

Cite this: *Digital Discovery*, 2025, 4, 3400

Screening Diels–Alder reaction space to identify candidate reactions for self-healing polymer applications

Maxime Ferrer,^{ID}* Bowen Deng, Javier E. Alfonso-Ramos^{ID} and Thijs Stuyver^{ID}*

Plastics are essential in modern society, but their susceptibility to damage limits their lifespan and performance, and results in unsustainable waste production. Self-healing polymers based on thermally reversible Diels–Alder (DA) reactions offer a potential solution by enabling heating controlled repair through bond-breaking and reformation. However, discovering new suitable DA monomer combinations has largely relied on intuition and trial-and-error so far. Here, we present a hierarchical workflow that integrates machine learning (ML) with automated reaction profile calculations to efficiently screen DA reactions for self-healing polymer applications. Using our in-house TS-tools software, we generate high-throughput profiles at the semi-empirical xTB level. Refining only a small fraction with DFT, we are able to train a robust ML model that predicts reaction characteristics with excellent accuracy. Adding a graph-based ML model to the workflow for pre-screening enables expansion to reaction spaces of hundreds of thousands of reactions, at a marginal cost. We first leverage our models to screen a comprehensive reaction space of synthetic diene–dienophile pairs, and subsequently use them to mine a database of commercially available natural products. Overall, this hybrid ML-computational chemistry approach enables data-efficient discovery of thermally responsive DA reactions, advancing the rational design of self-healing polymers with tunable properties.

Received 2nd August 2025
Accepted 17th October 2025

DOI: 10.1039/d5dd00340g

rsc.li/digitaldiscovery

Introduction

Plastics, and polymers more broadly, are essential in modern society, with applications spanning, among others, (bio)medicine,^{1,2} (micro)electronics,^{3,4} and industrial sectors such as the food or aerospace industries.^{5,6} Their widespread use is largely due to their low density, high durability, and cost-effectiveness. However, their susceptibility to mechanical damage oftentimes limits their lifespan and performance.⁷ Unlike biological materials, polymers lack intrinsic self-healing capabilities, meaning that fractures or microcracks can significantly compromise their durability. As a result, huge amounts of plastic waste are generated every year, posing significant environmental and sustainability challenges. In response, the past three decades have seen growing interest in self-healing/recyclable materials as a promising strategy to enhance the longevity and sustainability of polymer-based systems.^{8–10}

Self-healing materials can be classified into two main categories: autonomous and stimuli-responsive systems.⁷ As the name suggests, autonomous systems comprise materials that release a healing agent automatically upon detection of mechanical damage, *i.e.*, no external intervention is

required.^{11,12} In contrast, stimuli-responsive materials require an external factor to initiate the healing process.^{13,14} These systems typically rely on bond-breaking and bond-(re)forming mechanisms, which require well-characterized, reversible reactions that can be efficiently controlled through external stimuli, such as temperature variations. Given its practicality and broad applicability, thermal activation is particularly advantageous, making heat-induced reversible reactions a highly attractive approach for self-healing plastics.^{7,15}

Because of their inherently tunable thermodynamic and kinetic properties, Diels–Alder (DA) reactions in particular have been extensively investigated as a potential candidate for the development of thermally responsive self-healing polymers.^{16,17} The DA reaction is classified as a [4 + 2] cycloaddition reaction.^{18,19} It describes the concerted reaction between two unsaturated molecules, referred to as diene and dienophile (Fig. 1).

What makes the DA reaction particularly appealing for self-healing material applications is that the carbon–carbon bonds formed during the reaction are generally weaker than the other bonds within the system.^{21–23} These weaker and more flexible bonds are generally the first to be affected by mechanical deformation, *i.e.*, when stress is applied. However, for reversible DA reactions, these bonds can be re-formed, enabling the restoration of the polymer network.

Up to this point, the development of self-healing polymers has remained, by and large, a serendipitous process, whereby

Ecole Nationale Supérieure de Chimie de Paris, Université PSL, CNRS, Institute of Chemistry for Life and Health Sciences, 75005 Paris, France. E-mail: maxime.ferrer@chimieparistech.psl.eu; thijs.stuyver@chimieparistech.psl.eu



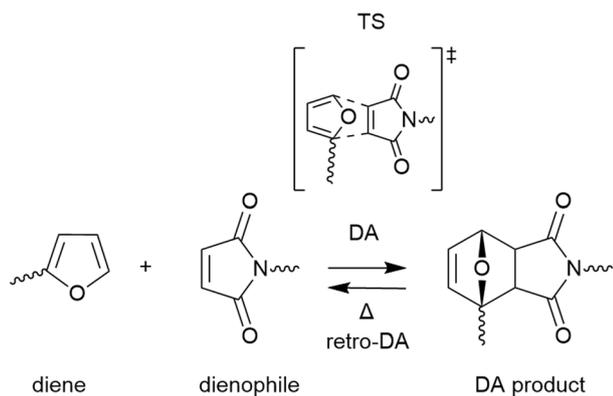


Fig. 1 Schematic representation of a prototypical self-healing polymer system based on a DA moiety.²⁰ The wavy bonds indicate the attachment points to the polymer backbone; TS stands for transition state.

new suitable DA monomer combinations are discovered primarily through intuition, in combination with a trial-and-error approach. To enable rational design, relationships between monomer chemical structure and macroscopic properties of the plastic need to be established. While the reaction characteristics of the DA monomers do not exclusively determine the properties of the resulting self-healing polymer (*e.g.*, the backbone, as well as the stoichiometric ratios, and their respective roles in determining the glass transition temperature, also play a major role),²¹ they clearly are essential design parameters that can be used to filter out unfavorable combinations, and identify promising ones for further investigation.

From a reactivity standpoint, a suitable DA reaction for thermally responsive self-healing applications must be (slightly) exothermic, or more precisely, must yield a product that is thermally unstable. At room temperature, the equilibrium should lie toward the product side, yet heating the system must allow the reaction to reverse—regenerating the precursors *via* the *retro*-DA reaction. This implies that the entropic contribution to the free energy must be sufficient to counteract the enthalpic drive toward product formation at elevated temperatures. Beyond thermodynamics, kinetics also play a crucial role: the forward and reverse activation barriers determine the rate and temperature at which free monomers can be (re)generated at the site of damage, as well as the conditions under which repolymerization—and thus self-healing—can occur.

It should be noted that the ability to fine-tune the healing temperature through reaction barrier modulation is highly desirable, as ideal operating temperatures are strongly dependent on the use case of the specific self-healing plastic. For daily-life applications, *e.g.*, soft robotics, low-temperature healing may be preferable from a safety perspective,²⁴ whereas plastics that may get exposed to more extreme conditions, *e.g.*, in aerospace applications, will ideally be more robust, *i.e.*, they should be able to retain their full integrity up to a higher temperature.²⁵

In principle, the reactive properties of a given diene–dienophile pair can be simulated with the help of quantum chemical methods, such as Density Functional Theory

(DFT).^{21,26,27} However, locating transition states (TS) on potential energy surfaces (PES) is a non-trivial task, and achieving accurate results that align with experimental observations necessitates the use of high-level computational methods, which incur a significant computational cost. As such, traditional DFT studies on DA reactivity have focused primarily on a limited number of model reactions, from which qualitative trends are extracted.²¹ Systematic exploration of DA reaction libraries, especially in the context of self-healing polymers, has consequently largely remained out of reach.

