

Cite this: *Chem. Sci.*, 2017, 8, 8419

Machine learning for quantum dynamics: deep learning of excitation energy transfer properties†

Florian Häse, Christoph Kreisbeck* and Alán Aspuru-Guzik *

Understanding the relationship between the structure of light-harvesting systems and their excitation energy transfer properties is of fundamental importance in many applications including the development of next generation photovoltaics. Natural light harvesting in photosynthesis shows remarkable excitation energy transfer properties, which suggests that pigment–protein complexes could serve as blueprints for the design of nature inspired devices. Mechanistic insights into energy transport dynamics can be gained by leveraging numerically involved propagation schemes such as the hierarchical equations of motion (HEOM). Solving these equations, however, is computationally costly due to the adverse scaling with the number of pigments. Therefore virtual high-throughput screening, which has become a powerful tool in material discovery, is less readily applicable for the search of novel excitonic devices. We propose the use of artificial neural networks to bypass the computational limitations of established techniques for exploring the structure-dynamics relation in excitonic systems. Once trained, our neural networks reduce computational costs by several orders of magnitudes. Our predicted transfer times and transfer efficiencies exhibit similar or even higher accuracies than frequently used approximate methods such as secular Redfield theory.

Received 13th August 2017
Accepted 23rd October 2017

DOI: 10.1039/c7sc03542j

rsc.li/chemical-science

1 Introduction

Studying excitation energy transport (EET) has been of great interest across different fields bridging evolutionary biology to solar cell engineering for many years. Especially natural light-harvesting has been the subject of intense research. Pigment–protein complexes exhibit remarkable transport properties which facilitate highly efficient excitation energy transfer across long distances.^{1–4} Thus, identifying working principles that ultimately transform into blueprints for novel nature-inspired excitonic devices is an active research frontier.^{5,6}

Mechanistic studies reveal valuable insight into the microscopic details of EET. Prominent examples are given by studies probing the impact of electronic coherence or non-trivial interactions between excitons and specific vibrational modes on transfer characteristics.^{7–13} However such investigations are tedious since they require sophisticated experimental setups,^{11–16} as well as computationally involved accurate simulations of open-quantum system dynamics.^{7–9,17–20} Further, there are only a few fundamentally different natural light-harvesting complexes from which alone we cannot extract the relation between the structure of an excitonic system and its dynamics in full detail.


In order to relate the dynamics to the underlying structure, it is desirable to investigate a large number of artificially designed excitonic systems. This has been recently addressed in several theoretical works.^{21–24} For example, analyzing perturbations on pigment geometries in the Fenna–Matthews–Olson (FMO) complex revealed that higher transport efficiencies tend to be realized by more compact structures.²⁵ The drawback of these statistical approaches is that they need to run exciton dynamics calculations for ten thousands of randomly generated physically-plausible multi-chromophoric structures. Due to the sheer number of performed dynamics simulations, such an analysis becomes quickly computationally exhaustive, even though less sophisticated methods such as Lindblad equations are used.²⁵

Here, we follow a novel path and leverage concepts from deep learning to bypass the computational demand of established techniques for exploring EET properties (see Fig. 1). Specifically, we train multi-layer perceptrons (MLPs), a class of fully connected feed-forward artificial neural networks to predict average exciton transfer times and overall transfer efficiencies. The input features to the MLPs are hereby given by the parameters of the corresponding Frenkel exciton Hamiltonians.^{26,27} For large scale screening of parameter space, only a fraction of all systems needs to be actually calculated to train the MLPs. Once trained, our neural networks evaluate transfer times just within a few milliseconds and thus bypass the computational demand of established techniques for exploring EET properties, while maintaining sufficiently high prediction accuracy.

Department of Chemistry and Chemical Biology, Harvard University, Cambridge, 02138, USA. E-mail: christophkreisbeck@gmail.com; aspuru@chemistry.harvard.edu; Tel: +1-617-384-8188

† Electronic supplementary information (ESI) available. See DOI: 10.1039/c7sc03542j





We assess the quality of our MLP predictions by comparing the relative error of our predicted transfer times to the relative error made by secular Redfield calculations. The latter is simple to implement and commonly used to avoid the numerical complexity of more accurate HEOM simulations. Our findings demonstrate that MLPs provide a computationally significantly cheaper alternative to secular Redfield computations at comparable or, in most of our examples, even higher accuracy. Results for the FMO complex are summarized in Fig. 1.

