

ORIGINAL RESEARCH

Open access

Generative Flow Network for Designing Novel Antimicrobial Peptides Targeting Multidrug-Resistant Gram-Negative Bacteria with Predicted Low Toxicity to Human Cells

Oliver Schmidt^{1*}, Lukas Weber¹, Jonas Richter²

Abstract

Antimicrobial resistance (AMR) is a major global health crisis, particularly due to multidrug-resistant gram-negative bacteria that are difficult to treat because of their impermeable outer membrane and strong efflux mechanisms, making antimicrobial peptides (AMPs) a promising alternative owing to their broad-spectrum activity and rapid bactericidal effects, although their clinical translation is limited by instability, production costs, and toxicity to human cells; meanwhile, traditional experimental discovery of AMPs is slow and expensive, and existing computational methods often optimize only antimicrobial activity while neglecting toxicity or diversity, leading to unsafe or narrow solutions, while reinforcement learning approaches may suffer from mode collapse and limited exploration of sequence space; to address these limitations, this work proposes a generative flow networks (GFlowNets)-based framework for multi-objective AMP design against gram-negative pathogens, in which a sequence generator constructs peptides stepwise and is guided by a reward function that integrates predicted antimicrobial activity and human cell toxicity from separate machine learning models, enabling simultaneous optimization of efficacy and safety; unlike conventional generative models, GFlowNets sample sequences proportional to reward, promoting diverse outputs that span the Pareto frontier of activity-toxicity trade-offs, while also allowing conditioning on desired properties and modular improvement of predictive components over time; overall, this framework provides a principled and scalable approach to antimicrobial peptide discovery that balances potency, safety, and diversity, potentially accelerating the identification of therapeutic candidates for combating antimicrobial resistance.

Keywords Generative flow networks, Antimicrobial peptides, Multidrug-resistant gram-negative bacteria, Multi-objective optimization, Peptide toxicity prediction, Deep generative models

*Correspondence:

Oliver Schmidt
oliver.schmidt@gmail.com

¹ Department of Artificial Intelligence in Healthcare, University of Stuttgart, Stuttgart, Germany

² Department of Clinical Data Intelligence, Technical University of Darmstadt, Darmstadt, Germany

Introduction

Antimicrobial resistance has escalated into a global health catastrophe, with an estimated 1.27 million deaths directly attributable to drug-resistant bacterial infections in 2019 alone, a figure projected to reach 10 million annually by 2050 without decisive intervention [1]. Among the most

alarming contributors to this crisis are the ESKAPE pathogens—*Enterococcus faecium*, *Staphylococcus aureus*, *Klebsiella pneumoniae*, *Acinetobacter baumannii*, *Pseudomonas aeruginosa*, and *Enterobacter* species—of which the gram-negative members (*K. pneumoniae*, *A. baumannii*, *P. aeruginosa*) together with *Escherichia coli* have been designated critical priority pathogens by the

World Health Organization [2]. These organisms possess a characteristic outer membrane enriched with lipopolysaccharide that acts as a formidable permeability barrier, coupled with multidrug efflux pumps and an expanding arsenal of acquired β -lactamases capable of hydrolyzing carbapenems and other last-resort antibiotics [3]. The dwindling pipeline of novel antibiotic classes, combined with regulatory and economic disincentives for pharmaceutical development, has created an urgent need for fundamentally new therapeutic modalities that can circumvent existing resistance mechanisms [4].

Antimicrobial peptides represent one of the most extensively studied alternative therapeutic classes, comprising evolutionarily ancient host defense molecules found across virtually all domains of life [5]. These typically cationic and amphipathic peptides, ranging from 10 to 50 amino acid residues, exert their bactericidal effects primarily through physical disruption of bacterial membranes—a mechanism that is inherently less susceptible to the target-site mutations conferring resistance to conventional antibiotics [6]. Beyond membrane lysis, many AMPs engage intracellular targets including nucleic acids and protein synthesis machinery, providing multiple, simultaneous mechanisms of action that further reduce resistance probability [7]. Several naturally occurring AMPs have progressed to clinical evaluation, including pexiganan (a magainin analog), omiganan (an indolicidin derivative), and the gram-negative-targeting polymyxins, yet widespread adoption has been limited by concerns over systemic toxicity, particularly hemolysis of red blood cells and cytotoxicity toward mammalian cell lines at or near therapeutic concentrations [8].

