

Rapid Prediction of Electron–Ionization Mass Spectrometry Using Neural Networks

Jennifer N. Wei,^{*,†,‡,§} David Belanger,^{*,†} Ryan P. Adams,[§] and D. Sculley[†]


[†]Google Brain, Cambridge, Massachusetts 02142, United States

[‡]Department of Chemistry and Chemical Biology, Harvard University, Cambridge, Massachusetts 02138, United States

[§]Department of Computer Science, Princeton University, Princeton, New Jersey 08540, United States

Supporting Information

ABSTRACT: When confronted with a substance of unknown identity, researchers often perform mass spectrometry on the sample and compare the observed spectrum to a library of previously collected spectra to identify the molecule. While popular, this approach will fail to identify molecules that are not in the existing library. In response, we propose to improve the library's coverage by augmenting it with synthetic spectra that are predicted from candidate molecules using machine learning. We contribute a lightweight neural network model that quickly predicts mass spectra for small molecules, averaging 5 ms per molecule with a recall-at-10 accuracy of 91.8%. Achieving high-accuracy predictions requires a novel neural network architecture that is designed to capture typical fragmentation patterns from electron ionization. We analyze the effects of our modeling innovations on library matching performance and compare our models to prior machine-learning-based work on spectrum prediction.



INTRODUCTION

Mass spectrometry (MS) is an important tool used to identify unknown molecular samples in a variety of applications, from characterization of organic synthesis products, to pharmacokinetic studies,¹ to forensic studies,² to analyzing gaseous samples on remote satellites.³ In electron–ionization mass spectrometry (EI-MS), molecular samples are ionized by an electron beam and broken into fragments. The resultant ions are separated by an electric field until they reach a detector. The mass spectrum is a distribution of the frequency or intensity of each type of ion, ordered by mass-to-charge (m/z) ratio.

A popular method for identifying a sample from its mass spectrum is to look up the sample's spectrum in a *reference library*.^{4,5} Here, a similarity function is used to measure the similarity between the query spectrum from the sample and each spectrum in the library. If the measurement noise when obtaining the query spectrum is reasonable, then the library spectrum with the highest similarity will correspond to the correct identification of the sample.^{6,7} A schematic of this process is shown in Figure 1a.

This library matching approach is very popular, but it suffers from a *coverage problem*: if the sample consists of a molecule that is not in the library, then correct identification is impossible. This is an issue in practice, since existing mass spectral reference libraries, such as the NIST/NIH/EPA MS database,⁴ Wiley Registry of Mass Spectral Data,⁵ and MassBank,⁸ only contain hundreds of thousands of reference spectra. The coverage problem could be reduced by recording

spectra for additional molecules, but this is time-consuming and expensive. For example, NIST releases updates to its library every 3 years, containing roughly 20 000 new spectra. Additionally, mass spectra of new molecules are only added to the library if the molecule is of common interest; molecules for newly synthesized compounds are typically not incorporated.^{4,9}

An alternative solution is to use *de novo* methods that input a spectrum and directly generate a molecule, without using a fixed list of molecules; we discuss some of these methods in the *Background* section. These approaches currently have low accuracy and are difficult for practitioners to incorporate into their existing work-flows.

Another method for alleviating the coverage problem is to augment existing libraries with synthetic spectra that are generated by a model. Thus far, this approach has not been practical, as existing spectrum prediction methods are very computationally expensive. These prediction models use quantum mechanics calculations^{10–12} or machine learning¹³ to estimate the probability of each bond breaking under ionization, and thus the frequency of each ion fragment. Since these methods must either compute molecular orbital energies with high accuracy using expensive calculations, or else stochastically simulate the fragmentation of the molecule, the time needed for each model to make a prediction scales with the size of the molecule, taking up to 10 min for large

Received: January 25, 2019

Published: March 19, 2019




Figure 1. Library matching task. (a) Depiction of how query spectra are matched to a collection of reference spectra as performed by mass spectrometry software. (b) Query spectra are compared against a library comprising spectra from the NIST 2017 main library and spectra predicted by our model (outlined in blue). Spectral images adapted from NIST Webbook.²⁹

molecules.^{10,13} For applications of identifying metabolites from a metabolomic spectra, much faster predictions of individual molecular spectra are required.¹⁴

In response, we present Neural Electron–Ionization Mass Spectrometry (NEIMS), a neural network that predicts the electron–ionization mass spectrum for a given small molecule. Since our model directly predicts spectra, instead of bond-breaking probabilities, it is dramatically faster than previously reported methods, making it possible to generate predictions for thousands of possible candidates in seconds. Furthermore, the approach does not rely on specific details of EI, and thus our model could be easily retrained to predict mass spectra for other ionization methods. We envision that this tool can expand coverage in areas of molecular space of interest to researchers that are likely candidates for the identity of the molecule.

