Check for updates

ACM DIGITAL LIBRARY | Association for Computing Machinery | acm open

Latest updates: https://dl.acm.org/doi/10.1145/3774816.3774832

RESEARCH-ARTICLE

Hierarchical Global-To-Local Feature Selection Architecture for HDLSS Datasets: A Computational Framework for Personalised Medicine in Oncology

MOSTAFA MOHIUDDIN JALAL, University of Technology Sydney, Sydney, NSW, Australia

PAUL J KENNEDY, University of Technology Sydney, Sydney, NSW, Australia

DANIEL ROBIN CATCHPOOLE, University of Technology Sydney, Sydney, NSW, Australia

Open Access Support provided by:

University of Technology Sydney

PDF Download
3774816.3774832.pdf
17 March 2026
Total Citations: 0
Total Downloads: 54

Published: 23 February 2026

Citation in BibTeX format

HIKM '25: Health Informatics Knowledge Management Conference 2025
September 16 - 17, 2025
Online, Australia

https://doi.org/10.1145/3774816.3774832
ISBN: 9798400715815

Hierarchical Global-To-Local Feature Selection Architecture for HDLSS Datasets: A Computational Framework for Personalised Medicine in Oncology

Mostafa Mohiuddin Jalal*
Children's Cancer Research Unit, Kids Research
Sydney Chilren's Hospitals Network
Westmead, New South Wales
Australia
School of Computer Science
University of Technology Sydney
Ultimo, NSW, Australia
MostafaMohiuddin.Jalal@uts.edu.au

Paul J Kennedy
Australian AI Institute, Faculty of Engineering & IT
The University of Technology Sydney
Ultimo, NSW, Australia
Biomedical Data Science Lab, School of Computer Science, FEIT
The University of Technology Sydney
Ultimo, NSW, Australia
paul.kennedy@uts.edu.au

Daniel Robin Catchpoole
Biomedical Research, Children's Cancer Research Unit
Kids Research, Sydney Children's Hospital Network
Westmead, NSW, Australia
Biomedical Data Science Lab, School of Computer Science, FEIT
The University of Technology Sydney
Ultimo, NSW, Australia
daniel.catchpoole@uts.edu.au

Abstract

Understanding individual patient characteristics is essential for personalized medicine. Computational approaches can help to identify patient-specific patterns similar to how clinicians learn from individual cases. Instance-based learning strategies are promising for this purpose, but applying them effectively to genomic tabular data remains challenging. Our work addresses this gap by developing a dynamic feature selection framework that identifies the most relevant genomic markers for each individual patient, rather than relying solely on population-level important features. While genomic datasets typically present High Dimension Low Sample Size (HDLSS) challenges where features vastly outnumber patients, our approach overcomes this limitation through a multi-metric feature selection integrated with a neural network classifier. The system employs a continuous feedback mechanism that refines feature relevance based on prediction outcomes, mimicking how clinicians iteratively improve their diagnostic understanding. In experiments with genomic data, our approach successfully identified both population level important features and patient-specific markers that conventional methods often overlook. This framework has significant clinical implications by enhancing treatment personalization through improved outcome explainability. It works as an aid to better understand the unique factors driving individual patient outcomes.

CCS Concepts

• **Computing methodologies** → **Feature selection**.

*Corresponding author.

This work is licensed under a Creative Commons Attribution 4.0 International License.
HIKM 2025, online, Australia
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1581-5/25/09
https://doi.org/10.1145/3774816.3774832

Keywords

Feature selection, adaptive thresholding, dimensionality reduction, machine learning, neural network, personalized outcome, model interpretability

ACM Reference Format:
Mostafa Mohiuddin Jalal, Paul J Kennedy, and Daniel Robin Catchpoole. 2025. Hierarchical Global-To-Local Feature Selection Architecture for HDLSS Datasets: A Computational Framework for Personalised Medicine in Oncology. In *Health Informatics Knowledge Management Conference 2025 (HIKM 2025), September 16–17, 2025, online, Australia.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3774816.3774832

1 Introduction

Understanding individual patient characteristics is central to personalized medicine. Just as clinicians learn from individual cases, computational approaches can identify patient-specific patterns crucial for effective treatment decisions. Instance-based learning strategies show promise in oncology where treatment responses vary significantly across patients. Despite this potential, applying IBL effectively to genomic tabular data remains challenging, especially when distinguishing universally important features from those significant only for specific patients.

