

AlphaPeptDeep: A modular deep learning framework to predict peptide properties for proteomics

Wen-Feng Zeng¹, Xie-Xuan Zhou¹, Sander Willems¹, Constantin Ammar¹, Maria Wahle¹,
Isabell Bludau¹, Eugenia Voytik¹, Maximillian T. Strauss², Matthias Mann^{1,2,*}

1. Department of Proteomics and Signal Transduction, Max Planck Institute of Biochemistry, Martinsried, Germany

2. Proteomics Program, NNF Center for Protein Research, Faculty of Health Sciences, University of Copenhagen, Copenhagen, Denmark

* E-mail: mmann@biochem.mpg.de

Abstract

Machine learning and in particular deep learning (DL) are increasingly important in mass spectrometry (MS)-based proteomics. Recent DL models can predict the retention time, ion mobility and fragment intensities of a peptide just from the amino acid sequence with good accuracy. However, DL is a very rapidly developing field with new neural network architectures frequently appearing, which are challenging to incorporate for proteomics researchers. Here we introduce AlphaPeptDeep, a modular Python framework built on the PyTorch DL library that learns and predicts the properties of peptides (<https://github.com/MannLabs/alphapeptdeep>). It features a model shop that enables non-specialists to create models in just a few lines of code. AlphaPeptDeep represents post-translational modifications in a generic manner, even if only the chemical composition is known. Extensive use of transfer learning obviates the need for large data sets to refine models for particular experimental conditions. The AlphaPeptDeep models for predicting retention time, collisional cross sections and fragment intensities are at least on par with existing tools. Additional sequence-based properties can also be predicted by AlphaPeptDeep, as demonstrated with a novel HLA peptide prediction model to improve HLA peptide identification for data-independent acquisition.

Introduction

The aim of MS-based proteomics is to obtain an unbiased view of the identity and quantity of all the proteins in a given system^{1,2}. This challenging analytical task requires advanced liquid chromatography – mass spectrometry (LC/MS) systems as well as equally sophisticated bioinformatic analysis pipelines³. Over the last decade, machine learning (ML) and in particular deep neural network (NN)-based deep learning (DL) approaches have become very powerful and are increasingly beneficial in MS-based proteomics^{4,5}.

Identification in proteomics entails the matching of fragmentation spectra (MS2) and other

properties to a set of peptides. Bioinformatics can now predict peptide properties for any given amino acid sequences so that they can be compared to actually measured data. This can markedly increase the statistical confidence in peptide identifications.

To do this, a suitable ML/DL model needs to be chosen which is then trained on the experimental data. There are a number of peptide properties that can be predicted from the sequence and for each of them different models may be most appropriate. For the peptide retention times in LC, relatively straightforward approaches such as iRT-calculator, RTPredict, and ELUDE have shown good results⁶⁻⁸. However, large volumes of training data are readily available in public repositories today and DL models currently tend to perform best⁹. This is also the case for predicting the fragment intensities in the MS2 spectra, where DL models such as our previous model pDeep^{10,11}, DeepMass:Prism¹², Prosit¹³ and many subsequent ones now represent the state-of-the-art. They mostly use long-short term memory (LSTM¹⁴) or gated recurrent unit (GRU¹⁵) models. Recently, transformers have been adopted in proteomics and show better performance^{16,17}. This illustrates the rapid pace of advance in DL and the need for updating proteomics analysis pipelines with them. However, the focus of existing efforts has not been on extensibility or modularity, making it difficult or in some cases impossible to change or extend their NN architectures.

Here we set out to address this limitation by creating a comprehensive and easy to use framework, termed AlphaPeptDeep. As part of the AlphaPept ecosystem¹⁸, we keep its principles of open source, community orientation as well as robustness and high performance. Apart from Python and its scientific stack, we decided to use PyTorch,¹⁹ one of the most popular DL libraries.

AlphaPeptDeep contains pre-trained models for predicting MS2 intensities, retention time (RT), and collisional cross sections (CCS) of arbitrary peptide sequences or entire proteomes. It also handles peptides containing post-translational modifications (PTMs), including unknown ones with user-specified chemical compositions. AlphaPeptDeep makes extensive use of transfer learning, drastically reducing the amount of training data required.

In this paper, we describe the design and use of AlphaPeptDeep and we benchmark its performance for predicting MS2 intensities, RT, and CCS on peptides with or without PTMs. On challenging samples like HLA peptides, AlphaPeptDeep dramatically boosts performance of peptide identification for data-dependent acquisition. We also describe how AlphaPeptDeep can easily be applied to build and train models for different peptide properties such as an HLA prediction model, which narrows the database search space for data-independent acquisition, and hence improves the identification of HLA peptides with the AlphaPeptDeep-predicted spectral library.

Results

AlphaPeptDeep overview and model training

For any given set of peptide properties that depend on their sequences, the goal of the AlphaPeptDeep framework is to enable easy building and training of deep learning (DL) models, that achieve high performance given sufficient training data (Fig. 1a). Although modern DL libraries are more straightforward to use than before, designing a neural network (NN) or developing a deployable DL model for proteomics studies is not as simple as it could be, even for biologists with programming experience. This is because of the required domain knowledge and the complexity of the different steps involved in building a DL model. The framework of AlphaPeptDeep is designed to address these issues (Fig 1b).

The first challenge is the embedding, which maps amino acid sequences and their associated PTMs into a numeric tensor space that the NN needs as an input. For each amino acid, a ‘one-hot encoder’ is customarily used to convert it into a 27-length fingerprint vector consisting of 0s and 1s (Online Methods). In contrast, PTM embedding is not as simple. Although recent studies also used one-hot encoding to embed phosphorylation for MS2 prediction via three additional amino acids¹⁶, this is not extendable to arbitrary PTMs. In pDeep2 (ref¹¹), the numbers of C, H, N, O, S, P atoms for a site-specific modification are prepended to the embedding vector which is flexible and can be applied to many different PTMs. AlphaPeptDeep inherits this feature from pDeep2 but adds the ability to embed all the other chemical elements. To make the input space manageable, we use a linear NN that reduces the size of the input vector for each PTM (Online Methods, Extended Fig. 1). This allows efficient embedding for most modification types, except for very complex ones such as glycans. The PTM embedding can be called directly from AlphaPeptDeep building blocks.

To build a new model, AlphaPeptDeep provides modular application programming interfaces (APIs) to use different NN architectures. Common ones like LSTM, convolutional NN (CNN) as well as many others are readily available from the underlying PyTorch library. Recently transformers – attention-based architectures to handle long sequences – have achieved breakthrough results in language processing but were then also found to be applicable to many other areas like image analysis²⁰, gene expression²¹ and protein folding²². AlphaPeptDeep includes a state-of-the-art HuggingFace transformer library²³. Our framework also easily allows combining different NN architectures for different prediction tasks.

The training and transfer learning steps are mostly generic tasks, even for different NNs. Therefore, we designed a universal training interface allowing users to train the models using just a single line of Python code – *model.train()*. In our training interface, we also provide a “warmup” training strategy to schedule the learning rate for different training epochs (Online Methods). This has proven very useful in different tasks to reduce the bias at the early training stage²⁴. Almost all DL tasks can be done on graphical processing units (GPUs) and training a model from scratch on a standard GPU usually take not more than hours in AlphaPeptDeep and is performed only once. Transfer learning from a

pre-trained model is feasible within minutes, even without GPU.

