



HAL
open science

Machine Learning Yield Prediction from NiCOLit, a Small-Size Literature Data Set of Nickel Catalyzed C–O Couplings

Jules Schleinitz, Maxime Langevin, Yanis Smail, Benjamin Wehnert, Laurence Grimaud, Rodolphe Vuilleumier

► **To cite this version:**

Jules Schleinitz, Maxime Langevin, Yanis Smail, Benjamin Wehnert, Laurence Grimaud, et al.. Machine Learning Yield Prediction from NiCOLit, a Small-Size Literature Data Set of Nickel Catalyzed C–O Couplings. *Journal of the American Chemical Society*, 2022, 144 (32), pp.14722-14730. 10.1021/jacs.2c05302 . hal-03790865

HAL Id: hal-03790865

<https://hal.sorbonne-universite.fr/hal-03790865v1>

Submitted on 28 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Machine Learning Yield Prediction from NiCOLit, a Small-Size Literature Dataset of Nickel Catalyzed C-O Couplings

J. Schleinitz,^{*,†,‡} M. Langevin,^{*,†,¶,§} Y. Smail,^{||} B. Wehnert,^{||} L. Grimaud,^{*,‡} and R.
Vuilleumier^{*,¶}

†Those authors contributed equally to this work

*‡LBM, Département de chimie, École Normale Supérieure, PSL University, Sorbonne
Université, CNRS, 75005, Paris, France*

*¶PASTEUR, Département de chimie, École Normale Supérieure, PSL University, Sorbonne
Université, CNRS, 75005, Paris, France*

*§Molecular Design Sciences - Integrated Drug Discovery, Sanofi R&D, 94400,
Vitry-sur-Seine, France*

||UPMC, PSL University, Sorbonne Université, CNRS, 75005, Paris, France

E-mail: jules.schleinitz@ens.psl.eu; maxime.langevin@sanofi.com; laurence.grimaud@ens.psl.eu;
rodolphe.vuilleumier@ens.psl.eu

Abstract

Synthetic yield prediction using machine learning is intensively studied. Previous work focused on two categories of datasets: High-Throughput Experimentation data, as an ideal case study and datasets extracted from proprietary databases, which are known to have a strong reporting bias towards high yields. However, predicting yields using published reaction data remains elusive. To fill the gap, we built a dataset on

nickel-catalyzed cross-couplings extracted from organic reaction publications, including scope and optimization information. We demonstrate the importance of including optimization data as a source of failed experiments and emphasize how publication constraints shape the exploration of the chemical space by the synthetic community. While machine learning models still fail to perform out-of-sample predictions, this work shows that adding chemical knowledge enables fair predictions in a low-data regime. Eventually, we hope that this unique public database will foster further improvements of machine learning methods for reaction yield prediction in a more realistic context.

Introduction

Machine learning (ML) algorithms learn complex functions from data. As it can leverage existing data to perform *in silico* approximations of costly experimental processes, ML applications have sparked strong interest in chemical sciences. While ML has already made a significant impact in drug development,^{1,2} synthetizability assessment of small molecules³ or Computer Aided Synthesis Planning,⁴ the ability of ML to predict a reaction yield from its experimental conditions remains a major challenge⁵ that is intensively studied.^{6,7} Advances on reaction yield prediction would have a major impact on organic synthesis by significantly reducing cost, time and resources necessary to synthesize new chemicals.

Progress in ML is markedly driven by the increasing access to data. Thus, currently available datasets shape the evolution of ML for reaction yield prediction. Despite this, there are very few publicly available and easily operable datasets of chemical reactions with associated yields (Table S1). One of those few public datasets is the United State Patent and Trademark Office (USPTO) dataset,⁸ that covers a wide range of chemical reactions extracted from patents. USPTO data is extremely diverse and suffers from a selection bias as only successful reactions tend to be reported in patents. ML has shown poor performance predicting yields on this dataset ($R^2 < 0.2$).⁷ In addition, two sets of High Throughput Experimentation (HTE) data, one of a Suzuki-Miyaura coupling,⁹ and one of a palladium-

catalyzed Buchwald-Hartwig cross-coupling,¹⁰ are available in the literature. State-of-the-art modeling performs extremely well on those high-quality datasets ($R^2 > 0.8$),^{7,10} but the extremely focused chemical reaction space covered by HTE limits the predictions to a narrow scope of experimental conditions and reactants.

While these datasets have enabled rapid progress of ML for yield prediction, there is a need for publicly available datasets^{11,12} more representative of published reaction data or used by chemists in their everyday work. To the best of our knowledge, the most recent works on predicting reaction yields^{5,13} rely on datasets extracted from Reaxys or Sci-Finderⁿ, which are not representative of the whole information contained in published reaction data. The main hurdle to gather a machine readable reaction database is the difficulty to automate data extraction from publications. One of the solution to overcome this issue would be a change toward a numerical data storage in the chemistry community, an option being the use of electronic laboratory notebooks¹⁴ (ELN) interfaced with open-access database.¹⁵ Nevertheless, the implementation of such tools requires significant time and investment. It also requires to convince the chemistry community of the merits of gathering data in a machine-readable format. Thus, we believe that showing the potentiality of a dataset derived from published reaction data to predict reaction yield would encourage chemists to embrace the new technologies available. Such a change would benefit the whole chemistry community.

To address this, we built a literature-mined, open-access reaction dataset that focuses on the Ni-catalyzed C–O bond activation to form C–C and C–N bonds: the NiCOLit dataset.¹⁶ It gathers more than two thousand peer-reviewed reactions with detailed experimental conditions. As a singular literature representative dataset, NiCOLit stands as a benchmark for machine learning prediction of chemical yields found in published reaction data.

