

Cite this: *Digital Discovery*, 2025, 4, 3623

Constructing and explaining machine learning models for the exploration and design of boron-based Lewis acids

Juliette Fenogli, * Laurence Grimaud * and Rodolphe Vuilleumier *

The integration of machine learning (ML) into chemistry offers transformative potential in the design of molecules with targeted properties. However, the focus has often been on creating highly efficient predictive models, sometimes at the expense of interpretability. In this study, we leverage explainable AI techniques to explore the rational design of boron-based Lewis acids, that activate a wide range of organic reactions. Using fluoride ion affinity as a proxy for Lewis acidity, we developed interpretable ML models based on chemically meaningful descriptors, including *ab initio* computed features and substituent-based parameters derived from the Hammett linear free-energy relationship. By constraining the chemical space to well-defined molecular scaffolds, we achieved highly accurate predictions (mean absolute error <math><6 \text{ kJ mol}^{-1}</math>), surpassing conventional black-box deep learning models in low-data regimes. Interpretability analyses of the models shed light on the origin of Lewis acidity in these compounds and identified actionable levers to modulate it through the nature and positioning of substituents on the molecular scaffold. This work bridges ML and chemist's way of thinking, demonstrating how explainable models can inspire molecular design and enhance scientific understanding of chemical reactivity.

Received 19th May 2025
Accepted 18th September 2025

DOI: 10.1039/d5dd00212e

rsc.li/digitaldiscovery

Introduction

Machine learning (ML) has become indispensable across various scientific domains for uncovering patterns in databases and modeling complex relationships among variables in high-dimensional spaces.^{1–3} In chemistry, ML has made significant strides in recent years, leveraging experimental and computed data made accessible by the development of extensive databases,^{4,5} high-throughput experiments,^{6,7} and supercomputers.^{8,9} While data-driven modeling and statistical analysis have a long-standing history in chemistry, building upon foundational concepts like the Hammett correlations in 1937,¹⁰ the adoption of ML techniques marks a notable advancement.¹¹ Indeed, chemical questions ranging from drug discovery,¹² molecular dynamics (MD) simulations,^{13,14} and chemical reaction outcome prediction^{5,7,8,15–17} to synthesis planning^{18–20} have been successfully addressed.

However, many models, particularly those based on deep neural networks, lack interpretability and are often referred to as “black-box” models,²¹ which poses challenges in trust-critical fields like medical diagnosis.²² Interpretability not only aids in rationalizing decision-making processes but also helps evaluate the validity of a model's predictions against domain knowledge,

guarding against spurious reasoning.²³ Furthermore, while some models can be highly accurate, their lack of interpretability limits the ability to extract scientific knowledge from them, thus diminishing their overall interest. The emergence of explainable artificial intelligence (XAI) seeks to address these issues by elucidating what ML algorithms have learned, fostering scientific knowledge, and inspiring new concepts and ideas.^{24–26} In chemistry, XAI techniques have been successfully applied to the study of quantitative structure–activity relationships (QSAR),^{27–30} drug discovery,³¹ the modeling of organic reactions,^{32,33} and to understanding what physical concepts models have learned when predicting interatomic potentials for MD simulations.^{34,35}

Here we investigate how ML can guide the design of molecules with a targeted specific property: the Lewis acidity, which is of high importance in a wide range of organic transformations.^{36,37} Strategies to boost the Lewis acidity of borane derivatives include introducing electron-withdrawing groups onto ligands^{38–41} or using constrained geometry ligands.^{42–44} While functional group decoration is common in QSAR organic chemistry studies, applying ML techniques allows for the exploration of a much broader chemical space, providing novel insights for molecule design. Therefore, predicting the Lewis acidity of compounds is highly valuable and has recently been explored by the Greb group,⁴⁵ who developed an advanced graph neural network (GNN) model trained on a dataset of nearly 49k Lewis acids with diverse central atoms and ligands.

CPCV, Département de Chimie École Normale Supérieure, PSL University, Sorbonne Université, CNRS, Paris 75005, France. E-mail: juliette.fenogli@ens.psl.eu; laurence.grimaud@ens.psl.eu; rodolphe.vuilleumier@ens.psl.eu



Complementing this large-scale study, we aim at exploring Lewis acidity using XAI, focusing on restricted chemical spaces, employing “white-box” self-explanatory models and post-hoc interpretability techniques. Indeed, interpreting ML models can enhance chemical understanding by acting as a “computational microscope”,^{24,25} to explore the qualitative concept of Lewis acidity and as a “source of inspiration”^{24,25} for designing novel boron compounds with targeted Lewis acidity.

In this context, we constructed and analyzed arrays of Lewis acids, annotated with their fluoride ion affinity (FIA), a quantitative measure of Lewis acidity (Fig. 1). By combining multiple molecular descriptors with a range of ML algorithms, we developed regression models to predict FIA and applied XAI techniques to decipher the intrinsic Lewis acidity of these compounds and identify actionable levers for molecular design.

Results and discussion

Lewis acidity scale

To construct a database of boron Lewis acids, a robust metric for quantifying Lewis acidity was essential. The chosen metric needed to be readily accessible to facilitate the fast labeling of molecules, while also ensuring consistency, as noisy datasets can hinder the learning process and potentially lead to over-fitting.⁴⁶ Several metrics were evaluated and benchmarked with conventional scales employed in experimental organic chemistry to determine Lewis acidity (see the SI, Section S1). FIA, representing the standard negative enthalpy change associated with the binding of a fluoride ion to the Lewis acid in the gas phase (Fig. 1), was retained as the most relevant, consistent and accessible quantity, as it is a computed metric. Various cost-




Fig. 1 Workflow towards the design of efficient boron Lewis acids: spanning from defining chemical space (four molecular scaffolds and substituents) and molecule features (chemo-informatics, *ab initio* computed or derived from the Hammett correlations) to labeling Lewis acidity by FIA and constructing predictive models. Models' predictions were then explained to get insight into the qualitative concept of Lewis acidity and design boron Lewis acids with specific Lewis acidity profiles.



effective *ab initio* methods were evaluated to compute FIA, ensuring alignment with values computed at a higher level of theory⁴⁷ (SI, Section S2). We used density functional theory (DFT) at the M06-2X/6-31G(d) level of theory in isodesmic calculations as a compromise between efficiency and precision to provide reliable FIA data. Throughout this manuscript, FIA has been used as a proxy for Lewis acidity. Therefore, any conclusions drawn about FIA also apply to Lewis acidity, and the terms have been used equivalently.

