# Harnessing Interpretable Deep Learning to Predict Meropenem Resistance in *Klebsiella pneumoniae*

**Nicolas da Matta Freire Araújo[1], Mateus Fernandes Santos[2], Renata Freire Alves Pereira[3,4], Rafaela Correia Brum[5], Felipe Ramos Pinheiro[3,6], Melise Chaves Silveira[7], Felicita Mabel Duré[1,8], Beatriz de Lima Alessio Müller[9], Audrien Alves Andrade de Souza[9], Alessandra Beatriz Santos Rondon Souza[4,10], Ágatha Ferreira de Souza[10], Ana Paula D'Alincourt Carvalho-Assef[11], Aline dos Santos Moreira[9], Alba Cristina Magalhaes Alves de Melo[12], Marcelo Trindade dos Santos[13], Adriano Maurício de Almeida Côrtes[2,14], Bruno de Araújo Penna[10], Thiago Pavoni Gomes Chagas[3,15], Fábio Aguiar-Alves[3,4,6,16], Fabrício Alves Barbosa da Silva[1,*]**

1 - Scientific Computing Program, Oswaldo Cruz Foundation (FIOCRUZ), Rio de Janeiro, Brazil
2 - Department of Applied Mathematics, Institute of Mathematics, Federal University of Rio de Janeiro (UFRJ), Rio de Janeiro, 21941-909, Brazil
3 - Laboratory of Molecular Epidemiology and Biotechnology, School of Pharmacy/Fluminense Federal University, Mario Viana, 523 - Niteroi/RJ, 2424.241-000, Brazil.
4 - Graduate Program in Microbiology and Parasitology. Biomedical Institute - Fluminense Federal University, Valonguinho Campus, Centro, Niteroi/RJ, 24020-140, Brazil.
5 -Department of System Engineering and Computation, State University of Rio de Janeiro, Rio de Janeiro, Brazil
6 - Graduate Program in Pathology, Fluminense Federal University, Niteroi /RJ, Brazil.
7 - Bioinformatics Laboratory, National Laboratory for Scientific Computing (LNCC), Petrópolis, Brazil
8 - Central Public Health Laboratory (LCSP), Ministry of Public Health and Social Welfare MSPyBS, Asunción, Paraguay
9 - Laboratory of Applied Genomics and Bioinnovations - IOC/FIOCRUZ, Next-Generation Sequencing Platforms IOC/RPT01J - Network of Technological Platforms/FIOCRUZ, Brazil
10 - Laboratory of Gram-Positive Cocci / Fluminense Federal University, Niteroi/RJ, Brazil.
11 - Laboratory of Bacteriology Applied to One Health and Antimicrobial Resistance, Oswaldo Cruz Foundation (FIOCRUZ), Rio de Janeiro, Brazil
12 - University of Brasilia (UnB), Brasilia, 70910-900, Brazil
13 - Graduate Program in Computational Modeling, National Laboratory for Scientific Computing (LNCC), Petrópolis, Brazil
14 - Systems Engineering and Computer Science Program, Coordination of Postgraduate Programs in Engineering (COPPE), Federal University of Rio de Janeiro (UFRJ), Rio de Janeiro, 21941-972, Brazil
15 - Department of Pathology, School of Medicine, Fluminense Federal University, Niteroi/RJ, Brazil.
16 - Department of Pharmaceutical Sciences, Lloyd L. Gregory School of Pharmacy, Palm Beach Atlantic University – USA.

*To whom correspondence should be addressed (fabricio.silva@fiocruz.br)

**ABSTRACT**

Antimicrobial resistance represents an escalating global healthcare threat, complicating treatment and increasing both morbidity and mortality. Deep learning offers promising solutions, particularly for bacterial profiling using omics data. For instance, classifying bacterial strains as resistant or susceptible to antibiotics depends on identifying genomic signatures associated with resistance mechanisms. This study introduces DeepMDC, a deep learning architecture for bacterial profiling that utilizes whole-genome data as input. In genomics, obtaining precise labels for each gene or mutation is costly and often ambiguous; therefore, labels are typically aggregated. Consequently, phenotypic classification is framed as a Multiple-instance learning problem, where training data comprise sets (bags) of instances with a single label per bag. The core component of DeepMDC is a Modern Hopfield Network, and its input encompasses all open reading frames (ORFs) identifiable in a genome, including small ORFs. A key feature of DeepMDC is its interpretability, which facilitates model validation and hypothesis generation regarding bacterial resistance by leveraging the attention mechanism inherent to Modern Hopfield Networks. Evaluation is centered on *Klebsiella pneumoniae* and meropenem due to their clinical significance. DeepMDC achieves values exceeding 0.9 for several performance metrics, including ROC AUC, balanced accuracy, and F1 score. Several well-known resistance-associated genes, such as the Tn*3*-like element Tn*4401* family transposase, received high attention scores during the inference process. The attention mechanism also suggests a potential role for small ORFs in meropenem resistance. Results are discussed in detail, including analysis of false positives and negatives.

## INTRODUCTION

Antimicrobial resistance is a growing global healthcare threat. It complicates treatment regimens and increases morbidity and mortality. Deep learning models can help address this issue, especially in bacterial profiling based on omics data. One relevant example is the characterization of bacterial strains' susceptibility to a given antibiotic using whole-genome data. Classifying bacterial strains as resistant or susceptible to specific antibiotics depends on identifying genomic signatures that encode resistance mechanisms. Traditionally, bioinformatics pipelines use genomic sequencing, curated databases, and rule-based methods such as PointFinder and ResFinder. These methods associate mutations and gene acquisitions with resistance phenotypes (Aytan-Aktug et al., 2020). More recently, machine learning (ML) has complemented these methods. Many studies report improved predictive performance using algorithms such as logistic regression, neural networks, and random forests (Aytan-Aktug et al., 2020; Peng et al., 2022).

Bacterial genomes are complex and high-dimensional. They comprise both highly conserved core regions and variable accessory regions that contribute to resistance phenotypes. Features of interest include single-nucleotide polymorphisms (SNPs), insertions and deletions, gene presence/absence patterns, and structural variations. Due to these diverse features, computational models should be able to handle both the sparsity and heterogeneity of genomic data. They must also distinguish between passenger mutations and those directly contributing to resistance (Lees et al., 2020).

Advances in high-throughput sequencing and omics technologies have significantly increased the volume of complex genomic data. Therefore, it is crucial to develop computational methods that efficiently infer relevant biological signatures from large volumes of weakly supervised data. Traditional classification paradigms assume high-quality, instance-level annotations. However, in genomics, acquiring precise labels for every nucleotide or gene is expensive and sometimes ambiguous. As a result, labels typically exist only at an aggregated level. These may include species, disease state, or a taxonomic classification

for a whole genome. The specific contributions of individual gene segments or k-mer patterns remain unannotated. Addressing this weak supervision is crucial for effective genomic classification.

Multiple-instance learning (MIL) is a supervised learning variant suited for data organized into sets (or bags) of instances. In MIL, each bag is assigned a single label, not the individual instances. In many biological settings, such as genome classification, the bag-level label captures a phenotype or taxonomic identity. The discriminative instances—such as sequence motifs, epigenomic signals, or gene expression profiles—remain unannotated (Kraus et al., 2016). Researchers have applied MIL to several problems, including computational pathology (Hou et al., 2016) and metagenomics (Rahman et al., 2020). Another application is immune repertoire classification (Widrich et al., 2020). This scenario features bags with hundreds of thousands of instances at extremely low witness rates. These successful adaptations strongly motivate treating genome classification as an MIL problem.

In genome classification, the entire genome is the bag of instances. Its constituent elements (e.g., gene segments, open reading frames (ORF), k-mer clusters, or other genomic markers) are instances. This bag-of-instances approach addresses the significant challenges of genomic data: the scarcity of instance-level labels and the heterogeneity of genomic regions. Through this method, MIL lets classification models learn significant patterns from aggregated weak supervision. Thus, it bypasses the need for exhaustive labeling of genomic segments.

In this work, we propose DeepMDC, a deep learning architecture for bacterial profiling using whole-genome data, framing the problem as MIL. Specifically, given a genome, the model predicts if the strain is resistant or susceptible to a predefined antibiotic. A key feature of DeepMDC is interpretability. During inference, DeepMDC assigns attention weights to each ORF in a genome, highlighting those most relevant for classification. This interpretability is crucial for validating the model and generating new insights into bacterial resistance.

In this study, we focused our evaluation of DeepMDC on the *Klebsiella pneumoniae*-meropenem pair due to its clinical importance. *K. pneumoniae* is a member of the ESKAPE group of multidrug-resistant

pathogens. It is associated with increased mortality rates, prolonged hospitalizations, and higher healthcare costs (De Oliveira et al., 2020; Loyola-Cruz et al., 2023). This opportunistic pathogen can cause urinary tract infections, bacteremia, pneumonia, and liver abscesses (WANG et al., 2020).

The World Health Organization classifies carbapenemase-producing *K. pneumoniae* (KPC) as a critical public health threat. This pathogen can inactivate penicillins, cephalosporins, and carbapenems and possesses intrinsic ampicillin resistance mediated by the $bla_{\mathrm{SHV}}$ gene (Holt et al., 2015; Wyres et al., 2020). Key virulence factors that enhance its pathogenicity include lipopolysaccharides, capsular polysaccharides, siderophores, and type 1 and 3 fimbriae (WANG et al., 2020). Meropenem is a carbapenem-class beta-lactam antibiotic with broad-spectrum activity and a good safety profile. It is effective against a diverse range of microbial pathogens. Given these factors, we selected the *K. pneumoniae*-meropenem pair as a relevant test case for DeepMDC. Below, we show that DeepMDC achieves values exceeding 0.9 for several performance metrics, including ROC AUC, balanced accuracy, and F1 score. Several well-known resistance-associated genes received high attention scores during the inference process. Finally, we also discuss in detail the interpretation of classification inferences made by DeepMDC.

## RESULTS

### DEEPMDC - INTERPRETABLE DEEP LEARNING ARCHITECTURE FOR BACTERIAL GENOME PROFILING

The central element of the deep learning architecture considered in this work is Modern Hopfield Networks (Ramsauer et al., 2020). Modern Hopfield networks (also known as dense associative memories) generalize classical binary Hopfield models to continuous states, thereby offering differentiability and compatibility with contemporary deep learning frameworks. A key theoretical

advantage is their update rule, which is mathematically equivalent to the key-value attention mechanism commonly used in transformers. This equivalence provides a natural interpretation of modern Hopfield networks as associative memories that retrieve stored patterns from noisy inputs. The energy function underlying these networks ensures that, for well-separated patterns, the retrieval error decreases exponentially with the input dimension, resulting in an exponential memory capacity that scales with the vector space dimension (Hu et al., 2023).

