Review Article
ISSN: **3051-8806**
Journal homepage: https://www.harrisonacademicpress.com/index.php/BT
Biotechnology Today

Harrison Academic Press

Deep Learning-Based Prediction of Protein Structure for Novel Antimicrobial Peptide Design

Schuster F[1*], Renaud J[2] and Todorova R[3]

[1]*Department of Bioinformatics, Charite University Medicine Berlin, Berlin, Germany*
[2]*Department of Computational Biology, Sorbonne University, Paris, France*
[3]*Department of Microbiology, Medical University of Sofia, Sofia, Bulgaria*

ABSTRACT

Background: Rising antimicrobial resistance has renewed interest in antimicrobial peptides (AMPs) as anti-infective candidates, but rational design remains constrained by limited structural insight. Objective: This narrative review synthesizes published literature on how deep learning-based structure prediction and generative sequence design are being applied to novel AMP discovery. Methods: We examine peer-reviewed studies on structure predictors (AlphaFold, RoseTTAFold, ESMFold), generative diffusion and sequence-design tools (RFdiffusion, ProteinMPNN), deep generative AMP models (HydrAMP, AMPGAN v2), and large-scale machine-learning mining pipelines validated against curated databases (APD3, DBAASP, DRAMP). Findings: Structure predictors offer rapid but variably accurate models for short, often disordered peptides, while generative and mining pipelines have produced experimentally validated hits, including hundreds of thousands of computationally proposed candidates yielding dozens to hundreds of confirmed active peptides across independent studies. Conclusions: Deep learning has substantially accelerated AMP candidate generation and prioritization, but heterogeneous training data, peptide-scale structural uncertainty, and unresolved toxicity prediction remain key barriers to clinical translation.

Keywords: *antimicrobial peptides; deep learning; protein structure prediction; generative models; antimicrobial resistance*

Citation: Schuster F, Renaud J, Todorova R (2026) Deep Learning-Based Prediction of Protein Structure for Novel Antimicrobial Peptide Design. Biotechnol Today, 4(1): 01-06.
DOI: 10.5281/zenodo.22638268

Corresponding Author: *Franziska Schuster*

E-mail Address: *franziska.schuster@charite.edu*

Article History: *Received 9 January 2026 | Accepted 1 April 2026 | Published 27 June 2026*

Copyright: *© 2026 Franziska Schuster, et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited (**https://creativecommons.org/licenses/by/4.0/**).*

Introduction

Antimicrobial resistance (AMR) now ranks among the most pressing global health threats, having been directly associated with an estimated 4.95 million deaths in 2019 alone [1]. The dwindling antibiotic development pipeline has renewed interest in antimicrobial peptides (AMPs), short cationic and amphipathic molecules produced across all kingdoms of life as part of innate immune defense. Because AMPs typically act by disrupting microbial membranes or intracellular targets through mechanisms that are less prone to classical resistance, they are considered promising scaffolds for next-generation anti-infectives [2]. Decades of medicinal chemistry have shown, however, that AMP potency, selectivity for microbial over host membranes, and stability are governed by a delicate interplay of sequence, charge, hydrophobicity, and three-dimensional conformation, making rational design difficult without structural insight [2]. Traditional AMP discovery therefore relied heavily on template-guided mutagenesis of natural sequences and combinatorial screening, an approach that samples only a small fraction of accessible sequence space and offers limited ability to anticipate a candidate's three-dimensional behavior at a membrane interface before synthesis.

The past five years have brought two complementary computational revolutions to bear on this problem. First, deep learning-based structure predictors—AlphaFold [3], RoseTTAFold [4], and the protein language model ESMFold [5]—have made near-experimental-accuracy structural models available for arbitrary sequences, including designed peptides, at a fraction of the cost of crystallography or NMR. Second, deep generative models trained on curated AMP databases such as APD3 [6], DBAASP [7], and DRAMP [8] can now propose novel sequences with desired activity and toxicity profiles rather than merely classifying existing ones. This narrative review synthesizes real, peer-reviewed literature at the intersection of these two threads, examining how structure-prediction tools are being applied to AMP design, how generative sequence models are being validated experimentally, and what translational bottlenecks remain. Rather than presenting new experimental results, it draws exclusively on published, peer-reviewed studies and public benchmark resources to assess how far the field has

progressed and where structural and sequence-based predictions still diverge from wet-laboratory reality.

