bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

HydrAMP: a deep generative model for antimicrobial peptide discovery

Paulina Szymczak[1,†], Marcin Możejko[1,†·], Tomasz Grzegorzek[1], Marta Bauer[2], Damian Neubauer[2], Michał Michalski[3], Jacek Sroka[1], Piotr Setny[3], Wojciech Kamysz[2], and Ewa Szczurek[1,*]

[1]*Faculty of Mathematics, Informatics and Mechanics, University of Warsaw*

[2]*Medical University of Gdańsk*

[3]*The Centre of New Technologies, University of Warsaw*

[†]*These authors contributed equally to this work*

[*]*Correspondence should be adressed to szczurek@mimuw.edu.pl*

Abstract

Antimicrobial peptides (AMP) emerge as compounds that can alleviate the global health hazard of antimicrobial resistance. Since the repertoire of experimentally verified AMPs is limited, there is a need for novel computational approaches to peptide generation. For such approaches, exploring the amino-acid peptide representation space is infeasible due to its sparsity and combinatorial complexity. Thus, we propose HydrAMP, a conditional variational autoencoder that learns a lower-dimensional and continuous space of peptides' representations and captures their antimicrobial properties. HydrAMP outperforms other approaches in generating peptides, either *de novo*, or by analogue discovery, and leverages parameter-controlled creativity. The model disentangles the latent representation of a peptide from its antimicrobial conditions, allowing for targeted generation. Wet-lab experiments and molecular dynamics simulation confirm the increased activity of a Pexiganan-based analogue produced by HydrAMP. HydrAMP proposes new promising AMP candidates, enabling progress towards a new generation of antibiotics.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

1 Introduction

Microbes pose a continuously increasing threat to human health, in particular by causing sepsis, post-surgical infections, and putting at risk patients with chronic conditions or immunodeficiency [1]. It is estimated that microbial infection will become the main cause of death by 2050, exceeding the currently dominating cancer and cardiovascular diseases [2]. The reason for the growing danger of microbes is their ability to gain resistance to antibiotics [1]. In the recent years, antimicrobial peptides (AMPs) are investigated as attractive alternatives to conventional antibiotic treatment. The acquisition of resistance to AMPs in microbes is much slower [3], moreover, they can be active against pathogens that are resistant to antibiotics [4].

Typically, AMPs are amphiphatic; cationic amino acids build the hydrophylic face of the peptide, while hydrophobic residues dominate the opposite side of the molecule. Amphiphacity together with high charge allow AMPs to invade and disrupt the negatively charged microbial cellular membrane [5]. Antimicrobial activity of peptides is measured experimentally by determining its Minimal Inhibitory Concentration (MIC). The most prominent peptides have low values of MIC, meaning that they remain active even in low concentrations, but their prevalence is limited. Given the high therapeutic promise of AMPs, it is critical to design novel peptides that are nonexistent in nature and could be synthesised and used to treat microbial infections in the clinic.

In biological labs, the process of identifying new antimicrobial peptides commonly proceeds by taking existing, known AMPs as prototypes, and adding or substituting amino acids, aiming at increasing the resulting amphiphacity and/or charge. Such generated peptides are subjected to synthesis attempt, and if synthesizable, their antimicrobial activity is experimentally verified. First, this process is tedious, and time and cost consuming. Second, it is difficult to improve existing AMPs, which already have good physicochemical properties. Finally, even if new candidates are obtained in this way, the novel peptides will be similar in their sequence to existing peptides, and as such their diversity is expected to be poor. Thus, there is a need to devise efficient and accurate *in silico* approaches to novel AMP generation.

The problem of modelling AMPs was undertaken by a number of different computational approaches. One group of these approaches are classifiers, which take a peptide as input and their task is to predict whether the peptide is an AMP or not [6, 7, 8, 9], whether it is toxic [10, 11], or whether it is active [12, 13]. A related group of methods are quantitative structure–activity relationship (QSAR) models [14, 15], which identify a set of structural features for a given peptide, all of which are associated with the peptides being AMP, for example helical structure, amphipathicity,

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

etc. Next, the model is applied to a peptide database and peptides with the highest scores for features associated with being AMP are chosen. As such, the QSAR methods can only score existing peptides and are not able to directly generate new ones. Another approach is to use autoregressive models trained on AMP sequences for peptide generation [16, 17, 18]. To generate new peptides, these models operate in an iterative manner. In each iteration, a subsequence of the peptide constructed so far is given at input, and the model is used to propose the consecutive amino acid in the sequence. Other approaches are based on the genetic algorithms [19, 20, 21, 12]. These methods iteratively evolve a population of peptide sequences by adding random mutations, evaluating their fitness, and performing crossover and other evolutionary operations. Their performance depends on the choice of the method for introducing the mutations and for evaluating their fitness. Finally, there are linguistic models [22], that consider peptides as a formal language with grammar and vocabulary. By inserting alpha-helical patterns into AMP sequences in a sliding window manner, they are able to generate novel sequences in few attempts.

Working in the peptide space (both amino acid sequence, as well as atomic composition, e.g. encoded with graph-based representations such as the Simplified Molecular Input Line Entry Specification, or shortly SMILES [23]), as it is the case of the QSAR, genetic algorithm-based methods, and linguistic approaches, has serious disadvantages. First, the sequence space is sparse. Second, it is combinatorial and discrete, and thus highly dimensional, causing these approaches to be computationally demanding and prone to quickly getting stuck in local minima [21]. On top of that, similarity in this sequence space does not imply similarity of peptide function. Specifically, it is likely that operations such as amino acid substitutions, deletions, or additions, making small changes to the sequence, have large impact on amphiphacity or charge, and thus also antimicrobial activity of the peptide. Thus, it is desirable for the computational approach to find a better, continuous, and reduced-dimensionality representation of peptides, and operate on such representations instead.

Such representation learning approaches to peptide generation include GANs [24, 25] and variational autoencoders (VAE) [26, 13] as well as their conditional variants cGANs [27] and cVAEs [28, 29]. The conditional variants enable generation of peptides satisfying a given condition. In contrast to VAE, training of GANs was reported to face substantial technical obstacles, such as training instabilities and mode collapse [27].

On top of that, the existing approaches are not explicitly trained to perform all desired tasks. Specifically, almost all of above generative models, except for the genetic algorithm and cVAEs, are suitable only for a generation mode, which corresponds to random generation of AMPs, and which we refer to as *unconstrained* generation. In fact,

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

generation by improving existing peptides, which we call the *analogue generation*, should also be optimized during training. Moreover, ideally, in the analogue generation, the peptides should be generable both from known active AMP (positives) and non-AMPs (negatives). Indeed, the former mode would allow to directly mimic the experimental approach, while the latter is expected to increase the diversity of the pool of generated peptides. To our knowledge, none of the existing approaches are trained to improve non-AMPs. Finally, the results of the generative models are rarely experimentally verified, and their functionality is usually not made available to non-technical users in a digestible manner, e.g. as a web service.

Here, we propose HydrAMP, a novel approach for peptide generation, designed to address these needs. HydrAMP is a cVAE-based model, which is specifically trained to perform analogue generation both from positives and negatives, as well as for unconstrained generation. It learns a hidden space of meaningful peptide representations, which is disentangled from the set of antimicrobial conditions that a generated peptide is expected to satisfy — whether it is supposed to be a highly active AMP or not. The model is available as a web-service at `www.hydramp.mimuw.edu.pl`, and its results were experimentally verified and investigated in detail using *in silico* molecular dynamics simulations. As such, HydrAMP is a step forward in the daunting task of generating novel, highly active AMPs and fighting the problem of antimicrobial resistance.

2 Results

2.1 HydrAMP — a conditional, generative model of peptide sequences

HydrAMP is a model for generation of novel peptide sequences satisfying given antimicrobial activity conditions. A pair of conditions, denoted $\mathbf{c} = (c_{AMP}, c_{MIC})$ specifies whether the generated peptide is supposed to be antimicrobial (condition c_{AMP}) and whether it is supposed to have high antimicrobial activity, or, equivalently, low MIC (condition c_{MIC}). Despite the fact that the feature of being AMP and being highly active are strongly related, we keep them as separate conditions, because of existence of peptides that are known to be antimicrobial but have low activity.

The training data for HydrAMP consists of a curated data set of peptide sequences, including sequences that are known to be AMP, sequences that are known to have a low MIC, and sequences collected from UniProt (Figure 1a). The model is trained in three modes: *reconstruction*, *analogue* and *unconstrained* (Figure 1b). Training in the reconstruction mode facilitates the model to properly capture peptide sequences distribution, as well as those properties that make them antimicrobial and active. This is achieved by ensuring that the reconstructed peptides are similar to

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

the input peptides from the training data and satisfy the same conditions. In the analogue mode, the model is trained to generate analogues based on the provided prototype peptide and satisfying a specified condition. Finally, in the unconstrained mode, the model is trained to generate peptides *de novo* that resemble training data and satisfy the specified condition.

More formally, HydrAMP is an extension of a conditional variational autoencoder (cVAE). The model is optimized to create a meaningful, latent, real-valued vector space representation of peptides, which is easier to sample from and has a lower dimension than the original, highly dimensional and combinatorial space of peptide sequences. Apart from standard neural network-based sub-models such as *Encoder* and *Decoder*, used in the cVAE framework to operate on the latent representation in a probabilistic manner, the model utilizes also a *Classifier*. The Classifier is also a neural network, which, unlike the Encoder and the Decoder, is pre-trained prior to HydrAMP training, and is used in order to classify whether any given peptide is AMP or not, and whether it has a low MIC or not.

HydrAMP utilizes a number of regularization terms: latent reconstruction regularization, KL divergence, and Jacobian disentanglement regularization (Figure 1b). The former two regularization terms are standard in the cVAE framework. The latter is specifically introduced in this work for obtaining a disentanglement between the latent representation of peptides and the condition. In this way, the latent space encodes the property of being a peptide, while the condition independently encodes whether this peptide is supposed to be AMP or active (have low MIC). See Methods for a detailed formal description of the model.

HydrAMP offers two generation modes: *analogue* and *unconstrained* (Figure 1c). The analogue mode improves upon the common practice of novel peptide discovery followed in experimental labs. In contrast to the tedious trial and error process of changing the original sequence, given the prototype and the desired condition, the model manipulates the latent representation of the prototype instead. Given the favorable nature of the latent space, which spans real-valued vectors that were trained as representations of valid peptides, and given that distances in the latent space should reflect the dissimilarities between the peptides, the points in the latent space in close proximity to the point representing the original peptide are good candidates for analogue samples. For such samples from the latent space, the role of the Decoder is then to generate a sequence of amino acids satisfying the desired condition. HydrAMP benefits from an additional temperature parameter τ that controls the creativity of analogue generation. Intuitively, the temperature influences the radius with respect to the prototype point in the latent space within which the analogues are searched for. Compared to the analogue mode, the unconstrained mode is more standard for cVAE and allows

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

generation of random peptides with either random, or desired condition. To this end, samples are generated from a prior distribution over the latent space, and Decoder is used to produce sequences with the fixed condition.