A number of studies across many fields in recent years have demonstrated how machine learning models can be utilized to accelerate a variety of computations by several orders of magnitude at a reasonable level of accuracy. For example, Gaussian processes were used to predict formation of free energies for catalyst surface chemistry.⁴² Kernel ridge regression methods were found to accurately predict atomization energies of small molecules.⁴³ Neural networks have been successfully employed for the construction of various forms of transferable and non-transferable atomistic potentials.^{44–46} Protein-ligand binding affinities were accurately predicted by atomic convolutional neural networks,⁴⁷ and multi-layer perceptrons were trained to predict excited state energies in the context of exciton dynamics,⁴⁸ as well as other electronic properties of small molecules.^{43,49}

This journal is © The Royal Society of Chemistry 2017

The study of excitation energy transport typically involves two steps: first, an effective Hamiltonian describing the system parameters needs to be constructed and second, transfer properties need to be computed from this effective Hamiltonian using open quantum system approaches. While some of the authors have successfully applied machine learning techniques to accelerate the construction of effective Hamiltonians by predicting excited state energies of excitonic sites from Coulomb matrices,⁴⁸ to our knowledge there has been no attempt to adapt machine learning models to predict transport properties of open-quantum systems.

In the subsequent sections, we develop a machine learning framework based on multi-layer perceptrons (MLPs) which predict excitation energy transfer properties of excitonic systems based on an effective Hamiltonian rather than obtaining them from computationally expensive quantum dynamics calculations. In future applications, this approach could facilitate large-scale screening such as the search for best-performing devices or studies on structure–function relationships in natural light-harvesting. MLPs have been shown to generally perform well in supervised regression problems in chemistry.^{48,49} Further, we choose MLPs since there is no informative relation between neighboring elements in the Frenkel exciton Hamiltonian which could be exploited by convolutional or recurrent neural networks, as excitonic sites can be numbered in arbitrary order.

Overall, our procedure can be summarized as follows. Based on the Frenkel exciton Hamiltonian we leverage standard open quantum system approaches to generate a database comprising of average transfer times and efficiencies for EET from a donor to a target pigment for a random set of Frenkel exciton Hamiltonians. The complete dataset is split into a training set, on which we train each MLP model, as well as a validation and a test set. For training data selection we will compare two strategies: (i) random selection of data points and (ii) selection of training data based on a principal component analysis (PCA) which allows us to extract those data points covering the most information sampled in the dataset. As we show in Section 2, the latter strategy is of particular relevance if the space of transfer properties is not evenly sampled and many representatives in the training set exhibit redundant information. We run a Bayesian optimization procedure to identify the best architecture for our MLP models. The performance of each architecture is quantified by the average relative absolute error made when predicting transfer properties for the validation set. Finally, we run predictions on the test set to assess the ability of the optimized architecture to generalize to realizations that were neither employed for training nor for validation during the Bayesian optimization. The source code for exciton transfer property predictions along with all trained MLP models as well as the datasets generated in this study are made available on GitHub.⁵⁰

2.1 Generating the excitation energy transfer database

To demonstrate the capabilities of our machine learning approaches, we investigate four datasets of randomly generated excitonic systems that are sampled around pigment–protein

complexes found in natural light-harvesting. For future reference, the generated database can be downloaded from a GitHub repository.⁵⁰

For our first dataset, we sample Hamiltonians around the FMO complex (Fig. 1), which serves frequently as the prototype light-harvesting complex. We construct three additional datasets that are motivated by the photosystem II of higher plants. For one set, we consider the eight pigments of the reaction center (RC) core, in which the primary step of charge separation is initiated through the electronically excited pigment Chl_{PD1}.^{30,51} For the other two sets, the reaction center core is extended by including either light-harvesting complex CP47 or CP43 of photosystem II into the exciton system. For simplicity, we refer to the dataset inspired by the CP43 + RC (CP47 + RC) complex as the CP43 (CP47) dataset from hereon. For each dataset, we generated 12 000 exciton Hamiltonians by uniformly sampling excited state energies and inter-site couplings from a fixed range of values, as is summarized in Table 1.

In the following, we are interested in transfer characteristics such as average transfer times from an initially excited pigment (donor) to a target pigment (acceptor). This model provides a simple description of the first step of photosynthesis, where energy is absorbed in the antenna pigments and subsequently transferred to the reaction center in which photochemical reactions are triggered. The energy transport in light-harvesting complexes is determined by coupled pigments which are embedded in a protein scaffold,^{52,53} and is typically modeled with an effective Frenkel exciton Hamiltonian. We include energy trapping in the acceptor pigment phenomenologically by introducing anti-Hermitian parts in the Hamiltonian. The exciton dynamics is expressed in terms of the reduced density matrix, which can be obtained from standard open quantum system approaches.

