

SOFT NON-DIAGONALITY PENALTY ENABLES LATENT SPACE-LEVEL INTERPRETABILITY OF PLM AT NO PER- FORMANCE COST

Anonymous authors

Paper under double-blind review

ABSTRACT

Emergence of large scale protein language models (pLMs) has led to significant performance gains in predictive protein modeling. However, it comes at a high price of interpretability, and efforts to push representation learning towards explainable feature spaces remain scarce. The prevailing use of domain-agnostic and sparse encodings in such models fosters a perception that developing both parameter-efficient and generalizable models in a low-data regime is not feasible. In this work, we explore an alternative approach to develop compact models with interpretable embeddings while maintaining competitive performance. With the Bidirectional Long Short-Term Memory Autoencoder (BiLSTM-AE) model as an example trained on positional property matrices, we introduce a soft weight matrix non-diagonality penalty. Through Jacobian analysis, we show that this penalty aligns embeddings with the initial feature space while leading to a marginal increase in performance on a suite of four common peptide biological activity classification benchmarks. Moreover, it was demonstrated that the use of one-hot encoded sequence clustering-based contrastive loss to produce semantically meaningful latent space allows to further improve benchmarking performance. The use of amino acid physicochemical properties and density functional theory (DFT) derived cofactor interaction energies as input features provides a foundation for intrinsic interpretability, which we demonstrate on fundamental peptide properties. The resulting model is over 33,000 times more compact than the state-of-the-art pLM ProtT5. It demonstrates performance stability across diverse benchmarks without task-specific fine-tuning, showcasing that domain-tailored architectural design can yield highly parameter-efficient models with fast inference and preserved generalization capabilities.

1 INTRODUCTION

Machine learning (ML) has achieved substantial progress in drug discovery (Jorner et al., 2021; Lee et al., 2020), supported by representations such as SMILES (Weininger, 1988), which can be converted into molecular graphs and used to calculate interpretable chemical descriptors using RDKit (Landrum et al., 2013) or Mordred (Moriwaki et al., 2018). However, these representations become limiting for biopolymers due to their higher structural complexity. Peptides, a class of biopolymers composed of <100 amino acids, exhibit various conformations governed by hydrogen bonding, hydrophobic interactions, metal coordination, and disulfide bridges (Gregorc et al., 2023; Rezai et al., 2006; Victorio & Sawyer, 2023). Their larger contact surface and rich interaction patterns make them highly selective and effective therapeutic agents (Henninot et al., 2018; Fisher et al., 2019; Peterson & Barry, 2018; Torres et al., 2019). At the same time, their biological activity depends on a sequence context that is difficult to capture with small-molecule-oriented descriptors.

ML methods are widely used in peptide research (Basith et al., 2020), and peptide representations typically follow two directions: property-based and sequence-based (Xu et al., 2020). The first relies on physicochemical descriptors and provides interpretability, but does not account for positional context. The second captures sequence-level information, yet lacks explicit chemical grounding and interpretability. In practice, models built on either approach tend to be task-specific. This

054 motivates the need for representations that combine contextual modeling with chemically meaningful,
 055 interpretable structure.

056 Using a Bidirectional Long Short-Term Memory Autoencoder (BiLSTM-AE) trained on positional
 057 property matrices, we introduce a soft penalty on off-diagonal weight-matrix elements together
 058 with diagonal initialization. This encourages feature disentanglement and aligns latent dimensions
 059 with the original physicochemical feature space. Jacobian-based analysis confirms that the penalty
 060 enforces this alignment while slightly improving performance across four peptide biological activity
 061 classification benchmarks.

062 We further show that incorporating one-hot encoded sequences clustering-based contrastive loss
 063 produces more semantically structured latent spaces, leading to additional performance gains. Because
 064 the model operates directly on amino acid physicochemical properties, including interaction energies
 065 derived from density functional theory (DFT), it naturally supports intrinsic interpretability. We
 066 illustrate this using regression tasks on fundamental peptide physicochemical properties.

067 To more fairly compare embeddings across benchmarking datasets, we introduce average Shannon
 068 entropy and Levenshtein distance based measures to approximate sequence and dataset complexity,
 069 since dataset size alone does not reliably predict model performance. Despite its compactness -
 070 being >4,000, >12,000, and >6,000 times more compact than ProtBERT (Elnaggar et al., 2022),
 071 Ankh-large (Elnaggar et al., 2023), and ESM-C (600M) (ESM Team, 2024), respectively - our
 072 model achieves competitive results. It surpasses all baselines on anti-inflammatory peptide prediction
 073 (AIP), matches Ankh-large on antimicrobial peptides (AMP), performs comparably to ESM-C
 074 on anti-oxidative peptides (AOP), and outperforms ProtBERT on AOP and anti-diabetic peptide
 075 prediction (ADP). Based on Matthew's correlation coefficient (MCC), it also shows the second-
 076 highest performance stability across all tasks after Ankh-large.

077 Overall, our model is more than 33,000 times more compact than the state-of-the-art (SOTA)
 078 protein language model (pLM) ProtT5-3B (Elnaggar et al., 2022) and achieves stable performance
 079 across diverse peptide tasks without task-specific fine-tuning. This illustrates that domain-tailored
 080 architectures can yield parameter-efficient models with fast inference and strong generalization. At the
 081 same time, scaling such interpretable models remains a challenge and warrants further investigation
 082 to close the performance gap with the largest protein language models.

083 **Key contributions:**

- 085 • We introduce a soft penalty on off-diagonal weight-matrix elements combined with diagonal
 086 initialization in a BiLSTM-AE, enabling feature disentanglement and alignment of latent
 087 dimensions with amino acid physicochemical properties.
- 088 • We demonstrate latent-space interpretability through feature importance and correlation
 089 analysis in four regression tasks involving fundamental peptide physicochemical properties.
- 090 • We show that the resulting model achieves competitive or superior performance to ProtBERT,
 091 Ankh-large, and ESM-C on four peptide biological activity classification benchmarks while
 092 remaining intrinsically interpretable.
- 093 • We find that embeddings from many existing protein language models still perform com-
 094 parably to simple one-hot baselines on peptide tasks, highlighting ongoing challenges in
 095 peptide-specific representation learning with domain-agnostic large models.
- 096 • We show that the proposed non-diagonality penalty contributes to a small but consistent
 097 improvement in benchmark performance within our architecture and training setup.
- 098 • We demonstrate that our final model is 3–4 orders of magnitude more parameter-efficient
 099 than ProtBERT, ESM-C, Ankh, and ProtT5-3B while trained on datasets at least two orders
 100 of magnitude smaller and without task-specific fine-tuning.

102 **2 RELATED WORKS**

103 Current SOTA pLMs are largely based on the Transformer architecture (Vaswani et al., 2017), adapted
 104 to biological sequences, and pre-trained on massive protein corpora. Representative models such as
 105 ProtBERT and ProtT5 extend the BERT (Devlin et al., 2019) and T5 (Raffel et al., 2020) architectures
 106 to enable effective transfer learning across diverse protein tasks. The ESM family further scales

108 this paradigm, with models like ESM-Cambrian leveraging hundreds of millions of parameters
 109 to capture evolutionary and structural regularities at large scale. The Ankh work systematically
 110 explores architectural design choices: span masking, activation functions, positional encodings, to
 111 optimize Transformer-based pLMs for protein modeling. Despite architectural differences, these
 112 models share an important limitation: their interpretability relies primarily on analyzing self-attention
 113 patterns. Although attention maps can highlight residue–residue interactions or structural contacts,
 114 they are often noisy, vary widely between layers and heads, and do not reliably correspond to causal
 115 importance (Jain & Wallace, 2019).

116 Motivated by this limitation, a variety of post-hoc interpretability methods have been developed.
 117 Attribution techniques, such as attention analysis, probing, and gradient-based methods, aim to
 118 uncover sequence elements that influence model outputs but do not alter or structure the latent space
 119 itself (Hunklinger & Ferruz, 2025). Sparse autoencoders (SAEs) offer a recent alternative by training
 120 sparse decoders on pLM activations to extract biologically meaningful latent factors (Gujral et al.,
 121 2025; Simon & Zou, 2025). These methods reveal that pLM representations encode useful but highly
 122 entangled structure. As a result, both attribution-based approaches and SAEs provide useful insight
 123 but remain fundamentally post-hoc and limited in their ability to shape embedding space.

124 In contrast, intrinsically interpretable peptide representation learning remains relatively underexplored.
 125 Existing approaches often sacrifice contextual modeling capacity or rely heavily on manual feature
 126 engineering. Our work addresses this gap by introducing an inductive bias directly into a sequence
 127 model: a soft non-diagonality penalty that encourages latent dimensions to align with predefined
 128 physicochemical feature axes. This design embeds interpretability into the model itself, producing
 129 disentangled chemically grounded representations while preserving the contextual expressiveness of
 130 recurrent architectures.