2.2 HydrAMP outperforms other models in the ability to generate antimicrobial and highly active peptide candidates

HydrAMP was compared to two alternative models: *Basic* and *PepCVAE* [28]. Basic is a standard cVAE, while PepCVAE is a state-of-the-art approach to peptide generation using the conditional variational autoencoder framework. In case of Basic, only reconstruction mode training was performed. For PepCVAE we used reconstruction and unconstrained modes. Both models lack analogue generation and Jacobian disentanglement regularization (see Methods for further details).

To evaluate the generation performance of the evaluated methods we inspected their creativity as the number of generated peptides satisfying given conditions. We also inspected their ability to obtain desired antimicrobial properties, as specified by setting the condition, in a number of different generation tasks. To this end, in each task, each generated peptide p was assessed by the Classifier, and its probability of being AMP $\mathbb{P}_{M_{AMP}}(p)$ or being active $\mathbb{P}_{M_{MIC}}(p)$ was recorded.

In order to assess the analogue generation mode, we first performed a series of experiments, in which HydrAMP and the compared models were asked to generate analogues of known, existing AMPs (referred to as *positives*; see Figure 2a). In these experiments, model was given a prototype sequence as input and the conditions were set to $c_{AMP} = 1$ (being antimicrobial) and $c_{MIC} = 1$ (being highly active). A 1319 AMP peptides from the test set (not used during training) were used as prototypes. We defined two distinct evaluation criteria: *discovery* and *improvement*, with the former corresponding to the ability of simply generating analogues with good antimicrobial properties, and the latter corresponding to generating analogues with properties strictly better than the input prototype. Specifically, a peptide met the discovery criterion if it had a probability of being AMP greater than 0.8 and probability of having low MIC greater than 0.5. A peptide met the improvement criterion if its probability of being AMP and the probability of having low MIC were both greater than of the original peptide. Meeting the discovery criterion turned out to be a relatively easy task, as novel analogues were accepted for nearly 50% of input positive prototypes for all models, with HydrAMP showing a 4 percent point advantage in the fraction of accepted analogues (0.52) compared to the next best model, Basic (0.48). In contrast, analogues that met the improvement criterion were produced only for a few percent

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

of prototypes. Here, based on the same set of prototype sequences, the HydrAMP model improved over twice (82) as many peptides the PepCVAE (31) and over three times more than the Basic model (23). These results show that the HydrAMP model is more creative than the compared models and has the capabilities of introducing new analogues for original prototypes.

For further analysis of the analogue generation, we gave HydrAMP and the compared models a harder task of generating analogues for 1253 known non-AMP peptides (negatives) from the test set (Figure 2b). Unsurprisingly, meeting the discovery criterion was more challenging where the negatives were given as input compared to the positives. Still, HydrAMP managed to discover analogues for 82 prototypes, which is over 60% more than PepCVAE (52) and Basic (51) models. Notably, HydrAMP was also able to improve over 50% of the test set negatives, three times more than the best of the competing models, PepCVAE. In order to compare the antimicrobial properties of the prototype negatives and their analogues that met the discovery criterion, we analyzed the probability of being AMP and being active of both prototypes and the analogues.

We also compared these probability distributions against the distributions for all peptides from the test data. HydrAMP was able to produce analogues that met discovery criterion for peptides that had lower *a priori* probability of being AMP (Figure 2c) and lower probability of being highly active (Figure 2d) than its competitors.

In contrast, Basic and PepCVAE models produced analogues that met the discovery criterion for peptides that were already initially likely to be antimicrobial and active. The outstanding performance of HydrAMP in generating highly antimicrobial and active novel peptides based on experimentally-validated non-AMP prototypes shows its potential to provide truly novel antimicrobial peptides and increase the diversity of the pool of AMP sequences. Indeed, in contrast to analogues produced from positive prototypes, those peptides are expected to have sequences that largely differ from the sequences of known AMPs.

In order to assess the creativity of HydrAMP model in the analogue generation, we evaluated the number of different analogues created using two popular AMP peptides as prototypes: Pexiganan and CAMEL (Figure 2ef). Here, we used an additional parameter of HydrAMP analogue generation — the temperature $\tau > 0$, which controls the level of peptide alternations with respect to the original peptide (Figure 1c). Indeed, the greater the temperature, the more the generated analogues differ from the prototype (Figure 2ef). The ability to control the creativity of the model is important. The more similar a newly created analogue is to the prototype one, the more likely it is to preserve the physicochemical properties of the original peptide. On the other hand, the more alternations introduced, the larger

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

the number of generated novel peptides.

We assessed the unconstrained generation abilities of HydrAMP in comparison to other methods by generation of peptides at random in two modes: *positive* and *negative*. In the case of the *positive* mode we randomly sampled the latent representation variable and set the condition of the new peptide p to have a high probability of being AMP ($\mathbb{P}_{M_{AMP}}(p)$) and a high probability of being active ($\mathbb{P}_{M_{MIC}}(p)$). In the negative mode, we set the peptide conditions to be non-AMP ($c_{AMP} = 0$) and not active ($c_{MIC} = 0$). Peptides generated by HydrAMP were more likely to follow the desired conditions than the competing models both for AMP (Figure 2g) and MIC (Figure 2h) conditions. The advantage of HydrAMP was the most noticeable in the case of the probability of being AMP: the median probability of being AMP for peptides generated by HydrAMP was around 1, while for the next best model, PepCVAE, it was around 0.5 (Figure 2g).

In order to confirm that HydrAMP is able to suggest new, synthesizable active AMP peptides when using the unconstrained mode we tested to what degree peptides generated in this mode match appropriate criteria (Figure 2i). For that, we generated 50k candidate peptides and run four experiments confirming their properties. First, we computed the fraction of such peptides p which had high probability of being AMP ($\mathbb{P}_{M_{AMP}} > 0.8$) and high probability of being active ($\mathbb{P}_{M_{MIC}}(p) > 0.5$). Next, we assumed their secondary structure as alpha-helical and tested their amphipathicity using hydrophobic moment [30]. Peptides with amphiphatic alpha-helical structure are known to be easily synthesizable and their mode of action has been thoroughly studied [31]. Accordingly, we computed the fraction of peptides with the hydrophobic moment > 0.4 (the mean hydrophobic moment of the positive peptides no longer than 26 amino acids). Finally, we evaluated the fraction of peptides that had charge > 4 (the mean charge of the positive peptides no longer than 26 amino acids; see Supplementary Figure S1 for details of the hydrophobic moment and charge thresholds selection). HydrAMP suggested the largest number of peptides that were confirmed as positives, were sufficiently amphipathic, and were positively charged, in each condition separately. Eventually, HydrAMP created almost 50% more peptide candidates that matched all three conditions than its best competitor (Basic).

2.3 HydrAMP generates peptides with desired physicochemical properties

Next, we evaluated HydrAMP by inspecting the physicochemical properties of peptides that it generates, in comparison to the properties of known peptides as well as peptides generated by PepCVAE and Basic (Figure 3). Physicochemical properties of antimicrobial (positive) peptides differ from peptides that were experimentally verified not to

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

be antimicrobial (negative peptides). Indeed, a comparison of the distributions indicated that isoelectric point (Figure 3a), charge (Figure 3d), hydrophobic moment (Figure 3g), and aromaticity (Figure 3j) are significantly larger (one-sided Mann-Whitney test p-value < 0.05) for known positives than for negatives.

We first inspected the performance of HydrAMP and other approaches in the task of finding positive peptides based on negative prototypes in the analogue generation. Given such significant differences between negatives and positives, this requires the introduction of a shift in the physicochemical properties from the non-AMP prototypes to the newly generated peptides. In this task, HydrAMP showed a capacity to generate analogues with the desired significant increase (one-sided Wilcoxon test p-value < 0.05) of all investigated properties: isoelectric point (Figure 3b), charge (Figure 3e), hydrophobic moment (Figure 3h), and aromaticity (Figure 3k). We tested three temperature parameter setups: a conservative $\tau = 1.0$, and more explorative $\tau = 2.0$ and $\tau = 5.0$ levels. For higher temperature levels, the level of improvement was much higher than for Basic and PepCVAE models. This implies that we can control the balance between physicochemical improvement and the degree of changes between new analogues and their prototypes. We confirmed that these results could not be obtained by chance by computing physicochemical properties of a randomly sampled subset of the UniProt dataset and peptides for which we randomly sampled a sequence of amino acids. In both of these cases, HydrAMP in explorative temperature setups produces peptides with better qualities. In contrast to HydrAMP, all four physicochemical properties of peptides generated by PepCVAE or Basic showed no significant difference to the properties of known negatives.

Second, we evaluated HydrAMP and other approaches in a task of improving the positives. Here the challenge is different than in the previous task, as it is hard to improve peptides that are already ”good” (are already AMP and active). Here, for all analyzed physicochemical properties, HydrAMP generated peptides with significantly better (one-sided Wilcoxon test p-value < 0.05) properties than the input positive prototypes (Figure 3c,f,i,l). The benefit of the model's creativity using a larger temperature parameter is most visible for aromaticity (Figure 3l). PepCVAE and Basic performed only slightly worse in this task, either producing peptides that did not improve the hydrophobic moment (PepCVAE; Figure 3i) or aromaticity (Basic; Figure 3l).

Finally, we evaluated the physicochemical properties of positive and negative peptides generated by HydrAMP in the unconstrained mode, in comparison to peptides generated by other methods (Supplementary Figure S2). For all four analyzed physicochemical properties, there was large and significant difference between their distributions for HydrAMP generated negatives and positives (one-sided Wilcoxon test p-value < 0.05). For PepCVAE and Basic, the

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

differences between negatives and positives were also significant, but the medians of the distributions did not differ by as large amount as for HydrAMP. We confirmed that the amino acid distribution of the peptides generated in the unconstrained mode is in agreement with the true amino acid frequencies (Supplementary Figure S3). High content of lysine (K) and arginine (R), as well as low content of negatively charged glutamic acid (E) and aspartic acid (D) contribute to the positive net charge.

Overall, these results illustrate the superior performance of HydrAMP in generating peptides with desired physicochemical properties, reflecting the properties observed for real peptides.

2.4 HydrAMP suggested valid novel peptides

Finally, we inspected the ability of HydrAMP to generate novel peptides using experimental wet lab validation (for details see Supplementary Material S5). To this end, we applied HydrAMP in the analogue mode, with temperature (creativity) parameter $\tau = 1.0$, treating Pexiganan and Temporin-A as prototypes. The low creativity parameter choice was justified by the fact that we intended to obtain novel peptides that were similar to the prototypes. For each prototype, we first generated a set of 900 positive analogues. Next, we applied stringent filtering criteria to both generated sets of peptides (see Biological filtering criteria) to increase the chance of peptides being synthesizable and were left with 92 sequences for Pexiganan and 84 for Temporin-A. Out of these candidates, five positive analogues of Pexiganan and five of Temporin-A were selected for experimental validation. One of the selected Pexiganan analogues did not synthesize. The remaining were investigated experimentally and their antimicrobial activity was tested against two *E. coli* strains (Table 1; Methods). Laboratory experiments validated the known and predicted by the Classifier very high antimicrobial activity of Pexiganan. Out of four synthesizable analogues of Pexiganan, one showed even higher activity than that of Pexiganan (MIC = 2 μg/mL). This novel validated antimicrobial peptide was called Hydraganan. Another predicted analogue was confirmed as AMP and of high activity (MIC = 16 μg/mL).

