



OPEN ACCESS

EDITED BY

Nathan Brown,
Healx Ltd., United Kingdom

REVIEWED BY

Zaid Anis Sherwani,
University of Karachi, Pakistan
Shubham Vishnoi,
University of Limerick, Ireland

*CORRESPONDENCE

Anthony Wong,
✉ awong16@illinois.edu

RECEIVED 18 August 2025

REVISED 17 October 2025

ACCEPTED 22 October 2025

PUBLISHED 17 November 2025

CITATION

Wong A, Guduri S, Chen T and Patel K (2025)
Machine learning-guided optimization of
triple agonist peptide therapeutics for
metabolic disease.
Front. Bioinform. 5:1687617.
doi: 10.3389/fbinf.2025.1687617

COPYRIGHT

© 2025 Wong, Guduri, Chen and Patel. This is
an open-access article distributed under the
terms of the [Creative Commons Attribution
License \(CC BY\)](#). The use, distribution or
reproduction in other forums is permitted,
provided the original author(s) and the
copyright owner(s) are credited and that the
original publication in this journal is cited, in
accordance with accepted academic practice.
No use, distribution or reproduction is
permitted which does not comply with
these terms.

Machine learning-guided optimization of triple agonist peptide therapeutics for metabolic disease

Anthony Wong^{1*}, Sanskruthi Guduri¹, TsungYen Chen^{1,2} and Kunal Patel^{1,2}

¹Carle Illinois College of Medicine, Urbana, IL, United States, ²Carle Foundation Hospital, Department of Medicine, Urbana, IL, United States

Introduction: Multi-target peptide therapeutics targeting glucagon receptor (GCGR), glucagon-like peptide-1 receptor (GLP1R), and glucose-dependent insulinotropic polypeptide receptor (GIPR) represent a promising approach for treating diabetes and obesity. Triple agonist peptides demonstrate promising therapeutic potential compared to single-target approaches, yet rational design remains computationally challenging due to complex sequence-structure activity relationships. Existing methods, primarily based on convolutional neural networks, impose limitations including fixed sequence lengths and inadequate representation of molecular topology. Graph Attention Networks (GAT) offer advantages in capturing molecular structures and variable-length peptide sequences while providing interpretable insights into receptor-specific binding determinants.

Methods: A dataset of 234 peptide sequences with experimentally determined binding affinities was compiled from multiple sources. Peptides were represented as molecular graphs with seven-dimensional node features encoding physicochemical properties and positional information. The GAT architecture employed a shared encoder with task-specific prediction heads, implementing transfer learning to address limited GIPR training data. Performance was evaluated using 5-fold cross-validation and independent validation on 24 literature-derived sequences. A genetic algorithm framework was developed for peptide sequence optimization, incorporating multi objective fitness evaluation based on predicted binding affinity, biological plausibility, and sequence novelty.

Results: Cross-validation demonstrated robust GAT performance across all receptors, with GCGR achieving high accuracy (AUC ROC: 0.915 $\pm$ 0.050), followed by GLP1R (AUC-ROC: 0.853 $\pm$ 0.059), and GIPR showing acceptable performance despite limited data (AUC-ROC: 0.907 $\pm$ 0.083). Comparative analysis revealed receptor-specific advantages: GAT significantly outperformed CNN for GCGR prediction (RMSE: 0.942 vs. 1.209, p = 0.0013), while CNN maintained superior GLP1R performance (RMSE: 0.552 vs. 0.723). Genetic algorithm optimization measurable improvement over baseline, with 4.0% fitness Enhancement and generation of 20 candidates exhibiting mean binding probabilities exceeding 0.5 across all targets. The GAT-based framework provides a computational approach in computational peptide design, demonstrating receptor-specific advantages and robust optimization capabilities.

Conclusion: Genetic algorithm optimization enables systematic exploration of sequence space within existing agonist scaffolds while maintaining biological constraints. This approach provides a rational framework for prioritizing experimental validation efforts in triple agonist development.

KEYWORDS

peptide design, machine learning, bioactivity prediction, drug discovery, graph attention networks

1 Introduction

The global obesity epidemic and the rising prevalence of type 2 diabetes mellitus (T2DM) represent major public health challenges, affecting over 650 million adults worldwide with obesity and 537 million individuals with diabetes (Müller et al., 2019; Baggio and Drucker, 2021). Metabolic syndrome, characterized by the clustering of insulin resistance, abdominal obesity, dyslipidemia, and hypertension, affects approximately 37.6%–41.8% of US adults and is associated with a 2-fold increased risk of cardiovascular disease and 1.5-fold increased risk of all-cause mortality (Li et al., 2023; Mottillo et al., 2010). Per-person healthcare costs for individuals with metabolic syndrome average \$5,732 annually compared to \$3,581 for those without the condition (Boudreau et al., 2009).

Traditional therapeutic approaches targeting single pathways have demonstrated limited long-term efficacy, highlighting the need for innovative multi-target strategies that address the complex pathophysiology underlying metabolic dysfunction (Brandt et al., 2018). The development of multi-target peptide therapeutics represents a paradigm shift in precision medicine, offering the potential for superior glycemic control, substantial weight reduction, and improved cardiovascular outcomes compared to conventional single-target approaches (Finan et al., 2016; Samms et al., 2020).

Recent clinical breakthroughs have validated the therapeutic potential of multi-receptor agonists in metabolic disease treatment. Tirzepatide, a dual glucose-dependent insulinotropic polypeptide receptor (GIPR) and glucagon-like peptide-1 receptor (GLP1R) agonist, demonstrated unprecedented efficacy in the SURPASS clinical trial program, achieving HbA1c reductions of up to 2.58% (Frias et al., 2018). The clinical development of retatrutide, a triple agonist targeting GCGR, GLP1R, and GIPR, has further demonstrated the potential of multi-target approaches, with Phase 2 results showing 24.2% weight reduction at 48 weeks (Jastreboff et al., 2022). These clinical successes underscore the therapeutic value of targeting multiple components of the incretin system simultaneously, leading to enhanced metabolic benefits through complementary mechanisms of action.

The glucagon receptor (GCGR), GLP1R, and GIPR represent critical nodes in metabolic homeostasis, each contributing distinct physiological effects that collectively address the multifaceted nature of metabolic disorders (Alexander et al., 2021). GLP1R agonism provides glucose-dependent insulin secretion, gastric emptying delay, and appetite suppression, while GIPR activation enhances insulin sensitivity and promotes beneficial effects on bone metabolism (Knerr et al., 2020). GCGR agonism contributes to increased energy expenditure, enhanced hepatic glucose production regulation, and potential benefits in non-alcoholic fatty

liver disease (Winther and Holst, 2024). The synergistic activation of these three receptors offers a comprehensive approach to metabolic regulation that addresses both the glycemic and weight management aspects of T2DM and obesity.

Traditional drug development approaches rely on iterative experimental trial-and-error cycles that are both time-intensive and prohibitively costly, with the average cost of bringing a new drug to market reaching up to \$2.23 billion (Wouters et al., 2020). Recent advances in machine learning have demonstrated the potential to accelerate peptide discovery through computational design platforms, enabling the exploration of vast chemical spaces that would be impractical to investigate experimentally (Yang et al., 2019). Puzkarska et al. employed deep multi-task convolutional neural networks (CNNs) to design GCGR/GLP1R dual agonists with superior biological potency, demonstrating up to sevenfold potency improvements compared to existing compounds in their training set (Puzkarska et al., 2024). However, this CNN-based approach imposed several methodological constraints, including fixed sequence lengths of 30 amino acids and limited flexibility in representing complex molecular topology and modifications.

Graph neural networks (GNNs) have emerged as a powerful paradigm for molecular representation learning, offering significant advantages over sequence-based methods in capturing three-dimensional molecular structures and chemical interactions (Wu et al., 2020; Zhang et al., 2022). Graph Attention Networks (GATs), introduced by Veličković et al., enable nodes to attend over their neighborhoods with learnable attention weights, providing interpretable insights into molecular interactions and functional relationships (Veličković et al., 2017). Unlike CNNs that require fixed-length inputs, GATs naturally accommodate variable-length peptide sequences while preserving molecular topology through explicit representation of chemical bonds, modifications, and spatial relationships (Lv et al., 2024; Gao et al., 2018).

The application of GNNs to peptide design has shown particular promise in capturing complex molecular features that are challenging to represent in sequence-based models (Kipf and Welling, 2016). Recent work by Xiong et al. demonstrated that graph attention mechanisms achieve state-of-the-art performance across molecular property prediction tasks, with attention weights revealing functionally important molecular regions and binding determinants (Xiong et al., 2020). Similarly, Strokach et al. successfully applied deep graph neural networks to protein design, demonstrating the capacity for *de novo* sequence generation with experimental validation (Strokach et al., 2020).

Despite these advances, existing neural network approaches for peptide therapeutics have primarily focused on dual-target optimization or single-receptor systems (Puzkarska et al., 2024). The extension to triple-agonist design presents unique

computational challenges, including the need for balanced multi-target optimization, the integration of transfer learning strategies to address limited experimental data availability, and the development of robust evaluation frameworks for multi-receptor binding prediction.

This study presents a methodological advancement that extends the CNN-based approach of Puskarska et al. to a GAT-based framework for triple-agonist peptide design targeting GCGR, GLP1R, and GIPR. Our approach addresses key limitations of previous methods by implementing flexible sequence length handling through graph representations that accommodate diverse peptide modifications, extending the target scope from dual agonist (GCGR/GLP1R) to triple agonist capability, incorporating transfer learning strategies to leverage limited GIPR experimental data, and employing attention mechanisms to provide interpretable insights into receptor-specific binding determinants.

We demonstrate that our GAT-based approach achieves improved predictive performance for GCGR and GIPR receptors, with comparable performance for GLP1R, while maintaining the ability to generate peptide sequences with high predicted binding affinities. The integration of genetic algorithm-based sequence optimization enables the systematic exploration of diverse sequence variants within the agonist design space, identifying promising candidates for experimental validation. This methodological framework establishes a foundation for enhanced computational tools in multi-target peptide therapeutics development, with immediate applications in obesity and diabetes drug discovery.

2 Methods

2.1 Dataset compilation and preprocessing

The training dataset was compiled from multiple established sources to ensure comprehensive coverage of GCGR, GLP1R, and GIPR binding data. Primary data sources included: (1) the open-access dataset from Puskarska et al. (Alexander et al., 2021), (2) curated ChEMBL entries for peptide-receptor interactions (Mendez et al., 2019), (3) experimental data from Knerr et al. investigating peptide-based polyagonists (Knerr et al., 2018), (4) dual agonist optimization studies by Evers et al. (Evers et al., 2018), and (5) structural characterization data from recent triple agonist studies (Zhao et al., 2022; Cock et al., 2009) (Supplementary Table S1).

Peptide sequences underwent standardized preprocessing to handle non-standard amino acids and chemical modifications. Modified residues including D-amino acids, unnatural amino acids, and lipidated variants, were systematically cataloged and encoded with their corresponding physicochemical properties using BioPython ProtParam v1.79 (Cock et al., 2009) (Supplementary Table S2).

Activity classification employed a binary high-affinity threshold of 1000 pM EC50. Given the limited availability of data across all three receptors, a more lenient threshold was adopted compared to the 10 pM cutoff used in previous studies (Alexander et al., 2021) to ensure sufficient positive examples for model training while maintaining pharmacologically relevant activity levels.

2.2 Graph representation of peptide sequences

Peptide sequences were converted to graph representations using a molecular topology-preserving approach implemented using PyTorch Geometric v2.0.4 (Fey and Lenssen, 2019). Each amino acid residue was represented as a node with a seven-dimensional feature vector encoding: (1) hydrophobicity (Kyte-Doolittle scale) (Kyte and Doolittle, 1982), (2) net charge at physiological pH, (3) molecular weight, (4) D-amino acid indicator, (5) lipidation status, (6) sine positional encoding, and (7) cosine positional encoding.

Node features were computed using the following protocol. For standard amino acids, hydrophobicity, charge and molecular weight values were obtained from established amino acid properties. D-amino acid and lipidation indicators were set to 0.0. For non-standard residues values were encoded using experimentally or computationally predicted physicochemical properties. D-amino acids (i.e., d-serine or d-alanine) used L-enantiomer properties with D-amino acid indicator sets to 1.0. Lipidated residues such as (K [(yE-C16)]) included lipid chain contributions to hydrophobicity and molecular weight calculation with lipidation status set to 1.0.