We test the performance of our model by predicting mass spectra for small molecules from the NIST 2017 Mass Spectral

Library. We find that the predictive capability of our model is similar to previously reported machine learning models, but requires much less time to make predictions. Additionally, we report the similarity of the spectra predicted by NEIMS. The code repository for NEIMS is publicly available at github.com/brain-research/deep-molecular-massspec.

■ BACKGROUND

De Novo Prediction of Molecules from Mass Spectra.

Several algorithms have been developed previously either for predicting spectra or for predicting the molecule's identity given the spectrum. One of the earliest efforts in artificial intelligence was a model used to identify molecules from their mass spectrum. Heuristic DENDRAL (Dentritic Algorithm) was a collaboration between chemists and computer scientists at Stanford in the 1960s.¹⁵ This algorithm used expert rules from chemistry to help identify patterns in the spectra and suggest possible identities for the molecule. A few years later,

Meta DENDRAL was introduced to learn the expert rules that originally been given to Heuristic DENDRAL.¹⁶

Since then, several models have been reported to predict identities of samples directly from the spectrum. Many have been developed for tandem mass spectrometry, where the task is to predict the original peptide sequences from digested fragments given the mass spectrum.¹⁷ Some of these methods use machine learning to achieve this task.^{18,19} Several previously published models use neural network models to analyze mass spectra to predict molecules directly. One approach predicts the fingerprint from the spectrum and looks up the fingerprint in a library of fingerprints.²⁰ Another approach predicts the molecule directly either as a SMILES representation^{21,22} or from a ranked list of possible structural conformers.²³

Prediction of Mass Spectra from Molecules. In this work, we focus on the prediction of spectra from molecules. The advantage of this approach over *de novo* approaches is that new libraries of synthetic spectra can be easily incorporated into the existing mass spectrometry software to improve the coverage of existing libraries.

The first prediction methods for EI-MS spectrum used quantum mechanical simulation techniques to predict fragmentation events. There are three methods of predicting the mass spectrum using first principles.¹⁰ The first is to use quasiequilibrium theory, also known as Rice–Ramsperger–Kassel–Marcus theory, to estimate the rate constants for ionization reaction.^{24–26} The second is to estimate the bond order energies within a molecule, and estimate where a molecule may fragment. A related method to this second method is to calculate the cross-section of molecular orbitals upon electron impact to predict the molecule's ionization behavior.^{12,27} The third method uses Born–Oppenheimer Molecular Dynamics. Quantum Chemistry Electron–Ionization Mass Spectrometry (QCEIMS) is a particularly recent example of the *ab initio* molecular dynamics method.^{10,11,28} The trajectories resulting from this simulation are then analyzed for the presence of ionic fragments. The distribution of the ion fragments aggregated from all the simulations is then renormalized to generate a calculated EI-MS spectrum. Each of these methods requires at least 1000 s per molecule,¹³ and may even take days or weeks for molecules of 50 atoms. While these methods may be fast for methods involving density functional theory, they do not have the speed needed to rapidly generate a collection of spectra of thousands of molecules. Furthermore, some of the basis sets used for the density functional theory might not support the presence of inorganic atoms.

Allen et al.¹³ introduced a machine learning model, Competitive Fragmentation Modeling–Electron Ionization (CFM–EI), to predict EI-MS spectra. This probabilistic model predicts the probability of breaking molecular bonds under electron ionization, and also predicts the charged fragment that is likely to form. To generate the spectra, it is necessary to run a stochastic simulation to determine the frequency of each molecular fragment. In the **Results and Discussion** section, we directly compare this method with our proposed model.

METHODS

Our goal is to design a model that will accurately predict the EI-MS spectrum for any molecule. This will be used to produce an augmented reference library containing both

predicted spectra and experimentally measured spectra. This task is outlined in **Figure 1b**.

We first discuss our choice of similarity metrics for mass spectra. Next, we describe our method for spectra prediction, describing our new architecture which is required for high accuracy of mass spectra. We then explain how we evaluate our model's impact on the library matching task.

Similarity Metrics for Mass Spectra. The ability to a match a query spectrum from a sample to the correct spectrum in the library depends on the choice of similarity metric between spectra.^{7,30} A weighted cosine similarity is commonly used by mass spectrometry software. The exact form of the cosine similarity is given below:⁷

$$\text{Similarity}(\mathbf{I}_q, \mathbf{I}_l) = \frac{\sum_{k=1}^{M_{\max}} m_k I_{qk}^{0.5} \cdot m_k I_{lk}^{0.5}}{\left\| \sum_{k=1}^{M_q} (m_k I_{qk}^{0.5})^2 \right\| \left\| \sum_{k=1}^{M_l} (m_k I_{lk}^{0.5})^2 \right\|} \quad (1)$$

Here, $\mathbf{I}_q$ and $\mathbf{I}_l$ are vectors of m/z intensities representing the query spectrum and the library spectrum, respectively; m_k and I_k are the mass-to-charge ratio and intensity found at $m/z = k$; M_l and M_q are the largest indices of $\mathbf{I}_q$ and $\mathbf{I}_l$ with nonzero values; and $M_{\max}$ is the larger of M_l and M_q . The motivation for the weighting by m/z is because the peaks in mass spectra corresponding to larger fragments are more characteristic and useful in practice for identifying the true molecule.

