

Cite this: *Chem. Sci.*, 2019, 10, 2298

All publication charges for this article have been paid for by the Royal Society of Chemistry

# How machine learning can assist the interpretation of *ab initio* molecular dynamics simulations and conceptual understanding of chemistry†

Florian Häse, <sup>a</sup> Ignacio Fdez. Galván, <sup>b</sup> Alán Aspuru-Guzik, <sup>cde</sup> Roland Lindh <sup>b</sup> and Morgane Vacher <sup>\*b</sup>

Molecular dynamics simulations are often key to the understanding of the mechanism, rate and yield of chemical reactions. One current challenge is the in-depth analysis of the large amount of data produced by the simulations, in order to produce valuable insight and general trends. In the present study, we propose to employ recent machine learning analysis tools to extract relevant information from simulation data without *a priori* knowledge on chemical reactions. This is demonstrated by training machine learning models to predict directly a specific outcome quantity of *ab initio* molecular dynamics simulations – the timescale of the decomposition of 1,2-dioxetane. The machine learning models accurately reproduce the dissociation time of the compound. Keeping the aim of gaining physical insight, it is demonstrated that, in order to make accurate predictions, the models evidence empirical rules that are, today, part of the common chemical knowledge. This opens the way for conceptual breakthroughs in chemistry where machine analysis would provide a source of inspiration to humans.

Received 10th October 2018  
Accepted 21st December 2018

DOI: 10.1039/c8sc04516j

rsc.li/chemical-science

## 1 Introduction

Computer simulations are a key complement to experiments in the laboratory, especially when the latter are expensive or challenging because of extreme conditions required. Simulations also provide much greater details of a molecular process than can be observed experimentally. For instance, studying the time evolution of matter with molecular dynamics simulations is essential for understanding the mechanism, rate and yield of

chemical reactions. Such simulations are also necessary to complement experiments and connect with time-resolved pump-probe measurements. At each time step of an *ab initio* molecular dynamics simulation, the energies and forces felt by the nuclei are calculated “on-the-fly” with an electronic structure method. With the growing complexity of the investigated chemical problems and the increasing need for improved accuracy, molecular dynamics simulations become very time-consuming. Typical time and length scales that are accessible with *ab initio* molecular dynamics are, with current computer systems, up to hundreds of femtoseconds (fs) to tens of picoseconds, and tens to few hundreds of atoms. As simulations become more complex, their usefulness for guidance and understanding may become obscured. Simple lessons are often lost among gigabytes or terabytes of data. The present work proposes to use machine learning methods to aid the interpretation of molecular dynamics simulations. The bigger goal here is in the future to allow machines to provide a source of inspiration to humans for the elaboration of new concepts in chemistry. This has been identified by some of us as one of the six grand challenges for the simulation of matter in the 21st century<sup>1</sup> and the present work demonstrates an approach for achieving this.

As a test application, the timescale of the chemiluminescent decomposition of 1,2-dioxetane is chosen (Fig. 1A). Chemiluminescence is the emission of light as a result of a chemical reaction. This process is called bioluminescence when occurring in living organisms as in the well-known example of the firefly. For a recent review on the topic, see ref. 2 and 3. Chemiluminescence

<sup>a</sup>Department of Chemistry and Chemical Biology, Harvard University, Cambridge, Massachusetts 02138, USA

<sup>b</sup>Department of Chemistry – Ångström, The Theoretical Chemistry Programme, Uppsala University, Box 538, 751 21 Uppsala, Sweden. E-mail: morgane.vacher@kemi.uu.se

<sup>c</sup>Department of Chemistry and Department of Computer Science, University of Toronto, Toronto, Ontario M5S 3H6, Canada

<sup>d</sup>Vector Institute for Artificial Intelligence, Toronto, Ontario, M5S 1M1, Canada

<sup>e</sup>Canadian Institute for Advanced Research (CIFAR), Senior Fellow, Toronto, Ontario, M5S 1M1, Canada

† Electronic supplementary information (ESI) available: Brief introduction to PCA. Further details on the implementation of the BNN models and on the hyperparameters optimisation including the scanned hyperparameter and obtained performances of different BNN architectures. Analysis of the BNN performance with the training set size. Predictions of dissociation half-times for vibrational states excited along two normal modes. Dynamics simulations of the dissociation of 1,2-dioxetane for different vibrational ground and excited states. Transition state structure and normal modes of the unmethylated 1,2-dioxetane. The two trained BNN can be found here: <https://github.com/FlorianHase/LearningMoleculeDissociations>. See DOI: 10.1039/c8sc04516j




Fig. 1 1,2-Dioxetane molecule. (A) Dark and chemiluminescent dissociation reactions. (B) Schematic representation of the relevant normal modes. Normal mode 1 corresponds mainly to the OCCO dihedral angle while normal mode 7 corresponds mainly to the HCCH dihedral angles.

and bioluminescence are increasingly used in biological and chemical analysis methods,<sup>4</sup> in various fields such as DNA sequencing,<sup>5</sup> immunoassays as an alternative to radioactive isotopes<sup>6</sup> and as a sensitive probe for mechanical stimulations.<sup>7,8</sup> The applications are of course limited by the light emission efficiency, *i.e.* the chemiluminescence yield. 1,2-dioxetane is known to be the smallest compound with chemiluminescent properties. Upon thermal activation, it decomposes into two formaldehyde molecules in a two-step process: first the O-O bond breaks (Fig. 1A, step 1) and then the C-C bond (Fig. 1A, step 2).<sup>9,10</sup> Non-adiabatic transitions from the electronic ground state to electronic excited states during the decomposition reaction lead to chemiexcitation. The resulting electronic excited states can relax back to the electronic ground state by emitting light. Previous *ab initio* molecular dynamics simulations have been performed for several methyl-substituted dioxetane molecules.<sup>11,12</sup> Those simulations have shown that the chemiexcitation yield is determined by the dissociation timescale: the longer the molecule stays in the so-called “entropic trap” region before dissociating, the higher the chemiexcitation yield. Being able to calculate and understand the dissociation time of dioxetanes is thus essential for rationalising chemiluminescence yields measured experimentally and for designing new and efficient chemiluminescent compounds. It is, however, *a priori* unclear which molecular modifications would eventually result in significant changes of the dissociation time. Testing possible hypotheses with *ab initio* methods is an arduous process due to the computational demand required by the methods and can easily exceed available computing resources.

