Digital Discovery



PAPER

View Article Online
View Journal | View Issue



Cite this: Digital Discovery, 2025, 4,

Received 23rd July 2025 Accepted 16th September 2025

DOI: 10.1039/d5dd00327j

rsc.li/digitaldiscovery

Machine learning driven design of spiropyran photoswitches

This study presents the development and application of a generative machine learning model for the design of novel spiropyran photoswitches with enhanced switching speed and absorption bands with small spectral overlap between the open and closed form (i.e. high addressability). Leveraging a scaffold decoration approach, we fine-tuned a general chemical recurrent neural network (RNN) model on a curated dataset of photoswitches. The fine-tuned model was evaluated against both the pretrained baseline and literature-reported spiropyran compounds, demonstrating superior performance in generating diverse and novel candidates. Notably, the fine-tuned model effectively mitigates common biases in decoration patterns and functional group selection observed in the literature. The study also outlines the synthesis and experimental characterization of several newly designed spiropyran photoswitches, validating the design principles derived from the generative model. These findings highlight the potential of generative models in accelerating the discovery of advanced molecular photoswitches with tailored properties.

1. Introduction

Photoswitches are molecules that isomerize upon exposure to light and whose initial state can subsequently be restored by irradiation with another wavelength, or via a thermal process with a characteristic lifetime τ (see Fig. 1(a)). Molecular photoswitches have been known for over 150 years. In this time a wide variety of classes of switches have been developed, including azobenzenes, ²⁻⁴ spiropyrans^{5,6} (SP), diarylethenes^{7,8} and others. Possible applications range from photosensitive devices and data storage, ^{12,13} over photopharmacology, ^{14,15} to molecular machines. ^{16,17}

This wide range of applications leads to different demands for the switching behavior. Thermal backreactions are for instance not desired in data storage applications, since the molecules need to be switched back and forth in a controlled manner. This class of switches are referred to as P-type photoswitches in the literature (P indicating a photochemical backreaction). The ideal timescale for thermal backreactions in other applications (using T-type thermal switches) can range from minutes in drug delivery to less than seconds in 3D printing applications. ^{18,19} Furthermore, the wavelengths used

for switching must often be in some desired range due to technical as well as material/biological boundary conditions (such as tissue penetration depth).

All of this implies that there is no single optimal photoswitch, but rather a range of optimal switches for different applications. Even for a given application, the existence of multiple competing design targets will usually lead to a Pareto front of candidates that represent different trade-offs between the target properties. Overall, photoswitch design is thus a highly challenging task. Unfortunately, this challenge is still most commonly addressed via trial and error. Here, a common strategy is to modify known photoswitches, for example by adding or exchanging side groups. This approach is prone to human bias towards specific chemistries, e.g. based on previous knowledge or synthetic accessibility. It is also inefficient, requiring a large number of time-consuming syntheses to cover a small volume of chemical space. Indeed, this is a general problem of molecular design and by no means limited to photoswitches.

More unbiased and efficient approaches for molecular design based on machine learning (ML) have therefore been at the focus of much attention in recent years, for example in drug discovery and materials design.^{20–26} Here, generative models, which implicitly learn molecular construction rules, are a particularly elegant approach. Unlike virtual screening or combinatorial approaches, generative models can transcend currently known chemical libraries, allowing for true *de novo* design. Different ML architectures have been used for this purpose, including Recurrent Neural Networks (RNN),²⁷

^aFritz-Haber-Institute of the Max-Planck-Society, Faradayweg 4-6, 14195 Berlin, Germany. E-mail: reuter@fhi.mpg.de

^bHumboldt-Universität zu Berlin, Brook-Taylor-Straße 2, 12489 Berlin, Germany. E-mail: sh@chemie.hu-berlin.de

University of Bayreuth, Bavarian Center for Battery Technology (BayBatt), Weiherstraße 26, 95448 Bayreuth, Germany. E-mail: johannes.margraf@uni-bayreuth.de




Fig. 1 (a) Potential energy surface scheme for a Spiropyran (SP) photoswitch. (b) Thermostability performance criterion described by a characteristic bond length of the merocyanine (MC) form. (c) Addressability performance criterion described by a ratio of the spectra of the closed SP (blue) and open MC (red) forms.

variational auto-encoders,²⁸ autoregressive deep neural networks,²⁹ and generative adversarial networks.³⁰ These approaches have also been used specifically in the context of the popular class of azobenzene photoswitches, including ML enhanced discovery,³¹ excitonic property learning,³² and tuning of thermal barriers with ML interatomic potentials.³³

In this paper, we explore the use of generative ML for the photoswitch design challenge. As a use case, we focus on the application of xolography, a volumetric 3D printing method, ¹⁹ based on SP photoswitches serving as spatially addressable dual-color photoinitiators. This enables the fast layer-wise printing of millimeter sized objects in a box of viscous resin, using an orthogonal laser-projector setup. Specifically, closed SP molecules are first opened to their merocyanine (MC) form by irradiation with wavelength λ_A (see Fig. 1(a)). In a second step, the MC form is photochemically excited with another light source (projector) to initialize a radical polymerization cascade. Because the MC form is only present in places that were previously irradiated with λ_A , the polymerization can thus be initiated in a highly localized fashion by the projector.

From a material standpoint, the speed of the printing process is primarily governed by the thermal lifetime of the open MC form of the switch. A perfect photoswitch for this application should display a fast thermal backreaction to the SP form, with a lifetime in the order of the illumination time of the second light source (projector). In other words, the thermostability of the open form should be low (see Fig. 1(b)). A second requirement for this application is that the closed form can be excited independently from the open form, in order to avoid unwanted polymerization initiated by λ_A *via* residual open MC molecules in the resin. It is thus desirable that there is a suitable wavelength λ_A with small spectral overlap between the SP and MC forms. This property is referred to as addressability in the following (Fig. 1(c)).