Physicochemical features (hydrophobicity, charge, and molecular weight) were normalized using dataset-specific statistics computed from all training sequences: mean-centered and scaled by standard deviation to ensure numerical stability during training. We used a simplified positional encoding approach with the following:

$$\sin\left(\pi\frac{pos}{L}\right), \cos\left(\pi\frac{pos}{L}\right)$$


to represent amino acid positions. This encoding normalizes position by sequence length and provides two dimensions of positional information, which we considered appropriate given our limited dataset size ($N = 125$) and the short length of peptide sequences in our study (Vaswani et al., 2017).

Edge connectivity was established through peptide bond relationships, creating bidirectional edges between sequential amino acid residues. For modified amino acids, feature vectors incorporated experimentally derived or computationally predicted physicochemical properties as cataloged in the preprocessing stage.

2.3 Graph attention network architecture

The GAT model employed a shared encoder architecture with task-specific prediction heads for multi-target optimization. The shared encoder comprised four GATv2Conv layers with 6-head attention mechanisms, 96-dimensional hidden representations, and ReLU activation functions. Each attention layer was followed by batch normalization and dropout ($p = 0.2$) to prevent overfitting, with residual connections applied every two layers (Figure 1).

Global graph representation was achieved using Set2Set pooling with three processing steps, followed by a representation layer consisting of linear transformation (192→96 dimensions), ReLU activation, batch normalization, and dropout ($p = 0.2$). Set2Set pooling was chosen over simple global pooling methods to better capture permutation-invariant graph-level representations while maintaining sensitivity to sequence order information.



Task-specific heads consisted of three-layer fully connected networks with ReLU activations, batch normalization, and dropout between layers, terminating in single-unit outputs for binary classification. The multi-task loss function employed weighted binary cross-entropy with equal weighting ($\alpha = 0.5$) across active receptors:

A three-stage transfer learning approach was implemented to address the limited availability of GIPR training data relative to GCGR and GLP1R. Stage 1 involved initial training on combined GCGR/GLP1R data (80 epochs) using all available sequences with valid labels for either receptor, with learning rate of 1×10^{-3} . Stage two employed encoder parameter freezing with exclusive GIPR head training (100 epochs) using maintained learning rate of 1×10^{-3} to prevent catastrophic forgetting. Stage 3 implemented unified fine-tuning (60 epochs) with encoder unfreezing and reduced learning rate of 1×10^{-4} for global optimization across all three receptors.

Model performance was evaluated using stratified 5-fold cross-validation. Stratification was performed based on multi-target activity patterns, creating stratification keys from the combination of receptor-specific activity states (high affinity, low affinity, or missing data) to ensure balanced representation across folds. Each fold employed an 80/20 train/validation split for hyperparameter optimization and early stopping, with validation loss monitoring and patience of 15 epochs. Training utilized the Adam optimizer with gradient clipping ($\text{max_norm} = 1.0$) for numerical stability. Focal loss was implemented with $\alpha = 0.25$ and $\gamma = 2.0$ to address class imbalance in receptor-specific datasets.

2.4 Model comparison with literature baseline

Model performance was compared against the established multi-task convolutional neural network (CNN) ensemble from Puskarska et al. (Alexander et al., 2021) using identical evaluation protocols. The GAT model employed a k-fold cross-validation ensemble approach, while the CNN baseline utilized the original 12

$\times 6$ ensemble architecture (72 total models) as reported in the source publication.

Both models were trained on the original 125-sequence dataset from Puskarska et al. and evaluated on an independent reference dataset from Day et al. containing 19 peptide sequences with experimentally validated GCGR and GLP1R activities (Cock et al., 2009). This evaluation protocol ensured direct comparability with published results while testing generalization to completely unseen sequences.

The GAT regression model architecture established through this comparison was subsequently adapted for binary classification tasks used throughout the remainder of this study, with sigmoid activation functions replacing linear outputs and binary cross-entropy loss replacing mean squared error.

Performance metrics were computed using identical preprocessing pipelines and evaluation criteria. Statistical significance was assessed using paired t-tests on prediction errors across the reference dataset, with p-values calculated for each receptor target.

2.5 Independent model validation

Independent validation sequences were obtained from previously published studies: Day et al., Finan et al., and Zhang et al. comprising 67 peptide sequences with experimentally determined EC50 values for GCGR, GLP1R, and GIPR (Day et al., 2009; Finan et al., 2015; Zhang et al., 2025) (Supplementary Table S3). These sequences were selected to evaluate model generalizability on data not used during training or cross-validation.

To assess model performance on novel sequences, we implemented a similarity-based filtering approach. Sequence identity was calculated between each validation sequence and all training sequences using token-level comparison, accounting for non-standard amino acid modifications. The maximum similarity to any training sequence was determined for each validation peptide.

Validation sequences were categorized into two groups: (1) novel sequences with $\leq 80\%$ similarity to any training sequence, and (2) the complete validation set. This threshold was selected to

distinguish sequences with substantial structural differences from the training data while maintaining adequate sample sizes for statistical evaluation.

In addition to the quantitative EC50-based validation, we performed a secondary validation using a curated dataset of known therapeutic peptides with established receptor activity profiles but without complete quantitative EC50 data across all three receptors. This dataset comprised nine clinically relevant peptides including FDA-approved therapeutics and compounds in clinical development with documented high or low affinity classifications for GCGR, GLP1R, and GIPR based on published pharmacological characterization studies (Supplementary Table S4) (Coskun et al., 2022; Whitley, 2025; Finan et al., 2025; Starling, 2022; Willard et al., 2020; Henderson et al., 2016; Yabut and Drucker, 2023; Chen et al., 2022; Zimmermann et al., 2022; Wiśniewski et al., 2016; Han et al., 2013; Lau et al., 2015; Meng et al., 2008; Knudsen et al., 2000; Miranda et al., 2008; Plisson et al., 2017; Murage et al., 2008; Mapelli et al., 2009; Knerr et al., 2022). Known receptor activities were manually curated from literature sources and regulatory documents, with activities classified as “high” (therapeutically relevant binding, typically EC50 < 1000 pM) or “low” (minimal or no significant binding activity) for each receptor. This binary classification approach was necessary due to the heterogeneous nature of activity reporting across different studies and the absence of standardized EC50 measurements for all peptide-receptor combinations.

Predictions were generated using the ensemble of k-fold trained GAT models (n = 5 folds). Each validation sequence was converted to a molecular graph representation using the same preprocessing pipeline employed during training, including feature normalization parameters computed from the original training dataset.

Binary classification performance was evaluated using an EC50 threshold of 1000 pM, where values < 1000 pM indicated high affinity (positive class) and ≥1000 pM indicated low affinity (negative class). Missing EC50 values were excluded from analysis. Performance metrics included accuracy, F1-score, precision, recall, area under the receiver operating characteristic curve (AUC-ROC), and area under the precision-recall curve (AUC-PR).

Model performance was assessed separately for each receptor (GCGR, GLP1R, GIPR) and validation subset (novel sequences vs. complete set). Ensemble predictions were obtained by averaging probability outputs across all fold models. All analyses were implemented in Python using scikit-learn v1.0.2 for metric calculations and PyTorch Geometric v2.0.4 for graph neural network operations.

2.6 Genetic algorithm framework for peptide design

The evolutionary optimization phase began by initializing a population of 100 randomly generated peptide sequences, each containing 25–35 amino acid residues consistent with the length distribution of known peptide hormones in this therapeutic class. All initial sequences were filtered using the biological plausibility scoring system to ensure only chemically reasonable candidates (scoring ≥0.3) proceeded to fitness evaluation. The optimization

process then iteratively evolved this population over a maximum of 50 generations, with each cycle involving comprehensive fitness assessment using the multi-objective function described in Equation 1, followed by parent selection through tournament-based competition among groups of three individuals. Selected parent sequences underwent single-point crossover with 80% probability to generate offspring, after which adaptive mutation was applied with 10% probability to introduce sequence variations. Throughout this process, the top 10% of individuals were preserved unchanged between generations to maintain high-quality solutions, while convergence monitoring assessed population improvement to determine optimal termination timing.

The fitness function was a summation of four weighted components to balance receptor binding activity, biological plausibility, and sequence novelty:

Equation 1: Multi – Objective Fitness Function

$$F(s) = H(s) + P(s) + Pmin(s) + M(s) + N(s) + D(s)$$

where: $H(s)$ = number of high affinity receptor

$P(s)$ = mean binding probability across receptor

$Pmin(s)$ = minimum binding probability across receptor

$B(s)$ = biological plausibility conservation score $\in [0, 1]$

$N(s)$ = sequence novelty score $\in [0, 1]$

$D(s)$ = population diversity score $\in [0, 1]$

2.6.1 Receptor binding activity

Binding probabilities for GCGR, GLP1R, and GIPR were predicted using an ensemble GAT model, with scoring incorporating the number of high-affinity targets achieving the target threshold of ($p \geq 0.5$), average predicted probability, and minimum predicted probability across receptors.

2.6.2 Motif preservation

Conserved motifs critical for receptor binding, identified from literature-based structure–activity relationships, were preserved using pattern-matching algorithms allowing conservative substitutions. Critical motifs included N-terminal his6-his8 positions, central FTSD tetrapeptide, and C-terminal amidation patterns. Motif scoring used position-specific matching with wildcard tolerance and conservative substitution allowances (Table 1).

2.6.3 Biological plausibility

A weighted score combined chemical constraint adherence, motif preservation, predicted proteolytic stability (trypsin/chymotrypsin cleavage site analysis), and similarity in amino acid composition to natural peptide hormones.

Equation 2: Biological Plausibility Score

$$B(s) = 0.3 C(s) + 0.35 M(s) + 0.2 P(s) + 0.15 A(s) \geq 0.3$$

TABLE 1 Literature Derived Critical Binding Motifs: Summary of experimentally validated and evolutionarily conserved amino acid motifs essential for incretin receptor binding and activation. Critical residues are identified from structure-activity relationship studies and comparative genomic analysis across species orthologs and protein paralogs (Vaswani et al., 2017; Fey and Lenssen, 2019).

Motif	Notes	Source
N-terminal: H ⁷ E ⁹ G ¹⁰ F ¹² T ¹³ D ¹⁵ C-terminal: F ²⁸ I ²⁹ L ³² R ³⁶	GLP-1 (7–36), H ⁷ G ¹⁰ T ¹³ D ¹⁵ F ²⁸ I ²⁹ form the most critical residues for receptor activation and binding	Manandhar and Ahn (2015)
GLP1 A ² E ³ T ⁵ T ⁷ Y ¹³ A ¹⁸ A ¹⁹ K ²⁰ E ²¹ G ²⁹ Glucagon S ² Q ³ T ⁵ Y ¹³ A ¹⁹ GIP Y ¹ E ³ S ⁸ S ¹¹ N ²⁴ L ²⁷	Residues conserved across Orthologs in mouse, anole, chicken, <i>Xenopus tropicalis</i> , medaka, fugu, tetraodon, stickleback, and zebrafish	Moon et al. (2012)
GLP1 H ¹ G ⁴ F ⁶ D ⁹ L ¹⁴ F ²² I ²³ W ²⁵ L ²⁶ Glucagon H ¹ G ⁴ F ⁶ D ⁹ L ¹⁴ D ¹⁵ D ²¹ F ²² V ²³ W ²⁵ L ²⁶ GIP G ⁴ F ⁶ D ⁹ D ¹⁵ D ²¹ F ²² V ²³ W ²⁵ L ²⁶	Residues conserved across Paralogs	Orskov et al. (1989)

where:C(s) = chemical constraint score (charge distribution,hydrophobic patches)

$M(s)$ = motif preservation score

$P(s)$ = proteolytic stability score (trypsin/chymotrypsin cleavage analysis)

$A(s)$ = amino acid composition similarity to natural hormones

In addition, sequences failing to achieve a biological plausibility score of 0.3 were rejected from the population.