Other similarity metrics besides cosine distance similarities are also employed. For example, one other similarity method involves estimating the relative importance of one peak given the other peaks.³⁰ Other methods uses a Euclidian difference between peaks, or use a variation of the Hamming distance.^{7,31} Another similarity metric accounts for neutral losses, or the intensity peaks corresponding to the loss of small, neutral fragments from the original molecular ion.³² It is also possible to use the same form of the similarity function as in **eq 1**, but with different weighting given to the intensity or the masses.⁷ In principle, machine learning could be also used to learn a parametrized similarity metric that yields improved library matching performance. However, this custom metric would be difficult to deploy, since it would require changing the software used by practitioners.

We develop our model with the assumption that **eq 1** will be used for the similarity metric in downstream library matching software that consumes an augmented library.

Spectral Prediction. We treat the prediction of mass spectrometry spectra as a multidimensional regression task. The output of our model is a vector that represents the intensity at every integral m/z bin. We use this discretization granularity for m/z because it is what is provided in the NIST data sets we use for training our model.

In the NEIMS model (**Figure 3**), we first map molecules to additive Extended Circular Fingerprints (ECFPs).³³ These fingerprints are similar to their binary counterparts³⁴ in that they record molecular subgraphs made up from local neighborhoods around each atom node in the molecule, but differ in that they count the occurrences for each subgroup. This information is then hashed into a vector representation. The difference is that additive fingerprints record the frequency that each bit is set, rather than just the presence. The RDKit Cheminformatics package³³ was used to generate the fingerprints. We use a fingerprint length of 4096 with a radius of 2. These features are then passed into a multilayer




Figure 2. Spectral prediction with MLP forward model (a) and MLP bidirectional model (b). For both spectra plots, the true spectrum is shown in blue on top, while the predicted spectrum is shown inverted in red. Note that the spectrum predicted by the bidirectional model shows fewer stray peaks than the forward model, particularly for larger m/z values. These peaks are much easier to predict with the reverse prediction mode.




Figure 3. Molecular representations are passed into a multilayer perceptron to generate an initial output. This output is used to make a forward prediction starting at $m/z = 0$ and $m/z = M$ and in reverse starting from $m/z = M$ and ending at $m/z = 0$. A sigmoid gating is applied to the inputs as shown in eq 5.

perceptron neural network (MLP). To account for some of the physical phenomena of ionization, we make some application-specific adjustments to the prediction from the MLP, described in the [Adjustments for Physical Phenomena](#) section.

In the [Library Matching Results](#) section, we compare the performance of NEIMS to that of a simple linear regression (LR) model. Here, we apply a linear transformation to the ECFP features.

To train the model, we use a modified mean-squared-error loss function. This loss function, shown below, follows the same weighting pattern as in eq 1:

$$L(I, \hat{I}) = \sum_{k=1}^{M(x)} \left(\frac{m_k \hat{I}_k^{0.5}}{\left\| \sum_{k=1}^M (m_k I_k^{0.5})^2 \right\|} - \frac{m_k \hat{I}_k^{0.5}}{\left\| \sum_{k=1}^M (m_k \hat{I}_k^{0.5})^2 \right\|} \right)^2 \quad (2)$$

where I is the ground truth spectrum, $\hat{I}$ is the predicted spectrum, and $M(x)$ is the mass of the input molecule. We used stochastic gradient descent to optimize the parameters of the MLP with the Adam optimizer.³⁵ We use Tensorflow³⁶ to construct and train the model.

Adjustments for Physical Phenomena. In practice, we have found that the conventional MLP described in the previous section struggles to accurately predict the right-hand side of spectra (Figure 2a). Errors in this region, which correspond to large m/z , are particularly damaging for library matching with the weighting in eq 1.

This section introduces a revised neural network architecture (Figure 3) designed to better model the underlying fragmentation process that occurs in mass spectrometry. We have found that it improves prediction in the high mass region of the spectrum (Figure 2b), which yields improvements in library matching as discussed in the [Results and Discussion](#) section.

As is standard for MLPs used for regression, the predictions of the above MLP model on an input molecule x are an affine transformation of a set of features $f(x)$, which are computed by all but the final layer of the network. For reasons that will become apparent, we refer to the above MLP as performing *forward* prediction. At bin $m/z = i$, we have the following predicted intensity:

$$p_i^f(x) = w_i^{\text{fT}} f(x) + b_i^f \quad (3)$$

where w_i^f and b_i^f are the model's weights and biases for forward prediction at bin i .