2. Methods

2.1 Design criteria

The thermostability and addressability of a SP photoswitch can be understood by considering the potential energy surface (PES) scheme in Fig. 1(a). First, the excitation wavelength λ_A must be energetic enough to allow excitation of the SP form, in order to induce switching to the MC form. The thermal backreaction then requires overcoming a free energy barrier ΔG^{\ddagger} , which defines the thermostability. Finally, the addressability depends on the nature of the ground and excited states of both the SP and MC form, from which the quotient of oscillator strengths at each wavelength can be computed. This ratio defines the addressability as the likelihood of exciting the SP form instead of the MC form, given equal concentrations of both forms. In order to train a generative ML model to design molecules with low thermostability and high addressability, we now need computationally affordable descriptors which predict these properties reasonably well.

Following Fig. 1(a), the most intuitive descriptor for thermostability would be the energetic barrier for the transition from the MC form back to the SP form. Unfortunately, this schematic depiction is overly simplistic since multiple conformers exist for both the SP and MC forms, so that in general several transition states lie between the SP and MC forms. This makes a reliable high-throughput prediction of reaction barriers extremely challenging. Furthermore, computationally efficient semiempirical methods like GFN2-xTB that work well for predicting molecular geometries are not sufficiently accurate for barriers.³⁴ As a consequence, the computational effort of predicting kinetic barriers accurately would be very high (Fig. 3).

We therefore define a heuristic descriptor for thermostability, which is physically motivated and cheap to compute. In a previous investigation of a prototypical SP photoswitch, it was Digital Discovery Paper

Fig. 2 Slicing scheme for a training molecule. Slicing along aliphatic bonds can lead to different definitions of the scaffold and the decoration. Stars indicate open decoration positions.

Fig. 3 6-Nitro-1'-3'-3'-trimethylspiro[2H-1-benzopyran-2,2'-indolin] as the most cited SP molecule.

found that the backreaction follows a rotation along the highlighted bond in Fig. 1(b).35 The MC structure has two resonance structures (a fully conjugated and a zwitterionic one), depending on which the rotating bond is formally either a double or single bond. We can thus infer that the thermostability is lower if the zwitterionic character of the structure is higher. We therefore use the length of the rotating bond in the MC form (predicted with the GFN2-xTB method, see Computational details) as a proxy for inverse thermostability of a given structure. While this approach is certainly not perfect, it yields a descriptor that is robust and fast to compute. For convenience, bond lengths are reported as relative deviations to the prototypical SP 6-nitro-1'-3'-3'-trimethylspiro[2H-1 benzopyran-2,2'-indolin] (6-nitro-BIPS, see Fig. 3)36 so that values above 0% indicate longer bonds (and therefore lower predicted thermostabilities).

As described above, the main ingredient needed to compute the addressability are ground and excited state energies of both forms, as well as the oscillator strengths σ for each transition. As for the barriers, semiempirical calculations proved to be of too low accuracy for this purpose in preliminary tests (see Fig. S1). We therefore used Time-Dependent Density Functional Theory (TD-DFT) calculations with the hybrid PBE0 functional and def2-TZVP basis set, from which Gaussian broadened UV/vis spectra are obtained (see Computational details). The quotient of SP and MC spectra, which is used to define the addressability, is shown in green in Fig. 1(c).

With this, computationally affordable descriptors for thermostability and addressability are defined, which can be used to train and validate the generative model.

2.2 Generative model

In contrast to general molecular design tasks, the current project requires preserving the photoswitchable SP core in all generated molecules. In addition, we need a data efficient model due to the scarcity of known SP reference molecules. To achieve this, we combine a transfer learning approach with the *REINVENT* scaffold-decoration architecture.³⁷ REINVENT is a long-short term memory RNN (LSTM-RNN), which can be trained on the SMILES molecular string representation. String representations show excellent performance in quality metrics for generation tasks and are competitive with graph based approaches.³⁸ More importantly, large chemical databases (like *ChEMBL*³⁹) contain millions of SMILES, which can be used for pretraining general chemical models.

The *REINVENT* scaffold-decoration model is trained in a self-supervised manner, based on scaffold-decoration pairs that are created with a slicing scheme applied to all training molecules (see Fig. 2). Each molecule in general yields several scaffold-decoration pairs depending on its size and the sliced bonds. This procedure allows the model to learn the grammar of chemical modifications from the examples in the training set. The trained model can subsequently be used to decorate any target scaffold on arbitrary predefined decoration positions. Importantly, it can also be applied to scaffolds that were not present in the training set and can generate completely novel decoration groups.

2.3 Transfer learning

A model that is pretrained on a general database like ChEMBL can already be used to generate novel SP molecules. However, the goal of molecular design is not just the generation of new, chemically reasonable candidates, but the generation of candidates with desired properties. To achieve this, we apply transfer learning (TL), which works by retraining a baseline model (e.g. trained on ChEMBL) on performant candidates with the desired properties. This general idea was first applied in chemistry by Segler et al.,40 and subsequently further adapted for a variety of systems in the fields of drug discovery, 41,42 catalysis43 and functional materials.44,45 Ideally, TL can guide the generation process towards systems of interest, while still maintaining the generalizability of the previously trained baseline model. Alternatively, generative models can also be trained in a conditional manner, but this requires large labelled datasets, which are not available here.46

The full generative TL workflow (see Fig. 4) starts with pretraining a *REINVENT* baseline model on a subset of *ChEMBL*, which is a manually curated database of over 2M bioactive molecules. Herein, training molecules have been filtered as in the work of Arús-Pous *et al.*³⁷ for size, number of rings and other aspects. The dataset was further tailored for the SP application by only selecting systems which feature a benzene or pyrrolidine moiety (the two substructures which occur in the SP photochromic core). In total 783 k molecules resulted from this filtering, yielding a training set of 36.8 M scaffold-decoration pairs after slicing. The resulting model is referred to as the baseline model in the following. Further details on the training




Fig. 4 Transfer learning workflow. From left to right: initial model is trained with a big chemical database (ChEMBL) and used to query all decoration pattern. The sampled molecules are analyzed according to the performance criteria and the hit candidates are used for retraining.

data and training process can be found in the Computational details.

