Machine learned calibrations to high-throughput molecular excited state calculations

Shomik Verma,^{1, a)} Miguel Rivera,² David O. Scanlon,³ and Aron Walsh¹

¹⁾*Department of Materials, Imperial College London, Exhibition Road, London SW7 2AZ, U.K.*

²⁾*Department of Chemistry, University College London, 20 Gordon Street, London WC1H 0AJ, U.K.*

³⁾*Department of Chemistry and Thomas Young Centre, University College London, 20 Gordon Street, London WC1H 0AJ, U.K.*

(*Electronic mail: a.walsh@imperial.ac.uk)

(Dated: 7 January 2022)

Understanding the excited state properties of molecules provides insights into how they interact with light. These interactions can be exploited to design compounds for photochemical applications, including enhanced spectral conversion of light to increase the efficiency of photovoltaic cells. While chemical discovery is time- and resource-intensive experimentally, computational chemistry can be used to screen large-scale databases for molecules of interest in a procedure known as high-throughput virtual screening. The first step usually involves a high-speed but low-accuracy method to screen large numbers of molecules (potentially millions) so only the best candidates are evaluated with expensive methods. However, use of a coarse first-pass screening method can potentially result in high false positive or false negative rates. Therefore, this study uses machine learning to calibrate a high-throughput technique (xTB-sTDA) against a higher accuracy one (TD-DFT). Testing the calibration model shows a ~ 5 -fold decrease in error in-domain and a ~ 3 -fold decrease out-of-domain. The resulting mean absolute error of ~ 0.14 eV is in line with previous work in machine learning calibrations and out-performs previous work in linear calibration of xTB-sTDA. We then apply the calibration model to screen a 250k molecule database and map inaccuracies of xTB-sTDA in chemical space. We also show generalizability of the workflow by calibrating against a higher-level technique (CC2), yielding a similarly low error. Overall, this work demonstrates machine learning can be used to develop a both cheap and accurate method for large-scale excited state screening, enabling accelerated molecular discovery across a variety of disciplines.

I. INTRODUCTION

Understanding the excited state properties of molecules helps describe how they interact with light. These photochemical interactions can include fundamental processes such as photosynthesis,¹ human vision,² or photostability³. Photochemistry is also important in designing new molecules with certain properties, for example spectral converters for photovoltaics,⁴ which are of particular interest in this study. Using an interplay between their excited states, certain molecules can up- or down-convert wavelengths of light to improve photovoltaic efficiency. Unfortunately, existing spectral conversion molecules have low efficiency,⁵ so further exploration is required. It is difficult to explore the excited state space of molecules with experimental methods alone, so researchers often turn to computational methods for more detailed study.^{6,7} Traditionally, time-dependent density functional theory (TD-DFT) has been the workhorse for excited state energy calculations. However, TD-DFT can be computationally intensive, and for applications in high-throughput screening or generative design, a faster method must be used.

TD-DFT often relies on DFT for ground-state calculations of charge density and structure optimization. Recently,

work has been done in tight binding as an approximation to DFT to improve its computation time while retaining most of its accuracy. Specifically, density functional tight binding (DFTB)⁸ was developed in the late 1990's⁹ and exhibited a combination of the accuracy of DFT and the efficiency of semi-empirical quantum chemistry methods. More recently, the eXtended Tight Binding (xTB)¹⁰ methods were developed to solve the issues with DFTB of extensive parameterization and low transferability.⁸ They differ from DFTB methods in that they utilize top-down parameterization, with semiempirical parameters fit to a large dataset rather than computed with first-principles calculations.¹⁰ The primary approximations are considering molecular orbitals to be a linear combination of atomic orbitals (LCAOs), using the local density approximation (LDA) for exchange-correlation energy, and using a truncated Taylor expansion to map density to total energy.¹⁰

To accelerate excited-state calculations, Grimme introduced the simplified Tamm-Dancoff density functional approach (sTDA)¹¹ as an approximation to TD-DFT. The key approximations of sTDA include simplifications to two-electron integrals and setting an upper limit to the excitation space, which improve computation time by 2 orders of magnitude.¹¹ As sTDA was developed to calculate excitation spectra, there is no excited state relaxation component, and only vertical excitation energies can be calculated. The differences between vertical excitation, vertical emission, and adiabatic energy are shown in Figure 1. In this work, only vertical excitations are considered. The

^{a)} Also at Department of Mechanical Engineering, Massachusetts Institute of Technology, 33 Massachusetts Ave, Cambridge, MA, USA

FIG. 1. Schematics depicting the basics of molecular excitation and emission. (a) Typical Jablonksi diagram from ground state (S_0) to excited state (S_1), with various vibrational levels (0-3) depicted for both states. (b) Demonstration of Frank-Condon principle of $0 \rightarrow 1$ vertical excitation (blue arrow) followed by nuclear re-configuration and $1 \leftarrow 0$ vertical emission (red arrow). Also shown is the $0 \rightarrow 0$ transition energy in yellow. (c) The expected experimental excitation/absorption curve (blue) and emission curve (red), along with the theoretical $0 \rightarrow 0$ energy difference (dashed yellow line), demonstrating the Stokes shift. Note that in computation, often the energy minimum is used instead of the lowest vibrational level, so the starting energies for excitation and emission may be different in experiment.

additional computational expense of excited state relaxation is prohibitively slow for high-throughput workflows. Further, the Stokes shift for rigid molecules should be small, on the order of 0.1 eV.¹² From now on, for concision, the vertical excitation energy will be referred to as the excitation energy or excited state energy.

xTB and sTDA can be combined in a workflow called xTB-sTDA, allowing ultrafast computation of excited states.¹³ This has been used extensively, with several studies using the method to screen large databases of materials such as copolymers,¹⁴ conjugated polymers,¹⁵ small aromatic molecules,¹⁶ photocatalysts,¹⁷ and organic dyes.¹⁸

However, due to the approximations presented above, there is a tradeoff between accuracy and computational speed. In Grimme and Bannwarth’s original paper introducing xTB-sTDA, they reported a mean absolute error (MAE) between xTB-sTDA and coupled-cluster/TD-DFT calculated excited state energies of 0.34 to 0.48 eV, depending on the complexity of the input structure.¹³ Even though xTB-sTDA is often used as a first-pass in high-throughput screening, with higher-quality computational methods used to evaluate properties of a screened subset of molecules, having an accurate first-pass method is essential to ensure all suitable candidates are

included in the suggested subset.

Therefore, a method of improving the accuracy of xTB-sTDA in a way that preserves its high-throughput characteristics is desired. One approach is to calibrate the results of xTB-sTDA against a higher-accuracy computational methods using machine learning (ML). This type of calibration of excited state energies has been done in the past for various computational techniques, including TD-DFT against coupled-cluster methods¹⁹ and TD-DFT against experiment.^{20–22} For xTB-sTDA, a few studies have used a linear calibration technique to correct excited state energies,^{16,18} instead of ML. However, the improvement from linear calibration was low, with a mean absolute error (MAE) of around 0.2 eV, compared to around 0.1 eV for the ML studies presented above. To the authors’ knowledge, no previous study has applied ML to calibrate xTB-sTDA.

Due to the promising potential of ML to increase the accuracy of baseline methods, this work presents a ML calibration of xTB-sTDA. We first provide an overview of the ML workflow, including information about the training and test sets, model architecture and characteristics, and computational settings used for data generation. Then, we test the generated model and apply it to several scenarios to show its applicability and generalizability.

II. METHODS

A. Reference computational technique

xTB was originally parameterized from the spin-component-scaled coupled cluster (SCS-CC2)²³ method and TD-DFT, so one of these would be a natural choice as the reference computational technique. As mentioned previously, TD-DFT is generally the workhorse for excited state calculation. However, TD-DFT often has relatively high errors, with an MAE of 0.2 to 0.4 eV depending on the functional compared to experiment or the theoretical best estimate.²⁴ A general-purpose exchange-correlation functional is B3LYP²⁵, which has been shown to have an MAE around 0.25 eV, while LDA, generalized gradient approximation, and other hybrid functionals have higher error.²⁶ While B3LYP performs well for localized densities, range-separated hybrid functionals such as CAM-B3LYP²⁷ are increasingly used for delocalized densities such as those in excitations as they include a long-range correction.²⁸ However, these range-separated hybrid functionals are more computationally expensive than B3LYP. Due to its accuracy and efficiency, B3LYP is the most commonly used functional in computational molecular chemistry.²⁹

Coupled-cluster (CC) methods are typically more accurate than TD-DFT, predicting excitation energies within 0.1 eV of experiment. However, because of the computational expense of CC, it is difficult to generate a large amount of data using these methods. Therefore, most of the existing molecular excited state databases use TD-DFT.^{19,30–33} Due to these existing databases and the computational efficiency of TD-DFT, the reference method was chosen to be TD-DFT with

B3LYP, as a large amount of excited state data is required for training purposes.

