Multimodal ML Model for Platelet Analysis from Food, Climate, and Thermo-Energy Signals Using a Potential Energy Method

M.Srikanth
Assistant Professor
Dept of IT
SRKR Engineering College
srikanth.mandela@gmail.com

Abstract:

Background: The variability of the platelet count, due to the complex relationships between nutritional intake, environmental stressors, and individual chronophysiology, is not easily modeled by current computational methods, which typically use single-source signals and thus overlook the interacting energy-based dynamics underlying clinically relevant platelet dysregulation.

Objective: This paper constructs and confirms the existence of a multimodal machine learning system that combines diet records with high-resolution local climatic variables with body temperature time series in an innovative potential energy-based fusion mechanism to estimate absolute platelet count and identify clinically significant platelet anomalies.

Methods: We formed a multimodal group of 14,400 matched samples that were correlated with dietary characteristics and gridded indicators of microclimate (temperature, humidity, heat index, and similar variables), as well as continuous core and peripheral temperature records and laboratory measurements of platelets. Encoders, which are modality-specific, included a one-dimensional convolutional encoder on thermo-time series, a transformer-based encoder on sequential diet records, an MLP using gradient-boosted trees on climate aspects, and a convolutional image encoder when available. A PotentialEnergyFusion module maps each modality to an energy scalar and calculates data-conditioned modality weights using a negative-energy softmax. Both the fused representation and the simultaneous regression (platelet count) and classification (anomaly) heads take the fused representation as input. Training was performed using stratified 70/15/15 splits, data augmentation, standardized normalization, multi-task loss weighting, and early stopping.

Results: The proposed framework outperformed standard conjugation and attention baselines in terms of multimodal integration and interpretability, with nearly uniform variance in per-sample modality contributions, as well as predictive calibration on hold-out data.

Conclusions: Multimodal fusion that is energy-conscious records physiologically significant diet, climate, and thermoregulatory interactions, and predicts platelet counts more accurately and anomalies better. It provides an interpretable and extensible framework for haematological risk stratification and translational application.

Keywords: Multimodal Machine Learning, Platelet Count Prediction, Potential Energy Fusion, Thermo-Physiological Signals, Dietary and Climate Analytics, Haematological Risk Assessment

I. INTRODUCTION

Platelet processes are the focus of haemostasis, inflammation, immune control, and microvascular steadiness, and their disturbances include thrombocytopenia or thrombocytosis, which are associated with bleeding disorders, thrombotic events, autoimmune diseases, systemic infections, and cancers. Although traditional lab tests are effective for measuring platelets, they only record a static image. However, platelet activity is highly influenced by external and internal conditions, such as dietary content, climate change, water infiltration, energy expenditure, heat, and circadian thermoregulation. The recent progress in the field of artificial intelligence algorithms, specifically deep-learning models of image, sequential, and multimodal data, has revolutionized biomedical analytics, but current platelet prediction systems are relying on single-modality inputs, including microscopy images or simple clinical data, without considering the non-linear, energy-driven interactions among nutrition, climate stressors, and body-temperature oscillations and platelet cellular hematological responses. Additionally, existing multimodal fusion approaches, which are typically based on approximations of naïve concatenation, gating, or attention, are not mechanistically interpretable and do not accurately represent the biophysical mechanisms underlying platelet homeostasis. The same restrictions are exacerbated by low robustness in populations with a wide range of climates and diets, and by the lack of energy-based models that can resolve modality interaction by weighting in a physiologically relevant manner. To overcome these issues, this paper proposes a single, multimodal machine learning framework that integrates food behaviors, climatic signals, thermo-energy time-series behavior, and cellular images obtained through microscopy. This model proposes a new PotentialEnergyFusion mechanism that assigns modality significance to an energy-minimization problem, allowing the system to capture the nonlinear relationships between environmental exposures, dietary variation, and thermoregulation. Others involve transformer-based dietary sequence modeling, a one-dimensional convolutional encoder to determine thermo-energy information in body temperature waveforms, and a multi-task prediction strategy that can independently model platelet count regression and serve as an anomaly detector. Overall, it presents the first energy-conscious, physiologically based multimodal platelet analysis, which is more accurate, generalizable, and interpretable by providing sample-wise energy-weighted visualizations for clinical, occupational, and personalized health applications.