The Classifier model predicted high probability of being antimicrobial, and a very low probability of being highly active for Temporin-A. The low activity was confirmed in the lab (MIC = 256 μg/mL). All five predicted positive analogues of Temporin-A also showed low antimicrobial activity (MIC $\geq$ 512 μg/mL). As controls, we randomly selected and validated two peptides generated as negative analogues by HydrAMP, one for Pexiganan, and one for Temporin-A. Both the negative analogues were validated as inactive (MIC $>$ 512 μg/mL).

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

2.5 Simulations confirmed that Hydraganan penetrates the cell membrane

To better understand the functional properties of Hydraganan and of the experimentally validated negative analogue of Pexiganan, we performed extensive, fully atomistic simulations of their behavior in the presence of a lipid membrane (Figure 4a, b). Each peptide was constructed as a regular α-helical structure and placed in aqueous compartment of the simulation box with its helical axis parallel to the membrane plane, at three different orientations. During the simulations we monitored the stability of peptides α-helical structure as well as the degree of their association with the membrane. The latter was assessed based on consecutive, 25 ns long simulation blocks, by considering a parameter $S \in [0,1]$ that indicates the fraction of peptide heavy atoms distribution along the membrane normal, that penetrates towards the bilayer core below a dividing plane corresponding to a depth at which the distribution of membrane heavy atoms reaches half of its maximum value (Figure 4c). Accordingly, $S = 0$ indicates no association, while $S = 1$ represents full peptide burial.

In the case of Hydraganan, irrespective of initial orientation, we observed peptide association with the membrane within the first 250 ns of the simulation time, as indicated by S values increasing above 0.5 (Figure 4d, lower plot). In all cases the association was followed by complete peptide burial ($S \to 1$) within 500 ns. In contrast, the negative Pexiganan analogue only loosely adhered to lipid bilayer surface and revealed no tendency to penetrate into its core, with S values remaining around 0.25 till the end of simulations. Notably, the active peptide consistently maintained ~ 0.75 fraction of helical geometry, that is considerably more than the inactive one, whose helical structure content dropped well below 0.5 already within the first 200 ns (Figure 4d, upper plot).

Together these findings underscore the importance of a stable α-helical scaffold which is necessary to secure properly aligned hydrophobic surface patch on one peptide side and oppositely facing cluster of hydrophilic or charged side chains. While the former provides a driving force for peptide insertion into the nonpolar membrane core, the latter apparently prevents membrane defect from closing. Notably, these features rely not merely on individual amino acids but rather on their appropriate distribution within the sequence that warrants correct placement upon folding as well as helical propensity of the resulting structure.

3 Discussion

In this work we have proposed HydrAMP, a generative model for antimicrobial peptides discovery. It leverages a conditional variational autoencoder to offer two functionalities: generating analogues of existing peptides with

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

improved antimicrobial properties (analogue generation) and generating peptides *de novo* (unconstrained generation). This is enabled by a continuous peptide representation of reduced dimensionality with disentangled antimicrobial conditions c_{AMP}, c_{MIC}. Additionally, this representation is directly optimized to not only properly represent the known peptides but also to efficiently generate new candidates.

We have evaluated the model's ability to improve existing antimicrobial peptides by producing analogues of Pexiganan and Temporin-A. We obtained Hydraganan, a potent antimicrobial peptide displaying better activity than Pexiganan, and another Pexiganan analogue of high activity. Hydraganan was verified experimentally and by molecular dynamics simulations. Importantly, the simulations were performed in the presence of a lipid membrane. To facilitate the usage of HydrAMP model, we developed a web service available freely at `https://hydramp.mimuw.edu.pl/`.

HydrAMP bears several novelties and advantages in comparison to existing approaches. First, in contrast to previous VAE/WAE-based approaches [26, 28, 29], HydrAMP was trained specifically for the task of analogue generation. Additionally, HydrAMP is the first model trained and used to identify active analogues of non-AMP prototypes. Third, leveraging a Gumbel-Softmax approximation [32] enabled a continuous approximation of sampling in the discrete space of amino acid sequences and thus a direct optimization of peptides generated by the model. Before, such optimizations required a complex multi-stage training [28]. Finally, HydrAMP is the only model controls in a parametrized way the model creativity understood as the number of modifications introduced to the query peptide.

Although HydrAMP conveys advancement over previous approaches, it still could be extended in several ways. First, HydrAMP could be enriched with a popular attention mechanism [33]. The architecture of the model allows easy replacement of both the Encoder and the Decoder with transformer modules [33]. It is not clear, however, that such an extension of the model would improve its performance, as initial evidence suggested that attention-based models might work worse for peptide modelling than recurrent models [34]. Additionally, we considered amino acid sequence only. Instead, we could leverage structure prediction models such as AlphaFold [35], and extend HydrAMP training procedure to directly optimize alpha-helix propensity and other structural features of AMPs. In this work, the loss term responsible for reconstruction objective treats each of the amino acids in the sequence as independent. However, the probability of each amino acid being present is dependent on proceeding and following amino acids. The local context can also affect the antimicrobial activity of a given peptide, as shown in [36].

Finally, the applicability of the model could be also extended. HydrAMP does not consider the host toxicity of

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

the peptides. Successfully optimizing both activity and host toxicity within the same model would provide a powerful tool for AMP design. Addressing the toxicity is an urgent matter as only highly selective peptides can be used in treatment of human subjects. Additionally, we trained the model to target only *E. coli* strains. Depending on the data availability, it could be retrained to target other strains, either Gram-positive or multi-drug resistant strains.

Antimicrobial resistance problem remains a serious threat that needs dire actions. HydrAMP is a novel, powerful tool capable of suggesting new antimicrobial peptides what has been shown both through wet-lab evaluation and molecular dynamics simulations. Thus, by quick indication of possible candidates HydrAMPs is making progress towards creation of a new generation of antibiotics.

4 Methods

4.1 Data collection

The training data for the model is a curated data set of a total of 247 506 peptide sequences collected from three sources: MIC data collected by [12], AMP data from known AMP databases, and UniProt data [37]). To speed up the training process and control the synthesis costs, all selected training sequences have length of at most 25 amino acids and contain only standard amino acids. Sequences shorter than 25 amino acids were padded using additional padding character.

MIC data The MIC data set consists of peptides with experimentally proven antimicrobial activity and measured MIC values downloaded from GRAMPA database [12]. Of 8049 entries only 4546 of them are unique sequences. We select those peptides which were experimentally tested for activity against E. coli strains. MIC values for peptides with multiple measurements are averaged resulting in 4546 entries. Only 3444 peptides have the required length of 25 amino acids. Peptides with MIC $<$ 30 μg/ml are labeled as active (positive), while peptides with MIC $\geq$ 30 μg/ml are labelled as inactive (negative). MIC data set contains 2126 positives and 1318 negatives.

AMP data The AMP data consists of positive and negative data set. Positive examples are sequences from manually curated AMP databases. We combine experimentally validated peptides from dbAMP [38], AMP Scanner [8], and DRAMP [39] databases. Duplicate sequences are removed. Negative examples are assumed to be biologically inactive and are obtained manually using UniProt search filters, requiring subcellular location: cytoplasm, and excluding the following properties: antimicrobial, antibiotic, antiviral, antifungal, effector, excreted. To increase data set diversity,

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

negative sequences sharing $\geq 40\%$ sequence identity are removed, and replaced with the representative sequences of the clusters using CD-HIT [40, 41]. The resulting set of representative sequences contains only sequence that are longer than 25 amino acids, thus negative sequences are randomly cropped to match the positive data set length distribution in order to avoid length bias. The positive and negative data sets are of equal size of 11,131 sequences each.

UniProt data To extend the biological diversity in the training process we downloaded sequences of unknown antimicrobial properties of the desired length from UniProt database. This gave 225 244 additional sequences with no duplicates.

4.2 AMP and MIC classifiers learning

In order to predict the properties of generated peptides, we trained a pair of classifiers, that for a given peptide p, predict its probabilities of being antimicrobial and active. For the prediction of a given peptide being antimicrobial (AMP), we implemented and trained the model from AMP Scanner [8], using AMP data. We refer to this model as M_{AMP}. For the prediction of peptide activity, we trained a model using MIC data, denoted M_{MIC}, with a custom architecture. The model consists of six layers. The first layer is an one-hot encoded input in form of a sequence of amino acids. The second layer is an 128-dimensional embedding layer of each of the amino acids. Next, we use the LSTM layer of 64 units [42]. Downsampling is performed by 1-dimensional pooling with kernel size of 5 and stride of 5. Another LSTM layer of 100 units is used where only the last output is return. The final output is fully connected layer with a sigmoid activation.

For a given peptide p, the M_{AMP} returns the probability of p being antimicrobial, denoted $\mathbb{P}_{M_{AMP}}(p)$. The M_{MIC} returns the probability of p being active, denoted $\mathbb{P}_{M_{MIC}}(p)$. From now on we assume that these probabilities are obtained only using M_{AMP} and M_{MIC} classifiers, and we refer to them jointly as Classifier. Both models are implemented in Keras [43].

Cross-validation results indicate highly accurate classification results for both M_{AMP} and M_{MIC} (Supplementary Table S1).

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

4.3 HydrAMP model

HydrAMP is an extended conditional VAE (cVAE) model. Its main generative part is a Decoder model that given a *latent variable* $z \in \mathbb{R}^{latent}$ and a pair of conditions $\mathbf{c} = (c_{AMP}, c_{MIC})$ produces a distribution over peptides $Dec(z, \mathbf{c})$. We refer to the likelihood of a peptide p w.r.t. to this distribution as $\mathbb{P}_{Dec(z,\mathbf{c})}(p)$. The latent variable is assumed to follow the *latent prior* distribution $P_z \sim \mathcal{N}(0, I)$.

The Decoder is trained so that the peptides sampled from $Dec(z, \mathbf{c})$ follow a given pair of conditions $\mathbf{c}$ and resemble valid peptides from the training dataset. As optimizing the Decoder directly is not feasible, we introduced an additional Encoder model trained to provide a variational approximation $q(z|p)$ of the posterior distribution $\mathbb{P}(z|p)$. The variational posterior approximation for a peptide p is set to be a normal distribution $\mathcal{N}\left(\mu_{q(\cdot|p)}, diag(\sigma_{q(\cdot|p)})\right)$, where $\mu_{q(\cdot|p)} \in \mathbb{R}^{latent}$ and $\sigma_{q(\cdot|p)} \in \mathbb{R}_{+}^{latent}$. Both the Decoder and Encoder are modeled as neural networks. For the detailed architecture see Supplementary Material S3.