After training, all learned NN parameters should be saved for persistent applications. This can be readily done using DL library functionalities, and is also implemented in AlphaPeptDeep – *model.save()*. In the latter case, AlphaPeptDeep will save the source code of the NN architectures in addition to the training hyperparameters. Thus, the NN code and the whole training snapshot can be recovered even if the source code was accidentally changed in the AlphaPeptDeep or developers' codebase. This is especially useful for dynamic computational graph-based DL libraries such as PyTorch and TensorFlow in 'eager mode' because they allow dynamically changing the NN architectures.

The most essential functionality of the AlphaPeptDeep framework is the prediction of a property of a given peptide of interest. When using only the CPU, one can choose multiprocessing (predicting with multiple CPU cores), making the prediction speed acceptable on regular personal computers (PCs) and laptops (nearly 2h for the entire reviewed human proteome). Prediction on GPU is still an order of magnitude faster. As PyTorch caches the GPU RAM in the first prediction batch, subsequent batches for the same model will be even faster. However, GPU random access memory (RAM) should be released after the prediction stage, thus making the RAM available for other DL models. These steps are automatically done in AlphaPeptDeep within the *model.predict()* functionality.

AlphaPeptDeep provides several templates in the "model shop" module to develop new DL models from scratch for classification or regression with very little code. All these high-level functionalities in AlphaPeptDeep give the user a quick on-ramp and they minimize the effort needed to build, train and use the models. As an illustrative example, we built a classifier to predict if a peptide elutes in the first or second half of the LC gradient using only several lines of code. Training took only ~16 minutes on nearly 350K peptide-spectrum matches (PSMs) on a standard HeLa dataset²⁵ and the model achieved 95% accuracy in the testing set (Extended Fig. 2).




Figure 1. Overview of the AlphaPeptDeep framework. (a) Measured peptide properties are encoded with the respective amino acid sequences and used to train a network in AlphaPeptDeep (left). Once a model is trained, it can be used on arbitrary sets of peptide sequences to predict the property of interest. This then improves the sensitivity and accuracy of peptide identification. (b) The AlphaPeptDeep framework reads and embeds the peptide sequences of interest. Its components include the build functionality in which the model can build. It is then trained, saved and used to predict the property of interest. The dial represents the different standard properties that can be predicted (RT, retention time; CCS, collision section; MS2, intensities of fragment spectra). Custom refers to any other peptide property of interest. The lower part lists aspects of the functionalities in more detail.

The MS2, RT, and CCS prediction models in AlphaPeptDeep are all publicly available in

our Python modules (Fig. 2). The MS2 prediction model was inherited from pDeep2 but reimplemented on transformers which have been shown very useful in MS2 prediction^{16,17}. The pre-trained MS2 model in AlphaPeptDeep is much smaller than other models without sacrificing accuracy (4M parameters vs 64M in the Prosit-Transformer¹⁷), making the prediction extremely fast (Extended Fig. 3). We also applied the same principle of light-weight models to our RT and CCS models (less than 1M parameters for each, Online Methods), which we built on previous LSTM models²⁵⁻²⁷.

We trained and tested the MS2 models with tens of millions of spectra from a variety of instruments, collision energies and peptides, and trained the RT and CCS models with about half a million RT and CCS values of peptides (Suppl. Data 1). The results of this initial training were then stored as pre-trained models for further use or as a basis for refinement with transfer learning.

Using these pre-trained models and specifically designed data structures (Online Methods), the prediction of a spectral library with MS2 intensities, RT, and ion mobilities (converted from CCS, Online Methods) for the human proteome took only 10 min on a regular GPU and 100 minutes on the CPU with multiprocessing (Extended Fig. 3). As this prediction only needs to be done at most once per project, we conclude that the prediction of libraries by DL is not a limitation in data analysis workflows.




Figure 2. The built-in and pre-trained MS2, RT, and CCS prediction models. The MS2 model is built on four transformer layers, and the RT/CCS models consist of a convolutional neural network (CNN) layer followed by two bidirectional long short-term memory (BiLSTM) layers. The pre-trained MS2 model currently supports predicting the intensities of backbone b/y ions as well as their modification-associated neutral losses if any (e.g. -98 Da loss of phosphorylation on Ser/Thr). However, the user can easily configure the MS2 model to train and predict water and ammonium losses from backbone fragments as well.

Prediction performance of the AlphaPeptDeep model for MS2 spectra

With the AlphaPeptDeep framework for prediction of MS2 intensities, RT and CCS in hand, we first benchmarked the MS2 model against datasets of tryptic peptides (phase 1 in Fig. 3a). The training and testing data were collected from various instruments and collisional energies (Suppl. Data 1), including ProteomeTools²⁸, which were derived from synthetic peptides with known ground truth. We split the data sets in two and trained on a LSTM model similar to pDeep or on the new transformer model. As expected, transformer performed better than the LSTM model on the test datasets (Extended Fig. 4). Overall, on ProteomeTools, 97% of all significantly matching PSMs had Pearson correlation coefficients (PCC) of the predicted vs. the measured fragment intensities of at least 90% (Fig. 3a), which we term ‘PCC90’ in this manuscript. Note that the experimental replicates also exhibit some variation, making the best possible prediction accuracy somewhat less than 100%. On the ProteomeTools replicates measured with the Lumos mass spectrometer, 99% had PCCs above 90% (Suppl. Data 1), meaning that our predicted intensities mirrored the measured ones almost within experimental uncertainties (99% experimental vs. 97% predicted). Next, we tested the model on the same ProteomeTools sample but measured on a trapped ion mobility Time of Flight mass spectrometer (timsTOF) in dda-PASEF mode^{25,29}, and achieved a PCC90 of 87.9% (Suppl. Data 1), showing that the prediction from the pre-trained model is already very good for timsTOF even without adaption.

As expected, our pretrained model performed equally well across different organisms. Interestingly, it did almost as well on chymotrypsin or GluC-digested peptides although it had not been trained on them (Fig. 3a).

HLA class 1 peptides are short pieces of cellular proteins (about 9 amino acids) that are presented to the immune system at the cell surface, which is of great interest to biomedicine³⁰. Because of their low abundance and non-tryptic nature, they are very challenging to identify by standard computational workflow, a task in which DL can help³¹. In a second training phase, we added a synthetic HLA dataset, which was also from ProteomeTools³², into the training set and trained the model for additional 20 epochs (‘fine tuning the model’). We first checked if the new model negatively impacted performance on the tryptic data sets, but this turned out not to be the case (phase 2 in Fig. 3a). On the HLA peptides, however, performance substantially increased the PCC90 from 79% to 92%.

Finally, we extended our model to predict phosphorylated and ubiquitylated peptides, which have spectra somewhat distinct from unmodified peptides. In this case, in addition to backbone fragmentation intensities, AlphaPeptDeep also needs to learn the intensities of fragments with or without modifications. For phosphopeptide prediction, performance of the pre-trained model was much worse, with PCC90 values of only around 30%. However, after training on PTM datasets at phase 3, the performance dramatically increased, almost to the level of tryptic peptides (Fig. 3a). The ubiquitylation prediction (rightmost in Fig. 3a) was already reasonable with the pre-trained model but increased further after phase 3 training (PCC90 from 75% to 93%).

Prediction performance of the AlphaPeptDeep models for RT and CCS

RT and CCS models are quite similar to each other as their inputs are the peptide sequences and PTMs, and outputs are scalar values. For both we used LSTM architectures. In the CCS prediction model, precursor charge states are considered in the model as well. Taking advantage of the PTM embedding in AlphaPeptDeep, the RT and CCS models naturally consider PTM information, and hence can predict peptide properties given arbitrary PTMs. We trained the RT model on datasets with regular peptides (unmodified and Met-oxidated) from our HeLa measurements²⁵.

