

Faraday Discussions

Accepted Manuscript



This is an Accepted Manuscript, which has been through the Royal Society of Chemistry peer review process and has been accepted for publication.

Accepted Manuscripts are published online shortly after acceptance, before technical editing, formatting and proof reading. Using this free service, authors can make their results available to the community, in citable form, before we publish the edited article. We will replace this Accepted Manuscript with the edited and formatted Advance Article as soon as it is available.

You can find more information about Accepted Manuscripts in the [Information for Authors](#).

Please note that technical editing may introduce minor changes to the text and/or graphics, which may alter content. The journal's standard [Terms & Conditions](#) and the [Ethical guidelines](#) still apply. In no event shall the Royal Society of Chemistry be held responsible for any errors or omissions in this Accepted Manuscript or any consequences arising from the use of any information it contains.

This article can be cited before page numbers have been issued, to do this please use: I. Kevlishvili, R. G. St. Michel, A. Garrison, J. Toney, H. Adamji, H. Jia, Y. Roman-Leshkov and H. Kulik, *Faraday Discuss.*, 2024, DOI: 10.1039/D4FD00087K.

Leveraging natural language processing to curate the tmCAT, tmPHOTO, tmBIO, and tmSCO datasets of functional transition metal complexes

Ilia Kevlishvili¹, Roland G. St. Michel^{1,2}, Aaron G. Garrison¹, Jacob W. Toney¹, Husain Adamji¹, Haojun Jia^{1,3}, Yuriy Román-Leshkov^{1,3}, and Heather J. Kulik^{1,3*}

¹*Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

²*Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

³*Department of Chemistry, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

ABSTRACT: The breadth of transition metal chemical space covered by databases such as the Cambridge Structural Database and the derived computational database tmQM is not conducive to application-specific modeling and the development of structure–property relationships. Here, we employ both supervised and unsupervised natural language processing (NLP) techniques to link experimentally synthesized compounds in the tmQM database to their respective applications. Leveraging NLP models, we curate four distinct datasets: tmCAT for catalysis, tmPHOTO for photophysical activity, tmBIO for biological relevance, and tmSCO for magnetism. Analyzing the chemical substructures within each dataset reveals common chemical motifs in each of the designated applications. We then use these common chemical structures to augment our initial datasets for each application, yielding a total of 21,631 compounds in tmCAT, 4,599 in tmPHOTO, 2,782 in tmBIO, and 983 in tmSCO. These datasets are expected to accelerate the more targeted computational screening and development of refined structure–property relationships with machine learning.



1. Introduction.

Transition metal complexes (TMCs) are used in a wide variety of applications, including as homogeneous catalysts¹⁻³ for fine chemical synthesis³⁻⁶ and in advanced sensor, display, and data storage devices.^{3,7} The electronic and steric properties a TMC are influenced by those of the constituent metal and ligands.^{7,8} However, the variety of metal identities and oxidation states that may be used in conjunction with ligands of different connectivity and charge results in a combinatorial design space too large to sample exhaustively.^{3,7,9,10} Density functional theory (DFT) calculations are often leveraged in high-throughput virtual screening (HTVS) campaigns to explore chemical space in search of new molecules with desired properties, though the cost of such calculations limits the number of complexes that may be investigated.^{1,3,7,9,11,12} Such campaigns can be further accelerated through machine learning, which rely on large datasets of experimental and computational results for training.^{7,9,11}

Prior efforts have been made to curate datasets of TMCs,¹³⁻¹⁶ their constituent ligands,¹⁷⁻²³ and relevant reactions.^{4,24-27} Many of these datasets are based on entries from the Cambridge Structural Database (CSD),²⁸ a digital repository of experimental crystal structures, including molecular crystals of thousands of TMCs. However, challenges exist with current datasets, which primarily fall into one of two classes. The first class of datasets are exceptionally large but contain properties of limited relevance to applications-based molecular discovery.^{10,13,24,29-36} The second class of datasets are highly focused, containing relevant information very specific to local regions of chemical space, and as such are not easily generalized to new chemical applications.^{14,19-23,27,37,38} The transition metal quantum mechanics (tmQM) dataset is an example of a large, nonspecific dataset of interest to transition metal chemistry, containing 86,665 mononuclear TMCs.¹⁶ These structures were extracted from the CSD and subjected to additional filtering by retaining only



structures with at least one C and H atom, and only those which contain allowed non-metal elements (i.e., B, Si, N, P, As, O, S, Se, F, Cl, Br, and I). Furthermore, oxidation states were assumed for metals to ensure closed-shell character where possible and only structures that had a net charge of no more than +1 or less than -1 were retained. Their geometries were optimized at the semiempirical extended tight binding (xTB) level of theory, and DFT single-point energy calculations were performed on the resulting structures. While the tmQM is a valuable dataset in computational chemistry workflows for investigating TMCs,^{17,39,40} a key limitation is the absence of a mapping between molecular structures and the relevant areas of chemistry. This hinders further investigation into structures that are particularly promising for applications in catalysis, photochemistry, or other fields of interest. In contrast, datasets such as the ligand knowledge base (LKB) curated in pioneering work by Fey et al.¹⁹⁻²² and *kraken* later developed by Gensch et al.²³ are examples detailed, applications-focused datasets with limited transferability. Both datasets primarily consist primarily of organophosphorus ligands, include relevant physicochemical descriptors useful in building quantitative structure–property relationships, and are based on commercial and virtual libraries.^{19-23,41} However, these do not generalize well to other areas of chemistry beyond organophosphorus ligands, exemplifying a limitation of such datasets and a need for large datasets linked to targeted chemical applications.

The curation of a chemically targeted and synthetically accessible TMC dataset relies on systematically reviewing literature on the TMCs that are contained within existing databases. The broad scope of TMC literature, however, would make manual processing arduous, prompting the use of natural language processing (NLP) techniques^{42,43} for efficient analysis. NLP has been utilized extensively in the extraction of material properties and material synthesis parameters from the literature.^{37,44-51} More recently, large language models (LLMs) coupled with prompt



engineering have gained increasing popularity in automating scientific text mining for chemical information due to their more user-friendly nature.⁵²⁻⁵⁶ A crucial aspect of text mining for classifying text based on chemical domain involves topic modeling,⁵⁷ which is the identification of underlying themes in large sets of scientific text. For tasks of this nature, prompt engineering typically requires *a priori* definition of the possible latent topics. Nevertheless, LLMs can still be leveraged to obtain contextualized embeddings of the text that capture semantic information^{58,59} and subsequently cluster text based on semantic similarity^{60,61} where each cluster corresponds to a latent topic, as facilitated by algorithms like bidirectional encoder representations from transformers for topic modeling (BERTopic).⁶² Simpler topic modeling approaches, such as latent Dirichlet allocation (LDA),⁶³ which utilizes bag-of-words and statistical patterns of co-occurring words to infer latent topics, can also be employed to cluster manuscripts in a corpus. While these unsupervised NLP methods have been leveraged for summarizing research trends in chemistry,⁶⁴ with an emphasis on biochemical and medicinal research⁶⁵⁻⁶⁸ as well as in the classification of large biomolecular datasets,⁶⁹ they have yet to be extended to the space of transition metal chemistry and in the development of application-specific TMC datasets.

To construct chemically targeted TMC datasets, we conduct text mining on manuscripts associated with synthesizable TMCs from the tmQM database, focusing only on their titles and abstracts, and leverage both simple NLP tools as well as transformer models to process the text. Using topic modeling, we segment the structures in the tmQM database based on distinct chemistry applications. Through this process, we introduce four new TMC datasets – tmCAT containing catalytically-relevant TMCs, tmPHOTO with photoactive TMCs, tmSCO comprising TMCs with magnetic properties, and tmBIO containing biologically-relevant TMCs. Additionally, we performed substructure analysis to compare trends in metal-local structures among tmQM TMCs



and the four curated datasets to subsequently enrich each chemistry-specific dataset by adding additional tmQM TMCs that could potentially be suitable for the given application.

2. Computational Details.

Corpus curation and text pre-processing

Manuscript titles and abstracts used for supervised learning as well as unsupervised clustering were obtained from a corpus that was curated in November 2020.^{70,71} Most manuscripts in this corpus were retrieved directly using the ArticleDownloader package.⁷² Manuscripts from the Royal Society of Chemistry (RSC), Wiley-VCH, the American Association for the Advancement of Science (AAAS), Springer, and Nature were obtained directly. Articles from the American Chemical Society (ACS) were obtained via a direct download agreement between ACS and the Massachusetts Institute of Technology. We created a secondary corpus of abstracts for manuscripts that could not be obtained with the ArticleDownloader package by scraping article URLs using the BeautifulSoup package v.4.12.2⁷³. We used the HTML title to retrieve manuscript titles. To obtain the manuscript abstract, we parsed the HTML paragraphs and retrieved the first paragraph that contained more than 400 characters. From a set of 100 randomly selected DOIs, we manually validated the abstract and title retrieval procedure, which shows that this approach can be used to retrieve titles and abstracts at a high rate (Supporting Information Table S1). However, this procedure can in rare cases (i.e., 1 case out 100 tested) lead to retrieval of an introduction paragraph instead of the abstract. Because the secondary corpus contains some degree of impurity, it was only used to identify more catalysis-relevant complexes based on the abstract with a trained classifier model (see Sec. 3a).



Abstracts were preprocessed using the NLTK v.3.8.1 package⁷⁴. The text was cleaned by lowercasing, removing punctuation, URLs, and numbers. The cleaned text was then tokenized using a regular-expression tokenizer, RegexpTokenizer, implemented in the NLTK package. Tokenized text was filtered using stop words with standard English stop words. For unsupervised clustering, an additional set of stop words was introduced to avoid clustering based on chemical languages (See Sec. 3c). The filtered text was stemmed with the Snowball Stemmer and lemmatized with the WordNet Lemmatizer, both implemented in the NLTK package.