The HydrAMP model is optimized using three objectives: reconstruction, analogue (see Supplementary Material S1.1), and unconstrained (see Supplementary Material S1.2). The reconstruction objective aims at teaching the Decoder how to generate valid peptide structures by reconstructing known peptides. The analogue and unconstrained objectives mimic the process of the analogue and unconstrained generation during training. Besides that, we also applied a two-fold regularization to the HydrAMP model: the Jacobian disentanglement regularization to encourage disentanglement between latent variable z and a pair of conditions $\mathbf{c}$, and latent reconstruction regularization for better latent variable preservation properties.

In the formulas below we use the following notation. Denote a constant $c \in [0, 1]$, p a peptide and $P_M(p)$ its probability in a classifier model M, with $M \in \{M_{AMP}, M_{MIC}\}$. Let

$$H_M(c, p) = \log\left(\mathbb{P}_M(p)^c (1 - \mathbb{P}_M(p))^{1-c}\right), \tag{1}$$

be the cross-entropy between the $Bernoulli(c)$ and $Bernoulli(\mathbb{P}_M(p))$. We define

$$H_\Sigma(\mathbf{c} = (c_{AMP}, c_{MIC}), p) = \Sigma_{cond \in \{AMP, MIC\}} H_{M_{cond}}(c_{cond}, p), \tag{2}$$

where $\mathbf{c} = (c_{AMP}, c_{MIC})$ is the pair of conditions, as the sum of the cross entropies for the two different conditions c_{AMP}, c_{MIC} and their probabilities $\mathbb{P}_{M_{MIC}}(p)$ and $\mathbb{P}_{M_{AMP}}(p)$, respectively.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

4.3.1 Reconstruction objective

The reconstruction objective forces the model to capture the structure of valid peptides collected from available databases. We achieve this by training the Decoder and Encoder maximizing the conditional evidence lower bound (ELBO) introduced in [44]. For each peptide $p \sim \mathbb{P}_{\mathcal{X}}$, we compute its pair of conditions $\mathbf{c}_p = (\mathbb{P}_{AMP}(p), \mathbb{P}_{MIC}(p))$ using the Classifier. Next, we maximize a conditional ELBO given by:

$$ELBO^{\beta}_{rec} = \mathbb{E}_{z\sim q(z|p)}\big(\log \mathbb{P}_{Dec(z,\mathbf{c}_p)}(p) + \tag{3}$$

$$\lambda^{class}_{rec}\mathbb{E}_{p`\sim Dec(z,\mathbf{c}_p)} H_{\Sigma}(\mathbf{c}_p, p`)\big) - \tag{4}$$

$$\beta \cdot KL(q(z|p)\|P_z), \tag{5}$$

where (3) is the expected log likelihood of reconstruction of the initial peptide p, (4) is the expected log likelihood of recovering the initial peptide pair of conditions $\mathbf{c}_p$ with parameter $\lambda^{class}_{rec} > 0$ and (5) is a β-VAE regularization term [45] with parameter $\beta > 0$ decaying in the process of training. Parameters β and λ^{class}_{rec} control the trade-off between reconstruction of the original peptide, satisfying the peptide condition and keeping the posterior approximation $q(z|p)$ close to the prior. The expectations w.r.t. $q(z|p)$ are obtained using a reparametrization trick [46] and are approximated using a single sample from $q(z|p)$.

4.3.2 Jacobian disentanglement regularization

For every generation process used by the HydrAMP model, a newly generated peptide should have properties provided by a pair of conditions $\mathbf{c}$. To measure that, let us define an average condition reconstruction function of the *Dec* distribution as:

$$ACR^{Dec}(z, c = (c_{AMP}, c_{MIC})) = \mathbb{E}_{p\sim Dec(z,c_{AMP})}(\mathbb{P}_{AMP}(p), \mathbb{P}_{MIC}(p)). \tag{6}$$

ACR^{Dec} measures the expected likelihood of the peptides sampled from the distribution *Dec* modeled by the Decoder, for a given z and pair of conditions $\mathbf{c}$, actually satisfying $\mathbf{c}$. Ideally $ACR^{Dec}(z, \mathbf{c}) = \mathbf{c}$, which means that on average peptides sampled from $Dec(z, \mathbf{c})$ have properties defined by the pair of conditions $\mathbf{c}$. This means that in an ideal scenario the ACR^{Dec} is constant w.r.t. z and in case when ACR^{Dec} is differentiable w.r.t. to z, the following condition

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

holds:

$$\frac{\partial ACR^{Dec}}{\partial z} \equiv \mathbf{0}, \tag{7}$$

where $\mathbf{0}$ is an all zero matrix. In that case z and $\mathbf{c}$ are disentangled w.r.t. to ACR^{Dec}, because any change of z does not affect the expected pair of conditions of newly generated candidates. In order to impose this property we introduce the following Jacobian disentanglement regularization function:

$$JDR^{Dec}(z,\mathbf{c}) = \sum_{i=1,j=1}^{2,latent} Huber\left(\frac{\partial ACR_i^{Dec}(z,\mathbf{c})}{\partial z_j}\right), \tag{8}$$

where $Huber : \mathbb{R} \to \mathbb{R}$ (sometimes also referred as a smooth $L1$) is given by

$$Huber(x) = \min(|x|, x^2). \tag{9}$$

We use the Huber function as it is less prone to be affected by outlier examples.

Accordingly, we extended reconstruction objective with the following term:

$$JDR_{rec} = \mathbb{E}_{z\sim q(z|p)} JDR^{Dec}(z,\mathbf{c}_p), \tag{10}$$

where p is a peptide being reconstructed and $\mathbf{c}_p$ is its pair of conditions. We approximate this expectation with a single sample from $q(z|p)$ using the reparametrization trick.

Analogue and unconstrained objectives were also extended with the Jacobian disentanglement regularization term. See Supplementary Material S2.1 for the details.

4.3.3 Latent reconstruction regularization

In the cVAE framework, the Decoder plays a role similar to the inverse function of the Encoder. Indeed, the Decoder aims to reconstruct the peptide fed to the Encoder that is sampled from a posterior distribution generated by the Encoder. To further impose that relation, similarly to [28], we introduced an additional latent reconstruction regularization objective. Consider a peptide p and its posterior mean $\mu_{q(z|p)}$ given by the Encoder. Peptide p' returned by the Decoder for a point sampled from that posterior can be given as input to the Encoder and will obtain its posterior mean $\mu_{q(z|p')}$. The latent reconstruction regularization objective enforces the two posterior means to be similar. To

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

this end, we minimize the following expectation for the reconstruction objective:

$$LRR_{rec} = \mathbb{E}_{z\sim q(z|p)}\mathbb{E}_{p'\sim Dec(z,\mathbf{c}_p)} \left|\left|\mu_{q(\cdot|p)} - \mu_{q(\cdot|p')}\right|\right|_2^2, \tag{11}$$

where p is a reconstructed peptide and $\mathbf{c}_p$ is its pair of conditions. Using a reparametrization trick, we approximate this expectation with a single sample from $q(z|p)$ and the expectation w.r.t. *Dec* is approximated using a Gumbel Softmax [32].

When the expectation above is low, the Encoder preserves the latent code of the average peptide sampled from the Decoder. This property is essential especially for the analogue generation, as we assume the similarity between the analogue and original prototype because we sample both from precisely the single posterior distribution over latent codes. If the Encoder preserves the latent code of a generated analogue, then its similarity to prototype is granted for continuous Decoder models.

Analogue and unconstrained objectives were also extended with the latent reconstruction regularization terms. See Supplementary Material S2.2 for the details.

4.4 Basic and PepCVAE models

We compared HydrAMP with two other models: *Basic* and *PepCVAE* [28]. We used our own implementation of the PepCVAE model, as its code was not made publicly available by the authors.

Basic model corresponds to the standard cVAE model (see [44]), and is trained in the same manner as HydrAMP, but is optimized only for reconstruction objective and uses only the latent reconstruction regularization. PepCVAE model was trained to optimize the same objectives as Basic, but was additionally optimized for the unconstrained objective. Both PepCVAE and Basic models lack optimization of analogue objective and Jacobian disentanglement regularization, which are incorporated in HydrAMP.

Since the Basic, PepCVAE and HydrAMP are increasingly complex, such a selection of models for comparison effectively implements an ablation study.

4.4.1 Training procedure

We trained all models (HydrAMP, Basic, and PepCVAE) using ADAM [47] optimizer. The batch size was equal to 384, and each batch consisted of 128 peptides from AMP, 128 from MIC data, and 128 from UniProt data. We trained every model for 40 epochs. Reconstruction and analogue objectives were optimized using all peptides in each

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

batch. Additionally, in each iteration, the unconstrained objective was optimized using 128 samples from P_z. The Gumbel temperature used in Gumbel Softmax [32], t, was scheduled using exponential decay from 2.0 to 0.1 for 24 epochs and then kept stable at 0.1. The β parameter was increased from 10^{-4} to 10^{-2} via exponential annealing.

For the loss function, evaluation metrics and model selection details for training of HydrAMP, Basic and PepCVAE models see Supplementary Material S4.

4.5 Post-training prior refinement

Following successful results presented in [48], after the end of the training, we refined our P_z prior distribution to better match an aggregated posterior:

$$\mathbb{P}_z^{agg} = \mathbb{E}_{p \in \mathcal{X}} q(\cdot|p), \tag{12}$$

which is an average of all posterior distributions of peptides from dataset $\mathcal{X}$. According to [49] this distribution is the latent prior distribution which maximizes likelihood of data from $\mathcal{X}$ when Encoder and Decoder models are fixed. However, this property makes it prone to over-fitting. Because of that we decided to use less complicated distribution to approximate the aggregated posterior. This new distribution $\hat{\mathbb{P}}_z^{agg}$ is set to be a normal distribution $\hat{\mathbb{P}}_z^{agg} = \mathcal{N}\left(\hat{\mu}, \hat{\Sigma}\right)$ optimized to maximize likelihood of the set of aggregated variational posterior means $\left\{\mu_{q(\cdot|p)} | p \in \mathcal{X}\right\}$. The distribution parameters $\hat{\mu}, \hat{\Sigma}$ are selected using a classical PCA algorithm [50].

4.6 Generation modes of peptides

After the model is trained we use it in order to generate novel peptides. There are the following modes of this process:

4.6.1 Unconstrained generation

The following algorithm is used for generation of active and antimicrobial peptides in an unconstrained manner where z is sampled from a refined prior $\hat{\mathbb{P}}_z^{agg}$:

The algorithm above refers to the positive mode, while in the negative mode we sample peptides with conditions $(c^{AMP} = 0, c^{MIC} = 0)$. We skip peptides with $\mathbb{P}_{M_{AMP}}(p) > 0.2$ and select the peptide with the lowest $\mathbb{P}_{M_{MIC}}(p)$.