Figure 1 depicts DeepMDC's overall architecture, inspired by the DeepRC architecture (Widrich et al, 2020). In the first stage, we divide the entire genome into ORFs, which we then translate into amino acid sequences. In this architecture, all ORF instances compose the DeepMDC input, including small ORFs (sORF). In the second stage, a predefined encoder processes the bag of ORFs and produces a latent-space representation. The encoder can be a 1D convolutional network, an LSTM, or a genomic Language Model (e.g., Zhou et al., 2023; Dalla-Torre et al., 2025). In the third stage, a dense associative memory (Modern Hopfield Network) processes fixed-size latent representations. The Modern Hopfield Network then generates a fixed-sized vector, which serves as input to an MLP classifier.

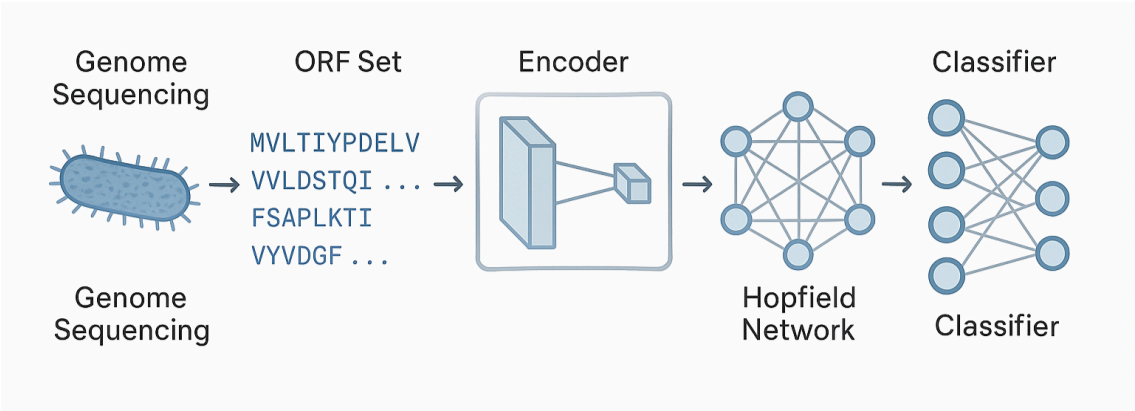


**FIGURE 1. DeepMDC architecture for processing bacterial whole-genome sequences.**

We evaluate this architecture for classifying bacterial strains into susceptible/resistant phenotypes to meropenem. We treated intermediate strains as susceptible, yielding a binary classifier. For the results presented in this work, the encoder was a 1D convolutional neural network (CNN).

In CNN, each filter (or kernel) is responsible for detecting subsequences of length k. This operation evolves each kernel so that they detect biological motifs that help distinguish the classes of an antimicrobial (resistant or sensitive). CNN processing converts each bacterial ORF into a fixed-length feature vector. This set of fixed-sized vectors serves as input to a modern Hopfield network. The update rule of the Modern Hopfield Networks is equivalent to the Transformers' self-attention mechanism (Ramsauer et al., 2020). In this work, we implemented a single attention head and utilized attention coefficients assigned to the ORFs during inference to interpret the model's output. The higher the attention weight attributed to an ORF, the higher the importance of that ORF in the inference decision process. It is worth noting that the sum of attention weights over all ORFs of a bacterial genome (> 5000 ORFs) is equal to one. In practice, we observe that attention weights tend to zero rapidly after the five highest attention weights of a genome. For this reason, we focus on the five highest-attention-weighted ORFs for each genome in the following. Additional details on the Modern Hopfield Network structure and implementation can be found in the supplementary text.

The classification network translates the representations into probabilities for each class (R, S). For training, we employed the Cross-Entropy Loss function and the Adam optimizer (Kingma & Ba, 2014) to adjust the model weights across all three networks, aiming to minimize error and enhance subsequent predictions.

In summary, the CNN learns to recognize patterns in ORFs, the modern Hopfield network identifies and prioritizes the most relevant ORFs, and the classification network checks whether the prediction approximates the actual label and fine-tunes all networks to improve performance.

## DEEPMDC ACHIEVES HIGH BALANCED ACCURACY FOR DIFFERENT TEST SETS

For model training, we used 1,550 *K. pneumoniae* clinical genomes collected by Lüftingeret et al. (2021), classified as follows: 775 are resistant to meropenem, 715 are susceptible, and 60 are intermediate. For training and evaluation, we considered the intermediate samples as susceptible. The set of parameters used for training, the methodology employed for hyperparameter screening, and additional details of the training process are provided in the Materials and Methods section and the Supplementary Material.

We used 5-fold cross-validation during training, i.e., we divided the training set into five subsets and repeated the training process 5 times. Each time we used a different subset for validation, to avoid overfitting. The validation results for each of the five folds are in Table I. One can observe that folds 1 and 2 achieved an area under the curve, balanced accuracy, and an F1-score of 0.9 or better. The Matheus Correlation Coefficients in these two models are 0.8 or better.

We used two datasets to test DeepMDC. Test set I comprises 155 additional clinical samples collected by Lüftingeret et al. (2021) and not included in the training/validation process. Test Set II shall consist of 100 unpublished genomes from various origins, including ambient, clinical, and environmental, that have been processed by our group. We processed the raw data for test set II using the CABGen server (Duré et al., 2022). Note that Test Set II is much more diverse than Test Set I with respect to sample origin. It is also worth noting that the training process used by DeepMDC employed only clinical samples. Additional details on the collection, classification, and data processing phases of these samples are in the Materials and Methods section.

For testing DeepMDC, we treated each fold generated during training/validation as a separate model. We submitted the test sets to all folds and considered the best result for each test set. Detailed results for all sets are available in the supplementary material.

For test set I, fold two stands out with a Receiver Operating Characteristic AUC of 0.93, a balanced accuracy of 0.86, a F1 score of 0.86, and a Matthews Correlation Coefficient of 0.72. Considering all folds, the best performance is achieved by folds 1 and 2, consistent with the results from the validation phase.

For test set II, once again, folds 1 and 2 obtained the best results. Receiver Operating Characteristic AUC of 0,86, a balanced accuracy of 0,87, a F1 score of 0,88, and a Matthews Correlation Coefficient of 0,74.

**Table I - Results in validation phase. Roc_auc stands for the Receiver Operating Characteristic Area Under the Curve, bacc stands for balanced accuracy, and mcc stands for Matthews Correlation Coefficient. Details on these metrics are in the Materials and Methods section.**

| Fold | roc_auc | bacc | f1_score | mcc |
|---|---|---|---|---|
| 1 | 0,93 | 0,91 | 0,9 | 0,81 |
| 2 | 0,95 | 0,9 | 0,9 | 0,8 |
| 3 | 0,89 | 0,83 | 0,83 | 0,67 |
| 4 | 0,87 | 0,83 | 0,83 | 0,67 |
| 5 | 0,94 | 0,87 | 0,88 | 0,75 |
| Average (Std Dev) | 0,916 (0.034) | 0,868 (0.038) | 0,868 (0.036) | 0,74 (0.068) |

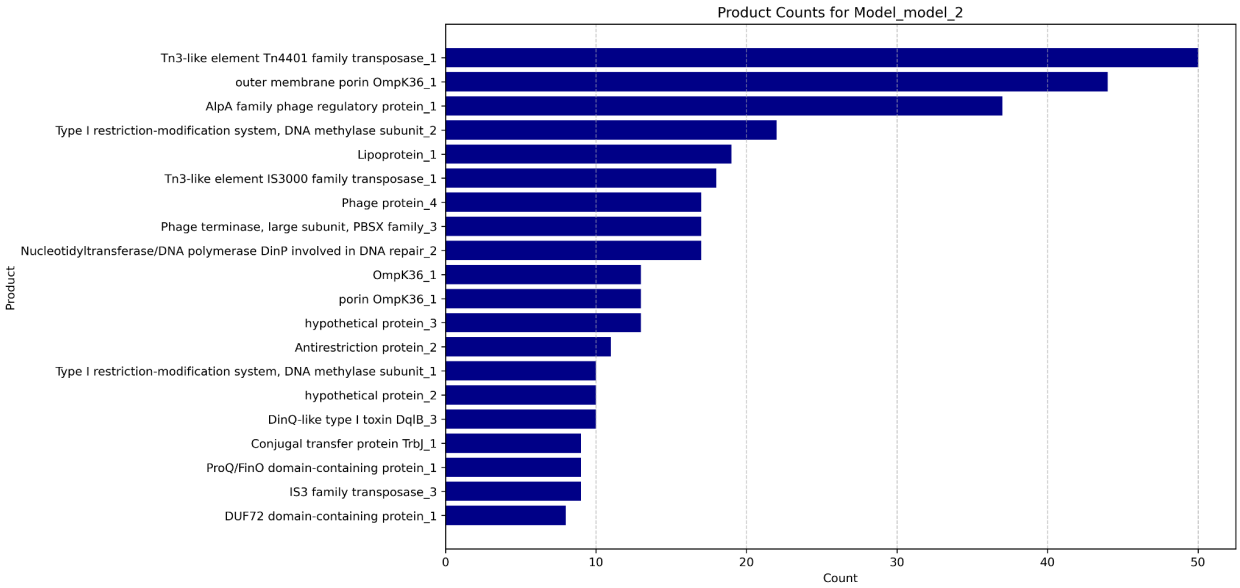

**Figure 2 - Top 20 Proteins Identified by Model 2 (Test set I). The chart displays the frequency of each product's appearance among the five products with the highest attention coefficients across all test samples.**