Chemical space

We targeted scaffolds possessing aromatic ligands, which are particularly effective at transmitting electronic effects over long distances through their electronic π systems, unlike aliphatic or non-conjugated ligands. This facilitates the combination of effects from multiple substituents on the same central atom. Additionally, as aromatic positions are not chemically equivalent, their diverse substitution allows to precisely modulate the effect on FIA. While the substitution of aromatic positions to modulate molecular reactivity is well-established in organic chemistry,³⁹ its systematic integration with ML to predict Lewis acidity is, to our knowledge, unprecedented. Previous databases, such as the one developed by Greb and coworkers,⁴⁵ have primarily focused on exploring the effects of the central atom type and ligand denticity. To enhance model performance and simplify interpretation, we limited the chemical space to four scaffolds: triarylboranes, which have been extensively studied for their catalytic activity,^{36,37} and three constrained Lewis acid scaffolds (ONO,⁴⁸ NNN,⁴³ OCO⁴⁹) featuring heteroatoms coordinated to boron. The planar geometry of these constrained ligands minimizes steric hindrance, while the pincer-like structure facilitates boron atom accessibility, enhancing Lewis acidity.⁴³ For each scaffold, the chemical space consists

of symmetric molecules, substituted with 13 possible electron-donating or -withdrawing substituents at specific aromatic positions, to minimize steric hindrance and ensure synthetic efficiency (Fig. 1). These constraints result in an accessible chemical space of 2197 molecules for the ONO scaffold for example (Fig. 2A). For each dataset, molecules intended for FIA computation were generated randomly by selecting one fragment from the set of 13 options at each substitutable position. However, we initially intended to study the effect of only one or two substituents on FIA. To that end, the random substitution pool was intentionally biased to generate more hydrogen fragments, which resulted in an overrepresentation of molecules featuring two non-hydrogen substituents (clusters in Fig. 2A, ONO dataset). To enhance diversity and better represent the chemical space for model development, the ONO dataset was expanded from 175 to 272 molecules. This expansion used k-means⁵⁰ and coverage⁵¹ algorithms applied to molecular fingerprint representations, to uniformly sample the ONO chemical space (Fig. 2A, additional details in SI, Section S3). Datasets of the OCO, NNN and triarylboranes scaffolds contain less molecules (61, 80 and 181, respectively) and were mainly used to assess the extrapolation capabilities of the developed models and to compare the effects of substituents on Lewis acidity across different scaffolds for molecular design. This brings the total number of molecules to 594, each of which was subsequently labeled with its computed FIA value.

The range of FIA values varies depending on the ligand scaffold (Fig. 2B). OCO scaffold exhibits a very narrow distribution of FIAs, whereas FIA is relatively normally distributed across ONO derivatives, which will be advantageous when constructing ML models. In contrast, triarylboranes show a more uniform distribution, offering a broader spectrum of accessible FIA values. Among these, the strongest Lewis acids are found in the ONO and triarylboranes scaffolds.




Fig. 2 Studied chemical space. (A) *t*-SNE map on fingerprint representation showing the ONO chemical space (2197 molecules, light brown circles), the initial random dataset (purple circles) and the dataset extension thanks to k-means (light blue circles) and coverage (light orange circles) algorithms (see SI, Section S3). (B) Kernel density estimation plot of the FIA distributions for each molecular scaffold.



Constructing models

Creating performant models is a prerequisite to any interpretability task to ensure the reliability of the interpretations. We aimed to design novel compounds by investigating the effects of varying the heteroatoms coordinated to boron. To this end, we focused on the ONO scaffold for building ML models, as the unsubstituted NNN borane scaffold had already been synthesized and studied by Martin and co-workers.⁴³ In general, oxygen ligated boron-based Lewis acids are more acidic than their nitrogen-ligated counterparts, so we expected the ONO ligand to yield more acidic compounds (refer to the NNN and ONO FIA distributions in Fig. 2B). A benchmark of ML models and molecular descriptors was realized to identify the most effective combinations to predict FIA (Fig. 3). Models were optimized on a training set following a grid search algorithm and 10-fold cross-validation (CV). They were subsequently

evaluated on a testing set (see the SI, Section S5, for details on dataset splitting). Among the evaluated molecular descriptors, some were already implemented in Python libraries (like RDKit⁵² or DeepChem⁵³), and are commonly used in drug design, such as Morgan fingerprints^{54,55} and RDKit descriptors. While Morgan fingerprints are beneficial for visualizing fragment diversity within chemical space (Fig. 2A), they are less effective at accurately predicting FIA, with a mean absolute error (MAE) higher than 10 kJ mol⁻¹. This limitation arises because they focus primarily on local connectivity, probably missing the effect of the delocalized π electrons on aromatic rings. RDKit descriptors yield highly effective models, such as linear models or SVR (Fig. 3A), yet out of the 208 features generated (*e.g.*, counts of fragments, partial charge or molecular weight), only a few dozen is readily interpretable. More physics-based parameters can be obtained through the Auto-QChem⁹ DFT




Fig. 3 Benchmarking machine learning models and descriptors in FIA prediction. (A) Heatmap of MAE scores calculated on the testing set for different combinations of models and descriptors (Linear: linear regression, LR: linear ridge, Bayes. Ridge: Bayesian ridge, LASSO: least absolute shrinkage and selection operator, SVR: support vector regression, Tree: decision tree, RF: random forest, Grad. Boost.: gradient boosting, GPR: gaussian process regressor, KNN: K nearest neighbors, MLP: multilayer perceptron), (B) Boxplots of MAE obtained on the training set for the linear ridge model, combined with different molecular descriptors, (C) Boxplots of MAE scores on the training set for the Hammett-extended descriptors, combined with different regressors (10-fold CV repeated 10 times with different folds).