ProtoGate, introduced by Jiang et al.[5], represents a significant advancement in this direction, combining global and instance-wise feature selection through a prototype-based neural network. It employs a disjoint selection method with a trainable feature selector and a non-trainable KNN-based predictor, using L1-regularization for global selection and L0-regularization for instance-wise sparsity. This approach addresses the co-adaptation problem, where trainable selectors and predictors might jointly learn to achieve high accuracy with low-fidelity features.

While promising, ProtoGate lacks critical mechanisms for clinical applications. It cannot dynamically adjust feature relevance based on prediction outcomes, limiting its adaptability. Its single-metric approach to global selection inadequately captures complex genomic patterns. Importantly, ProtoGate does not explicitly identify which features drive predictions, making interpretation difficult

for further validation. Additionally, it was not optimized for the severe class imbalance common in oncology datasets.

Our work enhances ProtoGate with three significant technical improvements. First, we implement a continuous feedback mechanism that dynamically refines feature relevance based on prediction outcomes, mimicking how clinicians iteratively improve their diagnostic understanding for each individual patients. Second, we replace ProtoGate's single-metric approach with a multi-metric selection system combining mutual information, random forest importance, and lasso regression. Third, we integrate components optimized for class imbalance through a Focal Loss function.

Our approach integrates feature selection and classification within a single computational graph while explicitly identifying and providing selected feature sets—both globally important and patient-specific—with validation metrics explaining each selection. Instead of binary masks, our system forwards weighted feature sets to the classifier, providing more nuanced representations of feature importance while maintaining safeguards against co-adaptation problems.

The modified framework provides more granular insights into which features drive predictions for specific patients, enhancing interpretability for further computational analysis.

The contributions of our work include: technical enhancements that improve ProtoGate's suitability for genomic data in HDLSS contexts, and a continuous feedback mechanism that refines feature selection over time. The remainder of this paper covers related work in instance-specific feature selection (Section 2), details our technical modifications (Section 3), presents experimental results (Section 4), discusses our contributions and limitations (Section 5), and concludes with a summary and future directions (Section 6). Through these advancements, we aim to bridge the gap between computational methods and individualized interpretability in patient genomic analysis.

2 Related Work

Analyzing High Dimension Low Sample Size (HDLSS) genomic datasets presents significant challenges due to the "curse of dimensionality," where features vastly outnumber samples. This paradigm is particularly prevalent in oncology, where tumors exhibit substantial molecular heterogeneity across relatively few patient samples. Effective feature selection becomes critical, as traditional machine learning approaches often overfit when features greatly outnumber observations. Recent research has increasingly focused on instance-based learning approaches to improve both interpretability and predictive accuracy in these challenging contexts.

Jiang et al.[5] introduced ProtoGate, a prototype-based neural network specifically designed for HDLSS datasets that combines l1-based global selection with l0-based instance-wise sparsity. This global-to-local approach is particularly well-suited for genomic data, where certain genes may be broadly relevant while others are significant only for specific patient subgroups. ProtoGate's disjoint learning paradigm prevents co-adaptation between selector and predictor, while its prototype-based non-parametric prediction aligns with the clustering assumption often valid in biomedical data. Despite these innovations, ProtoGate lacks mechanisms to dynamically adjust feature relevance based on prediction outcomes, limiting its adaptability in evolving clinical scenarios.

Several approaches have focused explicitly on instance-wise feature selection. Liu et al.[7] developed DIWIFT, which employs influence functions to measure how features impact validation loss for specific instances. While effective at identifying patient-specific features, DIWIFT does not dynamically adapt selection criteria during model operation. Masoomi et al.[8] proposed instance-wise feature grouping based on representation and relevant redundancies defined using information theory. Their method discovers instance-wise feature groups that are redundant internally but collectively relevant for prediction. However, their approach relies on fixed groupings that limit adaptability in dynamic clinical contexts. Panda et al.[9] emphasized causal rather than correlative feature identification, enhancing precision by selecting features based on their causal effect on model outputs. Though innovative, their method lacks mechanisms to adjust relevance based on outcomes across time.

Attention-based approaches have shown promise for interpretable feature selection. Arik and Pfister[13] developed TabNet, an attentive interpretable architecture for tabular data that uses sequential attention to choose features at each decision step. This enables both interpretability and efficient learning as the model capacity is used for the most salient features. While effective for general tabular data, TabNet was not optimized for HDLSS conditions or severe class imbalance common in oncology. Yasuda et al.[20] proposed Sequential Attention, which selects features one-by-one using attention scores in a greedy manner. Though theoretically connected to Orthogonal Matching Pursuit, this approach faces scalability challenges with extensive feature spaces typical in genomic data.