In the next step, the possible SP decoration patterns are provided to the model. For two side-groups and 10 decoration sites, a total number of 45 unique decoration patterns can be identified. Here, the two open positions at the double bond in the central pyran ring are not considered because they are synthetically inaccessible. Since the generation with the RNN model is probabilistic, each decoration pattern is queried multiple times to sample extensively from the chemical space of possible SP molecules. It is pertinent to emphasized that the REINVENT model does not merely memorize the decorations in the training set, but generates new side-groups, not present in the input data. Indeed, 52.7% of the unique decorations predicted by the final model were not part of the training data.

We next draw a sample from these new photoswitch candidates (144 molecules for each of the 45 decoration patterns) and compute the addressability and thermostability descriptors for each. From this, a set of performant SP molecules is selected. Since addressability and thermostability can be conflicting criteria, we cannot define a simple ranking of all candidates. Instead, we selected molecules according to the Pareto front of addressability and thermostability via a non-dominated sorting algorithm. To this end, we define the first three layers of the Pareto front as the set of performant SP candidates, in order to strike a balance between performance and chemical diversity.

In the last step of the TL workflow, the REINVENT model is fine-tuned using these performant SP candidates as a new training set. The fine-tuned model was then used to generate improved additional photoswitch candidates. Note that in addition to providing performant SP candidates for retraining, the first set of generated systems also provides important insights on which decoration patterns are favorable for our design criteria, by considering how frequently each pattern occurs in the performant set. Indeed, this pattern effect is substantial, so that we also used a correspondingly weighted sample of generated molecules to evaluate from the fine-tuned set. Nonetheless, at least five candidates were included for each decoration pattern, to ensure a diverse sample of candidates. Additional ablation studies on different steps of the workflow can be found in the SI. This confirms that both fine-tuning and reweighting promising decoration patterns is important to achieve optimal results.

Results & discussion 3.

3.1 Evaluation of generated molecules

In order to evaluate the molecular distributions generated by the baseline and fine-tuned models, we computed the performance descriptors for 17.3 k and 15.9 k unique molecules based on three independent TL runs, respectively. Here, the different total number of unique molecules is due to a higher number of duplicates generated by the fine-tuned model. This is consistent with a more focused exploration of chemical space.

In addition, a set of literature-known doubly substituted SP molecules was compiled via a SP substructure search on Sci-Finder and PubChem, resulting in 976 doubly decorated SP molecules (see Computational Details). Fig. 5 provides an overview of the addressabilities and relative bond lengths in both sets. Focusing on the extreme values, the central panel shows the convex hulls of the different distributions, including the respective Pareto fronts towards the top right. Additionally, marginal distribution plots for the addressability and bond lengths are included as separate panels on the top and right.




Fig. 5 Convex hull plot of the addressability (based on the spectral overlap of open and closed forms) and relative bond length (as a proxy for thermostability of the open form) for different sets of molecules. Green: literature known spiropyran molecules. Yellow: candidates generated by the general baseline model. Blue: candidates generated by the fine-tuned spiropyran model.

Overall, the molecules of the literature set show a bimodal distribution for the relative bond length criterion with maxima around 0% and 1%, with a somewhat higher density at the shorter bond length peak. In contrast, the generative datasets display unimodal distributions peaking around 1%. Interestingly, the fine-tuned model has an additional shoulder and tail towards even longer bond lengths, indicating that the finetuning indeed works as expected. For the addressability, all datasets are unimodal and peaked at the lower end of the range. Again, the fine-tuned model displays a pronounced tail towards higher addressability and significantly larger mean values (6.5 vs. 2.4 for the baseline). In terms of the Pareto fronts, it can be seen that the fine-tuned model dominates the baseline, which in turn dominates the literature data. Since the Pareto fronts correspond to the most extreme values in the sample (i.e. they are outliers by definition), the precise shape of the front somewhat varies between different runs of the workflow. Nonetheless, the trends described above are robust across multiple independent runs.

The bimodal distribution of the literature data in terms of bond lengths to some extent reflects the biases in traditional "trial and error" molecular design. Many of the reference molecules are decorated in *para*-position to the oxygen, and often feature a nitro group. This is due to the large number of 6-nitro-BIPS derivatives that have been reported in the literature. Unsurprisingly, this leads to a bond length distribution peaked around the 6-nitro-BIPS reference. Another factor that likely influences the trend towards shorter bond lengths in the literature distribution is the fact that long thermal backreaction times are actually desired in some applications, such as energy storage or drug delivery.

In contrast, the distribution generated by the baseline model has no bias towards 6-nitro-BIPS or other previously known photoswitches. This shows that SP molecules naturally tend towards longer bond lengths across chemical space. Assuming that the inverse relation between bond length and thermostability is robust, this indicates that a large range of fast-switching SP molecules is unexplored. On the other hand, the addressability distribution between the literature and baseline datasets are virtually identical. Literature bias thus does not seem to affect this property significantly.