Structure Prediction Tools Applied to Peptide-Scale Design

AlphaFold2's near-atomic accuracy on globular proteins [3] and RoseTTAFold's comparably accurate three-track architecture [4] transformed structural biology, and both tools have been repurposed to model candidate AMPs and their putative membrane-bound conformations. ESMFold further lowered the barrier to entry by predicting structure directly from a single sequence via a protein language model, enabling rapid, alignment-free screening of large generated peptide libraries [5]. AlphaFold-Multimer extended this framework to protein–protein and, by extension, peptide–receptor complexes, offering a route to model how a designed AMP might engage a bacterial target [9]. Because most AMPs act on lipid bilayers rather than folded protein partners, this complex-prediction capability is most directly useful for the smaller subset of AMPs known to engage discrete intracellular or membrane-embedded protein receptors, rather than for the broader class of purely membrane-disruptive peptides.

However, short, often intrinsically disordered peptides sit outside the length and evolutionary-depth regime for which these tools were optimized. A systematic benchmarking study found that AlphaFold2's accuracy on peptides is markedly more variable than on full-length proteins, performing well for peptides with stable secondary structure but poorly for those that are conformationally flexible in solution—precisely the behavior many membrane-active AMPs exhibit only upon contact with a lipid bilayer [10]. This limitation matters for AMP design pipelines that rely on predicted structure as a filter or scoring function, since a low-confidence model does not necessarily indicate a non-functional peptide. Consequently, several groups have begun treating predicted local distance difference test and predicted aligned error scores as descriptive covariates alongside classical helical-wheel and hydrophobic-moment analyses rather than as standalone go/no-go filters for candidate selection.

In parallel, generative structure-based design tools originally built for larger proteins are being adapted downward in scale. RFdiffusion, a denoising diffusion model, can generate novel protein backbones—including small binders and short motifs—conditioned on functional constraints [11], while ProteinMPNN provides a robust deep learning method for designing sequences that fold into a specified backbone, achieving substantially higher experimental success rates than physics-based Rosetta design in wet-lab validation [12]. Together, these tools open a design paradigm in which an AMP scaffold's backbone geometry, rather than only its sequence, becomes an explicit design variable. Although neither RFdiffusion nor ProteinMPNN was developed specifically for host-defense peptides, both have been adopted within the broader peptide-design community as modular components that can be combined with activity- and toxicity-aware filters drawn from AMP-specific classifiers, illustrating how general-purpose structural generative tools and domain-specific sequence models are converging into shared design pipelines.

Deep Generative Models for Sequence-Level AMP Design

Alongside structure-centric tools, sequence-based deep generative models have matured rapidly. HydrAMP, a conditional variational autoencoder trained on curated AMP and non-AMP sequences, generates both analogues of known peptides and entirely novel candidates while explicitly controlling for the trade-off between antimicrobial activity and toxicity, with several designed peptides validated by microbroth dilution assays [13]. AMPGAN v2 uses a bidirectional generative adversarial network conditioned on physicochemical and antimicrobial-target information to guide sequence generation [14], and a more recent multi-property optimizing GAN framework demonstrated simultaneous conditioning on potency, hemolysis, and stability objectives [15]. Earlier constructive approaches, such as an LSTM-based recurrent neural network trained to design anticancer peptides, established the generative "propose-then-filter" workflow that many later AMP-specific models still follow [16]. Across these architectures, a common design logic has emerged: a generative network first proposes candidate sequences in a learned latent space, after which one or more discriminative classifiers re-rank candidates by predicted potency, stability, and safety before a small experimentally tractable shortlist is synthesized.