calculation workflow yielding DFT-derived molecular (e.g., frontier orbitals energies) and atomic descriptors for the boron atom (e.g., partial charge), totaling 43 features. These descriptors, referred to as “quantum descriptors”, offer reasonably good performance across all ML models but require significantly more computational resources. While computing FIA requires two DFT calculations, these quantum descriptors can be extracted from the borane structure alone. Predicting FIA from this single structure remains valuable, particularly because the tetravalent boronate form is often more challenging to converge. Although quantum descriptors may not be the most efficient choice for FIA prediction due to their computational cost, their strong physical relevance makes them ideal to investigate the relationship between the quantum features of molecules and their Lewis acidity (see part Interpretability – insights in the Lewis acidity). To balance cost and accuracy in determining physically meaningful descriptors, an alternative approach could be to use so-called “surrogate models”; though this was not explored in the present study.^{56–58} However, quantum parameters are not directly manipulable by organic chemists for designing molecules. Instead, the typical strategy involves designing molecular structures by substituting various functional groups on a molecular scaffold, which play a significant role in modulating Lewis acidity. To better understand this influence, particularly the electronic effects of substituents, we introduced substituent-based molecular descriptors. This method builds upon the foundational work of Hammett,¹⁰ who derived the substituent constants σ_m and σ_p , for the *meta* and *para* substituents respectively, establishing the groundwork for what would become a pioneering approach to QSAR. However, Hammett constants are limited to *meta* and *para* substituents on aromatic rings, as *ortho* substituents complicate the analysis by introducing both electronic and steric effects.⁵⁹ Indeed, when considering only *meta*- and *para*-substituted molecules from the triarylboranes dataset, the trend of FIA follows a linear relationship with σ_m and σ_p ($\text{FIA} = 243 \sigma_m + 91 \sigma_p + 351$, determination coefficient $R^2 = 0.91$) (see Fig. S9A). Conversely, when considering also the triarylboranes possessing a substituent at the *ortho* position, the linear relationship is compromised (Fig. S9B, $R^2 = 0.63$), and this effect is intensified when considering other molecular scaffolds such as ONO ($R^2 = 0.54$), for which the labeling of *ortho*, *meta* and *para* substituents is less clear. To address this issue, we implemented “Hammett-extended descriptors”, as proposed by Sigman *et al.*⁵⁹ Additionally to σ_m and σ_p , these descriptors encompass steric and electronic parameters for substituents across all three aryl positions (*ortho*, *meta*, and *para*) of benzoic acid. Although these parameters are computationally derived too, they are in essence very different from the “quantum descriptors” introduced above. They include infrared (IR) carbonyl stretching ($\nu_{\text{C=O}}$) and COH bending (ν_{COH}) frequencies and intensities, natural bond orbital (NBO) charges of each atom in the carboxylic acid moiety, Sterimol B_1 , B_5 , and L of the substituent,⁶⁰ and the torsion angle between the carbonyl group and the aromatic ring plane when the substituent is positioned in *ortho*. However, these descriptors have limited applicability because they can only be applied to one molecular scaffold at a time, since they

do not describe the carbon backbone of the molecule (Fig. S10). For the ONO scaffold, the *ortho*, *meta*, and *para* positions were conventionally defined relative to the oxygen bonding atoms. This convention was extended to the NNN scaffold, but these descriptors were not implemented for the OCO. Using SMARTS substructure identifiers⁶¹ implemented in the RDKit Python library,⁵² the chemical nature of substituents at *ortho*, *meta*, and *para* positions was identified. Then Hammett-extended parameters derived by Sigman and co-workers⁵⁹ corresponding to the *ortho*, *meta* and *para* substituents were concatenated into a vector featuring the molecule (36 features). These descriptors demonstrate robust performance with advanced models such as gradient boosting (Grad. Boost.) and multi-layer perceptron (MLP) regressors. They also excel with simpler models like linear and linear ridge (LR) regressors, achieving a mean absolute error of approximately 6 kJ mol^{-1} (Fig. 3C). Given the straightforward nature of the molecule description and the impressive results obtained with these basic models, they provide an optimal balance between interpretability and prediction accuracy. These findings align with Hammett’s theory of linear free energy relationships (LFER), which suggests that functional groups on an aromatic ring near the reactive site linearly influence the molecule’s reactivity.

From this benchmark (Fig. 3A), we recognized the importance of both global descriptors and substituent-based descriptors in predicting FIA, which is why RDKit and Hammett-extended descriptors were concatenated in an attempt to create a high-performant predictive model, an oracle. Apart from the challenge of creating an accurate predictive model, having such a model would be helpful to investigate molecular design (see part Interpretability – Molecular design) chemometrics on ONO chemical space. For any given molecule, one could quickly determine its precise FIA, as these concatenated descriptors are also fast to compute. This yielded 244 features, from which we selected 126 based on their correlation with the target FIA, evaluated using the F-statistic.




Fig. 4 Optimized model for ONO evaluated on the testing set (MAE = 5.39 kJ mol^{-1}).



We selected LR as the machine learning algorithm due to its strong performance in predicting FIA from these two molecular descriptors. This allowed to reach a MAE of 6.39 kJ mol^{-1} on the training set (repeated 10-fold CV, Fig. 3B) and of 5.39 kJ mol^{-1} on the testing set (Fig. 4, $R^2 = 0.98$), representing less than 2% error relative to the average FIA value (around 450 kJ mol^{-1}). The performance on the ONO dataset, with a MAE between 5 and 10 kJ mol^{-1} across almost all ML algorithms and molecular descriptors (Fig. 3A), can be attributed to the limited chemical space it encompasses. In contrast, the optimized graph neural network model of Greb and co-workers, evaluated on a wide variety of molecular structures, struggles to achieve a MAE below 10 kJ mol^{-1} . This model gave a MAE of 23 kJ mol^{-1} ($R^2 = 0.51$) evaluated on the testing dataset of ONO scaffold (in comparison with the MAE of 5.39 kJ mol^{-1} obtained by the present oracle). In addition, such deep learning (DL) models necessitate huge amounts of data to perform and are not designed for tasks in low-data regime.