The advent of machine learning (ML) has made the screening of large chemical reaction spaces feasible in principle, and this has recently enabled the development of protocols for high-throughput reaction discovery.²⁸ In general terms, two main families of ML models for reactivity prediction can be distinguished: graph-based and geometry-based ones. The former are extremely fast, as they can generate predictions based on a simple reaction SMILES input, focusing only on connectivity, without the need for any conformational search.^{29–31} This means that predictions can be performed in milliseconds. While graph-based models can typically reach accuracies up to a couple of kcal mol^{−1}, they can, in principle, be outperformed by models that operate on geometries.^{32–34} Given enough training data, the latter can reach accuracies close to, or even below, chemical accuracy (± 1 kcal mol^{−1}) with respect to the reference method, though at a significant computational cost, since accurate geometries need to be determined as input first.³⁵

Effectively applying ML to explore chemical reaction spaces relies heavily on the availability of high-quality training data. Training data is typically generated with the help of DFT as well, and hence, the same issues as mentioned above are still encountered. Most reactivity screening studies in the literature so far indicate that several thousand reaction profiles are usually needed to build an accurate model, which corresponds to several million CPU hours of compute time.^{36–46}

Here, we propose a synergistic reaction screening approach, combining hierarchical ML-based reaction filtering with robust, and ultra-fast, reaction profile computation, facilitated by our in-house developed software tool for TS and reaction profile computation, TS-tools⁴⁷ (see Fig. 2). Focusing on a synthetically realistic reaction space of DA reactions, we demonstrate how the TS-tools package can be used to compute tens of thousands of reaction profiles at semi-empirical GFN2-*x*TB level of theory (hereafter referred to as *x*TB), on the fly, at a rate of over 100 reactions per hour on a single 72 CPU core workstation. Refining a small fraction of the so-generated reaction profiles at – a thoroughly benchmarked – DFT level subsequently enables the training of an (in-house designed) deep learning model, operating on the TS-tools generated *x*TB profiles (Fig. 4). The resulting model – DeepReaction – can reproduce the DFT quality reaction characteristics at near chemical accuracy.

Finally, we demonstrate how a rapid (yet less accurate and robust) secondary graph-based ML model can subsequently be used to expand the reaction search space beyond its original limits at marginal additional cost, by identifying promising patches of reaction space through reaction SMILES analysis, which can



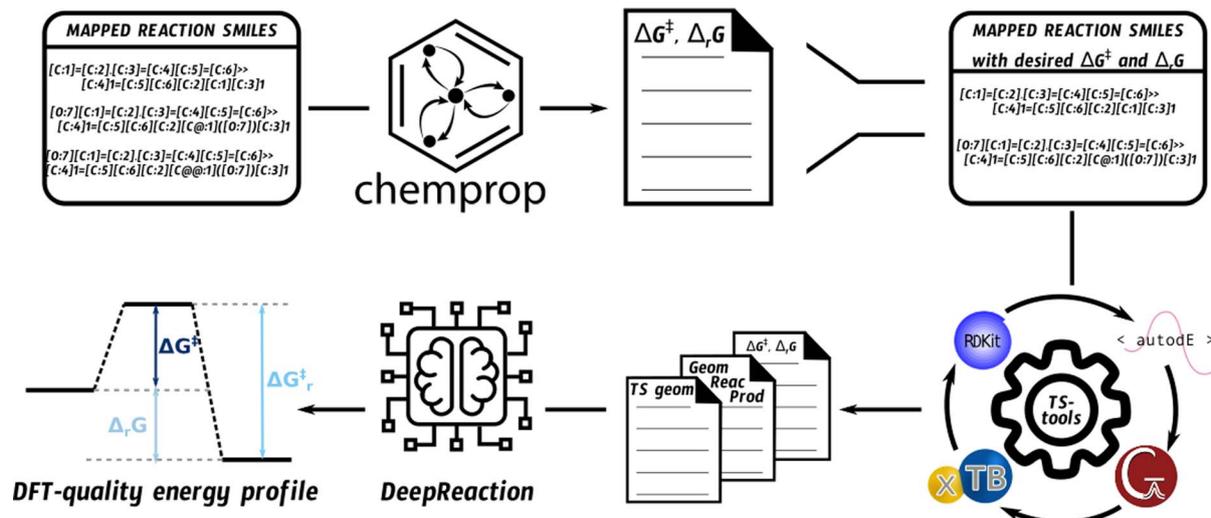


Fig. 2 A schematic representation of the proposed hierarchical workflow for data-efficient DA reaction screening for self-healing polymer applications.

subsequently be refined and validated with more expensive methods.

By combining different model architectures, leveraging their respective strengths, into a hierarchical model, we obtain an extremely data-efficient QM-augmented ML workflow²⁸ for DA reaction discovery, that can potentially be expanded towards the discovery of other types of reactions as well.

Methodology

Definition of the initial chemical reaction space

Starting from a recent review by Briou *et al.*,⁴⁸ an initial synthetic dataset (Fig. 3A and B) was constructed by combinatorially

matching previously experimentally considered diene and dienophile scaffolds with common substituents extracted from the SciFinder database using the scaffolds as substructure queries.⁴⁹ Only neutral, non-isotopically labeled, and commercially available compounds were considered, focusing exclusively on the most common and chemically reasonable substituents. This resulted in a total of 136 distinct dienes. For the dienophiles, 8 reactive scaffolds, derived from previous experiments, were considered (Fig. 3B). By symmetrically decorating these with common substituents, 53 unique dienophiles were obtained.

To underscore the realism of the diene and dienophile libraries generated in this manner, synthetic accessibility (SA)

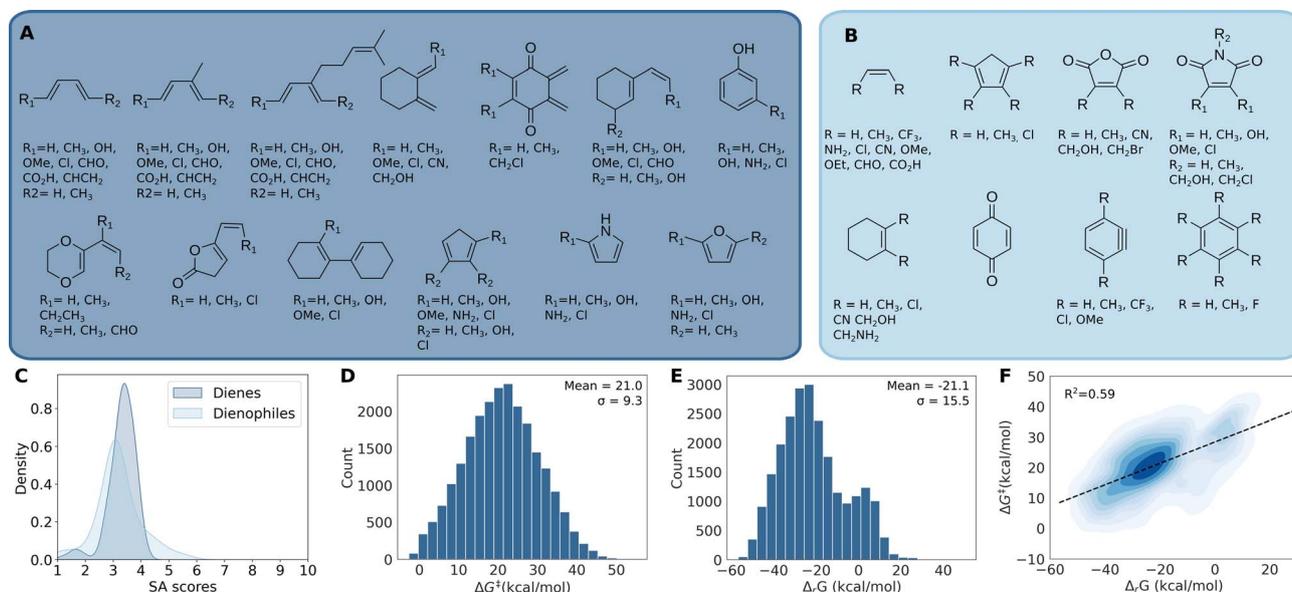


Fig. 3 Schematic overview of the chemical space explored. The dienes (A) are allowed to be asymmetric, while the dienophiles (B) are kept symmetric to limit the number of possible diastereoisomers. (C) SAScore computed for the considered dienes and dienophiles. (D) Distribution of the activation free energies of the reactions studied at xTB level of theory. (E) Distribution of the reaction free energies at xTB level. (F) Correlation plot between the activation free energy and the reaction free energy. All the energies reported are in kcal mol⁻¹.



scores were computed for all compounds, and as can be observed from (Fig. 3C), most fall in the range 1–4, and none exhibit scores above 6, indicating that they should be readily synthesizable.