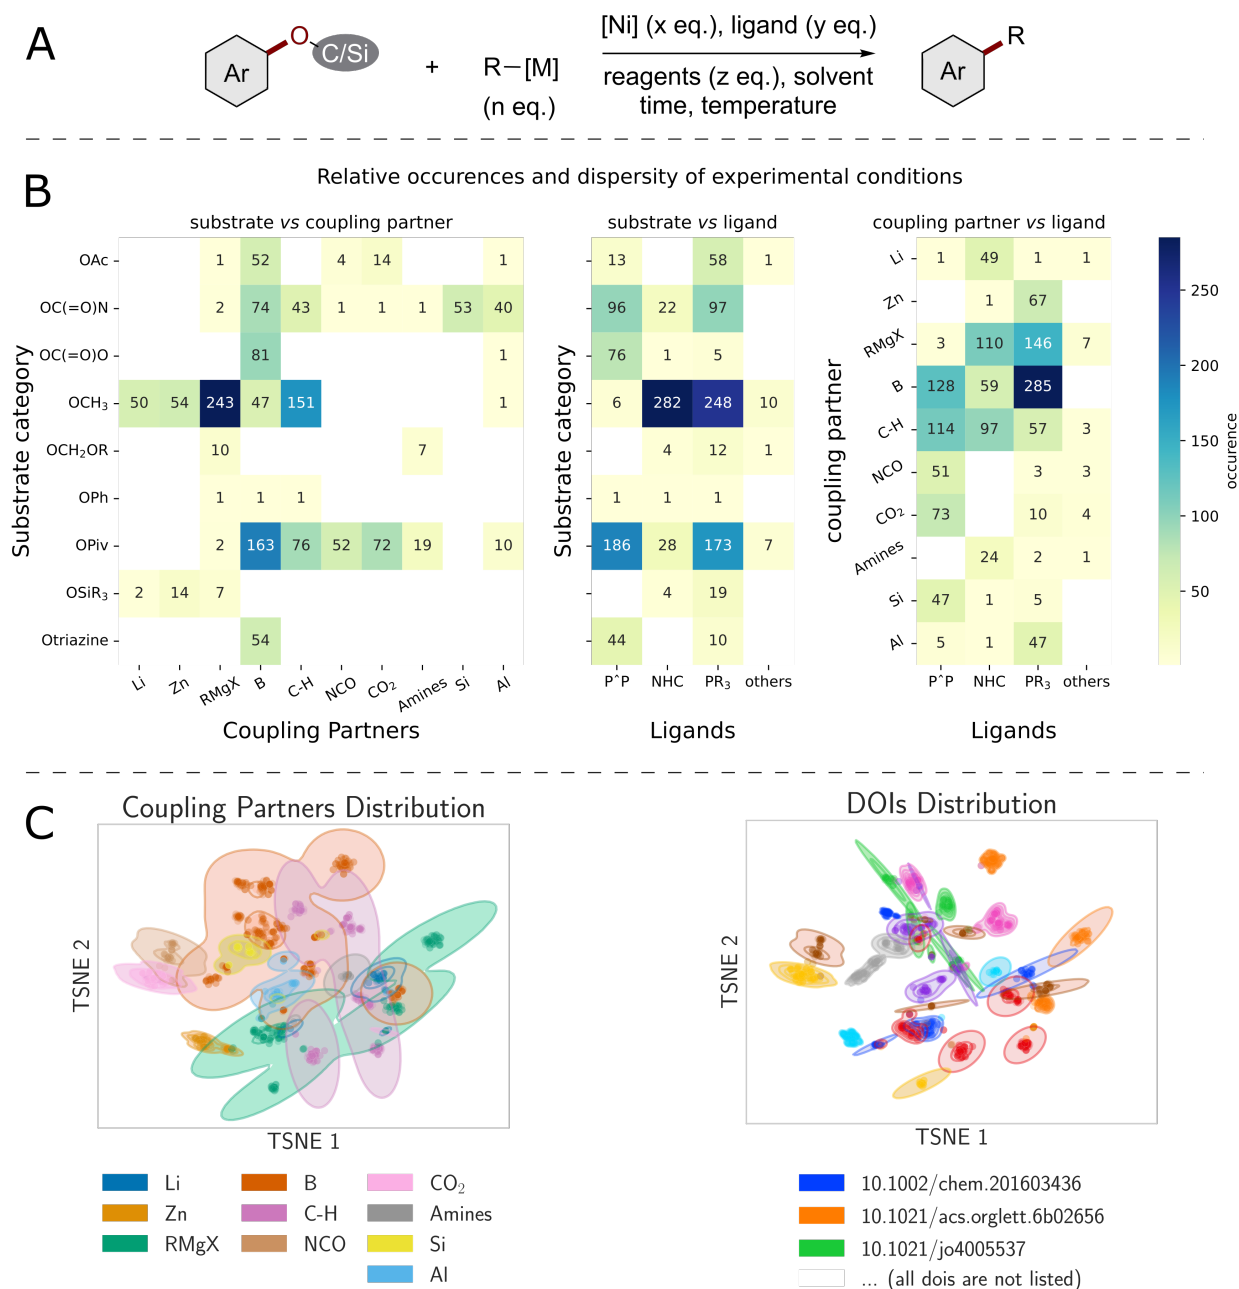


Figure 1: (A) Chemical space of NiCOLit, (B) Diversity of NiCOLit in terms of coupling partners, substrates and ligands combinations, the color map is proportional to the number of reactions encountered for each combination. (C) t-SNE projection of NiCOLit obtained with a DFT-featurization of the dataset, reaction data points are coloured by coupling partner category (left figure) and by publication origin (right figure).