Capacity to extrapolate to another molecular scaffold

After optimizing models to predict FIA within the restricted ONO chemical space, we aimed to assess their ability to predict FIA across different chemical spaces, particularly from a molecular scaffold to the other. We also explored methods, such as feature selection, to customize models for enhanced versatility. However, not all descriptors are suitable for this purpose. In addition to being difficult to interpret, fingerprints, represented as bit-vectors, have interdependent features that cannot be individually selected. Furthermore, Hammett-extended descriptors rely solely on the chemical properties of substituents. Molecules with identical *ortho*, *para*, and *meta* substituents are then assigned the same molecular descriptors, even though their FIAs can differ significantly across structures (refer to Fig. S10). For this reason, we focused on enhancing the extrapolation capabilities of models based on quantum or RDKit descriptors (see the SI, Section S5 for the latter).

Models trained on ONO dataset, featured by the quantum descriptors, can generally predict the FIA trend for the NNN, though with a model-dependent bias that impacts the MAE (Table S7). Simple models struggle to extrapolate, whereas the Multilayer Perceptron (MLP) achieves reasonable prediction performance (MAE = 24.1 kJ mol^{-1}), as expected given the inherent versatility of neural networks. This is illustrated in Fig. 5A, showing the prediction for the NNN from a LR model trained on ONO. While the trend is well captured (Pearson's $r = 0.96$), the MAE is high, reaching 188 kJ mol^{-1} . The bias in prediction observed in most models can be reduced by carefully selecting features. Indeed, reducing the complexity of models usually helps enhancing their extrapolation abilities. Following this general guideline, features can be removed along two axes. First, features that are not correlated with the target, here, FIA. This is especially relevant for linear models (like LR), which tend to overfit noise when uncorrelated features are included, thereby compromising extrapolation. Second, features displaying substantial distribution differences between the training and new test dataset warrant removal, regardless of




Fig. 5 Feature selection to extrapolate from ONO to NNN with the LR model and quantum descriptors. (A) No feature selection. (B) Features selected.

their apparent predictive importance for the model (like having strong correlation with the target for example). This selective approach prevents models from learning dataset-specific patterns that fail to generalize, thereby improving robustness when applied to new data contexts.

Therefore, we ranked features based on their differences between ONO and NNN structures (quantified by mean values of each feature, Table S9) and their correlation to the target FIA (Table S10). We assessed the prediction performance of the LR model by systematically removing features using these two criteria. Among the features not correlated with FIA, boron atom parameters such as its coordinates and buried volume led to a notable drop in MAE (from 188 to 158 kJ mol^{-1}) when removed. Eliminating the lower unoccupied molecular orbital (LUMO) energy, that is strongly correlated with FIA (Pearson's $r = 0.628$) but varies considerably between structures (mean value is -0.010 hartree for ONO and -0.002 hartree for NNN), decreased the MAE to 140 kJ mol^{-1} . Similarly, although not



correlated with FIA (Pearson's $r = 0.076$), the population of Rydberg atomic orbitals (NPA_Rydberg) changes significantly between the two structures. Removing this feature decreased the MAE to 14.5 kJ mol^{-1} , indicating a considerable contribution to the observed bias, albeit difficult to explain. When this feature alone was excluded, the MAE dropped directly from 227 to 35.5 kJ mol^{-1} . Further removal of the dipolar moment of the molecule reduced the MAE to 13.6 kJ mol^{-1} . The improvement in prediction performance of the LR model with the selected features is illustrated in Fig. 5B. Prediction performance remained high within the ONO (8.26 kJ mol^{-1}) and the NNN

(11.4 kJ mol^{-1}) chemical spaces using these selected features. While this feature selection approach is general, the optimal feature set remains task-specific, as evidenced by varying performance in predicting from ONO to OCO (MAE = 84.8 kJ mol^{-1}) and triarylboranes (MAE = 759 kJ mol^{-1}) using these features. This is due to inherent differences between scaffolds, notably in terms of FIA distribution (Fig. 2B) and electronic structure (Fig. 6). To explore these differences, we attempted training on three molecular structures while testing on the fourth, which partly improved the extrapolation performances (see the SI, Section S5).




Fig. 6 Principal component analysis (PCA) of the quantum descriptors. (A) Projection onto the first two principal components. Molecules are colored according to their FIA values. Triarylboranes (small crosses) appear clearly separated from other molecular scaffolds in terms of electronic structure of the boron atom. (B) Explained variance ratio for each component. (C) Contributions of the quantum descriptors features to the first two principal components.



Interpretability

Our goal here is to address two main questions by interpreting the developed models. First, we aim to gain insight into the intrinsic nature of Lewis acidity for boron derivatives. Second, we seek actionable explanations for a compound's Lewis acidity, exploring new routes to design molecular structures with targeted Lewis acidity.

Insights in the Lewis acidity. ML can serve as an instrument to unveil properties of a system that are challenging or even impossible to probe using traditional methods. Lewis acidity is inherently a qualitative concept. It can be accessed through the reaction enthalpy of a Lewis acid-Lewis base adduct formation, but unlike electrophilicity, that is associated with the energy level of the LUMO, the relationship between molecular properties and Lewis acidity remains elusive. Analyzing FIA predictive models may help to unravel the origins of Lewis acidity at the electronic scale.