Comparing the fine-tuned and baseline distributions, we find that both properties are significantly influenced by the retraining step, vindicating our TL approach. Here it appears that the addressability is a more elastic property, though this is likely due to its definition as a quotient of oscillator strengths, which can in principle be arbitrarily large. In contrast, bond lengths are naturally constrained by chemistry, and variations beyond a few percentage points would be highly unusual.

3.2 Decoration patterns

As mentioned above, a key feature of the TL workflow is that the fine-tuned model preferentially samples decoration patterns which are more likely to produce performant molecules. A comparison between the most common decoration sites in the fine-tuned and literature datasets is shown in Fig. 6. This reveals significant differences between both sets. As mentioned above, literature SPs show a strong preference for the site *para* to oxygen (index 9 in Fig. 6), with 45% of all reported structures being decorated at this position. Apart from this, only




Fig. 6 (a) Most favoured decoration positions for doubly decorated spiropyrans reported in the literature (b) Most favoured decoration positions for spiropyrans generated by the fine-tuned model. Red numbers are the indices of the position used in the main text.

decorations at the nitrogen (index 1), para to the nitrogen (index 4) and ortho to the oxygen (index 7) are observed. In contrast, the fine-tuned model prefers the position para to the nitrogen (index 4) with about 25% of all sampled structures being decorated here. The second most likely decoration sites are meta position to the oxygen (indices 8 & 10) and ortho to the nitrogen (index 2).

In general, the fine-tuned model thus samples a much wider variety of decoration sites, providing the possibility for novel discoveries. It is also interesting to note that in some cases the patterns are opposed to what is found in the literature: for example, decorations meta to the oxygen are not found in the literature at all but are among the most common for the finetuned model. These trends can again be related to biases in the literature and opposing design targets (e.g. long vs. short backreaction times). While a decoration of an electron withdrawing group ortho or para to the oxygen can help to stabilize the open form, an electron withdrawing decoration meta to the oxygen (and therefore para to the rotatable double bond) can weaken the double bond character via mesomeric effects.

In terms of biases, the influence of 6-nitro-BIPS derivatives on the literature dataset is clearly evident in Fig. 6. However, this bias should not simply be interpreted as a lack of creativity on the side of synthetic chemists. Instead, the synthetic accessibility of the site para to the oxygen is much higher, when SPs are synthesized by condensation of a substituted salicylaldehyde with a substituted indolin. The generative models used herein do not explicitly account for synthesizability, though concepts of chemical stability should implicitly be learned from the ChEMBL training data. Indeed, a recent review by Bilodeau et al. highlights the role of synthesizability for molecular generative models.47 Westermayr et al. showed that the structure-based synthetic complexity score (SCS) of Coley et al. could be used to bias generative ML models towards synthetically accessible systems. 45,48 However, we found that


this approach cannot easily be transferred to the current scaffold decoration setting, because the SCS distribution for decorated SP scaffolds is highly static (see Fig. S6). Furthermore, overemphasizing synthesizability can also unduly constrain the novelty of the generated molecules, in particular when synthesizability (an inherently subjective concept) is quantified based on literature-known molecules. At this point, a critical assessment by a trained synthetic chemist is therefore still the best option.

3.3 Pareto optimal molecules

So far, our discussion has focused on general properties and trends of the generated datasets. We now turn to specific molecules, namely those on the Pareto fronts of the baseline and fine-tuned datasets, respectively (see Fig. 7). These molecules represent extremes of the generated distributions and either optimize one criterion at the expense of the other (the edges of the front) or represent an optimal trade-off between both (the center of the front). The eighteen depicted molecules are a representative sample covering the full property range. All Pareto optimal molecules are shown in the SI.

In Fig. 7, the molecules in each box are arranged from highest to lowest addressability from top left to bottom right (and reverse for the relative bond length). Some general structural trends are similar between both datasets. For structures with high addressability (first row), conjugated chains attached meta to the SP oxygen (index 8 in Fig. 6) are prominent. For the fine-tuned model these decorations additionally feature heterocycles. These modifications are physically plausible, as they would be expected to lead to a red-shift of the absorption bands, in particular for the MC form. A second common decoration for highly addressable systems is a methoxy group ortho to the oxygen, which may have additional mesomeric effects.

To achieve long relative bond lengths, the fine-tuned model consistently decorates the position next to the double bond of






Fig. 7 Molecules of the pareto front with numerical values for the relative bond length on the top and addressability on the bottom. Left: results from the baseline model. Right: results from the fine-tuned model.

the central pyran ring (index 10 in Fig. 6) with bulky substituents. This likely has steric effects on the conformational ensemble of the MC form, favoring longer bonds at this position.

Whether this has the desired effect of lower thermostability of the MC form is unclear, however, since such substituents could potentially also lead to larger barriers for the backreaction or prevent switching altogether. This highlights a general risk in proxy-based optimization, namely that the descriptor can be optimized in a way that is detrimental to the intended goal. On the other hand, these molecules also consistently feature nitrogen-based substituents, which could feature beneficial electronic effects by extending the conjugated π -system, in addition to the electronic withdrawing effects of the NO2 moiety. Additionally, some functional groups predicted by the model can be highly sensitive to their chemical environment and thus solvent polarity or pH can have drastic influence on their photophysical properties. Thus, additional precautions must be taken when evaluating the predicted molecules. Comparing the baseline and fine-tuned datasets, it is also interesting to note that the former more frequently contains functional groups that are typical in medicinal chemistry, such as cyclopropane or trifluoromethyl groups, which are much less prevalent in the fine-tuned set. This indicates that the finetuning moves the model away from the medicinal chemistry focused ChEMBL training set.