2.6.4 Sequence novelty

Sequence similarity to the training set was penalized to encourage diversity, with population-level diversity tracking used to avoid premature convergence. Novelty scores were computed using sequence alignment and Levenshtein distance metrics, with penalties applied for high similarity (>80%) to existing training sequences.

2.7 Evolutionary operations

Parent selection for sequence reproduction employed tournament-based competition, where groups of three individuals competed based on their fitness scores, with the highest-scoring sequence from each tournament selected for breeding. This approach balanced selective pressure toward high-fitness individuals while maintaining sufficient population diversity to prevent premature convergence to local optima.

Sequence recombination was performed through single-point crossover applied to 80% of parent pairs. A random crossover point was selected within each parent sequence, and genetic material was exchanged to create two offspring sequences. Following crossover, offspring sequences were validated for length constraints and biological plausibility, with sequences exceeding acceptable limits or failing plausibility checks undergoing repair through motif-preserving substitution operators.

Sequence diversification was achieved through adaptive mutation applied with 10% probability across multiple mechanisms. Point mutations constituted 70% of mutation events, involving single amino acid substitutions at random positions. Conservative substitutions, representing 20% of mutations, introduced physiochemically similar amino acid replacements to maintain local structural properties. The remaining 10% of mutations involved introduction of modified residues including D-amino acids or lipidation modifications to expand chemical diversity. Mutation rates were dynamically adjusted based on population diversity metrics and convergence indicators to maintain optimal exploration-exploitation balance throughout the evolutionary process.

Population continuity between generations was ensured by preserving the top 10% of individuals unchanged, corresponding to the 10 highest-fitness sequences. This elitist strategy prevented loss of high-quality solutions while allowing the majority of the population to undergo evolutionary modification. Primary termination occurred upon reaching the maximum generation limit of 50 cycles, providing sufficient evolutionary time for convergence while maintaining computational feasibility. Premature termination was triggered when fitness improvement fell below 0.1 units over three consecutive generations, indicating population convergence to optimal solutions.

Additional monitoring detected population stagnation, defined as lack of meaningful diversity changes over 10 generations, prompting termination to prevent computational waste on converged populations. Optional user-defined fitness thresholds could also trigger early termination when target performance levels were achieved.

2.8 Sequence conservation and novelty assessment

Genetically optimized peptide sequences generated from the evolutionary algorithm were subjected to comprehensive

computational analysis. Conservation patterns were analyzed by comparing synthetic sequences against three native peptide hormone references: human glucagon, GLP-1, and GIP (Orskov et al., 1989; Bromer et al., 1957; Moody et al., 1984). Physicochemical properties were computed for both synthetic and native sequences using BioPython's ProteinAnalysis module (Cock et al., 2009). Calculated parameters included molecular weight, isoelectric point, instability index, and grand average of hydropathy (GRAVY) scores. Lipophilicity was estimated using position-specific hydrophobicity values from the Kyte-Doolittle scale, normalized by sequence length. Polar surface area (PSA) was approximated by counting polar amino acids and applying a conversion factor of 50 Å per polar residue (Ertl et al., 2000).

To evaluate the novelty and diversity of computationally generated peptide candidates, sequence similarity analysis was performed between the top-ranking optimized peptides and the original training dataset sequences. A flexible similarity scoring algorithm was implemented using dynamic programming to calculate edit distances between peptide sequences (Qiao et al., 2020), accounting for non-standard amino acid (NSAA) modifications through specialized tokenization. Each candidate sequence was tokenized using regular expression pattern matching to preserve NSAA bracket notation, then compared against all training sequences using a normalized edit distance metric based on the Levenshtein distance algorithm (Cheng et al., 2024). For each generated peptide candidate, the highest similarity score to any training sequence was recorded along with the corresponding best match peptide ID and sequence. This analysis enabled assessment of whether the evolutionary algorithm successfully explored novel sequence space beyond the training data or converged toward existing high-activity peptides.

2.9 Software implementation

All analyses were implemented in Python 3.8+ using PyTorch v1.9+ for deep learning frameworks and PyTorch Geometric v2.0+ for graph neural network operations (Fey and Lenssen, 2019). Additional computational tools included scikit-learn v1.0+ for performance metrics, BioPython v1.79+ for sequence analysis, and RDKit for cheminformatics operations. Complete software dependencies with specific version requirements are provided in Supplementary Methods.

3 Results

3.1 Database characterization

The curated dataset comprised 234 unique peptide sequences with experimentally determined binding affinities across three G-protein coupled receptors, representing an expansion from the original 125-sequence dataset (Table 2). GCGR measurements were available for 206 sequences (88.0%), with 101 sequences (49.0%). GLP1R data were complete across all 234 sequences, with 175 sequences (74.8%) demonstrating high affinity. GIPR measurements, representing the most limited subset, were available

TABLE 2 Database Summary Distribution of peptide sequences and binding affinity characteristics across three incretin receptors. High affinity defined as EC50 ≤ 1000 pM. GLP1R exhibits the highest proportion of high-affinity sequences (74.8%) with lowest median potency (22.6 pM), while GCGR shows intermediate performance (49.0% high affinity, 1265.0 pM median EC50). GIPR represents the smallest dataset with moderate high-affinity representation (57.1%, 469.5 pM median EC50).

Receptor	Total	High affinity	Median EC50
GCGR	206	101 (49.0%)	1265.0 pM
GLP1R	234	175 (74.8%)	22.6 pM
GIPR	56	32 (57.1%)	469.5 pM

for 56 sequences (23.9%), with 32 sequences (57.1%) showing high affinity (Figures 2A,B).

The dataset exhibited substantial multi-target activity, with 97 sequences (41.5%) demonstrating dual agonism for GCGR and GLP1R, 24 sequences (10.3%) for GCGR and GIPR, and 31 sequences (13.2%) for GLP1R and GIPR (Figure 2). Triple agonist sequences comprised 24 sequences with complete activity data across all three receptors.

3.2 Comparative analysis to existing models



Cross-validation performance of the ensemble GNN (4 × 4 configuration) was compared against previously reported results. The ensemble GNN achieved RMSE values of 0.52 ± 0.09 for GCGR and 0.79 ± 0.08 for GLP1R, compared to the original multi-task neural network ensemble RMSE values of 0.59 ± 0.05 for GCGR and 0.68 ± 0.04 for GLP1R (Supplementary Figures S1A–F).


Direct comparison between the 4-fold, 4-model ensemble Graph Attention Network (GAT) and the established multi-task Convolutional Neural Network (CNN) was performed using 19 literature-derived sequences (Cock et al., 2009). For GCGR prediction (EC50_LOG_T1), GAT outperformed CNN with lower RMSE (0.942, 95% CI: [0.298, 1.443] vs. 1.209, 95% CI: [0.735, 1.697]) and higher R² values (0.305 vs. -0.144). GAT also achieved superior Pearson correlation (0.683, 95% CI: [-0.385, 0.981] vs. 0.600, 95% CI: [-0.200, 0.904]), with statistically significant better performance (p = 0.0013) (Figures 3A–C).

For GLP1R prediction (EC50_LOG_T2), CNN demonstrated superior performance with lower RMSE (0.552, 95% CI: [0.461, 0.633] vs. 0.723, 95% CI: [0.540, 0.896]), higher R² (-0.208, 95% CI: [-1.454, 0.284] vs. -1.067, 95% CI: [-2.850, -0.284]), and higher Pearson correlation (0.686, 95% CI: [0.336, 0.922] vs. -0.037, 95% CI: [-0.473, 0.515]), though statistical testing revealed no significant difference between approaches (p = 0.136) (Figure 3).

3.3 Triple agonist cross-validation performance

The GAT model demonstrated robust performance across all three target receptors in 5-fold cross-validation experiments





(Figure 4). GCGR classification achieved the highest performance with an AUC-ROC of 0.915 ± 0.050 and F1-score of 0.882 ± 0.067 . GLP1R prediction yielded an AUC-ROC of 0.853 ± 0.059 and F1-score of 0.908 ± 0.027 . GIPR classification, despite limited training data, achieved an AUC-ROC of 0.907 ± 0.083 and F1-score of 0.818 ± 0.137 (Table 3).


The model was evaluated on an independently derived dataset from published literature. The independent validation dataset contained 58 total sequences. 9 sequences were removed due to exact matches with the training set. Similarity analysis revealed a bimodal distribution, with 37 sequences (67%) showing $\leq 80\%$ similarity to training data (novel sequences) and 19 sequences (34%) showing $> 90\%$ similarity (Supplementary Figure S2). The complete validation set provided larger sample sizes with 58 samples for GCGR and GLP1R, and 42 samples for GIPR.

Evaluation on novel sequences ($\leq 80\%$ similarity to training data) showed variable performance across receptors. GCGR demonstrated high discriminatory ability (AUC-ROC = 0.953, AUC-PR = 0.951) but moderate classification performance (F1-score = 0.679, $n = 37$). GLP1R showed moderate discrimination (AUC-ROC = 0.604, AUC-PR = 0.811) with acceptable classification metrics (F1-score = 0.787, $n = 37$). GIPR achieved good discriminatory performance (AUC-ROC = 0.818, AUC-PR = 0.989) with high classification accuracy (F1-score = 0.971, $n = 35$).

Evaluation on the complete validation set showed similar discriminatory patterns with some performance variations. GCGR maintained comparable performance (AUC-ROC = 0.950, AUC-PR = 0.957, F1-score = 0.795, $n = 58$). GLP1R showed reduced discriminatory ability (AUC-ROC = 0.358, AUC-PR = 0.779) while maintaining reasonable classification performance (F1-score = 0.863, $n = 58$). GIPR performance

TABLE 3 Cross-Validation Metrics Five-fold cross-validation performance metrics (mean $\pm$ standard deviation) for graph attention network models predicting peptide binding affinity across GCGR, GLP1R, and GIPR receptors. Metrics include area under receiver operating characteristic curve (AUC-ROC), area under precision-recall curve (AUC-PR), F1-score, precision, recall, and balanced accuracy. Total sample sizes shown for each receptor dataset.

Receptor	AUC-ROC (mean $\pm$ Std)	AUC-PR (mean $\pm$ Std)	F1-score (mean $\pm$ Std)	Precision (mean $\pm$ Std)	Recall (mean $\pm$ Std)	Balanced accuracy (mean $\pm$ Std)	Total samples
GCGR	0.915 $\pm$ 0.050	0.932 $\pm$ 0.037	0.882 $\pm$ 0.067	0.856 $\pm$ 0.108	0.920 $\pm$ 0.068	0.869 $\pm$ 0.079	194
GLP1R	0.853 $\pm$ 0.059	0.946 $\pm$ 0.034	0.908 $\pm$ 0.027	0.844 $\pm$ 0.039	0.983 $\pm$ 0.023	0.650 $\pm$ 0.097	222
GIPR	0.907 $\pm$ 0.083	0.955 $\pm$ 0.045	0.818 $\pm$ 0.137	0.871 $\pm$ 0.124	0.829 $\pm$ 0.229	0.773 $\pm$ 0.132	49



remained consistent when including the broader sequence set (AUC-ROC = 0.943, AUC-PR = 0.988, F1-score = 0.958, $n = 42$) (Figure 5; Supplementary Figure S3).