In the present work we illustrate how machine learning models can be used to solve this task. More specifically, we employ Bayesian neural networks (BNN) which are trained to predict the dissociation times of 1,2-dioxetane from initial

nuclear positions (or from initial nuclear positions and velocities). We demonstrate that the trained BNN can be used in two different ways: (i) to autonomously find physical correlations, and (ii) to test already formulated hypotheses. This is performed on both the unmethylated and tetramethylated dioxetanes, the latter presenting a longer dissociation timescale and thus a higher chemiexcitation yield than the former. High prediction accuracy for the dissociation time is achieved despite the limited size of the dataset, which consists of 250 trajectories simulating up to 250 fs with a time step of 0.24 fs. This presents an advancement over prior studies with datasets with orders of magnitude more data points.<sup>13,14</sup> More importantly, we provide an example for how to interpret the machine learning model itself and the changes it has undergone during the training procedure in order to learn how to reproduce the targeted outcome. In addition, the trained BNN is used to predict dissociation times of vibrationally excited states of the unsubstituted 1,2-dioxetane. Detailed understanding of the effect of specific nuclear distortions, that could be either enhanced or diminished *via* chemical modification, opens up possibilities for more efficient molecular design. Still aiming at gaining physical insight, it is demonstrated that while being trained to reproduce the dissociation times of 1,2-dioxetanes with high accuracy, the BNN has evidenced some chemical rules that connect the nuclear positions (and velocities) and the dissociation times.

It is noted that to reduce the computational cost of molecular dynamics simulations, a lot of efforts have recently been devoted to construct efficient models that “learn” or fit the potential energy surfaces many orders of magnitude faster than electronic structure calculations.<sup>13,15–22</sup> For example, machine learning techniques have been used to reproduce energies at the level of hybrid density-functional theory, from lower-level calculations and so-called molecular descriptors<sup>23,24</sup> or directly from just the Cartesian coordinates and the nuclear charges.<sup>25–27</sup> The present work goes to a higher level of abstraction and proposes to use a machine learning model to directly predict a specific outcome of the molecular dynamics simulation, bypassing the construction of potential energy surfaces and avoiding the computation of the time propagation. The focus here is however neither on reducing the computational cost of electronic structure in *ab initio* molecular dynamics simulations, nor on describing with quantitative accuracy chemiluminescent or bioluminescent reactions. It is rather on training a machine learning model on already simulated trajectories (of a given chemical reaction and at a given level of theory), and interpreting and using the trained machine learning model in order to gain physical insight into the studied chemical reaction. We describe two possible approaches to gain insights: studying the trained models themselves to find correlations, and using the models to test already formulated hypotheses. We note further that these approaches do not depend on the way the potential energy surfaces are generated along the dynamics trajectories. Our strategy could be used as well to interpret dynamics simulations based on machine-learned potential energy surfaces. In fact, the use of machine learning to reduce the cost of *ab initio* molecular dynamics



simulations is expected to increase the amount of data generated and to be interpreted, which will further increase the need for the tools proposed in the present work.

## 2 Methods

### 2.1 *Ab initio* molecular dynamics simulations

The approach used for the *ab initio* molecular dynamics simulations is the same as in previously published works.<sup>11,12</sup> The complete active space self-consistent field (CASSCF) method<sup>28,29</sup> state-averaging over the four lowest-energy singlet states equally is used to describe the electronic structure of the system. The active space chosen consists of 12 electrons and 10 orbitals: the four  $\sigma$  and four  $\sigma^*$  orbitals of the four-membered ring, plus the two oxygen lone-pair orbitals perpendicular to the ring. The ANO-RCC basis set with polarized triple-zeta contraction is used.<sup>30,31</sup> Born–Oppenheimer dynamics is simulated with a time step of 10 au ( $\approx 0.24$  fs) and all nuclear coordinates are taken into account. It is noted that, in the present simulations, only the electronic ground state is included and non-adiabatic transitions to electronic excited states are not allowed. The implementation of this method in the OpenMolcas package is used.<sup>32,33</sup> The trajectories are initialized and propagated from the transition state for the O–O bond breaking – since the latter controls the overall reaction rate. A small amount of kinetic energy (1 kcal mol<sup>−1</sup>) is given along the reaction coordinate (as suggested in previous theoretical studies of post-transition state dynamics<sup>34</sup>) toward the biradical region where the O–O bond is broken. The Newton-X package<sup>35</sup> was used to sample 250 initial positions and velocities along all normal modes (other than the reaction coordinate) from a Wigner distribution in order to reproduce the vibrational ground state. The normal modes were calculated at the transition state structure using the electronic structure method mentioned above. Dissociation is considered to occur when the C–C bond length exceeds 2.4 Å (two times the van der Waals radius of a carbon atom). It is noted that choosing a slightly smaller or larger value for the bond length dissociation threshold has been shown not to change any relative comparisons, nor the findings regarding the entropic trap and the effect of the singlet excited states.<sup>11</sup>