Featurization

For the classification model, we featurized the corpus using the term-frequency inverse-document-frequency (TF-IDF) vectorizer implemented in scikit-learn v.1.4.0 package⁷⁵. The TF-IDF vectorizer accounts for the frequency of a given token within a document and its frequency in a collection of documents, assigning lower weight to common tokens across the entire corpus. Only tokens that appeared in at least 10 documents were retained, and the vector included mono-, bi-, and trigrams. The feature vector was fit using only the training set to avoid data leakage from inverse document frequency weighting. To reduce the feature vector length, we evaluated the χ^2 score of each feature using the training set and retained only the 300 most important features as computed by the χ^2 test. For unsupervised learning using BERTopic, we used the Sentence Transformers v2.2.2⁷⁶ package for transforming abstracts into a feature vector. The semantic embedding was done using the sentence transformer sentence-BERT (SBERT) model⁷⁶, which converted the title and abstract into a 768-dimensional vector. Embedding was done using a pre-trained Siamese BERT network, all-mpnet-base-v2 transformer, which embeds paragraphs to a 768-dimensional vector. Embeddings can be compared using cosine similarity. We then reduced the dimensionality of this vector to a five-dimensional mapping using uniform manifold



approximation and projection (UMAP)⁷⁷, which was selected because dimensionality can be reduced in UMAP using the same cosine metric as in the SBERT embedding. We generated a CountVectorizer feature vector of the corpus for unsupervised learning with latent Dirichlet allocation (LDA)⁷⁸. LDA utilizes a bag of words vector, which consists of the overall token count for each document. Furthermore, LDA requires a predefined number of clusters. We decomposed the corpus into 20 clusters to maintain relative consistency with the 23-25-cluster size identified by BERTopic. The count vector was generated using scikit-learn and consists of a vector of term count per document.

While BERT-based models can, in principle, be applied without text preprocessing, this in practice leads to clustering by transition metal or material (Supporting Information Table S2). To avoid this, we introduced stop words before text processing to avoid dependence on specific materials or metals, with a full list of stop words provided on Zenodo⁷⁹.

NLP models

The catalysis classifier is a random forest classifier implemented in the scikit-learn v.1.4.0 package. A random 80:20 train/test split was used to evaluate the model performance on a set-aside test set. Hyperparameter optimization does not significantly affect the model performance. Full grid search cross-validation summary can be accessed through Zenodo repository⁷⁹. Default random forest hyperparameters were used for training without hyperparameter optimization, with the exception of the minimum samples required to split a node that was set to 10, and the number of trees that was set to 1,000 (Supporting Information Table S3). The minimum sample split was increased from the default to avoid overfitting, and the number of models in the ensemble was increased to improve model performance. Dimensionality reduction of high-dimensional feature



vectors for unsupervised clustering (5-dimensional vector) and visualization (2-dimensional vector) were carried out using uniform manifold approximation and projection for dimension reduction (UMAP) using UMAP package v0.5.5⁷⁷. We use the 5-dimensional vector to identify dense clusters, and we interpret topic assignments through cluster-level TF-IDF vectors. Clustering on the reduced 5-dimension vector was carried out using HDBSCAN⁸⁰, a hierarchical density-based clustering algorithm using the HDBSCAN v0.8.33 package. The topic assignment of dense clusters was achieved using the BERTopic package v0.16.0 with a modified class-based TF-IDF vector (c-TF-IDF)⁶². Clustering using LDA was carried out with the implementation in the scikit-learn package.

All machine learning models, scripts, Jupyter notebooks, datasets, and their associated structures are provided on Zenodo⁷⁹. Geometry assignment was carried out using functionality implemented in the molSimplify geometry_changes branch, and the script is included in the Zenodo repository⁷⁹.

3. Results and Discussion.

3a. Catalytic transition metal dataset

To curate a dataset of catalytically relevant transition metal complexes, we utilized the tmQM dataset⁸¹ consisting of 86,665 unique CSD refcodes as a starting point, and we obtained the manuscripts associated with each CSD refcode using the ArticleDownloader package⁷². Through this procedure, we curated a corpus consisting of 28,394 unique manuscripts, accounting for 50,968 crystals in the dataset. To utilize natural language processing models for identifying catalysis manuscripts, we focused on manuscript abstracts because they are information-dense texts that tend to avoid the discussion of broader topics and are therefore well suited for identifying



whether a manuscript will discuss catalysis (i.e., versus introductions). Furthermore, abstracts are publicly available, typically on the article's DOI-accessible HTML webpage, which enables retrieval of abstracts that could not be obtained using the ArticleDownloader package⁷². If abstracts could not be extracted automatically from the manuscript (here, this occurred for 4,682 of 28,394 manuscripts), we used titles instead.

To train a sentiment analysis model that could identify whether an abstract is associated with catalysis, we first identified the subset of manuscripts related to catalysis based solely on whether the manuscript titles contained the keyword “catal” but does not contain false-positive keywords (e.g., “uncatal”, “acid catal” or “base catal”, Supporting Information Table S3). These steps produced a 4,585-manuscript subset where titles were confidently labeled as addressing catalysis (Supporting Information Table S3). We then confidently identified non-catalytic manuscripts by excluding manuscripts with catalytic and associated keywords (i.e., “catal”, “turnover”, and “polymer”) that occurred at least once in either the title or abstract (Supporting Information Table S4). This non-catalytic set consists of 20,557 manuscripts, the majority of manuscripts not identified as catalytic in our initial step, leaving only 3,252 unlabeled manuscripts. To analyze the label assignment, we randomly sampled 50 catalysis and 50 non-catalysis manuscripts and checked them manually. We find that 48 of the positive and 49 of the negative labels were correctly assigned using our approach.

Despite our efforts to carefully label manuscripts confidently, simple pattern matching leads to a small number of incorrect label assignment and our approach to avoid false positive labels by only looking at patterns in the title is expected to miss true positive hits, motivating a more robust approach for identifying catalyst-focused manuscripts. We next pursued a more systematic approach by developing a classifier model using natural language processing. Prior to



training the classifier, we first created a balanced dataset of catalysis and non-catalysis manuscripts by subsampling the non-catalytic manuscripts to achieve a balanced set of 4,585 non-catalysis manuscripts to match the 4,585 catalysis manuscripts. Among the set of 9,170 manuscripts, 1,312 manuscripts (364 catalysis, 948 non-catalysis) did not have a defined abstract, and the title was instead used for training the classifier model. We then separated this set of 9,170 manuscripts into training and test sets using a stratified random split of 80:20. Using the abstract text of each of these papers, we preprocessed the text by elimination of uppercase letters, removal of punctuation, lemmatization, stemming, and removal of stop words to make the text suitable for natural language processing tasks (see Computational Details). We first featurized the papers based on the term frequency-inverse document frequency (TF-IDF) vectorizer⁸² from the training set (see Computational Details). Subsequently, a random forest classifier model was trained on the reduced TF-IDF feature vectors to predict whether an abstract is related to catalysis, achieving a high accuracy of 0.97 and strong separation between the two classes and ROC-AUC of 0.99 (Figure 1).



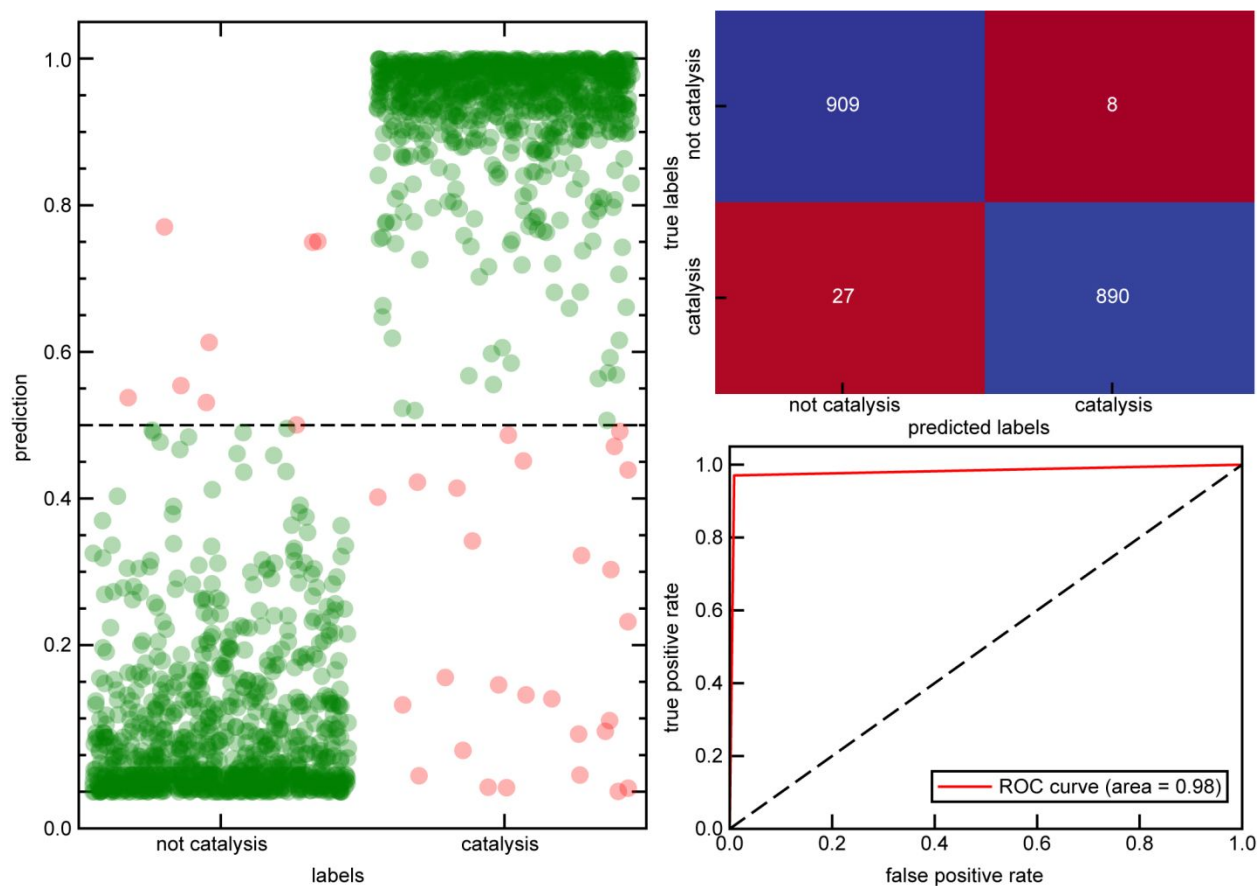


Figure 1. Prediction probability (left), confusion matrix (top right) and receiver operating characteristic curve (bottom right) of the catalysis classifier on the set-aside test set. All data points are represented as translucent circles to depict data density and colored by classification correctness: correct (green) and incorrect (red).