Integrative approaches combine classical methods with neural networks. Lemhadri et al.[6] introduced LassoNet, extending LASSO principles to neural networks through a skip connection from input to output with a hierarchical constraint: a feature can only be used in hidden layers if its skip-connection weight is non-zero. This enables feature sparsity to be directly controlled by the skip-layer weights. While effective for global feature selection, LassoNet overlooks patient-specific variations crucial for precision medicine. Song and Xiao[15] investigated variable selection with false discovery rate control in neural networks, making significant contributions to statistical feature selection through their SurvNet approach, though primarily addressing global rather than instance-specific selection patterns.

Several specialized methodologies target specific aspects of feature selection. Peng et al.[11] incorporated Gaussian copulas into instance-wise selection to model feature dependencies, addressing a limitation in prior methods that assumed feature independence. Imrie et al.[3] developed composite feature selection using deep ensembles to discover groups of interacting features without predefined groupings. Their approach trains multiple selection models that are encouraged to diversify their feature selections, revealing complementary predictive structures. Hassanieh et al.[2] proposed selective deep autoencoders for unsupervised feature selection, selecting features from the original set that can reconstruct the full feature space without predefining the number of features. These

approaches, while innovative, lack mechanisms for continuous refinement based on prediction outcomes.

For high-dimensional biological data specifically, Singh et al.[14] introduced FsNet, a DNN-based selection method using tiny predictor networks to generate large virtual weight matrices, preventing overfitting in biological datasets with low sample sizes. Their approach reduces parameter count from O(dK) to O(bK), making it computationally feasible for genomic applications. Yang et al.[19] presented locally sparse neural networks that learn sample-specific sparsity for tabular biomedical data. Their gating architecture determines which features should be "active" per instance, enhancing interpretability by providing per-instance explanations. While addressing individual-level sparsity, this approach inadequately handles class imbalance typical in oncology datasets where positive cases (e.g., metastasis, recurrence) are often substantially outnumbered by negative cases.

Dynamic selection approaches include VFDS by Ardywibowo et al.[1], a Bayesian framework for foresight dynamic selection that learns a policy to select features before observing them, based on previous context. VFDS optimizes a variational Bayesian objective that characterizes the trade-off between model performance and feature cost. Though innovative for sensor-based applications, VFDS was not specifically optimized for HDLSS genomic conditions. Sristi et al.[16] proposed conditional stochastic gates for contextual feature selection, enhancing interpretability through context-specific selection, but lacking dynamic feedback loops for continuous refinement based on clinical outcomes.

Several researchers have critically evaluated feature selection approaches. Passemiers et al.[10] conducted comprehensive benchmarks of neural network feature selection methods, revealing limitations in detecting non-linear signals in high-dimensional, noisy contexts—precisely the conditions encountered in genomic data. Wojtas and Chen[17] explored feature importance ranking for deep learning, though their dual-network architecture lacks patient-level granularity. Yamada et al.[18] proposed stochastic gates for embedded feature selection, effectively addressing feature redundancy but not specifically optimized for HDLSS conditions. Tonekaboni et al.[12] assessed interpretability challenges in time-series models, underscoring the importance of dynamic refinement mechanisms in feature selection methodologies.

Yoshikawa et al.[21] introduced neural generators of sparse local linear models, combining neural networks' predictive power with linear models' interpretability. Their framework uses a DNN-based weight generator to create dense, sample-specific linear weights, with a K-Hot Gate module enforcing sparsity. While enhancing sparse interpretability, their approach lacks explicit adaptive strategies for patient-specific profiles. Critically, Jethani et al.[4] identified "prediction leakage" in joint amortized explanation methods, where selectors encode predictions rather than identifying truly informative features. Their work introduced EVAL-X and REAL-X to address these issues, highlighting the need for disjoint learning approaches and evaluation methods that detect such failures.

Our approach addresses these limitations by enhancing ProtoGate with three significant improvements: (i) a continuous feedback mechanism that dynamically refines feature relevance based on prediction outcomes; (ii) an enhanced multi-metric selection system balancing population-level and patient-specific markers; (iii) and neural network integration optimized for class imbalance. These advancements enable dynamic feature refinement and interpretable patient-specific analysis, resulting in a powerful framework uniquely suited for personalized medicine in oncology, where both common pathways and rare variants contribute to disease progression and treatment response.