In order to explore the biases inherent in the ChEMBL training set, we explored the PubChem dataset as an alternative (see Fig. S3). This reveals that the training set of the generative model influences both the nature and the complexity of the decorations. For example, we observed a trend towards more hydrocarbon-based substituents in PubChem, while the ChEMBL decorations tend to contain more halogens and aryl substituents. Furthermore, the PubChem model yields overall

more complex decorations. This reveals that the underlying training data has significant implications on the chemical space accessible by the generative model. The increased complexity of the PubChem decorations, while potentially enhancing the performance of the resulting photoswitches, also presents challenges related to reduced synthetic accessibility.

3.4 Experimental validation

The ultimate test for a molecular design framework is the experimental validation of the predicted molecules. To this end a series of molecules designed by both the baseline and finetuned generative models were synthesized (see Fig. 8). We aimed to explore the effects of certain functional groups and decoration patterns that were commonly found in molecules predicted to be performant. The molecules were selected according to chemical diversity and synthetic accessibility. The manual selection process involved retrosynthetic analysis of the potential candidates and the commercial availability of the precursors. There were more than six synthetic sequences attempted, however, some of the candidate molecules could not be obtained as their synthesis was either unsuccessful or they could not be obtained in a purity sufficient for spectroscopic analysis. For these reasons, the synthesized molecules do not correspond to the Pareto optimal molecules designed by the generative model but rather reflect the general design principles which can be derived from the generative model. Nonetheless, all candidates are expected to have above average addressabilities and/or relative bond lengths.

Clearly, this experimental dataset is inherently limited by the challenge of synthesizing and isolating certain model-predicted molecules. A significant portion of the models' predictions involves molecules that, while chemically intriguing, are challenging to synthesize and isolate due to underlying thermal




Fig. 8 Scatterplot of descriptors of all photoswitches generated by the fine-tuned model. Experimentally synthesized molecules are shown in red and highlighted in the margins.

reactivity, not directly related to their photoswitching properties. To address this, candidate molecules were subjected to retrosynthetic analysis to identify synthetically approachable structures. Note, that even the elusive molecules remain valuable for advancing theoretical studies as they provide critical insights into the photophysical, structural, and thermochemical factors governing spiropyran switching behavior.

Among the synthesized compounds were the following derivatives. Spiropyran 1 featuring electron withdrawing nitro group on the indoline part of the molecule and methoxy group was found to have red-shifted absorption maxima of the open merocyanine form tailing to 700 nm and significantly lower activation barrier (9.85 \pm 2.99 kJ mol⁻¹) for thermal ring-closure as compared to the model derivative. Compound 2 with two nitro groups was measured to have an activation barrier of 63.84 \pm 1.22 kJ mol⁻¹. Compound 3 was a switch with significantly different absorption of the closed form compared to the model compound and the previously measured ones with its absorption maxima of only 310 nm. The activation barrier for ringclosure for this compound was measured as 39.94 \pm 3.44 kJ mol⁻¹. The spiropyran 4 was observed to undergo slight opening upon irradiation with 313 nm at -35 °C, however, due to the freezing point of acetonitrile, the exact activation barrier could not be determined. For compound 5, an activation barrier of 26.52 \pm 1.27 kJ mol⁻¹ was measured. This derivative shows the difference of alkoxy group influence compared to its regioisomer 1. Compound 6 was observed to have blue shifted lowest absorption maxima of merocyanine form and its activation barrier was found to be 42.62 ± 1.82 kJ mol⁻¹.

For reference, the activation barrier of 6-nitro-BIPS was measured as 101.77 ± 0.45 kJ mol $^{-1}$. A full comparison of all measured candidates can be found in the SI (Fig. S8). This also allows a post-hoc validation of the bond-length criterion for thermostability. We find a moderate negative correlation (with $R^2 = 0.67$) between the experimental activation barriers and the relative bond length. Though not perfect, this is certainly a useful accuracy for guiding molecular design, especially given the heuristic nature of the descriptor.

4. Conclusion

In this contribution, we have presented a generative TL framework for the *de novo* design of SP photoswitches. We find that the presented approach effectively pushes the Pareto front towards systems with the desired performance descriptors and avoids biases in the current literature. In particular, we find that the range of possible SP decoration patterns is underexplored and suggest promising sites for functionalization, as well as novel side group chemistries. This approach could readily be extended to other photoswitch classes, since the scaffold-decoration architecture used herein allows maintaining arbitrary photoactive backbones.

Beyond statistical analyses of the generated chemical spaces, we have also critically examined the individual structures designed by the generative models. In agreement with previous reports, we find that the synthesizability of some of the candidates is likely low. Nonetheless, consistent design rules

regarding the nature of the functional groups and the decoration patterns can be derived from the proposed molecules. These have been validated by synthesizing six molecules. All of the newly synthesized compounds are predicted to display shorter thermal half-lives of their merocyanine isomers when compared to the baseline 6-nitro-BIPS molecule. This was confirmed by spectroscopic measurements, which show that the observed activation barriers are significantly lower than for 6-nitro-BIPS in all cases. However, even though the qualitative prediction is correct, the criterion of relative bond length was found to not correlate perfectly with the measured barriers.