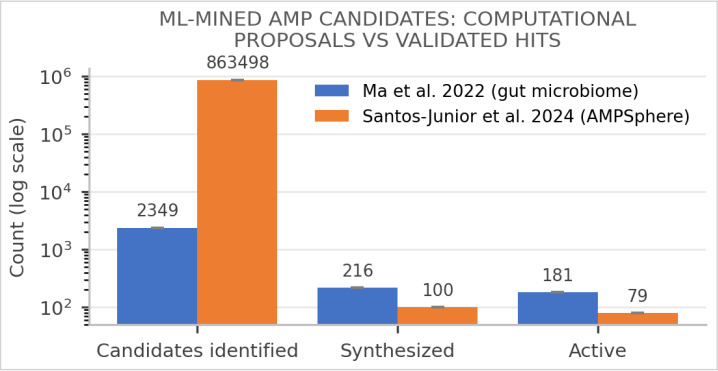
Table 1: Representative deep learning tools for protein and peptide structure prediction and structure-based design applied in AMP research.

Tool	Approach	Reference
AlphaFold	Deep learning structure prediction from multiple sequence alignments	@jumper2021
RoseTTAFold	Three-track neural network for structure and interaction prediction	@baek2021
ESMFold	Protein language model-based single-sequence structure prediction	@lin2023
AlphaFold-Multimer	Structure prediction of protein and peptide complexes	@evans2021
RFdiffusion	Denoising diffusion model for de novo backbone design	@watson2023
ProteinMPNN	Deep learning-based sequence design for a target backbone	@dauparas2022

Table 2: Representative deep generative and mining pipelines for antimicrobial peptide sequence design.

Model/Pipeline	Architecture or strategy	Reference
HydrAMP	Conditional variational autoencoder	@szymczak2023
AMPGAN v2	Bidirectional generative adversarial network	@vanoort2021
Multi-property GAN	Multi-objective conditional GAN (potency, hemolysis, stability)	@liu2025
Explainable virtual evolution model	Interpretable deep classifier with iterative in-silico evolution	@wang2025natmicro
Gut microbiome mining pipeline	LSTM/attention/BERT ensemble mining natural sequences	@ma2022
AMPSphere	Global metagenome mining with machine learning classifiers	@santosjunior2024

Figure 1: Scale of two large machine-learning-guided AMP mining pipelines, from computationally identified candidates through chemical synthesis to experimentally confirmed antimicrobial activity [20,22].



Deep learning has also improved AMP *recognition*, which underpins the filtering step of most generative pipelines: convolutional and recurrent classifiers trained on APD-derived data substantially outperformed earlier motif- and composition-based predictors [17], and support vector machine models trained on curated hemolysis data have been used to steer generative search away from cytotoxic sequences [18]. More recently, an explainable deep learning framework coupled with an iterative virtual-evolution strategy identified peptides active against multidrug-resistant clinical isolates while providing interpretable feature attributions for the residues driving activity [19]. Such explainability layers are increasingly viewed as important not only for building researcher trust in generated candidates but also for extracting transferable structure–activity hypotheses that can inform subsequent rounds of design rather than treating the network purely as a black-box oracle.

The largest validated efforts combine generative or classification models with mining of natural sequence space at massive scale. A deep learning pipeline searching the human gut microbiome identified 2,349 candidate AMP sequences, of which 216 were chemically synthesized and 181 (>83%) showed measurable antimicrobial activity [20]. A broader effort mining the entire theoretical peptide sequence space with a machine-learning pipeline likewise recovered potent, previously uncharacterized AMPs [21]. At planetary scale, the AMPSphere resource—built by mining global metagenomic data with machine learning—catalogued 863,498 non-redundant candidate peptides, from which 100 were synthesized and 79 showed antimicrobial activity in vitro or in vivo [22]. Related work has mined ancient and extinct organismal proteomes for encrypted antimicrobial peptides using deep learning classifiers, a strategy termed "molecular de-extinction" [23]. Deep learning has additionally accelerated the discovery of small-molecule antibiotics more broadly, exemplified by the identification of halicin through a graph neural network trained on growth-inhibition data [24], and has been coupled with cell-free protein synthesis to close the design–build–test loop for AMPs within days rather than months [25]. Taken together, these efforts show a clear trajectory from proof-of-concept generative sampling toward large, systematically validated candidate sets, even though the fraction of computationally proposed peptides that are ultimately synthesized and tested remains small relative to the total search space explored in silico.