### 2.2 Machine learning predictions

Two probabilistic models were used in order to estimate dissociation times of the dioxetane molecule: these are implemented as feedforward fully connected BNN. Neural networks are constructed as a set of nodes, called ‘neurons’, with connections between them (Fig. 2). Traditionally, each neuron is characterised by a weight  $w$ , a bias  $b$  and an activation function  $f_{\text{act}}$ . For a given input  $x$ , a single neuron performs the operation

$$y = f_{\text{act}}(w \cdot x + b), \quad (1)$$

where  $y$  denotes the output of the neuron. The architecture of a neural network is defined by a set of parameters, collectively referred to as ‘hyperparameters’, which define for instance the number of neuron layers, the number of neurons per layer and

#### A) Single Bayesian neuron



#### B) Bayesian neural network




Fig. 2 Illustration of a Bayesian Neural Network (BNN). (A): A Bayesian neuron defines a mathematical operation based on an activation function  $f_{\text{act}}$ , a distribution of weights  $w$  and a distribution of biases  $b$  intrinsic to the neuron. Every input  $x$  is processed by sampling one instance of weights and biases from the distributions and applying the activation function. (B) A BNN consists of a set of interconnected Bayesian neurons. The neurons in the network are organised in layers, and can differ in their activation functions as well as their weight and bias distributions.

the activation function of the neurons. In the case of a BNN, both weights  $w$  and biases  $b$  are modelled as random variables, sampled from a probability distribution, with conditional dependencies on either other model parameters or the input parameters, called ‘input features’ (Fig. 2A). As a consequence, a BNN produces an output distribution (Fig. 2B). The BNN can be trained so that the output distribution resembles a desired target distribution by adapting the distributions of weights and biases for each individual neuron. As such, BNN propagate information based on entire parameter distributions as opposed to traditional neural networks, which compute a target value based on single-value parameters. BNN therefore retain the flexibility of traditional neural networks, but provide a more robust framework for identifying relevant correlations between inputs and outputs, especially for small and medium sized datasets.<sup>36</sup>

The two models constructed in the present work use different sets of input features. The first model, hereafter noted ‘BNN1’, uses only the initial nuclear geometry of the molecule.



The second model, hereafter noted 'BNN2', also uses the initial nuclear velocities in addition to the geometry. The nuclear geometry and velocities are given in normal mode coordinates to account for translational and rotational invariances. The normal modes used are the ones calculated at the transition state structure as indicated above. For the unmethylated dioxetane molecule consisting of 8 atoms, BNN1 has therefore 18 input features while BNN2 has 36 input features. For the tetramethylated dioxetane molecule consisting of 20 atoms, BNN1 and BNN2 have thus 54 and 108 input features, respectively.

Datasets for the BNN models were extracted from the *ab initio* molecular dynamics simulations by sampling every fifth frame (up to dissociation) along each of the 250 trajectories, resulting in a total of 11 650 frames and the corresponding dissociation times. We split the total set of data into a training set (80%, 9320 frames), a validation set (10%, 1165 frames) and a test set (10%, 1165 frames). Data for the test set were selected as random samples from the entire set and were not used for any part of the training procedure other than for reporting the out-of-sample prediction accuracy after all BNN models have been trained. An informative and diverse training set was assembled from the remaining 90% of the dataset *via* principal component analysis (PCA).<sup>37</sup> (The reader is referred to ESI† for a brief introduction to PCA.) We selected the frames for the training set which are maximally separated in the reduced PCA space spanned by the most contributing principal components. This protocol has shown to improve prediction accuracies in the context of excitation transfer property predictions.<sup>18</sup> The frames, which were not selected with this protocol, were used as the validation set.

The BNN used for the present study were implemented in the probabilistic programming library Edward,<sup>38</sup> and model parameters were updated *via* variational inference using the Adam optimisation algorithm.<sup>39</sup> We parametrised the distributions of weights and biases as Laplace distributions. This choice is made in order to construct interpretable models. While the training set is used to optimise the model parameters  $\mathbf{w}$  and  $\mathbf{b}$  of a BNN model of a given architecture, the validation set serves as a benchmark to determine the best performing BNN architecture. For both BNN1 and BNN2, we conducted extensive hyperparameter searches to determine the best performing BNN models. The most accurate BNN model was selected as the model with the lowest prediction error on the validation set. We computed an upper bound on the overfitting error for both most accurate constructed models, BNN1 and BNN2, based on the Rademacher complexity<sup>40</sup> for k-layered networks:<sup>41</sup> it is 2.20 fs for BNN1 and 2.38 fs for BNN2. Details on the implementation of the BNN models and on the hyperparameter optimisation including the scanned hyperparameter and obtained performances of different BNN architectures, are reported in the ESI†. It is noted here that hyperparameters of BNN1 and BNN2 were found to be similar, despite the different number of input features. It is important to remember that the BNN are trained to learn the dissociation time of 1,2-dioxetane or tetramethyl-1,2-dioxetane from a particular set of initial conditions (nuclear positions, and velocities for BNN2) without *a priori* knowledge about the dynamics of the chemical reaction of interest.

## 3 Results

In this section, the machine learning predictions are used to study the mechanism of the dissociative reaction of 1,2-dioxetane. The BNN trained for the unmethylated dioxetane molecule is analysed in detail and unless explicitly stated otherwise, the results are presented for this compound. We begin our analysis with validating that the trained BNN are indeed capable of predicting dissociation times within reasonable accuracy. Then, we proceed with analysing the architecture of the trained BNN to identify nuclear coordinates relevant to the dissociation timescale. Finally, we use the BNN to test hypotheses about physical correlations.