The input ECFP features, from which $f(x)$ is computed, capture local structures in the molecule, so generally $f(x)$ will be more accurate in capturing the presence of small substructures of molecule x . Often, there is a direct correspondence between the presence of such substructures and spectral peaks with small m/z . For example, in Figure 2a a peak occurs at $m/z = 35$, due to the presence of chlorine. Therefore, an accurate forward prediction model will have a learned weight w_{35} that will output a high intensity at $i = 35$ if there is evidence in $f(x)$ for the presence of chlorine.

On the other hand, forward prediction often struggles to accurately predict intensities for large fragments that are the result of neutral losses.⁶ One reason for this is that the composition of large fragments is not captured well by the ECFP representation. Another reason is that information learned about the cleavage of a small group does not transfer well across molecules of different masses. For pentachlorobenzene, which has a molecular mass of 250 Da, the fragment that results from the loss of a neutral chlorine atom results in a peak at 215 Da. Meanwhile, for chlorobenzene, which has a mass of 112 Da, the fragment resulting from a loss of a chlorine atom would have a peak at 77 Da. Despite the clear relationship between these intensity peaks, the forward model is not parametrized to capture this pattern.

In response, following the physical phenomenon that created the fragments, we define larger ion peaks as a function of the residual groups that were broken off from the original molecule. Referring to our previous example of pentachlorobenzene ($M(x) = 250$), we can parametrize the m/z ratio of the fragment which lost a chlorine group as $m/z = 250 - 35 = 215$. The corresponding fragment in chlorobenzene would have a mass of $m/z = 215 - 35 = 77$. By defining the peaks in this way, it is possible for these predictions of spectral intensities to be linked by the prediction at index 35. This leads to the indexing scheme of our *reverse prediction* model:

$$p_{M(x)+\tau-i}^r(x) = w_i^{\text{rT}} f(x) + b_i^r \quad (4)$$

Here, $\tau > 0$ is a small shift that allows for peaks to occur at intensities greater than $M(x)$, due to isotopes. In practice, reverse prediction is implemented using a copy of the forward model, with separate sets of parameters for the final affine layer, but shared parameters for $f(x)$. The outputs of this model are postprocessed on a per-molecule basis to obey the indexing in eq 4, which depends on each molecule's mass.

Both the forward and reverse predictions are combined to form a *bidirectional* prediction. That is, the final prediction at index i is a combination of both p_i^f and p_i^r . In the case of pentachlorobenzene, the prediction of spectral intensity at $m/z = 215$ is a function of p_{215}^f from the forward mode and $p_{35+\tau}^r$ from the reverse mode. Instead of simply averaging the two prediction modes, we have found that small additional performance improvements can be obtained using a coordinate-wise *gate*. Here, the output $p_i(x)$ at position i is given by

$$p_i(x) = \sigma(\text{gate}_i) p_i^f(x) + (1 - \sigma(\text{gate}_i)) p_i^r(x) \quad (5)$$

where gate_i is an affine transformation of $f(x)$, and $\sigma(\cdot)$ is a sigmoid function. This approach echoes the formulation of the Hybrid Similarity Search designed by Moorthy et al.,³² which accounts for peaks that are created by small fragment ions and those which are created by large fragments which have lost smaller groups.

Finally, for all models, we zero out predicted intensities at m/z that are greater than $M(x) + \tau$.

By adding these features, we incorporate some of the physical phenomena that occur in mass spectrometry into our model while maintaining the overall simplicity of the MLP. In this way, we are able to predict the spectrum directly without resorting to sampling bond-breaking events within the molecule, which requires subsequent stochastic sampling to obtain a spectrum.

Library Matching Evaluation. We evaluate NEIMS using an *augmented reference library* consisting of a combination of observed spectra and model-predicted spectra, with library matching performance computed with respect to a *query set* of spectra. These are from the NIST 2017 replicates library, which is a collection of noisier spectra for molecules that are contained in the NIST main library. The inconsistencies in these spectra reflect experimental variation, and make an informative data set to test our model's performance.

To construct the augmented reference library, we edit the NIST main library, removing spectra corresponding to the query set molecules and replacing them with the predictions from NEIMS. We then perform library matching and calculate the similarity between each query spectrum and every spectrum from the augmented library. We record the rank of the correct spectrum, i.e., the rank of the predicted spectrum corresponding to the molecule which made the query spectrum. The similarity metric is eq 1.

For the purposes of tuning model hyperparameters, we chose to optimize $\text{recall}@10$, i.e., the percentage of our query set for which the correct spectra had a matching rank of less than or equal to 10 in the library matching task. Half of the replicates library was used for tuning hyperparameters, and the remaining half was used to evaluate test performance. All models were trained on the spectra prediction task for 100 000 training steps with a batch size of 100.