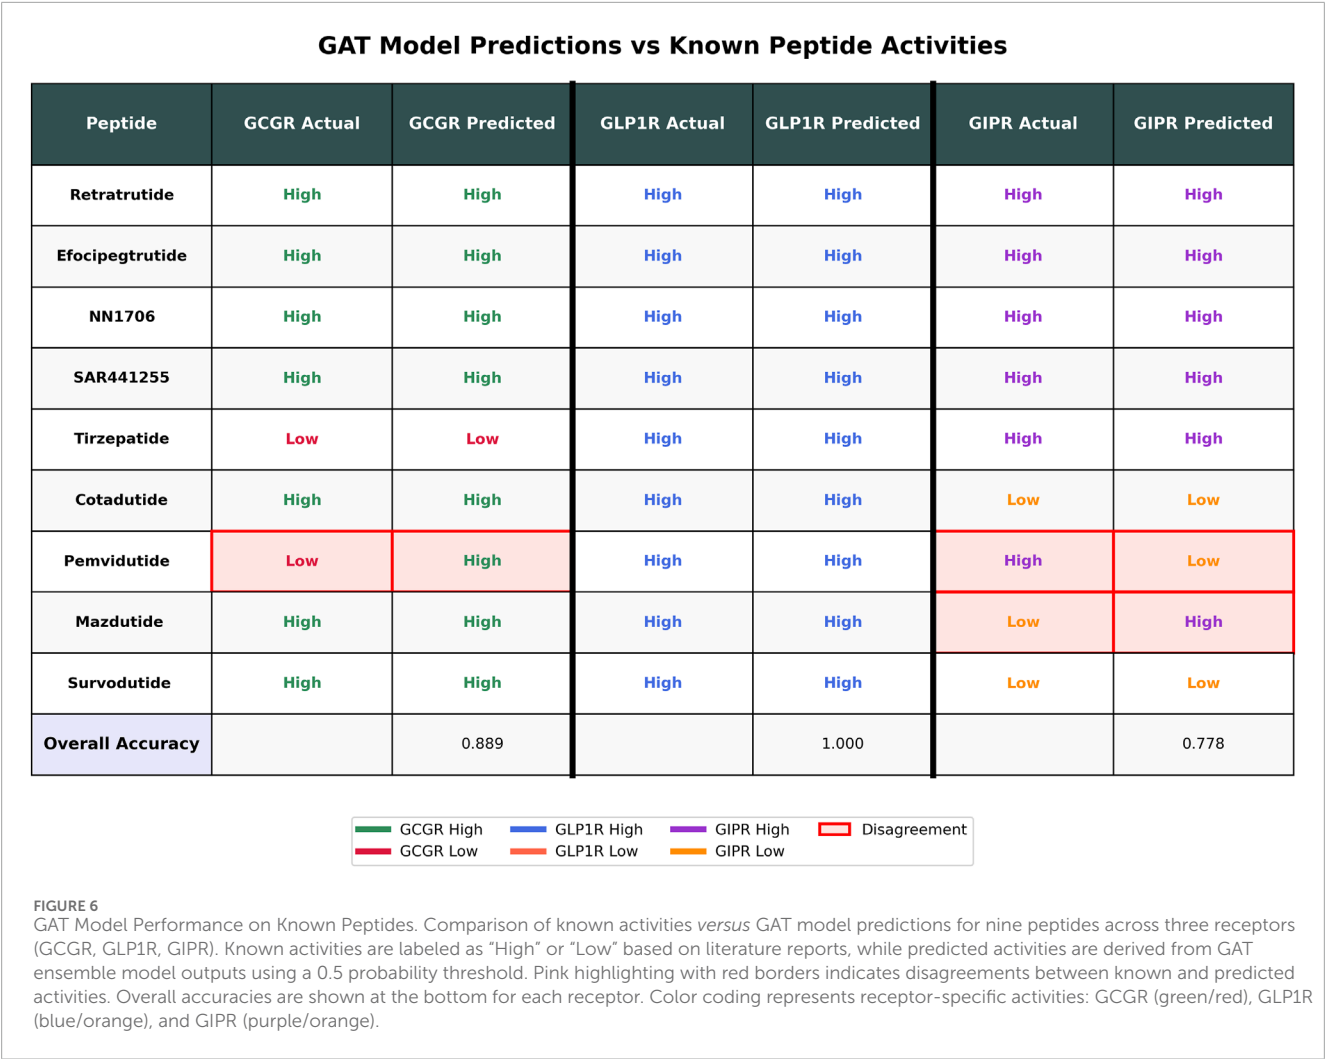
To validate the GAT ensemble model performance on clinically relevant peptides, we evaluated predictions against nine known peptides with established receptor activities (Figure 6). The model achieved accuracies of 88.9% for GCGR (8/9 peptides correct), 100% for GLP1R (9/9 peptides correct), and 77.8% for GIPR (7/9 peptides correct). Three disagreements were observed: Pemvidutide showed predicted high activity for GCGR despite known low activity, while both Pemvidutide and Mazdutide exhibited predicted low activity for GIPR contrary to their known high activities. The model correctly predicted the activities for six peptides (Retratrutide, Efocipegtrutide, NN1706, SAR441255, Tirzepatide, and Cotadutide) across all three receptors where known data

was available, demonstrating consistent performance on well-characterized dual and triple agonists.

3.4 Genetic algorithm performance

The genetic algorithm successfully optimized peptide sequences for multi-receptor binding affinity over six generations. The optimization process achieved a cumulative fitness improvement of 2.351 units, with the best-performing sequence reaching a final fitness score of 60.743, representing a 4.0% enhancement from the initial population (Figure 7).

The optimization showed a biphasic improvement pattern: an initial steep ascent with improvement rates of 1.369 and 0.893 fitness units per generation in generations 2 and 3,



respectively, followed by a stabilization phase with minor fluctuations ranging from -0.028 – 0.133 units per generation. The algorithm satisfied convergence criteria after generation 4, when the three-generation moving average improvement rate consistently fell below the predefined threshold of 0.1 fitness units per generation (Supplementary Figure S4).

Receptor-specific binding probability analysis demonstrated high performance across all targets, with mean binding probabilities greater than 0.6 for GCGR and GLP1R, and approximately 0.53 for GIPR among the best-performing sequences (Figure 7C).

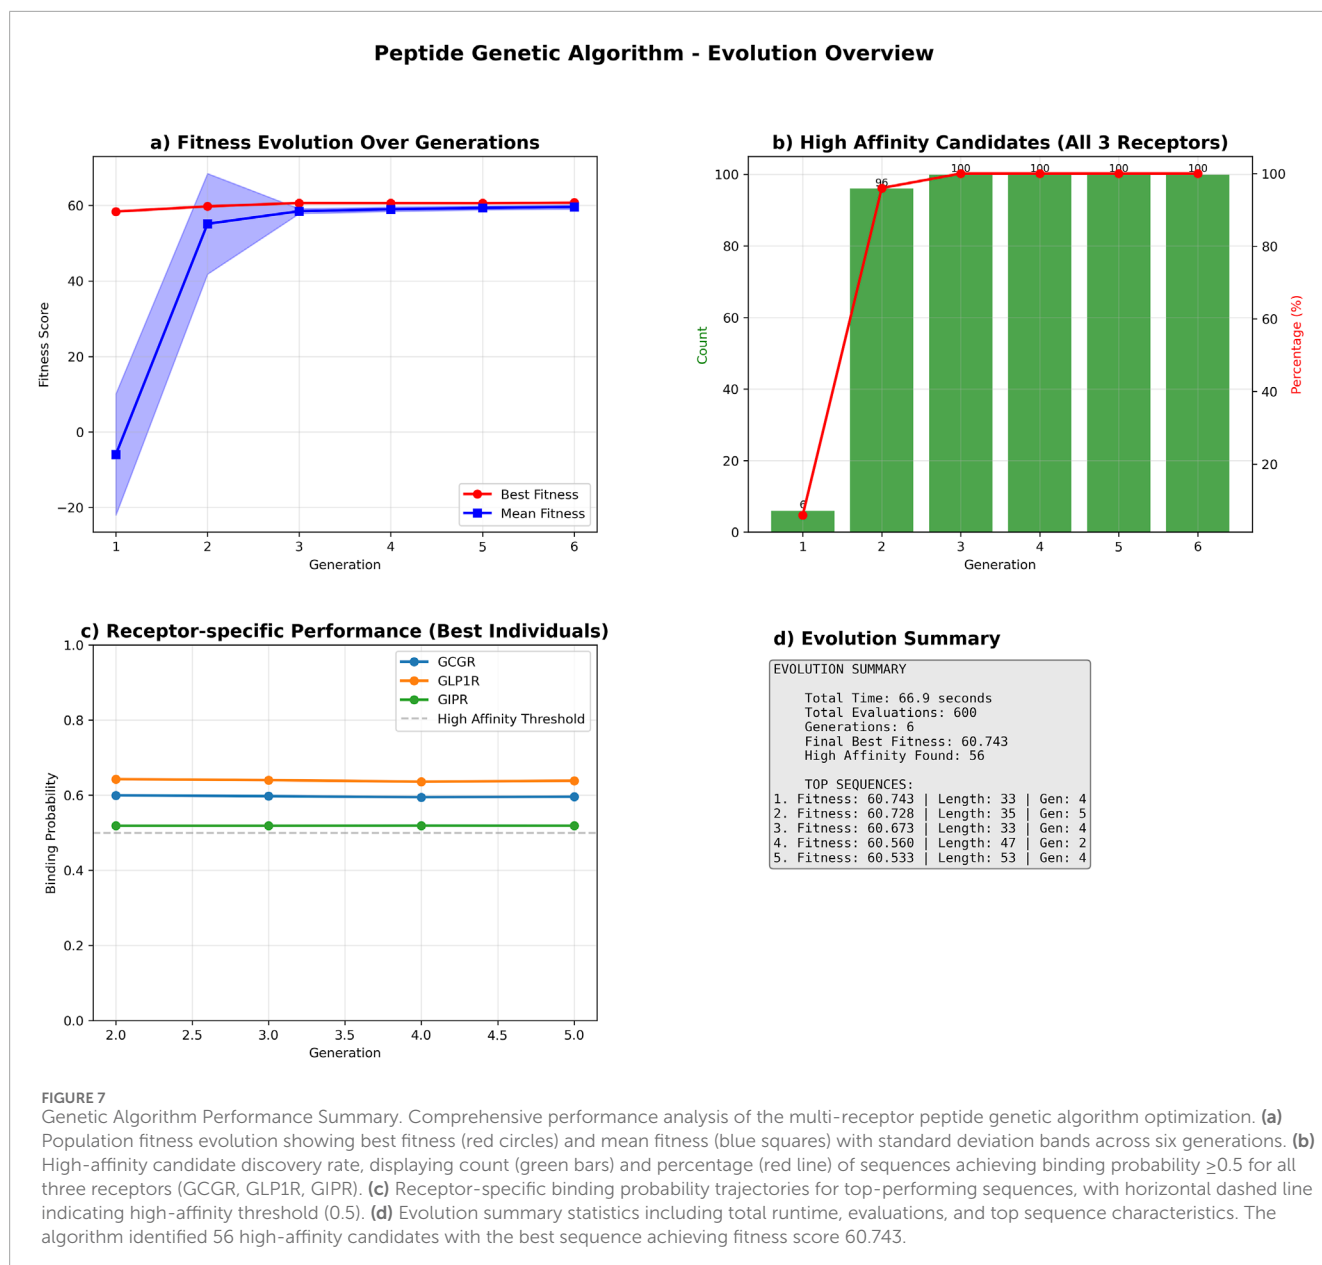
3.5 Synthetic peptide analysis

Analysis of the top 20 high-performing sequences revealed structural motifs and design principles underlying multi-receptor binding activity. The generated sequences maintained an average length of 32.7 ± 2.4 amino acids (range: 29–35), conforming to the length constraints of native incretin peptides. The sequences showed binding affinity predictions across all three target receptors, with mean binding probabilities of 0.596 ± 0.015 for GCGR, $0.638 \pm$

0.008 for GLP1R, and 0.519 ± 0.003 for GIPR. Similarity analysis against the training dataset revealed that the top 20 generated sequences maintained moderate divergence from known peptides, with sequence similarity ranging from 30.0% to 64.1% (mean: $47.0\% \pm 9.33\%$), ensuring both novelty and biological relevance in the designed multi-receptor agonists (Table 4).

Motif analysis identified conservation of biologically relevant sequence patterns. The essential core $E^2G^3T^4F^5$ motif, critical for incretin receptor binding, was preserved in 13 of 20 sequences (65%). A glucagon family C-terminus conserved ortholog A^{18} appeared in seven sequences (35%), while another essential core motif ($F^{21}W^{23}L^{24}$) was found in six sequences (30%) (Figure 8).