By systematically combining components, we generated 7208 unique diene–dienophile pairs. Considering all potential regio- and stereoisomers, this led to a total of 28 560 unique reaction SMILES.^{37,50}

Dataset generation and benchmarking performance

To construct an initial xTB level-of-theory⁵¹ dataset of activation and reaction (Gibbs free) energies for the reactions present in the chemical reaction space outlined in Fig. 2, a fully automated workflow based on the TS-tools⁴⁷ package was developed. TS-tools is an in-house developed Python package facilitating the localization of transition states (TS) based on textual reaction SMILES inputs.⁴⁷ An in-depth discussion of the code, with particular attention to some custom additions introduced within the context of this study, primarily aiming to correctly capture the stereochemistry of DA reactions, can be found in Section S1 of the SI.

As indicated in the Introduction, at xTB level-of-theory, TS-tools can process reactions at an almost negligible computational cost: 23 327 reaction profiles were successfully computed (success rate of more than 81%) on a single 72 CPU core workstation in 9 days, totaling less than 15 000 CPU hours in overall compute time, *i.e.*, less than 1 h per reaction profile. Benchmarking the xTB results against both high-level computational, as well as experimental, datasets (consisting of 19 and 100 activation energy values, respectively) reveals – as expected for semi-empirical calculations – significant quantitative disagreement for the barrier heights. More specifically, the MAE for the experimental dataset amounts to 9.11 kcal mol⁻¹ and 13.83 kcal mol⁻¹ for the computational benchmarking dataset. Nevertheless, it is important to note that the main reactivity

trends are nicely recovered (R^2 of 0.80 and 0.89 for the respective datasets), implying that the xTB labels nonetheless have significant predictive value.

Next to the xTB dataset covering all the reactions present in the full reaction space, we also computed a DFT refinement on a subset of 1580 reactions, which were sampled based on a diversity criterion (see Section S4 and Fig. S6 of the SI for additional details).^{28,52} We used the M06-2X functional⁵³ (in combination with the def2-TZVP basis set⁵⁴) in our calculations,^{55,56} both because of its robust performance on the experimental and computational benchmarking datasets (R^2 of 0.90 and 0.97 respectively – though with an apparent tendency to slightly overestimate the barriers systematically *cf.* Sections S2.1 and S2.2), and because M06-2X performs exceptionally well in most kinetics benchmarking studies, particularly those on cycloaddition reactions.^{21,57} Overall, all DFT refinements took close to 200 000 CPU hours in total, implying that refining xTB reaction profiles across the entire reaction space would have cost on the order of 3 000 000 CPU hours.

Geometry-based model for accurate prediction of reaction characteristics

The geometry-based model, which we decided to call DeepReaction, leverages all information available from the xTB -generated reaction profiles and makes use of a DimeNet++ module for structural embedding.⁵⁸ The full architecture is displayed in Fig. 4. In short, the model takes five inputs: (xTB level) geometries of reactants, products, and TSs, as well as the xTB level activation (ΔG^\ddagger) and reaction ($\Delta_r G$) energy. Each geometry is embedded separately by passing it through DimeNet++, after which these embeddings are concatenated, together with both ΔG^\ddagger and $\Delta_r G$ to yield a final representation of the xTB reaction profile information. This representation is then passed through a feed-forward neural network (FFNN) and mapped to the targets, *i.e.*, the DFT level ΔG^\ddagger and $\Delta_r G$.

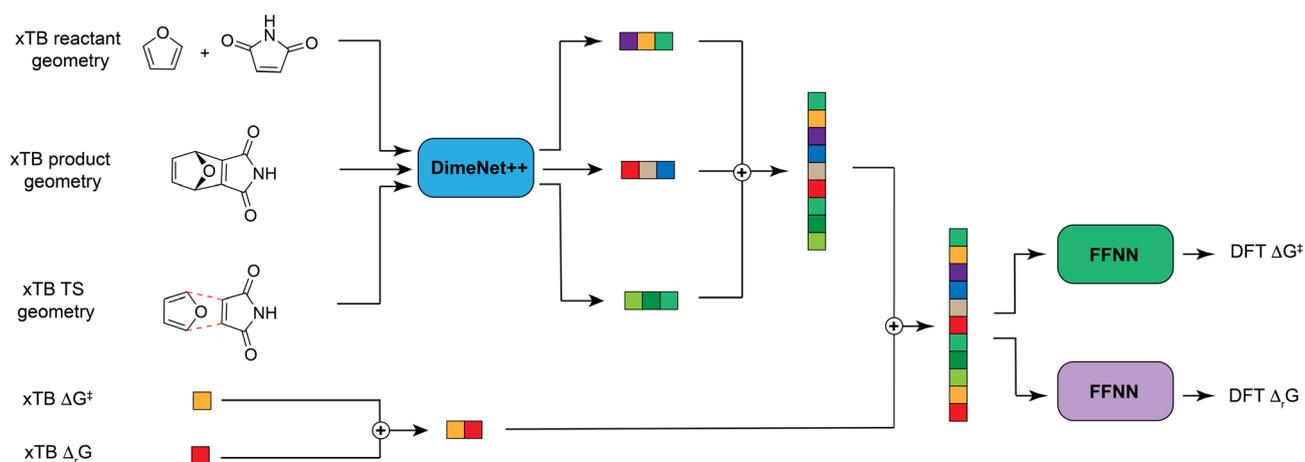


Fig. 4 A schematic overview of the geometry-based model for reaction characteristics prediction, DeepReaction. The model takes five inputs: (xTB level) geometries of reactants, products, and TSs, as well as the xTB level activation (ΔG^\ddagger) and reaction ($\Delta_r G$) energy. Each geometry is embedded separately by passing it through DimeNet++, after which these embeddings are concatenated, together with both ΔG^\ddagger and $\Delta_r G$ to yield a final representation of the xTB reaction profile information. This representation is then passed through a feed-forward neural network (FFNN) and mapped to the targets, *i.e.*, the DFT level ΔG^\ddagger and $\Delta_r G$.



A limited hyperparameter search was performed for this model; see SI Section S6 for details, as well as an overview of the final hyperparameters selected. A 10-fold cross-validation was used to estimate the performance of the model, with an 80/10/10 split for the training, validation, and test sets, in each fold. To avoid data leakage, all reactions involving a given diene–diphophile pair were kept grouped in a single subset in every fold.

Graph-based filtering model

To facilitate preliminary filtering of DA reactions and enable straightforward expansion of the search space (see Fig. 2), we decided to make use of the ChemProp model architecture.²⁹ Taking atom-mapped reaction SMILES strings as input, ChemProp generates Condensed Graphs of Reaction (CGR)⁵⁹ that are passed through a directed message-passing neural network. This graph-based model was set up as a multi-target model that outputs both predicted activation and reaction energies.

In the first instance, the model was trained on the small, DFT-level dataset. Additionally, we also trained a second version of the model on the larger 23 327 reaction dataset, where the (close to) DFT-quality labels, outputted by the geometry-based model, were used as the target. Details regarding hyperparameter search and selection can be found in Section S7 of the SI. The final performance of the model was again evaluated through 10-fold cross-validation, and similar precautions to avoid data leakage as for the geometry-based model were taken.