During the library match search, we have a *mass filtering* option. This feature reduces the library size so it only includes spectra from molecular candidates that have a molecular mass that differ by a few Daltons from the mass of the query molecule. If the EI-MS analysis is combined with mass spectrometry techniques using soft ionization methods, it is possible to determine the mass of the molecule being analyzed. In the CFM-EI model, the molecular formula is used to filter the search library.¹³ Using the molecular mass to filter the library allows more possible candidate spectra to be considered in the search than using a molecular formula filter.

RESULTS AND DISCUSSION

To analyze the performance of the models, we trained with 240 942 spectra from the NIST 2017 Mass Spectral Main Library. These spectra were selected so that no molecules in the replicates library have spectra in the training set.

After hyperparameter tuning using Vizier,³⁷ we found that the optimal MLP architecture has seven layers of 2000 nodes,




Figure 4. Performance of different model architectures. (a) Comparison of the recall@10 accuracy of the linear regression model and the multilayer perceptron model using the forward, reverse, and bidirectional architecture. (b) Performance of NEIMS at different recall levels, and its comparison against the performance of using the NIST main library itself.

with residual network connections between the layers,³⁸ using ReLU activation and a dropout rate of 0.25.

Library Matching Results. We first examine the effects of our various modeling decisions on performance. Figure 4a compares the performance of forward, reverse, and bidirectional versions of the linear regression and MLP models on the library matching task. For bidirectional prediction in the linear regression model, the forward and reverse predictions are simply averaged together, rather than applying the gate described in eq 5.

The top row of Figure 4a shows that it is not possible to achieve perfect recall accuracy on the library matching task even when using the full NIST main library as the reference library, without any model-predicted spectra. Observing Figure 4b we see that, using the NIST main library as the reference library, we have 86% recall@1 accuracy, and 98.3% recall@10 accuracy. This serves as a practical upper bound on achievable library matching accuracy and reflects the experimental inconsistencies between the main library spectra and replicates spectra.⁹

The forward prediction mode for both the linear regression model and the multilayer perceptron (MLP) has poor performance. The linear regression model is improved by 20% when switching to using reverse mode prediction. Using bidirectional prediction mode improves recall@10 accuracy by 30% for both the linear regression and the multilayer perceptron model. This finding suggests that the bidirectional prediction mode is more effective at capturing the fragmentation events than the forward-only model.

Figure 2 shows the improvement in spectral prediction for pentachlorobenzene using the bidirectional MLP model. Note that the bidirectional model on the right more accurately models intensities at larger m/z . The intensity peaks for larger m/z are critical for determining the identity of a molecule, and are more heavily weighted in eq 1.

NEIMS achieves 91.7% recall@10 after applying a mass filter. The mass filter was set to a tolerance of 5 Da of the query molecule's mass; this reduces the size of the library to a median of 6696 spectra for each query molecule. In practice, this tolerance window could be set to a larger window, depending on the uncertainty of the information about the molecular mass of the ion. For the rest of this report, we will refer to the bidirectional multilayer perceptron model with mass filtering of 5 Da as the default settings for NEIMS.

From Figure 4b we see that while NEIMS has decent performance for recall levels of 10 and above compared to the NIST spectral library, it has considerably worse performance for recall values of 1 and 5. This result is unsurprising given that the hyperparameters of the model were trained to maximize performance on recall@10. As many experimentalists will examine the top matches to find the library spectrum that best matches the sample's spectrum, we believe that tuning for recall@10 is sufficient for an initial approach.

Comparison to Previously Reported Models. We next compared our model's performance directly to the performance of the CFM-EI model.¹³ The setup of Allen et al. differs from our current setup in a few ways. First, they evaluate their model on the NIST 14 spectral library. Second, for the library matching task, their augmented reference library contains only spectra predicted by their model, and none from the original NIST collection. Third, the cosine similarity metric eq 1 used for evaluation in library matching in CFM-EI uses a different weighting scheme. In their analysis, the cosine similarity is weighted by $m_k^{0.5}$ instead of m_k to de-emphasize the larger peaks in the mass spectrum, as they ran their experiments on other data sets with a higher proportion of larger molecules.¹³

To compare the performance of NEIMS to that of CFM-EI, we match their setup identically. We retrain our NEIMS model on the NIST 14 data set, and evaluate the performance using the NIST 14 replicates as the query set. For library matching, we incorporate only predicted spectra into our augmented library, and using the same modified similarity metric.

The library matching performances for CFM-EI and NEIMS are compared against the NIST 14 library, as reported in Table 1. NEIMS performs slightly better than CFM-EI on the library matching task. More importantly, NEIMS is able to make spectral predictions orders of magnitude faster than CFM-EI.