3 Methodology

We propose a hierarchical feature selection framework tailored explicitly for genomic datasets characterized by high dimensionality and small sample sizes. Our approach significantly extends ProtoGate's global-to-local paradigm by incorporating three novel methodological advancements described above. Collectively, these innovations address key limitations identified within ProtoGate, resulting in enhanced robustness, greater clinical relevance, and improved predictive accuracy.

3.1 Consensus-Based Multi-Metric Selection

ProtoGate utilizes linear regularization, emphasizing sparsity and simplicity at the expense of capturing biologically relevant, complex feature interactions. However, genomic datasets typically exhibit intricate, nonlinear relationships that are critical for accurate outcomes yet remain undetected under strictly linear selection methods. Additionally, these datasets commonly exhibit severe class imbalance—where one class (e.g. good responders) vastly outnumbers another (e.g. poor responders)—making accurate prediction particularly challenging for minority class . To overcome this limitation, we introduce a consensus-based multi-metric feature selection strategy that integrates mutual information, random forest importance, and sparse linear relationships identified by LASSO regression.

$$\text{Score}_{\text{consensus}} = w_1 \cdot \text{MI} + w_2 \cdot \text{RF} + w_3 \cdot \text{LASSO}$$

MI quantifies mutual information between features and outcomes, capturing both linear and non-linear dependencies; RF represents random forest importance scores for hierarchical relationships; and LASSO identifies sparse linear relationships. This weighted combination provides a more comprehensive assessment than any single metric.

Features selected through this consensus approach undergo rigorous stability assessments via repeated bootstrap resampling. By retaining only those features that demonstrate consistent predictive relevance across subsets, our methodology ensures selection robustness and reproducibility, significantly improving ProtoGate's single-dimensional selection criterion. This integration of complementary metrics broadens feature detection capabilities, capturing richer biological insights crucial for clinically meaningful outcomes. The overall workflow of this consensus-based multi-metric feature selection framework is illustrated in Figure 1.

3.2 Patient-Specific Feature Adaptation

ProtoGate employs fixed, uniform feature-selection thresholds across all patients, disregarding individual genetic heterogeneity and class-specific considerations, especially problematic in datasets with significant class imbalance typical in clinical genomics. Such rigidity potentially overlooks critical patient-specific biomarkers

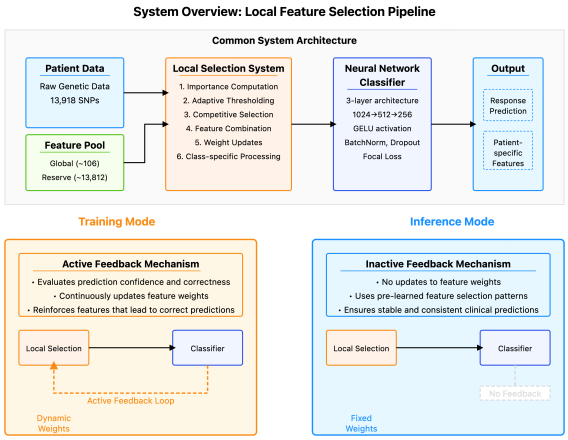

Figure 1: Overview of the hierarchical feature selection framework.

that deviate from population-level averages, severely limiting clinical applicability.

Further refining our method, we implement competitive selection among reserve features, adaptively ranking candidate features uniquely per patient to maintain computational efficiency while capturing the most informative, patient-specific markers. Features not meeting adaptive criteria undergo controlled attenuation rather than outright removal, thus preserving system-wide information without introducing excessive noise. These adaptive strategies explicitly overcome ProtoGate's inflexible, generalized selection methods, enabling precise and personalized genomic insights aligned with clinical decision-making needs.

3.3 Continuous Learning Mechanism

A significant shortcoming of ProtoGate's original design is the strict separation between feature selection and classification components, resulting in a static learning environment unable to iteratively refine predictive capabilities based on new clinical evidence. To resolve this limitation, our framework integrates continuous feedback-driven learning, directly coupling feature selection updates to classification outcomes.

Our proposed feedback mechanism dynamically adjusts feature importance weights based on prediction accuracy, leveraging a momentum-based update strategy that robustly integrates accumulated patterns with new evidence. Importantly, updates are gated by prediction confidence, ensuring that only high-confidence predictions with reliable outcomes influence feature recalibration. Such an adaptive approach closely mirrors clinical diagnostic processes, where understanding and relevance continually evolve with accumulating patient information. The architectural flow of this adaptive learning system, highlighting the continuous feedback loop between the feature selection and classification modules, is presented in Figure 2, which demonstrates the integrated patient-specific feature selection framework.