Despite the strong overall performance, we next analyzed the specific cases where the model failed to ensure that it was a suitable tool for confidently assigning a catalysis focus to manuscripts. Out of 1,834 unique abstracts in the test set, 35 were incorrectly labeled. Only eight of these abstracts were given a false-positive label and were manually inspected to prevent dataset contamination. Among these eight abstracts, five had labels that were inaccurately assigned by the rule-based method (i.e., missing the “catal” and additional keywords in the abstract/title) and thus



should have been labeled as catalysis manuscripts. Additionally, one entry lacked an abstract, and its title had been used for training instead, likely contributing to the failure of the model to correctly predict the non-catalytic label for the manuscript. Furthermore, 27 complexes were given false-negative labels. Among these, 14 had labels that were inaccurately assigned by the rule-based method. This analysis reveals that only a negligible fraction of manuscripts were wrongly labeled as catalysis-focused by the classifier model, and our method can detect complexes that the rule-based approach missed. Similarly, the model can detect true positive labels that were missed by the rule-based method. However a larger, but still negligible, number of catalysis manuscripts might be missed by the model. Detailed analysis of all false labels are provided on the Zenodo repository.⁷⁹ Unsurprisingly, the feature importance analysis of TF-IDF vectors showed that the most crucial features are keywords directly related to catalysis. However, several other significant word features were also identified, including activity, polymerization, hydrogenation, coupling, selectivity, Suzuki, and enantioselective, among others (Figure 2). To test the effectiveness of these additional tokens related to catalysis, we developed a separate random forest model that was trained on a TF-IDF feature vector that excluded the direct catalysis keywords (Supporting Information Table S5). This second random forest model still achieved 89% accuracy, demonstrating that other relevant keywords can still effectively identify catalysis-related manuscripts (Supporting Information Table S6 and Figure S1).



to this unlabeled set, we identified 6,208 additional manuscripts in the corpus as associated with catalysis (Supporting Information Figures S2–S3). With this added set, we identify the final tmCAT dataset, which consists of 10,793 unique manuscripts and the structures of 18,915 associated TMCs. While NLP will identify abstracts that are associated with catalysis, it is not necessary that all crystal structures associated with these papers are relevant for catalysis (i.e., they could correspond to non-catalytic cations or catalyst precursors), and further analysis of the chemical composition of the dataset is merited (see Sec. 4).

3b. Comparing catalysis and non-catalysis complexes

We next quantified differences in the composition between catalytically-relevant tmCAT complexes and the broader tmQM dataset. To make this comparison, we first examined the distribution of transition metals in the tmQM and the tmCAT subset. Based on their frequency in the tmCAT dataset relative to the whole tmQM set, group 4 metals (i.e., Ti, Zr, and Hf) are most disproportionately used for catalysis, which can likely be attributed to the large number of papers and structures related to olefin polymerization with homogeneous Ziegler-Natta catalysts in our set (Figure 3). Group 3 metals (i.e., Sc and Y) are also extensively studied for catalysis (Figure 3). These metals are present in compounds that are often used as post-metallocene olefin polymerization catalysts⁸³ and lactone ring-opening polymerization catalysts⁸⁴. These findings indicate the prevalence of polymerization catalysts in the tmCAT subset, which we also explore further with additional text analysis in Sec. 3. Furthermore, Pd, Ru, Ir, Rh, and Au metals are abundant in our set, as would be expected because of their relatively frequent use in catalysis applications.^{85–88} Conversely, more earth-abundant transition metals are less commonly used in catalysis despite recent efforts towards sustainable catalysis using these metals⁸⁹. While iron- and nickel-based catalysts have been studied relatively frequently, cobalt-based catalysts are notably



underrepresented, especially in comparison to isovalent Ir and Rh species that are more represented in the catalysis dataset. This analysis underscores the continued need to advance catalysis toward using more earth-abundant metals.

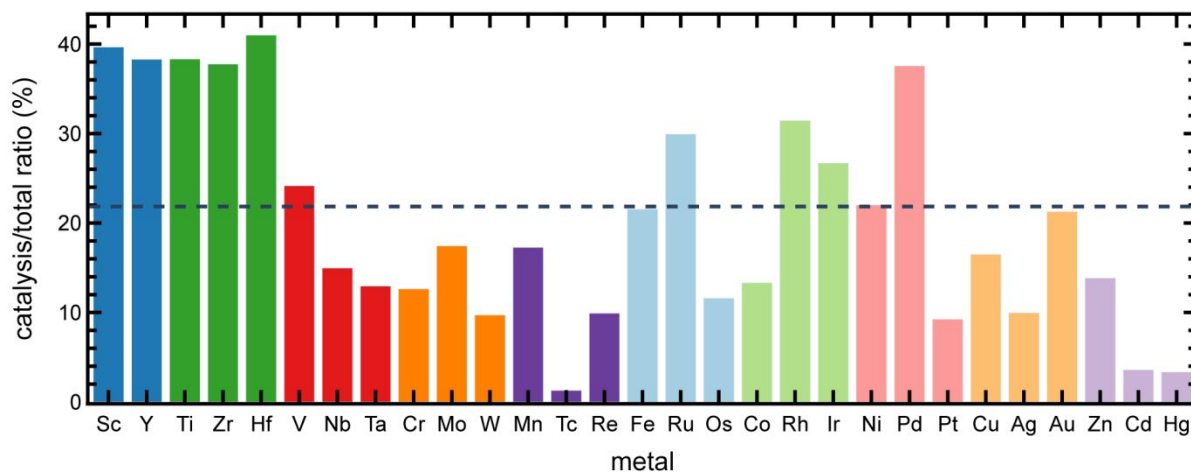


Figure 3. Relative frequency of transition metal complexes associated with the catalysis topic. Metals are grouped and colored by their group number. The ratio of the total size of tmCAT relative to tmQM is shown as a dashed line.

Next, we analyzed the differences and similarities between the tmCAT and the rest of the tmQM dataset (i.e., the non-catalytic portion) in terms of descriptors that had been computed during the curation of the tmQM dataset⁸¹. We first focus on electronic descriptors, such as molecular orbital energetics⁹⁰ and metal charges that are commonly employed in the screening of transition metal catalysts^{90,91}. The relevant descriptors available in the original tmQM set⁸¹ include the HOMO and LUMO energies, the HOMO–LUMO gap, and the transition metal center partial charge. These properties were computed using the TPSSh meta-GGA hybrid functional with empirical D3(BJ) dispersion correction and a def2-SVP basis set based on GFN2-xTB optimized



geometries. Interestingly, the distribution of these descriptors show no major differences between catalytic (tmCAT) subset and the non-catalytic subset of the tmQM dataset (Supporting Information Figures S4–S7). These observations suggest that when considering a diverse set of complexes with wide-ranging transition metals, oxidation states, coordination environments, and ligands, descriptors based on frontier orbitals or metal partial charges may be insufficient for inferring reactivity. These observations are consistent with past work showing that frontier orbital energies alone struggle to generalize to catalyst activities across multiple metals and oxidation states.⁹²

We next considered geometric descriptors that evaluate the steric environment defined by ancillary ligands⁹³ as another commonly employed class of descriptors for catalyst screening. We would expect an active catalyst (i.e., not a precatalyst) to feature an open site that can lead to the association of a reactant to the active site. However, active catalysts with open metal sites are usually not energetically stable intermediates and tend to have a sacrificial ligand or a solvent coordinated to the open site. To compare tmCAT structures to the non-catalytic subset of tmQM, we computed the percent buried volume of the metal for all complexes in the tmQM dataset. Analysis of the distribution of this parameter shows no significant difference in buried volume between the tmCAT and the rest of the tmQM dataset (Supporting Information Figure S8). We anticipate this lack of distinction is attributable to the fact that deposited crystal structures are likely precatalysts that need to undergo an activation process to form an active catalyst, meaning that geometric descriptors on CSD structures are unlikely to be useful for identifying catalysis-capable complexes.

Beyond steric metrics, one might anticipate other differences in the metal-local coordination that might distinguish the tmCAT and non-catalysis tmQM subsets. We hypothesized that even though



pre-catalyst complexes should be heavily featured in tmCAT, some noticeable differences could still be observed between catalysis and non-catalysis datasets when comparing coordination geometries because some geometries are less probable for precatalysts (e.g., those with six monodentate ligands or three bidentate ligands). We assigned metal coordination geometries by examining the geometric deviations from possible ideal transition metal geometries and assigning a geometric class with the lowest deviation (see Computational Details). When a haptic ligand is encountered (e.g., an alkene bound via its π bond to a metal center), a single occupancy was assigned at the geometric centroid of the haptic ligand. The most noticeable difference between tmCAT and the non-catalysis tmQM subset is due to the significant reduction in the number of octahedral complexes in tmCAT accompanied by a significant enhancement in the frequency of square planar complexes (Figure 4). Despite a lack of difference between tmCAT and the non-catalytic tmQM in terms of steric descriptors, this enhancement of square planar over octahedral structures is consistent with our expectation of enhancing coordinatively unsaturated complexes in the tmCAT dataset as well as the fact that more octahedral structures are likely to be incompatible with catalysis due to higher-denticity ligands. Furthermore, the relative frequencies of other coordinatively unsaturated geometries, such as square pyramidal and trigonal planar complexes are also enhanced in the tmCAT dataset.



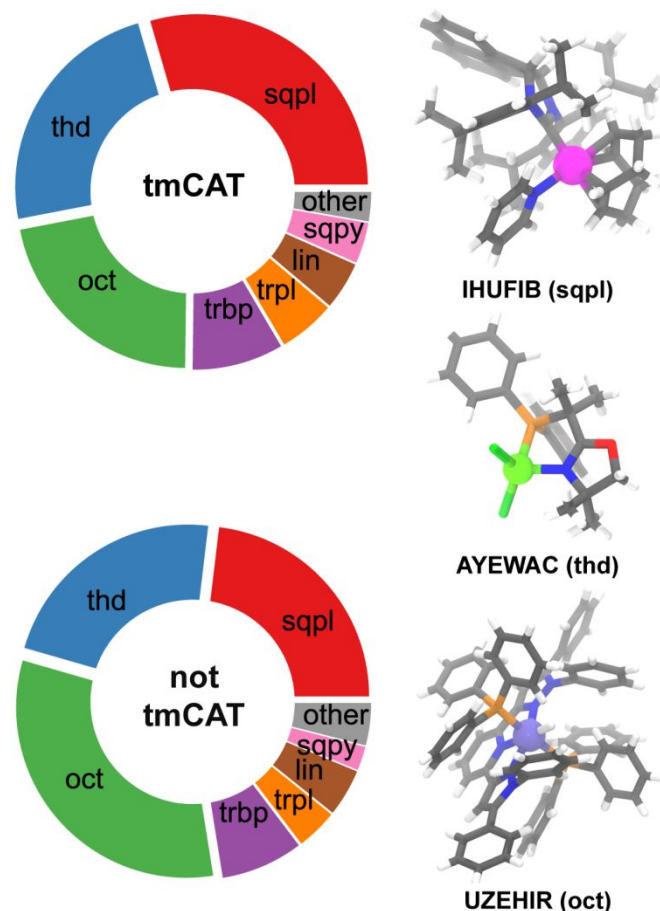
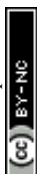


Figure 4. Relative frequency of different transition metal coordination geometries for the catalysis subset (tmCAT) and the rest of the tmQM dataset. Abbreviations: sqpl – square planar, thd – tetrahedral, oct – octahedral, trbp – trigonal bipyramidal, trpl – trigonal planar, lin – linear, sqpy – square pyramidal, other – all other coordination geometries. Randomly selected structures of the three most common geometries in the tmCAT set are shown on the right side. Transition metals are shown as spheres. Iridium is shown in pink, nickel in light green, ruthenium in purple, nitrogen in blue, oxygen in red, phosphorus in orange, chlorine in green, carbon in gray and hydrogen in white.