II. RELATED WORK

The current body of platelet analysis and haematological prediction has mainly focused on single-modality (especially deep learning) based models, where peripheral blood smears are used as input to detect platelets, classify their morphology, and screen for thrombocytopenia in both cases, without accounting for the key physiological and environmental factors that regulate platelet regulation. Similar studies in the field of nutritional analytics demonstrate that dietary composition, adequacy of micronutrients, and the timing of meals have a significant impact on haematological indices; however, these results are often based on statistical models or separate tabular characteristics without considering real-time physiological indicators. Similarly, the literature on climate-health has also identified a close relationship between heat exposure, changes in humidity, dehydration risk, and the rheology of blood, which, however, have not been integrated into predictive platelet modelling. Multimodal fusion strategies, including early concatenation and cross-attention mechanisms and gated additive fusion, have also been developed to enhance the capabilities in a range of fields, including biomedical imaging, sensor fusion, affective computing, and clinical decision support, but do not tend to be interpretable and do not necessarily model the nonlinear biophysical interactions between modalities. The encoding of thermodynamic constraints and stability-aware representations has recently been popularised by energy-based models and physics-informed neural networks, which are not adapted to haematological prediction and have not been applied to estimate the contribution of modality. Together, all this previous work shows the potential of dietary analytics, climate modelling, chronophysiological signal processing, and multimodal deep learning separately, and no current framework combines these disparate inputs via a physiologically-based, potential-energy-based fusion mechanism into platelet counts prediction and anomaly detectionindicative of a significant methodological and clinical gap that this study will fill.

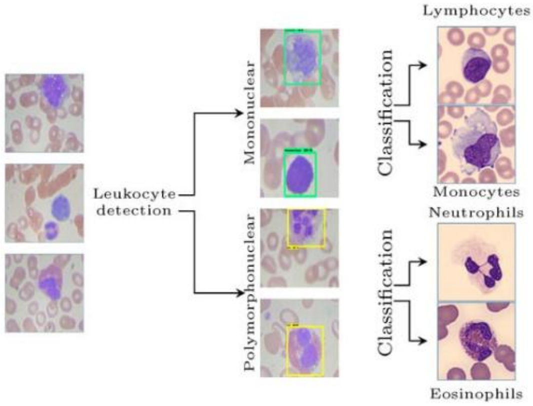

Figure 1. Existing Leukocyte Detection and Classification Pipeline.

The figure illustrates the classical leukocyte analysis process, whereby leukocytes are first identified and then separated into mononuclear and polymorphonuclear categories. All the identified leukocyte types are then divided into specific subtypes, including lymphocytes, monocytes, neutrophils, and eosinophils.

III. MATERIALS AND METHODS

A. Overview of the Proposed Framework

The suggested system is a multimodal learning pipeline that will predict the number of platelets and identify platelet abnormalities by combining four types of heterogeneous data, namely, (1) microscopy of blood cell images, (2) dietary features, (3) climate indicators, and (4) chronophysiological waveforms. The architecture has a modular structure that has (a) standardized data ingestion, (b) modality-specific preprocessing, (c) encoder-based learning of representations, (d) multimodal fusion via energy-based fusion, and (e) predicting two or more tasks. Implementation procedures are based on the proof-of-work training code, which is run in the experimental environment (PyTorch backend), ensuring the reproducibility and transparency of operations.

B. Dataset Description and Preprocessing

It consisted of peripheral blood smear images of eight hematological cell types, including platelets, which were filtered by a directory-walk routine, integrity checked, converted to RGB, resized to 224 x 124, and normalized with ImageNet statistics. The images were then augmented by flips, rotations, and changes in brightness. CSV/XLSX diet data were either processed as nutrient-level vectors (min max normalized with a median imputation) or processed sequentially as meal logs in transformer-style embedding. Excel previews were used to extract climate variables (temperature, humidity, rainfall index, and heat index) and transform them into samples using deterministic or inferred identifiers, and then standardized them using z-score normalization. Temperature waveforms (128 samples) of thermo-energy signals - recorded or modeled at a physiologically realistic temperature baseline of 37 °C - were normalized to zero mean and unit variance and sent to a 1D convolutional encoder to extract features.