We used quantum descriptors, that provide precise physical parameters for characterizing molecules, and examined the whole database (the four molecular scaffolds together) to find broader patterns and insights. Conducting a principal component analysis (PCA) to simplify complex data into two dimensions, we identified distinct groups of molecules (Fig. 6). PC1 primarily encompasses electronic parameters of the boron atom (ES_root_Mulliken_charge, ES_root_NPA_Rydberg, ES_root_NPA_charge, ES_root_NPA_core, ES_root_NPA_total, ES_root_NPA_valence, Mulliken_charge, NMR_anisotropy, NMR_shift, NPA_Rydberg, NPA_charge, NPA_core, NPA_total, NPA_valence, VBur, that are 15 out of the 27 features used), whereas none of these atomic features contribute to PC2. Triarylboranes appear thus separate from other structures along the PC1 axis due to the electronic structure of boron atom. This is unsurprising as triarylboranes lack heteroatoms bonded to boron, which can affect electronic population through electron-withdrawing effects. Additionally, the two PC derived from quantum descriptors reflect FIA evolution across the database. This confirms the relevance of these descriptors in capturing Lewis acidity. Additional projection onto PC1 and PC3 revealed a separation of the constrained scaffold molecules cluster (SI, Section S6, Fig. S13).

We then simplified the data by choosing twelve significant but uncorrelated features through a hierarchical clustering (SI, Section S6, Fig. S14.A). Most ML methods struggled with these simplified descriptors across the four molecular scaffolds, except for tree-based ensemble models like Random Forest (RF) and Grad. Boost., which performed better than the others (see Table S17). Our goal was to explore how these features impacted the FIA using two interpretable ML approaches. We first looked at a straightforward and simple model using linear regression to understand how much each feature contributed. Then, we used the permutation feature importance post-hoc technique, with a Grad. Boost. regressor, as this model was the most performant on uncorrelated features (MAE = 14.6 kJ mol⁻¹, Table S17). The learned coefficient of each feature in the linear model, as depicted in Fig. S14C, reveals its importance. Both the global

electronegativity of the molecule and the partial charge of the boron atom (NPA_charge) emerge as the most significant features, with approximately equal importance. This observation aligns with their ranking in the permutation feature importance, although the electronegativity being slightly more influential (Fig. S14B). When the regression is restricted to the ONO dataset (Fig. S16), electronegativity becomes the most significant parameter, with the partial charge coefficient being three times lower. Electronegativity is a linear combination of frontier orbital energies, suggesting that, for the boron derivatives studied, Lewis acidity is somewhat more influenced by molecular orbital interactions than by coulombic interactions. This is notable because FIA is traditionally considered as an index of hard acidity, where coulombic interactions would be expected to dominate.^{47,62}

The absolute electronegativity is then a key global factor in determining FIA magnitude, while local electronic parameters of boron atom more precisely adjust the predicted FIA. FIA can be roughly linearly predicted using only these two parameters for the ONO scaffold (FIA = 60.0 χ + 8.15 NPA charge + 161, R^2 = 0.88, Fig. S17). Steric parameters (e.g., molar volume) seem to play a less significant role, potentially due to the predominance of constrained-ligand boron derivatives in our database. Thus, controlling the electronic environment around the boron atom is essential. While chemists cannot directly manipulate quantum properties that are not "actionable", they can adjust Lewis acidity through molecular design by varying functional groups on the molecular scaffold. However, this approach is complicated by the interconnected effects of the substituents on the same reactive site, making the use of ML relevant.


Molecular design

Interpretable ML. This part will focus on the ONO molecular scaffold (see SI, Section S6 for a comparison with the triarylboranes). We used Hammett-extended descriptors, that are specifically tailored to unravel the effect of substituents.

Before any regression by a ML algorithm, we examined the correlation between *ortho*, *meta*, and *para* parameters, relative to the oxygens, and FIA (Fig. S18A), which shows that *para* substituents have the most significant effect on FIA. Steric parameters, such as the torsional angle θ_{tor} or sterimol parameters, poorly correlate with FIA. We then selected twelve uncorrelated features through hierarchical clustering (Fig. S19A). Using again the linear model (Fig. S18B) and a permutation importance algorithm (Fig. S19B) on a Grad. Boost. regressor, we identified the Hammett σ constants as the most significant parameters, σ_p being slightly more important than σ_m , with a lower contribution of parameters characterizing the electronic demand of the *ortho* substituent. This observation is consistent with the correlations with FIA (Fig. S18A).

These methods help identify which aromatic positions to prioritize when designing an ONO compound with a specific Lewis acidity, namely here, the *para* position. However, they do not inform on the electronic demand of the substituent or the interactions between substituents at different positions. We used decision trees to understand the decision criteria for





Scheme 1 Decision tree model for the ONO structure (0.73 accuracy on 10-fold CV). Nodes are split based on Hammett-extended feature thresholds. A left arrow indicates a “Yes” decision, while a right arrow indicates a “No” decision. The entropy (E) of each node reflects its purity, indicating whether it contains a unique class ($E = 0$) or not. Original tree from Scikit-learn⁶³ is provided in Fig. S20.

predicting FIA. We thus categorized FIA into six classes based on the ONO chemical space distribution (Fig. 2B), turning the original regression task into a classification problem. To simplify the analysis and the resulting tree structure, we focused on four Hammett-extended features capturing steric and electronic effects of *ortho*, *meta*, and *para* substituents. We included the partial charge of the oxygen atom of the C=O

bond of benzoic acid ($\text{NBO}_{\text{=O}}$) for each position, as the Sigman group showed that these parameters could effectively replace the empirical σ constants,⁵⁹ plus L_o , the bond length between the *ortho* substituent and the aromatic ring, to account for steric effects at the *ortho* position. The root node of the decision tree shows that the $\text{NBO}_{\text{=O}}$ charge for the *para* substituent is the critical criterion for classifying LA as good, strong or super, if it




Fig. 7 (A) Violin plots separating data between molecules possessing mesomeric electron-withdrawing groups (CN or NO_2) in *para* position or not. (B) Distributions of molecules with a mesomeric electron-withdrawing group in *para*, separated by the *ortho* group and highlighting variations in the *meta* group; dotted markers indicate molecules where the *meta* group matches the *ortho* group.



exceeds $-0.59e$, a threshold distinguishing mesomeric electron-withdrawing groups (such as CN and NO₂) from others (see Table S6); otherwise, it is classified as medium. Subsequent nodes evaluate the *ortho* and *meta* substituents for majority-medium and majority-good classes, respectively. The obtained tree diagram is depicted in Scheme 1, translated into chemical terms based on the feature thresholds from the original Scikit-learn⁶³ tree (see Fig. S20).