We first tested the trained RT model on regular peptides from the pan human library³³. As shown in Fig. 3b, the pre-trained model gave very good predictions in most of the RT range, but failed to accurately predict the last few minutes (iRT (ref⁷) values higher than 100) possibly due to the different flushing settings of the LC in training and testing data. These differences could be addressed by fine-tuning the model with experiment-specific samples. Few-shot fine-tuning with only 500 training samples improved the accuracies of the RT prediction from an R^2 of 0.927 to 0.986.

We also tested the RT model on a phosphopeptide dataset,³⁴ although the model had not been trained on such data. After fine-tuning on 500 peptides, the R^2 increased from 0.958 to 0.984 (Fig. 3b). As RT behavior of peptides varies with the LC conditions in different experiments, we highly recommend fine-tuning whenever possible. It turns out that few-shot fine-tuning worked well to fit short LC conditions as well (Extended Fig. 5). Finally, as expected, the more training peptides we used, the better the fine-tuning, and with many peptides our model reached R^2 values up to 0.99 (Fig. 3b).

While the CCS model was trained on regular human peptides from our prior HeLa dataset²⁵ (Suppl. Data. 1), we tested the trained model on *E. coli* and yeast peptides from the same instrument in the same publication. For these regular peptides the CCS model achieved an $R^2 > 0.98$ of the predicted and detected CCS values. Next, we searched the HeLa and drosophila data by the open-search mode in Open-pFind³⁵, to obtain experimental CCS values for modified peptides (Online Methods). For these peptides in the testing set, the R^2 was 0.965 and 0.953, respectively, a prediction accuracy quite close to the one for regular peptides, even for unexpected modifications. The predicted CCS values can be converted to ion mobilities of Bruker timsTOF using the Mason Schamp equation.³⁶




Figure 3. The performance of MS2/RT/CCS models. (a) The MS2 prediction accuracies of the three training phases on different testing datasets. The performance is evaluated by “PCC90” (percentage of PCC values larger than 0.9). The prefix ‘PT’ of each data set refers to ProteomeTools. PT1 and PT2 refer to ProteomeTools part I and II, respectively. The black bars are the PCC90 values of experimental spectra. (b) For RT prediction, few-shot learning can correct the RT bias between different LC conditions. (c) Our CCS model works well for both regular (top panels) and unexpectedly modified (bottom panels) peptides.

Prediction performance for 21 PTMs with transfer learning

To further demonstrate the powerful and flexible support for PTMs in AlphaPeptDeep, we tested the pre-trained tryptic MS2 (phase 1 in Fig. 3a) and RT models using the 21 PTMs, which were synthesized based on 200 template peptide sequences³⁷.

Interestingly, there is a group of modifications for which the prediction of MS2 spectra is as good as the values of unmodified peptides (Fig. 4a). These include Hydroxypro@P, Methyl@R, and Dimethyl@R for which the PCC90 was greater than 80%. This is presumably because these modifications do not change the overall fragmentation pattern much. In contrast, most of the other PTMs cannot be well predicted by the pre-trained models, for example, the PCC90 values were less than 10% for Malonyl@K and Citrullin@R. Remarkably, transfer learning for each PTM type using as few as ten peptides with different charge states and collisional energies greatly improved the

prediction accuracies on the testing data. The largest improvements of PCC90 were as high as 60% (Citrullin@R and Malonyl@K, Fig. 4a). Overall, compared with the pre-trained model, the ones tuned by ten peptides improved the PCC90 from a median of 48% to 87% (Fig. 4b). We speculate that this is because the fragmentation properties of amino acids at different collisional energies have been well learned by the pre-trained model after which transfer learning only needs to learn the properties of modified ones. Including 50 PTM bearing peptides improved this number to 93% whereas using 80% of all the identified peptides ($n \leq 200$) with these PTMs only improved prediction by another 2%. This demonstrates that our models can be adapted to novel situations with very little additional data, due to the power of transfer learning.

AlphaPeptDeep has been included in AlphaViz³⁸, a tool suite for RAW MS data visualization (<https://github.com/MannLabs/alphaviz>), which among other features allows users to visualize a mirrored plot between experimental and predicted spectra. As an example, the MS2 prediction of the peptide “AGPNASIISLKS DK-Biotin@K11” before and after transfer learning is displayed in Fig. 4c. The y12++ ion was first wrongly predicted by the pre-trained model, but this was corrected after transfer learning with only 50 other biotinylated peptides. AlphaPeptDeep also allows users to visualize the ‘attention’ weights— a key feature of transformer models – showing what data attributes were important for the prediction. To depict the attention changes between pre-trained and transfer learning transformer models, we used the BertViz package (<https://github.com/jessevig/bertviz>) (Extended Fig. 6).

Next, we tested the performance of our pre-trained RT model using the datasets of 21 PTMs. Although the model was never trained on any of these PTMs, the accuracy of RT prediction on these peptides exceeds that of DeepLC³⁹, an RT prediction model designed for unseen PTMs (R^2 of 0.95 of AlphaPeptDeep vs. 0.89 of DeepLC, Fig. 4d and 4e). In this case, transfer learning only slightly improves the results, presumably because some of these synthetic modified peptides elute in very broad peaks, which makes them hard to predict.




Figure 4. Model performance with transfer learning on 21 PTMs from ProteomeTools. (a) The accuracy of MS2 prediction with different numbers of peptides for transfer learning for each PTM. Each PTM is tested separately. “80% seqs” refers to using 80% of the identified modified sequences for transfer learning. (b) Overall accuracy without unmodified peptides from (a). (c) Transfer learning dramatically improves the MS2 prediction of the example peptide “AGPNASIISLSDK-Biotin@K11” (tuned by 50 other peptides). (d) Comparisons of RT prediction for each PTM on pre-trained and transfer learning (by 50% of all the identified peptides) models, as well as DeepLC models. (e) Overall R^2 distribution without unmodified peptides from (d).

Boosting Data-Dependent Acquisition (DDA) identification of HLA peptides

As explained above, HLA peptides are among the most challenging samples for MS-based proteomics. Given the excellent model performance of the transformers in AlphaPeptDeep, we hypothesized that prediction of their MS2 spectra could substantially improve their identification.

The non-tryptic nature of these peptides results in an extremely large number of peptides that need to be searched, leading to a decreased statistical sensitivity at a given false discovery rate (FDR) level (usually 1%). The key idea of using MS2, RT and CCS

prediction to support HLA peptide identification is that, for correct peptides of the searched spectra, the predicted properties should be very close to the detected ones, while the predicted properties of the irrelevant peptides tend to be randomly distributed. Therefore, the similarities or differences between the predicted and detected properties can be used as machine learning features to distinguish correct from false identifications using semi-supervised learning. Such an approach has long been implemented in Percolator and later in other tools to re-score PSMs⁴⁰, which increases the sensitivity at the same FDR level^{31,32}. However, due to the lack of support for arbitrary PTMs with DL models it has not been for open-search of HLA peptides. Modern protein open-search engines like pFind³⁵ can perform very fast unspecific peptide search without limiting the peptide mass window using the sequence tag technique⁴¹, enabling the identification of unexpected PTMs.

AlphaPeptDeep fully supports the Percolator algorithm for regular as well as open-search of HLA peptides (Online Methods). To accelerate the rescoring, we calculate the fragment intensity similarities between predicted and detected spectra on a GPU, making the rescoring process extremely fast (~1 hour to rescore 16,812,335 PSMs from 424 MS runs using a PC with a GeForce RTX 3080 GPU, where ~60% of the time was used for loading the RAW files). This means that the rescoring by AlphaPeptDeep is not a bottleneck for HLA peptide search.