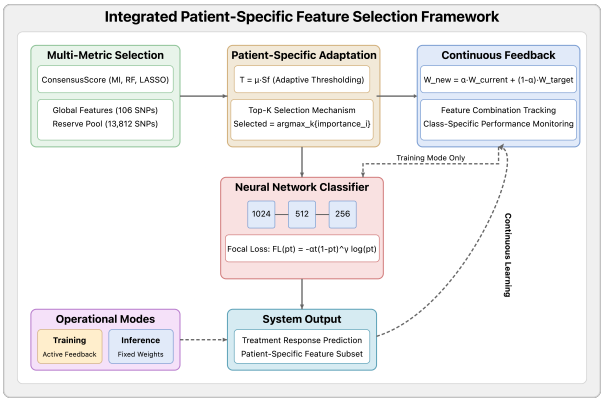

Figure 2: Integrated Patient-Specific Feature Selection Framework showing the three key components: multi-metric selection (left), patient-specific adaptation (center), and continuous feedback mechanism (right), connected to a neural network classifier with focal loss function for addressing class imbalance.

Additionally, our system actively monitors interactions between genomic features, employing adaptive tracking mechanisms to detect emerging synergistic combinations. By systematically recognizing and reinforcing these cooperative relationships over time, the continuous feedback loop provides a dynamic, real-time perspective on genomic interactions essential for accurate clinical predictions—capabilities entirely absent in ProtoGate's static approach.

3.4 Neural Architecture and Implementation

Our methodology is computationally realized through an integrated neural network architecture designed explicitly to facilitate end-to-end optimization between feature selection and classification processes. To address class imbalance, we incorporate a focal loss strategy that dynamically adjusts prediction contributions based on class-specific difficulties. This targeted loss formulation ensures balanced attention across diverse patient populations. Operationally, our framework maintains distinct training and inference phases. During training, feature selection mechanisms dynamically adapt in concert with classifier parameters through active feedback, while inference employs deterministic application of learned selection patterns, ensuring consistent and clinically interpretable results.

4 Experimental Results

We evaluate our framework through a systematic progression across three genomic datasets, each presenting distinct computational challenges that validate different aspects of our approach. This design progresses from discrete categorical values to continuous signatures, concluding with controlled synthetic validation.

4.1 Discrete categorical - ALL SNP Dataset

The Acute Lymphoblastic Leukemia (ALL) Single Nucleotide Polymorphisms(SNP) dataset presents computational challenges characteristic of genomic variant data. This dataset, obtained from the Children's Hospital at Westmead, comprises 139 pediatric patients with 13,918 germline SNP features derived from Illumina microarray technology. This dataset exhibits severe class imbalance (84% good responders, 16% poor responders).

4.1.1 Computational Characteristics. SNP data exhibits tri-modal clustering around discrete values 0, 0.5, 1.0, corresponding to homozygous reference, heterozygous, and homozygous alternative alleles respectively. This discrete distribution creates discontinuous feature spaces that challenge optimization algorithms designed for continuous variables. The categorical nature of genetic variants renders intermediate values biologically meaningless, requiring feature selection methods to operate effectively within these constraints.

4.1.2 Performance Results. Our framework achieves a 90. 9% accuracy and a F1 score of 0.66, compared to ProtoGate's 41. 3% accuracy and an F1 score of 0.66. The local feature selection analysis reveals significant differences in adaptability: our method selects 39-63 features per patient, while ProtoGate maintains a mean selection of approximately 3 features per patient. This increased adaptability correlates with improved minority class identification, which is critical for clinical applications requiring accurate identification of treatment-resistant cases.

4.2 Continuous Data - PAM50 Gene Expression

The METABRIC PAM50 gene expression dataset provides a computational environment with continuous feature distributions, enabling validation of our framework's generalizability beyond discrete optimization challenges. PAM50 (Prediction Analysis of Microarray 50) is a clinically validated 50-gene expression signature used to classify breast cancer into intrinsic molecular subtypes for prognostic and treatment decision-making. This dataset contains 200 patients with 4,160 gene expression features and maintains similar class imbalance (83.5% vs 16.5%).

4.2.1 Computational Characteristics. Gene expression data represents continuous biological processes with broader dynamic ranges than SNP data. Expression values span positive real numbers, providing smooth optimization landscapes that theoretically align with ProtoGate's original design assumptions. This dataset serves as a control to distinguish genuine algorithmic improvements from optimizations specific to discrete data distributions.