B. Training dataset

Specifically, the training sets for the ML models considered in this study were derived from the existing PubChemQC (PCQC)³¹ and QM-symex³³ databases. For concision, we will use the (functional/basis set) notation to describe the level of theory used in calculations. PCQC includes the first 10 singlet excited state energies (S_{1-10}) for 3.5M molecules computed using B3LYP/6-31G(d) for ground state optimization and B3LYP/6-31+G(d) for excitation. Similarly, QM-symex computes both S_{1-10} and the first 10 triplet excited state energies (T_{1-10}) for 173k molecules using B3LYP/6-31G(2df,p) for ground state optimization and B3LYP/6-31G for excitation.

We are interested in calibrating both singlet (S_1) and triplet (T_1) excited state energies output by xTB-sTDA, but PCQC does not include triplet excitations. Calculating triplet excitations for 3.5M molecules independently would be prohibitively expensive, so it was necessary to determine which molecules in PCQC would be relevant to spectral conversion applications and therefore have interesting excited state properties. To extract such molecules from PCQC, a literature scraping workflow was developed. We used the SCOPUS API³⁴ to obtain abstracts of articles tagged with "triplet-triplet annihilation" (TTA) or "singlet fission" (SF) keywords. Then, we used ChemDataExtractor³⁵ to extract molecule names from the abstracts. We then used the PubChem API³⁶ to convert molecule names into PubChem CIDs and conduct a 2D Tanimoto-coefficient based similarity search among PubChem molecules to expand the molecular space of interest. We then cross-referenced the identified molecules against PCQC to get the singlet energies, and triplet energies were independently generated with TD-DFT using equivalent settings to PCQC. Overall, this process allowed us to select 10k molecules of interest from PCQC, named SCOP-PCQC (after SCOPus-PCQC). As seen in Figure 2(a), plotting these molecules in the global chemical space of PCQC shows some densely populated areas indicating regions of interest. To balance the 10k molecules in SCOP-PCQC, a 10k subset of QM-symex was randomly selected and named QM-symex-10k.

However, this 10k molecule subset of PCQC may be too small for ML model training. In addition, a model trained on only molecules relevant to TTA/SF may have poor out-of-domain performance. Therefore, we turn to active learning (AL) to add diverse molecules to the training set with further sampling of PCQC.³⁷ Here, we use AL techniques to evaluate regions of chemical space where the ML model is uncertain. A 10-ensemble ML model was trained on the 10k SCOP-PCQC molecules to directly predict S_1 and T_1 energies. Then, the ensemble was used to predict S_1 and T_1 energies on the remaining 3.5M molecules in PCQC. The 100k molecules with the highest uncertainty (variance in ensemble prediction) were chosen as an expansion to SCOP-

PCQC, labeled SCOP-AL-Exp, for each of S_1 and T_1 . This process helps ensure broad applicability of the ML model. Figure 2(b) shows the map of the SCOP-AL-Exp molecules in the global chemical space of PCQC. As seen, the molecules have broader coverage of the space. However, there are still clearly more concentrated regions, suggesting that despite the chemical similarity of structures, their excited state energies may be significantly different. In the sparser regions, there is a clearer structure-property correlation.

C. Test datasets

To evaluate the generated models, various test datasets were used. First, we used 10-fold cross-validation (CV) to quantify in-domain accuracy. To prove broader applicability, however, external test sets were also compiled. 1143 molecules from Wilbraham et al.'s paper on Mapping the optoelectronic property space of small aromatic molecules (MOPSSAM) was used, 143 from their calibration training set and 1000 randomly selected from the remaining 250k molecules.¹⁶ 10k molecules from Fallon et al.'s paper on indolonaphthyridine thiophene (INDT) derivatives were also used as they are promising candidates for singlet fission.⁷ 1000 molecules from Abreha et al.'s VERDE database (VerdeDB) were used, as the classes of molecules identified (porphyrins, quinones, dibenzoperylenes) are relevant for various green chemistry excited state applications.³² Finally, to truly test the broader applicability of the model, another AL cycle was run on PCQC. Using a training set composed of the 10k molecules from SCOP-PCQC plus the 200k molecules from SCOP-AL-Exp, an ensemble ML model was generated and used to evaluate uncertainty on the remaining PCQC molecules. 100k of the highest uncertainty molecules for each of S_1 and T_1 were chosen as the last test set, labeled PCQC-AL.

D. Computational details

Since the supervised ML model takes in molecular structure and excited state data, we must obtain excited state data for all molecules. xTB-sTDA data was all independently generated. A 3D structure was first initialized using OpenBabel's gen3d function for a short conformer search and preliminary geometry optimization.³⁸ Full ground-state optimization was conducted with GFN2-xTB³⁹ with the tight threshold and a benzene generalized-Born surface-area (GBSA) solvation model to mimic a non-polar environment. xtb4stda⁴⁰ was then used to prepare the wavefunctions output by xTB for sTDA. Finally, sTDA was used to calculate excited-state properties, using an energy threshold of 10 eV. The -t flag was used for triplet excited state calculations. For TD-DFT data, database values were used where available. PCQC had S_1 TD-DFT data, but T_1 data was independently generated. MOPSSAM had S_1 TD-DFT data for the 143 calibration set, but not for the 1000 sampled molecules, so this was independently generated (see Supplementary Figure S2 for a comparison of MOPSSAM 143 S_1 data vs. S_1 data generated

FIG. 2. Embedding of (a) SCOP-PCQC data (red) in the global PCQC chemical space (grey), and (b) SCOP-PCQC expansions with active learning for S_1 (purple) and T_1 (blue) in the global PCQC chemical space (grey). A UMAP model created on 350k molecules of PCQC was used to predict the locations of the training set molecules. The Jaccard-Tanimoto similarity coefficient is calculated between each pair of molecules, and UMAP uses this metric for dimensional reduction to 2D space. The most common scaffolds of molecules in SCOP-PCQC is shown in Supplementary Figure S1.

with our workflow, showing virtually identical results). T_1 TD-DFT data was also independently generated. Both INDt and VerdeDB had S_1 and T_1 TD-DFT data available. However, VerdeDB used the M06 functional for calculations, so these molecules were re-calculated with B3LYP.

Note that while the xTB-sTDA portion of the workflow was standardized, the TD-DFT data was database-dependent. For consistency, only databases that used the B3LYP functional were included in this study, but initial coordinate generation technique and basis sets for (TD)DFT varied for each dataset. The specific settings for each database are shown in Supplementary Information Table S1. Once excited state values using both xTB-sTDA and TD-DFT were either compiled or calculated, they could be fed to the ML model.

E. Machine learning model

The type and architecture of the ML model must be optimized for performance. The 3 ML models considered were DeepChem’s⁴¹ GCN⁴² (DC GCN), DeepChem’s MPNN⁴³ (DC MPNN), and Chemprop’s MPNN⁴⁴ (CP MPNN). All 3 of these models use graph neural networks, a natural choice for molecules where nodes represent atoms and edges represent bonds, but use different architectures and methodologies for featurization and property prediction (more details can be found in the respective papers^{42–44}). We are interested in comparing each model’s performance in our application. The default, out-of-the-box settings for each ML model were used, as described in Supplementary Information

Section III. Calibration of the 1000 molecules in the VerdeDB³² database was used to compare the different ML models. The small size and relatively homogeneous nature of this dataset makes it suitable for quickly comparing different ML models. Only the SMILES (simplified molecular-input line-entry system⁴⁵) representation of the molecule was provided as input, and the goal of each model was to accurately predict the S_1 and T_1 error between xTB-sTDA and TD-DFT. Instead of predicting both S_1 and T_1 error simultaneously, two separate single-task models were generated, both using 10-fold cross-validation. For each fold, the trained ML model was used to predict error values of the test set. Then, each molecule’s predicted error was added to the xTB-sTDA output to give a calibrated energy, called the xTB-ML value. The xTB-ML values were compared to the TD-DFT reference results by calculating an R^2 score.

Figure 3(a) shows the results of comparison for S_1 and T_1 energies. As seen, all ML models vastly outperform the linear calibration method. Between the ML models, CP MPNN performs the best for both T_1 and S_1 . Note that the large variability in R^2 can be explained by the presence of outliers in the test set – since the test set was only composed of 100 molecules (10% of 1k), a few outliers can vastly impact performance.