3c. Unsupervised learning with natural language processing maps CSD structures to other applications



Given that the catalysis classifier only assigns approximately 25% of tmQM complexes to the tmCAT dataset, we next aimed to identify if the remaining 75% could be at least partly assigned to other application areas for TMCs by carrying out a more expansive review of the full corpus of papers in the tmQM set. Defining a topical subset within a corpus using supervised learning necessitates curation of a smaller subset, either manually or by defining heuristic rules for identifying positive and negative examples. On the other hand, by using unsupervised learning on the text, it is possible to identify clusters within a corpus that are semantically similar. Analyzing these clusters can lead to an improved understanding of different possible applications and topics covered by the corpus. A common approach for this purpose in natural language processing is topic modeling⁹⁴, an unsupervised learning approach that can both cluster a corpus and identify topics associated with each cluster. We applied topic modeling to our full corpus of manuscript titles and abstracts to identify latent application topics, offering a more comprehensive view of the underlying themes and connections within the corpus.

We utilized the BERTopic model⁶², a method that clusters bodies of text based on their semantic embedding, and then assigns topics using a modified, cluster-based TF-IDF (c-TF-IDF) vector or other interpreter (see Computational Details). Importantly, we identified several subtopics that could be associated with non-catalytic applications, including biological activity, photoactivity, magnetism, self-assembly, and X-ray crystal structure characterization (Figure 5 and Supporting Information Figure S9). Additionally, many of the topics uncovered by this analysis are complementary to our earlier labeling of catalyst papers, as other uncovered subtopics include several catalysis-related areas with more specific applications including polymerization, hydrogenation, chiral catalysis, and cross-coupling (Figure 6 and Supporting Information Figure S10). Even though the number of subtopic clusters and cluster composition is distinct across



different models (i.e., because there are multiple topic models that differ based on a random seed), all major subtopic clusters we identified are conserved. Using this additional information, we now introduce three additional datasets consisting of complexes associated with manuscripts consistently categorized across all five models as follows: i) tmPHOTO, which consists of photoactivity-associated complexes, ii) tmSCO, which consists of compounds exhibiting properties relevant to studies of magnetism, and iii) tmBIO, which consists of complexes with biologically relevant activities. We identified each of these datasets by inspecting most significant tokens associated with each of the cluster. For example, the tmPHOTO set is derived from the cluster that consists of manuscript abstracts that discuss phosphorescence, emission, and quantum yield, indicating that photophysical activity is discussed throughout these manuscripts. Similarly, we identified that tmSCO manuscript abstracts are associated with tokens that are related to spin-crossover, magnetic properties and hysteresis, all keywords that are related to changes in the spin state of a TMC. Likewise, we determined that manuscript abstracts associated with tmBIO cluster discuss properties such as cytotoxicity, cancer, cell, and apoptosis, which are all related to biological activity relevant to pharmaceutical applications.



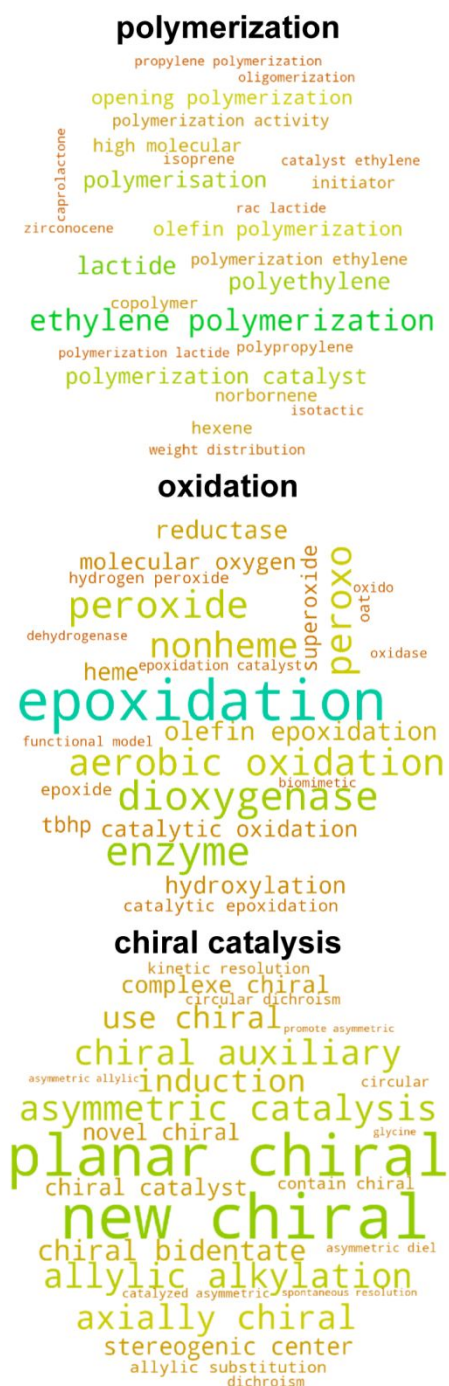


Figure 5. Selection of unique non-catalysis clusters identified using an unsupervised clustering approach showing the 25 most important tokens associated with each cluster based on the c-TF-IDF vector. Text is scaled and colored according to the token importance.



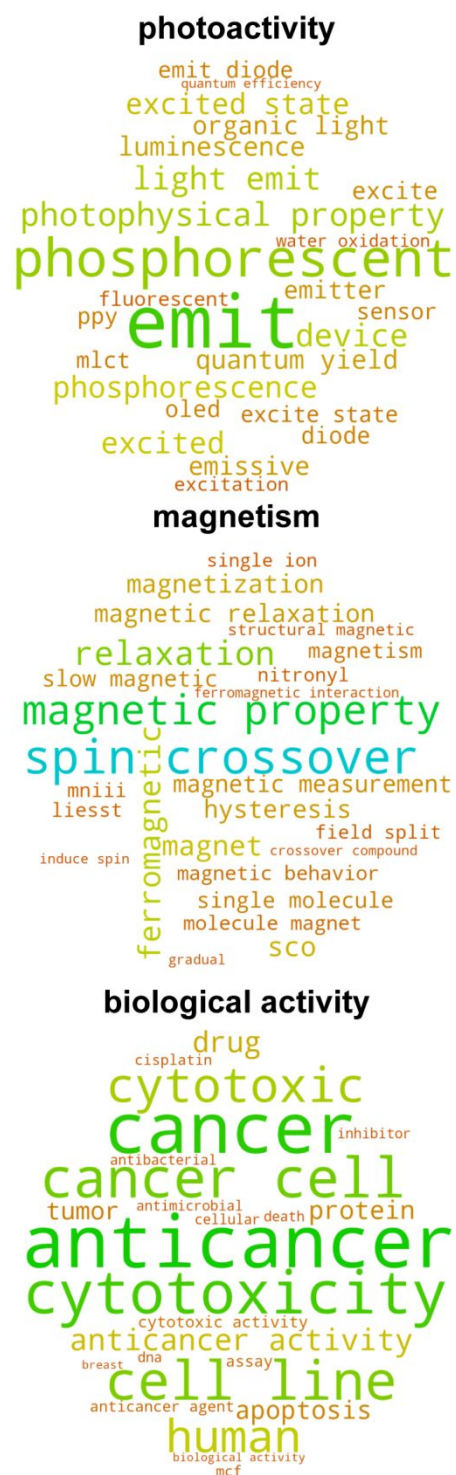


Figure 6. Selection of unique catalysis clusters identified using an unsupervised clustering approach showing the 25 most important tokens associated with each cluster based on the c-TF-IDF vector. Text is scaled and colored according to the token importance.



Based on how semantic embedding works, we expect it to place/arrange subtopics with greater similarity closer to each other in the reduced dimensional space. To analyze the performance of the unsupervised learning approach and visualize how different topics relate to each other, we employed UMAP for further dimensionality reduction on the SBERT embedding, reducing the embeddings to two dimensions better suited for visualization while retaining the global structure of the data for distance comparison. We find that catalysis sub-topics are predominantly clustered closely to each other, suggesting that the wording used throughout catalysis-associated abstracts is highly similar. Furthermore, catalysis topics that are more closely related to each other, such as polymerization and metathesis, or hydrogenation and hydroboration/boration, each of which is related to olefin functionalization, are more closely clustered (Figure 7). As expected, the other, non-catalysis topics are more distant in the UMAP-reduced space (Figure 7). The biologically active cluster, arguably the most different from other applications, due to its relevance in biological and pharmaceutical applications, stands out as the most distinct cluster. Manuscripts related to photoactivity and magnetism are comparatively clustered close to each other, which can be expected because both topics are associated with the transition to an excited state via external stimulus (Figure 7). Alternative dimensionality reduction techniques, such as t-distributed stochastic neighbor embedding (t-SNE)⁹⁵, lead to similar conclusions, although we avoided using t-SNE because it is less effective at preserving the global data structure (Supporting Information Figure S11).