The fundamental design challenge confronting AMP engineering is the tension between maximizing antimicrobial potency and minimizing collateral damage to human cells—a multi-objective optimization problem embedded within a vast sequence space whose size scales as 20^L , where L represents peptide length [9]. High-throughput experimental screening can explore only a minuscule fraction of this space, motivating the development of computational models capable of prioritizing sequences for synthesis [10]. Traditional machine learning approaches have focused on binary classification or regression tasks using hand-crafted sequence features such as amino acid composition, net charge, and hydrophobicity, but these methods do not inherently generate novel sequences [11]. More recently, generative deep learning models including variational

autoencoders, generative adversarial networks, and autoregressive language models have been applied to peptide design, yet these approaches typically optimize a single property without formal mechanisms for balancing multiple objectives or maintaining output diversity [12, 13]. Furthermore, reinforcement learning formulations that maximize a reward function through policy gradient methods often exhibit mode collapse, producing high-scoring but structurally homogeneous peptides that fail to explore potentially superior regions of property space [14].

This article proposes a conceptual framework that repositions antimicrobial peptide design as a generative flow network (GFlowNet) problem, where the objective is not to find a single optimal sequence but to learn a stochastic policy that samples diverse peptides from across the high-reward regions of a multi-objective landscape [15]. The GFlowNet paradigm, originally formulated for diverse candidate generation in drug discovery and biological sequence design, learns to produce objects with probability proportional to a scalar reward, thereby populating the full Pareto frontier between competing objectives rather than collapsing to a point estimate [16]. Within this framework, peptide sequences are constructed incrementally through a Markov decision process, with a forward policy trained via trajectory balance to match the target reward distribution, while independent neural network predictors supply activity and toxicity scores as reward components [17]. The envisioned system targets multidrug-resistant gram-negative bacteria specifically, leveraging curated databases of sequence-activity relationships to train property predictors, and incorporates explicit consideration of hemolytic and cytotoxic endpoints to guide generation toward safe therapeutic candidates [18, 19]. The subsequent sections delineate the biological context motivating this work, the mathematical foundations of GFlowNets, the architectural components of the proposed framework, and a structured evaluation strategy for assessing generated peptide quality—all presented as a generative design blueprint awaiting empirical implementation.

Background

Multidrug-resistant gram-negative bacteria

The gram-negative cell envelope constitutes an evolutionary masterpiece of defensive architecture,

comprising an inner cytoplasmic membrane, a thin peptidoglycan layer, and a distinctive outer membrane whose asymmetric bilayer positions lipopolysaccharide molecules at the extracellular face [8]. This outer membrane functions as a selective permeability barrier, excluding hydrophobic compounds that readily penetrate gram-positive organisms while permitting nutrient uptake through aqueous porin channels, the expression and selectivity of which can be modulated under antibiotic pressure [9]. Clinical isolates of *Escherichia coli*, *Klebsiella pneumoniae*, *Pseudomonas aeruginosa*, and *Acinetobacter baumannii* have accumulated multiple, often coexisting resistance determinants: extended-spectrum β -lactamases (CTX-M, SHV, TEM variants), carbapenemases (KPC, NDM, OXA-48 families), aminoglycoside-modifying enzymes, fluoroquinolone target mutations, and overexpression of resistance-nodulation-division efflux pumps such as AcrAB-TolC in Enterobacteriaceae and MexAB-OprM in *P. aeruginosa* [20]. The convergence of these mechanisms has produced extensively drug-resistant and pandrug-resistant phenotypes for which standard-of-care antibiotics exhibit no clinically meaningful activity, compelling reliance on toxic salvage regimens including colistin (polymyxin E), which itself faces erosion of efficacy due to the emergence of plasmid-mediated mcr resistance genes [21]. Against this backdrop, antimicrobial peptides capable of physically disrupting the gram-negative outer membrane—particularly those that competitively displace stabilizing divalent cations from lipopolysaccharide—represent an attractive therapeutic strategy independent of conventional resistance pathways [3].

Antimicrobial peptides

Antimicrobial peptides encompass a structurally diverse superfamily united by cationic charge, amphipathic character, and the ability to selectively compromise microbial membrane integrity [4]. Naturally occurring exemplars include the linear α -helical magainins from *Xenopus* skin, the cyclic β -sheet defensins, the proline-rich extended peptides such as apidaecins, and the lipopeptide polymyxins that have served as a last line of defense against gram-negative sepsis for decades [5]. Structural determinants of activity include net positive charge (typically +2 to +9), which facilitates electrostatic attraction to anionic bacterial surfaces rich in phosphatidylglycerol and cardiolipin, and a distinct hydrophobic face enabling membrane insertion following surface accumulation [6]. The ratio of hydrophobic to charged residues, often quantified through the hydrophobic moment, dictates the balance

between membrane affinity and aqueous solubility, with excessive hydrophobicity correlating with loss of selectivity and increased mammalian toxicity [7]. Standard assays quantify antimicrobial potency via minimum inhibitory concentration (MIC) determination in broth microdilution format, while bactericidal kinetics are assessed through time-kill curves yielding minimum bactericidal concentration (MBC) values [8]. Despite over four decades of intensive investigation, fewer than two dozen AMPs have entered advanced clinical trials, underscoring the formidable translational barriers that persist even for well-characterized natural sequences [9].