Results and Discussion

Description of NiCOLit

NiCOLit was manually extracted from published reaction data cited in a recent review from Diao and co-workers.¹⁷ In this review, the authors focus on the activation of carbon-oxygen bonds of phenol derivatives with nickel catalysts for coupling reactions. In order to reduce the size and the diversity of the dataset we arbitrarily restrained the study to challenging electrophiles towards the oxidative addition: sulfonates,¹⁸ phosphates and *in situ* activated phenols were left aside.

The reactions displayed in both the main articles and their SI were extracted. For each reaction, the Simplified Molecular-Input Line-Entry System (SMILES)¹⁹ chains of substrates, coupling partners, precursors, ligands, bases, additives, solvents, and products were gathered as well as experimental variables: reaction time, temperature and molar ratios of the different partners (see SI Section 2). This highlighted issues when harmonizing different data sources, such as disparities in yield measurement techniques, or information being reported in prose rather than machine-readable format. The resulting database is unique, it gathers 2003 reactions from 45 publications within the chemical space illustrated Fig 1A. The different types of substrates, coupling partners and ligands reported in this dataset are summarized in Fig 1B.

The coverage of the substrate-coupling partners combinations is sparse and strongly biased toward a few reactions such as the Kumada coupling²⁰ ($\text{RMgX} + \text{ArOCH}_3$), which is the most investigated with 243 reported reactions. This data-driven analysis²¹ of nickel-catalyzed couplings shows that published reaction data seems to focus on specific combinations. We were surprised to see that reactions were clustered not by reaction variables (e.g coupling partner) but almost perfectly by publications (Fig 1C). This minimal overlap between publications may be attributed to the poor reporting of reproduced experiments.

Chemical Diversity: HTE versus NiCOLit

Inherent differences in terms of chemical diversity and yield distributions between NiCOLit and HTE data were laid out. This allows further understanding of how performances displayed by ML on HTE could translate to NiCOLit.

The projections of the three datasets on a common 2-dimensional space using Multi-Dimensional Scaling (MDS)²² (Fig. 2A, see SI for details) shows that NiCOLit is more spread out than HTE data. The distributions of the pairwise Jaccard distances between the reactions of each dataset were also computed (Fig. 2B) and corroborate the MDS analysis: distances between NiCOLit reactions are on average higher than in HTE data.

Most strikingly, we calculated the accessible chemical space as the number of all possible combinations of discrete variables used for each category (e.g. reactants, catalysts, etc.). The proportion of accessible space explored²³ is an intensive metric that measures the ratio between the number of chemical reactions experimentally performed and the size of the accessible chemical space. Despite having roughly a similar number of reactions in the three datasets, the accessible chemical space of NiCOLit covers almost a trillion reactions versus less than 20k for both HTE datasets (Fig. 2C). This indicates that predicting yields on NiCOLit is more challenging than on HTE data. On the other hand, as most of the accessible space has been explored in HTE data (99% and 36% against only $2 \times 10^{-7}\%$ for NiCOLit), developing an accurate model for NiCOLit allows to predict yields for a much larger set of unperformed reactions (almost a trillion reactions for NiCOLit).

Scope/optimization structure of NiCOLit

The presence of reactions with low yields within a dataset is expected to be key for accurate predictions.^{11,24} Very recently, Glorius and co-workers demonstrated that data expansion strategy with artificial low yields could boost ML predictive performances on proprietary databases.¹³ However, this could introduce biases that have not yet been studied. HTE data display relatively homogeneous yield distribution, with many negative examples. Meanwhile,




Figure 2: Proportion of accessible chemical space observed and reaction diversities for HTE and NiCOLit data after preprocessing (see Section Data Preprocessing, SI). (A): Multi-Dimensional Scaling (MDS) projection of the three datasets. (B): Distribution of Jaccard distances between reactions. MDS and Jaccard distances were calculated on RXNFP featurizations of the datasets. (C): Proportion of accessible chemical space for HTE and NiCOLit data. Numbers indicated next to each variable corresponds to the number of different choices appearing in the dataset for this variable.

searching the proprietary database Sci-Finderⁿ (Fig. 3C) for reactions matching the NiCOLit chemical reaction space returns 2,203 reactions with a clear bias toward high yields : 60% of them have a yield above 70%. NiCOLit yield distribution lays between HTE and Sci-Finderⁿ data, with a significant amount of zero yields experiments but few reactions in the 20 to 40% yield range. This suggests a reporting bias in published reaction data, and an even stronger bias in proprietary databases. The fact that published reaction data contains significantly more negative examples than proprietary databases raises hope that such data could prove much more useful for machine learning model building.

Despite similarities in the reported yield distribution of NiCOLit and HTE data, the underlying structure of the reaction data drastically differs. In HTE, all possible combinations of reactants and reaction conditions are explored (Fig. 3A). Due to time and cost constraints, chemists tend to perform a sparser exploration of the chemical space to achieve a faster convergence. Indeed, published reaction data is reported in two categories of tables or schemes: optimization and scope. Optimization refers to the reaction conditions meaning that most variables except substrate and coupling partner are modified in order to achieve an efficient reaction (vertical dots arrays Fig. 3B). In a complementary fashion, scope refers to reactions with various substrates and coupling partners under optimized conditions (horizontal cross arrays Fig. 3B). Scope experiments are performed in order to demonstrate the robustness of the reaction. In the case of NiCOLit, we noticed that yield distribution of optimization data is similar to the HTE yield distribution and that scope data displays a distribution reminding that of Sci-Finderⁿ (Fig. 3C-D). Exploiting optimization tables during data extraction allows to bypass the lack of low yields reactions in proprietary datasets, and could offer improved predictive performances.

