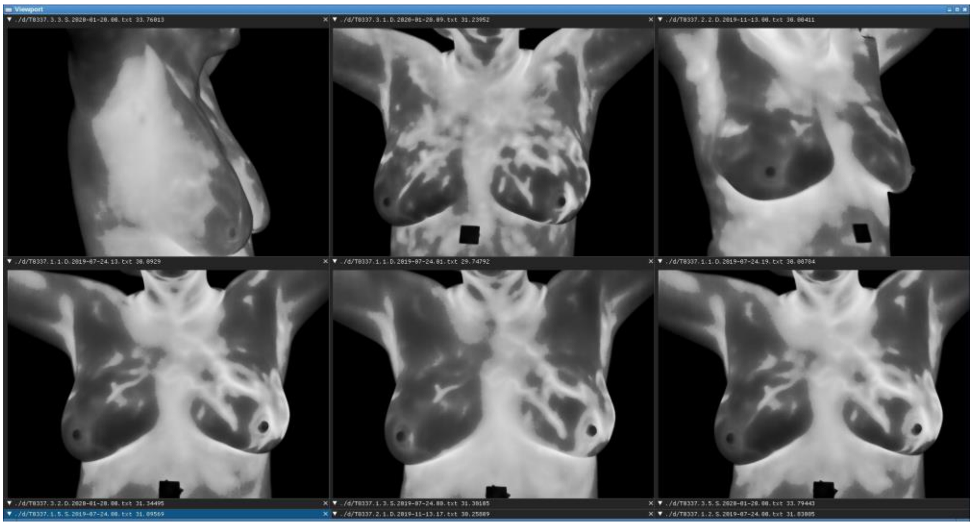

Source: Authors' work

Each patient's data includes several frontal images, along with side-view images captured at 90 and 45-degree angles, thus providing varied perspectives to enhance model accuracy and robustness in detecting thermal asymmetries that might signify underlying abnormalities.

With over 6,000 images in total, this dataset captures an extensive range of body spositions and thermal variations. Such diversity is intended to improve the model's generalizability by allowing it to learn from different anatomical views and thermal signatures unique to each patient. The wide range of samples also contributes to identifying subtle patterns and temperature discrepancies that might otherwise go unnoticed. By incorporating these different angles and multiple patient perspectives, the dataset establishes a solid foundation for a machine learning model aimed at achieving high diagnostic accuracy and reliability in breast cancer detection through thermal imaging.

Feature Extraction

In order to identify and extract areas of interest from an image, we first need to train a suitable model. This could be achieved using various machine learning or computer vision techniques. Once the area of interest is identified using the trained model, we will proceed to extract features from it for further analysis. The application that we made has the ability to mark the area of interest using circles. After we set the area of interest for some images, we apply machine learning techniques to learn the coordinates of the circles. For this we are using filtered data as described below.

Figure 2

Example of patient data with area of interest. The circles on the image represent the area of interest.