### 3.1 Validation of the dissociation time predictions

First we present, for the test set, the comparison of the C–C dissociation times obtained from the *ab initio* molecular dynamics simulations with the dissociation times predicted by the best performing BNN1 and BNN2 models in Fig. 3. Standard deviations of the predicted dissociation time distributions were used to indicate the uncertainty of the BNN model with respect to the predicted dissociation time. We find that BNN can predict dissociation times with high accuracy, despite the medium-sized dataset used for training (10 000 data points from only 250 trajectories). A detailed analysis of the sampling efficiency, *i.e.* the achieved performance with respect to different training set sizes is reported in the ESI†. The prediction accuracy is improved when supplementing the molecular geometry with the velocities of the atoms (Fig. 3B): the mean absolute deviation (MAD) is 2.40 fs for BNN2 while it is 6.55 fs for BNN1. This observation, combined with the observation that prediction accuracies on the validation set are similar for BNN1 or BNN2 (see ESI†) indicates that supplementing geometries with velocities improves the generalisation abilities of the BNN models. In addition, the BNN models generally predict dissociation times with higher uncertainties when the deviation between the predicted and the true dissociation time is larger. We note that the prediction accuracy can likely be improved when supplementing the presented model training procedure with strategies like cross-fold validation for an effective increase of the training set size as proposed in other studies,<sup>42,43</sup> or testing different models such as kernel-based methods.<sup>24,44,45</sup> However, in this study we aim to focus on the interpretability of the trained models and  $r^2 = 0.97$  for BNN2 is considered to be sufficient for this purpose.

Once trained, machine learning models can be extensively used to predict the targeted values with high accuracy and low computational cost. As a reference, it takes on one processor less than a minute to generate a set of 250 initial conditions,<sup>35</sup> a few seconds to predict the corresponding 250 dissociation times with the trained BNN and approximately 31.5 hours to simulate 160 fs of a single *ab initio* molecular dynamics simulation (for the unmethylated dioxetane, at the level of theory used). Here, in addition, we consider the trained BNN as a result by itself and analyse it to see if physical insight can be gained. In particular, we are interested in understanding the following points: How does the BNN reproduce the C–C dissociation





Fig. 3 Out-of-sample predictions for best performing (A) BNN1 and (B) BNN2 models. For all frames in the test set (1165 frames) we depict the dissociation times predicted by the BNN models in comparison to the true dissociation times. Root mean square deviations (RMSD), mean absolute deviations (MAD) and coefficients of determination ( $r^2$ ) are reported. Depicted points are coloured based on the predicted uncertainty of the BNN model. The grey dashed line indicates perfect agreement between predicted and true dissociation times.

timescales with such a good agreement? What correlations has the BNN identified to achieve high accuracies?

### 3.2 Analysis of the trained models to find correlations

The Laplace distribution (used here as priors for the weights and biases) has been shown to facilitate efficient pruning in BNN parameters,<sup>46</sup> and is the equivalent to L1-regularization in the Bayesian context.<sup>47</sup> As such, this choice is well suited to enable the BNN to identify the input features which are most relevant for accurate predictions of the target properties. Generally, the larger the magnitude of the coefficient, the more influential the feature. This relation can be used to determine input features which are relevant for accurate predictions of the dissociation time and therefore relevant to the physical process.

Fig. 4 shows the coefficient magnitude distributions of the input features for the Bayesian neural networks BNN1 and BNN2. The numbering of the normal modes goes from 0 to 17, 0 being the reaction coordinate;  $\{r_{ij}\}$  correspond to geometry coordinates along these normal modes while  $\{v_{ij}\}$  correspond to velocities. As just explained, an input feature that has a large coefficient magnitude does not imply that a distortion along this nuclear coordinate would lead to a longer dissociation time but rather, it indicates that this nuclear coordinate is relevant to predict accurately the dissociation time. It is noted that the choice of input features affects thus not only the learning ability of the BNN, but also the physical interpretation of the correlations that are identified. Trying other input features than the nuclear geometry and velocities in normal mode coordinates is beyond the scope of the present work, but it would be an interesting task to pursue.

Out of the 18 normal modes, BNN1 identifies four as more important for the decomposition reaction: normal modes 6, 7, 8 and 11. They are represented schematically in Fig. 1B. Normal mode 6 corresponds to the (symmetric) stretching of the two C–O bonds, associated with the stretching of the central C–C bond in an out-of-phase manner. Normal mode 7 corresponds to the HCCH dihedral angles. Normal modes 8 and 11 correspond to the stretching of the central C–C bond in phase with symmetric out-of-plane motions of the two OCH<sub>2</sub> moieties. In BNN2, the coefficients of the initial nuclear geometry (input features  $r_0$  to  $r_{17}$ ) are larger in magnitude, in general, than the ones of the initial velocities (input features  $v_0$  to  $v_{17}$ ). As in BNN1, the input features  $r_6$ ,  $r_7$ ,  $r_8$  and  $r_{11}$  are identified to be important. In addition,  $r_1$  is also identified to be important. Normal mode 1 corresponds to the OCCO dihedral angle (Fig. 1B). Among the remaining 18 input features corresponding to the initial velocities, the features  $v_1$  and  $v_8$  have the largest coefficient magnitudes. Normal modes 1 and 8 are already identified as important according to the coefficients of the initial position features.