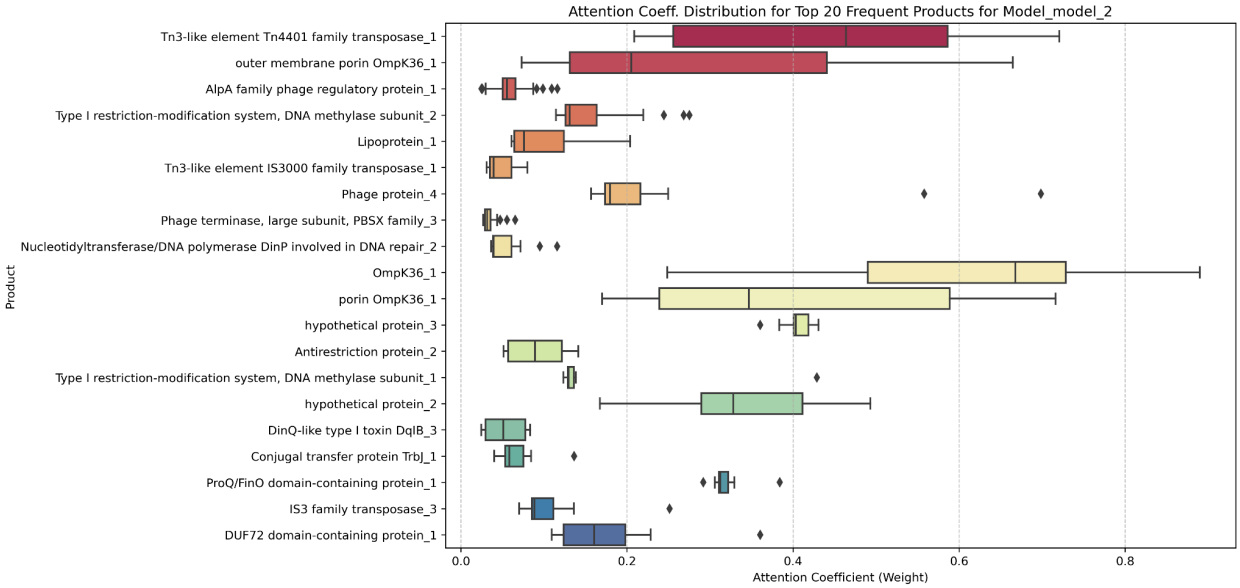

**Figure 3 - Top 20 Proteins Identified by Model 2 (Test set I). Variation of attention coefficients for each product.**

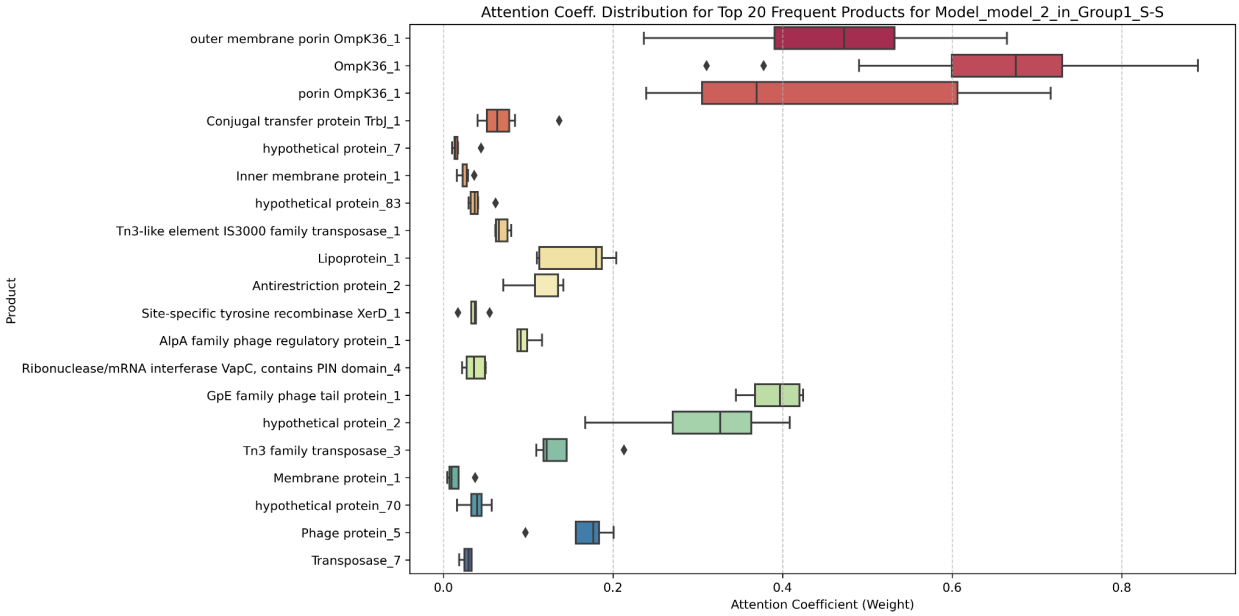


**Figure 4 - Top 20 Proteins Identified by Model 2 (Test set I) in the subgroup SG1 (S-S; susceptible samples correctly classified) and distribution of Attention Coefficients.**

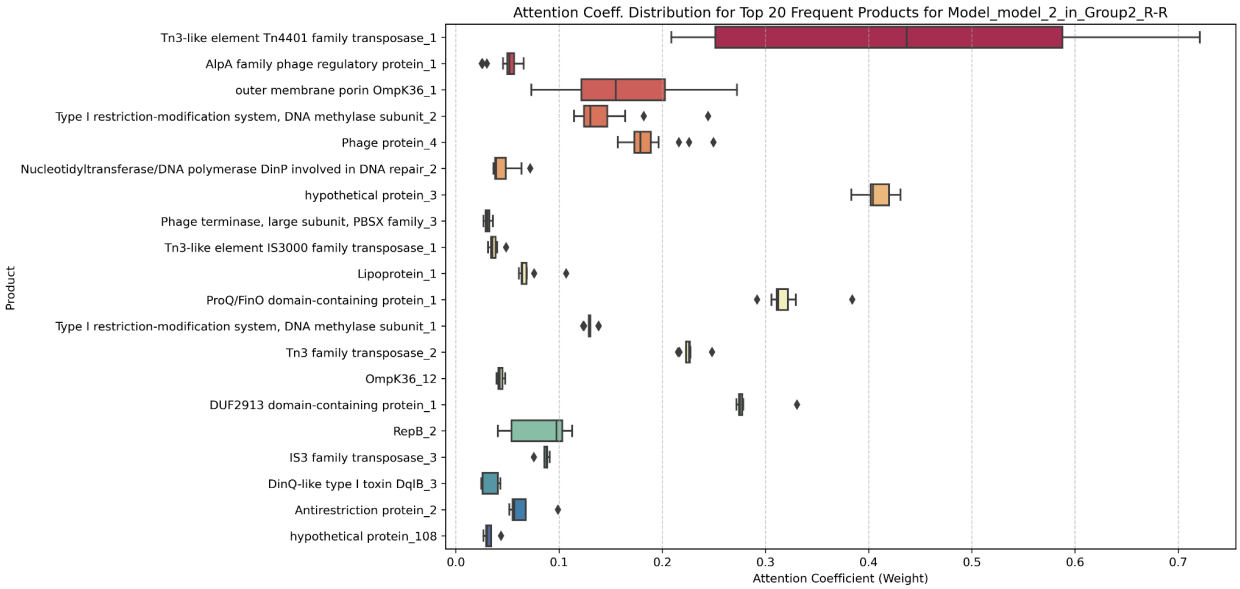

**Figure 5 - Top 20 Proteins Identified by Model 2 (Test set I) in the subgroup SG2 (R-R; resistant samples correctly classified) and distribution of Attention Coefficients.**

ATTENTION WEIGHTS CORRECTLY IDENTIFIES RESISTANT AND SUSCEPTIBILITY-RELATED GENES

For a thorough analysis of DeepMDC, we divided test samples into two groups: Test set I with 155 samples (of clinical origin), and Test set II with 100 samples recovered from the routine care of the Antônio Pedro University Hospital and veterinary and environmental samples - for a detailed description of test samples, see the Supplementary material. We evaluated the Test Set I and Test Set II samples five times, once per fold, corresponding to submitting the test samples to five separate models. We recovered the 5 ORFs with the highest attention coefficients among approximately 5,000 ORFs of this bacterium's genome for each test sample in each fold. See the supplementary materials for detailed results. For Test set I, analyzing all models in relation to accuracy, precision, recall, and F1-score, it was possible to notice that the Model corresponding to fold 2 surpassed the others in all metrics except recall, in which it was below fold 1 (See supplementary material), being fold two selected for attention weight analysis. For fold 2, the most frequent protein in the set of ORFs with the highest attention weights is the Tn*3*-like element Tn*4401* family transposase (Fig. 2). In this Model, the protein that received the highest absolute attention weight was *OmpK36* (Fig. 3).

We separated the test set samples into four subgroups: SG1 (evaluated as susceptible by both DeepMDC and AST), SG2 (evaluated as resistant by both DeepMDC and AST), SG3 (evaluated as resistant by DeepMDC but susceptible by AST), and SG4 (evaluated as susceptible by DeepMDC but resistant by AST). In Fold 2, SG1 was the largest group (67 samples), followed by SG2 (66 samples), while SG3 and SG4 included 11 samples each. The twenty most frequent proteins in the set of ORFs with the highest attention weights in each subgroup were then listed (see Figures 4-5 and Supplementary Figures 1 and 2). For this Model, the most frequent ORF with a high attention weight is the Tn*3*-like element Tn*4401*

family transposase (Figure 2), and *OmpK36* has the highest absolute attention weight (Figure 3). Regarding SG1, the most frequently occurring product amongst the highest-attention-weight lists and the one that received the absolute highest attention weight was outer membrane porin *OmpK36* (Figure 4). Regarding SG2, the most frequent product also received the highest absolute attention: the Tn*3*-like element Tn*4401* family transposase (Figure 5). In SG3, the most frequent was a hypothetical protein, and the AlpA family phage regulatory protein received the highest absolute attention weight (Supplementary Figure 1). Moreover, finally, in SG4, the most frequent product amongst the highest-attention-weight lists was again *OmpK36*, and the Replication initiation protein received the highest absolute attention weight (Supplementary Figure 2).

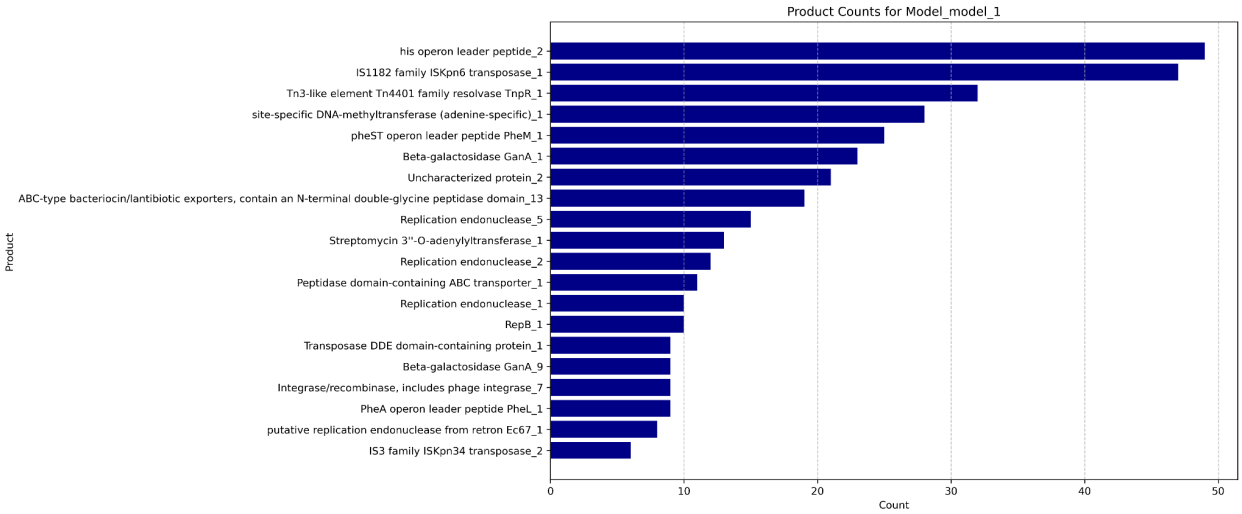


**Figure 6 - Top 20 Proteins Identified by Model 1 (Test set 2). The chart displays the frequency of each product's appearance among the five products with the highest attention coefficients across all test samples.**