Peptide toxicity

The therapeutic window of any antimicrobial peptide is defined by the ratio of its toxic concentration toward human cells to its effective concentration against the target pathogen, a quantity formalized as the therapeutic index [10]. Two principal toxicity endpoints dominate preclinical assessment: hemolytic activity, which measures lysis of human or animal erythrocytes quantified as the peptide concentration producing 50% hemoglobin release (HC50), and cytotoxicity, which assesses viability reduction in mammalian cell lines such as HEK293, HepG2, or primary human keratinocytes via MTT or resazurin assays yielding IC50 values [11]. The mechanistic basis of peptide toxicity parallels antimicrobial activity in many respects—both rely on membrane interactions—yet selectivity arises from fundamental differences between prokaryotic and eukaryotic membrane composition, including the presence of cholesterol in mammalian membranes, the asymmetric distribution of anionic lipids, and the lower magnitude of the transmembrane potential [12]. Computational toxicity prediction has matured substantially, with tools such as ToxinPred2 employing compositional and physicochemical descriptors trained on curated databases of experimentally characterized toxic and non-toxic peptides [13], while structure-aware deep learning models incorporating predicted secondary structure and residue-level properties have demonstrated improved discrimination [22]. Additional resources including ToxIBTL leverage information bottleneck theory and transfer learning to enhance generalization across peptide families with limited training data [12, 23].

Generative flow networks

Generative flow networks constitute a class of deep probabilistic models that learn to sequentially construct

compositional objects—graphs, sequences, sets—such that the probability of generating any given object is proportional to a pre-specified non-negative reward function [15]. Formally, a GFlowNet defines a directed acyclic graph whose vertices represent partially constructed objects (states) and whose edges correspond to incremental construction actions, with a designated initial state and terminal states corresponding to complete objects [16]. Learning proceeds by enforcing flow-matching or trajectory balance constraints: the flow into a state must equal the flow out of that state, with terminal flows equal to the reward of the corresponding completed object [17]. This formulation yields several theoretical advantages over reinforcement learning alternatives: the resulting policy samples proportionally to reward rather than concentrating probability mass at reward maxima, the learned flow network can be employed for both sampling and marginal inference, and the diversity of generated candidates emerges naturally from the proportional sampling property [24]. Extensions to multi-objective settings maintain these properties while enabling targeted exploration of specified regions of Pareto-optimal trade-off surfaces through reward decomposition and conditioning mechanisms [25].

Framework Overview

High-level architecture

The proposed framework comprises four interconnected computational modules arranged in a closed-loop generative cycle [18]. A GFlowNet-based peptide generator serves as the central engine, constructing amino acid sequences residue-by-residue through a learned autoregressive policy that maps partially assembled peptide chains to probability distributions over the twenty canonical amino acids and a termination action [19]. Complete sequences pass through two independent property prediction networks: an antimicrobial activity predictor trained to estimate minimum inhibitory concentrations against specified gram-negative target organisms, and a toxicity predictor that outputs expected hemolytic activity and mammalian cytotoxicity scores [26]. These predicted properties are combined through a scalar reward function—defined as the logarithm of the ratio between antimicrobial potency and toxicity, with tunable weighting parameters—that quantifies the desirability of each generated peptide according to the multi-objective design criteria [27]. Rewards flow backward through the GFlowNet training procedure, updating the forward policy

via trajectory balance loss to increase generation probability for high-reward sequences while preserving diversity across the reward landscape [16]. All components operate in silico without requiring feedback from wet-laboratory assays during training, though the architecture readily accommodates iterative refinement cycles when experimental data become available for predictor retraining [17].

Figure 1 presents the proposed GFlowNet-based architecture as a directional design framework linking peptide sequence generation, independent activity and toxicity prediction, composite reward construction, Pareto-aware sampling, and validation-ready candidate prioritization.




Figure 1. Generative Flow Network Framework for Multi-Objective Antimicrobial Peptide Design Against Multidrug-Resistant Gram-Negative Bacteria

Table 1 clarifies how the proposed framework separates peptide generation, property prediction, reward construction, Pareto-aware sampling, and validation screening into analytically distinct but interoperable modules.