### 3.3 Use of the trained models to test hypothesis

It is noted that the previous analysis, applied on the trained BNN able to make accurate predictions, allows the identification of the most relevant nuclear coordinates for the dissociation dynamics. However, it is not known how these affect the dynamics, *i.e.* whether they make the dissociation occur earlier or later. As an alternative way to gain physical insight, the two trained BNN have been used to predict the dissociation time-scales for 17 ensembles of 250 initial conditions, each ensemble representing a vibrational state that is excited to the first level along one particular normal mode. For example, the ensemble “3” corresponds to a vibrational state that is excited along normal mode 3, while it remains in the ground state along all other modes. The ensemble “0” is a reference vibrational ground state along all normal modes, different from the 250 ground state trajectories used for training the BNN. For each ensemble of 250 initial conditions, the corresponding 250 predicted dissociation times are sorted to then extract the overall dissociation half-time, time at which half of the trajectories





Fig. 4 Coefficient magnitude distributions of the input features of the trained (A) BNN1 and (B) BNN2. Input features  $\{r_i\}$  and  $\{v_i\}$  correspond to geometry coordinates and velocities along normal modes, respectively. The numbering of the normal modes goes from 0 to 17, 0 being the reaction coordinate.

have dissociated. The predicted dissociation half-times are presented in Fig. 5A: they vary from 49 fs to 63 fs, the reference being 60 fs. It is observed that accelerating the dissociation upon vibrational excitation is more probable than slowing it down, which makes sense since some kinetic energy is added. Excitations of the normal modes 4, 6, 8 and 11 lead to an acceleration of the dissociation while excitation along normal mode 3 is predicted to slow down the dissociation. Normal modes 6, 8 and 11 were already identified as important according to their high coefficient magnitude. Normal mode 3 corresponds to the antisymmetric stretching of the two C–O bonds; normal mode 4 corresponds to the central C–C bond stretching (Fig. 1B). So, according to the trained BNN, in order for the C–C bond to break earlier rather than later, the two formaldehyde moieties need to become planar at the same time as both C–O bonds need to shorten (*i.e.* excitation along normal modes 6, 8 and 11), while if the two C–O bonds stretch in an anti-symmetric fashion (*i.e.* excitation along normal mode 3), this delays dissociation. On the other hand, the high frequency modes (12 to 17) involving the hydrogen atoms seem to affect only mildly the dissociation timescale.

It is interesting to note that, although normal mode 7 is identified as important for predicting accurately the dissociation time (because of its large coefficient magnitude, Fig. 4), vibrational excitation along only this mode seems to barely affect the overall dissociation half-time compared to the reference one (Fig. 5A). To understand this fact, we have used the trained BNN2 to predict the dissociation times for ensembles of

initial conditions, where each ensemble now represents a vibrational state that is excited to the first level along two particular normal modes in order to detect second-order effects. The predicted dissociation half-times obtained for each pair of normal modes are given in ESI,<sup>†</sup> as well as the normal mode z-scores providing information about how excitation along a nuclear coordinate influences the dissociation half-time obtained with an excitation along another coordinate. We find that combining any vibrationally excited normal mode with an excitation along normal mode 7 slows down dissociation. We thus conclude that a sole excitation along normal mode 7 is not sufficient to influence the dissociation timescale (maybe because this mode mainly involves the motion of hydrogen atoms). However, excitation of normal mode 7 in combination with the excitation of another nuclear coordinate could impact the dissociation time significantly. This result would explain the large coefficient magnitude of normal mode 7 in Fig. 4 and also shows how the two proposed methods to gain physical insights from the trained BNN are complementary.

### 3.4 *Ab initio* molecular dynamics simulations of vibrational excited states as “numerical experiments”

To check the accuracy of the dissociation times predicted by the BNN and presented in Fig. 5A, we have simulated the *ab initio* molecular dynamics of four ensembles of 250 trajectories (in addition to the reference ensemble “0”) with excitations along normal modes 3, 4, 8 and 11. It is noted that the *ab initio* molecular dynamics simulations are computationally





Fig. 5 Dissociation half-times for 17 ensembles of 250 initial conditions representing different vibrationally excited states. Ensemble "0" is the reference vibrational ground state. Otherwise, ensemble " $n$ " corresponds to a vibrational excitation along normal mode  $n$ , and ground state along other normal modes. (A) Predicted dissociation half-times by the two trained BNN. The solid and dashed horizontal lines indicate the dissociation time for the ensemble "0", predicted and extracted from the simulations, respectively. (B) Comparison of the dissociation half-times predicted by the BNN and obtained from *ab initio* molecular dynamics simulations, for five ensembles of 250 trajectories. The error bars represent the 95% confidence intervals of the predictions.

expensive; they are run here for test purposes, in particular for checking the accuracy of the BNN predictions for initial vibrationally excited states not present in the training set.

The four normal modes 3, 4, 8 and 11 were chosen because excitations along modes 8 and 11 are predicted to decrease the most the dissociation half-time, while excitation along mode 3 is predicted to increase it. Normal mode 4 is also predicted to decrease the dissociation half-time and can be seen as the "naive" choice for enhancing C–C bond dissociation without any chemical knowledge, since it corresponds to additional energy along the C–C bond stretching solely. The comparison between the dissociation half-times predicted by the BNN and obtained from the *ab initio* molecular dynamics simulations is presented in Fig. 5B. The qualitative trend is correctly predicted by the BNN. Overall, the errors are on the same order of magnitude as for the reference ensemble "0" although the excited normal modes exhibit larger initial nuclear distortions that may not be present in the training set. In general, BNN2

Table 1 Average number of frustrated dissociations per trajectory for each ensemble of 250 trajectories

| Ensemble          | 0    | 3    | 4    | 8    | 11   |
|-------------------|------|------|------|------|------|
| $n_{\text{frus}}$ | 1.68 | 1.52 | 1.18 | 1.08 | 1.17 |

performs better than BNN1: the root-mean-square deviation is about 1.6 times smaller for the former than for the latter. BNN1 seems to overestimate the dissociation time. An interesting point is that the naive thought that exciting the C–C bond stretching would enhance dissociation is partially validated. It does indeed decrease the dissociation half-time compared to the reference by almost 5 fs, *i.e.* about 8%. However, the BNN successfully find other nuclear modes which decrease the dissociation time further (about 13%) thanks to its finding of physical correlations.