Biological plausibility assessment yielded overall scores of 0.796 ± 0.055 , with chemical plausibility scores of 0.960, motif preservation scores of 0.734 ± 0.107 , proteolytic stability scores of 0.837, and compositional scores of 0.560. Physicochemical characterization revealed the following properties: mean molecular weight of $3,794.5 \pm 185.8$ Da, average LogP value of -0.55 ± 0.22 , and polar surface area of 757.5 ± 116.2 Å. For comparison, native peptide values were: GLP-1 (3,298.6 Da, LogP -0.23), glucagon (3,482.7 Da, LogP -0.99), and GIP (4,983.5 Da, LogP -0.80) (Figures 9A–C). Isoelectric point values averaged 4.7 ± 0.59 , instability indices



averaged 29.2 ± 9.8 , and GRAVY scores confirmed hydrophilic character. Polar surface area values were lower than GIP (1,150.0 Å) but higher than GLP-1 (700.0 Å) (Figures 9D–F).

3.6 Computational performance

The GAT model training required approximately 2–3 h per fold on standard GPU hardware, representing a significant improvement in computational efficiency compared to structure-based design approaches. The genetic algorithm optimization completed within 1.5 h, enabling rapid exploration of sequence variants for experimental prioritization.

Memory requirements scaled linearly with sequence length, demonstrating the practical advantage of graph representations for variable-length peptide design. The attention mechanism provided

interpretable insights into residue importance, with attention weights correlating with known binding site interactions from crystallographic studies.

4 Discussion

The development of computational models capable of predicting peptide activity across multiple G-protein coupled receptors represents a critical advancement for metabolic disease therapeutics (Yang et al., 2019). Triple agonist peptides targeting glucagon receptor (GCGR), glucagon-like peptide-1 receptor (GLP1R), and glucose-dependent insulinotropic polypeptide receptor (GIPR) have demonstrated superior therapeutic efficacy compared to single or dual agonist approaches in treating type 2 diabetes and obesity (Samms et al., 2020). The synergistic

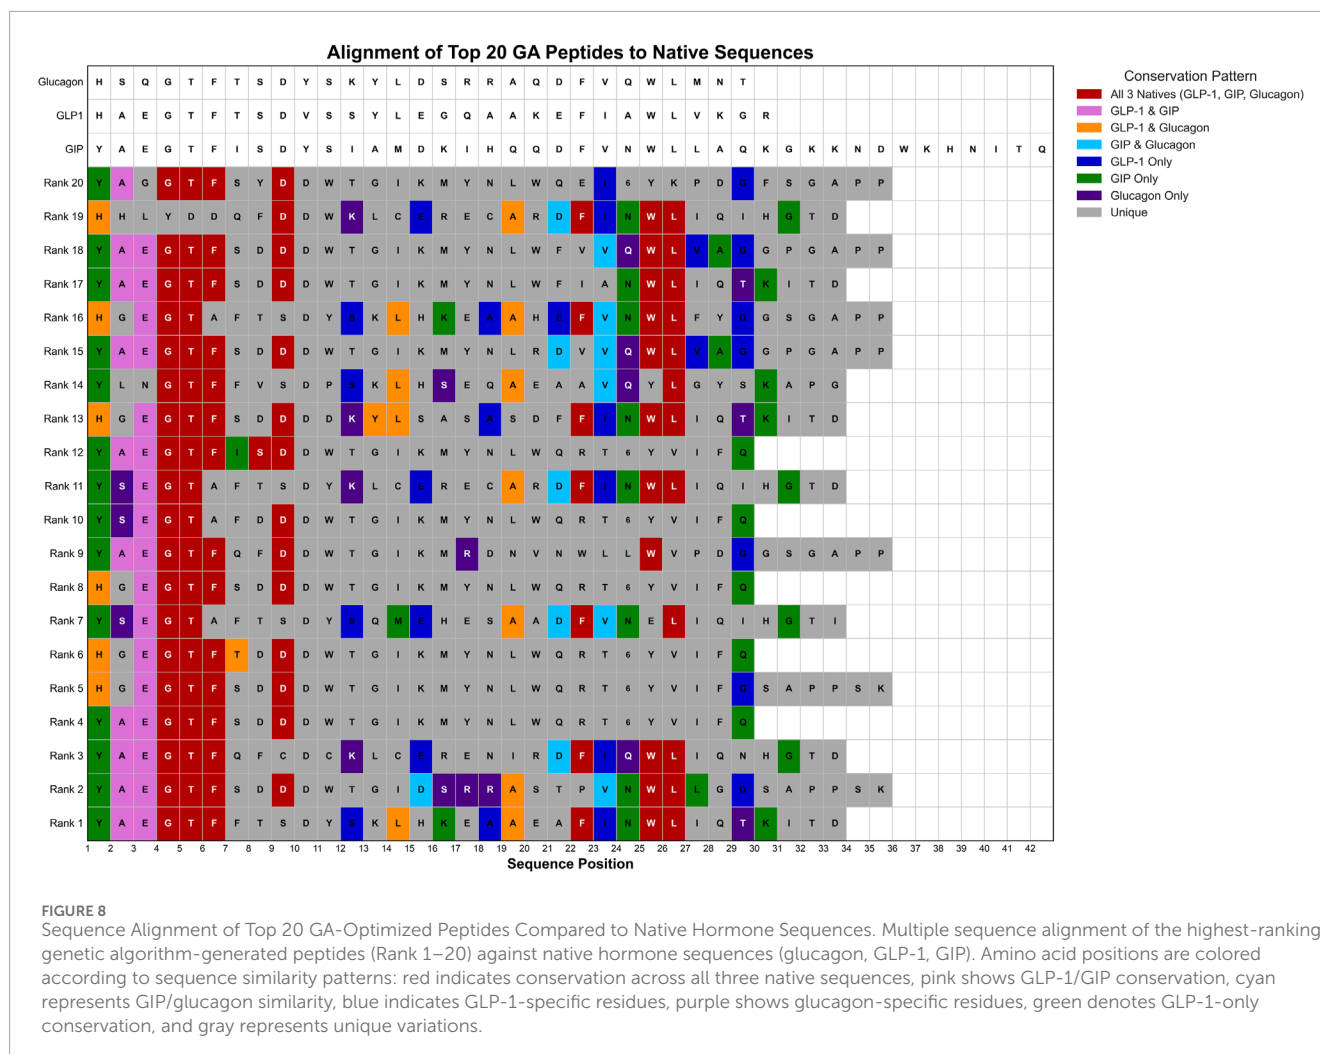
TABLE 4 Top 20 Generated Peptides by Predicted Multi-Receptor Activity and Physicochemical Properties. The table presents computationally generated peptides evaluated for their therapeutic potential against key metabolic hormone receptors. Sequence refers to the amino acid sequence in single-letter code. GCGR_Activity, GLP1R_Activity, and GIPR_Activity represent predicted binding activity (0–1 scale) against Glucagon Receptor, Glucagon-like Peptide-1 Receptor, and Glucose-dependent Insulinotropic Polypeptide Receptor, respectively, where one indicates highest predicted activity. Highest Similarity shows the maximum percentage similarity to training set sequences calculated using edit distance algorithm. Physicochemical properties include: mol_weight (estimated molecular weight in Da), isoelectric (predicted isoelectric point), instability index (computational measure of peptide stability *in vitro*), gravi (Grand Average of Hydropathy score indicating overall hydrophobicity), estimated_logP (predicted octanol-water partition coefficient for membrane permeability assessment), and estimated_PSA (estimated polar surface area in Å for bioavailability prediction).

Sequence	GCGR_ activity	GLP1R_ activity	GIPR_ activity	Highest similarity	mol_ weight	Isoelectric index	Gravi	Estimated_ logP	Estimated_ PSA
YAEGTFTSDYSKLEHKEAAEFIN WLIQTKITD	0.5941	0.6354	0.51907	54.5%	3839.22	5.01129	-0.3333	-0.3333	850
YAEGTFSDDDWTGIDRRASTPVNW LLGGSAPPSK	0.59582	0.63758	0.51879	53.8%	3754.98	4.36003	-0.7143	-0.7143	850
YAEGTFQFCDCKLCERENIRDFIQWLI QNHGTD	0.59285	0.6331	0.51893	45.5%	3994.4	4.69742	-0.6121	-0.6121	850
YAEGTFSDDDWTGKIMYNLWQRT [K (eK-eK-yE-C20DA)]YVIFQ	0.59932	0.64272	0.5185	34.5%	3576.94	4.68447	-0.6862	-0.6862	700
HGEGTFSDDDWTGKIMYNLWQRT [K (eK-eK-yE-C20DA)]YVIFGSAPPSK	0.59919	0.64196	0.51818	37.1% similarity	4033.44	6.75483	-0.7943	-0.7943	850
HGEGTFSDDDWTGKIMYNLWQRT [K (eK-eK-yE-C20DA)]YVIFQ	0.5991	0.6421	0.5183	37.9%	3550.91	5.4294	-0.8241	-0.8241	750
YSEGTAFTSDYSQMEHESAADFVNEL IQIHGTI	0.59367	0.63479	0.51908	48.5%	3691.89	4.07208	-0.3697	-0.3697	900
HGEGTFSDDDWTGKIMYNLWQRT [K (eK-eK-yE-C20DA)]YVIFQ	0.59959	0.64288	0.51834	41.4%	3536.88	5.4294	-0.8276	-0.8276	750
YAEGTFQFDDWTGKIMRDNVNWLWVPD GGSGAPP	0.59533	0.6369	0.5186	46.3%	3941.3	4.05003	-0.5086	-0.5086	650
YSEGTAFTDDDDWTGKIMYNLWQRT [K (eK-eK-yE-C20DA)]YVIFQ	0.59874	0.64213	0.51835	30.0%	3576.94	4.68447	-0.6862	-0.6862	700
YSEGTAFTSDYKLCERECARDFINWLI QIHGTD	0.59468	0.63679	0.51874	57.6%	3883.24	4.69742	-0.5091	-0.5091	850
YAEGTFSDDDWTGKIMYNLWQRT [K (eK-eK-yE-C20DA)]YVIFQ	0.59882	0.64252	0.5187	40.0%	3575.01	6.11897	-0.4103	-0.4103	650

(Continued on the following page)

TABLE 4 (Continued) Top 20 Generated Peptides by Predicted Multi-Receptor Activity and Physicochemical Properties. The table presents computationally generated peptides evaluated for their therapeutic potential against key metabolic hormone receptors. Sequence refers to the amino acid sequence in single-letter code. GCGR_Activity, GLP1R_Activity, and GIPR_Activity represent predicted binding activity (0–1 scale) against Glucagon Receptor, Glucagon-like Peptide-1 Receptor, and Glucose-dependent Insulinotropic Polypeptide Receptor, respectively, where one indicates highest predicted activity. Highest Similarity shows the maximum percentage similarity to training set sequences calculated using edit distance algorithm. Physicochemical properties include: mol_weight (estimated molecular weight in Da), isoelectric (predicted isoelectric point), instability index (computational measure of peptide stability *in vitro*), gravity (Grand Average of Hydropathy score indicating overall hydrophobicity), estimated_logP (predicted octanol-water partition coefficient for membrane permeability assessment), and estimated_PSA (estimated polar surface area in Å for bioavailability prediction).

sequence	GCGR_Activity	GLP1R_Activity	GIPR_Activity	Highest similarity	mol_weight	Isoelectric	Instability index	gravity	estimated_logP	estimated_PSA
HGECTFSDDDDKYLSASDFFINWLI QTKITD	0.59247	0.63218	0.51911	60.6%	3737.94	4.15507	-1.397	-0.5364	-0.5364	950
YLNGTFFVSPDKLHSEAEAAVQYLG YSKAPG	0.59509	0.63791	0.51859	45.5%	3575.89	5.44531	22.2242	-0.3515	-0.3515	700
YAEGTFSDDDDWTGIKMYNLRDVVQWLV AGPGGAPP	0.59467	0.63662	0.5189	53.8%	3827.19	4.05003	43.2229	-0.3286	-0.3286	600
HGEGTAFTSDYSKHLKEAAHEFVNWLF YGSGGAPP	0.59601	0.63832	0.51886	64.1%	3809.07	5.76946	23.6431	-0.5714	-0.5714	750
YAEGTFSDDDDWTGIKMYNLFIFIANWLI QTKITD	0.59495	0.63654	0.51917	54.5%	3957.38	4.19656	28.0212	-0.2667	-0.2667	750
YAEGTFSDDDDWTGIKMYNLFVFWQWLVA GGPGAPP	0.59485	0.63749	0.51924	56.4%	3889.3	4.05003	44.4857	-0.0457	-0.0457	500
HHLYDDQFDDWKLCECARDFFINWL IQIHGTD	0.59197	0.63246	0.51834	42.4%	4133.5	4.72022	25.4061	-0.8455	-0.8455	900
YAGGTFSYDDDDWTGIKMYNLIWQEI [K (eK-eK-yE-C20DA)]YKPDGFGGAPP	0.60014	0.64351	0.5183	35.9%	4004.39	4.68447	51.9971	-0.7314	-0.7314	650



effects observed with multi-receptor activation underscore the therapeutic potential of these molecules, yet their rational design remains computationally challenging due to the complex sequence-structure-activity relationships governing receptor selectivity and binding affinity (Sato et al., 2006).