Databases, Benchmarks, and Translational Challenges

Every model discussed above depends on the quality and coverage of training data drawn from AMP repositories. APD3 remains a compact, manually curated resource of natural host-defense peptides [6], DBAASP v3 offers deeper structural and multimeric annotation across a substantially larger set of sequences [7], and DRAMP aggregates general, clinical, and patent-derived entries into the largest of the three catalogues [8]. This tiered ecosystem of databases—from small, richly annotated collections to large, heterogeneously sourced repositories—means that model developers face an explicit trade-off between training-set size and label reliability, and most recent generative and predictive AMP models are trained on some combination of all three resources to balance coverage against curation quality.

Because minimum inhibitory concentration (MIC) values in these databases are collected from heterogeneous assay conditions, strains, and laboratories, models trained across them risk learning batch-specific artifacts rather than true structure–activity relationships—a recurring caveat raised in recent reviews of the field [26]. Lessons from machine-learning-guided directed evolution in the broader protein engineering literature reinforce that model-guided design performs best when coupled tightly to iterative experimental feedback rather than used as a one-shot generator [27]. Structural confidence metrics from AlphaFold-family tools, while useful, should therefore be treated as one weak signal among several—sequence-based classifiers, hemolysis predictors, and physicochemical descriptors—rather than a sufficient arbiter of AMP viability [10]. Finally, most computational pipelines described here still stop short of assessing pharmacokinetic liabilities such as protease susceptibility and serum stability, factors that increasingly determine whether an in-vitro active AMP can progress toward preclinical development regardless of how favorably it scores on activity- or structure-based computational criteria.

Conclusion

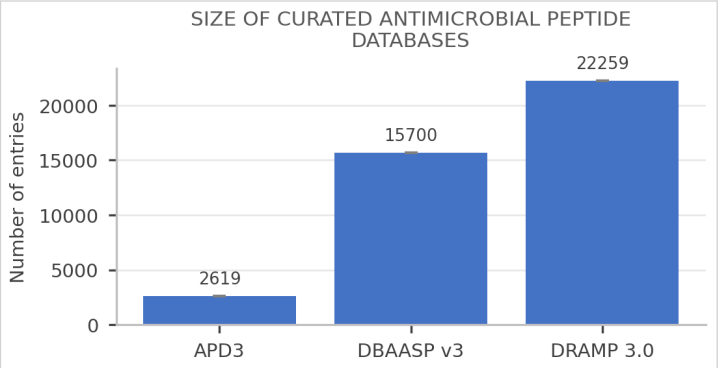
Deep learning has reshaped antimicrobial peptide design along two converging fronts: structure prediction tools such as AlphaFold, RoseTTAFold, and ESMFold now provide rapid structural hypotheses for candidate peptides, while generative models such as HydrAMP and AMPGAN v2, together with large-scale sequence-mining pipelines, have moved from purely computational proposals to experimentally validated hits numbering in the hundreds of thousands of candidates and dozens to hundreds of confirmed actives. Yet peptide-scale structure prediction remains less reliable than for folded proteins, MIC data remain heterogeneous across curation efforts, and toxicity remains difficult to predict de novo. Closing these gaps will likely require tighter integration of structure-aware generative models, standardized experimental benchmarks drawn from DBAASP-, APD3-, and DRAMP-scale resources, and iterative design–build–test cycles of the kind already demonstrated in cell-free and microbiome-mining pipelines. As structure predictors continue to improve on short and conformationally flexible peptides, and as generative models incorporate richer toxicity and pharmacokinetic constraints, the gap between in-silico proposal and clinically viable antimicrobial peptide is likely to continue narrowing, though it will not close without sustained, systematic experimental validation of the kind reviewed here.