131

132

3 METHODS

134

135

3.1 CONTEXT AND MOTIVATION

136

137

The proposed method constructs property matrices for peptides from physico-chemical descriptors
 138 of 20 canonical residues and 2 prevalent modifications. It uses 43 molecular features and 3 DFT-
 139 derived. Convolutional autoencoders (CAE) serve as a baseline, while recurrent and transformer
 140 variants are introduced to capture long-range dependencies. The embeddings are further refined using
 141 Information Noise-Contrastive Estimation (InfoNCE) based contrastive loss. Positive and negative
 142 pairs are generated via *MiniBatchKMeans* clustering of one-hot encoded sequences, ensuring that
 143 the interpretable property space corresponds to empirically effective sequence encodings. A soft
 144 non-diagonality penalty guides embedding features toward the original physicochemical space, with
 145 feature relationships quantified using the Jacobian matrix (see **Section 3.7**). An overview of the full
 146 processing pipeline is shown in **Figure 1**.

147

148

149

3.2 DATA COLLECTION

150

151

152

153

154

155

156

157

158

159

160

161

For model training, we retrieved unlabeled unique peptide sequences containing no extra monomers
 152 from the NCBI database (Sayers et al., 2023), yielding 6,749,334 sequences. A sampling procedure
 153 is implemented to reduce data volume and computational costs while preserving diversity: peptides
 154 are grouped by the most frequently occurring amino acids. This allowed us to maintain balanced
 155 composition over sequence length distribution. Consequently, small (**S**) and big (**B**) datasets are
 156 constructed through clustering peptide sequences based on amino acid frequency vectors, with the
 157 initial number of clusters equal to the unique amino acids. Then, the unique clusters are identified on
 158 the basis of the most represented residues for stratified sampling; the minimum number of samples is
 159 determined to prevent under-representation of smaller clusters. A small fraction (1%) of sequences
 160 forms dataset **S**; a medium fraction (10%) is sampled for dataset **B**. Finally, the subsets are combined
 161 to create the final datasets. Cluster-based stratified sampling facilitated the formation of manageable
 datasets, where the final model is trained using the dataset **S** comprising 155,920 sequences. Final
 data scaling experiment is performed on dataset **B** comprising 467,792 sequences.

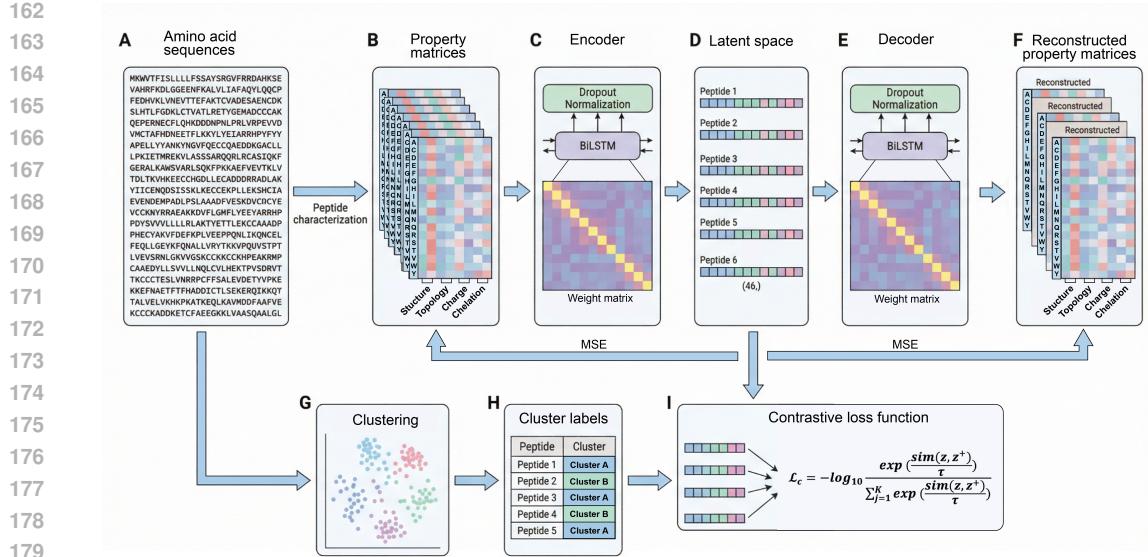


Figure 1: Overview of the proposed pipeline.

3.3 MOLECULAR DESCRIPTORS

A total of 43 molecular RDKit descriptors are used to characterize the properties of amino acids (see **Appendix A**), which were normalized to a range of [-1, 1]. Each peptide sequence is represented as a feature array based on these descriptors and DFT features (see **Section 3.4** for details), with padding applied to align all sequences to a uniform length of 96 amino acids.

3.4 DFT-DERIVED DESCRIPTORS

This group of descriptors includes three parameters derived from the interactions energies of amino acids with divalent metal cations: calcium (Ca^{2+}), magnesium (Mg^{2+}), and barium (Ba^{2+}), obtained through DFT calculations. Detailed methodology is described in Hu et al. (2022), with data available in the NOMAD database (Draxl & Scheffler, 2019). The free energy values are extracted using the NOMAD API and categorized into four groups: complexes with Mg^{2+} , Ca^{2+} , and Ba^{2+} cations and isolated amino acid conformers. For each amino acid, average free energy values across all conformers is calculated, and interaction descriptors are derived by normalizing the differences between complex and isolated free energy values.

3.5 AE-BASED ARCHITECTURES

Convolutional and Variational Autoencoders (CAE/VAEs). CAE-based models with per-feature convolutions were treated as a baseline to study the model performance when features were not mixed at all due to architectural restrictions. We then implemented a probabilistic framework by replacing the deterministic latent space of CAE with a variational latent space. This approach encourages smoother, more continuous latent representations, potentially mitigating sparsity. Experiments were conducted with multiple VAE variants, including β -VAE and InfoVAE. Hyperparameters, including the KL divergence weight (β), were chosen from a predetermined grid, this ensured an unbiased evaluation of embedding quality.

BiLSTM. Given the sequential nature of the data, BiLSTM-based architectures were used to capture long-range dependencies. Unlike unidirectional recurrent models, BiLSTM leverages both forward and backward contexts, enabling a more comprehensive understanding of sequence information. We hypothesized that this bidirectional approach would enhance the quality of the generated embeddings, making them more effective for downstream tasks.

216 **Transformer-Based Models.** Motivated by the success of transformer architectures in natural
 217 language processing, we tested a transformer-based AE. Transformers leverage self-attention mechanisms
 218 to capture global dependencies, typical in peptide sequences.

219 Each of the architectures described in this section was trained using the same dataset \mathbf{S} and evaluated
 220 on identical downstream tasks to ensure a fair comparison with the benchmark encoding strategies
 221 (see **Appendix B**).

223 3.6 CONTRASTIVE LEARNING IMPLEMENTATION

225 Contrastive learning was implemented to enhance embedding quality by drawing positive pairs closer
 226 in latent space and separating negative pairs.

227 **Positive and Negative Pair Construction.** Given the unique properties of peptide sequences, we
 228 designed a domain-specific strategy for constructing positive and negative pairs:

- 230 • Positive Pairs: Positive pairs were obtained by *MiniBatchKMeans* clustering of one-hot
 231 encoded peptides; k was selected *via* the Elbow criterion, and peptides sharing a cluster
 232 were deemed positives.
- 233 • Negative Pairs: Peptides from different clusters were treated as negatives.

235 **Contrastive Loss Function.** Utilized contrastive loss is based on the InfoNCE objective:

$$236 \quad \mathcal{L}_{\text{contrastive}} = -\log \frac{\exp(\frac{\text{sim}(\mathbf{z}, \mathbf{z}^+)}{\tau})}{\sum_{j=1}^K \exp(\frac{\text{sim}(\mathbf{z}, \mathbf{z}_j)}{\tau})}, \quad (1)$$

239 where $\mathbf{z} = E(\mathbf{x})$, $E(\cdot)$ is the encoder, $\text{sim}(\cdot, \cdot)$ denotes a similarity measure (e.g., cosine similarity),
 240 K is the number of negative samples in a batch, and $\tau > 0$ is a temperature hyperparameter.

242 3.7 SOFT NON-DIAGONALITY PENALTY

244 For the BiLSTM architecture described above, we introduce an additional regularization mechanism
 245 to reduce feature entanglement during training.

246 To achieve this, two complementary techniques are employed. First, we apply a soft penalty specifically
 247 to the off-diagonal elements of the BiLSTM weight matrices. For a weight matrix $M \in \mathbb{R}^{n \times n}$,
 248 the penalty term is defined as the mean of squared off-diagonal elements:

$$250 \quad \mathcal{L}_M = \frac{\|M \odot (\mathbf{1} - I)\|_F^2}{n \cdot n} \quad (2)$$

253 where \odot denotes the Hadamard product, I is the identity matrix, and $\mathbf{1}$ is a matrix of ones. The total
 254 penalty is the sum of all weight matrices, scaled by a coefficient λ :

$$256 \quad \mathcal{L}_{\text{off-diag}} = \lambda \sum_{M \in \mathcal{W}} \mathcal{L}_M \quad (3)$$

258 Here, \mathcal{W} comprises all *input-to-hidden* and *hidden-to-hidden* weight matrices of the BiLSTM layers.
 259 Second, we use forced diagonal weight initialization as an inductive bias. The combined effect
 260 encourages the model to maintain nearly diagonal weight transformations throughout training.