We compute exciton transfer times for all Hamiltonians in our datasets with the hierarchical equations of motion (HEOM)^{37–39} method, implemented in the *QMaster* software package, version 0.2.^{33,54,55} HEOM is a numerically exact method which accurately accounts for the reorganization process,^{56–59} in which the vibrational coordinates rearrange to their new equilibrium positions upon electronic transition from the ground to the excited potential energy surface. For all Hamiltonians we assumed identical Drude–Lorentz spectral densities

Table 1 Lower and upper limits in between which excited state energies ϵ and inter-site couplings V were sampled uniformly to generate the four datasets of this study. Each dataset consists of 12 000 Hamiltonians with excited state energies and inter-site couplings within the reported ranges. Note, that the labels CP43 (CP47) denote datasets which are inspired by the CP43 + RC (CP47 + RC) biological complexes

Label	#sites	ϵ_{low} [cm ⁻¹]	ϵ_{high} [cm ⁻¹]	V_{range} [cm ⁻¹]
RC	8	14 800	15 000	–50 to 50
FMO	8	12 000	12 800	–100 to 100
CP43	21	14 800	15 100	–60 to 60
CP47	24	14 500	15 300	–100 to 100



$J(\omega) = 2\lambda \frac{\omega\nu}{\omega^2 + \nu^2}$, describing the exciton–phonon interaction.

We do not use the parameters of the spectral density as input features for our neural networks. Extending our approach to predict transfer properties for various spectral densities goes beyond the present scope and is the aim of future work. More details on the Frenkel exciton Hamiltonian and the exciton dynamics methods, as well as the definition of the transfer time and transfer efficiencies, are given in the ESI Section A.†

Distributions of transfer times for all exciton Hamiltonians of each dataset are depicted in Fig. 2. The transfer times for the Hamiltonians of the biological complexes are highlighted in every distribution. Excited states and inter-site couplings for the exciton Hamiltonians of the biological complexes are taken from literature,^{29–31,60} and are uploaded to the GitHub repository.⁵⁰ All population dynamics simulations are initialized as a fully populated site 1, serving as a donor, while site 3 acts as acceptor that couples to an energy sink with trapping rate Γ_{trap} (see ESI Section A†). Note that the labeling of the donor and acceptor state is without loss of generality as rows and columns of the Hamiltonian can be permuted in a suitable way, which effectively corresponds to a relabeling of the pigments. Since excited state energies and inter-site couplings are drawn from the same distributions for all sites in one dataset we did not explicitly account for the ordering ambiguity which arises, for instance, in the case of Coulomb matrices for which matrix entries depend on the particular types of atoms to which they correspond.⁴³

We find large variations in the ranges of transfer times between the four datasets. The RC and CP43 datasets, both with relatively narrow ranges of excited state energies and site couplings, yield relatively small transfer times. In contrast, we observe a wider spread in transfer times for the FMO dataset

and the CP47 dataset which is consistent with the broader range of excited state energies and site couplings that were sampled.

The transfer times of the actual biological complexes lie close to the mode of the distributions for all four datasets. This suggests that natural systems may not be specifically selected for extraordinary transfer properties, as they exhibit transport characteristics that are just likely to occur, even for a random choice of the exciton Hamiltonian. We note that providing a conclusive answer goes beyond the scope of the present manuscript, but could be the subject of a future more detailed structure–function analysis. A recent evolutionary study for the FMO complex⁶¹ goes along a similar direction and suggests that the FMO complex has evolved towards stability to mutations rather than a selection of specific transfer characteristics.

2.2 Principal component analysis for improved training data selection

We select the training sets for our MLP models following two methods for dataset splitting. In the simplest ansatz, we select the training set randomly from our created dataset. However, due to the nature of how we randomly sampled our Hamiltonians, the transfer characteristics are not distributed homogeneously and many representations of our Hamiltonians might be very similar and thus are expected to carry redundant information. As can be seen in Fig. 2, Hamiltonians yielding longer transfer time-scales are for example underrepresented in all four datasets.

Therefore, we follow a different path and carry out a more sophisticated selection process. The idea is to add those Hamiltonians to our training set which give the most information. We perform a principal component analysis (PCA) on the 8000 Hamiltonians containing dataset (after separating 2000 Hamiltonians each for validation and testing). We project each Hamiltonian onto a reduced space spanned by the most relevant principal components. The Hamiltonians for the training set are then selected such that they are maximally separated in the reduced space. This procedure guarantees that our training set constitutes the most diverse entities.

2.3 Setup of the multi-layer perceptron architecture

The architectures of our multi-layer perceptrons (MLPs) are designed for supervised learning of exciton energy transfer properties. All exciton Hamiltonians were reshaped into vectors and provided as input features to the MLPs, which were used to predict exciton transfer times and transfer efficiencies simultaneously. Since, the input features of neural networks need to be of fixed size, we construct separate MLPs for each dataset in order to treat the different dimensionalities of the exciton Hamiltonians. Details on the rescaling of the input features and predicted output, as well as on the training procedure are provided in the ESI (see Section E†).