Results and discussion

Initial data analysis

Fig. 3D and E depict the distribution of the ΔG^\ddagger and $\Delta_r G$ for the 23 327 α TB level reaction profiles. Corresponding plots for the 1580 DFT-refined reactions can be found in Fig. S5 in the SI; similar distributions are obtained at both levels of theory, albeit with a shifted mean. In panel F, the correlation between the activation and reaction free energy is presented. With an R^2 value of 0.59, it is clear that the Bell–Evans–Polanyi principle^{60,61} is only partially recovered across the considered reaction space (for the DFT refined reactions, the R^2 value amounts to 0.58). The lack of a strong correlation suggests that kinetics and thermodynamics can be tuned fairly independently for DA reactions. As such, it should be possible to identify reactions that are thermoreversible and initiate self-healing over a broad temperature range within the designed chemical reaction space.

Another interesting preliminary observation is that computed reaction entropies, $\Delta S_{\text{rxn}}^\circ$, *i.e.*, the entropy differences (under standard conditions) between reactants and product, are distributed in a very narrow range at both α TB and DFT level of theory. This implies that fairly little control can be exerted over the entropy through molecular modulation across the defined reaction space. This is significant, since reaction entropy is one of the main determinants of whether a reaction will be thermoreversible or not, as it indicates the susceptibility of $\Delta_r G$ towards temperature changes.

Of course, the reaction entropy in an actual polymer environment will not be exactly the same as in the gas-phase environment assumed in our DFT simulations, and neither will this quantity stay perfectly constant across typical temperature ranges associated with self-healing plastics. Nonetheless, one can use the average $\Delta S_{\text{rxn}}^\circ$ to define a very crude rule-of-thumb to distinguish DA reactions that stand a reasonable chance of being thermoreversible from those that are highly unlikely to exhibit this behavior. A detailed discussion is provided in Section S5 of the SI. Our back-of-the-envelope analysis suggests that, focusing on temperature changes below 200 °C, reactions that are exothermic by more than 15–20 kcal mol⁻¹ ought to stand no realistic chance of being thermoreversible.

In the remainder of our analysis, this rule-of-thumb will be used as a guiding principle to gauge what fraction of the DA reactions screened are of potential interest for self-healing polymer applications.

Predicting DFT quality reaction characteristics across the search space

Trained on the 1580 DFT quality reaction profiles only, DeepReaction, *i.e.*, our geometry-based model, achieves an accuracy, expressed in terms of MAE, of 1.19 kcal mol⁻¹ for the ΔG^\ddagger target, and 1.72 kcal mol⁻¹ for $\Delta_r G$. The corresponding root-mean-square errors (RMSE) amount to 1.73 and 2.33 kcal mol⁻¹. To put this result in perspective, we also tested EquiReact, a state-of-the-art equivariant neural network for reactivity prediction, developed by van Gerwen and co-workers,⁶² on the same small dataset, and achieved MAEs of 3.20 and 4.09 kcal mol⁻¹, respectively (the corresponding RMSEs amount to 4.26 and 5.15 kcal mol⁻¹; see Section S6 in the SI). The graph-based model architecture ChemProp, in its turn, reaches an intermediate accuracy; 2.21 and 2.60 kcal mol⁻¹, respectively, in terms of MAE, and 3.03 and 3.78 kcal mol⁻¹ in terms of RMSE (*cf.* Section S7).²⁹

The underlying reason for the significantly superior accuracy of our model is the fact that we also provide the crude TS geometry, as well as the α TB level targets, as model inputs (see Section S6 for a concise ablation study). Generating these inputs doesn't come for free, of course, but since our in-house developed TS-tools code enables computation of full α TB-level reaction profiles at a cost of less than 1 CPU hour per reaction, we considered this a relatively small price to pay for the markedly enhanced accuracy. Especially when keeping in mind that computing an additional DFT quality training point – to expand the dataset and hence improve model accuracy of EquiReact/ChemProp – is two to three orders of magnitude more expensive on average (*vide supra*).

With our accurate geometry-based model, we predicted DFT quality ΔG^\ddagger and $\Delta_r G$ across our designated reaction space. In Fig. 5A, summarizing plots of the resulting data distribution are provided.

Validating chemical concepts across the reaction space

Now that accurate (predicted) reaction profiles are available for all reactions in the search space, the validity of some chemical



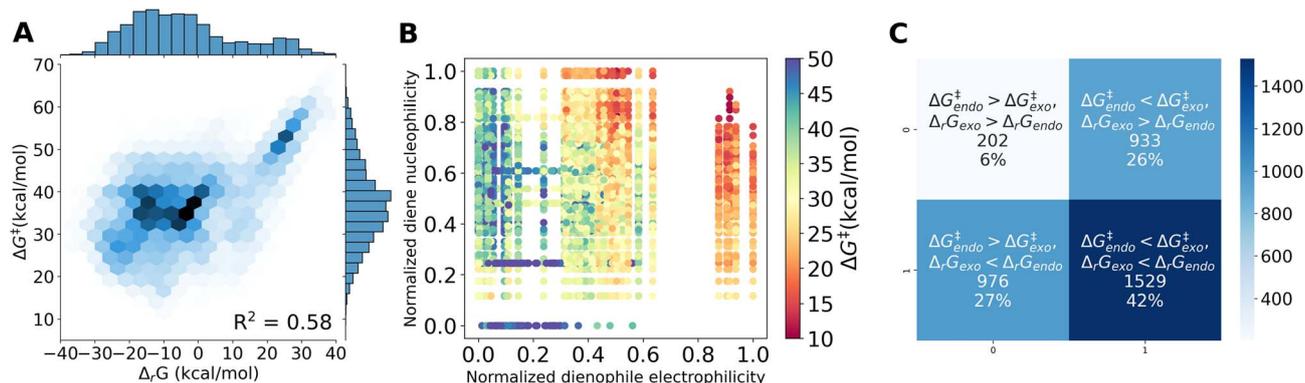


Fig. 5 (A) Bell–Evans–Polanyi plot for the synthetic reaction space, based on the energies predicted by our geometry-based model. (B) Representation of the activation free energies predicted as a scatter plot of the normalized dienophile electrophilicity and normalized diene nucleophilicity. (C) Confusion matrix for the heuristic rule that (i) *endo*-pathways tend to be kinetically preferred, and (ii) *exo*-pathways tend to be thermodynamically preferred. Diene–dienophile combinations for which the energy difference between the *endo*- and *exo*-stereoisomeric paths amounts to less than 1 kcal mol^{−1} were excluded from the analysis.

concepts across this reaction class can be evaluated. From Fig. 5A, one can already conclude that adherence to the Bell–Evans–Polanyi principle is only moderate, in line with the results presented for the *xTB* data in Fig. 3F.

Interestingly, we observe that reaction asynchronicity is a major driver of the observed deviations from the Bell–Evans–Polanyi principle, in line with previous, qualitative, observations.⁶³ Indeed, focusing exclusively on the most synchronous reactions, *i.e.*, those for which the bond distance ratio for the formed bonds in the TS is confined between 0.95 and 1.05, the correlation between ΔG^\ddagger and $\Delta_r G$ improves markedly, achieving an R^2 of 0.70 (Fig. S12 in the SI).

Additionally, we took a closer look at the predictive power of conceptual DFT reactivity descriptors (M06-2X/def2-TZVP).^{64–66} In Fig. 5B, a colormap of ΔG^\ddagger as a function of the dienophile electrophilicity ($\omega = \frac{\mu^2}{2\eta}$; $\mu = (\epsilon_{\text{LUMO}} + \epsilon_{\text{HOMO}})/2$; $\eta = (\epsilon_{\text{LUMO}} - \epsilon_{\text{HOMO}})/2$) and diene nucleophilicity ($N = \epsilon_{\text{HOMO}} - \epsilon_{\text{HOMO,ref}}$; the reference being tetracyanoethylene^{21,67}) is plotted. From this plot, one can observe, in line with previous observations derived from a handful of model reactions,²¹ that the electrophilicity of the dienophile appears to be correlated with the height of the reaction barrier. In contrast, no clear correlation is observed between the nucleophilicity of the diene and the activation barrier.