In future work, we aim to improve the generative design of photoswitches by developing more reliable performance descriptors and exploring 'human-in-the-loop' approaches. MLaccelerated quantum chemical calculations represent an intriguing route for predicting more accurate thermal backreaction barriers. Indeed, Axelrod et al. have shown that equivariant neural networks can be trained with active learning to predict diabatic ML potential energy surfaces for azobenzene photoswitches. 33 However, even for a somewhat smaller space of 3100 molecules (compared to the space considered herein), this required 570 000 training geometries, which would need to be added to the current cost of transfer learning and evaluation (13 000 DFT calculations in one run of the current workflow). Moreover, it should be noted that the switching of azobenzenes is less complex than that of spiropyrans, due to the presence of multiple conformers in the open merocyanine form of the latter. Therefore, to obtain an accurate diabatic ML potential for a broad chemical space of spiropyrans, even larger computational effort would be required. This is out of the scope of the current study yet certainly worth pursuing in future work. Beyond this, the establishment of dedicated experimental libraries for further benchmarking of computational models can further increase the quality of the descriptors and synthesizability should be considered in a more rigorous manner.

5. Computational & experimental details

5.1 Training data collection and curation

The training data for the REINVENT scaffold-decoration model is based on a subset of the ChEMBL dataset, tailored for the spiropyran core structure. The initial training set from the original REINVENT paper contains a filtered version of the ChEMBL25 database.39 The filtering steps include standardization, structural and QED filters as well as token filters. This cuts the initial ChEMBL set of 1.8 M structures down to 827 k structures. To further tailor the dataset for the spiropyran design task additional filtering was applied based on substructure matches to benzene or pyrrolidine with rdKit⁴⁹ leaving 783 k molecules. To generate scaffold decoration pairs, a slicing scheme was used to cut along every aliphatic bond with the constraints that each scaffold must have at least one ring and the maximum molecular weight of a decoration is below 300 g mol⁻¹. This slicing resulted in 36.8 M scaffold-decoration pairs, which were used to train the REINVENT RNN model.

This article is licensed under a Creative Commons Attribution 3.0 Unported Licence.

ben Access Article. Published on 16 September 2025. Downloaded on 11/21/2025 10:56:25 PM.

Digital Discovery

Paper

5.2 Training of the scaffold decorator model

The dataset from 5.1 was used to train the RE-INVENT scaffold-decorator model. The model was trained over 50 epochs with a batch size of 1600. The exponential learning rate decay started at 10^{-4} and stopped at 10^{-6} with a learning rate gamma of 0.95. This training procedure was used for the initial training of the ChEMBL model as well as the retraining within the TL workflow.

5.3 Query of the scaffold-decorator model

Since the RNN model has a probabilistic character, one must query each scaffold many times to obtain meaningful statistics. Herein, two strategies are used to sample diverse decorations: on one hand, the scaffold SMILES are randomized 36 times, leading to distinct but equivalent SMILES representations, which are processed differently by the RNN. On the other hand, each of these is queried 36 times. For the doubly decorated molecules discussed herein, this results in 36⁴ sampling attempts in total. The actual number of generated molecules is drastically lower, however, since many duplicate molecules are generated.

5.4 Semiempirical calculations with GFN2-xTB

To compute the performance descriptors, 3D coordinates for the candidate molecules were first embedded using the DGETK method as implemented in rdKit.⁵⁰ The resulting structures were prerelaxed with the MMFF and further refined using the GFN2-xTB method³⁴ with the ALPB implicit solvation model and acetonitrile as the solvent. This procedure was performed for both the MC as well as SP form of the candidate molecules for 10 generated conformers each.

5.5 DFT calculations

The GFN2-xTB optimized geometries were subsequently used in single-point TD-DFT calculations to obtain the excited states with ORCA 5.0.3.⁵¹ All calculation were carried out using the PBE0 functional⁵² with the def2-TZVP⁵³ basis set.

5.6 Spiropyran literature dataset

Literature known SP molecules were obtained *via* a substructure query with *SciFinder* and *PubChem*. In total 5.2 k unique SP molecules were found for both datasets combined. For a fair comparison further filtering steps were applied in rdKit: 2.1 k (approx. 41%) of the found SP molecules were doubly decorated. Some of the doubly decorated molecules feature condensed ring systems attached to the photochromic core and therefore change its general structure. Since only the "primitive" SP core was decorated in this work, systems with condensed rings as decorations were also excluded. Some of the literature known structures had decorations at the rotatable double bond in the MC form which was not decorated in the workflow and therefore also discarded. The result of this filtering was 976 SP molecules.

5.7 Synthetic procedure for the candidates

Spiropyranes presented in this work were synthesized by condensation of an appropriate indolinium salt (1 equiv.) with the corresponding derivative of salicylaldehyde (2 equiv.) using piperidine (\sim 1.1 equiv.) as a base. The reaction is performed in ethanol or DMF based on the specific derivative. Upon the completion of the reaction, the crude material was purified by crystallization from hot ethanol or by column chromatography in order to obtain samples of sufficient quality for spectroscopic experiments. Further details are available in the SI.

5.8 Spectroscopic analysis of the candidates

The spectroscopic measurements were performed using Agilent Cary 60 instrument connected to a cryostat. A short-arc mercury xenon lamp together with monochromator is used as a light source for the measurements with concurrent irradiation. All samples were prepared using spectroscopy grade acetonitrile and compounds were weight using Sartorius ME5 analytical microbalance. Further details are available in the SI.

Conflicts of interest

There are no conflicts of interest to declare.

Data availability

The TL workflow scripts, the trained models, the results for the three TL runs, and all datasets used in this work can be found at https://doi.org/10.5281/zenodo.17099758.

Supplementary information is available. See DOI: https://doi.org/10.1039/d5dd00327j.

Acknowledgements

This project was supported by the priority program SPP2363 "Molecular Machine Learning" of the German Science Foundation (DFG). RS was supported by a Kekulé-Scholarship of the FCI. The computational resources provided by the Max Planck Computing and Data Facility (MPCDF) are gratefully acknowledged. S.H. thanks the Einstein Foundation Berlin as well as Humboldt University for generous support. Open Access funding provided by the Max Planck Society.