Graph Attention Networks were selected for this task based on their demonstrated capability to capture complex relational dependencies within molecular structures (Veličković et al., 2017). Unlike traditional sequence-based approaches that treat peptides as linear strings of amino acids, GAT architectures can explicitly model spatial relationships and long-range interactions between residues through graph representations (Xiong et al., 2020). This offers potential advantages over conventional machine learning approaches that may fail to capture the three-dimensional nature of peptide-receptor interactions.

The GAT model demonstrated robust predictive performance across all three target receptors in cross-validation experiments. For GCGR, the model achieved excellent discrimination between high and low affinity peptides, with classification accuracy exceeding 88% and area under the receiver operating characteristic curve approaching 92%. GLP1R prediction showed similarly strong performance, with classification accuracy above 90% and

discrimination ability of approximately 85%. GIPR classification proved more challenging due to limited training data availability yet still achieved acceptable performance with classification accuracy around 82% and discrimination ability exceeding 90%.

Direct comparison with established computational approaches revealed distinct performance patterns across receptor targets. When evaluated against the multi-task convolutional neural network approach established by Puszkarska et al., the GAT model demonstrated improved performance for GCGR prediction, achieving significantly lower prediction errors and improved correlation with experimental values (Puszkarska et al., 2024). However, the traditional CNN approach maintained superior performance for GLP1R prediction, suggesting that optimal architectural choices may be receptor-dependent. These findings highlight the importance of comprehensive model comparisons rather than relying on single performance metrics or individual receptor assessments.

External validation on peptide sequences with limited similarity to training data provided insights into model generalizability. The GAT approach maintained excellent predictive capability for GCGR when applied to novel peptide sequences (AUC-ROC = 0.953 but with moderate classification performance (F1-score = 0.679).



GLP1R showed moderate discrimination ($\text{AUC-ROC} = 0.604$) with acceptable classification metrics ($\text{F1-score} = 0.787$). GIPR prediction on novel sequences demonstrated good discriminatory performance ($\text{AUC-ROC} = 0.818$) with high classification accuracy ($\text{F1-score} = 0.971$). This pattern suggests that GIPR shows robust generalizability to novel sequences, while GCGR may have more stringent classification requirements despite excellent discrimination ability.

The receptor-specific performance differences observed in this study reflect the underlying biological complexity of peptide-receptor interactions. GCGR demonstrated the most predictable binding patterns, potentially due to more stringent structural requirements for activation (Qiao et al., 2020). The improved GAT performance for this receptor suggests that graph-based representations effectively capture the key molecular features governing GCGR selectivity. In contrast, GLP1R showed different patterns where traditional CNN approaches maintained competitive performance, indicating that linear sequence features may be sufficient for this receptor class under certain conditions.

The robust GIPR performance on novel sequences, despite limited training data, suggests that the available data may be sufficient to capture key structure-activity relationships for this receptor. However, the smaller training dataset available for

GIPR compared to GCGR and GLP1R indicates that expanded datasets could further improve model confidence and performance consistency across all receptor types. Validation against clinically relevant peptides provided important insights into the model's performance on therapeutically important sequences. The GAT ensemble achieved strong overall accuracy across the three receptors when evaluated on nine established peptides, with particularly robust performance for GLP1R (100% accuracy) and good performance for GCGR (88.9% accuracy) and GIPR (77.8% accuracy). The model successfully predicted activities for six well-characterized dual and triple agonists including Retatrutide, Tirzepatide, and Cotadutide, demonstrating consistent performance on clinically advanced compounds.

The observed prediction errors offer valuable insights into model limitations. The misclassification of Pemvidutide's GCGR activity (predicted high vs. known low) and GIPR activities for both Pemvidutide and Mazdutide (predicted low vs. known high) suggests that certain structural features or receptor interaction modes may not be fully captured by the current training data. These discrepancies highlight the importance of expanding training datasets with diverse clinical candidates to improve model robustness across the full spectrum of therapeutic peptides. Nevertheless, the high accuracy achieved on the majority of

clinically relevant sequences supports the potential utility of this approach for prioritizing peptide candidates in drug discovery pipelines.

Analysis of computationally generated peptide sequences provides insights into the molecular determinants of multi-receptor binding activity. The generated peptides exhibited molecular characteristics consistent with known incretin hormone properties (Figure 9), with average sequence lengths and molecular weights falling within the expected range for bioactive peptide hormones (Müller et al., 2019). The predicted binding affinities across all three target receptors (mean probabilities: GCGR 0.596 ± 0.015 , GLP1R 0.638 ± 0.008 , GIPR 0.519 ± 0.003) suggest potential for balanced multi-receptor activation, though experimental validation remains necessary to confirm these computational predictions.

Motif analysis revealed partial conservation of established incretin receptor binding determinants. The preservation of the EGTF motif in approximately two-thirds of generated sequences aligns with its known importance for incretin receptor recognition (Kyte and Doolittle, 1982; Vaswani et al., 2017). However, the variable presence of other conserved regions, such as the glucagon family C-terminus motif (present in 35% of sequences), suggests that the model may identify alternative binding configurations that warrant experimental investigation. The biological plausibility scores, while generally favorable, indicate that computational optimization may generate sequences with non-natural characteristics that could affect stability or bioactivity.

Sequence similarity analysis revealed that the generated peptides maintained moderate divergence from the training dataset, with similarities ranging from 30.0% to 64.1% (mean: $47.0\% \pm 9.33\%$). This indicates that the optimization approach explores sequence variants within a reasonable distance from known agonists while avoiding excessive extrapolation beyond the model's training domain. The balance between sequence novelty and similarity to established agonists supports the approach as a systematic method for peptide optimization rather than *de novo* design.

The physicochemical properties of generated peptides fell within reasonable boundaries relative to native incretin hormones, though some parameters deviated from natural ranges. The intermediate molecular weights and hydrophobicity values suggest that the model attempts to balance the distinct physicochemical requirements of the three target receptors. However, the relatively low compositional scores indicate potential departures from natural amino acid distributions, which could impact peptide stability, immunogenicity, or pharmacokinetic properties in biological systems (Guruprasad et al., 1990). These results demonstrate a computational optimization pipeline that integrates sequence optimization with structural and physicochemical constraints to guide experimental validation efforts. The methodology provides a rational framework for candidate prioritization, with partially preserved key motifs and reasonable molecular properties suggesting these optimized sequences may warrant experimental evaluation to assess their functional characteristics.

This multi-target capability represents an alternative over single-receptor prediction tools, though the performance gains come with increased computational complexity and reduced interpretability for non-expert users. Existing simpler approaches may retain

advantages in specific applications. Rule-based peptide design tools offer greater transparency in decision-making processes and require substantially less computational resources (Cheng et al., 2024). Additionally, established pharmacophore-based methods may provide more reliable predictions for peptide modifications within well-characterized chemical space (Giordano et al., 2022). The GAT approach may be most suitable when exploring novel peptide sequences or optimizing across multiple targets simultaneously.

4.1 Limitations and challenges

Several limitations should be considered when interpreting these results. The training dataset comprises EC50 measurements from multiple laboratories using diverse assay conditions, cell lines, and experimental protocols. This inter-laboratory variability introduces systematic biases that may affect model predictions, as differences in receptor expression levels and measurement methodologies can significantly impact reported EC50 values. Additionally, the availability of publicly accessible datasets for triple agonist peptides remains limited, constraining our validation set size and highlighting the need for larger shared datasets to enable more robust validation of computational models in this specialized research area.

The graph-based molecular representation, while capturing local amino acid interactions, may not fully represent long-range conformational dependencies critical for receptor binding. The current feature encoding relies on static physicochemical properties but excludes dynamic structural information and context-specific amino acid interactions that influence binding affinity. Additionally, the model's reliance on sequence-derived features excludes critical three-dimensional structural information that governs binding specificity and selectivity. Finally, *in vitro* EC50 predictions do not encompass pharmacological properties essential for therapeutic development, including peptide stability, membrane permeability, proteolytic resistance, and pharmacokinetic profiles. These limitations highlight opportunities for future model enhancement through expanded training datasets, incorporation of structural features, and integration of pharmacokinetic modeling.

The computational methodology presented here primarily involves optimization of existing agonist sequence scaffolds rather than *de novo* creation of entirely novel peptide frameworks. The genetic algorithm systematically explores sequence variants within established incretin hormone design space, building upon known structural templates to identify improved variants. While this approach limits exploration to modifications around established peptide scaffolds, it provides a rational framework for systematic optimization that may identify therapeutically relevant improvements within well-characterized chemical space. This targeted optimization strategy balances computational tractability with biological relevance, though it may not uncover breakthrough therapeutic properties that could emerge from more radical structural innovations.

The absence of experimental validation in this computational study represents an important limitation for assessing practical therapeutic potential. While the models demonstrate robust predictive performance on available datasets and show promising predictions for clinically relevant peptides, the biological activity

and pharmacological properties of the computationally generated sequences remain to be confirmed. Experimental validation will be essential to establish the true therapeutic relevance of these computational predictions and to bridge the gap between *in silico* optimization and practical drug development applications.

Finally, *in vitro* EC50 predictions do not encompass pharmacological properties essential for therapeutic development, including peptide stability, membrane permeability, proteolytic resistance, and pharmacokinetic profiles. These limitations highlight opportunities for future enhancement through expanded training datasets, incorporation of structural features, integration of pharmacokinetic modeling, and experimental validation of computational predictions.

4.2 Implications and potential applications

Despite these limitations, the GAT-based predictor may offer practical applications for peptide drug discovery. The tool could facilitate initial screening of large peptide libraries to identify candidates with predicted multi-receptor activity, potentially reducing the experimental burden of comprehensive activity testing (Macarron et al., 2011). The multi-target prediction capability may prove particularly valuable for prioritizing experimental validation efforts. Rather than testing peptides sequentially against individual receptors, researchers could focus on candidates with favorable predictions across all target receptors (Anighoro et al., 2014). This approach could streamline the identification of balanced triple agonists while reducing resource requirements for preliminary screening.

Additionally, graph-based architecture could inform structure-activity relationship studies by highlighting amino acid positions critical for multi-receptor binding. This information could guide focused mutagenesis studies or assist in designing peptide analogs with improved pharmacological properties.

Several research directions could enhance the utility and reliability of computational triple agonist prediction. Expanding training datasets through systematic experimental characterization of peptide libraries would improve model robustness and generalizability. Integration of additional molecular descriptors, such as predicted secondary structure or dynamic conformational information, could enhance prediction accuracy.

The GAT architecture could be adapted for other multi-target therapeutic applications beyond incretin receptor agonists. Peptide hormones targeting multiple receptor families, such as opioid or neurotransmitter systems, may benefit from similar computational approaches (Stein and Machelska, 2011). Future model developments should address current limitations in capturing receptor dynamics and allosteric effects. Integration with molecular dynamics simulations or enhanced sampling techniques could provide more accurate representations of peptide-receptor interactions (Hospital et al., 2015).

5 Conclusion

The Graph Attention Network-based predictor represents a computational framework for identifying potential triple

agonist peptides targeting GCGR, GLP1R, and GIPR. The model demonstrated robust cross-validation performance and generated peptide sequences with biologically plausible characteristics and preserved functional motifs. However, significant limitations remain regarding data availability, model generalizability, and the translation from computational predictions to experimental validation. While the tool may assist in guiding initial peptide screening and rational design efforts, extensive experimental validation will be required to confirm biological activity and therapeutic utility. The approach establishes a foundation for computational multi-target peptide design that could be expanded and refined as additional training data and improved modeling techniques become available.