The 12 000 Hamiltonians of each dataset were split into three sets: a training set of up to 6000 Hamiltonians for training MLP model instances with particular hyperparameters, a validation set of 2000 Hamiltonians used to evaluate the MLP architecture during optimization of the hyperparameters and




Fig. 2 Distributions of exciton transfer times computed for all 12 000 generated exciton Hamiltonians for each dataset using the HEOM approach implemented in *QMaster*. Vertical red lines indicate the transfer time of the exciton Hamiltonian corresponding to the biological complex. In all calculations we use a trapping rate of $\Gamma_{\text{trap}}^{-1} = 1$ ps, an exciton life-times of $\Gamma_{\text{loss}}^{-1} = 0.25$ ns, and a temperature of $T = 300$ K. The parameters of the spectral density are set to $\lambda = 35 \text{ cm}^{-1}$, $\nu^{-1} = 50$ fs.



a test set of 2000 Hamiltonians to probe out-of-sample prediction accuracies. All constructed MLP models were trained with stochastic gradient descent with 200 data points per batch and the ADAM optimizer,⁶² until the average relative absolute error (see eqn (1)) on the validation set increased over three full consecutive training epochs. Neuron saturation was avoided with L2 regularization on all weights of all neurons but the output neurons.

An essential component in developing accurate machine learning models consists in choosing proper values for the model hyperparameters. For this MLP framework, we consider a total of six hyperparameters. The initial learning rate μ for the ADAM optimizer and the regularization parameter λ . We also included the number of MLP layers and the number of neurons per layer, as well as the activation functions for neurons in each layer, for which we allowed five different options to choose from. The only exception is the last layer, for which we always use the softplus activation function to constrain our MLP models to the prediction of always positive transfer times and efficiencies. Lastly, we treat the number of training points as a hyperparameter in order to study the effect of the variations in the number of training samples on the prediction accuracy. The set of hyperparameters to be optimized and their allowed ranges are summarized in the ESI in Table III.†

We employ a Bayesian optimization algorithm,⁶³ in order to scan the space of hyperparameters for the most accurate model. The model accuracy was defined as the average relative absolute error (see eqn (1)) in exciton transfer times predicted by the MLP and corresponding HEOM simulations for the validation set. All generated MLP models were constructed and trained with the same random seed. Bayesian optimization is a common tool in machine learning and balances exploration of parameter space and exploitation of previous information. The idea of this ansatz is to reduce the number of costly function evaluations under the assumption that the unknown function was sampled from a Gaussian process. In contrast to gradient or Hessian based optimization techniques, Bayesian optimization uses information of all previously evaluated points and can thus find a good approximation to the minimum of non-convex functions in relatively few iterations. We carried out the Bayesian optimization of MLP hyperparameters in the spearmint software package.⁴¹ MLP models were generated and trained using the Tensorflow package, version 1.0.⁶⁴

3 Results: prediction of transfer times with neural networks

In the subsequent discussion, we demonstrate the capabilities of our trained MLP models by analyzing the average relative absolute error

$$\Delta\tau = \left\langle \frac{|t_{\text{HEOM}} - t_{\text{model}}|}{t_{\text{HEOM}}} \right\rangle_{\text{dataset}}, \quad (1)$$

between predicted exciton transfer times and the ones obtained with the numerically exact HEOM calculations. Although we restrict our discussion to transfer times, we note that similar

conclusions hold for the analysis of the transfer efficiencies since both characteristics are strongly correlated. Table 2 summarizes the results for the predicted transfer times for our four generated datasets.

The predictions are carried out with the Bayesian optimized MLP architectures, which show slight variations in their best-performing hyperparameters depending on the dataset at hand. However, for all datasets, the neural networks tend to prefer shallow but broad architectures comprising of only a few layers with each layer containing a larger number of neurons. More details on the procedure and results for the hyperparameter optimization can be found in the ESI Section F.†

3.1 Prediction accuracies of trained multi-layer perceptrons

Our trained MLP models predict exciton transfer times for out-of-sample Hamiltonians at almost the same accuracy as for Hamiltonians on which MLP parameters and hyperparameters were optimized (see Table 2). This demonstrates the ability of our MLP models to generalize to previously unseen data and to provide accurate out-of-sample predictions. Noteworthy, there is no significant asymmetry in the distribution of the relative absolute errors for the individual Hamiltonians or the training/validation and test set (see Fig. 3). Therefore, the architectures of the neural networks are well-balanced and neither in the regime of over-fitting, which would result in a large discrepancy in errors between the training and validation sets nor did we over-optimize the neural network architecture during Bayesian optimization.