4.2.2 Performance Results. Our method achieves 84.6% accuracy and F1-score of 0.782, compared to ProtoGate's 91.6% accuracy and F1-score of 0.974. Local feature selection demonstrates substantial adaptability. Our method selects a range of 73-128 features per patient, while ProtoGate selects approximately 9.5 features per patient with minimal variation. These results indicate that our feedback mechanism provides consistent benefits across different data distribution types.

4.3 Controlled Synthetic Validation - Unbiased Benchmark

To systematically evaluate our feature selection algorithm under controlled conditions with known ground truth, we programmatically created a synthetic HDLSS dataset using Python. Unlike clinical datasets where true biological relevance may be ambiguous, this synthetic dataset eliminates dataset-specific confounders and provides definitive ground truth validation for feature selection accuracy. This controlled environment contains 400 samples with 2,000 features: 20 globally informative features, 100 locally informative features (20-30 selected per patient based on class), and 1,880 noise features.

4.3.1 Computational Characteristics. The synthetic dat-aset employs continuous value distributions similar to gene expression data while maintaining known ground truth for feature relevance. This design enables direct assessment of the accuracy of global versus patient-specific feature discovery under identical optimization conditions for both algorithms.

4.3.2 Performance Results. Performance advantages persist in this controlled environment, with our method selecting 17-28 local features with substantial inter-patient variability compared to ProtoGate's 21.95 mean selection with minimal variation.

As summarized in Table 1 our framework achieves higher predictive accuracy and lower loss across all datasets while maintaining competitive F1-scores relative to ProtoGate in most cases. Furthermore, the observed distribution of selected features across patients, visualized in Figure 3, reveals broader and more heterogeneous selection patterns under our method. This variability indicates enhanced adaptability to patient-specific genomic characteristics, aligning with the intended goal of individualized feature learning.

The consistency of these performance improvements across discrete categorical, continuous biological, and controlled synthetic datasets provides strong evidence that the improvements stem from inherent algorithmic robustness rather than dataset-specific tuning.

Table 1: Performance comparison across datasets. Our framework achieves higher accuracy and lower loss while maintaining competitive F1-scores compared to ProtoGate.

Dataset	Accuracy			F1 Score			Prediction Loss		
	Ours	ProtoGate	Δ	Ours	ProtoGate	Δ	Ours	ProtoGate	Δ
ALL SNP	0.909	0.413	**+0.496**	0.66	0.66	0	0.08	0.83	**-0.75**
Metabric PAM50	0.846	0.916	-0.07	0.782	0.974	-0.192	0.11	0.075	0.035
Synthetic	0.975	0.954	**+0.021**	0.974	0.954	**+0.02**	0.164	0.134	**-0.03**

4.4 Adaptibility Analysis

Figure 2 demonstrates the fundamental difference in local feature selection patterns between our framework and ProtoGate. Our approach exhibits high variability in feature selection counts (indicated by large intervals), suggesting adaptive responses to individual patient characteristics. This adaptability pattern remains consistent across all three datasets. The observed variability in local feature selection represents algorithmic responsiveness to patient-specific patterns rather than selection instability, as evidenced by improved predictive performance across all datasets.

This systematic evaluation across diverse data types establishes our framework's applicability to the full spectrum of genomic data

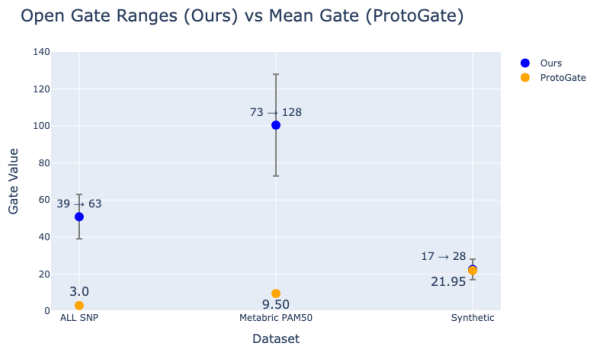

Figure 3: Local feature selection variability comparison. Our method (blue) exhibits broader selection ranges and higher variability compared to ProtoGate (orange), demonstrating enhanced adaptability to individual patient genomic patterns across all three datasets.

encountered in clinical decision-support systems, from static genetic risk assessment to dynamic molecular monitoring during treatment. The demonstrated capability to handle both discrete genetic variants and continuous molecular signatures positions our approach as a general-purpose solution for genomic-guided personalized medicine applications requiring patient-specific feature identification across heterogeneous data types.