In a previous theoretical study,<sup>11</sup> the longer dissociation times among the 1,2-dioxetane trajectories were explained by the presence of so-called "frustrated" dissociations – significant stretching of the central C–C bond followed by a shortening rather than a "successful" breaking of the bond. Frustrated dissociations would postpone in time the final successful dissociation. To understand further the link between the predicted/simulated dissociation times and the number of frustrated dissociations, the average number of frustrated dissociations per trajectory are calculated for each ensemble of 250 simulated trajectories (Table 1). As expected, the vibrationally excited ensembles present fewer frustrated dissociations than the reference ground state since more kinetic energy is available to overcome energy barriers. Interestingly, for the four vibrational excited ensembles, the fewer frustrated dissociations, the shorter the dissociation time. This is particularly true for ensemble 8 that exhibits the largest decrease in the average number of frustrated dissociations (35%) and the largest decrease in the dissociation time (13%). The BNN have thus identified nuclear coordinates that would affect the overall dissociation time through the number of these frustrated dissociations. It is noted, however, that ensemble "3" presents a lower average number of frustrated dissociations than the reference ensemble "0" despite its longer dissociation half-time. This indicates that the number of frustrated dissociations is not the only parameter determining the dissociation timescale, dissociations needing to be attempted in order to be frustrated or successful.

## 4 Discussion

In this section, the trained models are discussed with first interpreting the identified correlations in the context of chemistry. Finally, implications for chemiexcitation and chemical design are discussed, using the case of the tetramethylated dioxetane molecule.

### 4.1 Interpretation of the findings of the trained models

Interpreting the correlations identified by the BNN leads to the following fundamental concepts known in chemistry: the octet



rule,<sup>48–50</sup> the relation between bond order and bond length<sup>51,52</sup> and orbital hybridisation.<sup>53,54</sup> The octet rule states that atoms (with an atomic number  $Z \geq 5$ ) tend to combine in such a way that each atom has a full valence shell (*i.e.* with eight electrons). This electronic configuration, identical to the one of noble gases, is associated to a maximal stability. As a consequence, the carbon atom that has only four electrons in its valence shell will form four covalent bonds to acquire the four missing electrons. As the C–C bond breaks, each carbon atom has to form a new bond with one of the remaining neighbouring atoms in order to keep eight electrons in the valence shell: it achieves this by forming a double bond with the oxygen atom. The number of bonding electrons between a pair of atoms and thus the bond order determines then the length of a bond. When the single bond between the C and O atoms becomes a double bond, more electrons participate in forming the bond and the bond becomes shorter (*i.e.* excitation along normal mode 6). The molecular shape is also explained with orbital hybridisation, *i.e.* the fact that 2s and 2p atomic orbitals mix, by minimising the repulsion between the pairs of electrons. The theory used to rationalise the molecular shape is known as valence shell electron pair repulsion (VSEPR) model.<sup>55,56</sup> When each carbon atom is surrounded by four atoms (here before dissociation, an oxygen atom, the other carbon atom and two hydrogen atoms), the 2s and three 2p orbitals mix to form the four covalent bonds which adopt a tetrahedral arrangement: this is called  $sp^3$  hybridisation. It is noted that the tetrahedral conformation of the carbon atoms is constrained in 1,2-dioxetane by the four-membered ring structure. When the C–C bond breaks and each carbon atom becomes surrounded by only three atoms, then the 2s orbital mixes only with two 2p orbitals to form three  $sp^2$  hybridised orbitals; the remaining 2p orbital stays intact and will form the  $\pi$  bond by parallel overlap. Minimisation of the repulsion energy is obtained by a trigonal planar geometry. Upon C–C bond breaking, the two formaldehyde moieties become thus planar (*i.e.* excitation along normal modes 8 and 11). In summary, the BNN has identified the following correlation between nuclear coordinates and dissociation time, without any knowledge of electronic structure: the C–C bond breaking must be associated with the formation of a shorter (since double) C–O bond and the planarity of the formaldehyde moieties. This is chemistry!

## 4.2 Implications for chemiexcitation and chemical design

Because the dissociation timescale determines the chemiexcitation yield,<sup>12</sup> a machine learning model able to find nuclear coordinates that affect efficiently the dissociation dynamics could be used in design of chemiluminescent systems. Here, normal modes 3 and 8 are found to be the nuclear coordinates that induce the largest deviations from the reference ensemble 0 (Fig. 5B). The chemiexcitation yield is actually low in 1,2-dioxetane (0.3%).<sup>57</sup> A chemical substitution that would enhance the simultaneous planarisation of the two formaldehyde moieties together with the stretching of the central C–C bond would make the dissociation occur earlier and therefore decrease further the chemiexcitation yield. More

importantly, a chemical substitution that would induce an asymmetric stretching of the two C–O bonds would make the dissociation occur later and therefore increase the chemiexcitation yield. According to the kinetic model developed in a previous work,<sup>12</sup> a postponing of the dissociation time of (unmethylated) 1,2-dioxetane by 3.2 fs – as predicted by ensemble 3 compared to the reference ensemble 0 – enhances the chemiexcitation yield by more than a factor of six!

It was observed experimentally that the yield increases by two orders of magnitude upon substitution of the four hydrogen atoms by four methyl groups.<sup>57</sup> Initially, it was suggested that the greater number of nuclear coordinates in the methyl-substituted compounds would increase the depth of the “entropic trap” and thus the chemiexcitation yield.<sup>9</sup> Recently, the *ab initio* molecular dynamics of the methyl-substituted dioxetane has been simulated<sup>12</sup> although these simulations are more than 20 times more computationally expensive than the ones for the unmethylated compound. This later work demonstrated instead that the increase in chemiexcitation yield is due to a mass effect. When training a BNN based on an ensemble of trajectories of the tetramethylated dioxetane, we found that the nuclear coordinates presenting large coefficient magnitudes are the stretching of the C–O bonds, the stretching of the central C–C bond, a planarisation of the two ketone moieties and the O–C–C–O dihedral angle. These coordinates are the equivalent of the normal modes 1, 6, 8 and 11 in the unsubstituted 1,2-dioxetane molecule, which were identified as important for predicting the dissociation time of the unsubstituted dioxetane. This result confirms the similarity in the dissociation dynamics of the unsubstituted and methyl-substituted molecules and therefore that the importance of the methyl groups for increasing the chemiexcitation yield is not the additional nuclear coordinates, but the heavier mass. It also shows how machine learning models help interpreting the results of molecular dynamics simulations.