To investigate how much AlphaPeptDeep can boost the HLA peptide search, we applied it on two datasets, MSV000084172 from samples in which particular mono-allelic HLA-I types were enriched⁴², here referred to as the ‘mono-allelic dataset’ and our published dataset from tumor samples (PXD004894⁴³) referred to as the ‘tumor dataset’. These two datasets had been analyzed with a regular search engine (MaxQuant⁴⁴) by the Kuster group³² (Fig. 5a) and we here used pFind in the open-search mode (Fig. 5b).

First, we wanted to compare the AlphaPeptDeep results with MaxQuant as well as Prosit, a recently published DL based tool that has also been applied to HLA peptides³². Since Prosit only supports fixed iodoacetamide modification on alkylated peptides (IAA in Fig. 5a), we only used the results of the same IAA RAW files in rescoring. On the mono-allelic and the tumor datasets, AlphaPeptDeep covered 93% and 96% of the MaxQuant results while more than doubling the overall numbers at the same FDR of 1% (Fig. 5a). Compared to Prosit, AlphaPeptDeep captured 91% of their peptides and still improved the overall number on the mono-allelic dataset by 7%.

Next, we searched both datasets with the open-search mode of pFind (Fig. 6b), and rescored the results in AlphaPeptDeep. Here, both alkylated and non-alkylated peptides were analyzed. Interestingly, the open-search itself already identified similar numbers of peptides as the DL-boosted regular search, but AlphaPeptDeep further improved the total number of identified peptides by 29% and 42%, while retaining 99% and 98% of the pFind hits at the same FDR for the mono-allelic and tumor datasets, respectively (Fig. 5b). This demonstrates the benefits of AlphaPeptDeep’s support of open-search for HLA

peptide analysis.

AlphaPeptDeep with open-search identified PTMs such as phosphorylation, which are known to exist on HLA peptides but are very difficult to identify by regular unspecific search. For the mono-allelic dataset we identified a total of 490 phosphopeptides. To gauge the biological reasonability of these peptides, we searched for sequence motifs of both the phosphorylated and non-phosphorylated peptides. This revealed the expected HLA peptide motifs, dominated by the anchor residues for their cognate major histocompatibility complex proteins. Only the phospho-HLA peptides additionally had linear phospho-motifs, like the prominent SP motif common to proline directed kinases (Fig. 5c and Extended Fig. 7). We also identified 359 phospho-HLA peptides from the tumor dataset, with similar phospho-motifs (Extended Fig. 7). We further used AlphaPeptDeep to inspect retention time and MS2 spectrum similarities (Extended Fig. 8). Note that the MS2 and RT models were only fine-tuned by at most 100 phospho-PSMs from eight RAW files (Online Methods), so most of the phosphopeptides from other RAW files were not used in fine-tuning. Our method was also able to identify other PTMs associated with HLA peptides, such as cysteinylation⁴⁵ (Extended Fig. 9). Overall, most of the HLA peptides additionally identified by this method had modifications related to sample preparation, such as deamidation, N-terminal pyro-Glu, and N-terminal carbamidomethylation (Extended Fig. 9).




Figure 5. AlphaPeptDeep drastically improves upon regular **(a)** and open-search **(b)** results for DDA identification of HLA peptides. IAA refers to iodoacetamide alkylated peptides. **(c)** Logo plots of unmodified and phosphorylated peptides with nine amino acids identified by open-search for four different HLA types. Logo plots were generated by LogoMaker.⁴⁶

Building an HLA prediction model for HLA DIA Search

In recent years, DIA has become a method of choice to generate large-scale proteome datasets. DIA data analysis traditionally requires DDA experiments to generate a library to which the data is then matched⁴⁷. These libraries contain RT, ion mobility (if applicable)

and the most intense and specific fragments for each peptide. The generation of experimental libraries is laborious and sample consuming. With the development of DL in proteomics, libraries with predicted RT, CCS/ion mobilities and fragment intensities from whole proteome sequences are becoming more and more popular, although there is still a debate about whether measured or predicted libraries are preferable. This is because the large search space introduced by purely in silico libraries can make FDR control difficult.

DIA for HLA peptide analysis is also getting more attention^{48,49}. So far, these efforts have been restricted to experimental DIA libraries because analysis with a predicted HLA library is far more challenging than with an experimental one. This is mainly because HLA peptides are not tryptic, meaning they do not follow specific cleavage rules and do not necessarily have a favorable fragmentation pattern. The number of theoretical peptides with amino acid lengths between 8 and 14 from a reviewed human proteome is more than 70M, which is nearly two orders of magnitude more than that of tryptic peptides in the same length range (~900K). Due to this enormous search space, a predicted library is difficult or even impossible to search by state-of-the-art DIA search tools such as DIA-NN⁵⁰ and Spectronaut⁵¹.

Fortunately, HLA peptides follow certain sequence motifs guided by the HLA-types that are present. We reasoned that these motifs could be learned by DL for more efficient peptide identification. To test this hypothesis, we built an HLA prediction model using the model shop functionalities in our AlphaPeptDeep framework (Online Methods). This model - a binary LSTM classifier predicts if a given sequence is likely to be an HLA peptide presented to the immune system and extracts these peptides from the human proteome sequence. There are two main goals of the model: (1) keep as many actually presented HLA peptides as possible (i.e., high sensitivity); and (2) reduce the number of predicted peptides to a reasonable number (i.e., high specificity). Note that sensitivity is more important here as we hope that all measured HLA peptides are still in the predicted set.

Based on these goals, we developed a pipeline which enables predicted library search for DIA data (Fig. 6a). In brief, we first trained a pan-HLA prediction model with peptides from all known HLA types ('pre-trained model' in Fig. 6a). Normally, only a few HLA types are actually present in the samples from any given individual. Therefore, we used transfer learning to create a person-specific model with sample-specific peptides ('tuned model' in Fig. 6a). This model should then be able to predict whether an HLA peptide is potentially present in the sample or not, thus further reducing the number of peptides to be searched and increasing prediction accuracy. For this strategy, we need to identify a number of sample specific HLA peptides. This can be done directly from the already acquired DIA data by a 'direct-DIA search'⁵² obviating the need for a separate DDA experiment. This involves grouping eluting fragment detected peaks belonging to the same peptide signal into a pseudo-spectrum for DIA data, and then searching the pseudo-spectrum with conventional DDA search algorithms.

To test this pipeline, we used the HLA-I dataset of the RA957 cell line in PXD022950⁴⁸. We started with our pan-HLA prediction model from 94 known HLA types (Fig. 5). It reduced the number of sequences from 70M to 7M with 82% sensitivity. However, 7M peptides are still too many to search and the model would have lost 18% of true HLA peptides. Furthermore, the pre-trained model is not able to identify unknown HLA types as it is only trained on already known ones.

To enable transfer learning, we searched RA957 data with DIA-Umpire⁵². It identified 12,998 unique sequences with length from 8 to 14. We used transfer learning on 80% of this data to train the sample specific HLA model while keeping 20% for testing. This dramatically increased the specificity to 96% with 92% sensitivity (note that this is judged on the identifications by direct-DIA; thus our sensitivity may be even higher). The number of HLA peptides predicted by this model is 3M, which is comparable to the tryptic human proteome library.

Having predicted our sample-specific HLA peptides, including their MS2 fragment spectra and RTs, we used this as input for a DIA-NN search of the DIA data. Our workflow identified 36,947 unique sequences. PEAKS-Online⁵³ is a very recently published tool which combines searching a public library, direct-DIA, and *de novo* sequencing. It identified 30,733 unique sequences within the same length range. Our workflow almost tripled the number of unique sequences of DIA-Umpire and obtained 20% more than PEAKS-Online. As a reference, MaxQuant identified 14,563 sequences in the 8 to 14 aa range on DDA of the same sample in the original publication⁴⁸.