262 We additionally evaluated hard diagonalization - strict constraints enforcing zero off-diagonal weights.
 263 Although theoretically appealing, this formulation restricted the optimization to an extremely narrow
 264 feasible set, preventing stable convergence of the reconstruction and contrastive objectives. Because
 265 these difficulties occurred during optimization itself, we did not pursue the hard constraint further.
 266 The soft formulation preserves differentiability and allows the optimization to proceed while still
 267 encouraging near-diagonal structure.

268 To quantitatively evaluate the degree of feature disentanglement, we propose a diagonality metric
 269 based on the encoder’s Jacobian. We compute the mean absolute Jacobian $\bar{J} \in \mathbb{R}^{D_{in} \times D_{out}}$ over the
 dataset. The diagonality metric is then defined as:

270
271
272
273

$$\mathcal{D} = \frac{\sum_i \bar{J}_{ii}}{\sum_{i,j} \bar{J}_{ij}} \quad (4)$$

274 A value of \mathcal{D} close to 1 indicates a strong feature-wise separation.
275276 3.8 COMPUTATIONAL RESOURCES
277278 All calculations were performed on a server with a general configuration consisting of 6 A6000
279 GPUs, 256 cores, AMD EPYC 7763 64-Core Processor, 512 GB RAM. The training procedure was
280 performed using 2 GPUs and 50 GB RAM.281 4 EXPERIMENTS
282283 4.1 BENCHMARK DATASETS
284285 We performed benchmarking of all the models mentioned in the paper on four public peptide
286 biological activity classification datasets: antimicrobial (AMP) (Cao et al., 2023), anti-inflammatory
287 (AIP) (Raza et al., 2023), antidiabetic (ADP) (Chen et al., 2022), and antioxidant (AOP) (Qin et al.,
288 2023). Dataset statistics covering size, sequence complexity, and data imbalance are summarized in
289 **Table 1**. Sequence complexity is described here by several parameters including length statistics, as
290 well as average Shannon entropy characterizing the extent of non-equiprobability of amino acids per
291 sequence and average Levenshtein distance showing inter-sequence dissimilarity. Based on these
292 statistics, ADP appears to be the most challenging dataset based on dataset size and high values of
293 average Shannon entropy, Levenshtein distance, as well as length spread.
294295 Table 1: Main statistics of benchmark datasets. Average Levenshtein distance characterizes the
296 diversity of sequences by measuring the pairwise differences. Average Shannon entropy characterizes
297 the diversity of amino acids within a single sequence by quantifying the average level of uncertainty
298 associated with all types of amino acids.

Dataset	Num. of peptides	Avg. Levenshtein distance	Avg. Shannon entropy	Avg. length	Min. length	Max. length	Pos/Neg sample ratio
ADP	472	19.20 \pm 4.26	3.22 \pm 0.16	19.60 \pm 5.09	11	41	1.00
AIP	3790	15.48 \pm 2.36	3.15 \pm 0.11	16.39 \pm 2.62	11	30	0.76
AMP	8268	18.38 \pm 4.21	3.03 \pm 0.26	18.53 \pm 5.34	11	30	1.00
AOP	2120	7.12 \pm 3.58	2.03 \pm 0.48	5.92 \pm 3.66	2	20	1.00

307 4.2 EXPERIMENTAL SETUP
308309 This section presents the experimental framework for evaluating the quality and generalizability of
310 the embeddings generated by our model. The assessment was conducted on four peptide classification
311 datasets encompassing distinct biological activities (detailed in **Section 4.1**). To isolate the evaluation
312 on the embedding quality itself, rather than the complexity of a deep learning classifier, we employed
313 a classical gradient boosting model as a simple yet effective downstream predictor (Shwartz-Ziv &
314 Armon, 2022). The primary metric for comparison was the MCC, chosen for its informativeness and
315 robustness to class imbalance (Chicco & Jurman, 2020) — a property particularly relevant for datasets
316 like AIP, as evidenced by the statistics in **Table 1**. For reliable model assessment, we employed a
317 rigorous 5-fold cross-validation protocol, reporting both the mean metric values and their standard
318 errors (the latter indicated in parentheses).
319320 Our comparative analysis includes three conventional peptide encoding methods: one-hot encoding,
321 BLOSUM62, and 3-mer counts, and several SOTA pLMs: ProtBERT, Ankh-large, ESM-C (600M),
322 and ProtT5-3B. The following subsections detail the development and optimization experiments for
323 our model, with results provided in **Appendices C and F** (Tables 6 through 14). The final comparative
benchmarking results against all baselines are consolidated in **Section 4.5 (Table 2)**.

324 4.3 AE-BASED ARCHITECTURES BENCHMARKING
325326 We begin by establishing a reference using a CAE with per-feature convolutions, which serves as
327 a baseline where no feature mixing is possible due to the architectural structure. Results for all
328 encodings with the CAE model are shown in **Appendix C**.329 To perform a systematic comparison, we evaluated a pre-specified set of 80 models across four AE
330 families: CAEs, VAEs (incl. Info-VAE, β -VAE), BiLSTM-AEs, and transformer-AEs—each trained
331 under a fixed grid of hyperparameters (**Appendix B**). The embeddings were evaluated with the
332 same downstream predictor as the baseline and benchmark encodings (see **Section 4.2**). Subsequent
333 methodological developments were implemented using the BiLSTM-AE architecture due to its
334 suitability for sequential inputs and stable training dynamics (see **Table 7**).
335336 4.4 CLUSTERING-BASED CONTRASTIVE LEARNING
337338 We next incorporated a contrastive learning component (**Section 3.6**) to explicitly structure the latent
339 space by drawing similar peptides closer while pushing dissimilar ones apart. The comparison results
340 for BiLSTM-AE (cBiLSTM-AE) with the benchmark encodings appear in **Table 2**.
341342 4.5 SOFT NON-DIAGONALITY PENALTY
343344 To enforce feature disentanglement within the model, we applied a soft penalty to the off-diagonal
345 elements of the weight matrices (as detailed in **Section 3.7**). Our evaluation includes two main
346 parts: (1) a comprehensive benchmarking of cBiLSTM-AE model with soft non-diagonality penalty
347 (dcBiLSTM-AE) against all baselines (see **Section 4.2** for details), and (2) a scalability analysis where
348 the model was trained on a large-scale dataset **B**. The results of both the comparative benchmarking
349 and the scalability experiment are consolidated in **Table 2**.
350351 Table 2: dcBiLSTM-AE model benchmarking. No preliminary fine-tuning on tasks was performed
352 for any of the models.

353 Encoding 354 type	Model 355 size	Dataset for training	MCC (5-fold cross validation)				Avg.	Min	Max
			Anti 356 diabetic	Anti 357 inflammatory	Anti 358 microbial	Anti 359 oxidant			
One-hot	N/A	N/A	0.197 (0.015)	0.216 (0.015)	0.560 (0.004)	0.764 (0.013)	0.434	0.197	0.764
Blosum	N/A	N/A	0.026 (0.020)	0.290 (0.012)	0.337 (0.015)	0.189 (0.012)	0.211	0.026	0.337
Threemers	N/A	N/A	0.131 (0.060)	0.357 (0.009)	0.519 (0.003)	0.539 (0.015)	0.387	0.131	0.539
ProtBert	420M	217M	0.334 (0.031)	0.138 (0.021)	<u>0.658</u> <u>(0.007)</u>	0.580 (0.011)	0.428	0.138	0.658
ESM C	600M	-	<u>0.433</u> <u>0.017</u>	0.193 (0.012)	0.679 (0.005)	0.761 (0.019)	0.517	0.193	0.761
Ankh	1.15B	59M	0.574 (0.032)	0.335 (0.011)	0.614 (0.007)	0.863 (0.015)	0.596	<u>0.335</u>	0.863
BiLSTM-AE	<u>1.3M</u>	0.15M	-0.004 (0.053)	0.342 (0.009)	0.415 (0.005)	0.496 (0.031)	0.312	-0.004	0.496
cBiLSTM-AE	<u>1.3M</u>	0.15M	0.277 (0.065)	0.316 (0.009)	0.569 (0.002)	0.692 (0.012)	0.464	0.277	0.692
dcBiLSTM-AE	90K	0.15M	0.371 (0.039)	<u>0.355</u> <u>(0.004)</u>	0.611 (0.006)	<u>0.775</u> <u>(0.018)</u>	<u>0.528</u>	0.355	<u>0.775</u>
dcBiLSTM-AE	90K	<u>0.47M</u>	0.327 (0.042)	0.337 (0.009)	0.599 (0.009)	0.750 (0.016)	0.503	0.327	0.750
ProT5	3B	45M	0.659 (0.029)	0.402 (0.013)	0.686 (0.008)	0.893 (0.007)	0.660	0.402	0.893

374 4.6 DCBiLSTM-AE EMBEDDINGS CORRELATION ANALYSIS
375376 To assess representational fidelity, we computed four physicochemical properties on the dataset **S** test
377 split (see **Section 3.2**). Targets were chosen to be analytically computable for interpretation clarity,

378 yet being functionally relevant: instability index (ISI), theoretical net charge (TNC), isoelectric point
 379 (IEP), and molecular weight (MW). Correlating these clear whole-peptide properties with interpretable
 380 feature space reveals the extent to which the embeddings capture monomer-level physicochemical
 381 structure. (**Table 3**). The same protocol was applied to earlier encodings (**Table 14**).
 382

383 5 DISCUSSION

385 5.1 MODEL SCREENING AND OPTIMIZATION

387 We initially benchmarked common peptide encoding strategies along with the baseline CAE model on
 388 four datasets (**Table 1**) to characterize task complexity and establish baseline performance. AMP and
 389 AOP datasets were generally easier for the models to classify (MCC = 0.499 and 0.503, respectively;
 390 **Table 6**), while ADP and AIP proved more challenging (MCC = 0.165 and 0.267). This difficulty
 391 correlates with high Shannon entropies and large inter-sequence Levenshtein distances, combined
 392 with relatively small dataset sizes. The baseline CAE achieved moderate performance (**Table 6**):
 393 fourth out of five in AMP and AOP, third in ADP, and second in AIP. Notably, CAE performed
 394 relatively better on the more challenging datasets, achieving comparable rankings to ProtBERT.