5 Discussion

5.1 Key Technical Contribution

Our dynamic feature selection framework addresses two fundamental challenges in genomic learning: the HDLSS problem and severe class imbalance. The continuous feedback mechanism enables adaptive feature refinement during training, demonstrating measurable improvements in minority class performance (F1 = 0.66) while maintaining high overall accuracy (90.9%). This represents a significant advancement over static selection methods that cannot adapt to patient-specific patterns.

The two-stage architecture provides computational efficiency by reducing feature space from 13,918 to $\tilde{1}06$ globally selected features, followed by patient-specific refinement. The observed increase in reserve feature utilization (39 $\rightarrow$ 63 features during training) validates the framework's ability to discover instance-specific patterns that complement population-level markers.

5.2 Comparison to Existing Methods

Unlike ProtoGate's single-metric approach, our multi-metric global selection (combining mutual information, random forest importance, and LASSO) captures both linear and non-linear dependencies. The weighted feature forwarding mechanism provides more nuanced representations than binary masks, while the feedback loop enables dynamic adaptation that ProtoGate lacks.

6 Limitations and Future Work

The framework has several technical constraints that point toward future research directions. Multiple adaptive thresholds require dataset-specific calibration, and pairwise interaction tracking cannot capture higher-order genomic dependencies. More fundamentally, the approach identifies correlative rather than causal relationships, which may not generalize across populations or treatment protocols.

To address these limitations, we plan to integrate causal inference methods to identify features with direct causal effects rather than correlations, offering improved robustness and generalizability. Future technical enhancements include tensor-based methods for higher-order interactions, meta-learning for automatic hyperparameter adaptation, and specialized evaluation metrics for dynamic selection in imbalanced settings.

7 Conclusion

We present a dynamic feature selection framework addressing key limitations in genomic data analysis through a two-stage global-to-local architecture with continuous feedback refinement. Our approach enhances ProtoGate by incorporating multi-metric global selection, adaptive patient-specific feature refinement, and class imbalance optimization.

The framework demonstrates effective performance across genomic datasets, achieving 90.9% accuracy with strong minority class performance (F1 = 0.66) on the ALL SNP dataset. The dynamic adaptation mechanism successfully discovers patient-specific patterns, with reserve feature utilization ranging from 39 to 63 features during training, indicating the system's ability to identify instance-specific markers that complement population-level features.

Our technical contributions include: (1) continuous feedback loops dynamically refining feature relevance based on prediction outcomes, (2) multi-metric global selection combining mutual information, random forest importance, and LASSO regression, and (3) Focal Loss integration for severe class imbalance optimization. These advances enable interpretable and personalized genomic analysis while maintaining computational efficiency.

The framework provides a foundation for personalized medicine by bridging population-level and patient-specific feature importance. Future integration of causal inference methods will enhance clinical applicability, moving beyond correlative associations toward mechanistically-grounded feature identification.

Through dynamic adaptation and explicit interpretability, this work advances instance-based learning for high-dimensional genomic data, contributing to more effective computational tools for personalized healthcare.

References

[1] Randy Ardywibowo, Shahin Boluki, Zhangyang Wang, Bobak Mortazavi, Shuai Huang, and Xiaoning Qian. 2022. VFDS: Variational Foresight Dynamic Selection in Bayesian Neural Networks for Efficient Human Activity Recognition. doi:10.48550/arXiv.2204.00130 arXiv:2204.00130 [cs]

[2] Wael Hassanieh and Abdallah Chehade. 2024. Selective Deep Autoencoder for Unsupervised Feature Selection. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 11 (March 2024), 12322–12330. doi:10.1609/aaai.v38i11.29123

[3] Fergus Imrie, Alexander Norcliffe, Pietro Lio, and Mihaela van der Schaar. 2023. Composite Feature Selection Using Deep Ensembles. doi:10.48550/arXiv.2211.00631 arXiv:2211.00631 [cs]

[4] Neil Jethani, Mukund Sudarshan, Yindalon Aphinyanaphongs, and Rajesh Ranganath. 2021. Have We Learned to Explain?: How Interpretability Methods Can

Learn to Encode Predictions in Their Interpretations. doi:10.48550/arXiv.2103.01890 arXiv:2103.01890 [stat]

[5] Xiangjian Jiang, Andrei Margeloiu, Nikola Simidjievski, and Mateja Jamnik. 2024. ProtoGate: Prototype-based Neural Networks with Global-to-local Feature Selection for Tabular Biomedical Data. doi:10.48550/arXiv.2306.12330 arXiv:2306.12330 [cs]