References

- 1 K. Nakatani, J. Piard, P. Yu and R. Métivier, Introduction: Organic Photochromic Molecules, *Phot. Mat.*, 2016, 1–45.
- 2 O. Bozovic, B. Jankovic and P. Hamm, Using azobenzene photocontrol to set proteins in motion, *Nat. Rev. Chem.*, 2022, **6**, 112–124.
- 3 S. Crespi, N. A. Simeth and B. König, Heteroaryl azo dyes as molecular photoswitches, *Nat. Rev. Chem.*, 2019, 3, 133–146.
- 4 M. J. Fuchter, On the Promise of Photopharmacology Using Photoswitches: A Medicinal Chemist's Perspective, *J. Med. Chem.*, 2020, **63**, 11436–11447.

- 5 R. Klajn, Spiropyran-based dynamic materials, Chem. Soc. Rev., 2014, 43, 148-184.
- 6 L. Kortekaas and W. R. Browne, The evolution of spiropyran: fundamentals and progress of an extraordinarily versatile photochrome, Chem. Soc. Rev., 2019, 48, 3406-3424.
- 7 Y. Hattori, et al., Cyclization from Higher Excited States of Diarylethenes Having a Substituted Azulene Ring, Chem, 2020, 26, 11441-11450.
- 8 N. F. Konig, D. Mutruc and S. Hecht, Accelerated Discovery of alpha-Cyanodiarylethene Photoswitches, J. Am. Chem. Soc., 2021, 143, 9162-9168.
- 9 Z. L. Pianowski, Recent Implementations of Molecular Photo-switches into Smart Materials and Biological Systems, Chem, 2019, 25, 5128-5144.
- 10 D. Majee and S. Presolski, Dithienvlethene-Based Photoswitchable Catalysts: State of the Art and Future Perspective, ACS Catal., 2021, 11, 2244-2252.
- 11 J. Volaric, W. Szymanski, N. A. Simeth and B. L. Feringa, Molecular photoswitches in aqueous environments, Chem. Soc. Rev., 2021, 50, 12377-12449.
- 12 Y. Ru, et al., Recent progress of photochromic materials towards photocontrollable devices, Mater. Chem. Front., 2021, 5, 7737-7758.
- 13 S. Wiedbrauk, H. McKinnon, L. Swann and N. R. B. Boase, Diarylethene Photoswitches and 3D Printing to Fabricate Rewearable Colorimetric UV Sensors for Sun Protection, Adv. Mater. Technol., 2023, 8, 2201918.
- 14 S. Jia, W.-K. Fong, B. Graham and B. J. Boyd, Photoswitchable Molecules in Long-Wavelength Light-Responsive Drug Delivery: From Molecular Design to Applications, Chem. Mater., 2018, 30, 2873-2887.
- 15 P. Kobauri, F. J. Dekker, W. Szymanski and B. L. Feringa, Rational Design in Photopharmacology with Molecular Photoswitches, Angew Chem. Int. Ed. Engl., 2023, 62, e202300681.
- 16 D. Dattler, et al., Design of Collective Motions from Synthetic Molecular Switches, Rotors, and Motors, Chem. Rev., 2020, **120**, 310-433.
- 17 S. Kassem, et al., Artificial molecular motors, Chem. Soc. Rev., 2017, 46, 2592-2621.
- 18 S. Son, E. Shin and B. S. Kim, Light-responsive micelles of spiropyran initiated hyperbranched polyglycerol for smart drug delivery, Biomacromolecules, 2014, 15, 628-634.
- 19 M. Regehly, et al, Xolography for linear volumetric 3D printing, Nature, 2020, 588, 620-624.
- 20 G. Schneider and P. Wrede, Artificial neural networks for computer-based molecular design, Prog. Biophys. Mol, 1998, 70, 175-222.
- 21 J. Peng, et al., Human- and machine-centred designs of materials for molecules and sustainability decarbonization, Nat. Rev. Mater., 2022, 7, 991-1009.
- 22 B. Sanchez-Lengeling and A. Aspuru-Guzik, Inverse molecular design using machine learning: Generative models for matter engineering, Science, 2018, 361, 360-365.
- 23 S. Axelrod, D. Schwalbe-Koda, S. Mohapatra, J. Damewood, K. P. Greenman and R. Gómez-Bombarelli, Learning