Interestingly, for $\Delta_r G$, diametrically opposite conclusions can be drawn (*cf.* Fig. S13 in the SI for the corresponding plot). For this quantity, an (inverse) correlation with the nucleophilicity of the diene is observed, and no clear relationship with the dienophile electrophilicity emerges. In line with the conclusions from the Bell–Evans–Polanyi plot, we find that the reactions with the lowest barrier, *i.e.*, those in the upper right corner of Fig. 5B, are often times slightly endothermic, and many promising reactions that have optimal reactive properties for self-healing polymer applications correspond to diene–dienophile combinations that do not exhibit maximized nucleophilicity and electrophilicity values, respectively. This underscores the

merit of our ML based approach to identify suitable reactant combinations.

Finally, we also aimed to validate the rule of thumb that the formation of *endo*-stereomers tends to be kinetically preferred (due to favorable through-space orbital interactions between the approaching reactants), whereas the formation of *exo*-stereomers tends to be thermodynamically preferred (due to more favorable steric interactions in the product). In Fig. 5C, a confusion matrix is shown in which every diene–dienophile pair is classified based on these two criteria, *i.e.*, whether the reaction pathway with the lowest activation energy is *endo* ($\Delta G^\ddagger_{endo} < \Delta G^\ddagger_{exo}$), and whether the most exothermic pathway is *exo* ($\Delta_r G_{exo} < \Delta_r G_{endo}$). It can be observed that this heuristic rule has only moderate predictive power: both criteria are fulfilled a mere 42% of the time; 6% of diene–dienophile combinations violate both criteria. Note that the confusion matrix obtained for the entire search space, based on the predicted values, matches almost perfectly the matrix obtained exclusively for the subset explicitly computed at the DFT level (Fig. S14).

Identifying plausible reactions for self-healing polymer applications

The predictions generated using our geometry-based model indicate that a large fraction of the search space, 12 480 reactions in total, exhibits a $\Delta_r G$ between 0 and -20 kcal mol^{−1}. The corresponding distribution of ΔG^\ddagger is shown in Fig. 6A. Clearly, the range of activation free energies is broad, with the bulk of them falling between 30 and 40 kcal mol^{−1}. To zoom in on the discovery of diene–dienophile pairs that could result in polymers that self-heal at relatively low temperatures, we imposed an initial cut-off at 30 kcal mol^{−1}. This corresponds to the computed ΔG^\ddagger of the furane–maleimide pair, the diene–dienophile pair for which self-healing properties were first observed, and for which the *retro*-DA reaction was triggered at around 150 °C.^{20,21} 2892 reactions are below this threshold, corresponding to 1160 unique diene–dienophile pairs, representing 16% of the total search space.



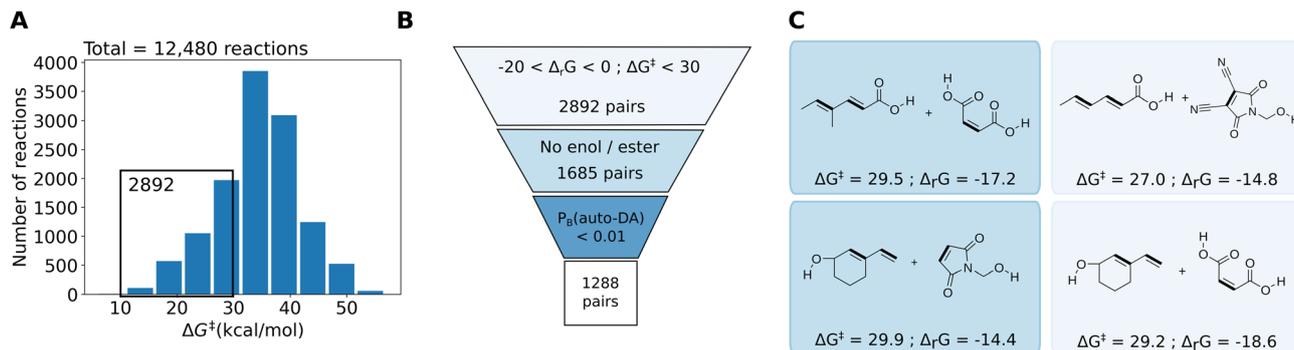


Fig. 6 (A) Distribution of the predicted activation energies (ΔG^\ddagger) for the reactions which satisfy $-20 < \Delta_r G < 0$ kcal mol $^{-1}$. (B) Schematic representation of the different filters applied to select the most relevant pairs. (C) Some representative examples of diene–dienophile pairs that have a predicted $\Delta_r G$ between -10 and 0 kcal mol $^{-1}$ and a predicted ΔG^\ddagger lower than 30 kcal mol $^{-1}$. All the energies reported are in kcal mol $^{-1}$.

Refining 100 randomly sampled examples from the set of retained reactions at DFT level confirms the excellent performance of our model, even when applied to reactions towards the edge of the training data distribution: for ΔG^\ddagger , an MAE of 1.36 kcal mol $^{-1}$ and an RMSE of 2.04 kcal mol $^{-1}$ is obtained between computed and predicted values; for $\Delta_r G$, an MAE of 2.06 kcal mol $^{-1}$ and an RMSE of 2.80 kcal mol $^{-1}$ is achieved.

Reducing the activation energy cut-off threshold rapidly reduces the number of retained reactions: at 25 kcal mol $^{-1}$, the number of reactions is reduced to 1357 (corresponding to 551 pairs); at 20 kcal mol $^{-1}$, only 476 reactions, *i.e.*, 209 pairs, are retained. This indicates that only a tiny fraction of the reaction search space involves diene–dienophile pairs that could give rise to polymers able to self-heal at relatively low temperatures.

To increase the odds of synthetic feasibility further, we subsequently applied an additional round of filtering to the reactions that fell within the suitable activation and reaction energy windows. Specifically, we removed dienes or dienophiles that could plausibly undergo ketone–enol tautomerism, that contained potentially hydrolysable motifs, or that exhibited competitive intramolecular DA reactions (see Fig. 6B; *cf.* Section S.5.3 for more information). We emphasize that these filters were applied in a deliberately conservative manner: in practice, the extent of tautomerization or hydrolytic instability is to some extent context- and substitution-dependent, and not all of the excluded motifs would necessarily be problematic under experimental conditions. However, by discarding the gray-zone cases, we ensured that the remaining subset represents a stringent lower bound of reactions that are unlikely to pose issues in a synthetic setting. After this filtering step, 1288 reactions survive; a sample of these promising candidate reactions is shown in Fig. 6C.

Setting up a fast graph-based filtering model

As indicated above, training the graph-based ChemProp architecture directly on the 1580 reaction DFT-level dataset results in an (interpolative) MAE of 2.21 and 2.60 kcal mol $^{-1}$ for ΔG^\ddagger and $\Delta_r G$, respectively. This is significantly worse than the performance achieved by our geometry-based model discussed above, and still relatively far from chemical accuracy. Nonetheless, this

model can be expected to robustly recover meaningful trends ($R^2 = 0.89$ and 0.94 for ΔG^\ddagger and $\Delta_r G$, respectively), and hence be employed, in principle, to screen previously uncharted territory – adjacent to the current search space – to identify reactions that could be of interest (*vide infra*).