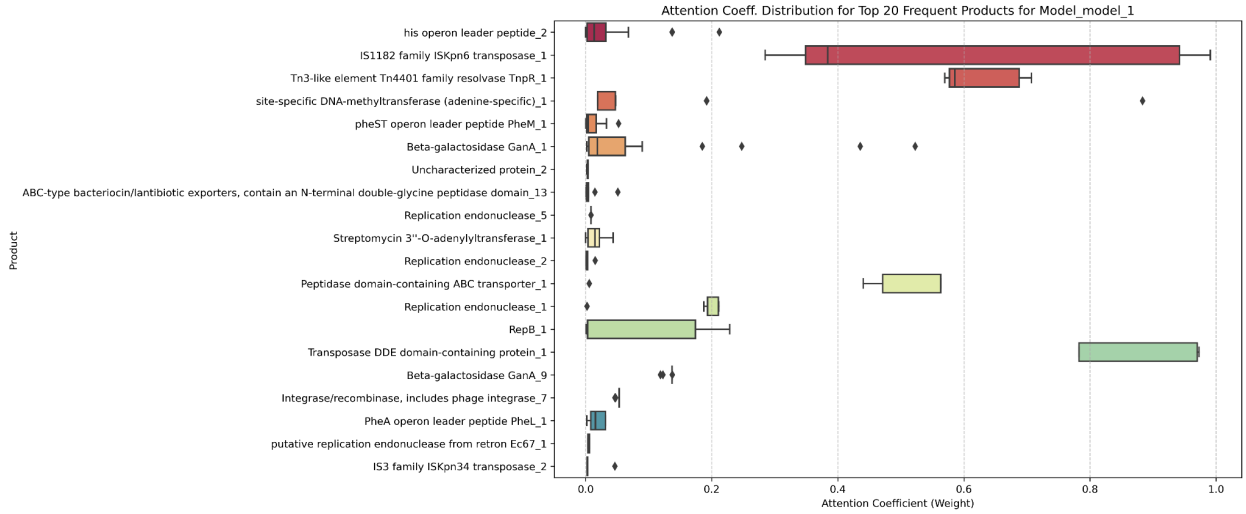


**Figure 7 - Top 20 Proteins Identified by Model 1 (Test set 2). Variation of attention coefficients for the top 20 proteins.**

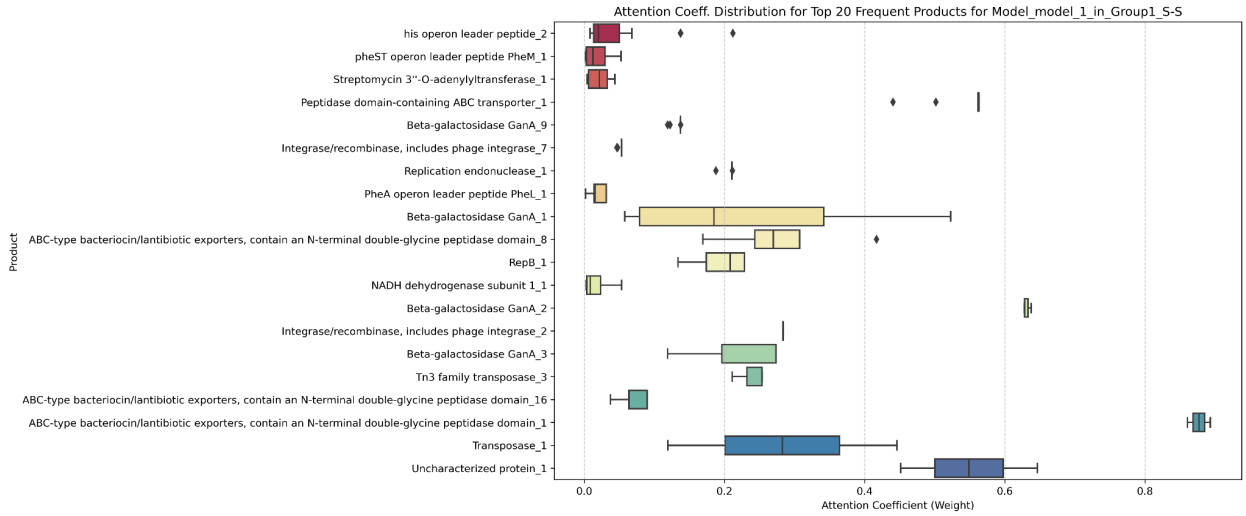


**Figure 8 - Top 20 Proteins Identified by Model 1 (Test set 2) in the subgroup SG1 (S-S; susceptible samples correctly classified) and distribution of Attention Coefficients.**

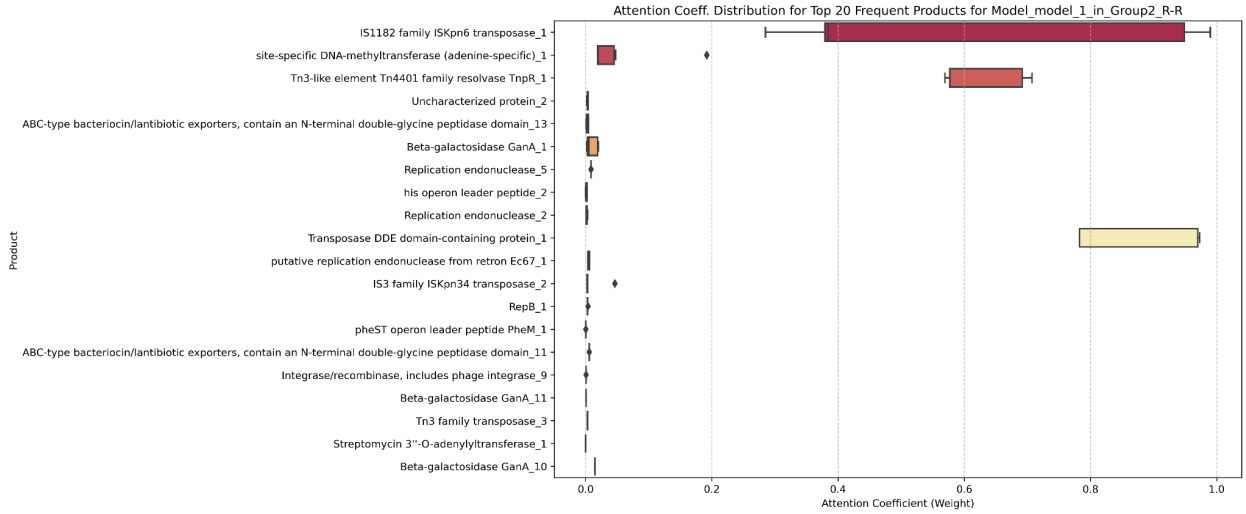


**Figure 9 - Top 20 Proteins Identified by Model 1 (Test set 2) in the subgroup SG2 (R-R; resistant samples correctly classified) and distribution of Attention Coefficients.**

For samples from Test Set 2 (100 samples, from several origins), Fold 1 performed best, surpassing other models in accuracy, precision, and F1-score. The most frequently occurring product among the highest attention-weight lists was the operon leader peptide (Figure 6). However, the IS*1182* family IS*Kpn6* transposase and a transposase DDE domain-containing protein received the highest absolute attention weights (Figure 7). We also divided test samples into four subgroups: SG1 (evaluated susceptible by DeepMDC/Susceptible by AST - 40 samples), SG2 (resistant DeepMDC/resistant AST- 47 samples), SG3 (predicted resistant by DeepMDC/susceptible by AST - 9 samples), and SG4 (predicted susceptible by DeepMDC/resistant by AST - 4 samples). In SG1, the most frequent protein was *his* operon leader peptide, and the highest absolute attention weight was assigned to an uncharacterized protein (Figure 8). In SG2, the most frequent protein was IS*1182* family IS*Kpn6* transposase, and the transposase DDE domain-containing protein received the highest absolute attention weight (Figure 9). For SG3, the most frequent protein was also the IS*1182* family IS*Kpn6* transposase, and the highest absolute attention weight was for the Tn*3*-like element Tn*4401* family resolvase TnpR (Supplementary Figure 3). Finally, for SG4,

the most frequent product was *his* operon leader peptide, and DeepMDC assigned a higher attention weight to integrase/recombinase, which includes phage integrase (Supplementary Figure 4).

It is noteworthy that some proteins receiving high attention weights contain mutations. In this case, the protein received a new index number to identify the mutation. For instance, in Figure 3 (Test set 1, SG1), the same hypothetical protein appears several times (hypothetical protein 2, 7, 70, 83), with various index numbers indicating different mutations.

## ATTENTION WEIGHTS RAISES NEW HYPOTHESIS

The interpretation of inference results may raise new hypotheses on potential targets and the mechanistic understanding of resistance. For instance, considering the subgroup SG2 (R-R; resistant samples correctly classified) in both test sets 1 and 2 (Figures 5 and 9), several interesting observations arise. In Figure 5, amongst the most frequent proteins are "Hypothetical Protein 3" and "Hypothetical Protein 108". Both hypothetical proteins represent small ORFs (sORFs—prokaryotic ORFs with fewer than 50 amino acids) and differ in length (see Supplementary Table IV). In Figure 9, an uncharacterized protein ("Uncharacterized Protein 2") also received high attention weights. This uncharacterized protein is not a sORF.

Regarding subgroup SG1 (S-S; susceptible samples correctly classified), the presence of hypothetical and uncharacterized proteins (Figures 4 and 8) warrants investigation of their functions and mechanisms of action to improve understanding of pathways related to susceptibility. In Figure 4, we have "Hypothetical Protein 2", "Hypothetical Protein 7", "Hypothetical Protein 70", and "Hypothetical Protein 83". These hypothetical proteins are all small ORFs, and "Hypothetical Protein 7" and "Hypothetical Protein 70" are mutated forms of the same hypothetical protein (sORF sequences of hypothetical proteins are available in Supplementary Table IV). In Figure 8, we have "Uncharacterized Protein 1" with high attention weight. It is worth noting that "Uncharacterized Protein 1" is not a sORF. Identifying the functions of the hypothetical and uncharacterized proteins will contribute to a better understanding of the bacterial

proteome. Indeed, several hypothetical and uncharacterized proteins exhibited high attention weights across models, suggesting potential roles for uncharacterized loci co-localized with resistance determinants. Complete sequences and attention weights of all proteins in Figures 2-9 are in the supplementary spreadsheets 1 and 2.

## DISCUSSION

The attention-based interpretability analysis of our deep learning model provided critical insights into the genomic signals influencing its predictions, highlighting both the biological coherence and the interpretive limitations in antimicrobial resistance (AMR) prediction. By examining the distribution of attention coefficients across gene products, we identified how the model differentially prioritized classical resistance determinants versus regulatory or structural features, particularly in misclassified isolates. Overall, the model prioritized biologically meaningful genomic features commonly associated with antimicrobial resistance, including porin proteins (*OmpK36*), transposases (Tn*3*-like, IS*3000*, IS*3* families), and phage-related or regulatory proteins (AlpA, terminases, methylases) (Partridge et al., 2018; Wyres & Holt, 2018). Interestingly, several hypothetical proteins also exhibited high attention weights, suggesting potential roles of uncharacterized loci co-localized with resistance determinants. Identifying the functions of the hypothetical proteins will contribute to a better understanding of the bacterial proteome and its role in virulence (Dey et al., 2022).