Table 1. Functional Separation of the Generative, Predictive, Reward, and Evaluation Components in the Proposed AMP Design Framework

Framework component	Primary function	Key inputs	Key outputs
Peptide representation layer	Defines the searchable molecular design space	Amino acid alphabet, length constraints, sequence encodings,	Valid variable-length peptide states

		physicochemical descriptors	
GFlowNet generator	Constructs peptide candidates residue by residue	Partial sequences, allowable actions, termination rule	Diverse complete peptide sequences
Antimicrobial activity predictor	Estimates potency against gram-negative targets	Peptide sequence, embeddings, MIC-labeled training data	Predicted MIC or activity score
Toxicity predictor	Estimates human-cell safety endpoints	Peptide sequence, hemolysis data, cytotoxicity data	Predicted hemolysis, cytotoxicity, uncertainty score
Composite reward module	Converts multiple predicted properties into a design objective	Activity score, toxicity score, uncertainty estimate, weighting parameters	Scalar reward or conditioned reward profile
Pareto sampling layer	Maintains candidate diversity across competing objectives	Reward-weighted generated sequences	Multiple trade-off profiles
Computational validation layer	Screens candidates before synthesis	Generated candidates, held-out predictors, validity rules, safety thresholds	Prioritized shortlist for wet-lab testing

Core assumptions

Implementation of this framework presupposes the availability of large, curated databases containing peptide sequences annotated with antimicrobial activity measurements—ideally quantitative MIC values against gram-negative organisms—and toxicity endpoints including hemolysis percentages and cell viability IC50 values [24]. Public repositories such as the Database of Antimicrobial Activity and Structure of Peptides (DAMP), the Antimicrobial Peptide Database (APD3), and the Database of Antimicrobial Peptides (dbAMP) collectively contain tens of thousands of sequences, though coverage of gram-negative-specific activity and quantitative toxicity data remains sparse and heterogeneous [25]. The framework further assumes that the structure-activity relationships encoded in these training corpora are sufficiently generalizable to support prediction for de novo sequences outside the training manifold, and that the reward signal derived from predicted properties correlates meaningfully with true wet-laboratory outcomes [20]. Critically, the design anticipates that predictor models can be pre-trained independently on available data before integration into the GFlowNet loop, with architectural separation between generative and predictive components enabling asynchronous model improvement as higher-quality data accumulate [21].

Design principles

The framework is governed by four cardinal design principles that distinguish it from existing generative approaches to antimicrobial peptide discovery [1]. First, multi-objectivity is treated as a first-class design constraint: rather than collapsing antimicrobial activity and toxicity into a single optimization target post hoc, the GFlowNet formulation explicitly represents the trade-off surface and learns to cover it with diverse candidate solutions [2]. Second, diversity is encoded structurally through the proportional sampling property of GFlowNets, which prevents mode collapse by assigning non-zero probability to all objects with positive reward rather than concentrating mass at a single maximum [15]. Third, the generative process is fully learnable in an offline setting, decoupling model training from the expensive and slow cycle of peptide synthesis and experimental characterization that constrains active learning and directed evolution approaches [7]. Fourth, the modular predictor architecture supports substitution and refinement of individual property models—for instance, upgrading an activity predictor as

additional gram-negative MIC data become available—without requiring retraining of the entire generative pipeline [8]. These principles collectively position the framework as a practical computational engine for hypothesis generation, prioritizing peptide sequences for subsequent synthesis and empirical validation.

Peptide Representation

Sequence encoding

Antimicrobial peptides are represented as variable-length sequences over an alphabet of twenty canonical amino acids, with allowable lengths constrained between 5 and 50 residues to encompass the size range of known active AMPs while maintaining tractability for both generative modeling and chemical synthesis [9]. Each sequence position is encoded through one of two complementary strategies: a one-hot vector representation of dimension 20 that provides a sparse, interpretable input amenable to shallow predictors, or a learned dense embedding that maps each residue to a continuous vector capturing biochemical similarity and evolutionary substitution patterns [10]. Pretrained protein language models such as ProtBERT, which have been fine-tuned on large corpora of natural protein sequences via masked language modeling objectives, can supply contextualized residue embeddings that implicitly encode local structural propensities and sequence motifs known to correlate with antimicrobial activity [11]. Sequences shorter than the maximum allowed length are right-padded with a designated padding token, while an extinction layer or attention mask ensures that generative and predictive architectures ignore padded positions during computation [12]. An alternative representation strategy treats peptides as position-specific scoring matrices or profiles rather than discrete sequences, enabling the GFlowNet to output continuous residue preferences that guide subsequent peptide synthesis at the library rather than single-sequence level [13].

Structural features

Beyond primary sequence, several physicochemical and structural descriptors provide complementary information that enhances property prediction and may be incorporated into the GFlowNet state representation or reward computation [14]. The mean hydrophobic moment—a vector sum of residue hydrophobicities weighted by their angular positions in an idealized α -helical wheel projection—quantifies the spatial segregation of polar and non-polar

residues that correlates with membrane activity [22]. Net charge at physiological pH, calculated from the counts of basic (Arg, Lys, His) and acidic (Asp, Glu) residues with appropriate pKa corrections, determines the electrostatic driving force for initial peptide association with anionic bacterial surfaces [23]. Additional computable features include the instability index weighting dipeptide frequencies against an empirical scale of in vivo half-life, the Boman index estimating protein binding potential from amino acid side-chain solvation energies, and the isoelectric point governing solubility and aggregation propensity under physiological conditions [28]. When structurally annotated training data are available, predicted secondary structure content (α -helix, β -sheet, coil percentages) and the corresponding amphipathicity scores can be concatenated with sequence embeddings to provide the GFlowNet forward policy network with richer state representations that condition generation on desired conformational properties [28].