Data availability statement

The datasets analyzed and generated for this study can be found in the Triple Agonist Peptide Therapeutics for Metabolic Disease <https://github.com/aw449/Triple-Agonist-Peptide-Therapeutics-for-Metabolic-Disease>.

Author contributions

AW: Writing – review and editing, Investigation, Supervision, Writing – original draft, Conceptualization, Validation, Funding acquisition, Visualization, Data curation, Project administration, Formal Analysis, Methodology, Software, Resources. SG: Investigation, Writing – original draft, Writing – review and editing, Visualization, Methodology. TC: Validation, Investigation, Writing – review and editing, Methodology, Writing – original draft. KP: Conceptualization, Writing – review and editing, Supervision, Writing – original draft, Methodology.

Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fbinf.2025.1687617/full#supplementary-material>

SUPPLEMENTARY FIGURE S1

Residual Analysis and Prediction Distribution Comparison Between GAT and Multi-task CNN Models. Comprehensive residual analysis comparing GAT (blue circles) and CNN ensemble (red circles) predictions for EC50 logarithmic values. (A–C) GCGR Log(EC50) target analysis: (A) True versus predicted values with unity line (dashed) showing prediction accuracy, (B) residual distribution versus true values with zero-residual reference line (dashed), and (C) residual frequency distribution. (D–F) GLP1 Log(EC50) target analysis: (D) true versus predicted scatter plot, (E) residuals versus true values, and (F) residual distribution histograms. Residual patterns indicate model performance characteristics and potential systematic biases in prediction accuracy across the EC50 value range.

SUPPLEMENTARY FIGURE S2

Sequence Similarity Distribution Between Validation and Training Datasets. Histogram depicting the maximum sequence similarity of each validation sequence ($n = 58$) to any sequence in the training set. The distribution reveals a bimodal pattern with 37 sequences (67%) exhibiting $\leq 80\%$ similarity (red dashed line), and a median similarity of 79.5% (orange dashed line at 90%).

SUPPLEMENTARY FIGURE S3

Comprehensive Classification Metrics for GAT Model Validation Performance. Extended performance analysis comparing novel sequences ($\leq 80\%$ similarity, red) with the complete validation set (blue) across multiple classification metrics. (A) Accuracy values showing GCGR: 0.541 (novel) and 0.690 (complete); GLP1R: 0.649 (novel) and 0.759 (complete); GIPR: 0.943 (novel) and 0.929 (complete). (B) Precision values showing GCGR: 0.514 (novel) and 0.660 (complete); GLP1R:

0.649 (novel) and 0.772 (complete); GIPR: 0.943 (novel) and 0.944 (complete). (C) Recall values showing GCGR: 1.000 (novel) and 1.000 (complete); GLP1R: 1.000 (novel) and 0.978 (complete); GIPR: 1.000 (novel) and 0.971 (complete). (D) AUC-ROC values showing GCGR: 0.953 (novel) and 0.950 (complete); GLP1R: 0.604 (novel) and 0.358 (complete); GIPR: 0.818 (novel) and 0.943 (complete). The dashed line represents random classifier performance (AUC-ROC = 0.5).

SUPPLEMENTARY FIGURE S4

Genetic Algorithm Convergence Analysis. Convergence metrics for the peptide optimization genetic algorithm across six generations. (A) Best fitness score progression showing rapid initial improvement followed by plateau behavior. (B) Fitness improvement rate per generation with convergence threshold (0.1, red dashed line) and 3-generation moving average (orange). The algorithm achieved convergence when the moving average improvement rate fell below the threshold at generation 6. (C) Cumulative fitness improvement from baseline, demonstrating total optimization gain of 2.2 fitness units. Convergence was achieved after 6 generations.

SUPPLEMENTARY TABLE S1

Complete Training Dataset. The dataset contains peptide identifiers (pep_ID), amino acid sequences with non-standard modifications in brackets, experimentally determined EC50 values in picomolar units for GCGR, GLP1R, and GIPR receptors, corresponding log-transformed EC50 values, and literature references.

SUPPLEMENTARY TABLE S2

Nonstandard Amino Acids and their calculated physicochemical properties with ProtParam. Physicochemical properties calculated using BioPython ProtParam v1.79 for non-standard amino acids and chemical modifications encountered in the dataset, including D-amino acids and unnatural amino acids (Aib, Nle, Sar, hSer, nVal, Dap, MetO) with corresponding hydrophobicity values, molecular weights, and net charges at physiological pH.

SUPPLEMENTARY TABLE S3

Validation Dataset. The validation dataset contains peptide identifiers (pep_ID), amino acid sequences, experimentally determined EC50 values in picomolar units (pM) for GCGR, GLP1R, and GIPR receptors, corresponding log-transformed EC50 values, and source publications.

SUPPLEMENTARY TABLE S4

FDA-Approved and Clinical-Stage Multi-Receptor Agonists. Current FDA-approved therapeutics and compounds in clinical development for metabolic diseases with documented receptor binding classifications for GCGR, GLP1R, and GIPR, including peptide sequences, therapeutic descriptions dual-, or triple-agonists.

References

- Alexander, S. P., Christopoulos, A., Davenport, A. P., Kelly, E., Mathie, A., Peters, J. A., et al. (2021). The concise guide to pharmacology 2021/22: g protein-coupled receptors. *Br. J. Pharmacol.* 178 (Suppl. 1), S27–S156. doi:10.1111/bph.15538
- Anighoro, A., Bajorath, J., and Rastelli, G. (2014). Polypharmacology: challenges and opportunities in drug discovery. *J. Med. Chem.* 57 (19), 7874–7887. doi:10.1021/jm5006463
- Baggio, L. L., and Drucker, D. J. (2021). Glucagon-like peptide-1 receptor co-agonists for treating metabolic disease. *Mol. Metab.* 46, 101090. doi:10.1016/j.molmet.2020.101090
- Boudreau, D. M., Malone, D. C., Raebel, M. A., Fishman, P., Ph.D., Nichols, G., Ph.D., Feldstein, A., M.D., M.S., et al. (2009). Health care utilization and costs by metabolic syndrome risk factors. *Metab. Syndr. Relat. Disord.* 7 (4), 305–314. doi:10.1089/met.2008.0070
- Brandt, S. J., Götz, A., Tschöpp, M. H., and Müller, T. D. (2018). Gut hormone polyagonists for the treatment of type 2 diabetes. *Peptides* 100, 190–201. doi:10.1016/j.peptides.2017.12.021
- Bromer, W. W., Sinn, L. G., Staub, A., and Behrens, O. K. (1957). The amino acid sequence of glucagon. *Diabetes* 6 (3), 234–238. doi:10.2337/diab.6.3.234
- Chen, T., Sun, T., Bian, Y., Pei, Y., Feng, F., Chi, H., et al. (2022). The design and optimization of monomeric multitarget peptides for the treatment of multifactorial diseases. *J. Med. Chem.* 65 (5), 3685–3705. doi:10.1021/acs.jmedchem.1c01456
- Cheng, Z., He, B. B., Lei, K., Gao, Y., Shi, Y., Zhong, Z., et al. (2024). Rule-based omics mining reveals antimicrobial macrocyclic peptides against drug-resistant clinical isolates. *Nat. Commun.* 15, 4901. doi:10.1038/s41467-024-49215-y
- Cock, P. J. A., Antao, T., Chang, J. T., Chapman, B. A., Cox, C. J., Dalke, A., et al. (2009). Biopython: freely available python tools for computational molecular biology and bioinformatics. *Bioinformatics* 25 (11), 1422–1423. doi:10.1093/bioinformatics/btp163
- Coskun, T., Urva, S., Roell, W. C., Qu, H., Loghini, C., Moyers, J. S., et al. (2022). LY3437943, a novel triple glucagon, GIP, and GLP-1 receptor agonist for glycemic control and weight loss: from discovery to clinical proof of concept. *Cell Metab.* 34 (9), 1234–1247.e9. doi:10.1016/j.cmet.2022.07.013
- Day, J. W., Ottaway, N., Patterson, J. T., Gelfanov, V., Smiley, D., Gidda, J., et al. (2009). A new glucagon and GLP-1 co-agonist eliminates obesity in rodents. *Nat. Chem. Biol.* 5, 749–757. doi:10.1038/nchembio.209
- Ertl, P., Rohde, B., and Selzer, P. (2000). Fast calculation of molecular polar surface area as a sum of fragment-based contributions and its application to the prediction of drug transport properties. *J. Med. Chem.* 43 (20), 3714–3717. doi:10.1021/jm000942e
- Evers, A., Bossart, M., Pfeiffer-Marek, S., Elvert, R., Schreuder, H., Kurz, M., et al. (2018). Dual glucagon-like peptide 1 (GLP-1)/glucagon receptor agonists specifically optimized for multidose formulations. *J. Med. Chem.* 61 (13), 5580–5593. doi:10.1021/acs.jmedchem.8b00292
- Fey, M., and Lenssen, J. E. (2019). Fast graph representation learning with PyTorch geometric.