We displayed the 20 most frequent proteins with high attention coefficients across the 5-fold cross-validation models, highlighting elements strongly associated with antimicrobial resistance and genomic plasticity in *K. pneumoniae*. Overall, our findings reveal that the model captured a multi-layered genomic signature, integrating mobility, envelope remodeling, and regulatory systems consistent with the evolutionary strategies underlying carbapenem resistance dissemination in Enterobacterales such as *K. pneumoniae* (Wyres & Holt, 2018; Partridge et al., 2018; Rodríguez-Beltrán et al., 2021; Lam et al., 2021).

To explore the biological relevance of the model's predictions, we analyzed the distributions of attention coefficients for the most frequent predicted products across different comparison groups. Each boxplot represents the relative attention weight assigned to a specific gene product, reflecting its contribution to the model's decision-making process. Products with higher attention coefficients indicate a more substantial influence on classification outcomes. The use of attention mechanisms as an interpretability layer in genomic deep learning models enables the identification of biologically relevant features associated with antimicrobial resistance (Hu et al., 2024; Su et al., 2019).

The increase of carbapenem-resistant *K. penumoniae* (CRKP) represents an unquestionable threat to the health of hospitalized patients worldwide, with a high mortality rate (Budia-Silva et al., 2024; Borer et al., 2009). Meropenem resistance in CRKP arises primarily from the production of carbapenemases, most commonly $bla_{\mathrm{KPC}}$, $bla_{\mathrm{NDM}}$, and $bla_{\mathrm{OXA\text{-}48}}$, often in combination with other mechanisms, such as porin loss and efflux pump overexpression. These mechanisms can act independently or synergistically, leading to varying levels of resistance (Nordmann et al., 2011; Pitout et al., 2015).

In the R–R group (predicted resistant, actually resistant), DeepMDC assigned high attention weights to Tn*4401*-like transposases, *OmpK36* porins, and AlpA family phage regulatory proteins. These elements are known to be involved in the mobilization of carbapenemase genes ($bla_{\mathrm{KPC}}$) and in porin modification, both of which are critical for carbapenem resistance (Doménech-Sánchez et al., 1999; Hernández-Allés et al., 1999; Sugawara et al., 2016). The focus on DNA methylases suggests that the model also captured epigenetic or regulatory elements relevant to horizontal gene transfer. DNA methylation has been shown to modulate gene expression, replication, and conjugative transfer of mobile genetic elements, thereby contributing to the dissemination of antimicrobial resistance determinants (Roberts et al., 2015; Oliveira et al., 2020).

Building on these findings, the transposon Tn*4401* emerges as a central player in the dissemination of the $bla_{\mathrm{KPC}}$ gene. This~10 kb Tn*3*-family transposon plays a pivotal role in disseminating the $bla_{\mathrm{KPC}}$ gene, which encodes a class A carbapenemase that hydrolyzes penicillins, cephalosporins, and carbapenems (NAAS et al., 2008; CUZON et al., 2011). Structurally, it comprises the $bla_{\mathrm{KPC}}$ gene flanked

by insertion sequences IS*Kpn7* and IS*Kpn6*, along with the transposase (tnpA) and resolvase (tnpR) genes required for transposition (CUZON et al., 2011; CHEN et al., 2014). Several isoforms (e.g., Tn*4401a*, Tn*4401b*, Tn*4401c*) have been described, differing by small deletions or rearrangements in the promoter region upstream of $bla_{KPC}$, which can affect expression levels and resistance phenotypes (NAAS et al., 2008; CUZON et al., 2011). Frequently carried on conjugative plasmids of different incompatibility groups (IncFII, IncN, IncL/M), Tn*4401* facilitates horizontal transfer among Enterobacterales, driving the global spread of carbapenem resistance (CHEN et al., 2014; MATHERS et al., 2015). Consequently, Tn*4401* represents a key mobile genetic element underpinning the epidemic dissemination of KPC-producing *K. pneumoniae* and other multidrug-resistant *Enterobacterales*.

Beyond gene mobilization, outer membrane permeability barriers also contribute substantially to the carbapenem resistance phenotype. The *OmpK36* is the primary channel for hydrophilic β-lactams, including meropenem, in *K. pneumoniae*. Reduced expression, complete loss, or structural mutations that narrow the pore of *OmpK36* significantly decrease outer membrane permeability and increase carbapenem MICs, particularly when combined with β-lactamases (Tsai et al., 2011; Wong et al., 2019; David et al., 2022). Some alterations reduce carbapenem influx and confer increased resistance to meropenem, with such variants now widely disseminated across epidemic STs (David et al., 2022; Wong et al., 2022).

Moreover, synergistic effects between porin loss and enzyme overexpression have been described. In OXA-48-like producers, both *OmpK36* loss and increased $bla_{\text{OXA-48}}$ copy numbers have been implicated in carbapenem resistance at clinically relevant levels (Meekes et al., 2025). While truncation of *OmpK35* alone has a limited impact on carbapenem susceptibility, it potentiates resistance when combined with alterations in *OmpK36*. In contrast, *OmpK36* is rarely completely truncated; the most common adaptations involve pore constriction or reduced abundance (Tsai et al., 2011).

Finally, the clinical and evolutionary relevance of Iof *OmpK36* modulation is underscored by *in vivo* studies demonstrating the selection of *OmpK36*-deficient mutants during prolonged antimicrobial therapy (Strahilevitz et al., 2024). Moreover, *OmpK36*-mediated resistance often entails a fitness or

virulence cost, suggesting adaptive trade-offs during infection (Wong et al., 2022). Thus, *OmpK36* modulation is a central determinant of intracellular meropenem exposure in *K. pneumoniae* and acts as a potent amplifier of resistance in β-lactamase-producing backgrounds.

In S–S Group (true sensitive isolates), attention was concentrated on porins (*OmpK35*/*OmpK36*), phage structural proteins, and transfer proteins (TrbJ), reflecting genomic stability and preserved permeability, typical of susceptible isolates lacking mobile resistance determinants (Davies & Davies, 2010; Shintani, Sanchez, and Kimbara, 2015). The model's high attention to porin, likely intact, in sensitive isolates, indicates it correctly recognized membrane permeability preservation as a hallmark of susceptibility. In contrast, resistant isolates typically exhibit disruptions, frameshift mutations, or downregulation of OmpK36, resulting in decreased carbapenem influx (Strahilevitz et al., 2024; Palomba et al., 2025).

The model also emphasized phage proteins and AlpA-like regulators, which reflect the presence of stable prophage remnants or non-lytic phage elements integrated into the bacterial chromosome. In susceptible isolates, these phage-associated genes often play neutral or regulatory roles rather than contributing directly to resistance (Touchon et al., 2017; Colavecchio et al., 2017). In contrast, in resistant isolates, phage-related elements are frequently found adjacent to or within mobile genetic islands carrying antibiotic resistance genes (ARGs) (Partridge et al., 2018; Botelho et al., 2021). Furthermore, the attention weight attributed to TrbJ, a component of the type IV secretion system, highlights its role in bacterial conjugation-dependent horizontal gene transfer. This level of attention weight may be unexpected in sensitive isolates. However, experimental studies show that Trb-like transfer systems are encoded in a wide variety of cryptic or non-conjugative plasmids, which may persist without actively mediating the dissemination of resistance (Smillie et al., 2010; Shintani et al., 2015; Gordils-Valentin et al., 2024).

Analysis of the attention coefficient distributions revealed distinct genomic signatures underlying the model’s misclassifications. This discrepancy between genotype and phenotype highlights a key limitation of deep machine learning AMR models that rely solely on genomic data. While they effectively

capture the presence of resistance markers, they may overestimate phenotypic resistance by failing to account for transcriptional or post-transcriptional regulation.

In the R–S group (predicted resistant but phenotypically sensitive), the model focused on Tn*4401*, *OmpK36*, and the AlpA regulator—elements commonly linked to β-lactam resistance—but the isolates were phenotypically sensitive. These findings indicate the presence of genotypic but non-functional determinants, such as truncated transposons, non-expressed porins, silenced β-lactamase genes, or regulatory repression of resistance expression (Martínez & Baquero, 2014; Rodríguez-Martínez et al., 2021; Li et al., 2023). Thus, the model detected latent genomic reservoirs of resistance that are not phenotypically expressed. The existence of these "silent resistomes" has been reported in both clinical and environmental isolates, where resistance genes or operons remain transcriptionally inactive under non-selective conditions (Perron et al., 2008; Allen et al., 2010). From a biological perspective, such latent genetic elements represent potential reservoirs that may become clinically relevant upon selective pressure or horizontal transfer mediated by mobile genetic elements (Wright, 2010; Martiny et al., 2011).

Alternatively, this group may correspond to false-positive predictions, where the model identified genomic contexts associated with resistance, such as truncated β-lactamase genes, nonfunctional promoters, or incomplete transposons that lack phenotypic expression (Gibson et al., 2015; Bengtsson-Palme et al., 2017). Previous genomic surveys have shown that fragments of resistance genes and inactive variants can persist in bacterial populations, acting as evolutionary substrates for future mobilization or reactivation under selective conditions (Forsberg et al., 2012; Lopatkin et al., 2016).

Conversely, in the S–R (predicted sensitive, actually resistant) group, the model failed to detect resistance despite the presence of the phenotype. Here, attention was directed toward genes such as RepB and structural porins, while classical markers received little weight. This observation suggests that alternative or context-dependent resistance mechanisms (such as upregulated efflux pumps, promoter mutations, small regulatory RNAs, or novel plasmid-borne determinants) were at play but not captured by the model's training set. This limitation is particularly problematic in real-world settings because it means the model may systematically undercall resistance in strains that rely on subtler or emergent pathways that

are not yet well characterized (Delcour, 2009; Fernández & Hancock, 2012; Nikaido, 2011; Lehner et al., 2023; Zhang et al., 2024). We expect that larger training datasets will mitigate some of the observed shortcomings, and we intend to investigate further the impact of training dataset size on model performance.

Regarding sORFs identified as hypothetical proteins by DeepMDC, a growing body of evidence underscores their diverse and indispensable roles in bacterial resistance. Their contributions span from direct interactions with drug efflux systems and the de novo emergence of resistance-enhancing peptides to participation in toxin–antitoxin systems that promote the formation of persister cells. Beyond these direct mechanisms, sORFs also serve as critical regulators of both transcriptional and translational responses during antibiotic stress, integrating environmental signals into adaptive cellular responses. The membrane-associated functions of many sORF-encoded proteins further add to the complexity, as they negatively regulate larger efflux pumps and virulence factors, ultimately fine-tuning bacterial survival strategies under challenging conditions (Khitun et al., 2019; Knopp et al., 2019; Garai & Blanc‑Potard, 2020). For example, sORFs have been implicated in modulating toxin-antitoxin (TA) systems that directly affect cell viability under antibiotic stress (Gray & Papenfort, 2022; Garai & Blanc‑Potard, 2020). These TA systems can serve as both protective and self-destructive mechanisms, enabling bacterial populations to balance survival with the elimination of damaged cells. The dual functionality of sORF-encoded peptides in this regard reflects their broader role as modulators of cellular homeostasis in response to environmental insults. Overall, with continued research and technological innovation, the detailed functional characterization of these sORFs will pave the way for novel antimicrobial strategies that target these small yet mighty contributors to bacterial survival (Gray & Papenfort, 2022; Prins & Billerbeck, 2025).