Property Prediction Models

Antimicrobial activity predictor

The antimicrobial activity predictor maps a complete or partial peptide sequence to a quantitative estimate of its potency against one or more target gram-negative organisms [26]. Architecturally, this module may adopt convolutional neural network (CNN) designs that scan sequence windows with learned motif filters, bidirectional long short-term memory (LSTM) networks that capture long-range residue dependencies, or transformer architectures with multi-head self-attention mechanisms that model pairwise interactions between all positions simultaneously [2]. Training data are derived from publicly available databases containing paired sequence-MIC measurements, with regression targets expressed as log-transformed MIC values ($\log \mu\text{g/mL}$) to normalize distributions spanning nanomolar to millimolar potency ranges [3]. Multi-task formulations that jointly predict activity against multiple bacterial species leverage shared sequence representations while learning organism-specific output heads, potentially improving generalization for pathogens with limited training data through transfer of knowledge from well-characterized organisms [4]. Outputs include both a point estimate of antimicrobial activity and, for binary classification applications, a probability of exceeding an activity threshold that simplifies integration with the GFlowNet reward signal [5].

Toxicity predictor

A dedicated toxicity prediction module estimates two complementary safety endpoints: hemolytic activity toward human erythrocytes and cytotoxicity toward mammalian cell lines [6]. The hemolysis predictor outputs a continuous value representing the predicted percentage of red blood cell lysis at a standardized peptide concentration or, alternatively, the concentration producing 50% hemolysis (HC50), trained on curated hemolysis datasets containing thousands of peptide measurements [18]. Parallel cytotoxicity prediction estimates mammalian cell viability—expressed as IC50 values from MTT, LDH release, or resazurin assays—using models trained on deposited data from cell lines relevant to preclinical safety assessment [8]. Both sub-modules may share a common sequence-encoding backbone with separate regression heads, reflecting the partially overlapping physicochemical determinants of membrane disruption across bacterial and mammalian membranes while preserving the capacity to learn endpoint-specific selectivity features [13]. Confidence quantification via Monte Carlo dropout or deep ensemble variance estimation accompanies each prediction, enabling the GFlowNet reward module to discount or reject sequences associated with high predictive uncertainty [14].

Confidence scores

Predictive uncertainty arising from limited training data coverage, noisy experimental measurements, and extrapolation to sequence regions distant from the training manifold can propagate through the reward signal and misdirect generative exploration [22]. To mitigate this risk, the framework incorporates uncertainty quantification through ensemble-based methods that train multiple predictor replicates with different random initializations or bootstrap-resampled training sets, computing the variance across ensemble member predictions as an epistemic uncertainty estimate [23]. Alternatively, Monte Carlo dropout applied at inference time generates a distribution of predictions whose spread reflects model uncertainty, requiring no architectural modification beyond standard dropout layers [29]. Sequences whose predicted activity or toxicity scores exhibit ensemble variance exceeding a calibrated threshold are either assigned a penalty term in the reward function or excluded from the candidate pool presented to the GFlowNet for policy updates, effectively establishing a trust region within predictor-accessible sequence space [28].

Gflownet Architecture

Action space

The GFlowNet operates over a discrete action space where peptides are built by sequentially adding one of twenty canonical amino acids or selecting a termination action, with trajectories starting from an empty sequence and proceeding until termination or reaching a maximum length of fifty residues [15, 16]. This sequential construction allows the forward policy to condition each residue choice on the entire preceding subsequence, enabling learned long-range dependencies characteristic of antimicrobial peptides [17]. The action space grows linearly with maximum length rather than combinatorially, preserving capacity to explore the 20^L sequence space through tree-structured rollouts [24].

Forward policy

The forward policy is parameterized by a transformer decoder with causal masking that processes partial sequences into contextualized representations, outputting probabilities over the twenty-one possible actions at each step [25, 26]. Policy training employs the trajectory balance objective, enforcing consistency among forward transition probabilities, a learned backward model, and a state flow function estimating achievable reward from each intermediate state [16]. Optimization proceeds via stochastic gradient descent on trajectories sampled from the current policy or an exploratory mixture with uniform action selection [15].