Overall we find a high accuracy of our predictions and small average relative errors on the test sets which are in the range between 3.35% for RC (PCA selected training set) and 13.59% for the largest considered exciton system CP47 attached to RC (random selected training set). The CP47 dataset exhibits the most diverse transfer properties (see Fig. 2), which explains the larger average relative absolute errors in the predictions when compared to the other datasets. Prediction accuracies for

Table 2 Average relative absolute error $\Delta\tau$ (see eqn (1)) of exciton transfer times computed with HEOM and either, predicted by the trained neural networks (with/without PCA selection) or computed with secular Redfield. For all four datasets, we show the results of the training, validation, and test set separately. Smallest errors for each dataset are printed in bold

Dataset	Model	$\Delta\tau_{\text{train}}$ [%]	$\Delta\tau_{\text{valid}}$ [%]	$\Delta\tau_{\text{test}}$ [%]
FMO	Network (PCA)	4.53	4.38	7.41
	Network	10.53	10.75	11.56
	Redfield	9.70	9.96	9.60
RC	Network (PCA)	2.71	2.73	3.35
	Network	3.61	3.58	3.76
	Redfield	8.62	8.67	8.60
CP43	Network (PCA)	4.42	4.47	4.72
	Network	4.66	4.71	4.86
	Redfield	4.71	4.66	4.73
CP47	Network (PCA)	12.36	12.32	12.59
	Network	13.36	13.34	13.59
	Redfield	10.48	10.47	10.51






Fig. 3 Normalized distributions of the average relative absolute error of predicted exciton transfer times and exciton transfer times computed with HEOM. The left (blue) side of the plots illustrate the distributions of average relative absolute errors for predictions on the training and the validation set, while the right (orange) side of the plots illustrates the errors for predictions on the test set.

exciton Hamiltonians with permuted rows and columns are reported in the ESI Section G.† We find prediction accuracies similar to those achieved on the test sets for Hamiltonians with permutations not involving the source or target sites. The observed prediction errors are also consistent with the distance distributions of Frenkel exciton Hamiltonians for each of the four datasets (see ESI Section C†), which indicates that MLP models generally benefit from a finer sampling of the input parameter space.

The accuracy of the predictions can be enhanced by a more sophisticated PCA selection of the training set without the need of generating additional computationally expensive data points. The level of improvement of the PCA selection over a random selection of the training set differs for the four complexes. In general, we find that MLPs can be trained almost equally accurate with either selection method. The highest benefit of the PCA selected training set is obtained for the FMO and CP47 dataset, which are not only the most diverse ones out of our four datasets but are biased towards Hamiltonians showing fast transfer. As intuitively expected, selecting training points based on PCA is most advantageous for datasets with an extremely unevenly sampled feature space.

3.2 Comparing multi-layer perceptron predictions to secular Redfield results

Next, we provide a context for the observed MLP prediction accuracies by comparing them to the errors made by the frequently employed secular Redfield method, which is essentially derived from second order perturbation theory in the system–bath interaction in combination with a Markov approximation. Accuracies of the transfer times for both, the secular Redfield calculations and the MLP predictions are evaluated according to eqn (1). Here, the HEOM calculations again serve as ground truth. For the datasets inspired by the smaller exciton systems FMO and RC, the trained MLPs outperform secular Redfield, even for out-of-sample

predictions, whereas for the datasets around larger systems both approaches are similarly accurate.

For example in the case of the biological exciton Hamiltonian of the FMO complex, HEOM reveals a transfer time of 7.95 ps. The trained MLP model predicts a transfer time of 7.52 ps which is slightly more accurate than secular Redfield calculations that result in 7.48 ps. Exciton transfer times obtained for all four biological complexes with all three approaches are reported in the ESI Table I.† However, while the MLP prediction takes about 5 ms, secular Redfield calculations took about 14.5 min on a single CPU (computation times are listed in the ESI in Table II†). We conclude that our trained MLP predictions are competitive to secular Redfield calculations in terms of their accuracy, but (once trained) come at a significantly reduced computational cost. Computational costs for all three approaches are summarized in the ESI Section D.†

Besides analyzing the accuracy in terms of averaging over all realizations in the datasets, we compare the relative errors in transfer time for secular Redfield and the MLP predictions in more detail on the level of individual Hamiltonians. Fig. 4 depicts scatter plots where the horizontal axes measure the accuracy of secular Redfield calculations and the vertical axes reflect the accuracy of MLP predictions for MLPs trained on the PCA selected datasets. We do not distinguish between training, validation, and test set and show the complete dataset. Almost all the Hamiltonians show a $\Delta t_{\text{Redfield}} = (t_{\text{HEOM}} - t_{\text{Redfield}})/t_{\text{HEOM}} > 0$, which demonstrates that secular Redfield systematically underestimates transfer time scales. On the other hand, the




Fig. 4 Relative errors in exciton transfer times computed with the hierarchical equations of motion (HEOM) approach and exciton transfer times computed with the secular Redfield approach and predicted by neural networks respectively. Displayed are relative deviations for all four datasets: the Fenna–Matthews–Olson (FMO) complex, the reaction center (RC) core, the RC with the CP43 complex and the RC with the CP47 complex. Regions in which the absolute of deviations of neural network predicted transfer times from HEOM computed transfer times are smaller than deviations for Redfield are shaded in green.