We can improve the DeepMDC architecture in several ways. The replacement of 1D convolutional networks in the encoder phase with LSTMs or transformer-based architectures warrants further investigation. The definition of a “Biology-Informed DeepMDC”, where biological information is incorporated into the architecture and training process (Wysocka et al., 2023), similarly to what is

implemented in Physics-Informed Neural Networks (Cuomo et al., 2022) by adjusting neural network architecture, gradient optimization techniques, activation functions, and loss function structures, may improve DeepMDC performance, training time, and reduce the size of training datasets.

It is worth noting that classifying bacterial strains as resistant or susceptible to antibiotics is just one example of bacterial profiling that DeepMDC can perform. Other relevant applications include classifying strains by virulence level or biofilm-forming capacity. We intend to explore other possibilities of bacterial profiling using DeepMDC in future work.

Another future direction of research is to augment the model to analyze genomic and transcriptomic data jointly. The architecture proposed in this work is well-suited for processing the augmented input set, thanks to the modern Hopfield network's exponential storage capacity and computational efficiency. One difficulty is the limited availability of both transcriptomic and genomic data for many strains.

In summary, the attention-based interpretability analysis provided valuable insight into the genomic signals shaping our deep learning model's decisions for carbapenem resistance prediction in *Klebsiella pneumoniae*. The model consistently prioritized biologically coherent features such as Tn*4401*-like transposases, *OmpK36* porins, and regulatory or phage-associated proteins. Nevertheless, there are interpretive limitations of attention-based models. Therefore, further studies are necessary to elucidate the role of these unannotated loci in resistance and virulence. Our findings reinforce the central role of Tn*4401* in the dissemination of carbapenemase genes and in modulating *OmpK36*-mediated permeability, a determinant of carbapenem resistance.

## MATERIALS AND METHODS

### TRAINING DATA SELECTION AND PRE-PROCESSING

We recovered biosample accessions and associated meropenem (MEM) resistance/susceptibility profiles for 1705 clinical *K. pneumoniae* strains collected by Lüftingeret et al. (2021). We used 1550 samples for training/validation and 155 samples for testing (Test Set 1). For these samples, we downloaded reads from NCBI SRA databases. We obtained two additional sets of unpublished genomes from two other studies. First, a study evaluated 100 *Klebsiella pneumoniae* samples (Test Set 2; NCBI BioProject: PRJNA1356496), approved by the Research Ethics Committee of the Faculty of Medicine, Universidade Federal Fluminense (CEP–FM/UFF), under CAAE registration number 26823614.2.0000.5243. Second, we obtained a set of 21 isolates from different animals. The sample collection was approved by the Animal Use Ethics Committee of Fluminense Federal University, and an authorization for the Capture, Collection, and Transport of Biological Material No. 755/2016 by the Brazilian Institute of the Environment and Renewable Natural Resources (*Portuguese: Instituto Brasileiro do Meio Ambiente e dos Recursos Naturais Renováveis*, IBAMA). Test Set II genomes of *Klebsiella pneumoniae* were sequenced using Illumina® NGS on a NextSeq 2000® system. We extracted genomic DNA at the Universidade Federal Fluminense (UFF) and processed it at the Next-Generation Sequencing Platform – RPT01J (Oswaldo Cruz Institute, Technological Platforms Network, FIOCRUZ).

We assessed DNA quality and purity using a NanoDrop 2000 (Thermo Fisher Scientific Inc.), with A260/280 ratios of 1.8–2.0 and A260/230 ratios of approximately 2.0 considered indicative of high-quality DNA. We quantified DNA concentrations using a Qubit 4.0 fluorimeter, and we prepared libraries with the Illumina DNA Prep kit according to the manufacturer's instructions. Library fragment sizes (~600 bp) were evaluated with the TapeStation 4200® (Agilent Technologies®), and pooled libraries were normalized to 2 nM and diluted to a loading concentration of ~650 pM.

We performed sequencing using the 300-cycle P1 Kit, yielding ~30 Gb of data and ~100 million reads per run, with a minimum coverage of 50× per sample. For the training and Test Set I samples, the reads were assembled using Unicycler version 0.4.9 (conservative mode, minimum length 500 bp) (Wick et al., 2017). A maximum of 500 contigs was used as the cut-off for the genomic assemblies. The annotation was made using Bakta version v1.11 (Schwengers et al., 2021). The protein annotations were used as input

for training the model. We utilized the CABGen server (https://cabgen.fiocruz.br/) to conduct downstream computational analyses (Duré et al., 2022) on Test Set II samples.

## HYPERPARAMETER SPACE EXPLORATION

We used the Python Hydra library (https://hydra.cc/) to explore optimal parameter settings. After several training runs, we defined a standard set of hyperparameters to generate the models (Table II).

**Table II – Hyperparameter space search used to generate the models. pos_weight: defines the relative weight of the positive class in the loss function; n_updates: corresponds to the total number of parameter updates during training; evaluate_at: indicates the evaluation frequency on the validation set; kernel_size: denotes the size of the convolutional filters applied to the sequences; n_kernels: represents the number of convolutional filters; sample_n_sequences: defines how many ORFs are sampled per genome; learning_rate: is the learning rate used by the optimizer; rnd_seed: sets the random seed for reproducibility; sequence_embedding_layers: specifies the number of sequence embedding layers; attention_layers: defines the number of layers in the attention module; attention_units: corresponds to the number of internal units in the attention mechanism; output_layers: indicates the number of layers in the output network; output_units: represents the number of units in these layers; sequence_embedding_dropout: is the dropout rate applied to the sequence embedding; attention_dropout: corresponds to the dropout applied in the attention module; and output_dropout: is the dropout rate applied to the model output.**

| **Run** | 1 | 2 | 3 |
|---|---|---|---|
| **Hyperparameter** | | | |
| pos_weight | 1 | 1 | 1 |
| n_updates | 15000 | 50000 | 50000 |
| evaluate_at | 200,300,500 | 500,1000,2000 | 500,1000,2000 |

| | | | |
|---|---|---|---|
| kernel_size | 48,51,60 | 48 | 96 |
| n_kernels | 32,64 | 64 | 64,72 |
| sample_n_sequences | 0 | 0 | 0 |
| learning_rate | 0.0001 | 0.0001 | 0.0001 |
| rnd_seed | 0 | 0 | 0 |
| sequence_embedding_layers | 1 | 1 | 1 |
| attention_layers | 2 | 2 | 2 |
| attention_units | 32 | 32 | 32 |
| output_layers | 1 | 1 | 1 |
| output_units | 32 | 32 | 32 |
| sequence_embedding_dropout | 0 | 0 | 0 |
| attention_dropout | 0 | 0 | 0 |
| output_dropout | 0 | 0 | 0 |

MODEL TRAINING

For model training, since there were few representatives of the intermediate class (I), we decided to include class I within the sensitive group (S) and define the resistant group (R) as the positive class (classifier output > 0.5).

The data we used consisted of a table with the sample ID and its classification regarding resistance to meropenem, as well as tables containing the amino acid sequences of the open reading frames (ORFs) and their counts within the corresponding genome. Once training began, we performed validation at the end of each epoch to assess the model's accuracy. If the loss was lower than that of the previous epoch, we saved the model until the next validation step. At the end of training, the model with the lowest loss was loaded and evaluated on the test dataset, which is the portion of the data reserved for evaluating the model on

data not used during training. The result of this analysis was considered the final performance of the training process. Subsequently, an independent test was conducted using a dedicated set of samples that had not participated in any stage of the training or internal tuning (i.e., the training, validation, or internal test sets). Both internal and independent test results were considered when selecting the best model.

We generated the models (folds) independently, that is, one model per training run. We selected the best-performing model, and we used its hyperparameters in a cross-validation procedure to assess its robustness. In this validation, we divided the data into four folds for training and one for validation/testing, with the validation fold alternating with each new model. Similarly, we evaluated the five newly trained models using external data to assess their performance. Finally, we calculated the mean and standard deviation for each metric and selected the best model.

For model training and inference, we utilized a server with the following configuration: 2x AMD EPYC 7543 2.8GHz 32-core processors (64 physical cores total), 256 GB RAM, 32 TB HDD, 480GB SSD, and an NVIDIA A30 24GB GPU. We also used Amazon Web Services (AWS) *p5.4xlarge* instances with the following configuration: AMD EPYC 7R13 processor with eight physical cores, 16 vCPUs, 256 GB RAM, and an NVIDIA H100 80GB GPU.

## EVALUATION METRICS

Evaluation metrics are important indicators of how well-trained models perform classification tasks. In this work, we have calculated four metrics: Receiver Operating Characteristic Area Under the Curve (ROC AUC), balanced accuracy, F1 score, and Matthews Correlation Coefficient (MCC). We describe these metrics below:

**1- Balanced accuracy:** Balanced accuracy is a metric that corrects for the misleading results that standard accuracy can produce when evaluating machine learning models on imbalanced datasets. Instead

of simply calculating the percentage of correct predictions, balanced accuracy ensures each class contributes equally to the final score.

In binary classification, balanced accuracy is the average of the recall for each class. Recall for a class is also known as its true positive rate.

- *Sensitivity* (True Positive Rate): The percentage of positive cases correctly identified.
- *Specificity* (True Negative Rate): The percentage of negative cases correctly identified.

The formula is:

$$bacc = \frac{sensitivity + specificity}{2} \quad (1)$$

**2- F1-score:** F1-score is the harmonic mean of Precision and Recall, and is defined as a single parameter that balances these other two metrics:

$$F1score = 2\frac{Precision*Recall}{Precision+Recall} \quad (2)$$

Precision is the proportion of correct positive predictions with respect to all (true or false) positive predictions returned by the model:

$$Precision = \frac{True\ Positives}{True\ Positives + False\ Positives} \quad (3)$$

Recall (or True Positive rate) is the proportion of actual positive measures correctly predicted by the model:

$$Recall = \frac{True\ Positives}{True\ Positives + False\ Negatives} \quad (4)$$

**3- Receiver Operating Characteristic Area Under the Curve (ROC AUC):** Receiver Operating Characteristic (ROC) Area Under the Curve (AUC) is a performance metric that summarizes a binary

classifier's ability to distinguish between positive and negative classes. It is a single scalar value ranging from 0.5 to 1.0, where 0.5 indicates no discriminative ability (equivalent to random guessing) and 1.0 signifies a perfect classifier. The AUC is interpreted as the probability that a randomly selected positive sample is ranked higher than a randomly selected negative sample. It is calculated by integrating the ROC curve. The ROC curve is created by plotting the True Positive Rate against the False Positive Rate at various threshold settings.