Figure 3(b,c) show plots of original vs. CP MPNN-calibrated xTB-sTDA data for (a) S_1 and (b) T_1 energies, with test data from all 10 folds compiled and with outliers removed. From this analysis, it is evident that CP MPNN performs well in calibrating xTB-sTDA results, even with its default settings.

FIG. 3. (a) Comparison of various ML models in accurately calibrating xTB-sTDA against TD-DFT, quantified by R^2 score. ‘orig’ = original xTB-sTDA data with no calibration, ‘lin calib’ = linear regression calibration of xTB-sTDA data. All others are ML models as presented above. Blue bars are xTB-ML T_1 energies while orange bars are xTB-ML S_1 energies. R^2 for original S_1 data is -1.84 ± 0.65 , the plot was truncated for clarity. (b,c) Plots of original xTB data (‘orig’, red) and CP MPNN ML-calibrated xTB data (‘fixed’, blue) against reference TD-DFT data generated with Gaussian, for (b) S_1 energies and (c) T_1 energies. Datapoints are all test data compiled across 10 non-overlapping folds in cross-validation. (d) R^2 scores of xTB-ML vs. TD-DFT for various improvements attempted to CP MPNN. ‘default’ bars use the out-of-the-box hyperparameter settings with no additional features. ‘100ep’ bars use 100 epochs instead of the usual 30. ‘hyperopt’ bars use hyperparameter optimization. Bars labeled ‘multi’ use multitasking so only 1 model predicts both S_1 and T_1 .

To see if the performance could be boosted further, various architectural improvements were attempted. These included increasing the number of epochs to 100, conducting hyperparameter optimization, conducting multi-task training, and adding RDKit-calculated features as input (in addition to the NN-calculated features). The results from these

improvements are shown in Figure 3(d).

As seen, there are only small differences in performance between the default settings and any potential improvements to the ML settings. For T_1 , hyperparameter optimization provides minimal improvement, while including additional features or adding multitasking reduces accuracy. For S_1 , both hyperparameter optimization and multitasking marginally improve performance. There is thus a tradeoff in using multitasking as it could reduce accuracy for T_1 predictions but improve accuracy for S_1 , while also reducing overall computation time. Because of the time savings of the multi-task model, this was used for ML for the following sections. Hyperparameter optimization was not performed for the following models, due to the only marginal improvement seen. Based on this analysis, a larger-scale calibration model can now be developed using CP MPNN.

F. Calibration workflow

Combining all of the above methodology, a workflow was developed to create ML models that calibrate the xTB-sTDA energy of molecules against TD-DFT reference. The workflow can be separated into three distinct steps: data generation, model training, and model testing. In the data generation step, S_1 and T_1 excited state energies using TD-DFT and sTDA were either extracted from existing databases or calculated, if necessary. The errors between the energies derived from the two techniques were calculated and used as the ground-truth values that the ML model tried to predict. The SMILES strings were also extracted for molecules and used as a representation of molecular structure.

In the model training step, ML models were trained to take input data and predict xTB-sTDA vs. TD-DFT error. Two classes of models were trained, one with only the SMILES string as input (Class 1) and another with both SMILES and sTDA energy as input (Class 2). During training, the SMILES string was converted to the graphical representation of the molecule, which was then featurized using an MPNN. If the sTDA energy was used (Class 2), it was concatenated as an extra feature at this step. Then, feature to property prediction was conducted using a feed-forward NN. This training process was repeated in 10-fold cross-validation with 80%/10%/10% train/validation/test splits and resulted in an optimized ensemble ML model for error prediction.

Once the ML model was trained, it was tested on various datasets. Using the respective inputs, the ML model predicted xTB-sTDA vs. TD-DFT errors for the test molecules. When the errors were added to the original xTB-sTDA values, the final calibrated energies were obtained. These were compared to the TD-DFT-calculated values to get a quantitative measure of accuracy of each ML model.

Figure 4 presents an overview of the calibration workflow. The following section details the results of applying the above calibration workflow to generate an optimized ML model. The model is then applied for high-throughput screening and chemical space mapping.

FIG. 4. Workflow for xTB-ML calibration. Blue boxes represent data, red boxes represent intensive calculations, orange boxes represent quick calculations, and green boxes represent final results. Starting with the training datasets, the TD-DFT and SMILES data are directly extracted. The SMILES strings are converted to 3D molecular structures with OpenBabel and xTB, and then excited state calculations are conducted with xTB-sTDA. Then the SMILES string (Class 1) or the SMILES string and sTDA energy (Class 2) are fed to the ML model, which is trained to predict the error between xTB-sTDA and TD-DFT. The resulting ML model can be used to predict values of various test datasets and quantitatively evaluate its accuracy.

III. RESULTS

A. Cross-validation

Before considering external datasets, ML model performance was evaluated on subsets of the training sets themselves. 10-fold cross-validation models were trained separately on the SCOP-PCQC and QM-symex-10k datasets. Class 1 models (only molecular structure as input) were first tested, giving an MAE of around 0.20 eV for both S_1 and T_1 in both datasets. Class 2 models (both molecular structure and sTDA energy as input) improved performance further, as shown in Figure 5. As seen, the original data has low accuracy when compared to TD-DFT results, and linear calibration improves the accuracy slightly. However, there is not a clear linear shift due to some groups of molecules located farther from the line of best fit. The MAE for ML-calibrated xTB-sTDA for SCOP-PCQC S_1 and T_1 is ~ 0.16 eV. QM-symex-10k ML-calibration performs better, with an S_1 MAE of 0.14 eV and a T_1 MAE of 0.09 eV. As seen, for both SCOP-PCQC and QM-symex-10k, these are significant improvements from uncalibrated xTB-sTDA values, especially for T_1 data. The increase in accuracy

with ML is likely because ML detects higher-order patterns, allowing groups of molecules to shift locally instead of having to follow a global calibration rule.

As seen, the ML models performed well in cross-validation. However, it is possible that the ML model only performed well because the datasets were homogeneous, so similar molecules to those in the test set were included in the training set. To evaluate the broad applicability of our model, we used external test sets of molecules not included in either of the datasets above.

B. External test datasets

To test broad applicability, a more general ML model is needed. Therefore, an overarching ML model was trained on the 10k SCOP-PCQC molecules combined with the 10k QM-symex-10k molecules, for a total training size of 20k molecules. The overarching ML model was first tested on the MOPSSAM 143 external dataset. As seen in Figure 6, the ML calibrated xTB-sTDA data matches TD-DFT values better than both the original data and the linearly calibrated data. While the data are sparse, there are a few regions where the improvement is clearly visible. For example, for

FIG. 5. Plots of xTB-sTDA calibration of (a) SCOP-PCQC S_1 , (b) SCOP-PCQC T_1 , (c) QM-symex-10k S_1 , and (d) QM-symex-10k T_1 energies, using Class 2 cross-validation (CV) models. 10-fold CV was conducted, meaning all data points shown are test points predicted by an ML model trained on the other 90% of data. Red dots are original data with no calibration, green dots are linearly calibrated data, and blue dots are calibrated with ML. Inlaid box shows quantitative measurements of accuracy for original, linearly calibrated, and ML calibrated data. (Best R^2 is 1, while best MAE is 0.)

high S_1 energies, the linear calibration tends to overcorrect, while for low S_1 energies the linear calibration undercorrects. In contrast, the ML model is more flexible and adequately corrects models in both regions. For T_1 energies, the ML model performs similarly to linear calibration with both MAE and R^2 metrics. This is likely because xTB-sTDA nearly always over-predicts the T_1 energy, so calibrating it only requires shifting in one direction, which makes linear calibration sufficient for the task. For S_1 energies, there are both instances of over- and under-prediction, which motivates the need for an ML model. However, there is clearly room for improvement in these results, as the ML MAE is still high.

Two avenues of improvement were pursued - first, using a larger training set, and second, adding additional input data to the ML model. As discussed in the Methods section, a larger training set was generated using AL to sample regions of chemical space not represented in the 20k training set. The expanded ML model uses 200k AL-chosen molecules plus all QM-symex molecules (120k), added to the initial 20k training set, for a total of approximately 300k molecules. To distinguish the two ML models generated, the 20k model is named xTB-ML-20k while the expanded 300k model is xTB-ML-300k. The second improvement explicitly included the

FIG. 6. Plot of xTB calibration of the 143 MOPSSAM molecules for (a) S_1 and (b) T_1 energies. Red dots are original data with no calibration, green dots are linearly calibrated data, and blue dots are calibrated with ML. Training data was the 20k molecules in SCOP-PCQC + QM-symex-10k, and test data was the 143 molecules shown here. Inlaid boxes show quantitative measurements of accuracy for original, linearly calibrated, and ML calibrated data.

xTB-sTDA calculated energy as an input to the ML model. As discussed previously, for these models, called Class 2, the xTB-sTDA energy was concatenated to the generated molecular features during the training process.