4.6.2 Analogue generation

The following algorithm was used for generation of active peptides similar to prototype peptide p_{proto}. In this

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Data: A number of tries for generation `nb_of_tries`.
Result: A new active and anti-microbial peptide sampled from $\hat{\mathbb{P}}_z^{agg}$.
$z \leftarrow sample(\hat{P}_z^{agg})$;
$best_candidate \leftarrow None$;
$best_candidate_mic \leftarrow 0$;
for $i \leftarrow$ `range(nb_of_tries)` **do**
  $new_candidate \leftarrow sample(Dec(z, (c^{AMP} = 1, c^{MIC} = 1)))$;
  if $\mathbb{P}_{M_{AMP}}(new_candidate) > 0.8$ ***and*** $\mathbb{P}_{M_{MIC}}(new_candidate) > best_candidate_mic$ **then**
    $best_candidate \leftarrow new_candidate$;
    $best_candidate_mic \leftarrow \mathbb{P}_{M_{MIC}}(new_candidate)$;
  end
end
return $best_candidate$;

Algorithm 1: Unconstrained generation in a positive mode

algorithm we introduce a creativity parameter $\tau \in (0, +\infty)$. Prototypes are sampled from a modified variational posterior $\mathcal{N}(\mu_{q(\cdot|p_{proto})}, \tau^2 \cdot \sigma_{q(\cdot|p_{proto})})$ with a covariance matrix rescaled by a factor of τ^2. This means that the closer τ is to 0 - the peptide sampled is similar to the one assigned to the posterior mode. On the other hand - for $\tau > 1$ the sampling probability distribution has the same mean but greater variance than in case of posterior approximation what encourages generating peptides differing from p_{proto} to a greater degree.

Data: A prototype peptide p_{proto}, creativity parameter $\tau \in (0, +\infty)$, number of tries for generation `nb_of_tries`.
Result: A new active and anti-microbial peptide sampled from similar to peptide p_{proto}.
$\mu_{q(\cdot|p_{proto})}, \sigma_{q(\cdot|p_{proto})} \leftarrow Enc(p_{proto})$;
$best_candidate \leftarrow None$;
$best_candidate_mic \leftarrow 0$;
for i *in* `range(nb_of_tries)` **do**
  $z_{proto} \sim \mathcal{N}(\mu_{q(\cdot|p_{proto})}, \tau^2 \cdot \sigma_{q(\cdot|p_{proto})})$;
  $new_candidate \sim Dec(z_{proto}, (c^{AMP} = 1, c^{MIC} = 1))$;
  if $\mathbb{P}_{M_{AMP}}(new_candidate) \geq 0.8$ ***and*** $\mathbb{P}_{M_{MIC}}(new_candidate) > best_candidate_mic$ **then**
    $best_candidate \leftarrow new_candidate$;
    $best_candidate_mic \leftarrow \mathbb{P}_{M_{MIC}}(new_candidate)$;
  end
end
return best_candidate;

Algorithm 2: Analogue generation for a positive mode

An analogue meets the improvement criteria when it increased $\mathbb{P}_{M_{AMP}}(p)$ and $\mathbb{P}_{M_{MIC}}(p)$ with respect to the input peptide. An analogue meets the discovery criteria with $\mathbb{P}_{M_{AMP}}(p) \geq 0.8$ and $\mathbb{P}_{M_{MIC}}(p) > 0.5$.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

4.7 Biological filtering criteria

In general, the biological criteria serve as approximation of expert selection that takes into account peptide synthesizability. First, we filter out all of the known AMPs we collected in our AMP data set. Then, we exclude sequences in which in a window of 5 amino acids there were more than 3 positively charged amino acids (K, R). Finally, we remove sequences in which occur three hydrophobic amino acids in a row. We consider as hydrophobic following amino acids based on Eisenberg scale [51]: F, I, L, V, W, M, A.

In case of selection of peptides for experimental validation we use more stringent criteria. We remove sequences of known AMPs, and sequences with accumulation of positive charge, as described above. Additionally, we remove any sequence in which three amino acids in a row are the same. We also exclude sequences containing cysteines (C).

4.8 Computer simulations of peptide-membrane systems

Given the amino acids sequence of interest, a peptide starting structure was modelled as a regular α-helix using the Discovery Studio Visualizer 2021 program (Dassault Systèmes, BIOVIA) [52]. It was then submitted to the CHARMM-GUI service [53] for peptide-membrane system construction. Standard protonation states were assigned to titratable peptide residues together with charged N-terminus and amidated C-terminus. The peptide was placed such that its centre of mass was 3.5 nm from the midplane of a rectangular lipid bilayer patch consisting of 120 1-palmitoyl-2-oleoyl-sn-glycero-3-phosphoethanolamine (POPE) and 60 1-palmitoyl-2-oleoyl-sn-glycero-3-(phospho-rac-(1-glycerol)) (POPG) molecules, and its helical axis was oriented parallel to the membrane. A random peptide rotation angle along the helical axis was assigned and two additional transformations by 120 and 240 degrees, respectively, were generated, resulting in three independent starting structures, each with different peptide side facing the membrane. A rectangular simulation box providing 1.5 nm of solvent margins on both membrane sides was filled with water molecules together with K^+ and Cl^- ions, whose number was chosen to achieve 0.15 mol/l concentration and to neutralise the total system charge. The fully atomistic CHARMM36 force field [54] was used for protein and ions, and the rigid TIP3P model [55] was used for water. All simulations were carried out using the Gromacs program [56] with the default simulation set up implemented in CHARMM-GUI for the chosen force field combination and periodic boundary conditions [57]. The protocol included potential energy minimisation, six rounds of equilibration with step-wise removal of positional restraints for peptide and lipids, and production runs. The latter were

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

conducted at ambient pressure and temperature of 310 K for 500 ns. Three independent simulations with individual starting structures were performed for each peptide-membrane system. The stability of peptides α-helical structure was assessed by the DSSP program [58].

Data and code availability

The HydrAMP source code, data used for training, and the scripts for generation of the results can be found at `https://github.com/szczurek-lab/hydramp`. Free web-service is available with all the functionalities can be accessed at `https://hydramp.mimuw.edu.pl/`.

References

[1] US Department of Health, Human Services, et al. "CDC. Antibiotic Resistance Threats in the United States, 2019". In: *Atlanta, GA, USA: US Department of Health and Human Services, CDC* (2019).

[2] Jim O'Neill. "Tackling drug-resistant infections globally: final report and recommendations". In: (2016).

[3] Maria Magana et al. "The value of antimicrobial peptides in the age of resistance". In: *The Lancet Infectious Diseases* (2020).

[4] Lloyd Czaplewski et al. "Alternatives to antibiotics—a pipeline portfolio review". In: *The Lancet infectious diseases* 16.2 (2016), pp. 239–251.

[5] Håvard Jenssen, Pamela Hamill, and Robert EW Hancock. "Peptide antimicrobial agents". In: *Clinical microbiology reviews* 19.3 (2006), pp. 491–511.

[6] Chenkai Li et al. "AMPlify: attentive deep learning model for discovery of novel antimicrobial peptides effective against WHO priority pathogens". In: (Dec. 2020). DOI: `10.21203/rs.3.rs-120780/v1`.

[7] Qinze Yu et al. "HMD-AMP: Protein Language-Powered Hierarchical Multi-label Deep Forest for Annotating Antimicrobial Peptides". In: *arXiv preprint arXiv:2111.06023* (2021).

[8] Daniel Veltri, Uday Kamath, and Amarda Shehu. "Deep learning improves antimicrobial peptide recognition". In: *Bioinformatics* 34.16 (2018), pp. 2740–2747.

[9] Ernest Y Lee et al. "Mapping membrane activity in undiscovered peptide sequence space using machine learning". In: *Proceedings of the National Academy of Sciences* 113.48 (2016), pp. 13588–13593.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

[10] Patrick Timmons and Chandralal Hewage. "HAPPENN is a novel tool for hemolytic activity prediction for therapeutic peptides which employs neural networks". In: *Scientific Reports* 10 (July 2020), p. 10869. DOI: 10.1038/s41598-020-67701-3.

[11] Fabien Plisson, Obed Ramirez, and Cristina Martinez-Hernández. "Machine learning - guided discovery and design of non-hemolytic peptides-annotated". In: *Scientific Reports* 10 (Oct. 2020). DOI: 10.1038/s41598-020-73644-6.

[12] Jacob Witten and Zack Witten. "Deep learning regression model for antimicrobial peptide design". In: *BioRxiv* (2019), p. 692681.

[13] Scott N. Dean et al. "PepVAE: Variational Autoencoder Framework for Antimicrobial Peptide Generation and Activity Prediction". In: *Frontiers in Microbiology* 12 (2021), p. 2764. ISSN: 1664-302X. DOI: 10.3389/fmicb.2021.725727. URL: https://www.frontiersin.org/article/10.3389/fmicb.2021.725727.

[14] Ernest Lee et al. "Mapping membrane activity in undiscovered peptide sequence space using machine learning". In: *Proceedings of the National Academy of Sciences of the United States of America* 113 (Nov. 2016). DOI: 10.1073/pnas.1609893113.

[15] Marlon H. Cardoso et al. "Computer-Aided Design of Antimicrobial Peptides: Are We Generating Effective Drug Candidates?" In: *Frontiers in Microbiology* 10 (2020), p. 3097. ISSN: 1664-302X. DOI: 10.3389/fmicb.2019.03097. URL: https://www.frontiersin.org/article/10.3389/fmicb.2019.03097.

[16] Christina Wang, Sam Garlick, and Mire Zloh. "Deep Learning for Novel Antimicrobial Peptide Design". In: *Biomolecules* 11.3 (2021), p. 471.

[17] Jeanne Trinquier et al. "Efficient generative modeling of protein sequences using simple autoregressive models". In: *arXiv preprint arXiv:2103.03292* (2021).

[18] Marwin HS Segler et al. "Generating focused molecule libraries for drug discovery with recurrent neural networks". In: *ACS central science* 4.1 (2018), pp. 120–131.

[19] Mari Yoshida et al. "Using evolutionary algorithms and machine learning to explore sequence space for the discovery of antimicrobial peptides". In: *Chem* 4.3 (2018), pp. 533–543.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

[20] Kyle Boone et al. "Combining genetic algorithm with machine learning strategies for designing potent antimicrobial peptides". In: *BMC Bioinformatics* 22 (May 2021). DOI: 10.1186/s12859-021-04156-x.

[21] William Porto et al. "In silico optimization of a guava antimicrobial peptide enables combinatorial exploration for peptide design". In: *Nature Communications* 9 (Apr. 2018). DOI: 10.1038/s41467-018-03746-3.

[22] William F Porto et al. "Joker: An algorithm to insert patterns into sequences for designing antimicrobial peptides". In: *Biochimica et Biophysica Acta (BBA)-General Subjects* 1862.9 (2018), pp. 2043–2052.

[23] David Weininger. "SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules". In: *Journal of Chemical Information and Computer Sciences* 28.1 (1988), pp. 31–36. DOI: 10.1021/ci00057a005. eprint: https://pubs.acs.org/doi/pdf/10.1021/ci00057a005. URL: https://pubs.acs.org/doi/abs/10.1021/ci00057a005.

[24] Donatas Repecka et al. "Expanding functional protein sequence spaces using generative adversarial networks". In: *Nature Machine Intelligence* 3.4 (2021), pp. 324–333.

[25] Andrejs Tucs et al. "Generating ampicillin-level antimicrobial peptides with activity-aware generative adversarial networks". In: *ACS omega* 5.36 (2020), pp. 22847–22851.