**4- Matthews Correlation Coefficient:** The Matthews Correlation Coefficient (MCC) is a metric that evaluates the quality of binary classification models, providing a balanced measure even for imbalanced datasets. Ranging from -1 to +1, the MCC considers all four outcomes from a confusion matrix (True Positives, True Negatives, False Positives, False Negatives). A score of +1 indicates a perfect prediction, 0 suggests random guessing, and -1 signifies a complete disagreement between prediction and actual values. The formula for the Matthews Correlation Coefficient (MCC) for binary classification derives from the four values in a confusion matrix: true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN):

$$MCC = \frac{TP.TN - FP.FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}} \quad (5)$$

ATTENTION COEFFICIENTS RETRIEVAL

For attention-weight retrieval after an inference, we used the scripts extract_attention_weights.py and prediction_and_att_weights.py, available in the companion repository (https://github.com/CABGenOrg/deepmdc). Those scripts extract the attention coefficients for each ORF in each sample. The *extract_attention_weights.py* script extracts all coefficients, resulting in extensive tables, while the *prediction_and_att_weights.py* script collects only the top 5 coefficients for each sample.

It also displays the model's prediction for the corresponding sample and whether the prediction is correct. The execution of either script follows the same parameters as the test script.

## DATA ANALYSIS

We used Python version 3.11, along with its modules matplotlib (version 3.10.6), NumPy (version 2.3.2), pandas (version 2.3.2), and seaborn (version 0.13.2) to analyze the data. The analysis script, data_analysis.py, is available in the companion repository (https://github.com/CABGenOrg/deepmdc). We divided the analysis into three parts: frequency analysis, attention coefficient analysis, and performance metrics.

Initially, we merged the datasets from the five folds into a single dataset, and an additional column, "model," was added to specify the originating model for each record. We identified each amino acid sequence as a unique product and divided the samples into four groups based on their combination of predicted and labeled classes. Our first group has samples predicted and labeled as sensitive (true negatives), and our second group has those with a resistance match (true positives). For our third group, we selected samples labeled as sensitive but for which our model predicted resistance (false positives). The fourth group comprises those predicted to be sensitive but labeled as resistant (false negatives).

The frequency analysis aimed to compute the frequency of each product's appearance across the models. We first analyze these frequencies within each group separately across all products, and then compute the global top 20 for better readability. We also analyzed the top 20 most frequent products in each model, in each group independently, and in each combination of model and group.

After frequency analysis, we proceeded to the attention coefficient analysis. We computed the average coefficient reported by the model across all samples for each of the top 20 globally expressed products, within each model and group. Finally, we computed performance metrics regarding the matches and mismatches in each model. We computed the accuracy, precision, recall, and F1 score. We computed accuracy as the ratio of the total number of matches (groups 1 and 2) to the total number of samples (all

groups). We calculated the precision as the ratio of resistance matches (group 2) to the total number of samples predicted as resistance (groups 2 and 3). The recall is the ratio of resistance matches (group 2) by the total number of samples labeled as positive (groups 2 and 4). Finally, we computed the F1 score as in Equation 2.

## SUPPLEMENTARY INFORMATION

The file *Supplementary_text.pdf* contains additional sections, figures, and tables. The Spreadsheet *Attention_weights_test_set_1.xlsx* contains complete attention weight analysis results for test set 1. The Spreadsheet *Attention_weights_test_set_2.xlsx* contains complete attention weight analysis results for test set 2.

## DATA AND CODE AVAILABILITY

DeepMDC source code and instructions are available at https://github.com/CABGenOrg/deepmdc. The training dataset and attention spreadsheets (supplementary spreadsheets 1 and 2) are available at https://github.com/CABGenOrg/deepmdc_mem_data. The Test Set II samples are available under the BioProject: PRJNA1356496 at https://www.ncbi.nlm.nih.gov/.

## ACKNOWLEDGEMENTS

We acknowledge the INOVA FIOCRUZ program (VPPCB-002-FIO-20-2-23) for financial support. We thank CNPq for funding Amazon AWS access (CNPq/AWS 421828/2022-6). We also acknowledge

FAPERJ for financial support (MEDICINA DE PRECISÃO - E_20/2021). We would also like to thank Bruna Jesuíno for her assistance in organizing and processing the training samples.

**REFERENCES**

Allen HK, Donato J, Wang HH, Cloud-Hansen KA, Davies J, Handelsman J. Call of the wild: antibiotic resistance genes in natural environments. Nat Rev Microbiol. 2010;8(4):251–9.

Aytan-Aktug, D., Clausen, P.T.L.C., Bortolaia, V., Aarestrup, F.M., Lund, O.: Prediction of acquired antimicrobial resistance for multiple bacterial species using neural networks. mSystems 5 (2020) https://doi.org/10.1128/msystems.00774-19

Bengtsson-Palme J, Kristiansson E, Larsson DGJ. Environmental factors influencing the development and spread of antibiotic resistance. FEMS Microbiol Rev. 2018;42(1):fux053.

Borer, A. et al. Attributable mortality rate for carbapenem-resistant *Klebsiella pneumoniae* bacteremia. Infect. Control Hosp. Epidemiol. 30, 972–976 (2009).

Botelho J, Schulenburg H. The role of integrative and conjugative elements in antibiotic resistance evolution. Trends Microbiol. 2021;29(9):849–861.

Budia-Silva, M., Kostyanev, T., Ayala-Montaño, S. et al. International and regional spread of carbapenem-resistant *Klebsiella pneumoniae* in Europe. Nat Commun 15, 5092 (2024). https://doi.org/10.1038/s41467-024-49349-z.

Chen, L., Mathema, B., Chavda, K. D., DeLeo, F. R., Bonomo, R. A., & Kreiswirth, B. N. (2014). Carbapenemase-producing *Klebsiella pneumoniae*: molecular and genetic decoding. Trends in Microbiology, 22(12), 686–696. https://doi.org/10.1016/j.tim.2014.09.003

Colavecchio A, Cadieux B, Lo A, Goodridge LD. Bacteriophages contribute to the spread of antibiotic resistance genes among foodborne pathogens of the *Enterobacteriaceae* family — a narrative review. Front Microbiol. 2017;8:1108.

Cuomo, S., Di Cola, V. S., Giampaolo, F., Rozza, G., Raissi, M., & Piccialli, F. (2022). Scientific machine learning through physics–informed neural networks: Where we are and what's next. Journal of Scientific Computing, 92(3), 88.

Cuzon, G., Naas, T., & Nordmann, P. (2011). Functional characterization of Tn*4401*, a Tn*3*-based transposon involved in $bla_{\mathrm{KPC}}$ gene mobilization. Antimicrobial Agents and Chemotherapy, 55(11), 5370–5373. https://doi.org/10.1128/AAC.05202-11

Dalla-Torre, H., Gonzalez, L., Mendoza-Revilla, J., Lopez Carranza, N., Grzywaczewski, A. H., Oteri, F., ... & Pierrot, T. (2025). Nucleotide Transformer: building and evaluating robust foundation models for human genomics. Nature Methods, 22(2), 287-297.

David, S.; et al. Widespread Emergence of *OmpK36* Loop 3 Insertions among Multidrug-Resistant *Klebsiella pneumoniae*. PLoS Pathog. 2022, 18, e1010334.

Davies, J., & Davies, D. (2010). Origins and evolution of antibiotic resistance. Microbiology and Molecular Biology Reviews, 74(3), 417–433.

De Oliveira, D. M., Forde, B. M., Kidd, T. J., Harris, P. N., Schembri, M. A., Beatson, S. A., ... & Walker, M. J. (2020). Antimicrobial resistance in ESKAPE pathogens. *Clinical microbiology reviews*, *33*(3), 10-1128.

Delcour AH. Outer membrane permeability and antibiotic resistance. Biochim Biophys Acta. 2009;1794(5):808–816. doi:10.1016/j.bbapap.2008.11.005

Dey S, Shahrear S, Afroj Zinnia M, Tajwar A, Islam ABMMK. Functional Annotation of Hypothetical Proteins From the Enterobacter cloacae B13 Strain and Its Association With Pathogenicity. Bioinform Biol Insights. 2022 Aug 6;16:11779322221115535.

Doménech-Sánchez, A., Hernández-Allés, S., Martínez-Martínez, L., Benedi, V. J., & Albertí, S. (1999). Identification and characterization of a new porin gene of *Klebsiella pneumoniae*: its role in β-lactam antibiotic resistance. Journal of Bacteriology, 181(9), 2726–2732.

Duré, F. M., Silveira, M. C., Rocha-de-Souza, C. M., Leão, R. S., de Oliveira Santos, I. C., Albano, R. M., ... & da Silva, F. A. B. (2022). CABGen: A web application for the Bioinformatic analysis of bacterial genomes. Frontiers in Microbiology, 13, 893474.

Fernández L, Hancock RE. Adaptive and mutational resistance: role of porins and efflux pumps in drug resistance. Clin Microbiol Rev. 2012;25(4):661–681. doi:10.1128/CMR.00043-12

Forsberg KJ, Reyes A, Wang B, Selleck EM, Sommer MOA, Dantas G. The shared antibiotic resistome of soil bacteria and human pathogens. Science. 2012;337(6098):1107–11.

Garai, P., & Blanc‐Potard, A. (2020). Uncovering small membrane proteins in pathogenic bacteria: Regulatory functions and therapeutic potential. Molecular Microbiology, 114(5), 710-720.

Gibson MK, Forsberg KJ, Dantas G. Improved annotation of antibiotic resistance determinants reveals microbial resistomes cluster by ecology. ISME J. 2015;9(1):207–16.

Gordils-Valentin, L., Ouyang, H., Qian, L. et al. Conjugative type IV secretion systems enable bacterial antagonism that operates independently of plasmid transfer. Commun Biol 7, 499 (2024). https://doi.org/10.1038/s42003-024-06192-8

Gray, T., Storz, G., & Papenfort, K. (2022). Small proteins; big questions. Journal of Bacteriology, 204(1), e00341-21.

Hernández-Allés, S., et al. (1999). Porin expression in clinical isolates of *Klebsiella pneumoniae*. Microbiology, 145, 673–679

Holt, K. E., Wertheim, H., Zadoks, R. N., Baker, S., Whitehouse, C. A., Dance, D., ... & Thomson, N. R. (2015). Genomic analysis of diversity, population structure, virulence, and antimicrobial resistance in *Klebsiella pneumoniae*, an urgent threat to public health. Proceedings of the National Academy of Sciences, 112(27), E3574-E3581.

Hou, L., Samaras, D., Kurc, T.M., Gao, Y., Davis, J.E., Saltz, J.H.: Patch-based convolutional neural network for whole slide tissue image classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2424–2433 (2016)

Hu, J.Y.-C., Yang, D., Wu, D., Xu, C., Chen, B.-Y., Liu, H.: On sparse modern Hopfield model. Advances in neural information processing systems 36, 27594–27608 (2023)

Hu, K., et al. Assessing computational predictions of antimicrobial resistance. Brief Bioinform. 2024 27;25(3):bbae206.

Khitun, A., Ness, T. J., & Slavoff, S. A. (2019). Small open reading frames and cellular stress responses. Molecular omics, 15(2), 108-116.

Kingma, D. P., Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

Knopp, M., Gudmundsdottir, J. S., Nilsson, T., König, F., Warsi, O., Rajer, F., ... & Andersson, D. I. (2019). *De novo* emergence of peptides that confer antibiotic resistance. MBio, 10(3), 10-1128.