An alternative version of this model was developed as well, where the more expansive 23 327 reaction dataset, with (close to) DFT quality labels predicted by the DeepReaction model, was used for training. On this extended dataset, the model achieves an (interpolation) accuracy of 1.05 and 1.40 kcal mol $^{-1}$, respectively. This result suggests that the graph-based ChemProp model can, in principle, reach accuracies similar to the ones of our geometry-based model, despite not having access to the geometric details of the reaction pathways. The main difference between the model architectures, however, lies in the data efficiency: ChemProp needs around an order of magnitude more – computationally expensive – reference data to reach a similar accuracy as DeepReaction (see Fig. S15 in the SI for the respective learning curves).

It is important to underscore that, in practice, the second version of the ChemProp model developed is not necessarily going to be more accurate than the first one, since the labels for the former inherently contain noise due to the prediction error of DeepReaction. At the same time, one might expect the latter to exhibit a better generalization potential, as it is trained on a broader reaction space. To test whether dataset expansion based on predicted labels can indeed be a viable tactic to increase model robustness, we will compare the performance of both ChemProp versions in the final sections below.

Expanding the covered chemical reaction space

As a first test of our envisioned hierarchical approach, where a graph-based model is used for pre-screening, after which a geometry-based model helps to fine-tune the predictions in regions of interest (*cf.* Fig. 2), we explicitly enumerated a new DA reaction space, adjacent to the one previously considered. 15 889 reactions were generated, by exhaustively combining 87 new dienes with the previously considered dienophiles (see Fig. 7A for an overview of the new dienes probed).



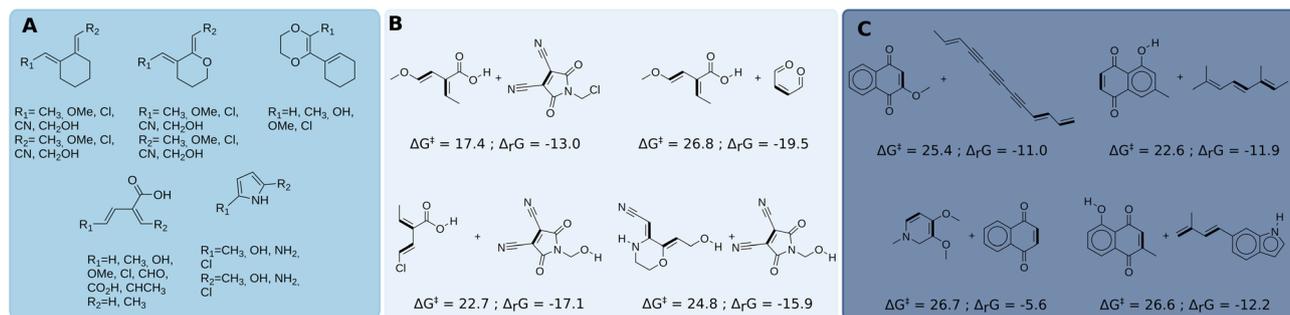


Fig. 7 (A) The 87 dienes included in the expansion of the chemical reaction space. (B) Representative examples of reactions found to be reversible by the DeepReaction model in the expansion of the chemical space. (C) Representative examples of reactions found to be reversible by the DeepReaction model in the ZINC NP database. In bold the bonds involved in the reaction. The energies reported are in kcal mol⁻¹ and they were computed at the M06-2X/def2-TZVP level.

Since the original synthetic reaction space, enumerated and analyzed before, contained fairly few diene–dienophile pairs that could potentially give rise to polymers that self-heal at mild temperatures, we focus in the first instance on discovering more of those. As such, we applied (the first version of) our ChemProp model, and selected all the reactions in the newly generated space predicted to be exothermic by up to 20 kcal mol⁻¹, and to exhibit barriers below 25 kcal mol⁻¹. For the resulting 197 reactions, a geometry-based validation, followed by DFT confirmation, was performed. Comparing the computed labels with the ones predicted by the different models, immediately reveals that the accuracies of all models deteriorate notably – which is not entirely surprising, as we are now unequivocally extrapolating well beyond the original training domain (*cf.* Section S12 of the SI). For DeepReaction, the deterioration in performance is limited in absolute terms (MAE and RMSE for ΔG^\ddagger amount to 1.9 and 2.5 kcal mol⁻¹, respectively; for $\Delta_r G$, 2.8 and 3.6 kcal mol⁻¹ are obtained). The ChemProp models, on the other hand, become markedly less reliable in extrapolation mode (MAEs for both targets fall in a range between 4 and 5.5 kcal mol⁻¹).

It is nonetheless noteworthy that the second version of the ChemProp model, *i.e.*, the one trained on all 23 327 labelled data points, fares significantly better than the first version. For $\Delta_r G$, the improvement in MAE amounts to 1 kcal mol⁻¹, for ΔG^\ddagger 0.5 kcal mol⁻¹. The superior performance of the second version is however most accentuated in terms of the correlation coefficient between predicted and computed ΔG^\ddagger : whereas the first model fails to capture any fine-grained trend and simply clusters the barriers for all selected reactions around 20 kcal mol⁻¹ ($R^2 = 0.03$), the second ChemProp model does manage to differentiate between relatively high- and low-barriers within the selection window reasonably well ($R^2 = 0.48$).

These observations nicely confirm our expectation regarding the relative generalization potential of these models, and demonstrate – at least for this first test case – that the gains from expanding the training set size trump the uncertainty induced by using (noisy) DeepReaction predicted labels.

In total, 107 reactions were validated with DFT to adhere to the criteria indicated above. Some representative examples are shown in Fig. 7B. Clearly, we can expand the pool of plausible

thermoreversible DA reactions with low barriers with our hierarchical approach, and this at a marginal computational cost.

Further expansion towards bio-based reactants

From a sustainability perspective, it would be particularly appealing if (self-healing) plastics could be synthesized starting from bio-based building blocks. As such, as a final application, we also aimed to mine ZINC NP,^{68,69} a database containing 76 615 commercially available natural products, for suitable DA reactions for self-healing polymer applications. A full overview of the filtering steps taken to obtain a list of plausible dienes and dienophiles can be found in Section S12 of the SI. In total, 84 potential dienes and 423 potential dienophiles are identified, which leads to 308 261 DA reactions in total. We emphasize that, in the present context, ‘natural products’ are primarily used to probe chemical diversity, not to imply direct feasibility as biomass-derived feedstocks; economic and production-scale considerations are a separate matter that would require further scrutiny.

As a first step, we applied the same strategy as above, *i.e.*, the first ChemProp model was used to pre-select promising reactions, after which DeepReaction validations, and DFT confirmations, were performed (see Section S12 in the SI for a detailed discussion). Similar trends are obtained as before, *i.e.*, DeepReaction retains its robustness the best, still reproducing the DFT labels with reasonable accuracy, and the second ChemProp model significantly outperforms the first one. Particularly noteworthy here is the extremely poor extrapolation performance of the first ChemProp version: for ΔG^\ddagger , its MAE amounts to a whopping 11 kcal mol⁻¹ (almost 4 kcal mol⁻¹ worse than the second version); for $\Delta_r G$, the MAE amounts to a slightly less dramatic 5.3 kcal mol⁻¹ (which is still almost 1 kcal mol⁻¹ worse than the second model).

Based on these results, we decided to perform the final pre-screening of the ZINC NP database with the second ChemProp model. Across the enumerated reaction space, 4245 reactions are identified as potentially suitable. Grouping these reactions by diene–dienophile pair, and selecting for each pair the reaction with the lowest predicted activation energy, 1277 reactions were selected for χ TB reaction profile computation, followed by DeepReaction refinement. Based on these final predictions, the



32 most promising reactions were selected for a final DFT validation. Gratifyingly, all these reactions are computed to adhere to the first criterion listed above, *i.e.*, $\Delta_r G$ between -20 and 0 kcal mol $^{-1}$; the second criteria, *i.e.* $\Delta G^\ddagger < 30$ kcal mol $^{-1}$, is respected by 28 of them, with the highest computed barrier amounting to 32.95 kcal mol $^{-1}$ (see Fig. 7C for some representative examples of validated reactions).