395 To improve embedding quality, we evaluated architectures designed for sequential data, including
 396 BiLSTM and transformer-based autoencoders. Additionally, we tested the InfoCVAE model (Zhao
 397 et al., 2019), which addresses potential latent space sparsity and mitigates the tendency of VAEs to
 398 ignore latent variables when using flexible decoders. The results of the AE-based models screening
 399 stage are summarized in **Table 7**. CAE and BiLSTM-AE showed comparable performance, while the
 400 transformer-based AE was less stable during training. The BiLSTM-AE was selected for subsequent
 401 experiments due to its architectural suitability for sequential inputs and stable training dynamics.

402 We incorporated contrastive learning into the training process (details in **Section 3.6**) to further
 403 structure the latent space, using one-hot encoded sequences as an initial reference. As demonstrated
 404 in **Table 2**, contrastive learning increased the model’s predictive metrics compared with both the
 405 one-hot baseline and the original BiLSTM-AE.

406 To mitigate feature entanglement, we implemented a feature disentanglement procedure (**Section**
 407 **3.7**), combining forced diagonal initialization with a soft penalty. Although this structural constraint
 408 could potentially reduce predictive capacity, the results in **Table 2** indicate that performance was
 409 slightly improved. This suggests that the imposed structural bias acts as a useful regularizer, effec-
 410 tively guiding the learning process. This procedure also enabled interpretability of the latent space,
 411 quantified by a diagonality metric of 49.4% for the dcBiLSTM-AE model (**Section 3.7**).

412 These results demonstrate that our model consistently delivers stable and robust performance across
 413 the evaluated cases including ADPs and AIPs being the most problematic for all encoding strategies
 414 under the study, while showing latent space-level interpretability. For instance, although dominance
 415 in raw performance is not the main aim of this research, in two out of four tasks our method achieved
 416 higher MCC than the ProtBERT model. This result is particularly remarkable given that our model
 417 has no preliminary fine-tuning step (thereby, excluding a trivial explanation that the model performs
 418 well due to overfitting) and operates with $>4,000$ times fewer trainable parameters. This drastic
 419 reduction in model complexity underscores the efficiency, while still achieving competitive and
 420 reliable performance across diverse tasks. Detailed discussion of the model’s limitations and potential
 421 improvements is provided in **Appendix D**.

422 5.2 EMBEDDING SPACE INTERPRETABILITY STUDIES

424 To evaluate the interpretability of the obtained peptide embeddings, we conducted a regression
 425 analysis aimed at predicting a set of basic peptide properties (see **Section 4.5**). Additionally, we
 426 assessed the relationships between individual embedding elements and several biologically significant
 427 physicochemical properties of peptides using Pearson’s correlation tests. The embeddings generated
 428 by our model were used as descriptors for an extreme gradient boosting regression model, which
 429 predicted the target properties based on these embeddings. The regression analysis demonstrated
 430 that the embeddings effectively captured relevant information about the peptide properties, where
 431 the results for the final dcBiLSTM-AE model were comparable with the other approaches. Detailed
 432 results of this analysis are provided in **Table 9 (Appendix C)**. To further explore the interpretability

432 of the embeddings, we examined the correlations between individual latent space features as well
 433 as the physicochemical properties. For each property, we identified the top 10 embedding elements
 434 with the highest absolute Pearson correlation coefficients (PCC). Among these, we highlighted the
 435 features with the most statistically significant correlations based on their p-values.

436 Overall, the analysis revealed that certain features consistently demonstrated strong correlations
 437 with specific physicochemical properties (**Table 3**; for a complete list see **Appendix F**), which
 438 can be interpreted as indirect yet a strong sign that the latent space features provide a substantial
 439 extent of interpretability. For instance, results for ISI show the strongest correlation of number of
 440 aromatic rings, saturated heterocycles and spiro atoms with the target value. Each of these parameters
 441 characterize connectivity within the molecules, therefore being related to branching, which directly
 442 influences peptide conformational flexibility and, therefore, the stability. Also, in case of TNC
 443 and IEP directly related to peptide charge, model shows that partition coefficient logP and Ca^{2+}
 444 interaction energy are highly relevant, where the former is a parameter describing protein solubility
 445 and the latter is strongly related to the presence and alignment of negatively charged amino acids in
 446 the protein. It is important to note, however, that not all the features have a clear and well-studied
 447 relationship with the properties predicted in these tasks, while for some of the tasks (e.g., MW), most
 448 of the parameters are directly related to the target value, and the model was unable to differentiate
 449 between them. Overall, these findings suggest that our model is able to capture meaningful features
 450 related to peptide properties, providing a foundation for interpreting its outputs in a biologically
 451 informed manner.

452 Table 3: Correlation between embedding dimensions and peptide properties.

454 Peptide property	455 Encoded input feature	456 PCC (r)
456 ISI	457 Num Arom Rings	-0.253
	458 Num Saturated Heterocycles	0.202
	459 Num Spiro Atoms	-0.193
460 TNC	461 CrippenClogP	-0.738
	462 Ca^{2+} interaction energy	0.542
	463 Num Atom Stereo Centers	-0.537
464 IEP	465 CrippenClogP	-0.620
	466 Ca^{2+} interaction energy	0.494
	467 Num Atom Stereo Centers	-0.392
468 MW	469 chi1v	-0.994
	470 chi2v	0.992
	471 kappa3	0.989

470 6 CONCLUSION

471 In this work, we propose a novel approach towards domain-tailored development of 3-4 orders more
 472 parameter-efficient pLMs trained on much smaller yet representative peptide datasets combining
 473 cross-task stability requiring no per-task fine-tuning and feature space-level interpretability for peptide
 474 representation learning. Although model scaling requires further research to compete with or surpass
 475 current SOTA ProtT5 model with input-level interpretability in raw performance, these findings
 476 hold great potential by introducing explainable latent spaces crucial for domain scientists building
 477 their own predictive models on specific downstream tasks, which is not provided by any of the pLM
 478 models existing to date, at no performance cost.

479 480 REPRODUCIBILITY STATEMENT

481 The code and datasets used in this study are available to ensure reproducibility of the presented results.
 482 They can be accessed at the following repository: <https://anonymous.4open.science/r/SeQuant-perfomance-1822/>.

486 REFERENCES
487

488 Shaherin Basith, Balachandran Manavalan, Tae Hwan Shin, and Gwang Lee. Machine intelligence
489 in peptide therapeutics: A next-generation tool for rapid disease screening. *Medicinal research*
490 *reviews*, 40(4):1276–1314, 2020.

491 Qiushi Cao, Cheng Ge, Xuejie Wang, Peta J Harvey, Zixuan Zhang, Yuan Ma, Xianghong Wang,
492 Xinying Jia, Mehdi Mobli, David J Craik, et al. Designing antimicrobial peptides using deep
493 learning and molecular dynamic simulations. *Briefings in Bioinformatics*, 24(2):bbad058, 2023.

494 Xue Chen, Jian Huang, and Bifang He. Antidmppred: a web service for identifying anti-diabetic
495 peptides. *PeerJ*, 10:e13581, 2022.

496 Davide Chicco and Giuseppe Jurman. The advantages of the matthews correlation coefficient (mcc)
497 over f1 score and accuracy in binary classification evaluation. *BMC genomics*, 21(1):6, 2020.

498 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
500 bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of*
501 *the North American chapter of the association for computational linguistics: human language*
502 *technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.

503 Claudia Draxl and Matthias Scheffler. The nomad laboratory: from data sharing to artificial intelli-
504 gence. *Journal of Physics: Materials*, 2(3):036001, 2019.

505 Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom
506 Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, Debsindhu Bhowmik, and Burkhard
507 Rost. Prottrans: Toward understanding the language of life through self-supervised learning.
508 *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):7112–7127, 2022. doi:
509 10.1109/TPAMI.2021.3095381.