Reward function

The reward function formalizes multi-objective criteria as $R(x)$
 $= \log_{10} \left(\frac{IC50_{antimicrobial}}{C50_{mammalian}} \right)$, symmetrizing the therapeutic window by assigning positive rewards to peptides with favorable activity-toxicity ratios and negative rewards otherwise [22, 23, 29]. An alternative weighted-sum

$$R(x) = \alpha \times activity_{score} - \beta \times toxicity_{score}$$

formulation, introduces tunable

hyperparameters adjusting relative emphasis on potency versus safety [28]. When confidence scores indicate high epistemic uncertainty, a penalty proportional to prediction variance discourages exploitation of predictor blind spots [14].

Multi-Objective Optimization

Pareto frontier sampling

The proportional sampling property of GFlowNets enables multi-objective optimization by preferentially generating sequences near the Pareto frontier—the set of peptides where no alternative achieves simultaneously higher activity and lower toxicity—without converging to a single point [2, 17]. Instead, the learned policy populates the full trade-off curve, producing candidates from highly potent peptides with moderate toxicity to exceptionally safe peptides with acceptable activity [25]. This coverage emerges because terminal rewards are matched proportionally rather than maximized, giving sampling probability commensurate with reward magnitude across diverse activity-toxicity trade-offs [15]. The resulting ensemble provides a menu of lead candidates with distinct property profiles for selection informed by clinical context [8].

Reward shaping

Reward shaping techniques guide the GFlowNet toward specific property regions: conditioning variables specifying desired activity or toxicity thresholds enable controlled generation without retraining [19, 24]. Adaptive weight adjustment progressively shifts emphasis between objectives during training, initially encouraging broad exploration before tightening focus toward clinically viable trade-offs [18]. Potential-based rewards introduce state-dependent bonuses that accelerate credit assignment across long trajectories, offsetting sparse terminal rewards while preserving the theoretical guarantee of proportional sampling when potentials satisfy consistency conditions [16, 17].

Evaluation Strateg

Evaluation strateg

Table 2 specifies how generated peptides should be evaluated not only by predicted antimicrobial potency but also by novelty, diversity, human-cell safety, uncertainty, Pareto-frontier coverage, and experimental readiness.

Table 2. Multi-Objective Evaluation Logic for Prioritizing Generated Antimicrobial Peptide Candidates

Evaluation dimension	What it assesses	Suggested metric or criterion	Why it matters for this framework
Sequence validity	Whether generated peptides obey design constraints	Standard amino acids only; length 5–50 residues; no invalid tokens	Ensures GFlowNet outputs are chemically interpretable and synthetically plausible
Novelty	Whether outputs differ from training peptides	Sequence identity threshold; absence from AMP training databases	Prevents simple memorization of known AMPs
Diversity	Breadth of generated sequence space	Unique sequence count; pairwise sequence distance; motif diversity	Tests whether GFlowNet avoids mode collapse
Antimicrobial potency	Predicted activity against MDR gram-negative pathogens	Predicted MIC, log-MIC, or probability of activity	Captures the therapeutic objective of bacterial inhibition
Human-cell safety	Predicted toxicity toward mammalian cells	Hemolysis percentage, HC50, cytotoxicity IC50	Addresses the main translational barrier for AMPs
Therapeutic window	Balance between potency and toxicity	Activity–toxicity ratio; weighted reward; safety-adjusted potency score	Converts competing objectives into a clinically meaningful trade-off
Predictive confidence	Reliability of model-	Ensemble variance; MC dropout	Reduces reward hacking and

	estimated properties	uncertainty; distance from training distribution	extrapolation-driven false positives
Pareto-frontier coverage	Breadth of activity–safety trade-off profiles	Hypervolume indicator; number of nondominated candidates	Demonstrates multi-objective advantage over single-objective generators
Experimental readiness	Suitability for synthesis and laboratory testing	Solubility, aggregation risk, protease susceptibility, human-protein similarity	Bridges computational generation and wet-lab feasibility

Distributional metrics

Distributional metrics assess global characteristics including the number of unique peptides across a fixed sampling budget and the fraction of generated sequences absent from training data, measuring effective diversity and innovation capacity [26, 27]. Sequence validity verifies outputs use only standard amino acids within specified length constraints [20]. Benchmarking against variational autoencoders, generative adversarial networks, and genetic algorithms using identical metrics over matched sample sizes enables rigorous comparison of diversity characteristics [4, 5, 28].

Reward metrics

Reward-centric evaluation includes mean and maximum reward over multiple sampling batches to assess generation quality and capacity for discovering favorable trade-offs [1, 21]. The hypervolume indicator quantifies Pareto frontier coverage by measuring the volume of property space dominated by generated candidates relative to a reference point [25]. Additional metrics count candidates exceeding activity thresholds while remaining below toxicity cutoffs, identifying sequences meeting minimal criteria for experimental progression [2].