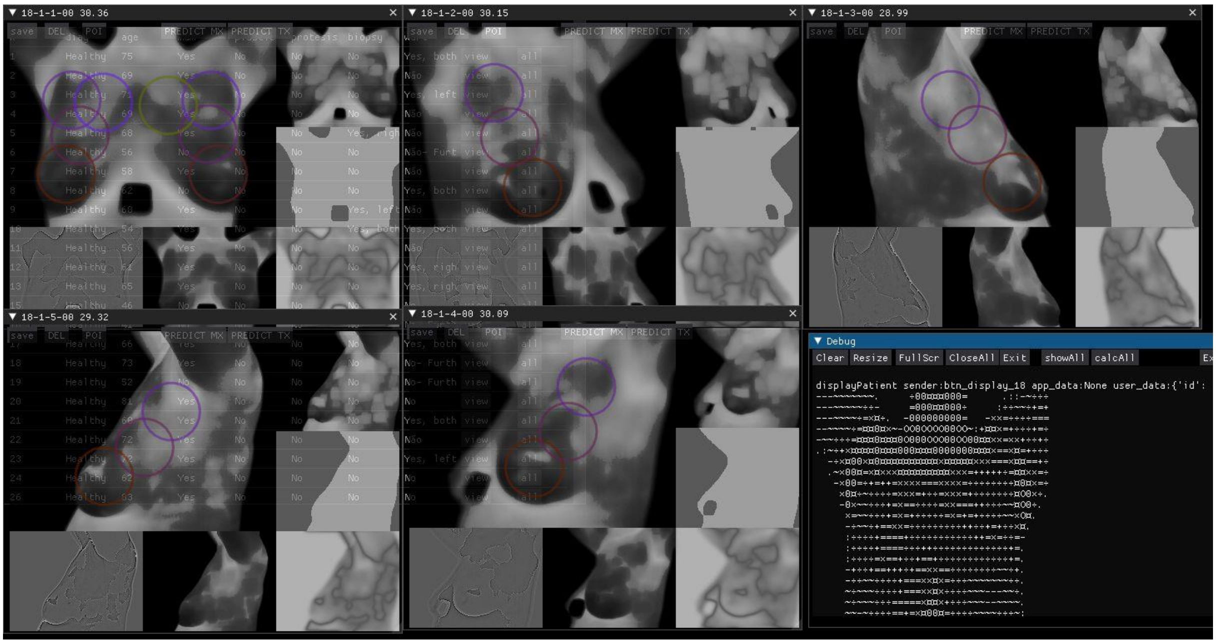

Source: Authors' work

Standard statistical features are numerical values calculated from the image data, such as mean, median, variance, standard deviation, histogram, and moment invariant features. These features help in describing the intensity distribution and color information of the area of interest.

To prepare the image for feature extraction, we apply several image preprocessing filters to enhance edges, suppress noise, and reduce unwanted data. Some of these filters include:

- The **Laplace filter** is a second-order derivative filter used for edge detection in images. It highlights regions of rapid intensity change and is commonly used to find edges and textures. The filter calculates the Laplacian of the image, emphasizing areas where the intensity gradient changes significantly. This filter is particularly effective for detecting edges and is often applied as a preprocessing step for further image analysis tasks. The Gaussian gradient filter combines Gaussian smoothing with gradient computation, providing a way to detect edges while reducing noise.
- The **Gaussian Gradient filter** is a powerful image processing technique that integrates Gaussian smoothing with gradient computation to detect edges while minimizing noise. Initially, it applies a Gaussian blur to the image, effectively reducing high-frequency noise, which enhances the reliability of the subsequent gradient calculation. By measuring changes in pixel intensity, typically using operators like Sobel, the filter highlights edges and transitions in the image. This dual-functionality makes the Gaussian Gradient filter particularly effective for tasks such as feature extraction, object detection, and segmentation, especially in applications where edge integrity and noise reduction are critical.
- The **Fourier Gaussian filter** operates in the frequency domain, utilizing the Fourier transform to apply a Gaussian filter to an image. This filter effectively reduces high-frequency noise while preserving low-frequency components, making it suitable for tasks such as image smoothing and noise reduction. The Gaussian shape in the frequency domain corresponds to a smooth blurring effect in the spatial domain, which can be beneficial in applications requiring enhanced image quality.
- The **Minimum filter** replaces each pixel's value with the minimum value within a specified neighborhood around that pixel. This filter is particularly useful for removing small-scale noise or bright spots from images while preserving the overall shape and features of larger objects. It is often used in conjunction with other filters to enhance the quality of the image by suppressing isolated high-intensity pixels.

These filters are kernel-based transformations that change the intensity distribution of an image based on specific mathematical operations, thereby bringing out the desired features.

Additionally, we manually label a subset of images (approximately 6000) to create a training dataset for our machine learning model. This labeled data will be used to teach the model how to identify and classify different areas of interest. Once the model is trained, it can be applied to a large dataset of unlabeled images to automatically extract and label their respective areas of interest based on the learned features.