It is noteworthy that even for the reactions, emerging as the most promising from our analysis, the computed barriers are all clustered within the 25–30 kcal mol $^{-1}$ window. This suggests that the potential to discover purely bio-based reactions for ultra-low temperature self-healing plastics may be limited, and that synthetic reaction spaces like the ones outlined above provide more opportunities to identify DA reactions with optimal reactive properties.

Nonetheless, the analysis above clearly demonstrates that, in general, our hierarchical workflow also enables identification of promising bio-based candidate reactions for self-healing polymer applications, even though we did not set up our models explicitly for this purpose. Due to the limited generalization ability of ChemProp, plenty of “false positive” reactions tend to survive the initial filtering step this time around, but the secondary filtering with our geometry-based model readily identifies those, so that computational resources can be focused on DFT validation of the reactions that are the most likely to have desirable properties.

Conclusion

In this study, we have presented a hierarchical, data-efficient workflow to efficiently discover suitable DA reactions for self-healing polymer applications. Leveraging our in-house TS-tools software to rapidly generate reaction profiles, we built two complementary ML models: a rapid, yet less accurate, graph-based model, and a slower, but more precise, geometry-based one. We demonstrated how the former can be used to crudely screen extensive chemical reaction spaces, consisting of hundreds of thousands, or even millions, of individual reactions on-the-fly, while the latter facilitates a more fine-grained, DFT-quality validation. Interestingly, we observe that using the latter model to generate synthetic training data for the former improves model robustness and accuracy, opening an interesting avenue to reduce the cost of training data generation in chemical reaction discovery/screening studies. We demonstrated the capabilities and accuracy of our approach by screening on the one hand a comprehensive synthetic reaction space, covering most of the building blocks previously considered within the context of DA reactivity for self-healing polymer applications, and on the other hand, mining a larger database of commercially available natural products, ZINC NP. We discovered a broad range of suitable reactions across both databases, at negligible computational cost. All the identified candidate DA reactions, as well as all the models and datasets generated as part of this study, have been made publicly available and will hopefully inspire the development of new self-healing plastics, tailored to specific applications.

Author contributions

M. F.: software, methodology, formal analysis, visualization, writing. B. D.: software, formal analysis. J. E. A. R.: software. T. S.: conceptualization, methodology, formal analysis, writing, supervision, funding acquisition.

Conflicts of interest

There are no conflicts to declare.

Data availability

The code of the geometry-based DeepReaction model is publicly available through GitHub: <https://github.com/chimie-paristech-CTM/DeepReaction>, <https://doi.org/10.5281/zenodo.17367023>. The xTB- and DFT-level datasets, as well as a list of all suitable DA reactions identified from the synthetic search space and the zinc NP database (with the respective activation and reaction energy predictions), can be downloaded from Figshare: https://figshare.com/projects/Screening_Diels-Alder_reaction_space_to_identify_candidate_reactions_for_self-healing_polymer_applications/244595. The xTB and DFT datasets are available at <https://doi.org/10.6084/m9.figshare.28930874.v1>, and <https://doi.org/10.6084/m9.figshare.28936460.v1>.

Supplementary information is available. See DOI: <https://doi.org/10.1039/d5dd00340g>.

Acknowledgements

This research was supported by the French National Agency for Research (ANR) through a CPJ grant (ANR-22-CPJ1-0093-01) and a JJC grant (ANR-24-CE29-5745). This work was granted access to the HPC resources of IDRIS under the allocation A0170810135 granted by GENCI.

References

- 1 J. Kurowiak, T. Klekiel and R. Będziński, *Int. J. Mol. Sci.*, 2023, **24**, 16952.
- 2 M. A. Ward and T. K. Georgiou, *Polymers*, 2011, **3**, 1215–1242.
- 3 M. Angelopoulos, *IBM J. Res. Dev.*, 2001, **45**, 57–75.
- 4 Q. Lin, *Polymer*, 2023, **286**, 126395.
- 5 V. T. Rathod, J. S. Kumar and A. Jain, *Appl. Nanosci.*, 2017, **7**, 519–548.
- 6 M. Bashir and P. Rajendran, *J. Intell. Mater. Syst. Struct.*, 2018, **29**, 3681–3695.
- 7 Y.-L. Liu and T.-W. Chuo, *Polym. Chem.*, 2013, **4**, 2194–2205.
- 8 S. Wang and M. W. Urban, *Nat. Rev. Mater.*, 2020, **5**, 562–583.
- 9 Y. Zhou, L. Li, Z. Han, Q. Li, J. He and Q. Wang, *Chem. Rev.*, 2023, **123**, 558–612.
- 10 S. Terryn, J. Langenbach, E. Roels, J. Brancart, C. Bakali-Hassani, Q.-A. Poutrel, A. Georgopoulou, T. George Thuruthel, A. Safaei, P. Ferrentino, T. Sebastian, S. Norvez, F. Iida, A. W. Bosman, F. Tournilhac, F. Clemens, G. Van