Computational validation

Computational validation uses held-out predictor models to estimate MIC values against target gram-negative organisms and hemolysis percentages at clinically relevant concentrations for each candidate [14, 26]. Sequences passing acceptability thresholds—MIC below 16 µg/mL and hemolysis below 10% at 128 µg/mL—are flagged for synthesis [3]. Additional checks assess aggregation propensity, protease susceptibility, and similarity to human proteins that could cause off-target effects, with all assessments explicitly qualified as predictions requiring experimental confirmation [9, 10, 22].

Limitations

Technical limitations

The framework's reliance on supervised predictors introduces vulnerability to error propagation, where inaccurate activity or toxicity estimates misdirect generative exploration toward sequences performing poorly experimentally [11, 12]. Peptide synthesizability remains unmodeled: generated sequences may exhibit poor solubility, aggregation, or cysteine patterns complicating synthesis [13]. The discrete action space precludes non-canonical amino acids, post-translational modifications, and D-amino acid substitutions known to enhance stability and activity in clinically relevant AMPs [14].

Clinical limitations

In vitro activity under standardized conditions may not translate to complex biological environments where serum proteins bind cationic peptides, divalent cations antagonize membrane interactions, and biofilms reduce susceptibility [22, 23]. Toxicity assessment limited to hemolysis and cell viability misses organ-level effects—nephrotoxicity, hepatotoxicity, immunogenicity—that have terminated clinical development of multiple AMP candidates [26]. Rapid renal clearance, serum proteolysis, and negligible oral bioavailability necessitate formulation strategies that may alter the sequence-activity relationships upon which the generative model was trained [27].

Conclusion

This article has delineated a conceptual framework applying generative flow networks to multi-objective antimicrobial peptide design targeting multidrug-resistant gram-negative bacteria. Peptide generation is formulated

as a sequential decision process with a learned stochastic policy trained via trajectory balance, producing diverse candidates whose antimicrobial potency and predicted safety jointly satisfy design criteria. The modular architecture separates generative, predictive, and evaluative components for independent refinement as data quality improves.

The GFlowNet paradigm offers advantages over existing generative approaches: proportional sampling maintains diversity and avoids mode collapse, the Pareto frontier is comprehensively explored, and the policy can be conditioned on desired property specifications at inference time without retraining. Probabilistic framing provides interpretable likelihoods supporting uncertainty-aware candidate prioritization.

Important limitations temper enthusiasm: output quality depends on predictor accuracy, itself constrained by training data skewed toward gram-positive pathogens and lacking systematic toxicity annotation. High-confidence predictions require wet-laboratory validation, and the gap between in vitro activity and in vivo efficacy remains the rate-limiting step in AMP development. This framework is a hypothesis-generation engine producing testable predictions, not validated therapeutics.

Realization requires implementation on curated databases with matched activity and toxicity annotations spanning WHO critical priority gram-negative organisms.

Computational development should proceed alongside wet-laboratory partnerships for peptide synthesis, purification, and susceptibility testing. The iterative cycle of generation, validation, and predictor refinement with newly acquired experimental data offers a principled pathway toward accelerated discovery of peptide therapeutics for infections increasingly defying conventional antibiotics.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 18 Feb 2025 Revised: 10 May 2025 Accepted: 25 Jun 2025
Published online: 20 January 2026

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Capecchi A, Cai X, Personne H, Köhler T, van Delden C, Reymond JL, et al. Machine learning designs non-hemolytic antimicrobial peptides. *Chem Sci*. 2021;12(26):9221-32.
- Nagarajan D, Nagarajan T, Roy N, Kulkarni O, Ravichandran S, Mishra M, et al. Computational antimicrobial peptide design and evaluation against multidrug-resistant clinical isolates of bacteria. *J Biol Chem*. 2018;293(10):3492-509.
- Vishnepolsky B, Gabrielian A, Rosenthal A, Hurt DE, Tartakovsky M, Managadze G, et al. Predictive model of linear antimicrobial peptides active against gram-negative bacteria. *J Chem Inf Model*. 2018;58(5):1141-51.
- Tucs A, Tran DP, Yumoto A, Ito Y, Uzawa T, Tsuda K. Generating ampicillin-level antimicrobial peptides with activity-aware generative adversarial networks. *ACS Omega*. 2020;5(36):22847-51.

Van Oort CM, Ferrell JB, Remington JM, Wshah S, Li J. AMPGAN v2: machine learning-guided design of antimicrobial peptides. *J Chem Inf Model*. 2021;61(5):2198-207.

Dean SN, Walper SA. Variational autoencoder for generation of antimicrobial peptides. *ACS Omega*. 2020;5(33):20746-54.

Veltri D, Kamath U, Shehu A. Deep learning improves antimicrobial peptide recognition. *Bioinformatics*. 2018;34(16):2740-7.

Li C, Sutherland D, Hammond SA, Yang C, Taho F, Bergman L, et al. AMPlify: attentive deep learning model for discovery of novel antimicrobial peptides effective against WHO priority pathogens. *BMC Genomics*. 2022;23(1):77.