510 Ahmed Elnaggar, Hazem Essam, Wafaa Salah-Eldin, Walid Moustafa, Mohamed Elkerdawy, Char-
511 lotte Rochereau, and Burkhard Rost. Ankh: Optimized protein language model unlocks general-
512 purpose modelling. *arXiv preprint arXiv:2301.06568*, 2023.

513 ESM Team. Esm cambrian: Revealing the mysteries of proteins with unsupervised learning, 2024.
514 URL <https://evolutionaryscale.ai/blog/esm-cambrian>.

515 Elizaveta Fisher, Kirill Pavlenko, Alexander Vlasov, and Galina Ramenskaya. Peptide-based thera-
516 peutics for oncology. *Pharmaceutical medicine*, 33(1):9–20, 2019.

517 Jure Gregorc, Nathalie Lensen, Grégory Chaume, Jernej Iskra, and Thierry Brigaud. Trifluo-
518 romethylthiolation of tryptophan and tyrosine derivatives: A tool for enhancing the local hydropho-
519 bicity of peptides. *The Journal of Organic Chemistry*, 88(18):13169–13177, 2023.

520 Onkar Gujral, Mihir Bafna, Eric Alm, and Bonnie Berger. Sparse autoencoders uncover biologically
521 interpretable features in protein language model representations. *Proceedings of the National*
522 *Academy of Sciences*, 122(34):e2506316122, 2025. doi: 10.1073/pnas.2506316122. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2506316122>.

523 Antoine Henninot, James C Collins, and John M Nuss. The current state of peptide drug discovery:
524 back to the future? *Journal of medicinal chemistry*, 61(4):1382–1414, 2018.

525 Xiaojuan Hu, Maja-Olivia Lenz-Himmer, and Carsten Baldauf. Better force fields start with better
526 data: A data set of cation dipeptide interactions. *Scientific Data*, 9(1):327, 2022.

527 Andrea Hunklinger and Noelia Ferruz. Toward the explainability of protein language models, 2025.
528 URL <https://arxiv.org/abs/2506.19532>.

529 Sarthak Jain and Byron C Wallace. Attention is not explanation. *arXiv preprint arXiv:1902.10186*,
530 2019.

531 Kjell Jorner, Anna Tomberg, Christoph Bauer, Christian Sköld, and Per-Ola Norrby. Organic reactivity
532 from mechanism to machine learning. *Nature Reviews Chemistry*, 5(4):240–255, 2021.

540 G Landrum et al. Rdkit: cheminformatics and machine learning software. *RDKIT. ORG*, 232, 2013.
 541

542 Byungju Lee, Jaekyun Yoo, and Kisuk Kang. Predicting the chemical reactivity of organic materials
 543 using a machine-learning approach. *Chemical science*, 11(30):7813–7822, 2020.

544 Hirotomo Moriwaki, Yu-Shi Tian, Norihito Kawashita, and Tatsuya Takagi. Mordred: a molecular
 545 descriptor calculator. *Journal of cheminformatics*, 10(1):4, 2018.

546

547 Shaylee C Peterson and Arden R Barry. Effect of glucagon-like peptide-1 receptor agonists on
 548 all-cause mortality and cardiovascular outcomes: a meta-analysis. *Current Diabetes Reviews*, 14
 549 (3):273–279, 2018.

550 Dongya Qin, Linna Jiao, Ruihong Wang, Yi Zhao, Youjin Hao, and Guizhao Liang. Prediction of
 551 antioxidant peptides using a quantitative structure- activity relationship predictor (anoxpp) based
 552 on bidirectional long short-term memory neural network and interpretable amino acid descriptors.
 553 *Computers in Biology and Medicine*, 154:106591, 2023.

554 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi
 555 Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text
 556 transformer. *Journal of machine learning research*, 21(140):1–67, 2020.

557

558 Ali Raza, Jamal Uddin, Abdullah Almuhaimeed, Shahid Akbar, Quan Zou, and Ashfaq Ahmad. Aips-
 559 sntcn: Predicting anti-inflammatory peptides using fasttext and transformer encoder-based hybrid
 560 word embedding with self-normalized temporal convolutional networks. *Journal of chemical
 561 information and modeling*, 63(21):6537–6554, 2023.

562 Taha Rezai, Jonathan E Bock, Mai V Zhou, Chakrapani Kalyanaraman, R Scott Lokey, and Matthew P
 563 Jacobson. Conformational flexibility, internal hydrogen bonding, and passive membrane perme-
 564 ability: successful in silico prediction of the relative permeabilities of cyclic peptides. *Journal of
 565 the American Chemical Society*, 128(43):14073–14080, 2006.

566

567 Eric W Sayers, Jeff Beck, Evan E Bolton, J Rodney Brister, Jessica Chan, Donald C Comeau, Ryan
 568 Connor, Michael DiCuccio, Catherine M Farrell, Michael Feldgarden, et al. Database resources of
 569 the national center for biotechnology information. *Nucleic acids research*, 52(D1):D33, 2023.

570 Ravid Shwartz-Ziv and Amitai Armon. Tabular data: Deep learning is not all you need. *Information
 571 Fusion*, 81:84–90, 2022.

572 Elana Simon and James Zou. Interplm: Discovering interpretable features in protein language models
 573 via sparse autoencoders. *Nature Methods*, pp. 1–11, 2025.

574

575 Marcelo DT Torres, Shanmugapriya Sothiselvam, Timothy K Lu, and Cesar de la Fuente-Nunez.
 576 Peptide design principles for antimicrobial applications. *Journal of molecular biology*, 431(18):
 577 3547–3567, 2019.

578 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 579 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing
 580 systems*, 30, 2017.

581 Clara G Victorio and Nicholas Sawyer. Folding-assisted peptide disulfide formation and dimerization.
 582 *ACS Chemical Biology*, 18(7):1480–1486, 2023.

583

584 David Weininger. Smiles, a chemical language and information system. 1. introduction to methodol-
 585 ogy and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36,
 586 1988.

587 Yuting Xu, Deeptak Verma, Robert P Sheridan, Andy Liaw, Junshui Ma, Nicholas M Marshall, John
 588 McIntosh, Edward C Sherer, Vladimir Svetnik, and Jennifer M Johnston. Deep dive into machine
 589 learning models for protein engineering. *Journal of chemical information and modeling*, 60(6):
 590 2773–2790, 2020.

591

592 Shengjia Zhao, Jiaming Song, and Stefano Ermon. Infovae: Balancing learning and inference
 593 in variational autoencoders. In *Proceedings of the aaai conference on artificial intelligence*,
 volume 33, pp. 5885–5892, 2019.

594
595
596
597
598
599 **A APPENDIX A**
600599
600 Table 4: Full list of chemical descriptors used for monomer-wise peptide sequence description.
601

602 Property	603 Descriptors
604 Mass	605 exactmw, amw
606 Lipophilicity and Solubility	607 CrippenClogP, CrippenMR
608 Hydrogen bonds	609 lipinskiHBA, lipinskiHBD, NumRotatableBonds, NumHBD, 610 NumHBA
611 Structural descriptors	612 NumHeavyAtoms, NumAtoms, NumHeteroatoms, Nu- 613 mAmideBonds, FractionCSP3
614 Ring Structures	615 NumRings, NumAromaticRings, NumAliphaticRings, NumSaturat- 616 edRings, NumHeterocycles, NumAromatic Heterocycles, NumSat- 617 uratedHeterocycles, NumAliphatic Heterocycles, NumSpiroAtoms, 618 NumBridgeheadAtoms
619 Topological indices	620 chi0v, chi1v, chi2v, chi3v, chi4v, chi0n, chi1n, chi2n, chi3n, chi4n, 621 hallKierAlpha, kappa1, kappa2, kappa3
622 Polar surface	623 tpsa
624 Surface area	625 labuteASA
626 Energy	627 Ca ²⁺ Interaction energy, Mg ²⁺ Interaction energy, Ba ²⁺ Interaction 628 energy
629 Other descriptors	630 Phi

631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

648
649 B APPENDIX B650
651 Table 5: Hyperparameter configurations for experimental model variants. 1-30 CAE; 31-49 VAE; 50
652 Transformer; 51 GRU; 52 LSTM; 53 BiLSTM; 54 BiLSTM + attention; 55-80 cBiLSTM. Horizontal
653 lines stands for different experimental stages, details are provided in Comment column.