In our image processing system, we carefully extract distinct features from the specified regions of an input image for comprehensive analysis. The process is carried out separately for both the left and right sides of the image. To clarify, the following are the specific features we extract:

1. **Fos (First Order Statistics):** This feature represents the first order statistical measures, including mean, standard deviation, variance, skewness, kurtosis, and energy,
2. extracted from the intensity values of the pixels in the region.

3. **Ngtdm (Neighborhood Gray Tone Difference Matrix):** The Neighborhood Gray Tone Difference Matrix is a texture feature that describes the spatial distribution of grayscaleintensity differences within a defined neighborhood. This feature helps us understand the coarseness and contrast of the texture in the region.
4. **Sfm (Statistical Feature Matrix):** Statistical Feature Matrix represents higher order statistical features, such as autocorrelation, entropy, correlation matrix, and energy spectral density, that provide more detailed information about the distribution of pixel intensities within a region.
5. **Fdta (Fractal Dimension Texture Analysis):** Fractal Dimension Texture Analysis is a feature that quantifies the irregularity or complexity of the texture by determining its fractal dimension. This feature helps in identifying features such as cracks, pores, and rough surfaces.
6. **Fps (Fourier Power Spectrum):** The Fourier Power Spectrum analyzes the image in terms of its frequency components by calculating the power spectral density of each pixel in the region. This feature is particularly useful for detecting periodic patterns or shapes.
7. **Glszm (Gray Level Size Zone Matrix):** Gray Level Size Zone Matrix is a texture feature that characterizes spatial relationships between pixels with similar grayscale intensity values within various-sized neighborhoods. It describes the distribution of such zones in the image, providing insights into the uniformity and coarseness of the texture.
8. **Lbp (Local Binary Pattern):** Local Binary Pattern is a feature that represents the spatial arrangement of pixel intensity differences within a defined window or neighborhood. This feature helps in identifying local structures and patterns in the image.
9. **Dwt (Discrete Wavelet Transform):** Discrete Wavelet Transform is a multi-resolution transform used to analyze images at different scales, preserving both spatial and frequency information. It can effectively capture features such as edges and corners.

For each specified region in the image, we extract all these features. In total, this results in a feature vector with 732 elements.

Results

Model Training

For our predictive model we are using Random Forest of Classification Trees. It is an ensemble learning method that combines multiple decision trees to create a powerful and robust classifier. The algorithm reduces overfitting by introducing randomness in both the selection of training data and the features used for splitting, ultimately improving the accuracy and generalization capabilities of the resulting model.

The algorithm operates by generating numerous decision trees, each trained on a random subset of the data with a unique selection of features at each split. This process, known as bootstrapping and feature randomization, reduces overfitting by decreasing correlations among trees. Each tree is trained independently to make accurate predictions on its subset. When classifying a new data point, all trees in the forest make independent predictions. For classification tasks, the final classification is determined by majority voting across the trees, where the most frequent class is selected. Alternatively, averaging probabilities across trees can yield a refined probability estimate, useful for providing nuanced results in both classification and regression tasks.

They offer several advantages: they are highly adaptable and can model nonlinear relationships between features and target classes, making them applicable to a wide range of problems. By utilizing multiple decision trees, they reduce overfitting risk through an averaging effect, which mitigates noise from individual weak learners. Suitable for both classification and regression tasks, Random Forests perform effectively even in high-dimensional spaces with large datasets. They also provide insights into feature importance, helping users identify the most influential variables in the predictions, and they are relatively easy to train.

Model Evaluation

We are using 10 fold cross-validation for evaluating the predictive performance of the model. Cross-validation is a statistical method used for evaluating models and assessing how well the model generalizes to unseen data. The main idea behind cross-validation is to partition our data into different subsets: one for training the model (the "training set") and another for testing it (the "test set"). For 10-Fold Cross-Validation the dataset is randomly divided into 10 subsets or folds. In each iteration, 9 folds are used for training, while the remaining fold is used as a testing set to evaluate the model's performance. The process is repeated 10 times. The primary advantage of cross-validation is its ability to provide a more robust estimate of our model's generalization error compared to using a single training/test split. This is because it helps to average over the variability introduced by different random train/test splits, leading to more reliable estimates of performance.