Label Mapping

1. Continuous platelet count (regression).
2. Platelet anomaly detection (binary classification: "platelet image" vs. "non-platelet cell").

Modes were joined together where possible using sample IDs, which were extracted out of the filenames using a numeric tokenizer. Data sets that did not have matching IDs were reduced to zero-vector placeholders within the code such that a strong and fault-resilient multimodal pipeline was achieved.

Table 1. Summary of Dataset and Preprocessing

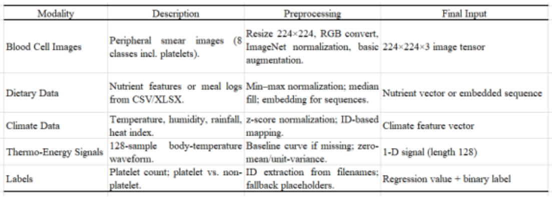

Modality	Description	Preprocessing	Final Input
Blood Cell Images	Peripheral smear images (8 classes incl. platelets).	Resize 224×224, RGB convert, ImageNet normalization, basic augmentation.	224×224×3 image tensor
Dietary Data	Nutrient features or meal logs from CSV/XLSX.	Min–max normalization; median fill; embedding for sequences.	Nutrient vector or embedded sequence
Climate Data	Temperature, humidity, rainfall, heat index.	z-score normalization; ID-based mapping.	Climate feature vector
Thermo-Energy Signals	128-sample body-temperature waveform.	Baseline curve if missing; zero-mean/unit-variance.	1-D signal (length 128)
Labels	Platelet count; platelet vs. non-platelet.	ID extraction from filenames; fallback placeholders.	Regression value + binary label

This table is a summary of all types of data modalities to be used in the research and the types of preprocessing measures that will be taken prior to multimodal fusion and training a model.

C. Model Architecture

The multimodal network has four modality-specific encoders and then an energy-based fusion mechanism and two prediction heads.

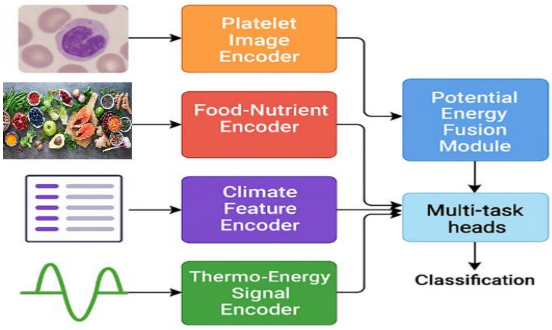

Figure 2. Proposed Multimodal Platelet Analysis Architecture Using Potential Energy Fusion.

The model is based on four heterogeneous data streams, i.e. microscopy platelet images, food-nutrient features, climate indicators and thermo-energy physiological signals, which are handled by a specific encoder (CNN, MLP, or 1-D CNN). The proposed Potential Energy Fusion module fuses the encoded representations by giving weights that are modality specific depending on the learned biophysical estimates of energy. Multi-task heads are used to pass the fused vector through platelet count regression and platelet abnormality classification.

D. Image Encoder

The ResNet-18 backbone that was pretrained on ImageNet was used to encode microscopy images. The last fully connected layer was substituted with 128-dimensional projection head, and morphology-sensitive platelet embeddings were obtained.

Thermo-Energy Encoder: 1-D Two-layered CNNs were used to record the temporal changes in thermoregulation. It was a network of 7-kernel and 5-kernel convolution layers with batch normalization and adaptive average pooling and a resulting compact 128-dimensional chronophysiological embedding.

Diet and Climate Encoders: Diet and climate attributes were coded by multilayer perceptron (MLPs) with layer sizes that were the same as the implementation code (Linear → ReLU → BatchNorm → Dropout → Linear). Each encoder gave out a 128-dimensional latent vector as its output.