Table 1. Performance on Library Matching Task for NIST 17^a

model	recall@1 (%)	recall@10 (%)	average run time (ms)
NIST 14 Reference Library	77	99 ^a	—
CFM-EI	42.6	89 ^a	300 000
NEIMS	54.3	92.7	0.47

^aValues were estimated from Figure 4 of Allen et al.¹³

With NEIMS, it would be possible to generate spectra for 1 million molecules in 90 min on a CPU, with potential for considerable speedup with using a GPU.

Distances between Predicted and Ground Truth Spectra. So far, we have evaluated the quality of the NEIMS predictions indirectly, by way of how they affect library matching with an augmented library. Next, we assess the prediction accuracy directly, by measuring the similarity (eq 1) between spectra in the NIST main library and the model's predictions. We refer to this similarity as the *predicted similarity*.

There is inherent noise in mass spectra due to stochasticity of the underlying physical process and also to experimental inconsistencies.⁹ The NIST replicates library provides multiple spectra for each molecule, and we can use these sets of spectra to characterize the scale of this noise for each molecule. Specifically, we define the inherent noise for a given molecule as the average pairwise similarity between all corresponding spectra, in both the NIST main library and the NIST replicates library, and refer to this as the *overall similarity*.

For each molecule, we compute the ratio of the predicted similarity to overall similarity as a normalized metric for the quality of our predictions. A ratio of 1.0 would suggest that there are limited available headroom for improvements using machine learning, since the model's errors are comparable to the variability in the data.

Figure 5 shows the improvement in this ratio for the MLP bidirectional model over the MLP forward model, confirming




Figure 5. Comparing the similarity between the predicted spectrum and the ground truth spectrum to the overall similarity between spectra for the same molecule.

that the bidirectional model has better spectral prediction performance. For the MLP bidirectional model, roughly half of the molecules have a predicted similarity to overall similarity ratio that is greater than 0.9, indicating that there is potential for further improvement to the model. Some of these molecules have ratios that are greater than 1, which is possible if there is more variation between the spectra (i.e., a lower overall similarity) than between the predicted spectrum and the main library spectrum (i.e., predicted similarity).

To analyze the ability of our model to extrapolate we analyzed the relationship between predicted similarity and similarity of the query molecule to the training set molecules. The molecules in the test set have limited structural similarity to the molecules in the training set. We observe that 22.7% of molecules in the test set have a Tanimoto similarity of greater

than 0.8 with at least one molecule in the training set. A plot of this relationship between predicted spectral similarity and closely related molecules in terms of Tanimoto similarity can be found in Figure S3 in the Supporting Information. Based on these results, we believe that our model is able to extrapolate to some areas of molecular space that were not fully covered by the training set. Future work will examine the limitations of the model's ability to extrapolate.

CONCLUSION

We demonstrate that NEIMS achieves high library matching performance on an augmented spectral library containing predictions for molecules in the query set. The performance of NEIMS is also slightly better than existing machine learning models for predicting EI-MS spectra, with significant boost in speed of prediction.

The high performance in library matching is attributable to the bidirectional prediction mode. The reverse mode in particular allows the model to more accurately predict intensities for larger fragments which result from the loss of small neutral subgroups. We observe that the improvement in the library matching task also corresponds with improvement in the similarity of the predicted spectra to the ground truth spectra.

Several adjustments could be made to further improve the predictive accuracy of NEIMS. For example, NEIMS currently does not have a method to model intensity peaks corresponding to isotopes in ion fragments. If we were to train on spectral data with greater precision in the peaks locations, we should be able to learn the exact identities of atoms based on the decimal values of the m/z peak locations.

Mass filtering improved the performance of NEIMS by 6%. This suggests that, for experimental setups where it is possible to know the molecular mass of the sample with some accuracy, it is possible to improve the accuracy of matching on the augmented spectral library. It would also be interesting to explore other settings for mass filtering, such as filtering out spectra which have a molecular mass that is much smaller than the position of the largest m/z peak.

Using graph-convolutional molecular representations,^{39,40} especially bond-centered representations,⁴¹ would likely improve predictive accuracy at a slightly higher computational cost. The predictions made from ECFP are limited by the descriptiveness of the fingerprint.⁴² In particular, the overlap in representation for different molecular features represents a huge limitation to the representation of the molecule.

Combining NEIMS with transfer learning methods could allow for spectral prediction specific to individual spectrometry machines. A library of such machine-specific spectra would improve matching.⁹

The lightweight framework of NEIMS makes it possible to rapidly generate spectral predictions for large numbers of molecular candidates. This collection of predicted spectra can then be used directly in mass spectrometry software to expand the coverage of molecules which can be identified by mass spectrometry. Because the requirements of NEIMS has limited dependence to EI mass spectrometry, it is likely that some of the principles used here could be extended to other types of mass spectrometry.

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge on the ACS Publications website at DOI: [10.1021/acscentsci.9b00085](https://doi.org/10.1021/acscentsci.9b00085).