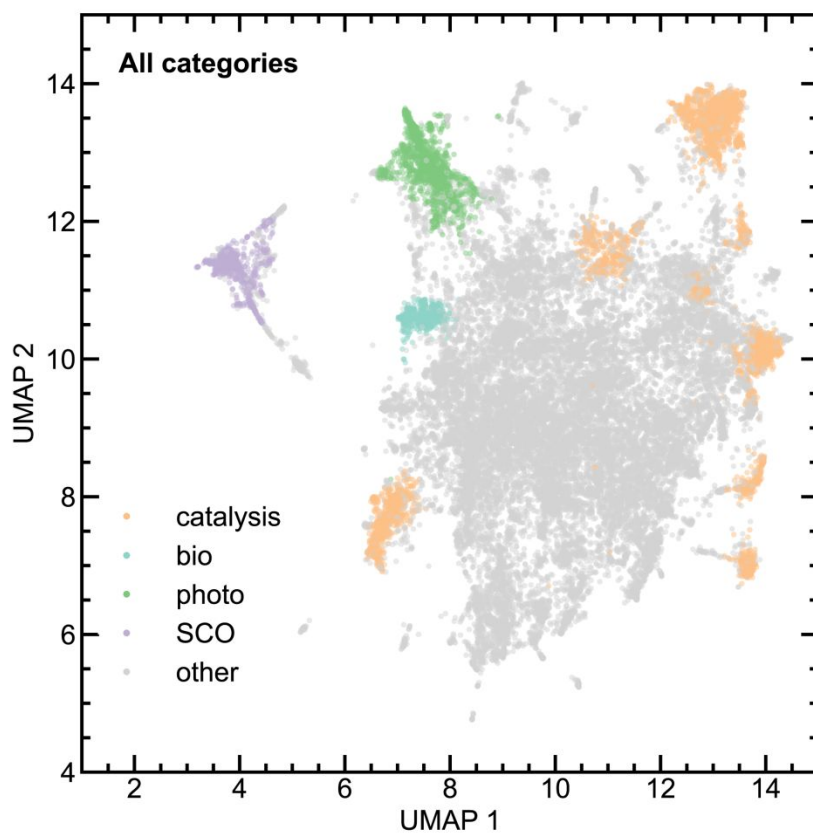
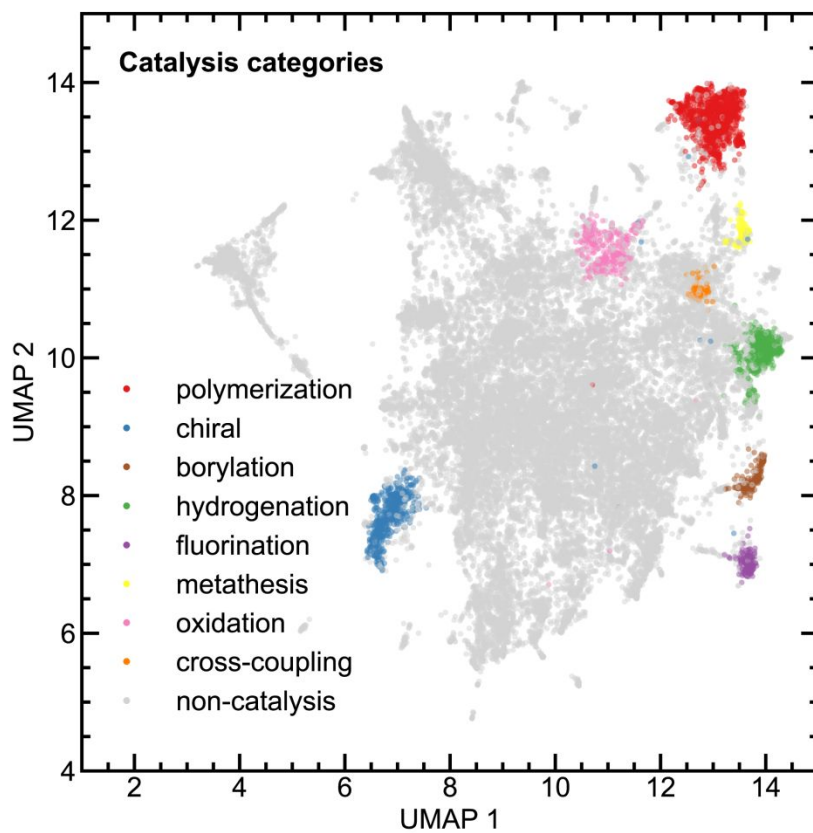


Figure 7. UMAP dimensionality reduction of SBERT embedding vectors colored by different cluster topics for different catalysis applications (top) and general applications (bottom).

To support the findings from the BERTopic model, we employed an additional topic modeling approach, latent Dirichlet allocation (LDA).⁷⁸ LDA is a Bayesian method that iteratively assigns two probabilities: one indicating that a given token belongs to a topic and the other indicating that a topic belongs to a document. This LDA analysis produces semantically similar clusters to BERTopic (Supporting Information Figure S12). Using the LDA approach, several catalysis clusters are identified, including polymerization, chiral catalysis, cross-coupling catalysis, olefin functionalization catalysis, and mechanism-focused catalysis. Some non-catalysis topics are also conserved, such as photoactivity, biological activity, magnetism, and X-ray crystal structure characterization. Feature reduction of the token count vector, obtained using UMAP reduction with Hellinger distance, leads to a similar mapping, demonstrating the close relationship between the identified clusters and the shorter distances between catalysis-related clusters (Supporting Information Figure S13).

Even with the expanded subsets identified by unsupervised clustering using BERTopic or using the LDA analysis, a significant portion of the tmQM complexes are either unlabeled or assigned to difficult-to-interpret clusters. Furthermore, these datasets curated using simple natural language processing methods are not fully context-aware and don't account for more detailed information present in the manuscript, such as discussion of failed experimental attempts, meaning that they may contain "negative" examples. That is, these subsets could contain complexes that were either used as counterexamples, found to be ineffective for a given activity, or represent a precursor structure that was crystallized before the in-situ assembly of chemically relevant species. Accordingly, we carried out analysis of the structures present in each dataset to both support the composition of these datasets and enhance them with chemically relevant species (see Sec. 4).



Such species might not have been originally used for a given application and therefore be missed by NLP approaches, but due to chemical similarity, they could have a complementary application.

3d. Analysis of common metal-centered motifs in curated datasets

It is intuitive that strong relationships should exist between the metal, its coordination environment in a TMC and the associated property or application. As a representative example, linear gold chlorides coordinated to phosphine or N-heterocyclic carbene (NHC)-type ligands are commonly utilized in catalysis^{96,97} but would not be expected to be relevant for spin-crossover because the Au metal is closed shell. We thus carried out an analysis of the metal-centered subgraphs to contrast trends in metal-local structure in the overall tmQM set to those in the tmCAT, tmPHOTO, tmBIO, and tmSCO subsets. We carried out this analysis after excluding lanthanide complexes because they are relatively poorly represented across all datasets. To capture the metal-local environment, we computed metal-centered subgraphs with a truncation of 2 bond paths away (i.e., $d = 2$) from the metal on the molecular graph. The radius of 2 was chosen because it captures metal-local electronic character, while also reducing the diversity of possible combinations. We obtained these subgraphs on the CSD-reported connectivity using the molSimplify package⁹⁸, and we then determined the uniqueness of these subgraphs by computing and comparing the edge-attributed molecular graph hashes of each metal-centered substructure. The Python script, resulting substructures, and graph hashes are all provided on Zenodo⁷⁹.

We first analyzed common structural motifs in the tmCAT dataset. Breaking down the dataset into structural motifs at the radius of 2 shows that while most complexes are unique when analyzing metal-local character, several structural motifs appear relatively frequently in the tmCAT dataset (Supporting Information Figure S14). Overall, 19,250 complexes in the tmCAT



dataset are represented by 10,094 unique $d = 2$ (i.e., metal-local) substructures. In particular, palladium dichlorides bound to either two phosphine ligands or to bidentate nitrogen-coordinating ligands are very common in the tmCAT dataset, appearing 131 and 89 times, respectively (Figure 8). These metal-local motifs are likely derived from complexes that have been widely studied for their ability to catalyze cross-coupling reactions⁹⁹. Similarly, linear gold chlorides bound to NHC or phosphine ligands are common in the tmCAT dataset, appearing 112 and 100 times, respectively (Figure 8). Linear gold complexes are known to catalyze various π -functionalization and annulation reactions, among others¹⁰⁰. The high abundance of these gold and palladium motifs is in line with their occurrence in the tmQM superset, where the popularity of these complexes has led to their widespread examination in many contexts. On the other hand, several frequently occurring motifs were identified that are almost exclusively studied for catalysis (Figure 8). These include nickel catalysts with four mixed N, C, O, and P coordinating ligands, i.e., where each ligand type coordinates the metal once; iron dichloride catalysts with tridentate nitrogen coordinating ligands; and ruthenium dichloride catalysts with an NHC ligand and carbene ligand with a chelating oxygen group¹⁰¹ (i.e., likely derived from the second-generation Hoveyda-Grubbs catalyst¹⁰²). Surprisingly, all these catalysts have been primarily studied for polymerization: Ni complexes are utilized for ethylene copolymerization with carbon monoxide¹⁰³, Fe complexes are used as catalysts for linear homo-polymerization of ethylene for the synthesis of high-density polyethylene¹⁰⁴, and Hoveyda-Grubbs catalysts are utilized for ring-opening metathesis polymerization (ROMP)¹⁰⁵. This highlights how some motifs, despite their limited range of application, have been the focus of a great deal of study as a result of their industrial relevance.



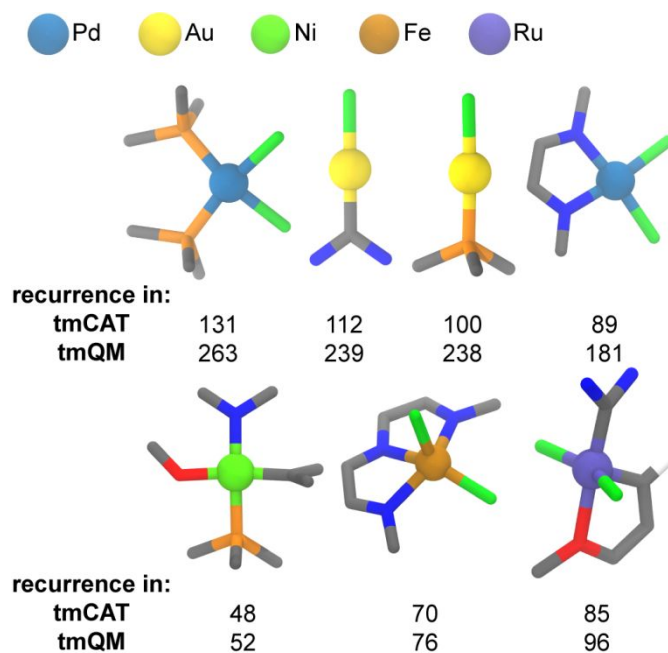


Figure 8. Representative substructures of the tmCAT set. The occurrence of each substructure in the tmCAT set and tmQM superset are displayed. Transition metal centers are shown as spheres. Palladium is shown in light blue, gold in yellow, nickel in light green, iron in brown, ruthenium in purple, carbon in gray, chlorine in green, nitrogen in blue, oxygen in red, phosphorus in orange, and hydrogen in white. The metal atom legend is shown at the top.