To judge the reliability of the identified HLA peptides, we used MixMHCpred⁵⁴ to deconvolute these identified peptides at the 5% rank level based on the HLA type list in the original publication of the datasets⁴⁸ (Fig. 6b). The overall peptide distribution identified by our pipeline for different HLA types was very similar to that of the original DDA data, indicating that our additionally identified HLA peptides were reliable at the same level.

Finally, we directly tested if our pipeline is also able to identify peptides with unknown HLA types. To simulate this situation, we removed all peptides of the dominant HLA-A*68:01 and used the rest to train a new pan-HLA model. This means that all HLA-A*68:01 peptides in the RA957 sample were now unknown. Then we used only 100 HLA-A*68:01 and all non-HLA-A*68:01 peptides identified by direct-DIA and deconvoluted by MixMHCpred for transfer learning. The resulting library then identified 29,331 peptides including 7,868 from HLA-A*68:01 (Transfer learning with 1000 HLA-A*68:01 peptides retrieved almost all of them) (Fig. 6b). This demonstrates that few-shot transfer learning is able to rescue many of the peptides of an unknown HLA type even if the peptide number is low after direct-DIA identification.




Figure 6. HLA prediction model built on AlphaPeptDeep functionalities. (a) The pipeline with the HLA prediction model to extract potential HLA peptides from the proteome databases. The HLA model is a binary classifier that predicts if a given sequence is a potentially presented HLA sequence. (b) Our HLA prediction model boosts the number of identified HLA-I peptides compared to other tools. Cell line HLA data from RA957 with sequence lengths from 8 to 14 were used. The top bar plots show the number of identified unique sequences of HLA types for each search method. The bottom bar plots the relative frequency of these HLA types. 'Trash' means the peptides cannot be assigned to any HLA types by MixMHCpred at 5% rank level. 'AlphaPeptDeep lib' (red) refers to the library predicted by the sample-specific HLA model and our MS2 and RT models. The bars represent DDA data analyzed by MaxQuant, and the DIA data analyzed by DIA-Umpire, or PEAKS-Online including de novo sequencing. AlphaPeptDeep with the sample-specific HLA library clearly outperforms these. The results of omitting the dominant HLA-A*68:01 (A6801) HLA type in the pan-model and using transfer learning with including 1000 or 100 of these peptides identified by direct-DIA from the data are shown in the last two bars of the A6801 type (see arrows in the panel).

Conclusion

We developed a deep learning framework called AlphaPeptDeep that unifies high-level functionalities to train, transfer learn and use the models for peptide property prediction. Based on these functionalities, we built MS2, RT and CCS models, which enabled the prediction for a large variety of different PTM types. These models can boost DDA identification of for example, HLA peptides, not only in regular search but also in open-search. We also provided a module called 'model shop' which contains generic models so that users can develop new ones from scratch with just a few lines of code. Based on the model shop, we also built an HLA prediction model to predict whether a

peptide sequence is a presented HLA peptide. With the HLA model and the MS2, RT and CCS models in AlphaPeptDeep, we predicted the HLA spectral libraries directly from the whole human proteome, and searched them using HLA DIA data. This is the first time that the predicted libraries at the proteome-level have been used to search DIA data for HLA peptides. Our predicted libraries out-performed other methods including recently published pipelines specifically designed for HLA DIA analysis.

Although AlphaPeptDeep is both powerful and easy to use, we note that traditional machine learning issues, such as overfitting in the framework, still need to be kept in mind. For instance, users still need to split the data, train and test the models on different sets. Trying different hyperparameters such as the number of training epochs is still necessary as well. Different mini-batch sizes and learning rates may also impact on the model training. However, the model shop at least provides baseline models for any property prediction problem.

We hope AlphaPeptDeep will minimize the challenges for researchers that are not AI experts to build their own models either from scratch or on top of our pre-trained models. As we pointed out in our recent review⁴, peptide property prediction can be involved in almost all steps to improve the computational proteomics workflow. Apart from specific properties of interest in MS-based proteomics, it can in principle be used to solve any problem where a peptide property is a function of the amino acid sequence, as we demonstrated by successfully predicting potential HLA peptides to narrow the database search. Therefore, with sufficient and reliable training data, we believe AlphaPeptDeep will be a valuable DL resource for proteomics.

Online Methods

Infrastructure development

To develop AlphaPeptDeep, we built an infrastructure package named AlphaBase (<https://github.com/MannLabs/alphabase>) which contains many necessary functionalities for proteins, peptides, PTMs, and spectral libraries. In AlphaBase, we use the pandas DataFrame as the base data structure, which allows transparent data processing in a tabular format and is compatible with many other Python packages. AlphaPeptDeep uses the AlphaBase DataFrames as the input to build models and predicts properties of peptides. Amino acid and PTM embedding is performed directly from ‘sequence’ (amino acid sequence), ‘mods’ (modification names), and ‘mod_sites’ (modification sites) columns in the peptide DataFrame.

Amino acid embedding

Each amino acid of a sequence is converted to a unique integer, for example, 1 for ‘A’, 2 for ‘B’, ..., and 26 for ‘Z’. Zero is used as a padding value for N- and C-terminals, and other “padding” positions. As a result, there are 27 unique integers to represent an amino acid sequence. A ‘one-hot encoder’ is used to map each integer into a 27-D vector with

zeros and ones. These vectors are mapped to an N-dimensional embedded vector using a linear layer (Extended Fig. 1). For this, we additionally make use of the '*torch.Embedding*' method, which is more efficient and flexible and can support more letters such as all the 128 ASCII codes.

PTM embedding

For each PTM, we use a 6-D embedding vector to represent the C, H, N, O, S, and P atoms. All other atoms of a PTM are embedded into a 2-D vector with a fully connected (FC) layer. The 6-D and 2-D vectors are concatenated into an 8-D vector to represent the PTM (Extended Fig. 1).

MS2 model

The MS2 model consists of an embedding layer, positional encoder layer, and four transformer layers followed by two FC layers. The embedding layer embeds not only amino acid sequences and modifications but also metadata (if necessary) including charge states, normalized collisional energies, and instrument type. All these embedded tensors are concatenated for the following layer.

We added an additional transformer layer to predict the ‘modloss’, which refers to neutral loss intensities of PTMs, for example, the –98 Da of the phospho-group. This modloss layer can be turned off by setting ‘mask_modloss’ as ‘True’. The output layer dimension is $(n - 1) \times 8$ for each peptide, where n is the length of the peptide sequence, and 8 refers to eight fragment types, i.e. b+, b++, y+, y++, b_modloss+, b_modloss++, y_modloss+, and y_modloss++. With ‘mask_modloss=True’, the modloss layer is disabled and the predicted modloss intensities are always zero. The hidden layer size of transformers is 256. The total number of the model parameters is 3,988,974.

All matched b/y fragment intensities in the training and testing datasets were normalized by dividing by the highest matched intensity for each spectrum. The MS2 models were trained based on these normalized intensities. For prediction, negative values will be clipped to zero, hence the predicted values will be between zero and one.

In training phase 1, we only used tryptic peptides in the training datasets. The training parameters were: epoch=100, warmup epoch=20, learning rate (lr)=1e–5, dropout=0.1. In training phase 2, HLA peptides were added to the training set and the parameters were: epoch=20, warmup epoch=5, lr=1e–5, dropout=0.1, mini-batch size=256. In phase 3, phosphorylation and ubiquitylation datasets were added for training, and only phosphorylation sites with >0.75 localized probabilities were considered. The training parameters were: epoch=20, warmup epoch=5, lr=1e–5, dropout=0.1, mini-batch size=256. For transfer learning of the 21 PTMs, the parameters were: epoch=10, warmup epoch=5, lr=1e–5, dropout=0.1, mini-batch size depends on the peptide length. L1 loss was used for all training phases. We used the “cosine schedule with warmup” method implemented in HuggingFace for warmup training of these models including all the following models.