[26] Scott Dean and Scott Walper. "Variational Autoencoder for Generation of Antimicrobial Peptides". In: *ACS Omega* XXXX (Aug. 2020). DOI: 10.1021/acsomega.0c00442.

[27] Colin M Van Oort et al. "AMPGAN v2: Machine Learning-Guided Design of Antimicrobial Peptides". In: *Journal of Chemical Information and Modeling* (2021).

[28] Payel Das et al. "Pepcvae: Semi-supervised targeted design of antimicrobial peptide sequences". In: *arXiv preprint arXiv:1810.07743* (2018).

[29] Payel Das et al. "Accelerated antimicrobial discovery via deep generative models and molecular dynamics simulations". In: *Nature Biomedical Engineering* 5.6 (2021), pp. 613–623.

[30] David Eisenberg, Robert M Weiss, and Thomas C Terwilliger. "The helical hydrophobic moment: a measure of the amphiphilicity of a helix". In: *Nature* 299.5881 (1982), pp. 371–374.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

[31] Igor Zelezetsky and Alessandro Tossi. "Alpha-helical antimicrobial peptides—Using a sequence template to guide structure–activity relationship studies". In: *Biochimica et Biophysica Acta (BBA) - Biomembranes* 1758.9 (2006). Membrane Biophysics of Antimicrobial Peptides, pp. 1436–1449. ISSN: 0005-2736. DOI: https://doi.org/10.1016/j.bbamem.2006.03.021. URL: https://www.sciencedirect.com/science/article/pii/S0005273606001088.

[32] Eric Jang, Shixiang Gu, and Ben Poole. *Categorical Reparameterization with Gumbel-Softmax*. 2017. arXiv: 1611.01144 [stat.ML].

[33] Ashish Vaswani et al. *Attention Is All You Need*. 2017. arXiv: 1706.03762 [cs.CL].

[34] Nicki Skafte Detlefsen, Søren Hauberg, and Wouter Boomsma. "What is a meaningful representation of protein sequences?" In: *ArXiv* abs/2012.02679 (2020).

[35] John Jumper et al. "Highly accurate protein structure prediction with AlphaFold". In: *Nature* 596.7873 (2021), pp. 583–589.

[36] Sam Clark et al. "The lexicon of antimicrobial peptides: a complete set of arginine and tryptophan sequences". In: *Communications biology* 4.1 (2021), pp. 1–14.

[37] Alex Bateman et al. "UniProt: the universal protein knowledgebase in 2021". In: *Nucleic Acids Research* (2020).

[38] Jhih-Hua Jhong et al. "dbAMP: an integrated resource for exploring antimicrobial peptides with functional activities and physicochemical properties on transcriptome and proteome data". In: *Nucleic acids research* 47.D1 (2019), pp. D285–D297.

[39] Xinyue Kang et al. "DRAMP 2.0, an updated data repository of antimicrobial peptides". In: *Scientific data* 6.1 (2019), pp. 1–10.

[40] Weizhong Li, Lukasz Jaroszewski, and Adam Godzik. "Tolerating some redundancy significantly speeds up clustering of large protein databases". In: *Bioinformatics* 18.1 (2002), pp. 77–82.

[41] Limin Fu et al. "CD-HIT: accelerated for clustering the next-generation sequencing data". In: *Bioinformatics* 28.23 (2012), pp. 3150–3152.

[42] Sepp Hochreiter and Jürgen Schmidhuber. "Long short-term memory". In: *Neural computation* 9.8 (1997), pp. 1735–1780.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

[43] Francois Chollet et al. *Keras*. 2015. URL: https://github.com/fchollet/keras.

[44] Zhiting Hu et al. "Toward controlled generation of text". In: *arXiv preprint arXiv:1703.00955* (2017).

[45] Irina Higgins et al. "beta-vae: Learning basic visual concepts with a constrained variational framework". In: (2016).

[46] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *arXiv preprint arXiv:1312.6114* (2013).

[47] Diederik P Kingma and Jimmy Ba. "Adam: A method for stochastic optimization". In: *arXiv:1412.6980* (2014).

[48] Partha Ghosh et al. *From Variational to Deterministic Autoencoders*. 2019. arXiv: 1903.12436 [cs.LG].

[49] Jakub Tomczak and Max Welling. "VAE with a VampPrior". In: *International Conference on Artificial Intelligence and Statistics*. PMLR. 2018, pp. 1214–1223.

[50] Karl Pearson F.R.S. "LIII. On lines and planes of closest fit to systems of points in space". In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2.11 (1901), pp. 559–572. DOI: 10.1080/14786440109462720.

[51] D Eisenberg et al. "Analysis of membrane and surface protein sequences with the hydrophobic moment plot". In: *Journal of molecular biology* 179.1 (1984), pp. 125–142.

[52] Dassault Systèmes BIOVIA. "Discovery Studio, version 21.1. 0". In: *San Diego: Dassault Systèmes* (2021).

[53] Sunhwan Jo et al. "CHARMM-GUI: a web-based graphical user interface for CHARMM". In: *Journal of computational chemistry* 29.11 (2008), pp. 1859–1865.

[54] Jeffery B. Klauda et al. "Update of the CHARMM All-Atom Additive Force Field for Lipids: Validation on Six Lipid Types". In: *J. Phys. Chem. B* 114.23 (June 2010), pp. 7830–7843. ISSN: 1520-6106. DOI: 10.1021/jp101759q. URL: https://pubs.acs.org/doi/10.1021/jp101759q.

[55] William L. Jorgensen et al. "Comparison of simple potential functions for simulating liquid water". In: *J. Chem. Phys.* 79.2 (1983), pp. 926–935. ISSN: 00219606. DOI: 10.1063/1.445869. URL: http://scitation.aip.org/content/aip/journal/jcp/79/2/10.1063/1.445869%7B%5C%%7D5Cnhttp://link.aip.org/link/?JCPSA6/79/926/1.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

[56] Mark James Abraham et al. "Gromacs: High performance molecular simulations through multi-level parallelism from laptops to supercomputers". In: *SoftwareX* 1-2 (2015), pp. 19–25. DOI: 10.1016/j.softx.2015.06.001.

[57] Jumin Lee et al. "CHARMM-GUI input generator for NAMD, GROMACS, AMBER, OpenMM, and CHARMM/OpenMM simulations using the CHARMM36 additive force field". In: *Journal of chemical theory and computation* 12.1 (2016), pp. 405–413.

[58] Wolfgang Kabsch and Christian Sander. "Dictionary of protein secondary structure: Pattern recognition of hydrogen-bonded and geometrical features". In: *Biopolymers* 22.12 (Dec. 1983), pp. 2577–2637. DOI: 10.1002/bip.360221211. arXiv: 83/122577-6 [0006-3525]. URL: http://onlinelibrary.wiley.com/doi/10.1002/bip.360221211/abstract%7B%5C%%7D5Cnhttp://onlinelibrary.wiley.com/doi/10.1002/bip.360221211/full%7B%5C%%7D5Cnhttp://onlinelibrary.wiley.com/store/10.1002/bip.360221211/asset/360221211%7B%5C_%7Dftp.pdf?v=1%7B%5C&%7Dt=hot47ujz%7B%5C&%7Ds=8927fb6c2450d37cfc0a3.

Acknowledgments

This work was supported by the OPUS grant no. 2019/33/B/NZ2/00956 to ES from the National Science Centre, Poland, https://www.ncn.gov.pl/?language=en.

5 Author contributions

P.Sz., M.Mo. T.G. and E.S. conceived the project and methodology. P.Sz. curated the data. P.Sz., M.Mo. and T.G. implemented the model. P.Sz. performed the computational analysis. M.B., W.K., and D.N. performed wet lab experiments. M. Mi. and P.Se. performed computer simulations and resulting summaries of the novel peptides. J.S. and T.G. created the web service. P.Sz. and M.Mi. prepared the figures. E.S. supervised the study. P.Sz., M.Mo, and E.S. wrote the manuscript, with contributions and critical feedback from all authors.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

6 Competing interests

Other projects in Ewa Szczurek's lab are cofunded by Merck Healthcare. Otherwise we do not declare other competing interests.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Sequence	$\mathbb{P}_{M_{AMP}}(p)$	$\mathbb{P}_{M_{MIC}}(p)$	MIC *E. coli* ATCC 43927 [μg/mL]	MIC *E. coli* ATCC 25922 [μg/mL]
GIGKFLKKAKKFGKAFVKILKK-NH2 (Pexiganan)	0.99	0.99	4	4
GIGKFLKFALKKGLGLVLKFKL-NH2	**0.99**	**0.99**	**16**	**16**
GVGKKLWFALKKPLGLVKFFKLL-NH2* (Hydraganan)	**0.99**	**0.99**	**2**	**2**
GVAKKLWIAAKKPAGAGSKFKLL-NH2	0.99	0.87	512	>512
GELKKLWQAGKLSEEDGGAFKAG-NH2*	**0.07**	**<0.01**	**>512**	**>512**
FLPLIGRVLSGIL-NH2 (Temporin-A)	0.99	<0.01	256	256
FLPLIGRVFSGIL-NH2	0.99	<0.01	512	512
FLPLIGRVFSGIK-NH2	0.99	0.97	>512	>512
FLPLIGRVLSGIA-NH2	0.99	0.01	512	512
FLPLIGRVKSGIK-NH2	0.99	0.99	>512	>512
FLPIKNRYASAAE-NH2	**0.08**	**<0.01**	**>512**	**>512**

Table 1 Pexiganan and Temporin-A analogues obtained in the analogue generation process. Each row corresponds to a single peptide. $\mathbb{P}_{M_{AMP}}(p)$ is a probability of a given peptide being an AMP. $\mathbb{P}_{M_{MIC}}(p)$ is a probability of a given peptide being active. Minimal Inhibitory Concentration (MIC) values (μg/mL) were measured against reference strains of microorganisms (*E. coli* ATCC 43927, *E. coli* ATC 25922). Peptides in bold were experimentally proven to show activity in accordance with Classifier prediction. Peptide whose sequences are marked with an asterisk were subjected to molecular dynamics simulations.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