In summary, the tree structure shows that adding or not a mesomeric electron-withdrawing group at the *para* position of an ONO molecule already determines the accessible range of FIA (Scheme 1). Then, the substituents at the *ortho* and *meta* positions must be examined to refine the targeted FIA value. To confirm these results, we need more data for statistical analyses on the diverse substituents. For that, we have used the oracle previously developed (Fig. 4) to screen the entire ONO chemical space and provide precise FIA values to enrich our database.

Chemometrics on ONO chemical space. Using the root node criterion from the tree in Scheme 1, we compared the FIA distributions for ONO molecules with a mesomeric electron-withdrawing group in the *para* position to those without. The 2197-molecules chemical space is effectively separated as shown by the weak overlap between violin plots of FIA distributions (Fig. 7A). Next, we divided the set of molecules possessing a mesomeric electron-withdrawing group in *para* by their *ortho* group (Fig. 7B). Since these distributions are still broad, merely fixing the *ortho* group while maintaining a *para* mesomeric electron-withdrawing group fails to standardize the FIA. Therefore, we put the same substituent at both the *meta* and *ortho* positions, resulting in a single molecule, with a fixed FIA value. The approach ensures that the electronic demand at the *ortho* and *meta* positions is nearly the same, reducing the complexity of the interconnected effects of the substituents. Varying this substituent, while keeping a mesomeric electron-withdrawing group in *para*, enables exploration of the full range of FIA values (from 400 to 600 kJ mol⁻¹) for the ONO scaffold (see the spread of ONO FIA distribution, Fig. 2B).

This analysis suggests that if an organic chemist needs a compound within a specific FIA range, say 450 to 500 kJ mol⁻¹, a practical approach using the ONO scaffold might involve adding a mesomeric electron-withdrawing group in *para* and halogens (such as -F, -Br, or -Cl) in *ortho* and *meta* positions, although synthetic challenges should be considered when implementing this strategy.

Conclusion

ML has proven highly effective in predicting Lewis acidity *via* FIA regression models. Advanced techniques like GNN, trained on extensive datasets of tens of thousands of molecules,⁴⁵ achieve MAE levels of 14 kJ mol⁻¹, enabling rapid reactivity assessment without costly *ab initio* calculations. However, these models require substantial computational resources, which are unaffordable for many experimental labs, but also extensive data, which is challenging to acquire in chemistry considering data scarcity, heterogeneity, and cost. Nonetheless, effective simple models can be developed in low-data regimes by

constraining the chemical space. We have successfully developed such a high-performing oracle, achieving a MAE of less than 6 kJ mol⁻¹ (*i.e.*, less than 2% error on average). Moreover, while simple classical ML models may be less versatile than DL models, they can be adapted to extrapolate across different chemical spaces through careful feature selection. DL models continue to serve as powerful tools for predicting quantities with profoundly nonlinear relationships with the features. However, in the case of FIA, we have demonstrated that linear models based on physically meaningful descriptors perform exceptionally well, suggesting a linear relationship between FIA and molecular parameters.

Apart from their simplicity, linear models and decision trees, that use chemical features in a way that resembles chemical thought, are also more interpretable than “black-box” DL models. Models that can learn and represent genuine chemical concepts are essential for integration into a broader scientific approach: building models, making predictions, interpreting results for deeper understanding, and eventually refining the models. Developing such ML approaches grounded in physical reality remains a challenge in chemistry.²¹

In this study, we leveraged explainable ML in two key ways. We employed quantum descriptors based on electronic structure parameters and revealed that the reactivity of the studied boron Lewis acids is governed by molecular orbital interactions, classifying them as soft Lewis acids. However, these quantum descriptors are not readily actionable for molecular design. To address this, we introduced “Hammett-extended” descriptors, which focus on the nature and positioning of substituents. Interpretations from these models align with the XAI attributes identified by Wellawatte *et al.*²⁵ They are actionable and succinct, and, as they align with the language and concepts of organic chemistry, they are domain applicable. These explanations are also correct, corresponding to the expected electronic demand of a substituent at a given position. By interpreting models based on these descriptors, we unraveled the interdependent effects of substituents and identified design rules for creating novel ONO compounds with targeted Lewis acidity.

Here, our aim was to provide a route for the design of novel organic compounds, bridging the gap between pure ML techniques and traditional intuition-based strategies in organic chemistry. While our focus was centered on Lewis acidity, this methodology holds promise for exploring other types of reactivity, provided relevant parameters can be readily accessed (either computationally or experimentally) to construct a dataset.

Author contributions

JF, RV and LG conceptualized the project. JF performed the work and wrote the manuscript. LG and RV directed the work, contributed to the writing of the manuscript, as well as secured funding for the research.

Conflicts of interest

There are no conflicts to declare.



Data availability

Supplementary information is available (see DOI: <https://doi.org/10.1039/d5dd00212e>) for additional details and extended results (Lewis acidity metrics benchmark, computational methods, database construction, molecular descriptors implementation, machine learning models building and interpretation). This project was implemented in Python leveraging the Scikit-Learn and RDKit libraries. The code, including optimized models and datasets, is publicly available on GitHub at https://github.com/jfenogli/XAI_boron_LA and on Zenodo at <https://doi.org/10.5281/zenodo.17019794>. This repository also includes Jupyter notebooks to reproduce the analyses presented in this manuscript and adapt workflows for new datasets or reactivity studies.