Kraus, O.Z., Ba, J.L., Frey, B.J.: Classifying and segmenting microscopy images with deep multiple instance learning. Bioinformatics 32(12), 52–59 (2016)

Lam, M. M. C., Wick, R. R., Watts, S. C., Cerdeira, L. T., Wyres, K. L., & Holt, K. E. (2021). A genomic surveillance framework and genotyping tool for *Klebsiella pneumoniae* and its related species complex. Nature Communications, 12(1), 4188.

Lees, J. A., Mai, T. T., Galardini, M., Wheeler, N. E., Horsfield, S. T., Parkhill, J., & Corander, J. (2020). Improved prediction of bacterial genotype-phenotype associations using interpretable pangenome-spanning regressions. MBio, 11(4), 10-1128.

Lehner A, et al. Integrating genomic features into machine learning models improves prediction of antimicrobial resistance. Nat Commun. 2023;14:5175. doi:10.1038/s41467-023-40814-7

Li, B., Li, Y., Eliceiri, K.W.: Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14318–14328 (2021)

Li, X., Ma, Q., Xu, X., et al. (2023). Hidden resistome: the role of non-expressed antibiotic resistance genes in bacterial adaptation. Frontiers in Microbiology, 14, 1213645. PMID: 37516732

Lopatkin AJ, Meredith HR, Srimani JK, Pfeiffer C, Durrett R, You L. Persistence and reversal of plasmid-mediated antibiotic resistance. Nat Commun. 2016;7:13304.

Loyola-Cruz, M. Á., Gonzalez-Avila, L. U., Martínez-Trejo, A., Saldaña-Padilla, A., Hernández-Cortez, C., Bello-López, J. M., & Castro-Escarpulli, G. (2023). ESKAPE and beyond: the burden of coinfections in the COVID-19 pandemic. Pathogens, 12(5), 743.

Martínez JL, Baquero F. Emergence and spread of antibiotic resistance: setting a parameter space. Ups J Med Sci. 2014 May;119(2):68-77. doi: 10.3109/03009734.2014.901444. Epub 2014 Mar 30. PMID: 24678768; PMCID: PMC4034563.

Martiny AC, Martiny JBH, Weihe C, Field A, Ellis JC. Functional metagenomics reveals previously unrecognized diversity of antibiotic resistance genes in coastal marine sediments. Environ Microbiol. 2011;13(5):1108–23.

Mathers, A. J., Peirano, G., & Pitout, J. D. (2015). The role of epidemic resistance plasmids and international high-risk clones in the spread of multidrug-resistant *Enterobacteriaceae*. Clinical Microbiology Reviews, 28(3), 565–591. https://doi.org/10.1128/CMR.00116-14

Meekes, L.M.; et al. Proteogenomic Analysis Demonstrates Increased $bla_{\text{OXA-48}}$ Copy Numbers and *OmpK36* Loss as Contributors to Carbapenem Resistance in *Klebsiella pneumoniae*. Antimicrob. Agents Chemother. 2025, 69(7): e00107-25.

Naas, T., Cuzon, G., Villegas, M. V., Lartigue, M. F., Quinn, J. P., & Nordmann, P. (2008). Genetic structures at the origin of acquisition of the $bla_{\text{KPC}}$ gene. Antimicrobial Agents and Chemotherapy, 52(4), 1257–1263. https://doi.org/10.1128/AAC.01451-07

Nikaido H. Multidrug efflux pumps of Gram-negative bacteria. Front Microbiol. 2011;2:13. doi:10.3389/fmicb.2011.00013

Oliveira, P. H., Touchon, M., & Rocha, E. P. C. (2020). Regulation of genetic flux between bacteria by restriction–modification systems. Proceedings of the National Academy of Sciences (PNAS), 117(12), 6611–6619.

Palomba, E., Comelli, A., Saluzzo, F. et al. Activity of imipenem/relebactam against KPC-producing *Klebsiella pneumoniae* and the possible role of *OmpK36* mutation in determining resistance: an Italian retrospective analysis. Ann Clin Microbiol Antimicrob. 2025, 24(1):23.

Partridge, S. R., et al. (2018). Mobile genetic elements associated with antimicrobial resistance. Clinical Microbiology Reviews, 31(4), e00088–17.

Peng, Z ., Maciel-Guerra, A., Baker, M., Zhang, X., Hu, Y., Wang, W., Rong, J.,Zhang, J., Xue, N., Barrow, P., Renney, D., Stekel, D., Williams, P., Liu, L., Chen, J., Li, F., Dottorini, T.: Whole-genome sequencing and gene sharing network analysis powered by machine learning identifies antibiotic resistance sharing between animals, humans, and environment in livestock farming. PLOS Computational Biology 18, 1010018 (2022) https://doi.org/10.1371/journal.pcbi.1010018

Perron GG, Whyte L, Turnbaugh PJ, Goordial J, Hanage WP, Dantas G, Desai MM. Functional characterization of bacteria isolated from ancient arctic soil exposes diverse resistance mechanisms to modern antibiotics. PLoS One. 2008;3(3):e1758.

Prins, R. C., & Billerbeck, S. (2025). Small proteins and peptides conferring protection against antimicrobial compounds. Trends in Microbiology.

Rahman, M.A., Rangwala, H.: Idmil: an alignment-free interpretable deep multiple instance learning (mil) for predicting disease from whole-metagenomic data. Bioinformatics 36(Supplement 1), 39–47 (2020)

Ramsauer, H., Schäfl, B., Lehner, J., Seidl, P., Widrich, M., Adler, T., Gruber, L., Holzleitner, M., Pavlović, M., Sandve, G.K., et al.: Hopfield networks is all you need. arXiv preprint arXiv:2008.02217 (2020)

Roberts, R. J., Vincze, T., Posfai, J., & Macelis, D. (2015). REBASE—a database for DNA restriction and modification: enzymes, genes and genomes. Nucleic Acids Research, 43(D1), D298–D299.

Rodríguez-Beltrán, J., DelaFuente, J., León-Sampedro, R., MacLean, R. C., & San Millán, Á. (2021). Beyond horizontal gene transfer: The role of plasmids in bacterial evolution. Nature Reviews Microbiology, 19(6), 347–359.

Rodríguez-Martínez, J. M., Machuca, J., Cano, M. E., Calvo, J., & Pascual, Á. (2021). Mechanisms of intrinsic and acquired resistance to β-lactams in Gram-negative bacteria. Current Drug Targets, 22(7), 646–662.

Schwengers, O., Jelonek, L., Dieckmann, M. A., Beyvers, S., Blom, J., & Goesmann, A. (2021). Bakta: rapid and standardized annotation of bacterial genomes via alignment-free sequence identification. Microbial genomics, 7(11), 000685.

Shintani M, Sanchez ZK, Kimbara K. Genomics of microbial plasmids: classification and identification based on replication and transfer systems and host taxonomy. Front Microbiol. 2015 Mar 31;6:242.

Smillie C, Garcillán-Barcia MP, Francia MV, Rocha EPC, de la Cruz F. Mobility of plasmids. Microbiol Mol Biol Rev. 2010;74(3):434–452.

Strahilevitz, J.; et al. In Vivo Selection of Carbapenem Resistance during Persistent *Klebsiella pneumoniae* Infection under Antimicrobial Therapy. Antimicrob. Agents Chemother. 2024, 68, e00663-24.

Su, M. Et al. Genome-Based Prediction of Bacterial Antibiotic Resistance. Journal of Clinical Microbiology, 2019. J Clin Microbiol57:10.1128/jcm.01405-18.

Sugawara, E., Kojima, S., Nikaido, H. (2016). *Klebsiella pneumoniae* major porins *OmpK35* and *OmpK36* allow more efficient diffusion of β-lactams than their Escherichia coli homologs *OmpF* and *OmpC*. Journal of Bacteriology, 198(23), 3200–3208.

Touchon M, Moura de Sousa JA, Rocha EP. Embracing the enemy: the diversification of microbial gene repertoires by phage-mediated horizontal gene transfer. Curr Opin Microbiol. 2017 Aug;38:66-73.

Tsai, Y.-K.; et al. Klebsiella pneumoniae Outer Membrane Porins *OmpK35* and *OmpK36*: Roles in Antimicrobial Resistance and Virulence. Antimicrob. Agents Chemother. 2011, 55, 1485–1493.

Wang, G., Zhao, G., Chao, X., Xie, L., & Wang, H. (2020). The characteristic of virulence, biofilm and antibiotic resistance of *Klebsiella pneumoniae*. International journal of environmental research and public health, 17(17), 6278.

Wick, R. R., Judd, L. M., Gorrie, C. L., & Holt, K. E. (2017). Unicycler: resolving bacterial genome assemblies from short and long sequencing reads. PLoS computational biology, 13(6), e1005595.

Widrich, M., Sch¨afl, B., Pavlovi´c, M., Ramsauer, H., Gruber, L., Holzleitner, M., Brandstetter, J., Sandve, G.K., Greiff, V., Hochreiter, S., et al.: Modern Hopfield Networks and attention for immune repertoire classification. Advances in neural information processing systems 33, 18832–18845 (2020)

Wong, J.L.C.; et al. *OmpK36*-Mediated Carbapenem Resistance Attenuates *Klebsiella pneumoniae* Virulence and Transmissibility. Nat. Commun. 2019, 10, 3957.

Wong, J. L., David, S., Sanchez-Garrido, J., Woo, J. Z., Low, W. W., Morecchiato, F., ... & Frankel, G. (2022). Recurrent emergence of Klebsiella pneumoniae carbapenem resistance mediated by an inhibitory *OmpK36* mRNA secondary structure. Proceedings of the National Academy of Sciences, 119(38), e2203593119.

Wright GD. The antibiotic resistome: the nexus of chemical and genetic diversity. Nat Rev Microbiol. 2010;8(12):959–68.

Wysocka, M., Wysocki, O., Zufferey, M., Landers, D., & Freitas, A. (2023). A systematic review of biologically-informed deep learning models for cancer: fundamental trends for encoding and interpreting oncology data. BMC bioinformatics, 24(1), 198.

Wyres, K. L., Lam, M. M., & Holt, K. E. (2020). Population genomics of *Klebsiella pneumoniae*. Nature Reviews Microbiology, 18(6), 344-359.

Wyres KL, Holt KE. *Klebsiella pneumoniae* as a key trafficker of drug resistance genes from environmental to clinically important bacteria. Curr Opin Microbiol. 2018 Oct;45:131-139. doi: 10.1016/j.mib.2018.04.004. Epub 2018 May 1. PMID: 29723841.

Zhang Y, et al. The plasmid replication initiator RepB contributes to copy number control and stability, affecting antibiotic resistance gene expression. Front Microbiol. 2024;15:1398124. doi:10.3389/fmicb.2024.1398124

Zhou, Z., Ji, Y., Li, W., Dutta, P., Davuluri, R., & Liu, H. (2023). Dnabert-2: Efficient foundation model and benchmark for multi-species genomes. arXiv preprint arXiv:2306.15006.