## 5 Conclusion

Bayesian neural networks have been optimised to predict a specific outcome of an *ab initio* molecular dynamics simulation: the dissociation time of the unmethylated and tetramethylated 1,2-dioxetane molecules from just the initial nuclear geometry (and velocities). This means replacing the *ab initio* molecular dynamics simulation in total and not only the calculation of the potential energy at every time step. Despite the medium size of the dataset used for training, a high prediction accuracy is obtained. More important than the ability to predict a number, we have analysed the trained Bayesian neural networks and demonstrated that machine learning models can help extracting conceptual information from the large amount of data produced by the simulations. Indeed, in order to make accurate predictions, the Bayesian neural networks have figured out that an earlier dissociation must be associated with the planarisation of the two formaldehyde moieties and the symmetric shortening of the C–O bonds. This is in connection with the octet rule, the relation between bond order and bond length and orbital hybridisation,



rules that are part of today's common knowledge of chemists. The present work is a step towards achieving one of the grand challenges in the 21st century<sup>1</sup> and opens thus the way for breakthroughs in chemistry where humans, inspired by the findings of machines, would develop new concepts.

## Conflicts of interest

There are no conflicts to declare.

## Acknowledgements

This work was supported by the Swedish Research Council (Grant 2016-03398). F. H. is supported by the Herchel Smith Graduate Fellowship. A. A.-G. thanks Dr Anders G. Frøseth for his generous support. Machine learning models were trained on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University. The *ab initio* molecular dynamics simulations and machine learning predictions were performed using resources provided by the Swedish National Infrastructure for Computing (SNIC) at UPPMAX and NSC centers.