Peripheral findings on the relationship between Tanimoto similarity and predicted spectra similarity, as well as the training and test loss curves (PDF)

■ AUTHOR INFORMATION

Corresponding Authors

*E-mail: weijennifer@google.com.

*E-mail: dbelanger@google.com.

ORCID

Jennifer N. Wei: 0000-0003-3567-9511

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

We thank Stephen Stein and Laura Castellanos for fruitful discussions about mass spectrometry and for their helpful feedback on the manuscript. We thank Steven Kearnes for his suggestions on the manuscript, and Lucy Colwell and Michael Brenner for their helpful conversations. The code for this work can be found at <https://github.com/brain-research/deep-molecular-masspec>.

■ REFERENCES

- (1) Hsieh, Y.; Korfmacher, W. A. Increasing speed and throughput when using HPLC-MS/MS systems for drug metabolism and pharmacokinetic screening. *Curr. Drug Metab.* **2006**, *7*, 479–489.
- (2) Zhou, Z.; Zare, R. N. Personal information from latent fingerprints using desorption electrospray ionization mass spectrometry and machine learning. *Anal. Chem.* **2017**, *89*, 1369–1372 PMID: 28194988.
- (3) Petrie, S.; Bohme, D. K. Ions in space. *Mass Spectrom. Rev.* **2007**, *26*, 258–280.
- (4) Stein, S. E. *Mass Spectral Database*; National Institute of Standards and Technology (NIST), 2017.
- (5) McLafferty, F. W. *Wiley Registry of Mass Spectral Data*, 11th ed.; John Wiley and Sons, 2016.
- (6) Stein, S. E. Chemical substructure identification by mass spectral library searching. *J. Am. Soc. Mass Spectrom.* **1995**, *6*, 644–655.
- (7) Stein, S. E.; Scott, D. R. Optimization and testing of mass spectral library search algorithms for compound identification. *J. Am. Soc. Mass Spectrom.* **1994**, *5*, 859–866.
- (8) Horai, H.; Arita, M.; Kanaya, S.; Nihei, Y.; Ikeda, T.; Suwa, K.; Ojima, Y.; Tanaka, K.; Tanaka, S.; Aoshima, K.; et al. MassBank: A public repository for sharing mass spectral data for life sciences. *J. Mass Spectrom.* **2010**, *45*, 703–714.
- (9) Stein, S. Mass spectral reference libraries: an ever-expanding resource for chemical identification. *Anal. Chem.* **2012**, *84*, 7274–7282.
- (10) Bauer, C. A.; Grimme, S. How to compute electron ionization mass spectra from first principles. *J. Phys. Chem. A* **2016**, *120*, 3755–3766.
- (11) Grimme, S. Towards first principles calculation of electron impact mass spectra of molecules. *Angew. Chem., Int. Ed.* **2013**, *52*, 6306–6312.
- (12) Guerra, M.; Parente, F.; Indelicato, P.; Santos, J. P. Modified binary encounter Bethe model for electron-impact ionization. *Int. J. Mass Spectrom.* **2012**, *313*, 1–7.
- (13) Allen, F.; Pon, A.; Greiner, R.; Wishart, D. Computational prediction of electron ionization mass spectra to assist in GC/MS compound identification. *Anal. Chem.* **2016**, *88*, 7689–7697.
- (14) Viant, M. R.; Kurland, I. J.; Jones, M. R.; Dunn, W. B. How close are we to complete annotation of metabolomes? *Curr. Opin. Chem. Biol.* **2017**, *36*, 64–69.
- (15) Buchanan, B. G.; Feigenbaum, E. A. *Readings in artificial intelligence*; Elsevier, 1981; pp 313–322.
- (16) Lindsay, R. K.; Buchanan, B. G.; Feigenbaum, E. A.; Lederberg, J. DENDRAL: a case study of the first expert system for scientific hypothesis formation. *Artificial intelligence* **1993**, *61*, 209–261.
- (17) Eng, J. K.; McCormack, A. L.; Yates, J. R. An approach to correlate tandem mass spectral data of peptides with amino acid sequences in a protein database. *J. Am. Soc. Mass Spectrom.* **1994**, *5*, 976–989.
- (18) Tran, N. H.; Zhang, X.; Xin, L.; Shan, B.; Li, M. De novo peptide sequencing by deep learning. *Proc. Natl. Acad. Sci. U. S. A.* **2017**, *114*, 8247–8252.
- (19) Schoenholz, S. S.; Hackett, S.; Deming, L.; Melamud, E.; Jaitly, N.; McAllister, F.; O'Brien, J.; Dahl, G.; Bennett, B.; Dai, A. M.; Koller, D. Peptide-spectra matching from weak supervision. 2018, arXiv:1808.06576. arXiv.org e-Print archive. <https://arxiv.org/abs/1808.06576>.
- (20) Dührkop, K.; Shen, H.; Meusel, M.; Rousu, J.; Böcker, S. Searching molecular structure databases with tandem mass spectra using CSI: FingerID. *Proc. Natl. Acad. Sci. U. S. A.* **2015**, *112*, 12580–12585.
- (21) Curry, B.; Rumelhart, D. E. MSnet: A neural network which classifies mass spectra. *Tetrahedron Comput. Methodol.* **1990**, *3*, 213–237.
- (22) Rabinowitz, T. Mass-Spectrometry-Prediction, 2017. <https://github.com/terryrabinowitz/Mass-Spectrometry-Prediction/blob/master/readme.pdf> (accessed on Oct 6, 2018).
- (23) Lim, J.; Wong, J.; Wong, M. X.; Tan, L. H. E.; Chieu, H. L.; Choo, D.; Neo, N. K. N. Chemical structure elucidation from mass spectrometry by matching substructures. 2018, arXiv:1811.07886. arXiv.org e-Print archive. <https://arxiv.org/abs/1811.07886>.
- (24) Lorquet, J. Whither the statistical theory of mass spectra? *Mass Spectrom. Rev.* **1994**, *13*, 233–257.
- (25) Lorquet, J.-C. Landmarks in the theory of mass spectra. *Int. J. Mass Spectrom.* **2000**, *200*, 43–56.
- (26) Rosenstock, H. M.; Wallenstein, M.; Wahrhaftig, A.; Eyring, H. Absolute rate theory for isolated systems and the mass spectra of polyatomic molecules. *Proc. Natl. Acad. Sci. U. S. A.* **1952**, *38*, 667–678.
- (27) Irikura, K. K. Ab initio computation of energy deposition during electron ionization of molecules. *J. Phys. Chem. A* **2017**, *121*, 7751–7760.
- (28) Asgeirsson, V.; Bauer, C. A.; Grimme, S. Quantum chemical calculation of electron ionization mass spectra for general organic and inorganic molecules. *Chem. Sci.* **2017**, *8*, 4879–4895.
- (29) Linstrom, P. J.; Mallard, W. G. The NIST Chemistry WebBook: A chemical data resource on the internet. *J. Chem. Eng. Data* **2001**, *46*, 1059–1063.
- (30) McLafferty, F.; Hertel, R.; Villwock, R. Probability based matching of mass spectra. Rapid identification of specific compounds in mixtures. *Org. Mass Spectrom.* **1974**, *9*, 690–702.
- (31) Hertz, H. S.; Hites, R. A.; Biemann, K. Identification of mass spectra by computer-searching a file of known spectra. *Anal. Chem.* **1971**, *43*, 681–691.
- (32) Moorthy, A. S.; Wallace, W. E.; Kearsley, A. J.; Tchekhovskoi, D. V.; Stein, S. E. Combining fragment-ion and neutral-loss matching during mass spectral library searching: A new general purpose algorithm applicable to illicit drug identification. *Anal. Chem.* **2017**, *89*, 13261–13268.
- (33) RDKit: Open-source cheminformatics. <http://www.rdkit.org> (accessed on Sept 23, 2018).
- (34) Rogers, D.; Hahn, M. Extended-connectivity fingerprints. *J. Chem. Inf. Model.* **2010**, *50*, 742–754 PMID: 20426451.
- (35) Kingma, D.; Ba, J. Adam: A method for stochastic optimization. 2014, arXiv:1412.6980. <https://arxiv.org/abs/1412.6980>.