[6] Ismael Lemhadri, Feng Ruan, Louis Abraham, and Robert Tibshirani. 2021. LassoNet: A Neural Network with Feature Sparsity. doi:10.48550/arXiv.1907.12207 arXiv:1907.12207 [stat]

[7] Dugang Liu, Pengxiang Cheng, Hong Zhu, Xing Tang, Yanyu Chen, Xiaoting Wang, Weike Pan, Zhong Ming, and Xiuqiang He. 2023. DIWIFT: Discovering Instance-wise Influential Features for Tabular Data. doi:10.48550/arXiv.2207.02773 arXiv:2207.02773 [cs]

[8] Aria Masoomi, Chieh Wu, Tingting Zhao, Zifeng Wang, Peter Castaldi, and Jennifer Dy. 2020. Instance-Wise Feature Grouping. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 13374–13386.

[9] Pranoy Panda, Sai Srinivas Kancheti, and Vineeth N. Balasubramanian. 2021. Instance-Wise Causal Feature Selection for Model Interpretation. doi:10.48550/arXiv.2104.12759 arXiv:2104.12759 [cs]

[10] Antoine Passemiers, Pietro Folco, Daniele Raimondi, Giovanni Birolo, Yves Moreau, and Piero Fariselli. 2024. A Quantitative Benchmark of Neural Network Feature Selection Methods for Detecting Nonlinear Signals. *Scientific Reports* 14, 1 (Dec. 2024), 31180. doi:10.1038/s41598-024-82583-5

[11] Hanyu Peng, Guanhua Fang, and Ping Li. 2023. Copula for Instance-wise Feature Selection and Ranking. doi:10.48550/arXiv.2308.00549 arXiv:2308.00549 [cs]

[12] Sana Tonekaboni, Shalmali Joshi, Kieran Campbell, David Duvenaud, and Anna Goldenberg. 2020. What Went Wrong and When? Instance-wise Feature Importance for Time-series Models. doi:10.48550/arXiv.2003.02821 arXiv:2003.02821 [cs]

[13] Sercan O. Arik and Tomas Pfister. 2020. TabNet: Attentive Interpretable Tabular Learning. doi:10.48550/arXiv.1908.07442 arXiv:1908.07442 [cs]

[14] Dinesh Singh, Héctor Climente-González, Mathis Petrovich, Eiryo Kawakami, and Makoto Yamada. 2020. FsNet: Feature Selection Network on High-dimensional Biological Data. doi:10.48550/arXiv.2001.08322 arXiv:2001.08322 [cs]

[15] Zixuan Song and Jun Li. 2019. Variable Selection with False Discovery Rate Control in Deep Neural Networks. doi:10.48550/arXiv.1909.07561 arXiv:1909.07561 [stat]

[16] Ram Dyuthi Sristi, Ofir Lindenbaum, Shira Lifshitz, Maria Lavzin, Jackie Schiller, Gal Mishne, and Hadas Benisty. 2024. Contextual Feature Selection with Conditional Stochastic Gates. doi:10.48550/arXiv.2312.14254 arXiv:2312.14254 [cs]

[17] Maksymilian Wojtas and Ke Chen. 2020. Feature Importance Ranking for Deep Learning. doi:10.48550/arXiv.2010.08973 arXiv:2010.08973 [cs]

[18] Yutaro Yamada, Ofir Lindenbaum, Sahand Negahban, and Yuval Kluger. 2020. Feature Selection Using Stochastic Gates. doi:10.48550/arXiv.1810.04247 arXiv:1810.04247 [cs]

[19] Junchen Yang, Ofir Lindenbaum, and Yuval Kluger. 2022. Locally Sparse Neural Networks for Tabular Biomedical Data. doi:10.48550/arXiv.2106.06468 arXiv:2106.06468 [cs]

[20] Taisuke Yasuda, MohammadHossein Bateni, Lin Chen, Matthew Fahrbach, Gang Fu, and Vahab Mirrokni. 2023. Sequential Attention for Feature Selection. doi:10.48550/arXiv.2209.14881 arXiv:2209.14881 [cs]

[21] Yuya Yoshikawa and Tomoharu Iwata. 2022. Neural Generators of Sparse Local Linear Models for Achieving Both Accuracy and Interpretability. *Information Fusion* 81 (May 2022), 116–128. doi:10.1016/j.inffus.2021.11.009 arXiv:2003.06441 [cs]