predictions under-as well as overestimate transfer time-scales yielding a more symmetrical distribution along the horizontal axis. For the RC (FMO) dataset, more than 95% (80%) of the Hamiltonians fall into regions marked as green, for which the neural networks provide higher accuracy than secular Redfield. For all other datasets, secular Redfield and the MLP predictions are equally likely to give better results, with about 59% (57%) of the Hamiltonians for CP43 (CP47) falling within the green shaded region. This is in agreement with our average relative absolute errors listed in Table 2. We did not observe any cases for which the MLPs show relative errors that significantly exceeded any of the secular Redfield ones.

4 Conclusion

In this study, we have outlined how machine learning approaches can be employed to bypass computationally costly simulations of open quantum system dynamics in the context of excitation energy transfer. Overall we find that MLPs are capable of predicting transfer times for excitonic systems at higher or comparable accuracy than the frequently used secular Redfield approach albeit at much lower computational costs. Therefore we conclude that MLP models are a promising alternative for extracting excitation energy transfer properties when compared to frequently used rate equation methods.

The presented approach is of particular interest for large-scale analyses of the structure–transport relationship in excitonic systems. An area of great interest in excitonics is the study of the dynamics of charge dissociation at the interface present in bulk heterojunction photovoltaics.^{65,66} We believe a tool like this will help in the rapid screening of material properties in the mesoscale and therefore help the search for high-performance OPV systems.⁶⁷

Once trained, evaluations of MLP models come at almost no additional cost. Our four generated MLP architectures (each optimized for one of the four datasets) predict transfer times for an aggregated set of 48 000 exciton Hamiltonians just within a few seconds, while the corresponding quantum dynamics simulations take several GPU (CPU) years for the HEOM (secular Redfield) calculations. Our trained MLP models extend well to out-of-sample predictions for exciton Hamiltonians that are close to the sampled parameter regime. However, to employ MLPs on parameter regimes beyond those probed in the existing database requires running computationally expensive exciton dynamics for a few thousand Hamiltonians in order to extend our training set. To avoid this bottleneck a potential strategy could be to leverage already existing data, *e.g.* produced by a user community of existing software packages such as *QMaster*. However such data can be quite diverse. To this end, future research needs to focus on novel more general neural network architectures that accurately predict transfer times for flexible spectral density parameters as well as for differently sized exciton systems.

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

F. H. is supported by the Herchel Smith Graduate Fellowship. C. K. is supported by the National Science Foundation under award number CHE-1464862. A. A.-G. acknowledges support from the Center for Excitonics and Energy Frontier Research Center funded by the U.S. Department of Energy under award DE-SC0001088. All computations reported in this paper were completed on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University.

References

- 1 S. Caffarri, R. Kouřil, S. Kereiche, E. J. Boekema and R. Croce, *EMBO J.*, 2009, **28**, 3052–3063.
- 2 N. R. Baker, *Annu. Rev. Plant Biol.*, 2008, **59**, 89–113.
- 3 C. Kreisbeck and A. Aspuru-Guzik, *Chem. Sci.*, 2016, **7**, 4174–4183.
- 4 K. Amarnath, D. I. G. Bennett, A. R. Schneider and G. R. Fleming, *Proc. Natl. Acad. Sci. U. S. A.*, 2016, **113**, 1156–1161.
- 5 G. D. Scholes, G. R. Fleming, L. X. Chen, A. Aspuru-Guzik, A. Buchleitner, D. F. Coker, G. S. Engel, R. van Grondelle, A. Ishizaki, D. M. Jonas, J. S. Lundeen, J. K. McCusker, S. Mukamel, J. P. Ogilvie, A. Olaya-Castro, M. A. Ratner, F. C. Spano, K. B. Whaley and X. Zhu, *Nature*, 2017, **543**, 647–656.
- 6 G. D. Scholes, G. R. Fleming, A. Olaya-Castro and R. van Grondelle, *Nat. Chem.*, 2011, **3**, 763–774.
- 7 S. M. Blau, D. I. G. Bennett, C. Kreisbeck, G. D. Scholes and A. Aspuru-Guzik, arXiv:1704.05449, 2017.
- 8 C. Kreisbeck and T. Kramer, *J. Phys. Chem. Lett.*, 2012, **3**, 2828–2833.
- 9 A. W. Chin, J. Prior, R. Rosenbach, F. Caycedo-Soler, S. Huelga and M. Plenio, *Nat. Phys.*, 2013, **9**, 113–118.
- 10 N. Christensson, H. F. Kauffmann, T. Pullerits and T. Mančal, *J. Phys. Chem. B*, 2012, **116**, 7449–7454.
- 11 J. C. Dean, T. Mirkovic, Z. S. D. Toa, D. G. Oblinsky and G. D. Scholes, *Chem*, 2016, **1**, 858–872.
- 12 E. Romero, R. Augulis, V. I. Novoderezhkin, M. Ferretti, J. Thieme, D. Zigmantas and R. van Grondelle, *Nat. Phys.*, 2014, **10**, 676–682.
- 13 A. De Sio, F. Troiani, M. Maiuri, J. Réhault, E. Sommer, J. Lim, S. F. Huelga, M. B. Plenio, C. A. Rozzi, G. Cerullo, E. Molinari and C. Lienau, *Nat. Commun.*, 2016, **7**, 13742.
- 14 E. Collini, C. Y. Wong, K. E. Wilk, P. M. G. Curmi, P. Brumer and G. D. Scholes, *Nature*, 2010, **463**, 644–647.
- 15 G. S. Engel, T. R. Calhoun, E. L. Read, T.-K. Ahn, T. Mancal, Y.-C. Cheng, R. E. Blankenship and G. R. Fleming, *Nature*, 2007, **446**, 782–786.
- 16 T. Brixner, J. Stenger, H. M. Vaswani, M. Cho, R. E. Blankenship and G. R. Fleming, *Nature*, 2005, **434**, 625–628.
- 17 J. Schulze and O. Kühn, *J. Phys. Chem. B*, 2015, **119**, 6211–6216.
- 18 B. Hein, C. Kreisbeck, T. Kramer and M. Rodríguez, *New J. Phys.*, 2012, **14**, 023018.