The performance of Class 1 and 2 xTB-ML-20k and xTB-ML-300k models are compared for several external datasets: MOPSSAM 143, MOPSSAM 1000, INDIT, VerdeDB, and PCQC-AL. Figures replicating Figure 6 for the various datasets are available in Supplementary Information Section IV, with primary accuracy metrics presented in Figure 7.

As seen, all ML models improve raw xTB-sTDA values, but to different extents depending on the training and test set considered. For the MOPSSAM and PCQC AL datasets, using larger datasets with more input data generally improves results. This result is intuitive, as more data allows the ML model to learn more about patterns in the datasets. The lowest MAE obtained was 0.08 eV using the Class 2 xTB-ML-300k model on T_1 energies of MOPSSAM 143.

For the INDIT and VerdeDB datasets, the results are less intuitive. In these cases, the Class 1 xTB-ML-20k ML model, i.e. the model trained on a smaller training set with less input data, performs better. There are several reasons this could be happening. The INDIT and VerdeDB datasets are composed specifically of molecules relevant to photon conversion or green chemistry applications. Similarly, the xTB-ML-20k dataset is composed primarily of literature scraped molecules intended for TTA/SF applications, so it is more likely to include molecules similar to those in INDIT or VerdeDB. Although the training set size in xTB-ML-300k is larger, the method of expansion through active learning specifically includes molecules significantly different than the 20k training set, so the model may not increase in accuracy for photon conversion molecules, i.e. is not backwards compatible. In terms of why the Class 2 models perform worse than Class 1 models, this could be related to the nature of the training sets used. Both the 20k and 300k training sets have only a few low-excited state energy molecules.

FIG. 7. Performance of xTB-ML-20k and xTB-ML-300k on various external test datasets, compared to raw xTB-sTDA energies. Two classes of model are tested, Class 1 with only molecular structure as input and Class 2 with both molecular structure and xTB-sTDA energy as input. Test datasets are completely different than the training datasets, and are split into 3 categories: MOPSSAM (for comparison to previous work in xTB-sTDA calibration), datasets relevant for TTA/SF (INDT and VerdeDB), and a broader applicability dataset (PCQC AL). Performance is measured with the MAE, which is calculated with TD-DFT (B3LYP) as reference.

Therefore, the calibration in this region may be inaccurate. As seen in Supplementary Figure S5, the INDT dataset is primarily composed of low-energy molecules. Therefore, if the energy is localized by providing the sTDA energy, the molecule may undergo an inaccurate calibration. This localization is minimized if the calibration is done solely based on molecular structure, allowing a more accurate calibration.

Most of the best-performing ML models result in an MAE of less than 0.20 eV. However, the one dataset with large MAEs is the PCQC-AL T₁ energies. As seen in Supplementary Figure S7, most of the calculated energies follow a general linear trend. However, there is a large cluster of molecules distinctly separated from the rest that is inflating the MAE. This is expected from the AL workflow, which selects molecules difficult to predict with the existing model. Naturally, if a subset of these molecules were included in the training set, the overall MAE would likely improve drastically. Regardless, while this test dataset has large errors, by design these are the largest errors one can obtain, and for general PCQC molecules the error should be lower.

Overall, these results show that the ML models significantly improve raw xTB-sTDA calculated values. In most cases, the best-performing ML model reduces the MAE by more than half. Further, while not shown in Figure 7, the ML models also consistently outperform linear calibration, showing the benefits of a higher-order calibration. We have thus shown that machine learned calibrations can help improve the accuracy of xTB-sTDA results over a wide variety of datasets, when compared to a TD-DFT (B3LYP) reference. We can now use these models for various applications.

C. Applications of xTB-sTDA calibration

1. Direct vs. calibration ML models

Intuitively, we would expect a ML model trained to directly predict TD-DFT results to perform worse than a calibration model where the baseline method does most of the work and the calibrator simply shifts the result in the right direction. This approach is particularly important for out-of-domain predictive performance. Because supervised ML is a data-driven method, it may have poor performance on molecules distinctly different than those in the training set. In contrast, xTB-sTDA is data-agnostic, so it should give reasonable results regardless, and ML should slightly improve results through calibration.

To prove this, we consider a Class 1 ML model trained on the 20k molecules in SCOP-PCQC and QM-symex-10k, but instead of being trained on the error between TD-DFT and xTB-sTDA (as xTB-ML-20k is), the new model, called TDDFT-ML, is trained directly on TD-DFT values. (Note that a direct Class 2 ML model would give equivalent results to a calibration Class 2 model, since the sTDA energy is provided as input.) xTB-ML-20k and TDDFT-ML are then tested on the MOPSSAM 143 dataset, with one directly predicting values, and the other predicting the errors which are added to the sTDA energies to get the final calibrated values. The results of this analysis are shown in Figure 8.

As seen, the performance of the directly trained ML model is vastly lower than the ML-calibrated xTB data. The ML model performs worse than xTB itself in some metrics, and the ML model performs worse than the linear calibration model in all metrics. Thus, calibrating xTB with

FIG. 8. Comparing direct vs. calibration ML models for MOPSSAM. ‘orig’ data is uncalibrated xTB-sTDA data, ‘direct ML’ indicates results of directly predicting TD-DFT data with ML, and ‘xTB ML’ indicates results of calibrating xTB-sTDA values with Class 1 xTB-ML-20k.

ML gives much higher accuracy than using ML to directly predict energies. The benefit to a direct ML model is computational speed, as it can screen approximately 2 orders of magnitude more molecules in a given time period than xTB-ML. However, our goal is to attain approximately the same accuracy as TD-DFT methods, so a direct ML model would not be useful. From the above analysis, our assumption of the improved performance of a calibration ML model is upheld. We now apply our generated calibration ML models for high-throughput screening and chemical space mapping.

2. High-throughput screening for spectral conversion

As discussed in the Introduction, one of the motivations of developing this ML calibration is fast and accurate high-throughput virtual screening (HTVS) of spectral conversion materials. The two spectral conversion techniques of interest are triplet-triplet annihilation (TTA) up-conversion and singlet fission (SF) down-conversion. In general terms, TTA involves two sensitizer molecules that absorb low-energy light and transfer their energy to a single emitter molecule which then re-emits the high-energy light. SF involves two emitter molecules, where one absorbs high-energy light and transfers half its energy to a neighboring emitter, and both re-emit low-energy light. In such molecules, the S_1 excited state is usually involved in absorption/emission while T_1 is typically used for energy transfer. The excited state energy levels of sensitizers and emitters must be well-aligned for efficient spectral conversion – figures of merit (FOMs) to evaluate this alignment are:

$$\text{FOM}_{\text{sens}} = \begin{cases} 0 & S_1 < T_1 \\ e^{-|1 - \frac{S_1}{T_1}|} & S_1 \geq T_1 \end{cases} \quad (1)$$

$$\text{FOM}_{\text{emit}} = e^{-|2 - \frac{S_1}{T_1}|} \begin{cases} \text{SF} & S_1 > 2T_1 \\ \text{TTA} & S_1 < 2T_1 \end{cases} \quad (2)$$

For sensitizers, the first check is if the energies are invalid, i.e. if the T_1 is greater than the S_1 . Then, molecules with S_1 close to T_1 are rated higher. Emitters can be separated into those suitable for SF (singlet more than twice triplet) or TTA (singlet less than twice triplet). The same FOM formula is used for both cases, where S_1 close to twice T_1 is desirable. By ensuring the ratios are as close as possible to ideal, we ensure there is minimal loss in energy. Note that properties related to absorption-to-emission likelihood such as oscillator strength, triplet-triplet energy transfer probability, TTA probability, and others are also important, but are not considered in the present analysis which focuses on optimizing excited state energy level alignment.

We screen the 250k molecules considered in Wilbraham et al.¹⁶ for sensitizers and emitters, to demonstrate the applicability of xTB-ML to HTVS. We use the Class 2 xTB-ML-300k model as it is the most accurate for the MOPSSAM dataset. First, we calibrate S_1 and T_1 energies using the ML model, and compare the results to the linear calibration done in the original work. The results are shown in Figure 9(a,b).