Li C, Warren RL, Birol I. Models and data of AMPlify: a deep learning tool for antimicrobial peptide prediction. *BMC Res Notes*. 2023;16(1):11.

Yan J, Bhadra P, Li A, Sethiya P, Qin L, Tai HK, et al. Deep-AmPEP30: improve short antimicrobial peptides prediction with deep learning. *Mol Ther Nucleic Acids*. 2020;20:882-94.

Meher PK, Sahu TK, Saini V, Rao AR. Predicting antimicrobial peptides with improved accuracy by incorporating the compositional, physico-chemical and structural features into Chou's general PseAAC. *Sci Rep*. 2017;7(1):42362.

Wei L, Ye X, Sakurai T, Mu Z, Wei L. ToxiBTL: prediction of peptide toxicity based on information bottleneck and transfer learning. *Bioinformatics*. 2022;38(6):1514-24.

Sharma N, Naorem LD, Jain S, Raghava GP. ToxinPred2: an improved method for predicting toxicity of proteins. *Brief Bioinform*. 2022;23(5):bbac174.

Rathore AS, Kumar N, Choudhury S, Mehta NK, Raghava GP. Prediction of hemolytic peptides and their hemolytic concentration. *Commun Biol*. 2025;8(1):176.

Bengio E, Jain M, Korablyov M, Precup D, Bengio Y. Flow network based generative models for non-iterative diverse candidate generation. *Adv Neural Inf Process Syst*. 2021;34:27381-94.

Jain M, Bengio E, Hernandez-Garcia A, Rector-Brooks J, Dossou BF, Ekbote CA, et al. Biological sequence design with gflownets. In: *International Conference on Machine Learning*. PMLR; 2022. p. 9786-801.

Jain M, Raparthy SC, Hernández-García A, Rector-Brooks J, Bengio Y, Miret S, et al. Multi-objective gflownets. In: *International Conference on Machine Learning*. PMLR; 2023. p. 14631-53.

Shen T, Seo S, Lee G, Pandey M, Smith JR, Cherkasov A, et al. TacoGFN: target-conditioned GFlowNet for structure-based drug design. *arXiv preprint arXiv:2310.03223*. 2023.

Koziarski M, Reakesh A, Shevchuk D, van der Sloot A, Gaiński P, Bengio Y, et al. Rgfn: Synthesizable molecular generation using gflownets. *Adv Neural Inf Process Syst*. 2024;37:46908-55.

Wang Y, Zhao L, Li Z, Xi Y, Pan Y, Zhao G, et al. A generative artificial intelligence approach for the discovery of antimicrobial peptides against multidrug-resistant bacteria. *Nat Microbiol*. 2025;10(11):2997-3012.

Jin S, Zeng Z, Xiong X, Huang B, Tang L, Wang H, et al. AMPGen: an evolutionary information-reserved and diffusion-driven generative model for de novo design of antimicrobial peptides. *Commun Biol*. 2025;8(1):839.

Ebrahimikondori H, Sutherland D, Yanai A, Richter A, Salehi A, Li C, et al. Structure-aware deep learning model for peptide toxicity prediction. *Protein Sci*. 2024;33(7):e5076.

Guan J, Xie P, Meng D, Yao L, Yu D, Chiang YC, et al. ToxiPep: peptide toxicity prediction via fusion of context-aware representation and atomic-level graph. *Comput Struct Biotechnol J*. 2025;27:2347-58.

Das P, Sercu T, Wadhawan K, Padhi I, Gehrmann S, Cipicigan F, et al. Accelerating antimicrobial discovery with controllable deep generative models and molecular dynamics. *arXiv preprint arXiv:2005.11248*. 2020.

Szymczak P, Możejko M, Grzegorzec T, Jurczak R, Bauer M, Neubauer D, et al. Discovering highly potent antimicrobial peptides with deep generative model HydrAMP. *Nat Commun*. 2023;14(1):1453.

Pandi A, Adam D, Zare A, Trinh VT, Schaefer SL, Burt M, et al. Cell-free biosynthesis combined with deep learning accelerates de novo-development of antimicrobial peptides. *Nat Commun*. 2023;14(1):7197.

Wang B, Lin P, Zhong Y, Tan X, Shen Y, Huang Y, et al. Explainable deep learning and virtual evolution identifies antimicrobial peptides with activity against multidrug-resistant human pathogens. *Nat Microbiol*. 2025;10(2):332-47.

Wang C, Garlick S, Zloh M. Deep learning for novel antimicrobial peptide design. *Biomolecules*. 2021;11(3):471.

Yan K, Chen S, Liu B, Wu H. Accurate prediction of toxicity peptide and its function using multi-view tensor learning and latent semantic learning framework. *Bioinformatics*. 2025;41(9):btaf489.