Benchmark of Machine Learning Models on NiCOLit

Then, we evaluated how existing methods for yield prediction perform on NiCOLit. As the representation of chemical reactions in a machine readable format is a crucial step in statistical




Figure 3: Analysis of scope-optimization dataset structures and yields distributions. **(A)**: Projection of the Buchwald-Hartwig dataset on the scope-optimization space, showing homogeneous coverage. The second Principal Component (PC) is displayed as the first PC is primarily driven by the 3 bases present in the dataset. **(B)**: Projection of NiCOLit on the scope-optimization space, showing a biased exploration. **(C)**: Yields distribution for the NiCOLit, HTE, and Sci-Finderⁿ data. Bias towards high yields is observed on the NiCOLit and especially the Sci-Finderⁿ datasets. **(D)**: Yields distribution for the scope and optimization data on NiCOLit.

modeling of reaction yields,²⁵ we selected four approaches representative of the state-of-the-art (see section III of SI). The first approach, RDKit FingerPrint (RDKit FP), is based on the RDKit’s (a cheminformatics tool) chemical reaction fingerprints^{26,27} and one-hot encodings of the remaining variables (ligand, additive, solvent and catalyst precursor). RDKit FPs were obtained through the concatenation of RDKit’s difference reaction fingerprint, computed by subtracting the molecular fingerprint of the reactant from the fingerprint of the product, and the one-hot encodings. The second approach, referred to as Density Functional Theory (DFT), follows the guidelines given by the Auto-QChem framework²⁸ to generate DFT-based molecular descriptors adapted to the selected reaction. For the third approach, RXNFP, we featurized chemical reactions using the deep-learning RXNFP method.²⁹ Eventually, the reactions were featurized using the differential reaction fingerprint (DRFP) which was shown to perform better than DFT-based featurization in reaction yield prediction on the HTE Buchwald-Hartwig dataset.³⁰ A preprocessing that removes successively reactions from reviews (190), reactions in 2 steps (76), reactions that could not be featurized with the DFT methods (205, see Table S6) and those that belong to publications with less than 20 reactions after previous steps (126) was applied, giving 1406 unique reactions. Yield prediction models were built for each featurization with Random Forest regression models³¹ that have shown excellent results on reaction yield prediction⁶ and outperform other machine learning approaches (i.e. support vector machines and neural networks, see Figure S21) on NiCOLit. All metrics reported are averaged over 10 experiments with different random seeds. The DFT method outperforms the three others, and reaches an R^2 of 0.54, even though the difference with the RDKit FP and DRFP is modest (R^2 of 0.50), while RXNFP showed the weakest performance (R^2 of 0.39) (Fig. S6). As expected, the performance of the model trained on optimization data performed better (R^2 of 0.48) than when trained on scope data (R^2 of 0.36) (Fig. S16).

The predictive performance on NiCOLit turns out to be far better than the one reported on the highly heterogeneous USPTO dataset ($R^2 < 0.2$)⁷ and on data extracted from

AstraZeneca’s Electronic Lab Notebooks ($R^2 < 0.3$).²² This shows the potential of machine learning applied to published reaction data. Performances remain nonetheless lower than reported on the HTE data ($R^2 > 0.8$). An explanation could be the highly biased structure of published reaction data described in Fig. 3B, while HTE data cover a narrow and homogeneous chemical space (Fig. 2A and 1A-B) and are devoid of experimental and reporting bias.^{13,24,32,33} Unlike published reaction data, HTE systematically reports yields for all reactions including low yields, and is comprised of reactions performed in the same experimental settings. This makes them a perfect case study for statistical learning compared to NiCOLit, at the cost of exploring a narrower chemical space. It is to be noted that while DFT outperforms other methods, it fails to featurize some reactions. There is thus a tradeoff between the accuracy of the DFT method and its ability to process every reaction. For some applications, other featurizations that show slightly lower accuracy but are faster to compute and more reliable could be more adapted. Nonetheless, the rest of the manuscript focuses on the results obtained with the DFT model.

Analysis of ML performance on out-of-sample predictions

Previous work on reaction yield prediction^{7,10} focused mainly on predictive yields on random splits of the data. However, the nature of scientific discovery pushes chemists to constantly explore new reaction chemical space. Thus, the reactions for which we want to make yield prediction are not sampled from a static distribution, but undergo continuous distribution shift.³⁴ For instance, chemists are often interested in reactions including a novel substrate or a novel coupling partner (Fig. 4). Therefore, validation on a random split is not necessarily informative of how a model would perform when used by chemists in a prospective fashion. This problem was underlined in recent publications,^{7,10,35} where machine learning algorithms showed far worse predictive performances when applied to out-of-sample data (e.g. on reactions with an additive not seen in the training set). While reported models performed very well on random splits (with an R^2 above 0.9, see Table S1), these performances dropped

significantly on some out-of-samples tests, with coefficient of correlations R^2 at best of 0.54 (obtained with a DFT model⁶).