For S_1 , the linear calibration is minimal. The ML calibration remains centered around the raw data for mid- to high- energies, but changes more drastically at low energies. This mirrors the previous discussion of Figure 6, where the linear calibration either over- or under-corrects, but the ML model is more flexible. For T_1 calibration, at mid- to high-energies, both linear and ML calibration shift the energy down, reflecting the tendency of xTB-sTDA to consistently overestimate T_1 energies. For low T_1 energies, the ML model increases the raw energy, suggesting sTDA tends to sometimes spuriously calculate low T_1 energies which can be corrected with ML. Note that because TD-DFT data was not calculated for these 250k molecules, we cannot compare the calibration to ground truth, but based on the metrics presented in Figure 7, it is likely the ML-calibrated values are more accurate.

Now that we have both S_1 and T_1 energies calculated for 250k molecules with xTB-ML, we can identify potential sensitizers and emitters, using the FOMs defined in Equations 1 and 2. Figure 9(c,d) shows the results of screening molecules for potential sensitizers and emitters.

As seen, there are several molecules that would function as potential sensitizers and emitters for photon conversion. The suggested molecules could then be verified with higher-accuracy techniques such as TD-DFT to confirm their suitability. Note here the importance of accuracy for a first-pass screening methodology such as xTB-sTDA. If the uncalibrated results were used, likely several suggested molecules would not be suitable (false positives), and several suitable molecules would not be suggested (false negatives).

FIG. 9. (a,b) Plots of 250k molecules showing difference in calibrated (a) S_1 and (b) T_1 with ML model compared to linear model. Red dots are without calibration, yellow dots are with linear calibration, and blue dots are with ML calibration. (c,d) Plots of 250k molecules showing S_1 and T_1 energies, colored with FOM for (c) sensitizers and (d) emitters. Target properties are correct ratio of S_1/T_1 as defined in Equations 1 and 2

Using xTB-ML improves the quality of suggestions by reducing both of these rates.

We have therefore used the Class 2 xTB-ML-300k model to make quick, relatively accurate calculations for S_1 and T_1 energies, and have used the results to screen for potential sensitizers and emitters. This screening was relatively fast as the dataset size was small (250k) and xTB-sTDA results were already provided by Wilbraham et al.¹⁶ For larger datasets on the order of millions (PubChem)⁴⁶ or billions of molecules (GDB-17)⁴⁷, running xTB-ML becomes expensive. A more intelligent sampling technique (such as active learning) could be used to screen such large databases, and this is an avenue of future work.

3. Mapping inaccuracies of xTB-sTDA in chemical space

Since our ML model predicts the error in xTB-sTDA, an interesting application is to map the error in S_1 and T_1 calculations in a global chemical space, to see if there are some areas where xTB systematically over- or under-estimates, or areas where xTB is projected to be fairly accurate. For this analysis, we use Class 1 xTB-ML-300k, as it is shown to be accurate in the general chemical domain and does not require xTB-sTDA computations, so large-scale computations can be done quickly with ML.

The first requirement is to generate a global chemical map.

FIG. 10. Global chemical space maps. Plots of xTB (a) S_1 and (b) T_1 errors in global chemical space. (c) Clustering of molecules in global chemical space. (d) Number of molecules per cluster in global chemical space.

We used UMAP⁴⁸ to embed the high-dimensional molecular data into 2 dimensions, using the Jaccard-Tanimoto similarity between Morgan fingerprints of molecules for proximity. We then colored the global chemical space map in 3 different ways, as shown in Figure 10.

For the first two plots (Figure 10(a) and (b)), we used our ML model to predict the error in xTB-sTDA. Here, the error is defined as:

$$\text{error} = \text{true energy} - \text{xTB-sTDA energy} \quad (3)$$

so a negative error implies xTB-sTDA is over-predicting the excited state energy. As seen in Figure 10(a), there are distinct regions where xTB-sTDA over-predicts S_1 (right side), regions where xTB-sTDA has reasonable accuracy (top left and center), and regions where it under-predicts (bottom left and top). In general, most molecules are within ± 0.5 eV error.

In contrast, for the T_1 energy, xTB-sTDA over-predicts for almost all molecules, as seen in Figure 10(b). Note that the scale in this plot is shifted from -0.5 – 0.5 eV (as in S_1) to 0 – -1.0 eV, to make the distribution of errors clearer. Only a few scattered molecules are under-predicted by xTB-sTDA and are colored red, and all other molecules are over-predicted. Similar to S_1 , xTB-sTDA over-predicts T_1 for most molecules on the right side, and gets reasonable accuracy on molecules in the middle and top left. T_1 is also over-predicted on a cluster of molecules on the bottom left and top.

Next, we used HDBSCAN⁴⁹ to cluster the molecules based on proximity, as shown in Figure 10(c). HDBSCAN is a soft clustering, not creating distinct categories but instead giving molecules a rating between 0 and 10 (or -1 for no cluster, as approximately 1/3 of the molecules were unable to be clustered). We can see that HDBSCAN effectively clusters molecules in space, with most molecules in close proximity included in the same cluster. Some of the clusters themselves

FIG. 11. Mean errors for S_1 and T_1 energies of molecules in global chemical space. (a) Mean S_1 error, (b) mean T_1 error, (c) mean absolute S_1 error, and (d) mean absolute T_1 error. "Absolute" error takes the absolute value of errors before averaging them.

are spread out across space, such as the purple cluster that includes many molecules along the edge of the global space. Note that this is a dataset-agnostic clustering, as the clustering algorithm only sees molecular information and no labelled data. More details about the HDBSCAN algorithm can be found in their paper⁴⁹ and website.⁵⁰

The natural question is whether each cluster as defined by HDBSCAN has a particular error associated with it. For example, it seems that xTB-sTDA does a relatively good job for the red cluster, but over-predicts energies for molecules in light blue, purple, and orange clusters. In contrast, the dark green and red clusters seem to have low errors. Although HDBSCAN is a soft clustering, we can categorize molecules into 100 distinct clusters based on the number assigned to them, as well as 2 additional clusters (1 each for unclustered molecules and for outliers). Figure 11 quantifies the mean errors for S_1 and T_1 energies for each cluster.

Subplots (a) and (b) show the mean errors (ME) of S_1 and T_1 , while (c) and (d) show mean absolute errors (MAE). As seen, the red/yellow/green clusters are likely to have low error, while the purple/dark green clusters have high error. While this analysis is generally useful, the mapping and clustering approach requires knowing the location and cluster categorization of a specific molecule in global chemical space. Oftentimes, this is not known, or would require significant computation.

Instead, it would be beneficial to have some chemical intuition of accuracy based on the molecular structure, to have greater confidence in xTB-sTDA calculations, or to know to use the ML model or consider other computational techniques.

FIG. 12. Grid of molecular substructures over-represented in molecules in each error category as predicted by the ML model, for (a) low S_1 error, (b) low T_1 error, (c) high S_1 overestimation, (d) high T_1 overestimation, (e) high S_1 underestimation, and (f) high T_1 underestimation. According to RDKit, blue atoms are the center atoms, yellow atoms are aromatic atoms, dark gray atoms are aliphatic ring atoms, and light gray atoms/bonds are connectivity invariants.

To this end, we can identify substructures that are more likely to be present in low-error or high-error molecules.

We first use our ML model generated above to predict the S_1 and T_1 error between xTB and TD-DFT for 1M molecules randomly subsampled from PCQC. We then categorize the molecules based on the predicted error as follows:

$$\text{Cat}_{S_1} = \begin{cases} \text{Low} & |S_{1, \text{err}}| < 0.05 \\ \text{High}_{\text{Under}} & S_{1, \text{err}} > 0.5 \\ \text{High}_{\text{Over}} & S_{1, \text{err}} < -0.5 \end{cases} \quad (4)$$

$$\text{Cat}_{T_1} = \begin{cases} \text{Low} & |T_{1, \text{err}}| < 0.05 \\ \text{High}_{\text{Under}} & T_{1, \text{err}} > 0 \\ \text{High}_{\text{Over}} & T_{1, \text{err}} < -1.0 \end{cases}$$

where "under" refers to xTB underestimating the energy while "over" refers to overestimating (note the error definition in Equation 3). For both T_1 and S_1 error, we define low error as $< \pm 0.05$ eV. However, for defining high error, for T_1 we shift the bounds down by 0.5 to reflect the distribution of errors, as seen in Figure 10(c).

We can conduct substructure analysis on the low and high overestimation categories (ignoring high underestimation due

to low fraction of molecules) to know when to trust the xTB-sTDA results, or when to expect exceptionally high errors. We use molZ⁵¹ to analyze which substructures are over-represented in each category. The results of this substructure analysis are shown in Figure 12.