- Assche and B. Vanderborght, *Mater. Today*, 2021, **47**, 187–205.
- 11 Y. J. Tan, G. J. Susanto, H. P. Anwar Ali and B. C. K. Tee, *Adv. Mater.*, 2021, **33**, 2002800.
- 12 S. Kashef Tabrizian, S. Terryn and B. Vanderborght, *Adv. Intell. Syst.*, 2025, 2400790.
- 13 A. J. R. Amaral and G. Pasparakis, *Polym. Chem.*, 2017, **8**, 6464–6484.
- 14 B. Ma, Y. Zhang, J. Li, D. Chen, R. Liang, S. Fu and D. Li, *Chem. Eng. J.*, 2023, **466**, 143420.
- 15 N. I. Khan, S. Halder, S. B. Gunjan and T. Prasad, *IOP Conf. Ser.: Mater. Sci. Eng.*, 2018, **377**, 012007.
- 16 M. A. Tasdelen, *Polym. Chem.*, 2011, **2**, 2133–2145.
- 17 A. Sanyal, *Macromol. Chem. Phys.*, 2010, **211**, 1417–1425.
- 18 J. Sauer, *Angew. Chem., Int. Ed.*, 1967, **6**, 16–33.
- 19 J. G. Martin and R. K. Hill, *Chem. Rev.*, 1961, **61**, 537–562.
- 20 X. Chen, M. A. Dam, K. Ono, A. Mal, H. Shen, S. R. Nutt, K. Sheran and F. Wudl, *Science*, 2002, **295**, 1698–1702.
- 21 L. Vermeersch, R. Verhelle, N. Van den Brande and F. De Vleeschouwer, *Macromolecules*, 2025, **58**, 32–44.
- 22 M. Diaz, G. Van Assche, F. Maurer and B. Van Mele, *Polymer*, 2017, **120**, 176–188.
- 23 G. Scheltjens, J. Brancart, I. De Graeve, B. Van Mele, H. Terryn and G. Van Assche, *J. Therm. Anal. Calorim.*, 2011, **105**, 805–809.
- 24 A. M. Peterson, R. E. Jensen and G. R. Palmese, *ACS Appl. Mater. Interfaces*, 2010, **2**, 1141–1149.
- 25 T. S. Coope, D. H. Turkenburg, H. R. Fischer, R. Luterbacher, H. van Bracht and I. P. Bond, *Smart Mater. Struct.*, 2016, **25**, 084010.
- 26 J. Sauer and R. Sustmann, *Angew. Chem., Int. Ed.*, 1980, **19**, 779–807.
- 27 P. Vermeeren, T. A. Hamlin, I. Fernández and F. M. Bickelhaupt, *Angew. Chem.*, 2020, **132**, 6260–6265.
- 28 N. Casetti, J. E. Alfonso-Ramos, C. W. Coley and T. Stuyver, *Chem.–Eur. J.*, 2023, **29**, e202301957.
- 29 E. Heid, K. P. Greenman, Y. Chung, S.-C. Li, D. E. Graff, F. H. Vermeire, H. Wu, W. H. Green and C. J. McGill, *J. Chem. Inf. Model.*, 2023, **64**, 9–17.
- 30 L.-Y. Chen, T.-W. Hsu, T.-C. Hsiung and Y.-P. Li, *J. Phys. Chem. A*, 2022, **126**, 7548–7556.
- 31 B. Deng and T. Stuyver, *J. Chem. Inf. Model.*, 2025, **65**, 649–659.
- 32 S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt and B. Kozinsky, *Nat. Commun.*, 2022, **13**, 2453.
- 33 P. van Gerwen, A. Fabrizio, M. D. Wodrich and C. Corminboeuf, *Mach. Learn.: Sci. Technol.*, 2022, **3**, 045005.
- 34 E. H. Farrar and M. N. Grayson, *Chem. Sci.*, 2022, **13**, 7594–7603.
- 35 K. A. Spiekermann, T. Stuyver, L. Pattanaik and W. H. Green, *Mach. Learn.: Sci. Technol.*, 2023, **4**, 048001.
- 36 T. Stuyver and C. W. Coley, *Chem.–Eur. J.*, 2023, **29**, e202300387.
- 37 T. Stuyver, K. Jorner and C. W. Coley, *Sci. Data*, 2023, **10**, 66.
- 38 J. E. Alfonso-Ramos, R. M. Neeser and T. Stuyver, *Digital Discovery*, 2024, **3**, 919–931.
- 39 P. Friederich, G. dos Passos Gomes, R. De Bin, A. Aspuru-Guzik and D. Balcells, *Chem. Sci.*, 2020, **11**, 4584–4601.
- 40 S. Heinen, G. F. von Rudorff and O. A. von Lilienfeld, *J. Chem. Phys.*, 2021, **155**, 064105.
- 41 G. F. von Rudorff, S. N. Heinen, M. Bragato and O. A. von Lilienfeld, *Mach. Learn.: Sci. Technol.*, 2020, **1**, 045026.
- 42 C. A. Grambow, L. Pattanaik and W. H. Green, *Sci. Data*, 2020, **7**, 137.
- 43 C. A. Grambow, L. Pattanaik and W. H. Green, *J. Phys. Chem. Lett.*, 2020, **11**, 2992–2997.
- 44 Q. Zhao, S. M. Vaddadi, M. Woulfe, L. A. Ogunfowora, S. S. Garimella, O. Isayev and B. M. Savoie, *Sci. Data*, 2023, **10**, 145.
- 45 S. G. Espley, E. H. E. Farrar, D. Buttar, S. Tomasi and M. N. Grayson, *Digital Discovery*, 2023, **2**, 941–951.
- 46 T. Lewis-Atwell, P. A. Townsend and M. N. Grayson, *WIREs Comput. Mol. Sci.*, 2022, **12**, e1593.
- 47 T. Stuyver, *J. Comput. Chem.*, 2024, **45**, 2308–2317.
- 48 B. Briou, B. Améduri and B. Boutevin, *Chem. Soc. Rev.*, 2021, **50**, 11055–11097.
- 49 S. W. Gabrielson, *J. Med. Libr. Assoc.*, 2018, **106**, 588.
- 50 C. W. Coley, W. H. Green and K. F. Jensen, *J. Chem. Inf. Model.*, 2019, **59**, 2529–2537.
- 51 C. Bannwarth, S. Ehlert and S. Grimme, *J. Chem. Theory Comput.*, 2019, **15**, 1652–1671.
- 52 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, *et al.*, *J. Mach. Learn. Res.*, 2011, **12**, 2825–2830.
- 53 Y. Zhao and D. G. Truhlar, *Theor. Chem. Acc.*, 2008, **120**, 215–241.
- 54 F. Weigend and R. Ahlrichs, *Phys. Chem. Chem. Phys.*, 2005, **7**, 3297–3305.
- 55 M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, G. Scalmani, V. Barone, G. A. Petersson, H. Nakatsuji, X. Li, M. Caricato, A. V. Marenich, J. Bloino, B. G. Janesko, R. Gomperts, B. Mennucci, H. P. Hratchian, J. V. Ortiz, A. F. Izmaylov, J. L. Sonnenberg, D. Williams-Young, F. Ding, F. Lipparini, F. Egidi, J. Goings, B. Peng, A. Petrone, T. Henderson, D. Ranasinghe, V. G. Zakrzewski, J. Gao, N. Rega, G. Zheng, W. Liang, M. Hada, M. Ehara, K. Toyota, R. Fukuda, J. Hasegawa, M. Ishida, T. Nakajima, Y. Honda, O. Kitao, H. Nakai, T. Vreven, K. Throssell, J. A. Montgomery Jr, J. E. Peralta, F. Ogliaro, M. J. Bearpark, J. J. Heyd, E. N. Brothers, K. N. Kudin, V. N. Staroverov, T. A. Keith, R. Kobayashi, J. Normand, K. Raghavachari, A. P. Rendell, J. C. Burant, S. S. Iyengar, J. Tomasi, M. Cossi, J. M. Millam, M. Klene, C. Adamo, R. Cammi, J. W. Ochterski, R. L. Martin, K. Morokuma, O. Farkas, J. B. Foresman and D. J. Fox, *Gaussian 16 Revision C.01*, Gaussian Inc., Wallingford CT, 2016.
- 56 G. Luchini, J. V. Alegre-Requena, I. Funes-Ardoiz and R. S. Paton, *F1000Research*, 2020, **9**, 291.
- 57 J. E. Alfonso-Ramos, C. Adamo, É. Brémond and T. Stuyver, *J. Chem. Theory Comput.*, 2025, **21**, 1752–1761.
- 58 J. Gasteiger, S. Giri, J. T. Margraf and S. Günemann, Fast and uncertainty-aware directional message passing for



- non-equilibrium molecules, *arXiv*, 2020, arXiv:2011.14115, DOI: [10.48550/arXiv.2011.14115](https://doi.org/10.48550/arXiv.2011.14115).
- 59 E. Heid and W. H. Green, *J. Chem. Inf. Model.*, 2022, **62**, 2101–2110.
- 60 R. P. Bell and C. N. Hinshelwood, *Proc. R. Soc. London, Ser. A*, 1936, **154**, 414–429.
- 61 M. G. Evans and M. Polanyi, *Trans. Faraday Soc.*, 1936, **32**, 1333–1360.
- 62 P. van Gerwen, K. R. Briling, Y. Calvino Alonso, M. Franke and C. Corminboeuf, *Digital Discovery*, 2024, **3**, 932–943.
- 63 C. Barrales-Martínez and P. Jaque, *Phys. Chem. Chem. Phys.*, 2022, **24**, 14772–14779.
- 64 P. Geerlings, F. De Proft and W. Langenaeker, *Chem. Rev.*, 2003, **103**, 1793–1874.
- 65 P. Geerlings, S. Fias, Z. Boisdenghien and F. De Proft, *Chem. Soc. Rev.*, 2014, **43**, 4989–5008.
- 66 D. Chakraborty and P. K. Chattaraj, *Chem. Sci.*, 2021, **12**, 6264–6279.
- 67 M. Ríos-Gutiérrez, A. Saz Sousa and L. R. Domingo, *J. Phys. Org. Chem.*, 2023, **36**, e4503.
- 68 M. Sorokina, P. Merseburger, K. Rajan, M. A. Yirik and C. Steinbeck, *J. Cheminf.*, 2021, **13**, 2.
- 69 V. Chandrasekhar, K. Rajan, S. R. S. Kanakam, N. Sharma, V. Weißenborn, J. Schaub and C. Steinbeck, *Nucleic Acids Res.*, 2024, **53**, D634–D643.