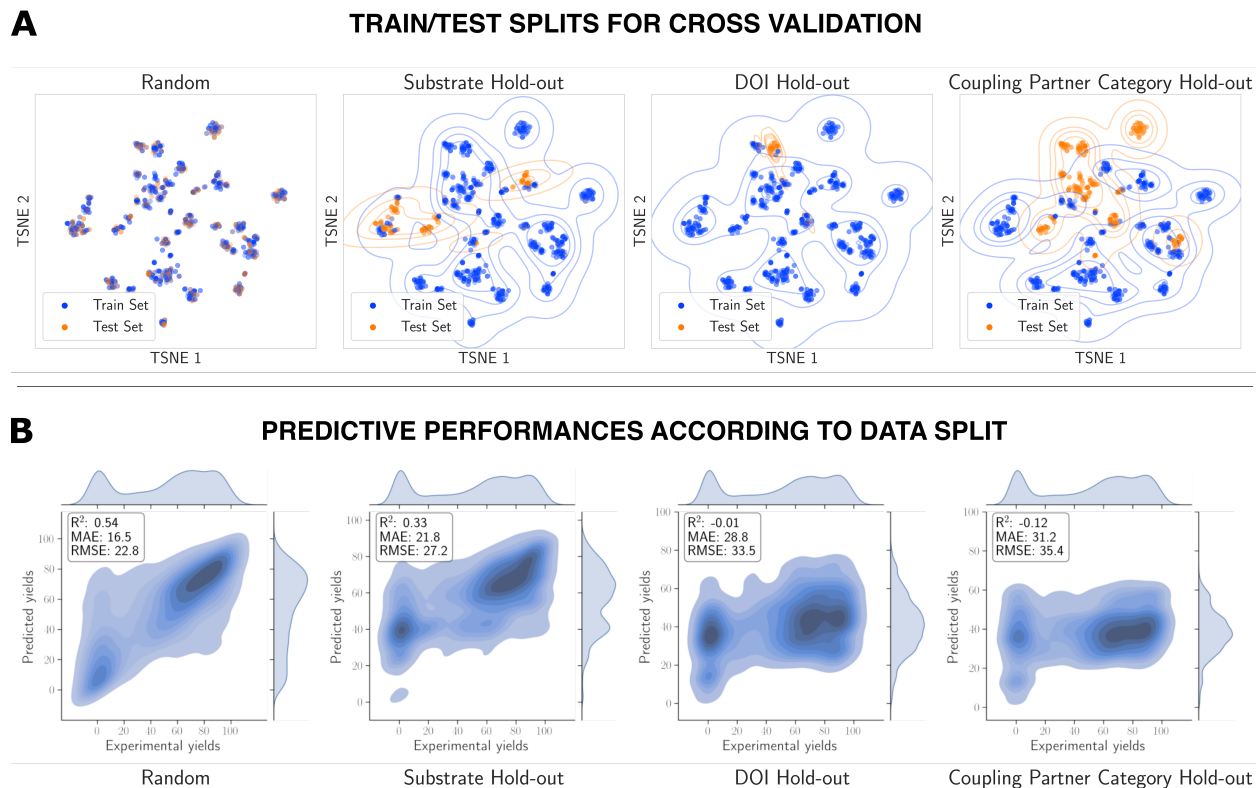


Figure 4: (A): Examples of train-test splits from left to right: random split, split according to substrate (one substrate in the test set and all others in the training set), split according to publication (one Digital Object Identifier (DOI) is taken as the test set and all others as training set) and split according to coupling partner category (all reactions of one coupling partner category are taken as test set and all other coupling partners as training set). (B) DFT model performances for the different splits displayed in A. The performances displayed represent an average result over 10 random splits for the random task and all the possible substrates, DOI or coupling partner categories splits for the three remaining tasks.

We chose the prediction of reaction yield for an unseen yields substrate (Fig. 4) as a task of practical interest. All reactions that feature a given substrate were held-out; after training on the rest of the dataset, the model predicts the yields of the held-out reactions. Those results were aggregated over all substrates in the dataset. While the DFT model showed

encouraging results on the substrate split task, the question of whether the reported predictive performance ($R^2 = 0.33$) is of practical interest is not clear. Therefore, we designed a realistic classification task, where the model classifies reactions using an unseen substrate in two classes, high yields (> 50) and low yields (≤ 50). This use case corresponds to the situation where a chemist wants to explore reactions with a new substrate, and relies on the model’s prediction to discard low yield substrates, and to prioritize efficient ones. On this task, the DFT method reached high predictive performance (with a ROC-AUC, a performance metric for classifiers, of 0.74, see Fig. S18). This highlights a practical application of yield prediction models that can be achieved with existing methods.

Researchers have incentives to explore novel chemical space (Fig. 1C). Moreover, for each publication, reaction yield is biased by the chemist’s skills and the way it is measured. This leads to a high heterogeneity between reactions from different publications. We evaluated how ML predicts yields on data from a new publication. A train-test split of the data, where the test set is comprised of all reactions from one publication, and the train set of all other reactions that do not appear in this publication, is used to assess the yield prediction performance of a model on a new publication (Fig. 4A - DOI Hold-out). Our results, computed over every different publication in NiCOLit, showed the inability (R^2 of -0.01) of ML to generalize to data from new publications.

While the reactions in NiCOLit are all extracted from publications referenced in the same review, they cover a wide range of possible mechanisms. As most of the publications extracted do not provide detailed mechanistic study of the reaction performed, a discrimination was made according to the nature of the coupling partner and substrate (as depicted Fig. 1B). All reactions of the same coupling partner category were held-out (e.g. boronic derivatives, see Fig. 1C). Models were trained on the rest of the dataset, and used to predict the yields of the held-out reactions (Fig. 4A - Coupling Partner Hold-out). Predictive performance, computed over every different coupling partner in NiCOLit, reported for these experiments indicates whether the model is able to predict yields on coupling reactions using a different category of

partners than the reactions of the training set. Again, the models fail to extrapolate to new coupling partners categories.

From a mechanistic point of view, the failure of the models to extrapolate is not very surprising as nickel catalysts can react through very different mechanistic pathways.³⁶ In this perspective, we analyzed the importance of the different features used by the models to perform predictions on more focused subsets of the NiCOLit. All reactions belonging to the same coupling partner category were regrouped and a predictive model was trained on each of these subsets. Then the importance of the features used by the models to perform the predictions were analyzed and displayed Fig. 5. The differences between the models are striking. As an example, the model trained on the RMgX coupling partner is highly dependent on the ligand, the substrate and the molar ratios used while the couplings with RLi mainly depend on substrate characteristics. Nevertheless, these features importance cannot be understood in the light of chemical rules (a more detailed analysis of the Si and Zn cases is available in the SI). Considering that widely different features are used to predict yields on the different coupling partner categories, the poor generalizability of the models is not surprising.




Figure 5: Analysis of feature importance in predictive models trained on coupling partner category NiCOLit subsets. Each subset contains all the reactions belonging to a category of coupling partners.