654 655 656 657 658 659 660 661 662 663 664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680 681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697 698 699 700 701 Architecture	654 655 656 657 658 659 660 661 662 663 664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680 681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697 698 699 700 701 Trainable parameters	654 655 656 657 658 659 660 661 662 663 664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680 681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697 698 699 700 701 Hyperparameters	654 655 656 657 658 659 660 661 662 663 664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680 681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697 698 699 700 701 Comment
1	511943	batch_size = 64 filters = 46 pool_size = 3-2-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Number of filters is constant
2	511943	batch_size = 64 filters = 46 pool_size = 2-2-2-2-2-3 depth (number of layer blocks) = 6 loss_function = MSE	Number of filters is constant
3	61536	batch_size = 64 filters = 46 pool_size = 4-3-2-2-2 depth (number of layer blocks) = 5 loss_function = MSE	Number of filters is constant
4	61536	batch_size = 64 filters = 46 pool_size = 2-2-2-3-4 depth (number of layer blocks) = 5 loss_function = MSE	Number of filters is constant
5	42500	batch_size = 64 filters = 46 pool_size = 8-3-2-2 depth (number of layer blocks) = 4 loss_function = MSE	Number of filters is constant
6	42500	batch_size = 64 filters = 46 pool_size = 4 depth (number of layer blocks) = 2-2-3-8 loss_function = MSE	Number of filters is constant
7	735000	batch_size = 64 filters = 368 pool_size = 3-2-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Filters decrease: 368-184-92-46-23-1
8	1809052	batch_size = 64 filters = 368 pool_size = 3-2-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Filters decrease: 368-368-184-92-46-1. Optimal architecture in the first stage

702	9	461482	batch_size = 64 filters = 184 pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Filters decrease: 184-184-92-46-23-1
703	10	257000	batch_size = 64 filters = 184 pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Filters decrease: 184-92-92-46-23-1
704	11	148000	batch_size = 64 filters = 92 pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Filters decrease: 92-92-46-46-23-1
705	12	205000	batch_size = 64 filters = 184 pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Filters decrease: 184-92-46-23-11-1
706	13	2901000	batch_size = 64 filters = 736 pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Filters decrease: 736-368-184-92-46-1
707	14	133 463 996	batch_size = 64 filters = 368 pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Add attention layers to encoder
708	15	265118940	batch_size = 64 filters = 368 pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Add attention layers to encoder and decoder
709	16	1809052	batch_size = 48 filters = 368 pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Decrease batch size
710	17	1809052	batch_size = 32 filters = 368 pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Decrease batch size
711	18	1809052	batch_size = 24 filters = 368 pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Decrease batch size
712				
713				
714				
715				
716				
717				
718				
719				
720				
721				
722				
723				
724				
725				
726				
727				
728				
729				
730				
731				
732				
733				
734				
735				
736				
737				
738				
739				
740				
741				
742				
743				
744				
745				
746				
747				
748				
749				
750				
751				
752				
753				
754				
755				

756	19	1809052	batch_size = 16 filters = 368 pool_size = 3-2-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Decrease batch size
757	20	4337672	batch_size = 32 filters = 368 pool_size = 3-2-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Additional Conv2D layers in each block
758	21	8982890	batch_size = 32 filters = 460 pool_size = 3-2-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Increase number of filters
759	22	94469672	batch_size = 32 filters = 368 pool_size = 3-2-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Add 7 Dense layers to latent space
760	23	62295982	batch_size = 32 filters = 368 pool_size = 3-2-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Add 2 Dense layers to latent space
761	24	30 122 292	batch_size = 32 filters = 368 pool_size = 3-2-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Add 1 Dense layer to latent space
762	25	4685330	batch_size = 32 filters = 46 pool_size = 96 depth (number of layer blocks) = 1 loss_function = MSE	
763	26	4702534	batch_size = 32 filters = 46 pool_size = 48-2 depth (number of layer blocks) = 2 loss_function = MSE	
764	27	4719738	batch_size = 32 filters = 46 pool_size = 24-2-2 depth (number of layer blocks) = 3 loss_function = MSE	
765	28	4736942	batch_size = 32 filters = 46 pool_size = 12-2-2-2 depth (number of layer blocks) = 4 loss_function = MSE	

810	29	4754146	batch_size = 32 filters = 46 pool_size = 6-2-2-2-2 depth (number of layer blocks) = 5 loss_function = MSE	
811	30	178490	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE	Dependence of the number of filters on the depth - Baseline CAE
812	31	176328	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + KL divergence	Add KL divergence
813	32	176328	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + Wasserstein distance	Wasserstein distance
814	33	176328	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + Shannon divergence	Shannon divergence
815	34	176328	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + KL divergence	
816	35	176328	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + Gaussian noise	Add Gaussian noise
817	36	176328	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + 100 * KL divergence	
818	37	176328	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + 1000 * KL divergence	
819				
820				
821				
822				
823				
824				
825				
826				
827				
828				
829				
830				
831				
832				
833				
834				
835				
836				
837				
838				
839				
840				
841				
842				
843				
844				
845				
846				
847				
848				
849				
850				
851				
852				
853				
854				
855				
856				
857				
858				
859				
860				
861				
862				
863				

864			
865	38	483010	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2-1-1 depth (number of layer blocks) = 8 loss_function = MSE + KL divergence
866			
867			
868			
869			
870	39	1027236	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2-1-1-1-1 depth (number of layer blocks) = 10 loss_function = MSE + KL divergence
871			
872			
873			
874			
875	40	1876718	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2-1-1-1-1-1 depth (number of layer blocks) = 12 loss_function = MSE + KL divergence
876			
877			
878			
879			
880	41	2028048	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2-1-1 depth (number of layer blocks) = 8 loss_function = MSE + KL divergence
881			
882			
883			
884			
885			
886	42	3086554	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2-1-1-1-1 depth (number of layer blocks) = 10 loss_function = MSE + KL divergence
887			
888			
889			
890			
891	43	4 165 944	batch_size = 32 filters = filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2-1-1-1-1-1 depth (number of layer blocks) = 12 loss_function = MSE + KL divergence
892			
893			
894			
895			
896	44	176328	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + KL divergence
897			Change layers order
898			
899			
900			
901	45	9681446	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 20 loss_function = MSE + KL divergence
902			
903			
904			
905			
906	46	176328	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + KL divergence
907			Remove tf.Data.Dataset from data preprocessing
908			
909			
910			
911			
912	47	48882	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + KL divergence
913			convolution kernel = 2
914			
915			
916			
917			

918	48	516184	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + KL divergence	convolution kernel = 6
919	49	176328	batch_size = 32 filters = height * ((depth - i) / 2) pool_size = 3-2-2-2-2 depth (number of layer blocks) = 6 loss_function = MSE + KL divergence + MMD	InfoVAE loss function
920	50	388022	batch_size = 32 depth (number of layer blocks) = 3 feed forward dim = 128 embeddings dim = 46 loss_function = MSE	Transformer-based model
921	51	515676	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE gru_units = 256	GRU-based model
922	52	295452	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE lstm_units = 160	LSTM-based model
923	53	256220	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE bilstm_units = 96	BiLSTM-based model
924	54	696038	batch_size = 32 depth (number of layer blocks) = 3 feed forward dim = 128 embeddings dim = 46 loss_function = MSE	BiLSTM-based model with attention layers
925	55	1293788	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Add contrastive loss function Clustering on one-hot with PCA data (2D)
926	56	1293788	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE + 0.1 * InfoNCE bilstm_units = 96	Add decreasing coefficient for contrastive loss Clustering on one-hot with PCA data (2D)
927	57	1293788	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Clustering on one-hot data (without dimensionality reduction)

972	58	1293788	batch_size = 48 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size
973	59	1293788	batch_size = 64 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size
974	60	1293788	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Dropout rate = 0.2
975	61	1293788	batch_size = 64 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size Dropout rate = 0.2
976	62	1293788	batch_size = 16 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Decrease batch size
977	63	1293788	batch_size = 8 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Decrease batch size
978	64	1293788	batch_size = 80 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size
979	65	1293788	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Dropout rate = 0.15
980	66	1293788	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Dropout rate = 0.05
981	67	1293788	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Dropout rate = 0.25
982	68	1293788	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Dropout rate = 0.30
983	69	1293788	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Dropout rate = 0.35
984				
985				
986				
987				
988				
989				
990				
991				
992				
993				
994				
995				
996				
997				
998				
999				
1000				
1001				
1002				
1003				
1004				
1005				
1006				
1007				
1008				
1009				
1010				
1011				
1012				
1013				
1014				
1015				
1016				
1017				
1018				
1019				
1020				
1021				
1022				
1023				
1024				
1025				

1026				
1027	70	1293788	batch_size = 48 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size Dropout rate = 0.05
1028				
1029				
1030				
1031	71	1293788	batch_size = 48 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size Dropout rate = 0.05 Learning rate = 0.0005
1032				
1033				
1034				
1035				
1036	72	1293788	batch_size = 320 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size
1037				
1038				
1039				
1040	73	1293788	batch_size = 320 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size Dropout rate = 0.05
1041				
1042				
1043				
1044	74	1293788	batch_size = 480 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size
1045				
1046				
1047				
1048				
1049	75	1293788	batch_size = 320 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size Dropout rate = 0.15
1050				
1051				
1052				
1053	76	1293788	batch_size = 320 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size Dropout rate = 0.20
1054				
1055				
1056				
1057	77	1293788	batch_size = 320 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size Dropout rate = 0.25
1058				
1059				
1060				
1061				
1062	78	1293788	batch_size = 320 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Increase batch size Dropout rate = 0.30
1063				
1064				
1065				
1066	79	1293788	batch_size = 320 depth (number of layer blocks) = 6 loss_function = MSE + InfoNCE bilstm_units = 96	Increase depth
1067				
1068				
1069				
1070	80	1293788	batch_size = 32 depth (number of layer blocks) = 3 loss_function = MSE + InfoNCE bilstm_units = 96	Exponential decrease of learning rate (decay rate = 0.9)
1071				
1072				
1073				
1074				
1075				
1076				
1077				
1078				
1079				