To ensure robust validation and prevent overfitting, we carefully assign different patients to separate folds when partitioning the data. This approach mitigates the risk of "fingerprinting," a phenomenon where a model learns to recognize specific individuals' unique characteristics rather than generalizable patterns indicative of broader conditions, like cancer. By keeping each patient's data exclusive to a single fold, the model is tested on entirely new individuals in each validation cycle, simulating real-world conditions where the model must generalize to new patients it has never encountered before. This technique avoids data leakage, where the model might otherwise exploit similarities within the same patient's data across multiple folds, leading to artificially inflated accuracy scores during cross-validation.

This patient-specific partitioning strategy is particularly important in medical imaging, as even slight personal thermal variations between patients can bias the model if not carefully managed. By ensuring that no patient appears in more than one fold, we reduce the chance that the model will rely on unique individual traits, such as anatomy or natural thermal patterns, rather than learning meaningful indicators of disease. This process enhances the reliability of performance metrics and helps create a model that is both accurate and generalizable to new, unseen patients in practical diagnostic settings.

In the final model predictions, we use a simple voting technique to aggregate the results from multiple images of the same patient, improving diagnostic consistency and accuracy. This means that, instead of relying on a single thermal image for each diagnosis, we combine predictions from several images. By aggregating these predictions, the model reduces the risk of errors due to image-specific variations, such as slight temperature inconsistencies or minor positioning differences, which might otherwise mislead a single-image prediction.

This voting approach leverages the diversity of images to stabilize and enhance the model's overall diagnostic reliability. It ensures that the final classification is based on a more comprehensive view of the patient's thermal profile, effectively minimizing the impact of any anomalies or noise in individual images. This process helps reduce false positives and false negatives, as the collective vote emphasizes patterns consistent across multiple images rather than outliers. As a result, aggregating predictions from multiple images per patient leads to a more dependable outcome, especially important in sensitive medical diagnoses like breast cancer detection, where accuracy and precision are critical.

Discussion

Using a Random Forest model composed of classification trees, we achieved an accuracy rate of 92.4%, demonstrating the method's strong capability in distinguishing between healthy and cancerous thermal images. This high accuracy is attributed to the ensemble approach of Random Forest, where each tree independently classifies the data, and the final decision is made through a voting mechanism that minimizes individual tree errors. By employing techniques such as bootstrapping and feature randomization, the model reduces overfitting, leveraging the diversity within the forest to create a more robust predictive framework.

This performance suggests that Random Forests are highly effective for complex image-based diagnostic tasks, particularly in medical imaging where subtle differences may indicate significant clinical outcomes. The 92.4% accuracy level also highlights the model's potential as a diagnostic tool, paving the way for its application in real-world screening scenarios, where consistency and precision are critical. This result underscores the promise of machine learning in supporting early detection efforts and enhancing the reliability of non-invasive diagnostic methods like thermal imaging.

Conclusion

This research showcases the promise of machine learning methods, specifically a Random Forest classifier, to enhance the precision of breast cancer identification via thermography. By concentrating on feature extraction techniques that capture geometric and structural details, we attained a robust classification accuracy of 92.4%, underscoring the effectiveness of machine learning in medical image analysis. With procedures such as bootstrapping and feature randomization, the model prevents overfitting while preserving high precision, particularly for complex tasks involving subtle thermal discrepancies related to cancer.

The use of 10-fold cross-validation with patient-specific data segmentation further bolsters its predictive validity, ensuring that performance metrics are both accurate and transferable to new patients. Additionally, the aggregation of multiple image predictions via a voting mechanism enhances diagnostic consistency, reducing potential errors from image-specific variations. This comprehensive strategy highlights thermography's value as an auxiliary tool in breast cancer screening.

Our findings demonstrate the feasibility of integrating machine learning-enhanced thermography into breast cancer diagnostics, paving the way for further advancements in non-invasive, AI-driven screening tools that could complement mammography, improve early detection, and potentially save lives.