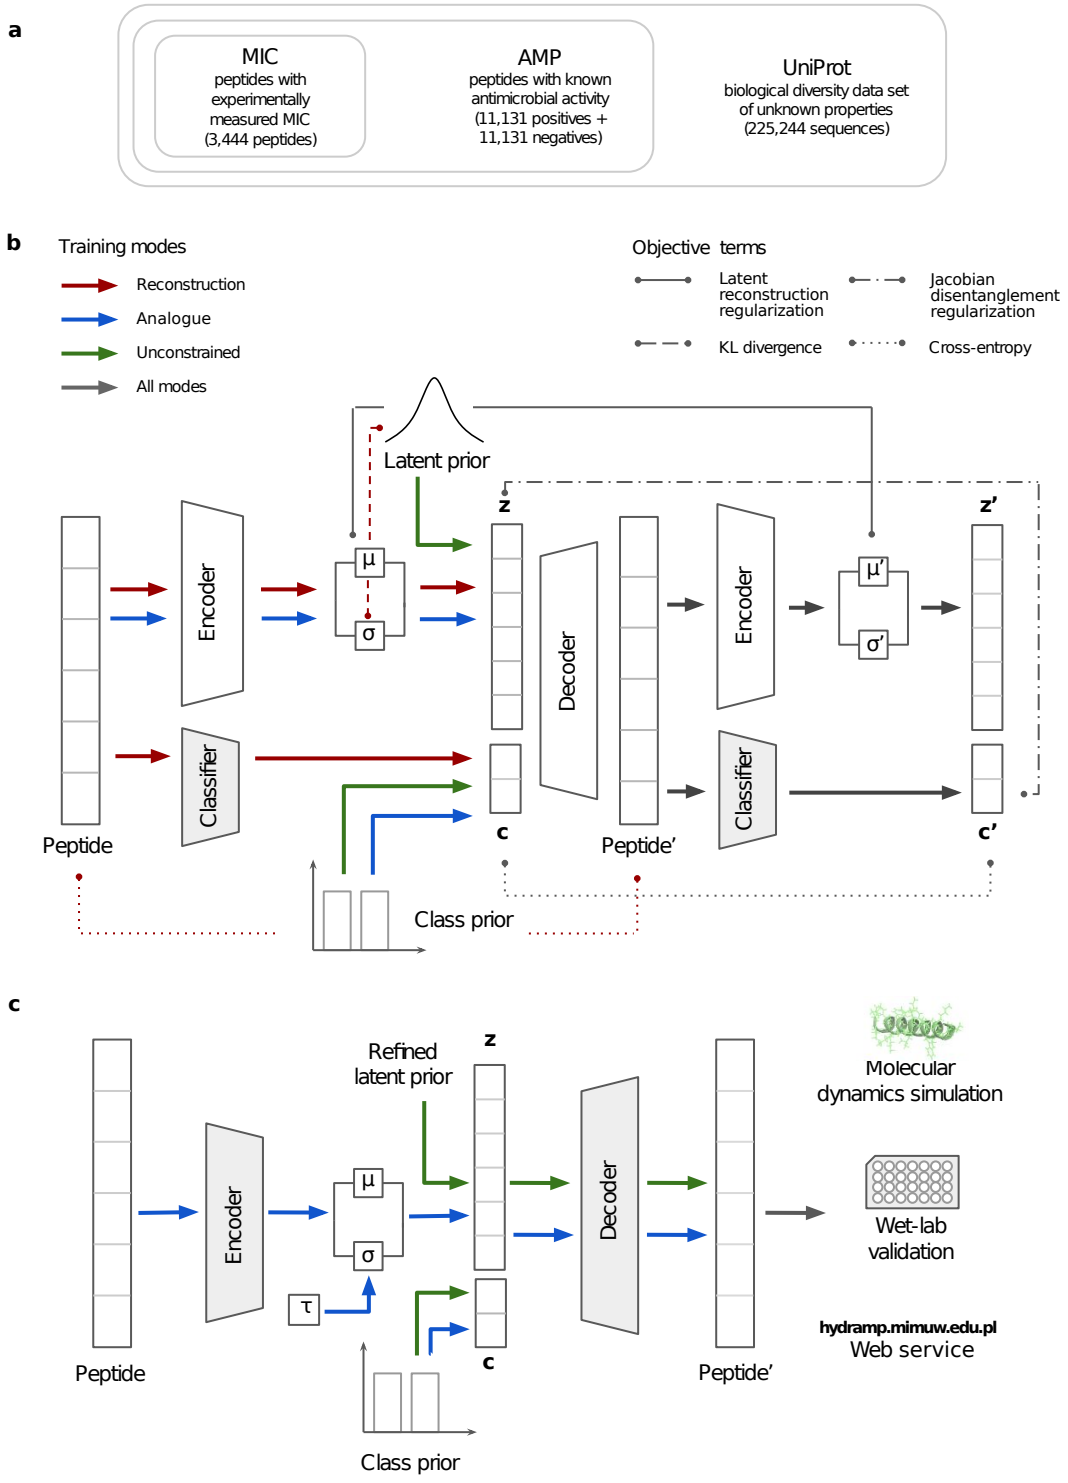

Fig. 1 HydrAMP architecture and data traversion overview.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Fig. 1 Continued caption. **a** Compositional structure of the training data set. **b** Data flow and optimization setup during training. Colors indicate training modes and show the path each peptide traverses within a mode. Line styling indicates the objective terms. Shaded areas indicate components with frozen weights. **c** Data flow during generation. In the final step molecular dynamics simulation and wet-lab validation are performed. HydrAMP functionality is available via a web service `https://hydramp.mimuw.edu.pl/`.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

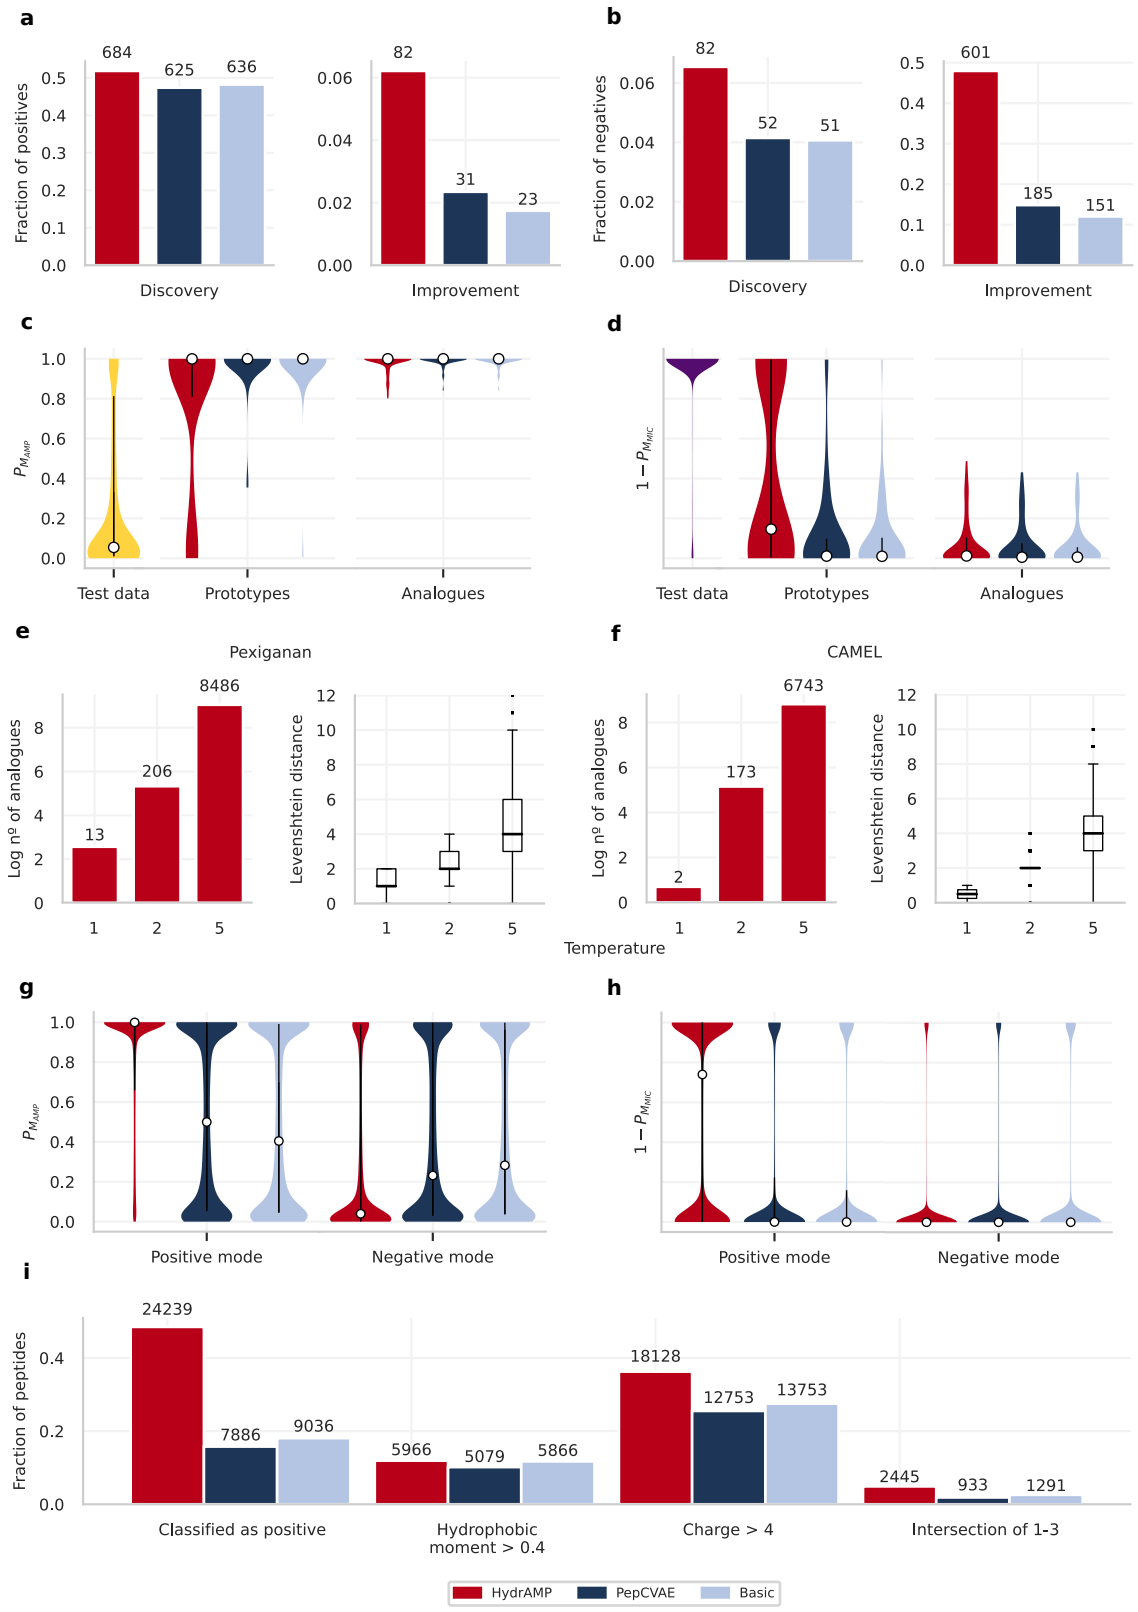

Fig. 2 Generative performance of HydrAMP (red), in comparison to PepCVAE (dark blue), and Basic (light blue).