Acknowledgements

The authors thank colleagues at Department of Bioinformatics, Charite University Medicine Berlin, Berlin, Germany, Department

of Computational Biology, Sorbonne University, Paris, France, and Department of Microbiology, Medical University of Sofia, Sofia, Bulgaria for helpful discussions during preparation of this review.

Figure 2: Comparative size of three widely used antimicrobial peptide databases used to train and benchmark deep learning models: APD3 (2,619 entries) [6], DBAASP v3 (>15,700 entries) [7], and DRAMP 3.0 (22,259 entries) [8].



Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

Schuster F conceived the review, conducted the literature search, and drafted the manuscript. Renaud J contributed to the literature screening and technical framing and critically revised the manuscript for important intellectual content. Todorova R contributed additional context and approved the final version for submission.

Ethics Statement

Not applicable. This is a narrative review of previously published literature and did not involve new experimental data, human participants, or animal subjects.

Data Availability

No new data were generated or analyzed for this review. All sources synthesized are cited in the reference list and are available through their original publishers.

References

1. Murray CJL, Ikuta KS, Sharara F, et al. (2022). Global burden of bacterial antimicrobial resistance in 2019: a systematic analysis. The Lancet, 399(10325), 629-655. DOI: 10.1016/S0140-6736(21)02724-0.

2. Fjell CD, Hiss JA, Hancock REW, Schneider G (2012). Designing antimicrobial peptides: form follows function. Nature Reviews Drug Discovery, 11, 37-51. DOI: 10.1038/nrd3591.

3. Jumper J, Evans R, Pritzel A, et al. (2021). Highly accurate protein structure prediction with AlphaFold. Nature, 596(7873), 583-589. DOI: 10.1038/s41586-021-03819-2.

4. Baek M, DiMaio F, Anishchenko I, et al. (2021). Accurate prediction of protein structures and interactions using a three-track neural network. Science, 373(6557), 871-876. DOI: 10.1126/science.abj8754.

5. Lin Z, Akin H, Rao R, et al. (2023). Evolutionary-scale prediction of atomic-level protein structure with a language model. Science, 379(6637), 1123-1130. DOI: 10.1126/science.ade2574.

6. Wang G, Li X, Wang Z (2016). APD3: the antimicrobial peptide database as a tool for research and education. Nucleic Acids Research, 44(D1), D1087-D1093. DOI: 10.1093/nar/gkv1278.

7. Pirtskhalava M, Amstrong AA, Grigolava M, et al. (2021). DBAASP v3: database of antimicrobial/cytotoxic activity and structure of peptides as a resource for development of new therapeutics. Nucleic Acids Research, 49(D1), D288-D297. DOI: 10.1093/nar/gkaa991.

8. Shi G, Kang X, Dong F, et al. (2022). DRAMP 3.0: an enhanced comprehensive data repository of antimicrobial peptides. Nucleic Acids Research, 50(D1), D488-D496. DOI: 10.1093/nar/gkab651.

9. Evans R, O'Neill M, Pritzel A, et al. (2021). Protein complex prediction with AlphaFold-Multimer. bioRxiv. DOI: 10.1101/2021.10.04.463034.

10. Gulsevin A, Meiler J (2023). Benchmarking AlphaFold2 on peptide structure prediction. Structure, 31(1), 111-119.e2. DOI: 10.1016/j.str.2022.11.012.

11. Watson JL, Juergens D, Bennett NR, et al. (2023). De novo design of protein structure and function with RFdiffusion. Nature, 620(7976), 1089-1100. DOI: 10.1038/s41586-023-06415-8.