Next, we expanded our substructure analysis to the photochemistry-relevant tmPHOTO subset. Analysis of metal identity in tmPHOTO reveals that iridium, platinum, and copper complexes are significantly amplified in this dataset (Supporting Information Figure S15). Iridium¹⁰⁶ and platinum¹⁰⁷ complexes have been common targets for photophysical applications due to spin-orbit coupling that allows intersystem-crossing, which leads to high quantum yields. On the other hand, copper complexes¹⁰⁸ have been explored as an earth-abundant alternative to more rare and expensive iridium and platinum complexes. Substructure mapping shows that 3,043 complexes in tmPHOTO dataset are represented by 1,150 unique $d = 2$ structural motifs. Several commonly recurring substructures can be observed in the tmPHOTO set (Supporting Information Figure S16). These include iridium complexes with a coordination number of six with two bidentate C^N coordinating ligands and two oxygen-coordinating ligands (i.e., structural analogs



of Ir(ppy)₂(acac)), which is a substructure that appears in tmPHOTO 101 times (Supporting Information Figure S17). Platinum complexes with a coordination number of four, a bidentate C^N type ligand, and two oxygen-coordinating ligands are other commonly recurring structural motifs (i.e., structural analogs of Pt(ppy)(acac)), with this substructure appearing 89 times in tmPHOTO. Another similarly commonly recurring motif is a 4-coordinate copper complex with mixed nitrogen and phosphorus coordinating atoms, including a bidentate nitrogen ligand, which appears 89 times in the tmPHOTO dataset (Supporting Information Figure S17). Examples of structural motifs that are almost exclusively studied for photophysical properties include Ir complexes with two bidentate C^N type ligands and a bidentate N^N type ligand, as well as platinum complex with one bidentate C, NHC type ligand and two oxygen-coordinating ligands (Supporting Information Figure S17).

Moving on to the spin-crossover relevant subset, we note that there are necessary differences for TMCs that exhibit switchable magnetic behavior. Here, iron, manganese, nickel, and cobalt complexes occur with higher relative frequency in the tmSCO set than in the tmQM superset (Supporting Information Figure S15). These metals are all third-row transition metals that tend to have relatively low *d*-orbital splitting energy (Δ) and, depending on the oxidation state, they are expected to have multiple accessible spin states. Substructure mapping shows that 834 complexes in the tmSCO dataset are represented by 534 unique *d* = 2 structural motifs. A few *d* = 2 structural motifs are representative of this dataset through multiple recurrences (Supporting Information Figure S18). These recurring motifs include manganese complexes with four nitrogen-coordinating ligands, including a bidentate *sp*³ hybridized ligand and two oxygen-coordinating ligands, which appear 39 times in tmSCO and only six additional times (i.e., 45 in total) across the entire tmQM dataset. This highlights how these complexes are nearly exclusively targeted for



magnetic properties. Another common structural motif includes an iron center with four sp^2 hybridized nitrogen coordinating ligands and two nitrogen coordinating ligands that could be either isocyanides or cyanates (Supporting Information Figure S19). These complexes appear in the tmSCO set 27 times, and in the tmQM superset a total of 44 times. Interestingly, a relatively common structural motif includes iron bound to a tetradentate nitrogen-coordinating ligand with two additional isocyanides/cyanate ligands, which appears in the tmQM dataset 9 times, all of which are in the tmSCO subset, suggesting that these motifs have only been studied for applications related to magnetism (Supporting Information Figure S19).

Finally, we analyzed the substructures in the biological activity subset, which we expect to be the most diverse due to the broad nature of this set. Ruthenium and platinum are the most representative metals in the tmBIO dataset (Supporting Information Figure S15). Overall, substructure mapping shows that 1,808 complexes in the tmBIO dataset are represented by 974 unique $d = 2$ structural motifs. A high interest in platinum for biological applications can be attributed to cisplatin¹⁰⁹, i.e., cis-diamminedichloroplatinum(II), the first inorganic small molecule approved as a pharmaceutical anti-cancer drug. Cisplatin-resistant cancers¹¹⁰ have led to the search for alternate platinum complexes as anti-cancer drugs¹¹¹ and are represented in the tmBIO dataset. Furthermore, the high toxicity of cisplatin has led to the search for alternate inorganic and organometallic complexes that could be used as anti-cancer medications. In particular, the high promise of ruthenium arene 1,3,5-triaza-7-phosphaadamantane (RAPTA)¹¹² compounds has led to a significant effort in screening ruthenium-based piano stool complexes as potential anti-cancer drugs¹¹³. These efforts are consistent with the makeup of the tmBIO dataset. In fact, the three most commonly recurring structural motifs feature ruthenium arene complexes with one chloride and two additional ancillary ligands (Supporting Information Figures S20–21). These complexes,



cumulatively, appear in the tmBIO dataset 82 times and 208 in the tmQM dataset. Furthermore, a motif that is closely related to cisplatin, with platinum, two chloride ligands, a single ammonia, and an organic N-coordinating ligand, appears in tmBIO set 17 times and has been exclusively studied for biological applications (Supporting Information Figure S21).

Finally, we analyzed if any structural motifs appear frequently across different datasets to identify if they are frequently studied for multiple applications. To achieve this, we first created a subset of each dataset, consisting of commonly recurring motifs in each of the dataset (i.e. five or more recurrence for tmCAT and 3 or more recurrence for other datasets). We then analyzed overlaps between each dataset. These motifs are mostly exclusive to a given set, with more than 82% of metal-centered motifs only appearing in one of the four sets exclusively (Figure 9). However, a single metal-centered substructure appears across all datasets. This motif consists of a nickel metal center with mixed N^N,O,O coordinating atoms and is reminiscent of salen complexes, despite the inability of radius 2 subgraphs to capture the entirety of tetradentate salen ligands (Figure 9, left inset). Salen complexes are a common enzyme mimic often used in asymmetric catalysis¹¹⁴ but could be used as targets for photoactive complexes due to the rigid nature of the ligand, as magnetic complexes with the nickel as the metal center, and the square planar coordination environment of salen ligands with Ni complex can be targets for DNA intercalation. Furthermore, several structural motifs were identified to reside among tmCAT, tmPHOTO, and tmBIO sets but not tmSCO. A noteworthy example includes an iridium motif with two bidentate C^N-type ligand and single bidentate N^N type ligand (i.e., Ir(ppy)₂(bpy) analogs). These complexes are commonly used as triplet sensitizers, which can be applied for photocatalysis¹¹⁵ to access excited states and for biological applications to target reactive singlet oxygen formation¹¹⁶ (Figure 9, right inset).



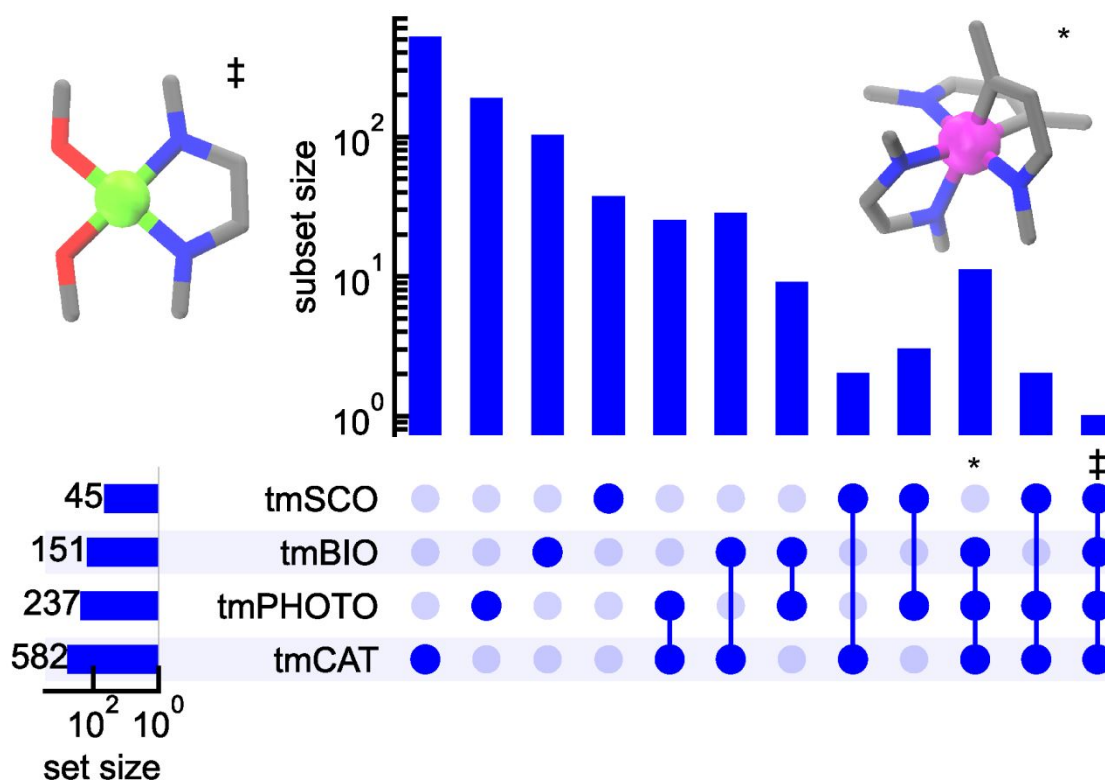


Figure 9. UpSet plot showing how frequently recurring substructures in each of the tmCAT, tmPHOTO, tmBIO and tmSCO datasets intersect. Each set is defined by only retaining substructures with high recurrence. Each subset includes motifs that appear only in highlighted sets. The structure that appears in all four sets is shown as an inset on the left. A structure that appears in all but tmSCO sets is shown as an inset on the right. Nickel in green, iridium in pink, carbon in gray, nitrogen in blue, oxygen in red. A representative motif indicated by a * is found in the intersection of the tmBIO, tmPHOTO, and tmCAT datasets, and a motif indicated by a ‡ is found in all four subsets.