From these plots, a few patterns become evident. Low error molecules are more likely to be aromatic, potentially with sequential attached rings, for both S_1 and T_1 . In contrast, S_1 high overestimation, S_1 high underestimation, and T_1 high underestimation molecules are likely to be not aromatic, with some unconventional molecular structures included in these groups. In particular, the S_1 overestimation group includes 5 and 7 C-ring molecules, and both S_1 and T_1 underestimation include charged N atoms. There are some aromatic substructures in the T_1 overestimation molecules, but they are attached to the bulk structure with a rotatable bond. This overestimation could be a result of the 3D structure generation, since only limited conformer analysis is conducted and potentially the lowest energy conformer was not achieved. To clarify the effect of this versus an inherent inaccuracy in the excited state energy calculation of xTB-sTDA, a more intensive conformer search could be an avenue of future work.

Overall, these predictions can be used as guides for the accuracy of xTB-sTDA in calculating excited state energies.

4. CC2 calibration of xTB-sTDA

Finally, to show the generalizability of the methodology presented here, we choose a different reference technique beyond TD-DFT, namely CC2.⁵² CC2 is known to better predict excitation energies than TD-DFT, but its computational cost is often prohibitively expensive.²⁴ We use the CC2 S_1 values compiled in QM8,^{19,53} randomly sampling 10k molecules as the training set and using the other 11.5k as the test set. xTB-sTDA values were generated using the same methodology as before. For ML model training, because of the smaller dataset, we use 20-fold cross-validation with 95%/5% train/validation splits. This helps ensure most of the data is used in training. The new model is termed xTB-CC-ML to distinguish it from the previously generated xTB-ML models.

Figure 13(a) shows the results of the comparison, with measurements of accuracy for both methods presented in the inlaid box. As seen, adding the ML calibration to xTB-sTDA results vastly improves results, reducing the MAE by 66%. For comparison, Figure 13(b) shows the results of TD-DFT calculations using PBE0⁵⁴ and CAM-B3LYP on the same test set of molecules. As seen, xTB-CC-ML has higher accuracy than TD-DFT calculations for the 11.5k test set, using either R2 or MAE as the metric.

Note that this figure also justifies the main calibration methodology presented in this section, of calibrating xTB-sTDA against TD-DFT. While xTB-sTDA was initially parameterized against mostly CC2 calculations, its accuracy is lower than TD-DFT with hybrid functionals such as PBE0, motivating the calibration presented.

FIG. 13. (a) Plot of xTB calibration against CC2 with 10k training set and 11.5k test set (shown). Red dots indicate original xTB calculations while green dots indicate calibrated xTB data. (b) Plot of TD-DFT calculated values against CC2 values, for accuracy comparison. Black dashed line in both plots indicates $x = y$ line. (c) Plot of xTB calibration accuracy as a function of training size, tested on the remaining molecules in the 21.5k dataset. MAE of the test set is calculated with CC2 values as reference. The 8 red dots correspond to the 8 ML models generated for xTB-sTDA to CC2 calibration.

To test the impact of training size on accuracy, 8 different ML models were generated with training sizes ranging from 100 to 15,000. The models were then predicted on the remaining molecules in QM8 not used in the training set. The MAE of the test set (against CC2 values) was compared to the MAE of PBE0/def2-TZVP and CAM-B3LYP/def2-TZVP, as shown in Figure 13(c). As seen, a training size of approximately 750 molecules allows xTB-sTDA to achieve similar accuracy to PBE0. It is more difficult to match CAM-B3LYP, but this is achieved at a training size of around 3750. At the largest training size considered (15k), xTB-CC-ML vastly outperforms both TD-DFT techniques, with a 47% reduction in MAE for PBE0 and 27% reduction in MAE for CAM-B3LYP.

These are promising results; however, the xTB-CC-ML model may not be as generalizable as xTB-ML, due to the smaller (10k), less diverse (only small molecules with up to 8 heavy atoms) training set. For this reason, xTB-ML was used in the prior sections instead of xTB-CC-ML. Regardless, xTB-CC-ML serves as an interesting proof of concept that can be expanded further in the future, perhaps with additional CC2 calculations on more diverse molecules.

IV. CONCLUSIONS

We have presented a methodology for calibrating a high-throughput computational chemistry technique (xTB-sTDA) against a high-accuracy one (TD-DFT) using machine learning. We first decided on Chemprop’s MPNN as the ML architecture of choice, then generated a training set using literature scraping of relevant molecules from abstracts (SCOP-PCQC) and an existing excited state database (QM-symex-10k). We also generated an expanded training set using active learning. We built two models based on these training sets (xTB-ML-20k and xTB-ML-300k).

We then generated blind test sets from: a study on linear calibration of xTB-sTDA results (MOPSSAM),¹⁶ a study on singlet fission materials (INDT),⁷ a database of molecules relevant to green chemistry (VerdeDB),³² and a dataset to test the broad applicability of the model (PCQC AL).³¹ On these external datasets, the ML calibration models outperformed both raw xTB-sTDA and linear calibration, oftentimes significantly. Averaging the best MAE over all external test sets (both S_1 and T_1) excluding PCQC AL gave an MAE of 0.14 eV, compared to 0.38 eV for xTB-sTDA. Including PCQC AL gave an average MAE of 0.57 eV, compared to 0.83 eV for xTB-sTDA. If xTB-ML is used as the first step in a high-throughput screening process instead of raw xTB-sTDA outputs, its low error can help ensure that all relevant molecules are selected, and vice versa.

After evaluating the performance of xTB-ML, we then used the model for four applications. First was comparing the xTB-ML model against directly predicted energies with ML, showing the xTB-ML model had better accuracy (0.17 eV vs. 0.25 eV MAE). Second was rapidly screening 250k molecules for suitability as TTA sensitizers and TTA/SF emitters. Third was mapping inaccuracies of xTB-sTDA in chemical space, using the ML model to predict errors. This was used to see which regions of chemical space xTB-sTDA has high errors in. S_1 errors were small, with most molecules being within 0.5 eV. There were clear regions where xTB-sTDA overpredicted S_1 , but only a few for underprediction. T_1 energies were generally overpredicted, with most molecules being between 0 and 1 eV below TD-DFT values. Global chemical space mapping provides another method of predicting xTB-sTDA error, by calculating which cluster a molecule belongs to and referencing the MAE of that cluster. Properties of low-error molecules were also evaluated, finding non-aromatic molecules are likely to have higher error. The final application was generalizing the methodology to calibrate xTB-sTDA against coupled cluster theory, and generating a new xTB-CC-ML model. The calibrated xTB-CC-ML values had high accuracy (0.15 eV MAE), out-performing TD-DFT values calculated with PBE0 (0.26 eV MAE) and CAM-B3LYP (0.19 eV MAE).

There are several avenues for future work. First is improving the ML model architecture. While Chemprop’s MPNN outperformed other ML models, primarily due to its advanced featurization, only the 2D molecular structure and xTB-sTDA energy were provided as input. Since the 3D structure is known, including this information would likely

improve performance. Another improvement to the ML workflow would be to conduct a more intensive conformer search. While OpenBabel’s gen3D function includes a search for 200 conformers, these may not include the lowest energy conformer, thus reducing the accuracy of the xTB portion of the workflow. Using a conformer searching tool such as CREST⁵⁵ would be more comprehensive, although the computation time added may detract from the high-throughput nature of the xTB-ML process.

Beyond higher level first-principles data, the calibration models could be further extended to experimental data. However, this would be time-consuming due to the requirement of real-world measurements. There have been a few previous studies in calibrating TD-DFT against experimental values,^{21,22} but these used only small experimental datasets. There is a potential here to apply techniques such as text mining to extract experimental excited state data from published papers, though the differences in reporting may make this difficult.

The models can also be extended to other applications. For example, the graph-based genetic algorithm (GB-GA) developed by Jensen currently uses uncalibrated xTB-sTDA for excitation energies.⁵⁶ xTB-ML models could increase accuracy, and if calculations are sufficiently parallelized, this process could result in thousands of high-quality candidates being rapidly generated.

Lastly, although xTB-ML is significantly faster than first-principles methods, it is still too slow to screen millions of molecules. As stated previously, with our setup xTB-ML can calculate excited states of approximately 1500 molecules per hour (parallelized over 4 computer nodes), for molecules with < 50 heavy atoms such as those in PCQC. Therefore it would take over 3 months to calculate all 3.5M molecules in PCQC. While a definite improvement over TD-DFT (over 3 years), this is still slow. Expanding to larger databases with bigger molecules would increase runtime even further. Therefore, as a final direction, an optimized workflow using active learning could be implemented, intentionally searching for molecules with certain desired properties.