C APPENDIX C

1082 Appendix C provides additional experiments' results.
 1083

1084 Table 6: CAE model benchmarking against existing peptide encoding strategies.
 1085

1086 1087 1088 1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1100 1101 1102 1103 1104 1105 1106 1107 1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133	1086 1087 1088 1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1100 1101 1102 1103 1104 1105 1106 1107 1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133						
	ADP	AIP	AMP	AOP	Avg.	Min	Max
One-hot	0.197 (0.015)	0.216 (0.015)	0.560 (0.004)	0.764 (0.013)	0.434	0.197	0.764
	0.026 (0.020)	0.290 (0.012)	0.337 (0.015)	0.189 (0.012)	0.211	0.026	0.337
Blosum	0.131 (0.060)	0.357 (0.009)	0.519 (0.003)	0.539 (0.015)	0.387	0.131	0.539
	0.334 (0.031)	0.138 (0.021)	0.658 (0.007)	0.580 (0.011)	0.428	0.138	0.658
Threemers	0.137 (0.040)	<u>0.336</u> (0.023)	0.421 (0.014)	<u>0.442</u> (0.016)	0.334	0.137	0.442
MCC (5-fold cross validation)							
MCC performance per task							
Avg.	0.165	0.267	<u>0.499</u>	0.503	-	-	-
Min	0.026	0.138	0.337	<u>0.189</u>	-	-	-
Max	0.334	0.357	<u>0.658</u>	0.764	-	-	-

Table 7: AE-based architectures benchmarking. Results are presented for best-in-class models.

1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133	1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133						
	ADP	AIP	AMP	AOP	Avg.	Min	Max
CAE	0.137 (0.040)	<u>0.336</u> (0.023)	<u>0.421</u> (0.014)	0.442 (0.016)	0.334	0.137	0.442
	0.035 (0.033)	<u>0.019</u> (0.009)	<u>0.025</u> (0.008)	-0.024 (0.028)	0.014	-0.024	0.035
InfoCVAE	<u>0.059</u> (0.024)	0.122 (0.009)	0.465 (0.012)	<u>0.448</u> (0.019)	0.273	<u>0.059</u>	<u>0.465</u>
	-0.004 (0.053)	0.342 (0.009)	0.415 (0.005)	0.496 (0.031)	<u>0.312</u>	-0.004	0.496

Table 8: Classification metrics for four benchmark tasks.

Peptide property	Encoding type	Accuracy	Precision	Recall	F1 score	ROC AUC	MCC
ADP	One-hot	0.598 (0.008)	0.596 (0.012)	0.622 (0.033)	0.606 (0.009)	0.597 (0.008)	0.197 (0.015)
	Threemers	0.559 (0.030)	0.559 (0.031)	0.559 (0.041)	0.558 (0.034)	0.559 (0.030)	0.119 (0.060)
	Blosum	0.513 (0.010)	0.511 (0.008)	0.560 (0.032)	0.533 (0.018)	0.513 (0.010)	0.026 (0.020)
	ProtBERT	0.665 (0.016)	0.663 (0.020)	0.682 (0.038)	0.669 (0.019)	0.665 (0.016)	0.334 (0.031)
	cBiLSTM-AE	0.638 (0.032)	0.638 (0.034)	0.652 (0.039)	0.643 (0.031)	0.638 (0.032)	0.277 (0.065)
AIP	One-hot	0.656 (0.005)	0.608 (0.013)	0.284 (0.015)	0.386 (0.016)	0.585 (0.007)	0.216 (0.015)
	Threemers	0.704 (0.004)	0.637 (0.006)	0.531 (0.014)	0.579 (0.009)	0.671 (0.005)	0.358 (0.009)
	Blosum	0.680 (0.005)	0.622 (0.013)	0.420 (0.008)	0.501 (0.007)	0.631 (0.005)	0.290 (0.012)
	ProtBERT	0.627 (0.008)	0.527 (0.022)	0.245 (0.012)	0.335 (0.015)	0.554 (0.008)	0.138 (0.021)
	cBiLSTM-AE	0.689 (0.004)	0.625 (0.008)	0.468 (0.010)	0.535 (0.007)	0.647 (0.004)	0.316 (0.009)
AMP	One-hot	0.779 (0.002)	0.810 (0.003)	0.728 (0.004)	0.767 (0.002)	0.779 (0.002)	0.560 (0.004)
	Threemers	0.759 (0.001)	0.777 (0.003)	0.726 (0.002)	0.751 (0.001)	0.759 (0.001)	0.519 (0.003)
	Blosum	0.667 (0.007)	0.693 (0.010)	0.601 (0.007)	0.643 (0.007)	0.667 (0.007)	0.337 (0.015)
	ProtBERT	0.828 (0.003)	0.863 (0.005)	0.779 (0.004)	0.819 (0.003)	0.828 (0.003)	0.658 (0.007)
	cBiLSTM-AE	0.783 (0.001)	0.807 (0.006)	0.746 (0.007)	0.775 (0.002)	0.783 (0.001)	0.569 (0.002)
AOP	One-hot	0.881 (0.007)	0.889 (0.019)	0.875 (0.011)	0.881 (0.005)	0.881 (0.007)	0.764 (0.013)
	Threemers	0.769 (0.008)	0.789 (0.008)	0.737 (0.016)	0.761 (0.009)	0.769 (0.008)	0.540 (0.015)
	Blosum	0.592 (0.006)	0.615 (0.006)	0.494 (0.020)	0.547 (0.014)	0.592 (0.006)	0.189 (0.012)
	ProtBERT	0.790 (0.006)	0.789 (0.010)	0.792 (0.010)	0.790 (0.005)	0.790 (0.006)	0.580 (0.011)
	cBiLSTM-AE	0.846 (0.006)	0.844 (0.006)	0.848 (0.006)	0.846 (0.006)	0.846 (0.006)	0.692 (0.012)

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

Table 9: Regression metrics for four benchmark tasks.

Peptide property	Encoding type	MAE	RMSE	R ²
ISI	One-hot	15.778 (0.059)	21.263 (0.177)	0.445 (0.004)
	Threemers	15.811 (0.049)	21.279 (0.122)	0.444 (0.006)
	Blosum	19.637 (0.036)	26.458 (0.143)	0.141 (0.010)
	ProtBERT	16.406 (0.059)	22.578 (0.099)	0.374 (0.007)
	cBiLSTM-AE	18.021 (0.043)	24.472 (0.113)	0.265 (0.011)
TNC	One-hot	1.157 (0.007)	1.872 (0.019)	0.823 (0.004)
	Threemers	1.596 (0.009)	2.443 (0.006)	0.698 (0.006)
	Blosum	2.830 (0.006)	4.173 (0.011)	0.119 (0.011)
	ProtBERT	1.132 (0.005)	1.731 (0.032)	0.849 (0.004)
	cBiLSTM-AE	1.200 (0.007)	1.716 (0.011)	0.851 (0.003)
IEP	One-hot	1.268 (0.002)	1.557 (0.004)	0.694 (0.002)
	Threemers	1.397 (0.007)	1.745 (0.009)	0.616 (0.004)
	Blosum	2.272 (0.005)	2.660 (0.003)	0.107 (0.003)
	ProtBERT	0.930 (0.003)	1.231 (0.003)	0.809 (0.001)
	cBiLSTM-AE	1.069 (0.005)	1.375 (0.004)	0.761 (0.002)
MW	One-hot	495.466 (1.437)	738.626 (3.331)	0.937 (0.001)
	Threemers	118.957 (1.394)	180.648 (3.629)	0.996 (0.0002)
	Blosum	248.439 (1.389)	363.888 (2.251)	0.985 (0.0002)
	ProtBERT	552.071 (1.058)	763.360 (1.980)	0.932 (0.0004)
	cBiLSTM-AE	75.067 (0.547)	104.754 (0.715)	0.999 (0.00002)

1242 D APPENDIX D
12431244 D.1 LIMITATIONS
12451246 **Sequence length constraints.** During training we capped sequences at 96 residues (more than
1247 enough for peptides but small for full-scale proteins) to satisfy the CAE fixed input size. Although
1248 subsequent BiLSTM and transformer variants can process variable lengths, the cap was retained for
1249 experimental consistency, where markedly longer inputs can render the embedding quality; sequences
1250 beyond 96 residues therefore require additional architectural changes.1251 **Handling of peptide modifications.** The model supports only 20 canonical residues plus a small
1252 set of predefined modifications. Extending to user-defined or rare modifications would necessitate
1253 (i) quantum-chemical calculations for new DFT descriptors and (ii) updates to the RDKit monomer
1254 dictionary.1255 **Peptide structure limitations.** The framework is restricted to linear peptides; branched or cyclic
1256 topologies were excluded. Supporting such structures would entail substantial changes to both
1257 preprocessing and model design and is left for future work.1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

1296 E APPENDIX E
12971298 LLM USAGE STATEMENT
12991300 Large language models were used during the preparation of this manuscript solely for text refinement
1301 and stylistic improvement. Specifically, LLMs assisted with phrasing adjustments and ensuring
1302 adherence to academic writing conventions. All scientific content, ideas, and conclusions remain
1303 entirely our own.