Analysis of common application-specific structural motifs reveals that, for a given motif, there are complexes within the tmQM dataset that contain the motif but the associated manuscripts do not indicate the complex has been assessed for that specific application. Therefore, we supplemented each dataset using structural mapping to identify chemically similar structures to add to our data subsets. To complete this augmentation, we note that for most of these application-specific datasets, multidentate ligands can play an important role, such as in defining the steric environment and introducing added stability for catalysis or inhibiting thermal relaxation pathways



for photoactive compounds. This is consistent with several common multidentate motifs (i.e., five-membered rings) when considering $d = 2$ substructures. However, $d = 2$ substructures cannot capture bidentate ligands that form six-membered metallocycles. Therefore, to avoid introduction of less relevant complexes in each dataset, we identify matching metal-centered substructures in tmQM dataset with $d = 3$. Unsurprisingly, increasing the radius of metal-centered substructures leads to a reduction in the number of recurring substructures, with, e.g., 13,696 unique $d = 3$ motifs (vs. 19,250 for $d = 2$) in the tmCAT dataset (Supporting Information Figure S22). To only introduce additional structures with high relevance to a given application, only motifs with high recurrence (i.e., 5 or more for tmCAT, 3 or more for the other subsets) were supplemented. Using structural mapping, we augmented the tmCAT, tmPHOTO, tmSCO, and tmBIO datasets with 2,381, 1,556, 149, and 974 additional chemically relevant complexes, respectively. The final, application-specific datasets we curated can be accessed on Zenodo⁷⁹.

4. Conclusions

In summary, we employed natural language processing techniques, both supervised and unsupervised, to link experimentally synthesized compounds in the large and diverse tmQM database, which consists of 86,665 datasets, to their respective applications. Using the manuscript abstracts, we first trained a classifier model to identify manuscripts that are related to catalysis with an accuracy of 0.97. Using this model we curated a dataset of catalysis-related transition metal complexes, called tmCAT, which initially consisted of 18,915 unique complexes. Analysis of common electronic and geometric descriptors revealed that commonly used descriptors fail to distinguish between catalytic and non-catalytic sets. However, the analysis of coordination geometry of catalytically relevant complexes showed that geometries with open metal sites were significantly enhanced in the tmCAT set.



Using topic modeling, an unsupervised clustering method often used in natural language processing, we further curated three additional initial datasets: tmPHOTO, a dataset consisting of 3,043 unique complexes with photophysical relevance, tmBIO, a dataset consisting of 1,808 unique complexes with biological relevance, and tmSCO, a dataset consisting of 834 unique complexes with relevance to magnetism. Analyzing the chemical substructures within each dataset identified frequently targeted complexes for their designated applications, such as bidentate N^N palladium dichlorides for catalysis, iridium complexes with two C^N ligands or platinum complexes with one C^N ligands for photophysics, and platinum dichlorides or ruthenium piano-stool complexes for biologically relevant complexes. By mapping these substructures to their applications, we identified previously synthesized complexes that had strong chemical similarity to those already identified for each application. We used these additional complexes to supplement the textually curated datasets, leading to 2,716 additional tmCAT complexes, 1,556 additional tmPHOTO complexes, 974 tmBIO additional complexes, and 149 additional tmSCO complexes in the final data sets.

The curated tmCAT, tmPHOTO, tmBIO, and tmSCO datasets are expected to enable more focused high-throughput computational screening and development of predictive machine learning models while still allowing for exploration across diverse chemical spaces. The language models employed in this study also have the potential for broader application, such as to curate subsets of other classes of materials, such as metal-organic frameworks.

ASSOCIATED CONTENT

Supporting Information.

Validation of abstracts and titles obtained from HTML; All keywords used to screen for catalysis manuscripts in the title; All keywords used to screen for catalysis manuscripts in the abstract and



title; Additional stop words included in the secondary TF-IDF vectorizer; Receiver operating curve of the secondary catalysis classifier; The performance metrics of the secondary catalysis classifier; Distribution of the predicted probabilities of the manuscript related to catalysis in the unlabeled set in the corpus; Distribution of the predicted probabilities of the manuscript related to catalysis in the HTML-mined corpus; Density distribution of HOMO orbital energies; Density distribution of LUMO orbital energies; Density distribution of HOMO-LUMO gap energies; Density distribution of the metal charge; Density distribution of buried volume of metal; The five largest clusters of the BERTopic model using text without the introduction of stop words; Wordclouds of two additional non-catalysis topics; Wordclouds of five additional catalysis-related topics; t-SNE embedding of SBERT embedding vectors; Topic breakdown using Latent Dirichlet Allocation; UMAP embedding of count vectorizer feature vector; Frequency of molecular subgraph recurrence in the tmCAT dataset; Frequency of the five most common transition metals in tmPHOTO, tmBIO and tmSCO datasets; Frequency of molecular subgraph recurrence in the tmPHOTO dataset; Representative substructures of the tmPHOTO dataset; Frequency of molecular subgraph recurrence in the tmSCO dataset; Representative substructures of the tmSCO dataset; Frequency of molecular subgraph recurrence in the tmBIO dataset; Representative substructures of tmBIO dataset; Frequency of $d = 3$ molecular subgraph recurrence in the tmCAT dataset. (PDF)

AUTHOR INFORMATION

Corresponding Author

*email:hjkulik@mit.edu

Notes

The authors declare no competing financial interest.

AUTHOR CONTRIBUTIONS

Conceptualization - I. K. and H. J. K.; Software - I. K., A. G. G., and R. G. S. M.; Formal Analysis- I. K., A. G. G., R. G. S. M., H. A., H. J., J. W. T.; Data Curation I. K. and H. A.; Writing - Original Draft - I. K., H. A., and J. W. T; Writing - Review and Editing - I. K., A. G. G., R. G. S. M, H. A., H. J, J. W. T., and H. J. K.; Supervision - H. J. K; Funding Acquisition - H. J. K. and Y. R.-L..

ACKNOWLEDGMENT

This work was primarily supported by the National Science Foundation under grant number CBET-1846426. The authors also acknowledge funding from the Office of Naval Research



(N00014-20-1-2150). H.J.K. holds an Alfred P. Sloan Fellowship in Chemistry and a Simon Family Faculty Research Innovation Fund, which supported this work. The authors thank Adam H. Steeves for providing a critical reading of the manuscript.

REFERENCES

- (1) Fey, N. Lost in Chemical Space? Maps to Support Organometallic Catalysis. *Chem. Cent. J.* **2015**, *9*.
- (2) Jover, J.; Fey, N. The Computational Road to Better Catalysts. *Chem. Asian J.* **2014**, *9*, 1714-23.
- (3) Swart, M.; Gruden, M. Spinning around in Transition-Metal Chemistry. *Acc. Chem. Res.* **2016**, *49*, 2690-2697.
- (4) Krska, S. W.; DiRocco, D. A.; Dreher, S. D.; Shevlin, M. The Evolution of Chemical High-Throughput Experimentation to Address Challenging Problems in Pharmaceutical Synthesis. *Acc. Chem. Res.* **2017**, *50*, 2976-2985.
- (5) Takale, B. S.; Kong, F.-Y.; Thakore, R. R. Recent Applications of Pd-Catalyzed Suzuki–Miyaura and Buchwald–Hartwig Couplings in Pharmaceutical Process Chemistry. *Organics* **2022**, *3*, 1-21.
- (6) Allardyce, C. S.; Dorcier, A.; Scolaro, C.; Dyson, P. J. Development of Organometallic (Organo-Transition Metal) Pharmaceuticals. *Appl. Organometal. Chem.* **2005**, *19*, 1-10.
- (7) Nandy, A.; Duan, C.; Taylor, M. G.; Liu, F.; Steeves, A. H.; Kulik, H. J. Computational Discovery of Transition-Metal Complexes: From High-Throughput Screening to Machine Learning. *Chem. Rev.* **2021**, *121*, 9927-10000.
- (8) Tolman, C. A. Steric Effects of Phosphorus Ligands in Organometallic Chemistry and Homogeneous Catalysis. *Chem. Rev.* **1977**, *77*, 313-348.
- (9) Fey, N.; Lynam, J. M. Computational Mechanistic Study in Organometallic Catalysis: Why Prediction Is Still a Challenge. *WIREs Comput. Mol. Sci.* **2021**, *12*.
- (10) Reymond, J.-L.; Ruddigkeit, L.; Blum, L.; van Deursen, R. The Enumeration of Chemical Space. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **2012**, *2*, 717-733.
- (11) Janet, J. P.; Ramesh, S.; Duan, C.; Kulik, H. J. Accurate Multiobjective Design in a Space of Millions of Transition Metal Complexes with Neural-Network-Driven Efficient Global Optimization. *ACS Cent. Sci.* **2020**, *6*, 513-524.
- (12) Hageman, J. A.; Westerhuis, J. A.; Frühauf, H. W.; Rothenberg, G. Design and Assembly of Virtual Homogeneous Catalyst Libraries –Towards in Silico Catalyst Optimisation. *Adv. Synth. Catal.* **2006**, *348*, 361-369.
- (13) Nandy, A.; Taylor, M. G.; Kulik, H. J. Identifying Underexplored and Untapped Regions in the Chemical Space of Transition Metal Complexes. *J. Phys. Chem. Lett.* **2023**, *14*, 5798-5804.