References

1. Ahmad, S., et al. (2014). Thermography as an adjunct screening method to mammography: A systematic review. *PLoS ONE, 9*(11).
2. Chatterjee, S., et al. (2022). Texture analysis in thermograms using a deep feature selection method. *Computers in Biology and Medicine, 141*.
3. Chatterjee, S., et al. (2022). Unique Grunwald-Letnikov-based feature selection approach to improve classification accuracy in thermal images. *Journal of Medical Systems, 46*(8).

4. Czerwinski, A., et al. (2015). Thermal imaging and breast cancer: A systematic review. *BioMed Research International*.
5. Goncalves, R., et al. (2022). CNNs optimized with bio-inspired algorithms to enhance diagnostic accuracy. *Computers in Biology and Medicine, 142*.
6. Kim, J., et al. (2015). Thermography in the detection of breast cancer: A meta-analysis. *Breast Cancer Research and Treatment, 151*(3), 345-352.
7. Mohamed, A., et al. (2022). Deep learning model for breast cancer detection using thermal images, achieving high accuracy through automated classification techniques. *PLOS ONE, 17*(1).
8. Mok, W., et al. (2013). Thermography in breast cancer screening: A systematic review and meta-analysis. *PLoS ONE, 8*(4).
9. Pramanik, S., et al. (2019). Segmentation methods for suspicious regions in thermographic breast images, using advanced statistical analysis to distinguish malignant areas. *IEEE Transactions on Medical Imaging, 38*(2), 480-490.
10. Pramanik, S., et al. (2019). Level-set segmentation techniques to isolate suspicious regions in breast thermograms. *Journal of Digital Imaging, 32*(3), 362-370.
11. Rastghalam, R., & Pourghassem, H. (2013). Spectral feature extraction methods for breast cancer detection in thermography. *Computers in Biology and Medicine, 43*(12), 2024-2033.
12. Selamat, H., et al. (2018). Comprehensive overview of infrared digital imaging techniques for breast cancer detection. *Infrared Physics & Technology, 91*, 273-286.
13. Zhang, Y., et al. (2017). Infrared thermography for breast cancer screening: A systematic review. *Medicine, 96*(2), e7223

About the authors

Sijche Pechkova is a teaching assistant at the Faculty of Technology and Metallurgy at the University of "Ss. Cyril and Methodius" in Skopje, Macedonia. She graduated at the Faculty of Science and Mathematics in Skopje. She has a master's degree in the field of statistical methods in business and economics. Pechkova's scientific research work is in the areas of applied mathematics, mathematical modeling, programming, statistics, machine learning, artificial intelligence and computer engineering. Sijce Pechkova has participated in several national projects and conferences. She can be contacted at **sijche@gmail.com**

Lyudmyla Venger leads iPectus, a cutting-edge company that offers a portable, radiation-free tool for breast cancer risk assessment. With a commitment to advancing non-invasive healthcare solutions, she drives iPectus's mission to make breast cancer screening safer and more accessible. Through innovative technology, iPectus provides an alternative to traditional imaging methods, enhancing early detection and empowering patients with critical health information. Lyudmyla Venger can be contacted at **millawenger@gmail.com**

Dragana Andonovski has a Bachelor's Degree in Biology with a minor in Chemistry from Park University, Missouri, United States. She is aspiring to further her academic studies in the field of medicine. She currently works in the laboratory of North Kansas City Hospital in Missouri, United States.The author can be contacted at **andonovski.dragana@gmail.com**

Beti Andonovic, PhD is an Full Professor at the Faculty of Technology and Metallurgy, Skopje, Macedonia. She obtained her PhD in mathematics at the Faculty of Mathematics and Natural Sciences, University St. Cyril and Methodius, Skopje, Macedonia, in 2009. She was Head of the Department of Chemical and Control Engineering at the Faculty of Technology and Metallurgy 2012-2016. She is the author of many scientific articles in mathematics and mathematical modelling, as well as in management, and is the author of two University books on the subjects of Mathematics and Communication skills. She has presented her scientific research at numerous international conferences and has given invited talks at Universities in Macedonia and abroad. The author can be contacted at **beti@tmf.ukim.edu.mk**