V. DATA AVAILABILITY

The raw data and code to reproduce the figures presented in this paper are available in a repository on GitHub.⁵⁷ Trained ML models and prediction data are also available on GitHub.⁵⁷

VI. CODE AVAILABILITY

The code for the calibration workflow presented here is available in a repository on GitHub.⁵⁸ The scripts for running TD-DFT, running xTB-sTDA, training the ML model, and using it for predictions are available.

ACKNOWLEDGMENTS

S.V. was supported by the Marshall Scholarship. This work made use of the CX1/2 high-performance computing clusters at Imperial College London.⁵⁹ The authors would like to thank Daniel Davies for useful discussions surrounding machine learning, Kyle Swanson for help and advice about using Chemprop, Jiali Li for helpful discussions regarding active learning, and Martijn Zwijnenburg and Stefano Angioletti-Uberti for fruitful discussions about computational chemistry. We are also grateful to the UK Materials and Molecular Modelling Hub for computational resources, which is partially funded by EPSRC (EP/P020194/1 and EP/T022213/1).

VII. REFERENCES

- ¹Y.-C. Cheng and G. R. Fleming, "Dynamics of Light Harvesting in Photosynthesis," *Annual Review of Physical Chemistry* **60**, 241–262 (2009), [_eprint: https://doi.org/10.1146/annurev.physchem.040808.090259](https://doi.org/10.1146/annurev.physchem.040808.090259).
- ²Q. Wang, R. W. Schoenlein, L. A. Peteanu, R. A. Mathies, and C. V. Shank, "Vibrationally Coherent Photochemistry in the Femtosecond Primary Event of Vision," *Science* **266**, 422–424 (1994), publisher: American Association for the Advancement of Science.
- ³T. Schultz, E. Samoylova, W. Radloff, I. V. Hertel, A. L. Sobolewski, and W. Domcke, "Efficient Deactivation of a Model Base Pair via Excited-State Hydrogen Transfer," *Science* **306**, 1765–1768 (2004), publisher: American Association for the Advancement of Science.
- ⁴M. B. Smith and J. Michl, "Singlet Fission," *Chemical Reviews* **110**, 6891–6936 (2010), publisher: American Chemical Society.
- ⁵J. Zhao, S. Ji, and H. Guo, "Triplet–triplet annihilation based upconversion: from triplet sensitizers and triplet acceptors to upconversion quantum yields," *RSC Advances* **1**, 937–950 (2011), publisher: The Royal Society of Chemistry.
- ⁶O. El Bakouri, J. R. Smith, and H. Ottosson, "Strategies for Design of Potential Singlet Fission Chromophores Utilizing a Combination of Ground-State and Excited-State Aromaticity Rules," *Journal of the American Chemical Society* **142**, 5602–5617 (2020), publisher: American Chemical Society.
- ⁷K. J. Fallon, P. Budden, E. Salvadori, A. M. Ganose, C. N. Savory, L. Eyre, S. Dowland, Q. Ai, S. Goodlett, C. Risko, D. O. Scanlon, C. W. M. Kay, A. Rao, R. H. Friend, A. J. Musser, and H. Bronstein, "Exploiting Excited-State Aromaticity To Design Highly Stable Singlet Fission Materials," *Journal of the American Chemical Society* **141**, 13867–13876 (2019), publisher: American Chemical Society.
- ⁸F. Spiegelman, N. Tarrat, J. Cuny, L. Dontot, E. Posenitskiy, C. Martí, A. Simon, and M. Rapacioli, "Density-functional tight-binding: basic concepts and applications to molecules and clusters," *Advances in Physics: X* **5**, 1710252 (2020), publisher: Taylor & Francis [_eprint: https://doi.org/10.1080/23746149.2019.1710252](https://doi.org/10.1080/23746149.2019.1710252).
- ⁹M. Elstner, D. Porezag, G. Jungnickel, J. Elsner, M. Haugk, T. Frauenheim, S. Suhai, and G. Seifert, "Self-consistent-charge density-functional tight-binding method for simulations of complex materials properties," *Physical Review B* **58**, 7260–7268 (1998), publisher: American Physical Society.
- ¹⁰C. Bannwarth, E. Caldeweyher, S. Ehlert, A. Hansen, P. Pracht, J. Seibert, S. Spicher, and S. Grimme, "Extended tight-binding quantum chemistry methods," *WIREs Computational Molecular Science* **11**, e1493 (2021), [_eprint: https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.1493](https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.1493).
- ¹¹S. Grimme, "A simplified Tamm-Dancoff density functional approach for the electronic excitation spectra of very large molecules," *The Journal of Chemical Physics* **138**, 244104 (2013), publisher: American Institute of Physics.
- ¹²C. Fang, B. Oruganti, and B. Durbbeej, "How Method-Dependent Are Calculated Differences between Vertical, Adiabatic, and 0–0 Excitation Energies?" *The Journal of Physical Chemistry A* **118**, 4157–4171 (2014), publisher: American Chemical Society.
- ¹³S. Grimme and C. Bannwarth, "Ultra-fast computation of electronic spectra for large systems by tight-binding based simplified Tamm-Dancoff approximation (sTDA-xTB)," *The Journal of Chemical Physics* **145**, 054103 (2016), publisher: American Institute of Physics.
- ¹⁴L. Wilbraham, R. S. Sprick, K. E. Jelfs, and M. A. Zwijnenburg, "Mapping binary copolymer property space with neural networks," *Chemical Science* **10**, 4973–4984 (2019), publisher: The Royal Society of Chemistry.
- ¹⁵L. Wilbraham, E. Berardo, L. Turcani, K. E. Jelfs, and M. A. Zwijnenburg, "High-Throughput Screening Approach for the Optoelectronic Properties of Conjugated Polymers," *Journal of Chemical Information and Modeling* **58**, 2450–2459 (2018), publisher: American Chemical Society.
- ¹⁶L. Wilbraham, D. Smajli, I. Heath-Apostolopoulos, and M. A. Zwijnenburg, "Mapping the optoelectronic property space of small aromatic molecules," *Communications Chemistry* **3**, 1–9 (2020), number: 1 Publisher: Nature Publishing Group.
- ¹⁷I. Heath-Apostolopoulos, L. Wilbraham, and M. A. Zwijnenburg, "Computational high-throughput screening of polymeric photocatalysts: exploring the effect of composition, sequence isomerism and conformational degrees of freedom," *Faraday Discussions* **215**, 98–110 (2019), publisher: The Royal Society of Chemistry.
- ¹⁸I. Heath-Apostolopoulos, D. Vargas-Ortiz, L. Wilbraham, K. E. Jelfs, and M. A. Zwijnenburg, "Using high-throughput virtual screening to explore the optoelectronic property space of organic dyes; finding diketopyrrolopyrrole dyes for dye-sensitized water splitting and solar cells," *Sustainable Energy & Fuels* **5**, 704–719 (2021), publisher: The Royal Society of Chemistry.
- ¹⁹R. Ramakrishnan, M. Hartmann, E. Tapavicza, and O. A. von Lilienfeld, "Electronic spectra from TDDFT and machine learning in chemical space," *The Journal of Chemical Physics* **143**, 084111 (2015), publisher: American Institute of Physics.
- ²⁰R. Pollice, P. Friederich, C. Lavigne, G. d. P. Gomes, and A. Aspuru-Guzik, "Organic molecules with inverted gaps between first excited singlet and triplet states and appreciable fluorescence rates," *Matter* **4**, 1654–1682 (2021).
- ²¹X. Wang, L. Wong, L. Hu, C. Chan, Z. Su, and G. Chen, "Improving the Accuracy of Density-Functional Theory Calculation: The Statistical Correction Approach," *The Journal of Physical Chemistry A* **108**, 8514–8525 (2004), publisher: American Chemical Society.
- ²²H. Li, L. Shi, M. Zhang, Z. Su, X. Wang, L. Hu, and G. Chen, "Improving the accuracy of density-functional theory calculation: The genetic algorithm and neural network approach," *The Journal of Chemical Physics* **126**, 144101 (2007), publisher: American Institute of Physics.
- ²³S. Grimme, L. Goerigk, and R. F. Fink, "Spin-component-scaled electron correlation methods," *WIREs Computational Molecular Science* **2**, 886–906 (2012), [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.1110](https://onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.1110).
- ²⁴C. Suellen, R. G. Freitas, P.-F. Loos, and D. Jacquemin, "Cross-Comparisons between Experiment, TD-DFT, CC, and ADC for Transition Energies," *Journal of Chemical Theory and Computation* **15**, 4581–4590 (2019).
- ²⁵P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch, "Ab Initio Calculation of Vibrational Absorption and Circular Dichroism Spectra Using Density Functional Force Fields," *The Journal of Physical Chemistry* **98**, 11623–11627 (1994), publisher: American Chemical Society.
- ²⁶D. Jacquemin, V. Wathelet, E. A. Perpète, and C. Adamo, "Extensive TD-DFT Benchmark: Singlet-Excited States of Organic Molecules," *Journal of Chemical Theory and Computation* **5**, 2420–2435 (2009), publisher: American Chemical Society.
- ²⁷T. Yanai, D. P. Tew, and N. C. Handy, "A new hybrid exchange–correlation functional using the Coulomb-attenuating method (CAM-B3LYP)," *Chemical Physics Letters* **393**, 51–57 (2004).
- ²⁸A. M. Grabar and B. Ośmiałowski, "Benchmarking Density Functional Approximations for Excited-State Properties of Fluorescent Dyes," *Molecules* **26**, 7434 (2021), number: 24 Publisher: Multidisciplinary Digital Publishing Institute.
- ²⁹I. Y. Zhang, J. Wu, and X. Xu, "Extending the reliability and applicability of B3LYP," *Chemical Communications* **46**, 3057–3070 (2010), publisher: The Royal Society of Chemistry.