- 19 D. Suess, A. Eisfeld and W. T. Strunz, *Phys. Rev. Lett.*, 2014, **113**, 150403.
- 20 P. Nalbach, D. Braun and M. Thorwart, *Phys. Rev. E: Stat., Nonlinear, Soft Matter Phys.*, 2011, **84**, 041926.
- 21 T. Scholak, T. Wellens and A. Buchleitner, *J. Phys. B: At., Mol. Opt. Phys.*, 2011, **44**, 184012.
- 22 S. Mostarda, F. Levi, D. Prada-Gravia, F. Mintert and F. Rao, *Nat. Commun.*, 2013, **4**, 2296.
- 23 S. Baghbanzadeh and I. Kassal, *Phys. Chem. Chem. Phys.*, 2016, **18**, 7459–7467.
- 24 S. Baghbanzadeh and I. Kassal, *J. Phys. Chem. Lett.*, 2016, **7**, 3804–3811.
- 25 G. C. Knee, P. Rowe, L. D. Smith, A. Troisi and A. Datta, *J. Phys. Chem. Lett.*, 2017, **8**, 2328–2333.
- 26 J. A. Leegwater, *J. Phys. Chem.*, 1996, **100**, 14403–14409.
- 27 V. May and O. Kühn, *Charge and energy transfer dynamics in molecular systems*, John Wiley & Sons, 2008.
- 28 R. E. Fenna and B. W. Matthews, *Nature*, 1975, **258**, 573–577.
- 29 F. Müh, M. Madjet and T. Renger, *Photosynth. Res.*, 2012, **111**, 87–101.
- 30 G. Raszewski and T. Renger, *J. Am. Chem. Soc.*, 2008, **130**, 4431–4446.
- 31 G. Raszewski, B. A. Diner, E. Schlodder and T. Renger, *Biophys. J.*, 2008, **95**, 105–119.
- 32 M. Mohseni, P. Rebentrost, S. Lloyd and A. Aspuru-Guzik, *J. Chem. Phys.*, 2008, **129**, 174106.
- 33 C. Kreisbeck, T. Kramer, M. Rodriguez and B. Hein, *J. Chem. Theory Comput.*, 2011, **7**, 2166–2174.
- 34 P. Rebentrost, M. Mohseni, I. Kassal, S. Lloyd and A. Aspuru-Guzik, *New J. Phys.*, 2009, **11**, 033003.
- 35 F. Caruso, A. W. Chin, A. Datta, S. F. Huelga and M. B. Plenio, *J. Chem. Phys.*, 2009, **131**, 105106.
- 36 F. Fassioli and A. Olaya-Castro, *New J. Phys.*, 2010, **12**, 085006.
- 37 Y. Tanimura and R. Kubo, *J. Phys. Soc. Jpn.*, 1989, **58**, 101–114.
- 38 A. Ishizaki and G. R. Fleming, *J. Chem. Phys.*, 2009, **130**, 234111.
- 39 Y. Tanimura, *J. Chem. Phys.*, 2012, **137**, 22A550.
- 40 C. Kreisbeck and T. Kramer, *Exciton Dynamics Lab for Light-Harvesting Complexes (GPU-HEOM)*, 2013, <http://nanohub.org/resources/16106>.
- 41 J. Snoek, H. Larochelle and R. P. Adams, *NIPS*, 2012, pp. 2951–2959.
- 42 Z. W. Ulissi, A. J. Medford, R. Bligaard and J. K. Nørskov, *Nat. Commun.*, 2017, **8**, 14621.
- 43 K. Hansen, G. Montavon, F. Biegler, S. Fazli, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko and K. R. Müller, *J. Chem. Theory Comput.*, 2013, **9**, 3404–3419.
- 44 J. Behler and M. Parrinello, *Phys. Rev. Lett.*, 2007, **98**, 146401.
- 45 J. S. Smith, O. Isayev and A. E. Roitberg, *Chem. Sci.*, 2017, **8**, 3192–3203.
- 46 K. Yao, J. E. Herr and J. Parkhill, *J. Chem. Phys.*, 2017, **146**, 014106.