For Thermo Orbitrap instruments, the fragment intensities of each identified PSM are directly extracted from the raw data. For this, we imported the centroided MS2 spectra with Thermo's RawFileReader API that is integrated in AlphaPept, hence the extracted intensities are reproducible across different search engines. For dda-PASEF data, the b/y ion intensities are extracted directly from the msms.txt file of MaxQuant results. Note that different search engines may have different centroiding algorithms for dda-PASEF, resulting in quite different fragment intensities, so fine-tuning is highly recommended for dda-PASEF data analyzed by different software.

A fragment DataFrame is designed to store the predicted intensities. Its columns are fragment ion types (e.g., 'b_z1' for b+ and 'y_z2' for y++ ions), and the rows refer to the different fragmented positions of peptides from which the fragments originate. The start and end pointers of the rows ('frag_start_idx' and 'frag_end_idx') belonging to peptides are stored in the peptide DataFrame to connect between peptides and their fragments. The fragment DataFrame is pre-allocated only once for all peptides before prediction. While predicting, the predicted values of a peptide are assigned to the region of the peptide located by 'frag_start_idx' and 'frag_end_idx'. The fragment DataFrame allows fast creation and storage of the predicted intensities. The tabular format further increases human readability and enables straightforward access by programming.

RT model

The RT model consists of an embedding layer for sequences and modifications, and a CNN layer followed by two LSTM layers with a hidden layer size of 128. The outputs of the last LSTM layer are summed over the peptide length dimension and processed by two FC layers with output sizes of 64 and 1. The total number of the model parameters is 708,224.

All RT values of PSMs in the training datasets were normalized by dividing by the time length of each LC gradient, resulting in normalized RT values ranging from 0 to 1. As a result, the predicted RTs are also normalized. The training parameters were: epoch=300, warmup epoch=30, lr=1e-4, dropout=0.1, mini-batch size=256. The fine-tuning parameters are: epoch=30, warmup epoch=10, lr=1e-4, dropout=0.1, mini-batch size=256. L1 loss was used for training.

To compare predicted RT values with experimental ones, each value is multiplied with the time length of each LC gradient. For testing on peptides with iRT values, we used 11 peptides with known iRT values⁷ to build a linear model between their iRT and predicted RT values. Then all the predicted RTs in the testing sets are converted to iRT values using the linear model.

CCS model

The CCS model consists of an embedding layer for sequence, modifications and charge states, and a CNN layer followed by two LSTM layers with a hidden layer size of 128. The

outputs of the last LSTM layer are summed over the peptide length dimension and processed by two FC layers with output sizes 64 and 1. The total number of the model parameters is 713,452.

The training parameters are: epoch=300, warmup epoch=30, lr=1e-4, dropout=0.1, mini-batch size=256. L1 loss was used for training. The predicted CCS values are converted to mobilities of Bruker timsTOF using the Mason Schamp equation.³⁶

Rescoring

Rescoring includes three steps:

1. Model fine-tuning. 5,000 PSMs are randomly sampled from at most eight RAW files at 1% FDR to fine-tune the MS2, RT and CCS (if applicable) models to obtain project-specific models. The top-10 frequent modifications are also selected for fine-tuning from the eight RAW files. At most 100 PSMs are sampled for each modification. Therefore, the fine-tuning covers not only unmodified peptides, but also modified peptides.
2. Deep learning feature extraction. The tuned MS2, RT and CCS models are used to predict MS2, RT and CCS values for all the reported PSMs including decoys. All PSMs are matched against the MS2 spectra in the RAW files to obtain detected fragment intensities. Then the predicted and detected values are used to calculate 61 score features, which include correlations of fragments, RT differences, mobility differences, and so on (Suppl. Data. 2).
3. Percolator for rescoring. We use the cross-validation schema⁵⁵ to perform the semi-supervised Percolator algorithm to reduce the chance of overfitting. All the peptides are divided into K folds (K=2 in the analyses of this work) and rescored by 5 iterations in Percolator. In each iteration, a Logistic regression model from scikit-learn⁵⁶ is trained with the 61 features on the K-1 folds, and the model is used to re-score on the remainder. All the K folds will be re-scored after repeating this for K times on each of the folds.

Multiprocessing is used in step 2 for faster rescoring. Because GPU RAM is often limited, it can become a bottleneck meaning that only one process is allowed to access the GPU space at a time for prediction. We developed a producer-consumer schema to schedule the tasks with different processes (Extended Fig. 10). The PSMs are matched against MS2 spectra in parallel with multiprocessing grouped by RAW files. Then, they are sent back to the main process for prediction in GPU. At last, the 61 Percolator features are extracted in parallel again. All correlation values between matched and predicted MS2 intensities are also calculated in GPU for acceleration. As this is not memory intensive, the GPU RAM can be shared and used in parallel from different processes. For multiprocessing without GPU, all predictions are done with separate processes and results merged into the main process to run Percolator.

HLA prediction model

The HLA prediction model consists of an embedding layer for sequences, a CNN layer

followed by two LSTM layers with a hidden layer size of 256. The outputs of the last LSTM layer are summed over the sequence length dimension and processed by two linear layers with output sizes of 64 and 1. The sigmoid activation function is applied for last linear layer to obtain probabilities. The total number of the model parameters is 1,669,697.

For training and transfer learning, identified HLA peptides with sequence lengths from 8 to 14 are regarded as positive samples. Negative samples were randomly picked from the reviewed human protein sequences. The sequence number and length distribution were the same for the positive and negative samples. These samples were then split 80% for training and 20% for testing. The parameters for training the pre-trained model were: epoch=100, warmup epoch=20, lr=1e-4, dropout=0.1. For transfer learning, the DIA data were searched by DIA-Umpire and MSFragger⁵⁷ in HLA mode at 1% FDR with reviewed human protein sequence. The parameters for transfer learning were: epoch=50, warmup epoch=20, lr=1e-5, dropout=0.1, mini-batch size=256. Binary cross-entropy loss was used for training.

To predict HLA peptides from fasta files, we first concatenate protein sequences into a long string separated by the “\$” symbol. Next, we use the longest common prefix (LCP) algorithm⁵⁸ to accelerate the unspecific digestion for the concatenated sequence. Only the start and end indices of the peptides in concatenated sequence are saved, thus minimizing the usage of RAM. These indices are used to generate peptide sequences on the fly for prediction. The LCP functionalities have been implemented in AlphaBase. All sequences with a predicted probability larger than 0.7 were regarded as potential HLA peptides.

Open-search for Orbitrap and dda-PASEF data

We performed an open search on the Thermo RAW data with Open-pFind. For HLA DDA data, the reviewed human protein sequences from UniProt (<https://www.uniprot.org/>) were searched with the following parameters: open-search mode=True, enzyme=Z at C-terminal (i.e., unspecific enzyme), specificity=unspecific. The search tolerance was set to ± 10 ppm for MS1 and ± 20 ppm for MS2. All modifications marked as ‘isotopic label’ in UniMod (www.unimod.org) were removed from the searched modification list. The FDR was set as 1% at the peptide level.