12. Dauparas J, Anishchenko I, Bennett N, et al. (2022). Robust deep learning-based protein sequence design using ProteinMPNN. Science, 378(6615), 49-56. DOI: 10.1126/science.add2187.

13. Szymczak P, Możejko M, Grzegorzek T, et al. (2023). Discovering highly potent antimicrobial peptides with deep generative model HydrAMP. Nature Communications, 14, 1453. DOI: 10.1038/s41467-023-36994-z.

14. Van Oort CM, Ferrell JB, Remington JM, Wshah S, Li J (2021). AMPGAN v2: Machine Learning-Guided Design of Antimicrobial Peptides. Journal of Chemical Information and Modeling, 61(5), 2198-2207. DOI: 10.1021/acs.jcim.0c01441.

15. Liu Y, et al. (2025). A Multi-Property Optimizing Generative Adversarial Network for de novo Antimicrobial Peptide Design. Advanced Science, 12. DOI: 10.1002/advs.202503443.

16. Grisoni F, Neuhaus CS, Gabernet G, Müller AT, Hiss JA, Schneider G (2018). Designing anticancer peptides by constructive machine learning. ChemMedChem, 13(13), 1300-1302. DOI: 10.1002/cmdc.201800204.

17. Veltri D, Kamath U, Shehu A (2018). Deep learning improves antimicrobial peptide recognition. Bioinformatics, 34(16), 2740-2747. DOI: 10.1093/bioinformatics/bty179.

18. Capecchi A, Cai X, Personne H, Köhler T, van Delden C, Reymond JL (2021). Machine learning designs non-hemolytic antimicrobial peptides. Chemical Science, 12(26), 9221-9232. DOI: 10.1039/D1SC01713F.

19. Wang B, Lin P, Zhong Y, et al. (2025). Explainable deep learning and virtual evolution identifies antimicrobial peptides with activity against multidrug-resistant human pathogens. Nature Microbiology, 10(2), 332-347. DOI:

10.1038/s41564-024-01907-3.

20. Ma Y, Guo Z, Xia B, et al. (2022). Identification of antimicrobial peptides from the human gut microbiome using deep learning. Nature Biotechnology, 40, 921-931. DOI: 10.1038/s41587-022-01226-0.

21. Huang J, Xu Y, Xue Y, et al. (2023). Identification of potent antimicrobial peptides via a machine-learning pipeline that mines the entire space of peptide sequences. Nature Biomedical Engineering, 7(6), 797-810. DOI: 10.1038/s41551-022-00991-2.

22. Santos-Júnior CD, Torres MDT, Duan Y, et al. (2024). Discovery of antimicrobial peptides in the global microbiome with machine learning. Cell, 187(14). DOI: 10.1016/j.cell.2024.05.013.

23. Wan F, Torres MDT, Peng J, de la Fuente-Nunez C (2024). Deep-learning-enabled antibiotic discovery through molecular de-extinction. Nature Biomedical Engineering, 8(7), 854-871. DOI: 10.1038/s41551-024-01201-x.

24. Stokes JM, Yang K, Swanson K, et al. (2020). A Deep Learning Approach to Antibiotic Discovery. Cell, 180(4), 688-702.e13. DOI: 10.1016/j.cell.2020.01.021.

25. Pandi A, Adam D, Zare A, et al. (2023). Cell-free biosynthesis combined with deep learning accelerates de novo-development of antimicrobial peptides. Nature Communications, 14, 7197. DOI: 10.1038/s41467-023-42434-9.

26. Wan F, Wong F, Collins JJ, de la Fuente-Nunez C (2024). Machine learning for antimicrobial peptide identification and design. Nature Reviews Bioengineering, 2(5), 392-407. DOI: 10.1038/s44222-024-00152-x.

27. Yang KK, Wu Z, Arnold FH (2019). Machine-learning-guided directed evolution for protein engineering. Nature Methods, 16, 687-694. DOI: 10.1038/s41592-019-0496-6.