- ³⁰G. Montavon, M. Rupp, V. Gobre, A. Vazquez-Mayagoitia, K. Hansen, A. Tkatchenko, K.-R. Müller, and O. A. v. Lilienfeld, "Machine learning of molecular electronic properties in chemical compound space," *New Journal of Physics* **15**, 095003 (2013), publisher: IOP Publishing.
- ³¹M. Nakata and T. Shimazaki, "PubChemQC Project: A Large-Scale First-Principles Electronic Structure Database for Data-Driven Chemistry," *Journal of Chemical Information and Modeling* **57**, 1300–1308 (2017), publisher: American Chemical Society.
- ³²B. G. Abreha, S. Agarwal, I. Foster, B. Blaiszik, and S. A. Lopez, "Virtual Excited State Reference for the Discovery of Electronic Materials Database: An Open-Access Resource for Ground and Excited State Properties of Organic Molecules," *The Journal of Physical Chemistry Letters* **10**, 6835–6841 (2019), publisher: American Chemical Society.
- ³³J. Liang, S. Ye, T. Dai, Z. Zha, Y. Gao, and X. Zhu, "QM-symex, update of the QM-sym database with excited state information for 173 kilo molecules," *Scientific Data* **7**, 400 (2020), number: 1 Publisher: Nature Publishing Group.
- ³⁴"Elsevier Developer Portal," (2021).
- ³⁵M. C. Swain and J. M. Cole, "ChemDataExtractor: A Toolkit for Automated Extraction of Chemical Information from the Scientific Literature," *Journal of Chemical Information and Modeling* **56**, 1894–1904 (2016), publisher: American Chemical Society.
- ³⁶S. Kim, P. A. Thiessen, T. Cheng, B. Yu, and E. E. Bolton, "An update on PUG-REST: RESTful interface for programmatic access to PubChem," *Nucleic Acids Research* **46**, W563–W570 (2018).
- ³⁷J. S. Smith, B. Nebgen, N. Lubbers, O. Isayev, and A. E. Roitberg, "Less is more: Sampling chemical space with active learning," *The Journal of Chemical Physics* **148**, 241733 (2018), publisher: American Institute of Physics.
- ³⁸N. M. O'Boyle, M. Banck, C. A. James, C. Morley, T. Vandermeersch, and G. R. Hutchison, "Open Babel: An open chemical toolbox," *Journal of Cheminformatics* **3**, 33 (2011).
- ³⁹C. Bannwarth, S. Ehlert, and S. Grimme, "GFN2-xTB—An Accurate and Broadly Parametrized Self-Consistent Tight-Binding Quantum Chemical Method with Multipole Electrostatics and Density-Dependent Dispersion Contributions," *Journal of Chemical Theory and Computation* **15**, 1652–1671 (2019), publisher: American Chemical Society.
- ⁴⁰"sTDA-xTB for ground state calculations," (2021), original-date: 2019-11-27T12:03:12Z.
- ⁴¹B. Ramsundar, P. Eastman, P. Walters, and V. Pande, *Deep learning for the life sciences: applying deep learning to genomics, microscopy, drug discovery and more*, first edition ed. (O'Reilly Media, Sebastopol, CA, 2019) oCLC: on1051083869.
- ⁴²D. K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams, "Convolutional Networks on Graphs for Learning Molecular Fingerprints," in *Advances in Neural Information Processing Systems*, Vol. 28 (Curran Associates, Inc., 2015).
- ⁴³O. Vinyals, S. Bengio, and M. Kudlur, "Order Matters: Sequence to sequence for sets," arXiv:1511.06391 [cs, stat] (2016), arXiv: 1511.06391.
- ⁴⁴K. Yang, K. Swanson, W. Jin, C. Coley, P. Eiden, H. Gao, A. Guzman-Perez, T. Hopper, B. Kelley, M. Mathea, A. Palmer, V. Settels, T. Jaakkola, K. Jensen, and R. Barzilay, "Analyzing Learned Molecular Representations for Property Prediction," *Journal of Chemical Information and Modeling* **59**, 3370–3388 (2019), publisher: American Chemical Society.
- ⁴⁵D. Weininger, "SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules," *Journal of Chemical Information and Computer Sciences* **28**, 31–36 (1988), publisher: American Chemical Society.
- ⁴⁶S. Kim, J. Chen, T. Cheng, A. Gindulyte, J. He, S. He, Q. Li, B. A. Shoemaker, P. A. Thiessen, B. Yu, L. Zaslavsky, J. Zhang, and E. E. Bolton, "PubChem in 2021: new data content and improved web interfaces," *Nucleic Acids Research* **49**, D1388–D1395 (2020).
- ⁴⁷L. Ruddigkeit, R. van Deursen, L. C. Blum, and J.-L. Reymond, "Enumeration of 166 Billion Organic Small Molecules in the Chemical Universe Database GDB-17," *Journal of Chemical Information and Modeling* **52**, 2864–2875 (2012), publisher: American Chemical Society.
- ⁴⁸L. McInnes, J. Healy, and J. Melville, "UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction," arXiv:1802.03426 [cs, stat] (2020), arXiv: 1802.03426.
- ⁴⁹R. J. G. B. Campello, D. Moulavi, and J. Sander, "Density-Based Clustering Based on Hierarchical Density Estimates," in *Advances in Knowledge Discovery and Data Mining*, Lecture Notes in Computer Science, edited by J. Pei, V. S. Tseng, L. Cao, H. Motoda, and G. Xu (Springer, Berlin, Heidelberg, 2013) pp. 160–172.
- ⁵⁰"scikit-learn-contrib/hdbscan," (2022), original-date: 2015-04-22T13:32:37Z.
- ⁵¹L. Wilbraham, "molZ," (2021), original-date: 2020-12-06T13:18:42Z.
- ⁵²O. Christiansen, H. Koch, and P. Jørgensen, "The second-order approximate coupled cluster singles and doubles model CC2," *Chemical Physics Letters* **243**, 409–418 (1995).
- ⁵³R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. von Lilienfeld, "Quantum chemistry structures and properties of 134 kilo molecules," *Scientific Data* **1**, 140022 (2014), number: 1 Publisher: Nature Publishing Group.
- ⁵⁴C. Adamo and V. Barone, "Toward reliable density functional methods without adjustable parameters: The PBE0 model," *The Journal of Chemical Physics* **110**, 6158–6170 (1999), publisher: American Institute of Physics.
- ⁵⁵P. Pracht, F. Bohle, and S. Grimme, "Automated exploration of the low-energy chemical space with fast quantum chemical methods," *Physical Chemistry Chemical Physics* **22**, 7169–7192 (2020), publisher: The Royal Society of Chemistry.
- ⁵⁶J. H. Jensen, "A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space," *Chemical Science* **10**, 3567–3572 (2019), publisher: The Royal Society of Chemistry.
- ⁵⁷shomikverma, "xTB-ML-data," (2022), original-date: 2021-12-17T17:42:28Z.
- ⁵⁸shomikverma, "xTB-ML-workflow," (2022), original-date: 2021-12-17T17:53:32Z.
- ⁵⁹"Imperial College Research Computing Service, DOI: 10.14469/hpc/2232," (2021).