## Notes and references

- 1 A. Aspuru-Guzik, R. Lindh and M. Reiher, *ACS Cent. Sci.*, 2018, **4**, 144–152.
- 2 I. Navizet, Y.-J. Liu, N. Ferré, D. Roca-Sanjuán and R. Lindh, *ChemPhysChem*, 2011, **12**, 3064–3076.
- 3 M. Vacher, I. Fdez. Galván, B.-W. Ding, S. Schramm, R. Berraud-Pache, P. Naumov, N. Ferré, Y.-J. Liu, I. Navizet, D. Roca-Sanjuán, W. J. Baader and R. Lindh, *Chem. Rev.*, 2018, **118**, 6927–6974.
- 4 C. Dodeigne, *Trends Biotechnol.*, 2000, **51**, 415–439.
- 5 M. Ronaghi, M. Uhlen and P. Nyren, *Science*, 1998, **281**, 363–365.
- 6 A. Mayer and S. Neuenhofer, *Angew. Chem., Int. Ed. Engl.*, 1994, **33**, 1044–1072.
- 7 Y. Chen, A. J. H. Spiering, S. Karthikeyan, G. W. M. Peters, E. W. Meijer and R. P. Sijbesma, *Nat. Chem.*, 2012, **4**, 559–562.
- 8 J. M. Clough, A. Balan, T. L. J. van Daal and R. P. Sijbesma, *Angew. Chem., Int. Ed.*, 2016, **55**, 1445–1449.
- 9 L. De Vico, Y.-J. Liu, J. W. Krogh and R. Lindh, *J. Phys. Chem. A*, 2007, **111**, 8013–8019.
- 10 P. Farahani, D. Roca-Sanjuán, F. Zapata and R. Lindh, *J. Chem. Theory Comput.*, 2013, **9**, 5404–5411.
- 11 M. Vacher, A. Brakestad, H. O. Karlsson, I. Fdez. Galván and R. Lindh, *J. Chem. Theory Comput.*, 2017, **13**, 2448–2457.
- 12 M. Vacher, P. Farahani, A. Valentini, L. M. Frutos, H. O. Karlsson, I. Fdez. Galván and R. Lindh, *J. Phys. Chem. Lett.*, 2017, **8**, 3790–3794.
- 13 J. S. Smith, O. Isayev and A. E. Roitberg, *Chem. Sci.*, 2017, **8**, 3192–3203.
- 14 K. Yao, J. E. Herr, S. N. Brown and J. Parkhill, *J. Phys. Chem. Lett.*, 2017, **8**, 2689–2694.
- 15 V. Botu and R. Ramprasad, *Int. J. Quantum Chem.*, 2015, **115**, 1074–1083.
- 16 Z. Li, J. R. Kermode and A. De Vita, *Phys. Rev. Lett.*, 2015, **114**, 096405.
- 17 F. Häse, S. Valleau, E. Pyzer-Knapp and A. Aspuru-Guzik, *Chem. Sci.*, 2016, **7**, 5139–5147.
- 18 F. Häse, C. Kreisbeck and A. Aspuru-Guzik, *Chem. Sci.*, 2017, **8**, 8419–8426.
- 19 F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke and K.-R. Müller, *Nat. Commun.*, 2017, **8**, 872.
- 20 E. Schneider, L. Dai, R. Q. Topper, C. Drechsel-Grau and M. E. Tuckerman, *Phys. Rev. Lett.*, 2017, **119**, 150601.
- 21 P. O. Dral, M. Barbatti and W. Thiel, *J. Phys. Chem. Lett.*, 2018, **9**, 5660–5663.
- 22 W.-K. Chen, X.-Y. Liu, W.-H. Fang, P. O. Dral and G. Cui, *J. Phys. Chem. Lett.*, 2018, **9**, 6702–6708.
- 23 R. M. Balabin and E. I. Lomakina, *J. Chem. Phys.*, 2009, **131**, 074104.
- 24 R. Ramakrishnan, P. O. Dral, M. Rupp and O. A. von Lilienfeld, *J. Chem. Theory Comput.*, 2015, **11**, 2087–2096.
- 25 J. Behler and M. Parrinello, *Phys. Rev. Lett.*, 2007, **98**, 146401.
- 26 M. Rupp, A. Tkatchenko, K.-R. Müller and O. A. von Lilienfeld, *Phys. Rev. Lett.*, 2012, **108**, 058301.
- 27 M. Gastegger, J. Behler and P. Marquetand, *Chem. Sci.*, 2017, **8**, 6924–6935.
- 28 B. O. Roos, P. R. Taylor and P. E. Siegbahn, *Chem. Phys.*, 1980, **48**, 157–173.
- 29 B. O. Roos, *Adv. Chem. Phys.*, 1987, **69**, 399–446.
- 30 B. O. Roos, R. Lindh, P.-Å. Malmqvist, V. Veryazov and P.-O. Widmark, *J. Phys. Chem. A*, 2004, **108**, 2851–2858.
- 31 F. Aquilante, L. Gagliardi, T. B. Pedersen and R. Lindh, *J. Chem. Phys.*, 2009, **130**, 154107.
- 32 F. Aquilante, J. Autschbach, R. K. Carlson, L. F. Chibotaru, M. G. Delcey, L. De Vico, I. Fdez. Galván, N. Ferré, L. M. Frutos, L. Gagliardi, M. Garavelli, A. Giussani, C. E. Hoyer, G. Li Manni, H. Lischka, D. Ma, P. Å. Malmqvist, T. Müller, A. Nenov, M. Olivucci, T. B. Pedersen, D. Peng, F. Plasser, B. Pritchard, M. Reiher, I. Rivalta, I. Schapiro, J. Segarra-Martí, M. Stenrup, D. G. Truhlar, L. Ungur, A. Valentini, S. Vancollie, V. Veryazov, V. P. Vysotskiy, O. Weingart, F. Zapata and R. Lindh, *J. Comput. Chem.*, 2016, **37**, 506–541.
- 33 *OpenMolcas*, <https://gitlab.com/Molcas/OpenMolcas>, 2017.
- 34 U. Lourderaj, K. Park and W. L. Hase, *Int. Rev. Phys. Chem.*, 2008, **27**, 361–403.
- 35 M. Barbatti, G. Granucci, H. Lischka, M. Persico and M. Ruckebauer, *Newton-X: a package for Newtonian dynamics close to the crossing seam*, version 0.11b, 2006, <http://www.univie.ac.at/newtonx>.
- 36 C. K. Williams, *Advances in neural information processing systems*, 1997, pp. 295–301.
- 37 K. Pearson, *London, Edinburgh Dublin Philos. Mag. J. Sci.*, 1901, **2**, 559–572.
- 38 D. Tran, A. Kucukelbir, A. B. Dieng, M. Rudolph, D. Liang and D. M. Blei, 2016, arXiv preprint arXiv:1610.09787.
- 39 D. P. Kingma and J. Ba, *CoRR*, 2014, arXiv:abs/1412.6980.



- 40 S. Shalev-Shwartz and S. Ben-David, *Understanding machine learning: from theory to algorithms*, Cambridge university press, 2014.
- 41 W. Gao and Z.-H. Zhou, *Sci. China Inf. Sci.*, 2016, **59**, 072104.
- 42 M. Rupp, *Int. J. Quantum Chem.*, 2015, **115**, 1058–1073.
- 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 M. Rupp, R. Ramakrishnan and O. A. von Lilienfeld, *J. Phys. Chem. Lett.*, 2015, **6**, 3309–3313.
- 45 K. Vu, J. C. Snyder, L. Li, M. Rupp, B. F. Chen, T. Khelif, K.-R. Müller and K. Burke, *Int. J. Quantum Chem.*, 2015, **115**, 1115–1128.
- 46 P. M. Williams, *Neural Comput.*, 1995, **7**, 117–143.
- 47 S. Mohamed, K. A. Heller and Z. Ghahramani, *CoRR*, 2011, arXiv:abs/1106.1157.
- 48 R. Abegg, *Z. Anorg. Chem.*, 1904, **39**, 330–380.
- 49 G. N. Lewis, *J. Am. Chem. Soc.*, 1916, **38**, 762–785.
- 50 I. Langmuir, *J. Am. Chem. Soc.*, 1919, **41**, 868–934.
- 51 L. Pauling, L. O. Brockway and J. Beach, *J. Am. Chem. Soc.*, 1935, **57**, 2705–2709.
- 52 C. A. Coulson, *Proc. R. Soc. London, Ser. A*, 1939, **169**, 413–428.
- 53 L. Pauling, *J. Am. Chem. Soc.*, 1931, **53**, 1367–1400.
- 54 L. Pauling, *The Nature of the Chemical Bond*, Cornell university press, Ithaca, NY, 1960, vol. 260.
- 55 R. J. Gillespie and R. S. Nyholm, *Q. Rev., Chem. Soc.*, 1957, **11**, 339–380.
- 56 R. J. Gillespie, *J. Chem. Educ.*, 1970, **47**, 18.
- 57 W. Adam and W. J. Baader, *J. Am. Chem. Soc.*, 1985, **107**, 410–416.