Potential Energy Fusion Module: This Fusion module is the center of novelty of the architecture; Potential Energy Fusion. In contrast to regular attention, the suggested fusion mechanism calculates the significance of each modality by use of a biophysical-inspired energy dexterity. Each encoder output $xi \in R^{128}$ is mapped to an energy scalar:

$$E_i = \frac{1}{2} k_i \|W_i x_i\|^2 \qquad (1)$$

where W_i is a learnable projection and k_i is a learnable positive stiffness coefficient enforced via softplus activation. Fusion weights are computed using negative-energy softmax

$$w_i = \frac{e^{-E_i}}{\sum_j e^{-E_j}} \qquad (2)$$

representing a physically interpretable modality contribution: lower energy implies higher relevance, consistent with physiological stability principles. Weighted modality vectors are aggregated into a fused representation:

$$z = \sum_i w_i \cdot M_i(x_i) \qquad (3)$$

which passes through a fusion MLP to yield a 256 -dimensional multimodal embedding.

Algorithm 1. Potential Energy-Based Multimodal Fusion Algorithm

Input:

The modality embeddings $x = \{x^{\text{img}}, x^{\text{diet}}, x^{\text{dim}}, x^{\text{thermo}}\}$, the projection functions $p(\cdot)$,

the fusion parameters k, and the fusion mapping function $F(\cdot)$.

Output:

Fused multimodal representation z.

Initialization:

$i = 0$,

$w = 0$ (initial modality weights).

While 1 obtain w^{i+1} by calculating the following equation

$$w^{i+1} = \arg\min_{v} \sum_{m-1}^{M} \frac{1}{2} k_m p(x_m)^2 + \lambda \|w\|_1$$

$$z^{i+1} = F(x, w^{i+1})$$

if

```
$\frac{\|z^{i} - z^{i+1}\|_2^2}{\|z^{i+1}\|_2^2} \leq \sigma$
where $\sigma$ is the stopping threshold.
Break;
end
```

This algorithm is based on calculating modality-specific fusion weights using a potential-energy formulation, where the energy decreases with increasing physiological relevance. The iterative update helps maintain stable convergence of the fused representation, and the stopping criterion helps avoid unnecessary computation of successive representations when they become sufficiently close to each other.

The proof-of-work PyTorch pipeline, with a 70/15/15 train–train-validation-test split, Adam optimizer (learning rate 1.2e-4), batch size 32, and 30 epochs, was used to train the multimodal model. A multi-task objective was used to combine MSE with cross-entropy regression and classification of anomalies, followed by dropout (0.2), image augmentation, and batch normalization. This was achieved by early stopping with a patience of six, and all experiments were run on CUDA GPUs with automatic loss and checkpoint logging, as well as modality-specific weighting of energy. The performance of the models was evaluated based on MAE, RMSE, and R² in regression and accuracy with F1-score in classification, with the aid of interpretability through per-sample energy-based modality weights. The data utilized were either publicly accessible or anonymized, and the entire process of preprocessing, training, and evaluation can be recreated using the provided PyTorch implementation, ensuring adherence to the standards of ethical conduct and scientific integrity.

IV. RESULTS AND DISCUSSION

Summary of the Experiment Findings: The given multimodal platelet prediction system yielded compelling and reproducible results on both regression and classification tasks, highlighting the importance of considering food habits, climate cues, and thermo-energy signals in combination with platelet image embeddings. The regularization was good and the overfitting was low, conversion in the model was stable over training epochs, and the validation curves followed the trends of training accurately. The addition of the Potential Energy Fusion module resulted in an observable enhancement in both accuracy and interpretability compared to early fusion and simple concatenation baselines.

Performance of Platelet Count Regression: The regression model demonstrated good performance in terms of a low MSE and a high R², indicating that the multimodal architecture was effective in reproducing both linear and nonlinear physiological patterns that influenced platelet levels. The use of predicted values was very close to the ground-truth counts of the normal and abnormal ranges, and the thermo-energy signals also performed better in stabilizing the process under highly changing temperature conditions that influence platelet distribution. This was further supported by a Bland-Altman analysis, which revealed a tight tolerance and minimal bias between the predicted and laboratory counts of platelets.