Chemist’s Expertise Enables ML Predictions in a Low-Data Regime

Based on those results, we hypothesize that models trained on a restricted dataset of reactions sharing similar coupling partners would lead to equivalent predictive performances than models trained on the entire NiCOLit dataset. If true, this would give a precious guideline when gathering published reaction data in order to perform yield prediction.

Table 1: Performances of DFT and baselines models on each coupling partner and substrate categories subsets.^a

Coupling partner	R^2		R^2 k-NN baseline	R^2 scope/optimization baseline	Number of reactions	Number of publications	Accessible reactions
	R^2	restricted					
B	0.47	0.45	0.17	0.26	472	11	1×10^8
C-H	0.59	0.56	0.31	0.11	271	3	3×10^7
RMgX	0.51	0.48	0.23	0.18	266	5	6×10^5
CO ₂	0.52	0.51	0.12	0.36	87	1	7×10^3
Zn	0.54	0.57	-0.16	0.09	68	2	1×10^3
NCO	0.39	0.30	-0.21	0.18	57	1	2×10^4
Al	0.20	0.18	-0.35	0.26	53	1	1×10^4
Si	0.64	0.57	0.10	0.42	53	1	3×10^4
Li	-0.13	0.05	-0.29	-0.03	52	1	2×10^4
Amines	0.17	-0.05	-0.58	0.32	27	1	5×10^2

Substrate	R^2		R^2 k-NN baseline	R^2 scope/optimization baseline	Number of reactions	Number of publications	Accessible reactions
	R^2	restricted					
OR	0.55	0.57	0.3	0.13	546	11	5.4×10^9
OPiv	0.55	0.56	0.31	0.33	394	12	3.2×10^9
OC(=O)N	0.41	0.35	-0.06	0.24	215	14	1.4×10^8
OC(=O)O	0.66	0.64	0.44	0.15	82	4	4.5×10^5
OAc	0.36	0.40	-0.15	0.31	72	7	4.2×10^6
Otriazine	0.42	0.53	0.01	0.39	54	1	1.8×10^5
OSiR ₃	-0.25	-0.25	-1.77	0.00	23	5	1.0×10^4
OCOR	0.06	-0.39	-0.74	0.31	17	4	3.2×10^3
OPh	-0.30	-1.34	-0.41	-0.01	3	3	4.9×10^2

^a R^2 corresponds to the predictions of a model trained on 80% of NiCOLit for a specific substrate or coupling partner category. R^2 restricted corresponds to the performances on the same target for a model trained on 80% of NiCOLit restricted to the targeted substrate or coupling partner. k-NN baseline is a nearest neighbor search based on DFT features and scope/optimization baseline returns the average yield of scope or optimization data according to whether the reaction is in a scope or optimization table. Metric reported is Pearson’s coefficient of correlation R^2 . The best values (with regard to predictive performance) are reported in bold.

To test this hypothesis, we trained the models on NiCOLit restricted to a given category of coupling partner (e.g. all reactions with a boron-based coupling partner) or substrate, and

compared the results with those obtained with a model trained on the full dataset (similar results for NiCOLit restricted to single DOIs are available Table S13). Indeed, for most of the coupling partners or substrates, the model trained on the restricted dataset performs as well as the model trained on the full data (Table 1).

We also compared the predictive performance of the model trained on random splits of the full NiCOLit and on NiCOLit restricted to a specific coupling partner or substrate categories with two baselines. When confronted with a novel reaction with unknown yield, the Nearest Neighbor baseline (k-NN baseline) searches for the closest reaction in the database with a known yield and return its value as the yield prediction for the query reaction. This baseline assesses the performances of one of the natural strategy to extrapolate an unseen yield from the closest known reaction, that mimics the way a human chemist could proceed. The second baseline (scope/optimization baseline) returns the mean yield of reactions extracted from scope (resp. optimization) tables in the database if the query reaction is itself extracted from a scope (resp. optimization) table. This baseline accounts for the fact that a stratification by scope/optimization explains a part of the variability in yields observed within the database. As recent work suggests machine learning doesn’t outperform simple baselines in related prediction tasks,⁵ our goal was to make sure that there was clear added value from using the machine learning approach. For substrates or coupling partners types with more than 70 data points, the machine learning approach systematically outperformed both baselines.

Furthermore, the performance is highly variable according to coupling partner or substrate categories. The most straightforward explanation for this behavior is the disparity between the number of reactions documented for each coupling partner category. The models exhibit poor or modest performances for coupling partners with less than 70 reactions reported. Adequate coverage of chemical space is crucial for building ML models.³⁷ Those results show that a dataset of a much smaller size than NiCOLit can be used to build a predictive model, provided that all reactions belong to the same coupling partner or substrate category (Fig. 6). This study also shed light on the approximate number of reactions needed to reach satisfying




Figure 6: ML performances when trained on full or restricted chemical space. The restricted dataset represented correspond to the boronic derivatives category of the coupling partners. Similar results are obtained with other coupling partner categories.

predictive performance (roughly one hundred reactions). We compared the performances on the restricted sets with between 70-500 reactions and retrieved comparable performances ($R^2 \approx 0.5$) than what was obtained on Buchwald HTE data with a similar number of training points ($R^2 = 0.59$ for 98 data points).¹⁰ Interestingly, good performances can be attained in low-data regimes by injecting domain specific knowledge. Furthermore, for similar predictive performances, most accessible reaction spaces (last column, Table 1) are larger than their HTE counterparts.