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

F APPENDIX F

Table 10: Correlation between embedding dimensions and instability index.

Peptide property	Encoded input feature	PCC (r)
1356	exactmw	-0.034
1357	amw	0.071
1358	lipinskiHBA	0.131
1359	lipinskiHBD	0.124
1360	NumRotatableBonds	0.010
1361	NumHBD	0.017
1362	NumHBA	0.101
1363	NumHeavyAtoms	0.040
1364	NumAtoms	-0.120
1365	NumHeteroatoms	0.141
1366	NumAmideBonds	-0.068
1367	FractionCSP3	0.101
1368	NumRings	-0.071
1369	NumAromaticRings	-0.253
1370	NumAliphaticRings	0.102
1371	NumSaturatedRings	0.137
1372	NumHeterocycles	0.138
1373	NumAromaticHeterocycles	-0.013
1374	NumSaturatedHeterocycles	0.202
1375	NumAliphaticHeterocycles	0.059
1376	NumSpiroAtoms	-0.193
1377	NumBridgeheadAtoms	-0.020
1378	NumAtomStereoCenters	-0.084
1379	ISI	0.022
1380	NumUnspecifiedAtomStereoCenters	-0.005
1381	labuteASA	0.142
1382	tpsa	-0.097
1383	CrippenClogP	0.047
1384	CrippenMR	-0.057
1385	chi0v	-0.049
1386	chi1v	0.059
1387	chi2v	0.069
1388	chi3v	0.076
1389	chi4v	0.007
1390	chi0n	-0.036
1391	chi1n	0.007
1392	chi2n	0.001
1393	chi3n	0.016
1394	chi4n	-0.017
1395	hallKierAlpha	0.045
1396	kappa1	0.063
1397	kappa2	0.056
1398	kappa3	-0.008
1399	Phi	0.068
1400	Ca ²⁺ interaction energy	0.076
1401	Mg ²⁺ interaction energy	0.084
1402	Ba ²⁺ interaction energy	

Table 11: Correlation between embedding dimensions and theoretical net charge.

Peptide property	Chemical descriptor	PCC (r)
1404	exactmw	0.054
1405	amw	-0.063
1406	lipinskiHBA	-0.293
1407	lipinskiHBD	0.446
1408	NumRotatableBonds	0.065
1409	NumHBD	0.012
1410	NumHBA	0.006
1411	NumHeavyAtoms	-0.469
1412	NumAtoms	0.206
1413	NumHeteroatoms	-0.212
1414	NumAmideBonds	-0.150
1415	FractionCSP3	0.109
1416	NumRings	-0.332
1417	NumAromaticRings	-0.320
1418	NumAliphaticRings	-0.218
1419	NumSaturatedRings	-0.267
1420	NumHeterocycles	-0.301
1421	NumAromaticHeterocycles	-0.055
1422	NumSaturatedHeterocycles	-0.131
1423	NumAliphaticHeterocycles	-0.156
1424	NumSpiroAtoms	-0.238
1425	NumBridgeheadAtoms	-0.006
1426	NumAtomStereoCenters	-0.537
1427	NumUnspecifiedAtomStereoCenters	-0.033
1428	labuteASA	0.074
1429	tpsa	0.012
1430	CrippenClogP	-0.738
1431	CrippenMR	0.075
1432	TNC	
1433	chi0v	0.179
1434	chi1v	-0.055
1435	chi2v	0.013
1436	chi3v	0.032
1437	chi4v	0.022
1438	chi0n	0.254
1439	chi1n	0.147
1440	chi2n	-0.140
1441	chi3n	-0.109
1442	chi4n	-0.140
1443	hallKierAlpha	0.190
1444	kappa1	0.035
1445	kappa2	0.395
1446	kappa3	0.042
1447	Phi	0.402
1448	Ca^{2+} interaction energy	0.542
1449	Mg^{2+} interaction energy	0.369
1450	Ba^{2+} interaction energy	0.375

1457

1458
1459

Table 12: Correlation between embedding dimensions and isoelectric point.

1460
1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511

Peptide property	Chemical descriptor	PCC (r)
	exactmw	0.012
	amw	-0.005
	lipinskiHBA	-0.150
	lipinskiHBD	0.387
	NumRotatableBonds	0.028
	NumHBD	-0.060
	NumHBA	0.022
	NumHeavyAtoms	-0.334
	NumAtoms	0.197
	NumHeteroatoms	-0.091
	NumAmideBonds	-0.082
	FractionCSP3	0.147
	NumRings	-0.335
	NumAromaticRings	-0.321
	NumAliphaticRings	-0.121
	NumSaturatedRings	-0.216
	NumHeterocycles	-0.185
	NumAromaticHeterocycles	0.030
	NumSaturatedHeterocycles	-0.064
	NumAliphaticHeterocycles	-0.162
	NumSpiroAtoms	-0.215
	NumBridgeheadAtoms	0.061
	NumAtomStereoCenters	-0.392
IEP	NumUnspecifiedAtomStereoCenters	0.043
	labuteASA	0.015
	tpsa	0.079
	CrippenClogP	-0.620
	CrippenMR	0.029
	chi0v	0.110
	chi1v	-0.112
	chi2v	0.076
	chi3v	0.062
	chi4v	0.033
	chi0n	0.220
	chi1n	0.084
	chi2n	-0.130
	chi3n	-0.069
	chi4n	-0.135
	hallKierAlpha	0.150
	kappa1	0.095
	kappa2	0.359
	kappa3	0.105
	Phi	0.367
	Ca^{2+} interaction energy	0.494
	Mg^{2+} interaction energy	0.365
	Ba^{2+} interaction energy	0.371

Table 13: Correlation between embedding dimensions and molecular weight.

Peptide property	Chemical descriptor	PCC (r)	
1512	exactmw	-0.601	
1513	amw	0.825	
1514	lipinskiHBA	0.725	
1515	lipinskiHBD	0.192	
1516	NumRotatableBonds	-0.281	
1517	NumHBD	-0.948	
1518	NumHBA	0.726	
1519	NumHeavyAtoms	0.730	
1520	NumAtoms	-0.090	
1521	NumHeteroatoms	0.828	
1522	NumAmideBonds	0.472	
1523	FractionCSP3	0.154	
1524	NumRings	-0.294	
1525	NumAromaticRings	-0.196	
1526	NumAliphaticRings	0.546	
1527	NumSaturatedRings	-0.368	
1528	NumHeterocycles	0.327	
1529	NumAromaticHeterocycles	0.927	
1530	NumSaturatedHeterocycles	0.176	
1531	NumAliphaticHeterocycles	-0.740	
1532	NumSpiroAtoms	-0.556	
1533	NumBridgeheadAtoms	0.819	
1534	NumAtomStereoCenters	0.138	
1535	MW	NumUnspecifiedAtomStereoCenters	0.979
1536	labuteASA	-0.451	
1537	tpsa	0.774	
1538	CrippenClogP	-0.257	
1539	CrippenMR	-0.272	
1540	chi0v	-0.687	
1541	chi1v	-0.994	
1542	chi2v	0.992	
1543	chi3v	0.766	
1544	chi4v	0.386	
1545	chi0n	-0.350	
1546	chi1n	-0.488	
1547	chi2n	-0.040	
1548	chi3n	0.434	
1549	chi4n	-0.006	
1550	hallKierAlpha	-0.257	
1551	kappa1	0.849	
1552	kappa2	0.685	
1553	kappa3	0.989	
1554	Phi	0.732	
1555	Ca ²⁺ interaction energy	0.734	
1556	Mg ²⁺ interaction energy	0.888	
1557	Ba ²⁺ interaction energy	0.883	

1566

1567

1568

1569

1570

1571

1572

1573

1574

1575

1576

1577

1578

1579

Table 14: Regression tasks results.

Encoding type	R^2 (5-fold cross validation)						
	ISI	TNC	IEP	MW	Avg.	Min	Max
One-hot	0.445 (0.004)	0.823 (0.004)	0.694 (0.002)	0.937 (0.001)	<u>0.725</u>	0.445	0.937
Blosum	0.141 (0.010)	0.119 (0.011)	0.107 (0.003)	0.985 (0.000)	0.338	0.107	0.985
Threemers	<u>0.444</u> <u>(0.006)</u>	0.698 (0.006)	0.616 (0.004)	<u>0.996</u> <u>(0.000)</u>	0.689	<u>0.444</u>	<u>0.996</u>
cBiLSTM-AE	0.265 (0.011)	0.851 (0.003)	0.761 (0.002)	0.999 (0.000)	0.719	0.265	0.999

1580

1581

1582

1583

1584

1585

1586

1587

1588

1589

1590

1591

1592

1593

1594

1595

1596

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619