To enable Open-pFind search for dda-PASEF data, the spectra were loaded by AlphaPept APIs¹⁸ and exported as pFind compatible MGF files using our in-house Python script. The reviewed drosophila and human sequences were used to search the respective tryptic DDA data with parameters: open-search mode=True, enzyme=KR at C-terminal, enzyme specificity=specific. The search tolerance was set to ± 30 ppm for both MS1 and MS2.

Spectral libraries

Functionalities for spectral libraries are implemented in AlphaBase. When providing DataFrames with sequence, modification and charge columns, the fragment m/z values and intensities are calculated and stored in fragment DataFrames. AlphaBase also

integrates functionalities to load and save DataFrames in a single Hierarchical Data Format (HDF) file for fast access. For subsequent use with DIA-NN or Spectronaut, all the DataFrames are then converted into a tab-separated values file (*.tsv) which is compatible with these tools.

For HLA DIA analysis, we used reviewed human protein sequences to predict HLA peptides. We considered charge states from one to three for each peptide. All RT, CCS, and MS2 were predicted using the model from training phase 3. The 12 most abundant b/y ions with 1+ and 2+ charge states were written to the *.tsv file. Fragment m/z range was set to be from 200 to 1800, precursor m/z range was from 300 to 1800.

In DIA-NN, the mass tolerance for MS1 and MS2 were set to 20 and 10 ppm respectively, with a scan window of 8. All other parameters were the default values of DIA-NN. The results identified from the first pass were used for post-search analysis.

Data availability

The reviewed protein sequence databases of human, E. coli, fission yeast, and drosophila were downloaded from uniprot (<https://www.uniprot.org/>). The training and testing data were from PRIDE with IDs: PXD010595, PXD004732, PXD021013, PXD009449, PXD000138, PXD019854, PXD019086, PXD004452, PXD014525, PXD017476, PXD019347, PXD021318, PXD026805, PXD026824, PXD029545, PXD000269, and PXD001250.

The mono-allelic HLA DDA dataset was downloaded from MassIVE with ID MSV000084172. The tumor HLA dataset was downloaded from PRIDE with ID PXD004894.

HLA DIA data and the MaxQuant results of DDA data from the RA957 cell line were downloaded from PRIDE with ID PXD022950. HLA DIA results of PEAKS-Online were downloaded from the PEAKS-Online publication.⁵³ Only results from RAW files '20200317_QE_HFX2_LC3_DIA_RA957_R01.raw' and '20200317_QE_HFX2_LC3_DIA_RA957_R02.raw' from RA957 were used to compare different methods.

Result files and Python notebooks to reproduce the analysis results in this study (total of 7 GByte) can be found in <https://doi.org/10.6084/m9.figshare.20260761>.

Code availability

The source code of AlphaBase and AlphaPeptDeep are fully opened on GitHub:

<https://github.com/MannLabs/alphabase> and

<https://github.com/MannLabs/alphapeptdeep>. They are also available through PyPI with

"pip install alphabase" and "pip install peptdeep". The versions used in this study of AlphaBase and AlphaPeptDeep are 0.1.2 and 0.1.2 respectively. All the pre-trained MS2,

RT, and CCS models can be found in https://github.com/MannLabs/alphapeptdeep/releases/download/pre-trained-models/pretrained_models.zip. These models will be automatically downloaded when using the AlphaPeptDeep package for the first time.

The versions of other software are displayed in the Reporting Summary.

Acknowledgements

We thank Marvin Thielert for the testing of the spectral libraries. We thank Mario Oroshi and Igor Paron for help with retrieval of MS RAW data. This study was supported by The Max-Planck Society for the Advancement of Science and by the Bavarian State Ministry of Health and Care through the research project DigiMed Bayern (www.digimed-bayern.de). I.B. acknowledges funding support from her Postdoc.Mobility fellowship granted by the Swiss National Science Foundation [P400PB_191046].

Author contributions

W.-F.Z. developed AlphaBase, AlphaPeptDeep and models, and analyzed the data. X.-X.Z. developed the models, contributed to AlphaPeptDeep code, and analyzed the data. S.W. designed the template code structure on GitHub and developed HDF functionalities in AlphaBase. C.A. reviewed the code and contributed to the model shop functionalities. M.W. came up with the idea of HLA prediction. I.B. contributed to AlphaBase. E.V. developed functionalities in AlphaViz for mirrored MS2 plots and testing the integration with AlphaPeptDeep. M.T.S. reviewed almost all the source code of AlphaBase and AlphaPeptDeep, and provided a lot of suggestions. M.M. supervised this project. W.-F.Z., M.T.S. and M.M. wrote the manuscript. All the authors revised the manuscript.

Ethics declarations

Competing interests

The authors declare no competing interests.

References

1. Aebersold, R. & Mann, M. Mass-spectrometric exploration of proteome structure and function. *Nature* vol. 537 Preprint at <https://doi.org/10.1038/nature19949> (2016).
2. Meissner, F., Geddes-McAlister, J., Mann, M. & Bantscheff, M. The emerging role of mass spectrometry-based proteomics in drug discovery. *Nature Reviews Drug Discovery* (2022) doi:10.1038/s41573-022-00409-3.
3. Li, S. & Tang, H. Computational methods in mass spectrometry-based proteomics. in *Advances in Experimental Medicine and Biology* vol. 939 (2016).
4. Mann, M., Kumar, C., Zeng, W. F. & Strauss, M. T. Artificial intelligence for

proteomics and biomarker discovery. *Cell Systems* vol. 12 Preprint at <https://doi.org/10.1016/j.cels.2021.06.006> (2021).

- 5. Wen, B. *et al.* Deep Learning in Proteomics. *Proteomics* vol. 20 Preprint at <https://doi.org/10.1002/pmic.201900335> (2020).
- 6. Moruz, L., Tomazela, D. & Käll, L. Training, selection, and robust calibration of retention time models for targeted proteomics. *Journal of Proteome Research* **9**, (2010).
- 7. Escher, C. *et al.* Using iRT, a normalized retention time for more targeted measurement of peptides. *Proteomics* **12**, (2012).
- 8. Pfeifer, N., Leinenbach, A., Huber, C. G. & Kohlbacher, O. Statistical learning of peptide retention behavior in chromatographic separations: A new kernel-based approach for computational proteomics. *BMC Bioinformatics* **8**, (2007).
- 9. Ma, C. *et al.* Improved Peptide Retention Time Prediction in Liquid Chromatography through Deep Learning. *Analytical Chemistry* **90**, (2018).
- 10. Zhou, X. X. *et al.* PDeep: Predicting MS/MS Spectra of Peptides with Deep Learning. *Analytical Chemistry* **89**, (2017).
- 11. Zeng, W. F. *et al.* MS/MS Spectrum prediction for modified peptides using pDeep2 Trained by Transfer Learning. *Analytical Chemistry* **91**, (2019).
- 12. Tiwary, S. *et al.* High-quality MS/MS spectrum prediction for data-dependent and data-independent acquisition data analysis. *Nature Methods* **16**, (2019).
- 13. Gessulat, S. *et al.* Prosit: proteome-wide prediction of peptide tandem mass spectra by deep learning. *Nature Methods* **16**, (2019).
- 14. Hochreiter, S. & Schmidhuber, J. Long Short-Term Memory. *Neural Computation* **9**, (1997).
- 15. Cho, K., van Merriënboer, B., Bahdanau, D. & Bengio, Y. On the Properties of Neural Machine Translation: Encoder-Decoder Approaches. (2014).
- 16. Lou, R. *et al.* DeepPhospho accelerates DIA phosphoproteome profiling through in silico library generation. *Nature Communications* **12**, (2021).
- 17. Ekvall, M., Truong, P., Gabriel, W., Wilhelm, M. & Käll, L. Prosit Transformer: A transformer for Prediction of MS2 Spectrum Intensities. *Journal of Proteome Research* (2021) doi:10.1021/acs.jproteome.1c00870.
- 18. Strauss, M. T. *et al.* AlphaPept, a modern and open framework for MS-based proteomics. *bioRxiv* (2021).
- 19. Paszke, A. *et al.* PyTorch: An imperative style, high-performance deep learning library. in *Advances in Neural Information Processing Systems* vol. 32 (2019).
- 20. Dosovitskiy, A. *et al.* An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. (2020).
- 21. Avsec, Ž. *et al.* Effective gene expression prediction from sequence by integrating long-range interactions. *Nature Methods* **18**, (2021).
- 22. Tunyasuvunakool, K. *et al.* Highly accurate protein structure prediction for the human proteome. *Nature* **596**, (2021).
- 23. Wolf, T. *et al.* HuggingFace's Transformers: State-of-the-art Natural Language Processing. (2019).
- 24. Goyal, P. *et al.* Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour.