Conclusion

To reach its full potential, machine learning relies on high quality data. In the future, we expect that initiatives such as the Open Reaction Database¹² will provide the community with the data needed. In the meantime, there is currently a lack of public unbiased datasets of chemical reactions with detailed experimental conditions. By releasing NiCOLit, we hope to foster the development of impacting data-driven approaches for yield prediction. Our

results highlight the specificities of published reaction data. Comparison with HTE data shows a reporting bias towards high yields. Nonetheless, this bias is weaker in NiCOLit than in its counterpart extracted from proprietary database. Our analysis reveals that it is the extraction of optimization tables that allows the presence of negative examples and limits this reporting bias. In the light of recent results,¹³ this could lead to a significant improvement of machine learning models trained on published data. We also highlight a clear experiment selection bias, where reactions are explored along the two orthogonal optimization and scope directions. A benchmark of several machine learning approaches confirms those observations. Indeed, predictive performance on NiCOLit lies between the one obtained on HTE data and on other available data sources of chemical reactions (proprietary databases or USPTO). A key step in evaluating predictive models remains experimental validation of the model’s predictions, a practice that we highly encourage even though it is outside the scope of this work. Nonetheless, we designed several experiments to test the limitations of machine learning models on out-of-sample prediction tasks. While predictive models cannot currently extrapolate to new reactions or coupling partner categories, we showed that the model is able to generalize on reactions with new substrates. Leveraging this observation, we explain how domain-specific chemistry knowledge can help synthetic chemists select relevant data points to build models in low-data regimes. With only a hundred data points, building predictive models is within reach. We hope that these findings and our open-access database will stimulate the adoption of machine learning for yield predictions in the chemistry community, and encourage systematic data sharing in machine readable formats.

Acknowledgement

The authors thank Maxime R. Vitale, Marc Bianciotto and Hervé Minoux for their thoughtful remarks and constructive feedback.

Funding: CNRS and ENS are gratefully acknowledged for supporting J.S., L.G. and

R.V. The French National Association of Research and Technology (ANRT) is gratefully acknowledged for supporting M.L. (contract 2019/0821). M.L. is employed by Sanofi.

Authors contributions: M.L. and J.S. equally contributed to this project.

Competing interests: M.L. is a Sanofi employee and may hold shares and/or stock options in the company. J.S., B.W., Y.S., R.V. and L.G. declares that they have no competing interests.

Supporting Information Available

Data and materials availability:

Additional details, materials, and methods are provided in the Supporting Information.

All code and data used to produce the reported results can be found online at:

<https://github.com/truejulosdu13/NiCOLit>.

References

- (1) Brogi, S.; Ramalho, T. C.; Kuca, K.; Medina-Franco, J. L.; Valko, M. Editorial: In silico Methods for Drug Design and Discovery. *Front. Chem.* **2020**, *8*, 612.
- (2) Gallego, V.; Naveiro, R.; Roca, C.; Insua, D. R.; Campillo, N. E. AI in drug development: a multidisciplinary perspective. *Mol. Diversity* **2021**, *25*, 1461–1479.
- (3) Patel, L.; Shukla, T.; Huang, X.; Ussery, D. W.; Wang, S. Machine Learning Methods in Drug Discovery. *Molecules* **2020**, *25*, 5277.
- (4) Coley, C. W.; Green, W. H.; Jensen, K. F. Machine Learning in Computer-Aided Synthesis Planning. *Acc. Chem. Res.* **2018**, *51*, 1281–1289.
- (5) Beker, W.; Roszak, R.; Wołos, A.; Angello, N. H.; Rathore, V.; Burke, M. D.; Grzybowski, B. A. Machine Learning May Sometimes Simply Capture Literature Popularity

- Trends: A Case Study of Heterocyclic Suzuki–Miyaura Coupling. *J. Am. Chem. Soc.* **2022**, *144*, 4819–4827.
- (6) Żurański, A. M.; Martinez Alvarado, J. I.; Shields, B. J.; Doyle, A. G. Predicting Reaction Yields via Supervised Learning. *Acc. Chem. Res.* **2021**, *54*, 1856–1865.
- (7) Schwaller, P.; Vaucher, A. C.; Laino, T.; Reymond, J.-L. Prediction of chemical reaction yields using deep learning. *Mach. learn.: sci. technol.* **2021**, *2*, 015016.
- (8) Lowe, D. M. Extraction of chemical structures and reactions from the literature. Ph.D. thesis, University of Cambridge, **2012**.
- (9) Perera, D.; Tucker, J. W.; Brahmabhatt, S.; Helal, C. J.; Chong, A.; Farrell, W.; Richardson, P.; Sach, N. W. A platform for automated nanomole-scale reaction screening and micromole-scale synthesis in flow. *Science* **2018**, *359*, 429–434.
- (10) Ahneman, D. T.; Estrada, J. G.; Lin, S.; Dreher, S. D.; Doyle, A. G. Predicting reaction performance in C–N cross-coupling using machine learning. *Science* **2018**, *360*, 186–190.
- (11) Pflüger, P. M.; Glorius, F. Molecular Machine Learning: The Future of Synthetic Chemistry? *Angew. Chem., Int. Ed.* **2020**, *59*, 18860–18865.
- (12) Kearnes, S. M.; Maser, M. R.; Wlekliniski, M.; Kast, A.; Doyle, A. G.; Dreher, S. D.; Hawkins, J. M.; Jensen, K. F.; Coley, C. W. The Open Reaction Database. *J. Am. Chem. Soc.* **2021**, *143*, 18820–18826.
- (13) Strieth-Kalthoff, F.; Sandfort, F.; Kühnemund, M.; Schäfer, F. R.; Kuchen, H.; Glorius, F. Machine Learning for Chemical Reactivity: The Importance of Failed Experiments. *Angew. Chem., Int. Ed.* **2022**, *61*, e202204647.
- (14) Higgins, S. G.; Nogiwa-Valdez, A. A.; Stevens, M. M. Considerations for implementing electronic laboratory notebooks in an academic research environment. *Nat. Protoc.* **2022**, *17*, 179–189.