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Fig. 2 Continued caption. **a** Fraction (y-axis) and number (over each bar) of 1319 positive (AMP and highly active) peptides from the test set (y-axis), that produced analogues that met discovery or improvement criteria (x-axis) in the analogue generation. **b** As in **a**, but for 1253 negative peptides from the test set. **c** The distribution of the probability of being AMP for 1253 negative prototypes and their analogues that met the discovery criteria, compared to the distribution for the test data (yellow). **d** The distribution of probabilities of not being highly active $(1-\mathbb{P}_{M_{MIC}})$ for 1253 negative prototypes and their analogues that met the discovery criteria, compared to the distribution for the test data (violet). **e, f** Left: The relation between the creativity parameter *temperature* (x-axis) and the log number of generated unique analogues that met the discovery criteria, out of 10k attempts (y-axis; the actual number of analogues shown above each bar). Right: the distribution of the Levenshtein distances between generated analogues and the prototype sequence of Pexiganan (**e**) and CAMEL (**f**) AMPs. **g** Probability of being antimicrobial ($\mathbb{P}_{M_{AMP}}$) for 50,000 peptides generated in unconstrained mode, for both positive and negative modes (x-axis). **h** The distribution of probabilities of not being highly active $(1-\mathbb{P}_{M_{MIC}})$ for 50,000 peptides generated in unconstrained mode for both positive and negative modes. In panels **c, d, g, h**, the white dots mark the median of each distribution. **i** Fraction of 50,000 peptides generated in the unconstrained positive mode classified as positives (first bar plot), fraction of peptides have hydrophobic moment > 0.4 (second bar plot), fraction of peptides with charge > 4 (third bar plot) and fraction of peptides that satisfy all previous criteria: classified as positive, high hydrophobic moment, and high charge (fourth bar plot). The number over each bar: the actual number of peptides with the condition.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

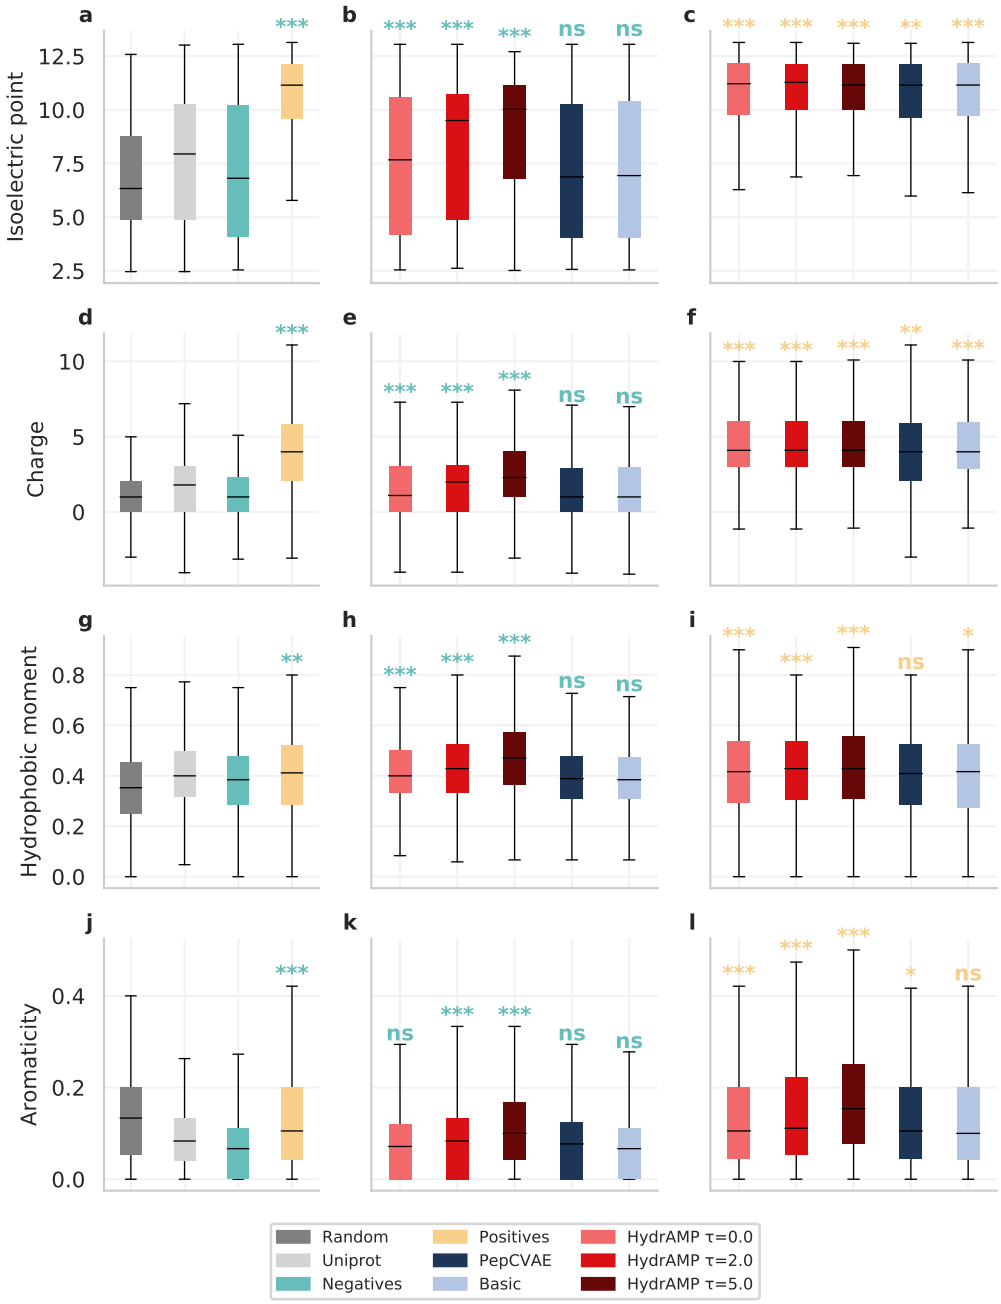

Fig. 3 Physicochemical properties of analogues generated by HydrAMP and compared methods in analogue generation mode for non-AMP or AMP templates in comparison with real and random data. Boxplots of properties (**a, b, c** Isoelectric point, **d, e, f** Charge, **g, h, i** Hydrophobic moment, **j, k, l** Aromaticity) of randomly generated peptides (dark gray), peptides sampled from UniProt (light gray), true negatives (green), and true positives (yellow), in comparison with AMP analogues generated from negatives (**b, e, h, k**), and positives (**c, f, i, l**) improved by different models: HydrAMP with various creativity parameter *temperature* values: $\tau = 1.0$ (light red), $\tau = 2.0$ (red), $\tau = 5.0$ (dark red), PepCVAE (dark blue), Basic (light blue). Significance levels are denoted as: ns - P$\geq$ 0.05; * - P$\leq$ 0.05; ** - P$\leq$ 0.01; *** - $P \leq 0.001$.

bioRxiv preprint doi: https://doi.org/10.1101/2022.01.27.478054; this version posted January 28, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

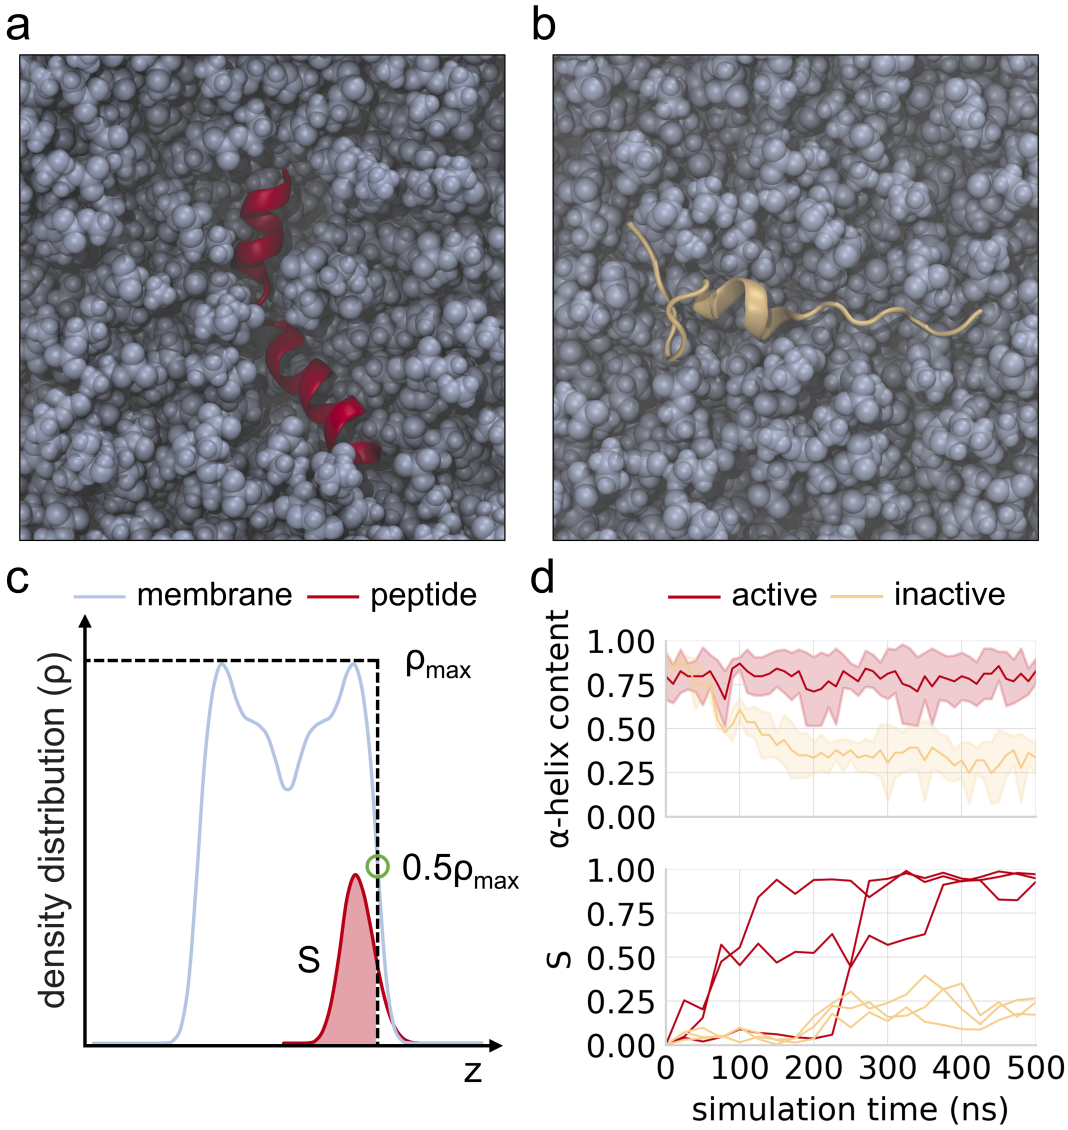

Fig. 4 Summary of atomistic molecular dynamics simulations of active antimicrobial peptide Hydraganan and experimentally verified non-AMP peptide within a model bacterial lipid membrane. a, b, late simulation snapshots of active (red) and inactive (beige) peptides, respectively, in the membrane (blue); a top view on membrane surface; water molecules not depicted for clarity; **c**, a scheme illustrating the evaluation of the *S* parameter describing the level of peptide burial within the membrane; *z* – membrane normal axis; **d**, upper plot: an average (lines) and standard deviation (shaded areas) of alpha helix fraction within peptide residues; lower plot: the evolution of S parameter for each of three initial peptide placements.