- (36) Abadi, M. et al. TensorFlow: A system for large-scale machine learning. 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), 2016; pp 265–283.
- (37) Golovin, D.; Solnik, B.; Moitra, S.; Kochanski, G.; Karro, J.; Sculley, D. Google vizier: A service for black-box optimization. Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2017; pp 1487–1495.
- (38) He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. 2015, arXiv:1512.03388. <https://arxiv.org/abs/1512.03388>.
- (39) Duvenaud, D. K.; Maclaurin, D.; Iparraguirre, J.; Bombarell, R.; Hirzel, T.; Aspuru-Guzik, A.; Adams, R. P. *Convolutional networks on graphs for learning molecular fingerprints*; 2015; pp 2224–2232.
- (40) Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; Dahl, G. E. Neural message passing for quantum chemistry. 2017, arXiv:1704.01211. <https://arxiv.org/abs/1704.01212>.
- (41) Kearnes, S.; McCloskey, K.; Berndl, M.; Pande, V.; Riley, P. Molecular graph convolutions: moving beyond fingerprints. *J. Comput.-Aided Mol. Des.* **2016**, *30*, 595–608.
- (42) Landrum, G. Colliding bits, 2014. <http://rdkit.blogspot.com/2014/02/colliding-bits.html> (accessed on Nov 11, 2019).