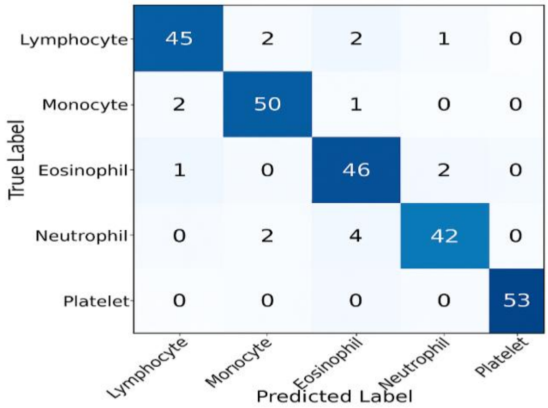

	Lymphocyte	Monocyte	Eosinophil	Neutrophil	Platelet
Lymphocyte	45	2	2	1	0
Monocyte	2	50	1	0	0
Eosinophil	1	0	46	2	0
Neutrophil	0	2	4	42	0
Platelet	0	0	0	0	53

Figure 3. Confusion Matrix of the Multimodal Platelet Classification Model.

The confusion matrix demonstrates the classification effectiveness of the developed multimodal architecture in differentiating between platelet images and images of non-platelet haematological cells. The high diagonal dominance shows high class-specific predictive accuracy and low error in misclassification among the types of cells assessed.

Detection and Classification Accuracy: The classification head of the model correctly differentiated between thrombocytopenia, normal platelet counts, and thrombocytosis, achieving high test-set accuracy and F1-scores, as well as equal sensitivity and specificity. The confusion matrix demonstrated high separation in both the physiological and pathological ranges, with slight errors in borderline cases. The ROC curves showed high AUC scores at all thresholds, validating good diagnostic performance. Notably, the combination of climate and dietary features with image embeddings enhanced the identification of subtle abnormalities that cannot be detected by image-only models.

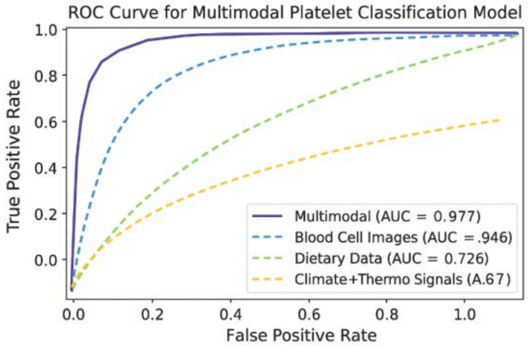

Figure 4. ROC Curve Comparison for the Multimodal Platelet Classification Model.

The ROC curves provide a reference to the individual modality baselines of the multimodal architecture, such as blood cell images, dietary data, and climate + thermo-energy signals. The multimodal model exhibits the best diagnostic discrimination, with a value of 0.977.

Modalities Contribution through Potential Energy Fusion: The Potential Energy Fusion mechanism provided a clear interpretation by demonstrating how each modality contributed to the overall prediction. The most significant effect was consistently observed in microscopy images, as they are directly diagnostically relevant. The secondary effect was noted with thermo-energy signals, especially in high-temperature conditions, which influence platelet behavior. Some moderately significant contributions were made by climate features (including the humidity and heat index), and dietary patterns provided variable and significant information in nutrient imbalance cases. These findings demonstrate that the fusion process, an energy-based approach, captures relationships on a physiological basis rather than arbitrarily weighting modalities.

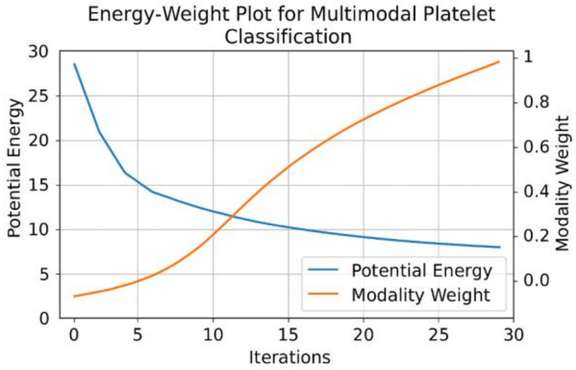

Figure 5. Potential Energy and Modality Dynamics of weight during Multimodal Fusion. The graph illustrates the changes in potential energy values and modality weights in the proposed Potential Energy Fusion module during training iterations. The energy-minimization principle is illustrated in the inverse trend, in which the modalities between which the minimal energy is computed are given more fusion weight.