Acknowledgements

The authors would like to thank Dr. Jules Schleinitz for his valuable guidance at the beginning of this project and his feedback on the manuscript. We also acknowledge Pablo Mas for his assistance with database extension *via* clustering algorithms and Maxime R. Vitale for fruitful discussion. JF acknowledges ENS-PSL for the doctoral fellowship support. We are grateful to GENCI for the allocation of computational resources (allocation No. A0170802309), as well as CNRS and Sorbonne Université for institutional and financial support.

References

- J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Židek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli and D. Hassabis, *Nature*, 2021, **596**, 583–589.
- R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirnsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, A. Merose, S. Hoyer, G. Holland, O. Vinyals, J. Stott, A. Pritzel, S. Mohamed and P. Battaglia, *Science*, 2023, **382**, 1416–1421.
- N. Tomašev, X. Glorot, J. W. Rae, M. Zielinski, H. Askham, A. Saraiva, A. Mottram, C. Meyer, S. Ravuri, I. Protsyuk, A. Connell, C. O. Hughes, A. Karthikesalingam, J. Cornebise, H. Montgomery, G. Rees, C. Laing, C. R. Baker, K. Peterson, R. Reeves, D. Hassabis, D. King, M. Suleyman, T. Back, C. Nielson, J. R. Ledsam and S. Mohamed, *Nature*, 2019, **572**, 116–119.
- S. M. Kearnes, M. R. Maser, M. Wlekinski, A. Kast, A. G. Doyle, S. D. Dreher, J. M. Hawkins, K. F. Jensen and C. W. Coley, *J. Am. Chem. Soc.*, 2021, **143**, 18820–18826.
- J. Schleinitz, M. Langevin, Y. Smail, B. Wehnert, L. Grimaud and R. Vuilleumier, *J. Am. Chem. Soc.*, 2022, **144**, 14722–14730.
- J. A. G. Torres, S. H. Lau, P. Anchuri, J. M. Stevens, J. E. Tabora, J. Li, A. Borovika, R. P. Adams and A. G. Doyle, *J. Am. Chem. Soc.*, 2022, **144**, 19999–20007.
- D. T. Ahneman, J. G. Estrada, S. Lin, S. D. Dreher and A. G. Doyle, *Science*, 2018, **360**, 186–190.
- A. F. Zahrt, J. J. Henle, B. T. Rose, Y. Wang, W. T. Darrow and S. E. Denmark, *Science*, 2019, **363**, eaau5631.
- A. M. Žurański, J. Y. Wang, B. J. Shields and A. G. Doyle, *React. Chem. Eng.*, 2022, **7**, 1276–1284.
- L. P. Hammett, *J. Am. Chem. Soc.*, 1937, **59**, 96–103.
- W. L. Williams, L. Zeng, T. Gensch, M. S. Sigman, A. G. Doyle and E. V. Anslyn, *ACS Cent. Sci.*, 2021, **7**, 1622–1637.
- J. Vamathevan, D. Clark, P. Czodrowski, I. Dunham, E. Ferran, G. Lee, B. Li, A. Madabhushi, P. Shah, M. Spitzer and S. Zhao, *Nat Rev Drug Discov*, 2019, **18**, 463–477.
- B. Deng, P. Zhong, K. Jun, J. Riebesell, K. Han, C. J. Bartel and G. Ceder, *Nat. Mach. Intell.*, 2023, **5**, 1031–1041.
- R. David, M. De La Puente, A. Gomez, O. Anton, G. Stirnemann and D. Laage, *Digital Discovery*, 2025, DOI: [10.1039/D4DD00209A](https://doi.org/10.1039/D4DD00209A).
- W. Beker, E. P. Gajewska, T. Badowski and B. A. Grzybowski, *Angew. Chem., Int. Ed.*, 2019, **58**, 4515–4519.
- E. Caldeweyher, M. Elkin, G. Gheibi, M. Johansson, C. Sköld, P.-O. Norrby and J. F. Hartwig, *J. Am. Chem. Soc.*, 2023, **145**, 17367–17376.
- P. Schwaller, A. C. Vaucher, T. Laino and J.-L. Reymond, *Mach. learn.: sci. technol.*, 2021, **2**, 015016.
- M. H. S. Segler, M. Preuss and M. P. Waller, *Nature*, 2018, **555**, 604–610.
- M. Suvarna, A. C. Vaucher, S. Mitchell, T. Laino and J. Pérez-Ramírez, *Nat. Commun.*, 2023, **14**, 7964.
- P. Schwaller, B. Hoover, J.-L. Reymond, H. Strobelt and T. Laino, *Sci. Adv.*, 2021, **7**, eabe4166.
- J. A. Kammeraad, J. Goetz, E. A. Walker, A. Tewari and P. M. Zimmerman, *J. Chem. Inf. Model.*, 2020, **60**, 1290–1301.
- J. G. Richens, C. M. Lee and S. Johri, *Nat. Commun.*, 2020, **11**, 3923.
- A. J. DeGrave, J. D. Janizek and S.-I. Lee, *Nat. Mach. Intell.*, 2021, **3**, 610–619.
- M. Krenn, R. Pollice, S. Y. Guo, M. Aldeghi, A. Cervera-Liarta, P. Friederich, G. dos Passos Gomes, F. Häse, A. Jinich, A. Nigam, Z. Yao and A. Aspuru-Guzik, *Nat. Rev. Phys.*, 2022, **4**, 761–769.
- G. P. Wellawatte, H. A. Gandhi, A. Seshadri and A. D. White, *J. Chem. Theory Comput.*, 2023, **19**, 2149–2160.
- Z. Wu, J. Chen, Y. Li, Y. Deng, H. Zhao, C.-Y. Hsieh and T. Hou, *J. Chem. Inf. Model.*, 2023, **63**, 7617–7627.
- G. P. Wellawatte, A. Seshadri and A. D. White, *Chem. Sci.*, 2022, **13**, 3697–3705.
- G. P. Wellawatte and P. Schwaller, *arXiv*, 2023, preprint, arXiv:2311.04047, <http://arxiv.org/abs/2311.04047>.
- H. A. Gandhi and A. D. White, *ChemRxiv*, 2022, preprint, DOI: [10.26434/chemrxiv-2022-v5p6m-v3](https://doi.org/10.26434/chemrxiv-2022-v5p6m-v3).