- Matter: Materials Design with Machine Learning and Atomistic Simulations, Acc. Mater. Res., 2022, 3, 343-357.
- 24 C. Kunkel, J. T. Margraf, K. Chen, H. Oberhofer and K. Reuter, Active discovery of organic semiconductors, Nat. Commun., 2021, 12, 2422.
- 25 H. Türk, E. Landini, C. Kunkel, J. T. Margraf and K. Reuter, Assessing Deep Generative Models in Composition Space, Chem. Mater., 2022, 34, 9455-9467.
- 26 K. Chen, C. Kunkel, K. Reuter and J. T. Margraf, Reorganization energies of flexible organic molecules as a challenging target for machine learning enhanced virtual screening, Digital Discovery, 2022, 1, 147-157.
- 27 A. Gupta, et al., Generative Recurrent Networks for De Novo Drug Design, Mol. Inf., 2018, 37, 1700111.
- 28 R. Gómez-Bombarelli, et al., Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules, ACS Cent. Sci., 2018, 4, 268-276.
- 29 N. Gebauer, M. Gastegger and K. Schütt, Symmetry-adapted generation of 3d point sets for the targeted discovery of molecules, Adv. Neural Inf. Process. Syst., 2019, 32, 8024-8035.
- 30 G. L. Guimaraes, B. Sanchez-Lengeling, C. Outeiral, P. L. C. Farias and A. Aspuru-Guzik, Objective-reinforced generative adversaryial networks (organ) for sequence generation models, arXiv, 2017, preprint, arXiv:1705.10843, DOI: 10.48550/arXiv.1705.10843.
- 31 F. Mukadum, Q. Nguyen, D. M. Adrion, G. Appleby, R. Chen, H. Dang, R. Chang, R. Garnett and S. A. Lopez, Data-driven discovery o molecular photoswitches with multioutput Gaussian processes, J. Chem. Inf. Model., 2021, 61, 5524-5534.
- 32 S. Vela, A. Fabrizio, K. R. Briling and C. Corminboeuf, Learning the Excitation Properties of Azo-dyes, J. Phys. Chem. Lett., 2021, 12, 5957-5962.
- 33 S. Axelrod, E. Shakhnovich and R. Gómez-Bombarelli, state non-adiabatic dynamics Excited of photoswitchable molecules using a chemically transferable machine learning potential, Nat. Commun., 2022, 13, 3440.
- 34 C. Bannwarth, S. Ehlert and S. Grimme, GFN2-xTB An Accurate and Broadly Parametrized Self-Consistent Tight-Binding Quantum Chemical Method with Multipole Electrostatics and Density-Dependent Dispersion Contributions, J. Chem. Theory Comput., 2019, 15, 1652-1671.
- 35 Y. Sheng, et al., Comprehensive Theoretical Study of the Conversion Reactions of Spiropyrans: Substituent and Solvent Effects, J. Phys. Chem. B, 2004, 108, 16233-16243.
- 36 C. F. Koelsch and W. R. Workman, Some Thermochromic Spirans, JACS Au, 1952, 74, 6288-6289.
- 37 J. Arús-Pous, et al., SMILES-based deep generative scaffold decorator for de-novo drug design, J. Cheminf., 2020, 12, 1-18.
- 38 N. Brown, M. Fiscato, M. H. S. Segler and A. C. Vaucher, GuacaMol: Benchmarking Models for de Novo Molecular Design, J. Chem. Inf. Model., 2019, 59, 1096-1108.

- 39 B. Zdrazil, *et al.*, The ChEMBL Database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods, *Nucleic Acids Res.*, 2024, 52, D1180–D1192.
- 40 M. H. S. Segler, T. Kogej, C. Tyrchan and M. P. Waller, Generating Focused Molecule Libraries for Drug Discovery with Recurrent Neural Networks, *ACS Cent. Sci.*, 2018, 4, 120–131.
- 41 M. Moret, L. Friedrich, F. Grisoni, D. Merk and G. Schneider, Generative molecular design in low data regimes, *Nat. Mach. Intell.*, 2020, 2, 171–180.
- 42 M. Moret, M. Helmstadter, F. Grisoni, G. Schneider and D. Merk, Beam Search for Automated Design and Scoring of Novel ROR Ligands with Machine Intelligence, *Angew Chem. Int. Ed. Engl.*, 2021, **60**, 19477–19482.
- 43 S. Singh and R. B. Sunoj, A transfer learning protocol for chemical catalysis using a recurrent neural network adapted from natural language processing, *Digital Discovery*, 2022, **1**, 303–312.
- 44 K. Sattari, D. Li, B. Kalita, Y. Xie, F. B. Lighvan, O. Isayev and J. Lin, De novo molecule design towards biased properties via a deep generative framework and iterative transfer learning, *Digital Discovery*, 2024, 3, 410–421.
- 45 J. Westermayr, J. Gilkes, R. Barrett and R. J. Maurer, High throughput property-driven generative design of functional organic molecules, *Nat. Comput. Sci.*, 2023, 3, 139–148.
- 46 N. W. A. Gebauer, M. Gastegger, S. S. P. Hessmann, K. R. Muller and K. T. Schutt, Inverse design of 3d

- molecular structures with conditional generative neural networks, *Nat. Commun.*, 2022, **13**, 973.
- 47 C. Bilodeau, W. Jin, T. Jaakkola, R. Barzilay and K. F. Jensen, Generative models for molecular discovery: Recent advances and challenges, *Wiley Interdiscip. Rev.:Comput. Mol. Sci.*, 2022, **12**, e1608.
- 48 C. W. Coley, L. Rogers, W. H. Green and K. F. Jensen, SCScore: Synthetic Complexity Learned from a Reaction Corpus, J. Chem. Inf. Model., 2018, 58, 252–261.
- 49 RDKit: Open-source cheminformatics, https://www.rdkit.org.
- 50 D. C. Spellmeyer, A. K. Wong, M. J. Bower and J. M. Blaney, Conformational analysis using distance geometry methods, *J. Mol. Graph. Model.*, 1997, **15**, 18–36.
- 51 F. Neese, F. Wennmohs, U. Becker and C. Riplinger, The ORCA quantum chemistry program package, *J. Chem. Phys.*, 2020, **152**, 224108.
- 52 J. P. Perdew, M. Ernzerhof and K. Burke, Rationale for mixing exact exchange with density functional approximations, *J. Chem. Phys.*, 1996, **105**, 9982–9985.
- 53 A. Hellweg and D. Rappoport, Development of new auxiliary basis functions of the Karlsruhe segmented contracted basis sets including diffuse basis functions (def2-SVPD, def2-TZVPPD, and def2-QVPPD) for RI-MP2 and RI-CC calculations, *Phys. Chem. Chem. Phys.*, 2015, 17, 1010–1017.