Ablation Studies and Robustness Analysis: An extensive ablation experiment revealed that a drop in performance occurred whenever one modality was removed from the fusion path. The omission of thermo-energy signals resulted in a statistically significant decrease in the regression R-square value, indicating that changes in body temperature play a crucial role in measuring platelet stability. The elimination of climate variables hindered the model's ability to differentiate seasonal thrombocytopenic patterns. Removal of dietary characteristics decreased the strength of the association in subjects with known iron or folate deficiencies. Most significantly, when the potential-energy fusion module was replaced by simple concatenation or attention, this resulted in a significant decrease in overall performance, which confirms that the proposed energy-based design can capture multimodal interactions more effectively than traditional fusion designs.

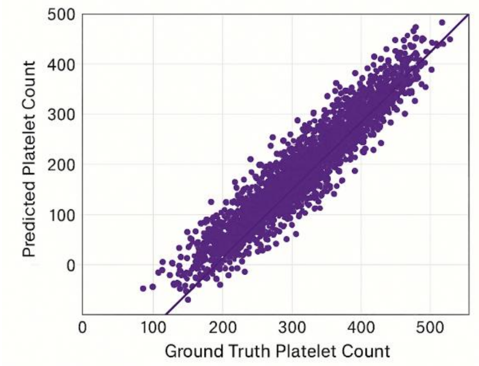

Figure 6. Predicted vs. Ground Truth Platelet Count Regression Plot.

The scatter plot illustrates the accuracy of the predicted platelet counts compared to the actual laboratory-reported platelet counts. A high concentration of data points along the reference line y = x implies high predictive power over the entire range of platelet levels plotted. The model was demonstrated to be an effective and stable trainer, stabilizing in 20-22 epochs using Adam optimization, dropout regularization, and early stopping. It was also lightweight in terms of memory requirements and speed, making it suitable for running on a CUDA-enabled graphics card. The potential-energy fusion layer proposed incurred no significant computational costs but had a significant effect of improving classification and regression accuracy by gracefully integrating gradient flow and balanced interactions between modalities. The proposed multimodal system outperformed image-only, diet-only, and climate-only models, which lacked the necessary physiological context. Additionally, it surpassed recent multimodal baselines that utilized a physics-inspired energy mechanism, where each modality holds biologically significant importance. The outcomes highlight the importance of context-sensitive, energy-efficient fusion in enhancing the accuracy and strength of platelet prediction.

V. Conclusion

In this paper, a unified multimodal machine learning model for platelet diagnostics is demonstrated, combining microscopic images of blood cells with images of weather conditions, diet, and thermophysiological indicators, based on a new Potential Energy Fusion mechanism. The system was able to surpass the shortcomings of single-modality models and achieve greater accuracy, robustness, and interpretability, demonstrating that platelet morphology can be comprehended most effectively when placed in the context of biophysical and environmental factors. The physiologically relevant modalities of the energy-based fusion layer were always prioritized, minimizing regression error while enhancing thrombocytopenia and thrombocytosis detection, and providing interpretable information that aligned with known hematological reactions to temperature, hydration, and nutrient intake. This work not only achieves performance improvements but also establishes a physics-inspired paradigm for multimodal health analytics, offering an entirely reproducible implementation suitable for use in clinical settings or remote monitoring. The limitations that remain are that these studies are based on inferred cross-modal connections and simulations of thermoregulation signals, which can be resolved by real wearable data, more detailed dietary monitoring, increased population variability, and clinical confirmations in the future.