- 30 T. Tian, S. Li, M. Fang, D. Zhao and J. Zeng, *J. Chem. Inf. Model.*, 2024, **64**, 2236–2249.
- 31 J. Jiménez-Luna, F. Grisoni and G. Schneider, *Nat. Mach. Intell.*, 2020, **2**, 573–584.
- 32 C. B. Santiago, J.-Y. Guo and M. S. Sigman, *Chem. Sci.*, 2018, **9**, 2398–2412.
- 33 K. Wu and A. G. Doyle, *Nat. Chem.*, 2017, **9**, 779–784.
- 34 M. Esders, T. Schnake, J. Lederer, A. Kabylda, G. Montavon, A. Tkatchenko and K.-R. Müller, *J. Chem. Theory Comput.*, 2025, **21**, 714–729.
- 35 K. Bonneau, J. Lederer, C. Templeton, D. Rosenberger, K.-R. Müller and C. Clementi, *arXiv*, 2024, preprint, arXiv:2407.04526, DOI: [10.48550/arXiv.2407.04526](https://doi.org/10.48550/arXiv.2407.04526).
- 36 M. M. Shmakov, S. A. Prihod'ko, V. N. Panchenko, M. N. Timofeeva, V. V. Bardin, V. N. Parmon and N. Yu. Adonin, *Catal. Rev.*, 2024, 1–47.
- 37 T. P. L. Cosby, A. Bhattacharjee, S. K. Henneberry, J. LeBlanc and C. B. Caputo, *Chem. Commun.*, 2024, **60**, 5391–5394.
- 38 G. J. P. Britovsek, J. Ugoletti and A. J. P. White, *Organometallics*, 2005, **24**, 1685–1691.
- 39 A. E. Ashley, T. J. Herrington, G. G. Wildgoose, H. Zaher, A. L. Thompson, N. H. Rees, T. Krämer and D. O'Hare, *J. Am. Chem. Soc.*, 2011, **133**, 14727–14740.
- 40 M. M. Morgan, A. J. V. Marwitz, W. E. Piers and M. Parvez, *Organometallics*, 2013, **32**, 317–322.
- 41 M. M. Alharbi, Y. Van Ingen, A. Roldan, T. Kaehler and R. L. Melen, *Dalton Trans.*, 2023, **52**, 1820–1825.
- 42 A. Ben Saida, A. Chardon, A. Osi, N. Tumanov, J. Wouters, A. I. Adjieufack, B. Champagne and G. Berionni, *Angew. Chem., Int. Ed.*, 2019, **58**, 16889–16893.
- 43 K. Huang, J. L. Dutton and C. D. Martin, *Chem.–Eur. J.*, 2017, **23**, 10532–10535.
- 44 W. Lv, Y. Dai, R. Guo, Y. Su, D. A. Ruiz, L. L. Liu, C. Tung and L. Kong, *Angew. Chem.*, 2023, **135**, e202308467.
- 45 L. M. Sigmund, S. Sowndarya, A. Albers, P. Erdmann, R. S. Paton and L. Greb, *Angew. Chem., Int. Ed.*, 2024, e202401084.
- 46 J. P. Janet and H. J. Kulik, *Machine Learning in Chemistry, Am. Chem. Soc.*, 2020, DOI: [10.1021/acs.infocus.7e4001](https://doi.org/10.1021/acs.infocus.7e4001).
- 47 P. Erdmann, J. Leitner, J. Schwarz and L. Greb, *ChemPhysChem*, 2020, **21**, 987–994.
- 48 T. P. Robinson, D. M. De Rosa, S. Aldridge and J. M. Goicoechea, *Angew. Chem., Int. Ed.*, 2015, **54**, 13758–13763.
- 49 Y. Miyamoto, Y. Sumida and H. Ohmiya, *Org. Lett.*, 2021, **23**, 5865–5870.
- 50 J. MacQueen, in *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, 1967, vol. 51, pp. 281–298.
- 51 D. J. Woodward, A. R. Bradley and W. P. van Hoorn, *J. Chem. Inf. Model.*, 2022, **62**, 4391–4402.
- 52 RDKit, <https://www.rdkit.org/>.
- 53 DeepChem, <https://deepchem.io/>.
- 54 D. Rogers and M. Hahn, *J. Chem. Inf. Model.*, 2010, **50**, 742–754.
- 55 S. Zhong and X. Guan, *Environ. Sci. Technol.*, 2023, **57**, 18193–18202.
- 56 Y. Guan, C. W. Coley, H. Wu, D. Ranasinghe, E. Heid, T. J. Struble, L. Pattanaik, W. H. Green and K. F. Jensen, *Chem. Sci.*, 2021, **12**, 2198–2208.
- 57 T. Stuyver and C. W. Coley, *J. Chem. Phys.*, 2022, **156**, 084104.
- 58 J. E. Alfonso-Ramos, R. M. Neeser and T. Stuyver, *Digital Discovery*, 2024, **3**, 919–931.
- 59 C. B. Santiago, A. Milo and M. S. Sigman, *J. Am. Chem. Soc.*, 2016, **138**, 13424–13430.
- 60 A. Verloop, W. Hoogenstraaten and J. Tipker, in *Drug Design*, ed. E. J. Ariëns, Academic Press, Amsterdam, 1976, vol. 11, pp. 165–207.
- 61 Daylight Inc. 4, SMARTS—a language for describing molecular patterns, <http://www.daylight.com/dayhtml/doc/theory/theory.smarts.html>.
- 62 R. G. Pearson, *Proc. Natl. Acad. Sci. U.S.A.*, 1986, **83**, 8440–8441.
- 63 Scikit-learn, <https://scikit-learn.org/>.