(2017).

- 25. Meier, F. *et al.* Deep learning the collisional cross sections of the peptide universe from a million experimental values. *Nature Communications* **12**, (2021).
- 26. Wen, B., Li, K., Zhang, Y. & Zhang, B. Cancer neoantigen prioritization through sensitive and reliable proteogenomics analysis. *Nature Communications* **11**, (2020).
- 27. Müller, J. B. *et al.* The proteome landscape of the kingdoms of life. *Nature* **582**, (2020).
- 28. Zolg, D. P. *et al.* Building ProteomeTools based on a complete synthetic human proteome. *Nature Methods* **14**, (2017).
- 29. Meier, F., Park, M. A. & Mann, M. Trapped ion mobility spectrometry and parallel accumulation–serial fragmentation in proteomics. *Molecular and Cellular Proteomics* vol. 20 Preprint at <https://doi.org/10.1016/j.mcpro.2021.100138> (2021).
- 30. Chong, C., Coukos, G. & Bassani-Sternberg, M. Identification of tumor antigens with immunopeptidomics. *Nature Biotechnology* vol. 40 Preprint at <https://doi.org/10.1038/s41587-021-01038-8> (2022).
- 31. Li, K., Jain, A., Malovannaya, A., Wen, B. & Zhang, B. DeepRescore: Leveraging Deep Learning to Improve Peptide Identification in Immunopeptidomics. *Proteomics* **20**, (2020).
- 32. Wilhelm, M. *et al.* Deep learning boosts sensitivity of mass spectrometry-based immunopeptidomics. *Nature Communications* **12**, (2021).
- 33. Rosenberger, G. *et al.* A repository of assays to quantify 10,000 human proteins by SWATH-MS. *Scientific Data* **1**, (2014).
- 34. Wang, S. *et al.* NAguideR: Performing and prioritizing missing value imputations for consistent bottom-up proteomic analyses. *Nucleic Acids Research* **48**, (2020).
- 35. Chi, H. *et al.* Comprehensive identification of peptides in tandem mass spectra using an efficient open search engine. *Nature Biotechnology* **36**, (2018).
- 36. Mason, E. A. & McDaniel, E. W. *Transport Properties of Ions in Gases*. *Transport Properties of Ions in Gases* (1988). doi:10.1002/3527602852.
- 37. Paul Zolg, D. *et al.* Proteometools: Systematic characterization of 21 post-translational protein modifications by liquid chromatography tandem mass spectrometry (lc-ms/ms) using synthetic peptides. *Molecular and Cellular Proteomics* **17**, (2018).
- 38. Voytik, E. *et al.* AlphaViz: Visualization and validation of critical proteomics data directly at the raw data level. *bioRxiv* 2022.07.12.499676 (2022) doi:10.1101/2022.07.12.499676.
- 39. Bouwmeester, R., Gabriels, R., Hulstaert, N., Martens, L. & Degroeve, S. DeepLC can predict retention times for peptides that carry as-yet unseen modifications. *Nature Methods* **18**, (2021).
- 40. Käll, L., Canterbury, J. D., Weston, J., Noble, W. S. & MacCoss, M. J. Semi-supervised learning for peptide identification from shotgun proteomics datasets. *Nature Methods* **4**, (2007).
- 41. Mann, M. & Wilm, M. Error-Tolerant Identification of Peptides in Sequence Databases by Peptide Sequence Tags. *Analytical Chemistry* **66**, (1994).

42. Sarkizova, S. *et al.* A large peptidome dataset improves HLA class I epitope prediction across most of the human population. *Nature Biotechnology* **38**, (2020).
43. Bassani-Sternberg, M. *et al.* Direct identification of clinically relevant neoepitopes presented on native human melanoma tissue by mass spectrometry. *Nature Communications* **7**, (2016).
44. Cox, J. & Mann, M. MaxQuant enables high peptide identification rates, individualized p.p.b.-range mass accuracies and proteome-wide protein quantification. *Nature Biotechnology* **26**, (2008).
45. Sturm, T. *et al.* Mild Acid Elution and MHC Immunoaffinity Chromatography Reveal Similar Albeit Not Identical Profiles of the HLA Class i Immunopeptidome. *Journal of Proteome Research* **20**, (2021).
46. Tareen, A. & Kinney, J. B. Logomaker: Beautiful sequence logos in Python. *Bioinformatics* **36**, (2020).
47. Ludwig, C. *et al.* Data-independent acquisition-based SWATH - MS for quantitative proteomics: a tutorial . *Molecular Systems Biology* **14**, (2018).
48. Pak, H. S. *et al.* Sensitive immunopeptidomics by leveraging available large-scale multi-HLA spectral libraries, data-independent acquisition, and MS/MS prediction. *Molecular and Cellular Proteomics* **20**, (2021).
49. Ritz, D., Kinzi, J., Neri, D. & Fugmann, T. Data-Independent Acquisition of HLA Class I Peptidomes on the Q Exactive Mass Spectrometer Platform. *Proteomics* **17**, (2017).
50. Demichev, V., Messner, C. B., Vernardis, S. I., Lilley, K. S. & Ralser, M. DIA-NN: neural networks and interference correction enable deep proteome coverage in high throughput. *Nature Methods* **17**, (2020).
51. Martinez-Val, A., Bekker-Jensen, D. B., Hogrebe, A. & Olsen, J. V. Data Processing and Analysis for DIA-Based Phosphoproteomics Using Spectronaut. in *Methods in Molecular Biology* vol. 2361 (2021).
52. Tsou, C. C. *et al.* DIA-Umpire: Comprehensive computational framework for data-independent acquisition proteomics. *Nature Methods* **12**, (2015).
53. Xin, L. *et al.* A streamlined platform for analyzing tera-scale DDA and DIA mass spectrometry data enables highly sensitive immunopeptidomics. *Nature Communications* **13**, 3108 (2022).
54. Gfeller, D. *et al.* The Length Distribution and Multiple Specificity of Naturally Presented HLA-I Ligands. *The Journal of Immunology* **201**, (2018).
55. Granholm, V., Noble, W. S. & Käll, L. A cross-validation scheme for machine learning algorithms in shotgun proteomics. *BMC Bioinformatics* **13 Suppl 16**, (2012).
56. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, (2011).
57. Kong, A. T., Leprevost, F. v., Avtonomov, D. M., Mellacheruvu, D. & Nesvizhskii, A. I. MSFragger: Ultrafast and comprehensive peptide identification in mass spectrometry-based proteomics. *Nature Methods* **14**, (2017).
58. Zhou, C. *et al.* Speeding up tandem mass spectrometry-based database searching

by longest common prefix. *BMC Bioinformatics* **11**, (2010).