- Finan, B., Yang, B., Ottaway, N., Smiley, D. L., Ma, T., Clemmensen, C., et al. (2015). A rationally designed monomeric peptide triagonist corrects obesity and diabetes in rodents. *Nat. Med.* 21 (1), 27–36. doi:10.1038/nm.3761
- Finan, B., Müller, T. D., Clemmensen, C., Perez-Tilve, D., DiMarchi, R. D., and Tschöp, M. H. (2016). Reappraisal of GIP pharmacology for metabolic diseases. *Trends Mol. Med.* 22 (5), 359–376. doi:10.1016/j.molmed.2016.03.005
- Finan, B., Douros, J. D., Goldwater, R., Hansen, A. M. K., Hjerpested, J. B., Hjøllund, K. R., et al. (2025). A once-daily GLP-1/GIP/glucagon receptor tri-agonist (NN1706) lowers body weight in rodents, monkeys and humans. *Mol. Metab.* 96, 102129. doi:10.1016/j.molmet.2025.102129
- Frias, J. P., Nauck, M. A., Van, J., Kutner, M. E., Cui, X., Benson, C., et al. (2018). Efficacy and safety of LY3298176, a novel dual GIP and GLP-1 receptor agonist, in patients with type 2 diabetes: a randomised, placebo-controlled and active comparator-controlled phase 2 trial. *Lancet* 392 (10160), 2180–2193. doi:10.1016/s0140-6736(18)32260-8
- Gao, H., Wang, Z., and Ji, S. (2018). “Large-scale learnable graph convolutional networks,” in *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery and data mining*, 1416–1424.
- Giordano, D., Biancanello, C., Argenio, M. A., and Facchiano, A. (2022). Drug design by pharmacophore and virtual screening approach. *Pharm. (Basel)* 15 (5), 646. doi:10.3390/ph15050646
- Guruprasad, K., Reddy, B. V., and Pandit, M. W. (1990). Correlation between stability of a protein and its dipeptide composition: a novel approach for predicting *in vivo* stability of a protein from its primary sequence. *Protein Eng.* 4 (2), 155–161. doi:10.1093/protein/4.2.155
- Han, J., Sun, L., Chu, Y., Li, Z., Huang, D., Zhu, X., et al. (2013). Design, synthesis, and biological activity of novel dicoumarol glucagon-like peptide 1 conjugates. *J. Med. Chem.* 56 (24), 9955–9968. doi:10.1021/jm4017448
- Henderson, S. J., Konkar, A., Hornigold, D. C., Trevaskis, J. L., Jackson, R., Fritsch Fredin, M., et al. (2016). Robust anti-obesity and metabolic effects of a dual GLP-1/glucagon receptor peptide agonist in rodents and non-human Primates. *Diabetes Obes. Metab.* 18 (12), 1176–1190. doi:10.1111/dom.12735
- Hospital, A., Goñi, J. R., Orozco, M., and Gelpi, J. L. (2015). Molecular dynamics simulations: advances and applications. *Adv. Appl. Bioinform Chem.* 8, 37–47. doi:10.2147/aabc.s70333
- Jastreboff, A. M., Aronne, L. J., Ahmad, N. N., Wharton, S., Connery, L., Alves, B., et al. (2022). Tirzepatide once weekly for the treatment of obesity. *N. Engl. J. Med.* 387 (3), 205–216. doi:10.1056/nejmoa2206038
- Knerr, P. J., Finan, B., Gelfanov, V., Perez-Tilve, D., Tschöp, M. H., and DiMarchi, R. D. (2018). Optimization of peptide-based polyagonists for treatment of diabetes and obesity. *Bioorg Med. Chem.* 26 (10), 2873–2881. doi:10.1016/j.bmc.2017.10.047
- Kipf, T. N., and Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *arXiv Prepr. arXiv:1609.02907*. doi:10.48550/arXiv.1609.02907
- Knerr, P. J., Finan, B., Gelfanov, V., Perez-Tilve, D., Tschöp, M. H., and DiMarchi, R. D. (2020). Selection and progression of unimolecular agonists at the GIP, GLP-1 and glucagon receptors as drug candidates. *Peptides* 125, 170225. doi:10.1016/j.peptides.2019.170225
- Knerr, P. J., Mowery, S. A., Douros, J. D., Premjee, B., Hjøllund, K. R., He, Y., et al. (2022). Next generation GLP-1/GIP/glucagon triple agonists normalize body weight in Obese mice. *Mol. Metab.* 63, 101533. doi:10.1016/j.molmet.2022.101533
- Knudsen, L. B., Nielsen, P. F., Huusfeldt, P. O., Johansen, N. L., Madsen, K., Pedersen, F. Z., et al. (2000). Potent derivatives of glucagon-like peptide-1 with pharmacokinetic properties suitable for once daily administration. *J. Med. Chem.* 43 (9), 1664–1669. doi:10.1021/jm9909645
- Kyte, J., and Doolittle, R. F. (1982). A simple method for displaying the hydropathic character of a protein. *J. Mol. Biol.* 157 (1), 105–132. doi:10.1016/0022-2836(82)90515-0
- Lau, J., Bloch, P., Schäffer, L., Pettersson, I., Spetzler, J., Kofoed, J., et al. (2015). Discovery of the once-weekly glucagon-like Peptide-1 (GLP-1) analogue semaglutide. *J. Med. Chem.* 58 (18), 7370–7380. doi:10.1021/acs.jmedchem.5b00726
- Liang, X., Or, B., Tsoi, M. F., Cheung, C. L., and Cheung, B. M. Y. (2023). Prevalence of metabolic syndrome in the United States national health and nutrition examination survey 2011–18. *Postgrad. Med. J.* 99 (1175), 985–992. doi:10.1093/postmj/qgad008
- Lv, Q., Chen, G., Yang, Z., Zhong, W., and Chen, C. Y. (2024). Meta learning with graph attention networks for low-data drug discovery. *IEEE Trans. Neural Netw. Learn. Syst.* 35 (8), 11218–11230. doi:10.1109/TNNLS.2023.3250324
- Macarron, R., Banks, M. N., Bojanic, D., Burns, D. J., Cirovic, D. A., Garyantes, T., et al. (2011). Impact of high-throughput screening in biomedical research. *Nat. Rev. Drug Discov.* 10 (3), 188–195. doi:10.1038/nrd3368
- Manandhar, B., and Ahn, J. M. (2015). Glucagon-like peptide-1 (GLP-1) analogs: recent advances, new possibilities, and therapeutic implications. *J. Med. Chem.* 58 (3), 1020–1037. doi:10.1021/jm500810s
- Mapelli, C., Natarajan, S. I., Meyer, J. P., Bastos, M. M., Bernatowicz, M. S., Lee, V. G., et al. (2009). Eleven amino acid glucagon-like peptide-1 receptor agonists with antidiabetic activity. *J. Med. Chem.* 52 (23), 7788–7799. doi:10.1021/jm900752a
- Mendez, D., Gaulton, A., Bento, A. P., Chambers, J., De Veij, M., Félix, E., et al. (2019). ChEMBL: towards direct deposition of bioassay data. *Nucleic Acids Res.* 47 (D1), D930–D940. doi:10.1093/nar/gky1075
- Meng, H., Krishnaji, S. T., Beinborn, M., and Kumar, K. (2008). Influence of selective fluorination on the biological activity and proteolytic stability of glucagon-like peptide-1. *J. Med. Chem.* 51 (22), 7303–7307. doi:10.1021/jm8008579
- Miranda, L. P., Winters, K. A., Gegg, C. V., Patel, A., Aral, J., Long, J., et al. (2008). Design and synthesis of conformationally constrained glucagon-like peptide-1 derivatives with increased plasma stability and prolonged *in vivo* activity. *J. Med. Chem.* 51 (9), 2758–2765. doi:10.1021/jm701522b
- Moody, A. J., Thim, L., and Valverde, I. (1984). The isolation and sequencing of human gastric inhibitory peptide (GIP). *FEBS Lett.* 172 (2), 142–148. doi:10.1016/0014-5793(84)81114-x
- Moon, M. J., Park, S., Kim, D. K., Cho, E. B., Hwang, J. I., Vaudry, H., et al. (2012). Structural and molecular conservation of glucagon-like Peptide-1 and its receptor confers selective ligand-receptor interaction. *Front. Endocrinol. (Lausanne)* 3, 141. doi:10.3389/fendo.2012.00141
- Mottillo, S., Filion, K. B., Genest, J., Joseph, L., Pilote, L., Poirier, P., et al. (2010). The metabolic syndrome and cardiovascular risk: a systematic review and meta-analysis. *J. Am. Coll. Cardiol.* 56 (14), 1113–1132. doi:10.1016/j.jacc.2010.05.034
- Müller, T. D., Finan, B., Bloom, S. R., D'Alessio, D., Drucker, D. J., Flatt, P. R., et al. (2019). Glucagon-like peptide 1 (GLP-1). *Mol. Metab.* 30, 72–130. doi:10.1016/j.molmet.2019.09.010
- Murage, E. N., Schroeder, J. C., Beinborn, M., and Ahn, J. M. (2008). Search for alpha-helical propensity in the receptor-bound conformation of glucagon-like peptide-1. *Bioorg Med. Chem.* 16 (23), 10106–10112. doi:10.1016/j.bmc.2008.10.006
- Orskov, C., Bersani, M., Johnsen, A. H., Højrup, P., and Holst, J. J. (1989). Complete sequences of glucagon-like peptide-1 from human and pig small intestine. *J. Biol. Chem.* 264 (22), 12826–12829. doi:10.1016/s0021-9258(18)51561-1
- Plisson, F., Hill, T. A., Mitchell, J. M., Hoang, H. N., de Araujo, A. D., Xu, W., et al. (2017). Helixconstraints and amino acid substitution in GLP-1 increase cAMP and insulin secretion but not beta-arrestin 2 signaling. *Eur. J. Med. Chem.* 127, 703–714. doi:10.1016/j.ejmech.2016.10.044
- Puszkarska, A. M., Taddese, B., Revell, J., Davies, G., Field, J., Hornigold, D. C., et al. (2024). Machine learning designs new GCGR/GLP-1R dual agonists with enhanced biological potency. *Nat. Chem.* 16 (9), 1436–1444. doi:10.1038/s41557-024-01532-x
- Qiao, A., Han, S., Li, X., Li, Z., Zhao, P., Dai, A., et al. (2020). Structural basis of Gs and Gi recognition by the human glucagon receptor. *Science* 367 (6484), 1346–1352. doi:10.1126/science.aaz5346
- Samms, R. J., Coghlan, M. P., and Sloop, K. W. (2020). How May GIP enhance the therapeutic efficacy of GLP-1? *Trends Endocrinol. Metab.* 31 (6), 410–421. doi:10.1016/j.tem.2020.02.006
- Sato, A. K., Viswanathan, M., Kent, R. B., and Wood, C. R. (2006). Therapeutic peptides: technological advances driving peptides into development. *Curr. Opin. Biotechnol.* 17 (6), 638–642. doi:10.1016/j.copbio.2006.10.002
- Starling, S. (2022). A new GLP1, GIP and glucagon receptor triagonist. *Nat. Rev. Endocrinol.* 18 (3), 135. doi:10.1038/s41574-021-00631-w
- Stein, C., and Machelska, H. (2011). Modulation of peripheral sensory neurons by the immune system: implications for pain therapy. *Pharmacol. Rev.* 63 (4), 860–881. doi:10.1124/pr.110.003145
- Strokach, A., Becerra, D., Corbi-Verge, C., Perez-Riba, A., and Kim, P. M. (2020). Fast and flexible protein design using deep graph neural networks. *Cell Syst.* 11 (4), 402–411.e4. doi:10.1016/j.cels.2020.08.016
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention is all you need. *Adv. Neural Inf. Process. Syst.* 30. doi:10.48550/arXiv.1706.03762
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. (2017). Graph attention networks. *arXiv Prepr. arXiv:1710.10903*. doi:10.48550/arXiv.1710.10903
- Whitley, H. P. (2025). Triple agonist therapy: a new frontier in treating type 2 diabetes and obesity. *Clin. Diabetes* 43 (3), 439–441. doi:10.2337/cd25-0018
- Willard, F. S., Douros, J. D., Gabe, M. B., Showalter, A. D., Waincott, D. B., Suter, T. M., et al. (2020). Tirzepatide is an imbalanced and biased dual GIP and GLP-1 receptor agonist. *JCI Insight* 5 (17), e140532. doi:10.1172/jci.insight.140532
- Winther, J. B., and Holst, J. J. (2024). Glucagon agonism in the treatment of metabolic diseases including type 2 diabetes mellitus and obesity. *Diabetes Obes. Metab.* 26 (9), 3501–3512. doi:10.1111/dom.15693
- Wiśniewski, K., Sueiras-Diaz, J., Jiang, G., Galyean, R., Lu, M., Thompson, D., et al. (2016). Synthesis and pharmacological characterization of novel glucagon-like Peptide-2 (GLP-2) analogues with low systemic clearance. *J. Med. Chem.* 59 (7), 3129–3139. doi:10.1021/acs.jmedchem.5b01909
- Wouters, O. J., McKee, M., and Luyten, J. (2020). Estimated research and development investment needed to bring a new medicine to market, 2009–2018. *JAMA* 323 (9), 844–853. doi:10.1001/jama.2020.1166

- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Philip, S. Y. (2020). A comprehensive survey on graph neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* 32 (1), 4–24. doi:10.1109/tnnls.2020.2978386
- Xiong, Z., Wang, D., Liu, X., Zhong, F., Wan, X., Li, X., et al. (2020). Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. *J. Med. Chem.* 63 (16), 8749–8760. doi:10.1021/acs.jmedchem.9b00959
- Yabut, J. M., and Drucker, D. J. (2023). Glucagon-like Peptide-1 receptor-based therapeutics for metabolic liver disease. *Endocr. Rev.* 44 (1), 14–32. doi:10.1210/endrev/bnac018
- Yang, K. K., Wu, Z., and Arnold, F. H. (2019). Machine-learning-guided directed evolution for protein engineering. *Nat. Methods* 16 (8), 687–694. doi:10.1038/s41592-019-0496-6
- Zhang, Z., Chen, L., Zhong, F., Wang, D., Jiang, J., Zhang, S., et al. (2022). Graph neural network approaches for drug-target interactions. *Curr. Opin. Struct. Biol.* 73, 102327. doi:10.1016/j.sbi.2021.102327
- Zhang, J., Xu, H., Dong, Y., Feng, J., and Li, R. (2025). Design, synthesis, and structure-activity relationship study of novel GLP-1/GIP/GCG triple receptor agonists. *Eur. J. Med. Chem.* 298, 118024. doi:10.1016/j.ejmech.2025.118024
- Zhao, F., Zhou, Q., Cong, Z., Hang, K., Zou, X., Zhang, C., et al. (2022). Structural insights into multiplexed pharmacological actions of tirzepatide and peptide 20 at the GIP, GLP-1 or glucagon receptors. *Nat. Commun.* 13 (1), 1057. doi:10.1038/s41467-022-28683-0
- Zimmermann, T., Thomas, L., Baader-Pagler, T., Haebel, P., Simon, E., Reindl, W., et al. (2022). BI 456906: discovery and preclinical pharmacology of a novel GCGR/GLP-1R dual agonist with robust anti-obesity efficacy. *Mol. Metab.* 66, 101633. doi:10.1016/j.molmet.2022.101633