- 47 J. Gomes, B. Ramsundar, E. N. Feinberg and V. Pande, arXiv preprint arXiv:1703.10603, 2017.
- 48 F. Häse, S. Valleau, E. Pyzer-Knapp and A. Aspuru-Guzik, *Chem. Sci.*, 2016, **7**, 5139–5147.
- 49 G. Montavon, M. Rupp, V. Gobre, A. Vazquez-Mayagoitia, K. Hansen, A. Tkatchenko, K. R. Müller and O. A. von Lilienfeld, *New J. Phys.*, 2013, **15**, 095003.
- 50 F. Häse, C. Kreisbeck and A. Aspuru-Guzik, *Deep Learning of Excitation Energy Transfer Properties at Redfield Accuracy*, 2017, <https://github.com/FlorianHase/LearningExcitonTransfer>.
- 51 A. R. Holzwarth, M. G. Müller, M. Reus, M. Nowaczyk, J. Sander and M. Rögner, *Proc. Natl. Acad. Sci. U. S. A.*, 2006, **103**, 6895–6900.
- 52 V. May and O. Kühn, *Charge and Energy Transfer Dynamics in Molecular Systems*, Wiley-VCH, Weinheim, 2004.
- 53 Y. C. Cheng and G. R. Fleming, *Annu. Rev. Phys. Chem.*, 2009, **60**, 241–262.
- 54 C. Kreisbeck and T. Kramer, *J. Phys. Chem. Lett.*, 2012, **3**, 2828–2833.
- 55 C. Kreisbeck, T. Kramer and A. Aspuru-Guzik, *J. Chem. Theory Comput.*, 2014, **10**, 4045–4054.
- 56 Y. Yan, F. Yang, Y. Liu and J. Shao, *Chem. Phys. Lett.*, 2004, **395**, 216–221.
- 57 R. Xu, P. Cui, C. Li, Y. Mo and Y. Yan, *J. Chem. Phys.*, 2005, **122**, 041103.
- 58 A. Ishizaki and Y. Tanimura, *J. Phys. Soc. Jpn.*, 2005, **74**, 3131–3134.
- 59 A. Ishizaki and G. R. Fleming, *J. Chem. Phys.*, 2009, **130**, 234111.
- 60 J. Adolphs and T. Renger, *Biophys. J.*, 2006, **91**, 2778–2797.
- 61 S. Valleau, R. Struder, F. Häse, C. Kreisbeck, R. G. Saer, R. E. Blankenship, E. I. Shakhnovich and A. Aspuru-Guzik, *ACS Cent. Sci.*, 2017, **3**, 1086–1095.
- 62 D. Kingma and J. Ba, arXiv preprint arXiv:1412.6980, 2014.
- 63 L. C. W. Dixon and G. P. Szegő, *Towards global optimisation*, North-Holland Amsterdam, 1978.
- 64 M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu and X. Zheng, *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*, 2015, <http://tensorflow.org/>.
- 65 A. E. Jailaubekov, A. P. Willard, J. R. Tritsch, W. L. Chan, N. Sai, R. Gearba, L. G. Kaake, K. J. Williams, K. Leung, P. J. Rossky and X. Y. Zhu, *Nat. Mater.*, 2013, **12**, 66.
- 66 D. A. Vithanage, A. Devizis, V. Abramavicius, Y. Infahsaeng, D. Abramavicius, R. C. I. MacKenzie, P. E. Keivanidis, A. Yartsev, D. Hertel, J. Nelson, V. Sundstrom and V. Sundstrom, *Nat. Commun.*, 2013, **4**, 2234.
- 67 J. Hachmann, R. Olivares-Amaya, S. Atahan-Evrenk, C. Amador-Bedolla, R. S. Sanchez-Carrera, A. Gold-Parker, L. Vogt, A. M. Brockway and A. Aspuru-Guzik, *J. Phys. Chem. Lett.*, 2011, **2**, 2241–2251.