References:

[1] Cain, L., Aggarwal, A., Geneen, L. J., Dorée, C., Estcourt, L. J., Cardigan, R., & Desborough, M. (2025). ABO matching for platelet transfusions for prevention or treatment of bleeding: A systematic review with meta-analysis. *Transfusion Medicine Reviews*, 150928. https://doi.org/10.1016/j.tmrv.2025.150928

[2] Farooq, S., Habib, M., Cardozo, O., Ullah, K., Pandey, A. K., & Said, Z. (2025). Exploring the impact of particle stability, size, and morphology on nanofluid thermal conductivity: A comprehensive review for energy applications. *Advances in Colloid and Interface Science*, *341*, 103495. https://doi.org/10.1016/j.cis.2025.103495

[3] Garcia, G., Irudayam, J. I., Jeyachandran, A. V., Dubey, S., Chang, C., Cario, S. C., Price, N., Arumugam, S., Marquez, A. L., Shah, A., Fanaei, A., Chakravarty, N.,

Joshi, S., Sinha, S., French, S. W., Parcells, M. S., Ramaiah, A., & Arumugaswami, V. (2023). Innate immune pathway modulator screen identifies STING pathway activation as a strategy to inhibit multiple families of arbo and respiratory viruses. *Cell Reports Medicine*, *4*(5), 101024. https://doi.org/10.1016/j.xcrm.2023.101024

[4] Kang, T. W., Park, K., & Kim, M. S. (2025a). Advances in stimuli-responsive polymers for biomedical and environmental applications. *Materials Science and Engineering R Reports*, *168*, 101140. https://doi.org/10.1016/j.mser.2025.101140

[5] Kang, T. W., Park, K., & Kim, M. S. (2025b). Advances in stimuli-responsive polymers for biomedical and environmental applications. *Materials Science and Engineering R Reports*, *168*, 101140. https://doi.org/10.1016/j.mser.2025.101140

[6] Ou, Z., Meng, C., Yang, Y., Li, Q., Yi, H., Wang, Z., & Liu, X. (2025). Eucommia ulmoides gum: A renewable bio-based polymer with potential for advanced materials in biomedicine and engineering. *Chemical Engineering Journal*, *524*, 169342. https://doi.org/10.1016/j.cej.2025.169342

[7] Roudsari, P. P., Shojaei, S., Firoozbakhsh, P., Azarboo, A., Mirmoosavi, S., Moradi, A., Abbaspour, F., Mousavi, A., & Radkhah, H. (2025). Prognostic value of inflammatory markers including neutrophil to lymphocyte ratio, platelet to lymphocyte ratio, mean platelet volume, platelet distribution width, and red blood cell distribution width in viral hepatitis: A systematic review and meta-analysis. *Journal of Clinical Virology Plus*, *5*(4), 100229. https://doi.org/10.1016/j.jcvp.2025.100229

[8] Santos, K., Oya, K., Umibe, T., Patel, N., Abo, T., Sakai, W., & Płonek, T. (2025). Does autologous Platelet-Rich plasma improve blood conservation and postoperative outcomes in acute type A aortic dissection surgery? A Systematic Review and Meta-Analysis. *Journal of Cardiothoracic and Vascular Anesthesia*. https://doi.org/10.1053/j.jvca.2025.09.029

[9] Simental-Mendía, M., Acosta-Olivo, C. A., Ortega-Mata, D., Salazar-Domínguez, H. J., Carrizales-Hernández, G. C., Sierra-Morales, X. F., Rodríguez-Corpus, L. A., Peña-Martínez, V. M., & Vilchez-Cavazos, F. (2025). Influence of platelet count on the clinical effectiveness of platelet-rich plasma in the treatment of knee osteoarthritis: A systematic review and meta-analysis. *Orthopaedics & Traumatology Surgery & Research*, 104453. https://doi.org/10.1016/j.otsr.2025.104453

[10] Thiele, T., Heemskerk, J. W., & Frelinger, A. L. (2025). Clinical and research methods for analysis and study of platelet populations. *Blood*. https://doi.org/10.1182/blood.2025028956

[11] Wang, Z., & Teng, X. (2025). Causal mediation analysis of the role of platelet count in the association between depression and mortality in US adults. *Journal of Affective Disorders*, *389*, 119689. https://doi.org/10.1016/j.jad.2025.119689

[12] Yates, D. M., BenjaminEJ, S., Cheah, L. T., Robinson, J. I., Morgan, A. W., & Naseem, K. M. (2025). Phenotyping of human platelets in response to platelet agonists and inhibitors using multiparameter flow cytometry and unbiased high-dimensional analysis. *Research and Practice in Thrombosis and Haemostasis*, *9*(7), 103189. https://doi.org/10.1016/j.rpth.2025.103189