- (15) Jablonka, K. M.; Patiny, L.; Smit, B. Making the collective knowledge of chemistry open and machine actionable. *Nat. Chem.* **2022**, *14*, 365–376.
- (16) Schleinitz, J.; Langevin, M.; Smail, Y.; Wehnert, B. truejulosdu13/NiCOLit: NiCOLit dataset first release. 2022; <https://zenodo.org/record/6380442>, accessed 2022-03-23.
- (17) Diao, H.; Shi, Z.; Liu, F. Biaryl Construction Based on Nickel-Catalyzed C–O Bond Activation. *Synlett* **2021**, *32*, 1494–1512.
- (18) Bajo, S.; Laidlaw, G.; Kennedy, A. R.; Sproules, S.; Nelson, D. J. Correction to Oxidative Addition of Aryl Electrophiles to a Prototypical Nickel(0) Complex: Mechanism and Structure/Reactivity Relationships. *Organometallics* **2017**, *36*, 1880–1880.
- (19) Weininger, D. SMILES-A Language for Molecules and Reactions. *Handbook of Chemoinformatics* **2008**, 80 – 102.
- (20) Tamao, K.; Sumitani, K.; Kumada, M. Selective carbon-carbon bond formation by cross-coupling of Grignard reagents with organic halides. Catalysis by nickel-phosphine complexes. *J. Am. Chem. Soc.* **1972**, *94*, 4374–4376.
- (21) Pereira, A.; Alborno, C.; Trofymchuk, O. S. Data-Driven Analysis of Reactions Catalyzed by [CoCp*(CO)I₂]. *Organometallics* **2022**, *41*, 1158–1166.
- (22) Saebi, M.; Nan, B.; Herr, J.; Wahlers, J.; Guo, Z.; Zurański, A.; Kogej, T.; Norrby, P.-O.; Doyle, A.; Wiest, O.; Chawla, N. On the Use of Real-World Datasets for Reaction Yield Prediction. *ChemRxiv* **2021**, 2x06r, DOI: 10.26434/chemrxiv-2021-2x06r-v3, accessed 2021-10-01.
- (23) Reker, D.; Hoyt, E. A.; Bernardes, G. J.; Rodrigues, T. Adaptive Optimization of Chemical Reactions with Minimal Experimental Information. *Cell Rep. Phys. Sci.* **2020**, *1*, 100247.

- (24) Raccuglia, P.; Elbert, K. C.; Adler, P. D. F.; Falk, C.; Wenny, M. B.; Mollo, A.; Zeller, M.; Friedler, S. A.; Schrier, J.; Norquist, A. J. Machine-learning-assisted materials discovery using failed experiments. *Nature* **2016**, *533*, 73–76.
- (25) Jorner, K.; Tomberg, A.; Bauer, C.; Sköld, C.; Norrby, P.-O. Organic reactivity from mechanism to machine learning. *Nat. Rev. Chem.* **2021**, *5*, 240–255.
- (26) Landrum, G. RDKit: Open-source cheminformatics. 2020; <http://www.rdkit.org>, accessed 2021-03-06.
- (27) Schneider, N.; Lowe, D. M.; Sayle, R. A.; Landrum, G. A. Development of a Novel Fingerprint for Chemical Reactions and Its Application to Large-Scale Reaction Classification and Similarity. *J. Chem. Inf. Model.* **2015**, *55*, 39–53.
- (28) Żurański, A. M.; Wang, J. Y.; Shields, B. J.; Doyle, A. G. Auto-QChem: an automated workflow for the generation and storage of DFT calculations for organic molecules. *React. Chem. Eng.* **2022**, *7*, 1276–1284.
- (29) Schwaller, P.; Probst, D.; Vaucher, A. C.; Nair, V. H.; Kreutter, D.; Laino, T.; Reymond, J.-L. Mapping the space of chemical reactions using attention-based neural networks. *Nat. Mach. Intell.* **2021**, *3*, 144–152.
- (30) Probst, D.; Schwaller, P.; Reymond, J.-L. Reaction classification and yield prediction using the differential reaction fingerprint DRFP. *Digit. Discov.* **2022**, *1*, 91–97.
- (31) Breiman, L. Random Forests. *Mach. Learn.* **2001**, *45*, 5–32.
- (32) Strieth-Kalthoff, F.; Sandfort, F.; Segler, M. H. S.; Glorius, F. Machine learning the ropes: principles, applications and directions in synthetic chemistry. *Chem. Soc. Rev.* **2020**, *49*, 6154–6168.
- (33) Williams, W. L.; Zeng, L.; Gensch, T.; Sigman, M. S.; Doyle, A. G.; Anslyn, E. V. The

- Evolution of Data-Driven Modeling in Organic Chemistry. *ACS Cent. Sci.* **2021**, *7*, 1622–1637.
- (34) Quiñonero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N. D., Eds. *Dataset Shift in Machine Learning*; The MIT Press, **2008**.
- (35) Sandfort, F.; Strieth-Kalthoff, F.; Kühnemund, M.; Beecks, C.; Glorius, F. A Structure-Based Platform for Predicting Chemical Reactivity. *Chem* **2020**, *6*, 1379–1390.
- (36) Chernyshev, V. M.; Ananikov, V. P. Nickel and Palladium Catalysis: Stronger Demand than Ever. *ACS Catal.* **2022**, *12*, 1180–1200.
- (37) Kariofillis, S. K.; Jiang, S.; Żurański, A. M.; Gandhi, S. S.; Alvarado, J. I. M.; Doyle, A. G. Using Data Science To Guide Aryl Bromide Substrate Scope Analysis in a Ni/Photoredox-Catalyzed Cross-Coupling with Acetals as Alcohol-Derived Radical Sources. *J. Am. Chem. Soc.* **2022**, *144*, 1045–1055.

TOC Graphic