- (14) DiLuzio, S.; Mdluli, V.; Connell, T. U.; Lewis, J.; VanBenschoten, V.; Bernhard, S. High-Throughput Screening and Automated Data-Driven Analysis of the Triplet Photophysical Properties of Structurally Diverse, Heteroleptic Iridium(III) Complexes. *J. Am. Chem. Soc.* **2021**, *143*, 1179-1194.
- (15) Motz, R. N.; Lopato, E. M.; Connell, T. U.; Bernhard, S. High-Throughput Screening of Earth-Abundant Water Reduction Catalysts toward Photocatalytic Hydrogen Evolution. *Inorg. Chem.* **2021**, *60*, 774-781.
- (16) Balcells, D.; Skjelstad, B. B. Tmqm Dataset-Quantum Geometries and Properties of 86k Transition Metal Complexes. *J. Chem. Inf. Model.* **2020**, *60*, 6135-6146.
- (17) Chen, S. S.; Meyer, Z.; Jensen, B.; Kraus, A.; Lambert, A.; Ess, D. H. Realigands: A Ligand Library Cultivated from Experiment and Intended for Molecular Computational Catalyst Design. *J. Chem. Inf. Model.* **2023**, *63*, 7412-7422.
- (18) Kevlishvili, I.; Duan, C.; Kulik, H. J. Classification of Hemilabile Ligands Using Machine Learning. *J. Phys. Chem. Lett.* **2023**, *14*, 11100-11109.
- (19) Fey, N.; Tsipis, A. C.; Harris, S. E.; Harvey, J. N.; Orpen, A. G.; Mansson, R. A. Development of a Ligand Knowledge Base, Part 1: Computational Descriptors for Phosphorus Donor Ligands. *Chem. Eur. J.* **2006**, *12*, 291-302.
- (20) Mansson, R. A.; Welsh, A. H.; Fey, N.; Orpen, A. G. Statistical Modeling of a Ligand Knowledge Base. *J. Chem. Inf. Model.* **2006**, *46*, 2591-2600.
- (21) Jover, J.; Fey, N.; Harvey, J. N.; Lloyd-Jones, G. C.; Orpen, A. G.; Owen-Smith, G. J. J.; Murray, P.; Hose, D. R. J.; Osborne, R.; Purdie, M. Expansion of the Ligand Knowledge Base for Monodentate P-Donor Ligands (Lkb-P). *Organometallics* **2010**, *29*, 6245-6258.
- (22) Jover, J.; Fey, N.; Harvey, J. N.; Lloyd-Jones, G. C.; Orpen, A. G.; Owen-Smith, G. J.; Murray, P.; Hose, D. R.; Osborne, R.; Purdie, M. Expansion of the Ligand Knowledge Base for Chelating P,P-Donor Ligands (Lkb-Pp). *Organometallics* **2012**, *31*, 5302-5306.
- (23) Gensch, T.; Dos Passos Gomes, G.; Friederich, P.; Peters, E.; Gaudin, T.; Pollice, R.; Jorner, K.; Nigam, A.; Lindner-D'Addario, M.; Sigman, M. S.; Aspuru-Guzik, A. A Comprehensive Discovery Platform for Organophosphorus Ligands for Catalysis. *J. Am. Chem. Soc.* **2022**, *144*, 1205-1217.
- (24) Kearnes, S. M.; Maser, M. R.; Wlekinski, 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.
- (25) 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.
- (26) Fitzner, M.; Wuitschik, G.; Koller, R. J.; Adam, J. M.; Schindler, T.; Reymond, J. L. What Can Reaction Databases Teach Us About Buchwald-Hartwig Cross-Couplings? *Chem. Sci.* **2020**, *11*, 13085-13093.
- (27) 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.



- (28) Groom, C. R.; Bruno, I. J.; Lightfoot, M. P.; Ward, S. C. The Cambridge Structural Database. *Acta. Crystallogr. B Struct. Sci. Cryst. Eng. Mater.* **2016**, *72*, 171-179.
- (29) Montavon, G.; Rupp, M.; Gobre, V.; Vazquez-Mayagoita, A.; Hansen, K.; Tkatchenko, A.; Müller, K.-R.; von Lilienfeld, O. A. Machine Learning of Molecular Electronic Properties in Chemical Compound Space. *New J. Phys.* **2013**, *15*, 095003.
- (30) Virshup, A. M.; Contreras-Garcia, J.; Wipf, P.; Yang, W.; Beratan, D. N. Stochastic Yoyages into Uncharted Chemical Space Produce a Representative Library of All Possible Drug-Like Compounds. *J. Am. Chem. Soc.* **2013**, *135*, 7296-303.
- (31) Ramakrishnan, R.; Dral, P. O.; Rupp, M.; von Lilienfeld, O. A. Quantum Chemistry Structures and Properties of 134 Kilo Molecules. *Sci. Data.* **2014**, *1*, 140022.
- (32) Rupp, M.; Tkatchenko, A.; Muller, K. R.; von Lilienfeld, O. A. Fast and Accurate Modeling of Molecular Atomization Energies with Machine Learning. *Phys. Rev. Lett.* **2012**, *108*, 058301.
- (33) Reymond, J.-L.; van Deursen, R.; Blum, L. C.; Ruddigkeit, L. Chemical Space as a Source for New Drugs. *Med. Chem. Comm.* **2010**, *1*.
- (34) Ramakrishnan, R.; Hartmann, M.; Tapavicza, E.; von Lilienfeld, O. A. Electronic Spectra from Tddft and Machine Learning in Chemical Space. *J. Chem. Phys.* **2015**, *143*, 084111.
- (35) Ruddigkeit, L.; van Deursen, R.; Blum, L. C.; Reymond, J. L. Enumeration of 166 Billion Organic Small Molecules in the Chemical Universe Database Gdb-17. *J. Chem. Inf. Model.* **2012**, *52*, 2864-75.
- (36) Kneiding, H.; Nova, A.; Balcells, D. Directional Multiobjective Optimization of Metal Complexes at the Billion-Scale with the Tmqmg-L Dataset and PI-Moga Algorithm. *ChemRxiv* **2023**, DOI:10.26434/chemrxiv-2023-k3tf2-v2 10.26434/chemrxiv-2023-k3tf2-v2.
- (37) Schwaller, P.; Vaucher, A. C.; Laino, T.; Reymond, J.-L. Prediction of Chemical Reaction Yields Using Deep Learning. *Mach. Learn.: Sci. Technol.* **2021**, *2*.
- (38) Schwaller, P.; Vaucher, A. C.; Laino, T.; Reymond, J.-L. Data Augmentation Strategies to Improve Reaction Yield Predictions and Estimate Uncertainty. *NeurIPS* **2020**, DOI:10.26434/chemrxiv.13286741.v1 10.26434/chemrxiv.13286741.v1.
- (39) Jin, H.; Merz, K. M., Jr. Modeling Zinc Complexes Using Neural Networks. *J. Chem. Inf. Model.* **2024**, *64*, 3140-3148.
- (40) Garrison, A. G.; Heras-Domingo, J.; Kitchin, J. R.; Dos Passos Gomes, G.; Ulissi, Z. W.; Blau, S. M. Applying Large Graph Neural Networks to Predict Transition Metal Complex Energies Using the Tmqm_Wb97mv Data Set. *J. Chem. Inf. Model.* **2023**, *63*, 7642-7654.
- (41) Gensch, T.; Smith, S. R.; Colacot, T. J.; Timsina, Y. N.; Xu, G.; Glasspoole, B. W.; Sigman, M. S. Design and Application of a Screening Set for Monophosphine Ligands in Cross-Coupling. *ACS Catal.* **2022**, *12*, 7773-7780.
- (42) Chowdhary, K. R. In *Fundamentals of Artificial Intelligence*; Chowdhary, K. R., Ed.; Springer India: New Delhi, 2020, DOI:10.1007/978-81-322-3972-7_19 10.1007/978-81-322-3972-7_19,603-649



- (43) Krallinger, M.; Rabal, O.; Lourenço, A.; Oyarzabal, J.; Valencia, A. Information Retrieval and Text Mining Technologies for Chemistry. *Chem. Rev.* **2017**, *117*, 7673-7761.
- (44) Kim, E.; Huang, K.; Saunders, A.; McCallum, A.; Ceder, G.; Olivetti, E. Materials Synthesis Insights from Scientific Literature Via Text Extraction and Machine Learning. *Chem. Mater.* **2017**, *29*, 9436-9444.
- (45) Kim, E.; Huang, K.; Jegelka, S.; Olivetti, E. Virtual Screening of Inorganic Materials Synthesis Parameters with Deep Learning. *npj Comput. Mater.* **2017**, *3*, 53.
- (46) Swain, M. C.; Cole, J. M. Chemdataextractor: A Toolkit for Automated Extraction of Chemical Information from the Scientific Literature. *J. Chem. Inf. Model.* **2016**, *56*, 1894-1904.
- (47) Mavracic, J.; Court, C. J.; Isazawa, T.; Elliott, S. R.; Cole, J. M. Chemdataextractor 2.0: Autopopulated Ontologies for Materials Science. *J. Chem. Inf. Model.* **2021**, *61*, 4280-4289.
- (48) Park, S.; Kim, B.; Choi, S.; Boyd, P. G.; Smit, B.; Kim, J. Text Mining Metal–Organic Framework Papers. *J. Chem. Inf. Model.* **2018**, *58*, 244-251.
- (49) Jensen, Z.; Kim, E.; Kwon, S.; Gani, T. Z. H.; Román-Leshkov, Y.; Moliner, M.; Corma, A.; Olivetti, E. A Machine Learning Approach to Zeolite Synthesis Enabled by Automatic Literature Data Extraction. *ACS Cent. Sci.* **2019**, *5*, 892-899.
- (50) Kim, E.; Jensen, Z.; van Grootel, A.; Huang, K.; Staib, M.; Mysore, S.; Chang, H.-S.; Strubell, E.; McCallum, A.; Jegelka, S.; Olivetti, E. Inorganic Materials Synthesis Planning with Literature-Trained Neural Networks. *J. Chem. Inf. Model.* **2020**, *60*, 1194-1201.
- (51) Nandy, A.; Duan, C.; Kulik, H. J. Using Machine Learning and Data Mining to Leverage Community Knowledge for the Engineering of Stable Metal-Organic Frameworks. *J. Am. Chem. Soc.* **2021**, *143*, 17535-17547.
- (52) Zheng, Z.; Zhang, O.; Borgs, C.; Chayes, J. T.; Yaghi, O. M. Chatgpt Chemistry Assistant for Text Mining and the Prediction of Mof Synthesis. *J. Am. Chem. Soc.* **2023**, *145*, 18048-18062.
- (53) Jablonka, K. M.; Schwaller, P.; Ortega-Guerrero, A.; Smit, B. Leveraging Large Language Models for Predictive Chemistry. *Nat. Mach. Intell.* **2024**, *6*, 161-169.
- (54) Polak, M. P.; Morgan, D. Extracting Accurate Materials Data from Research Papers with Conversational Language Models and Prompt Engineering. *Nat. Commun.* **2024**, *15*, 1569.
- (55) Liu, S.; Wen, T.; Pattamatta, A.; Srolovitz, D. J. A Prompt-Engineered Large Language Model, Deep Learning Workflow for Materials Classification. *arXiv preprint arXiv:2401.17788* **2024**.
- (56) Hatakeyama-Sato, K.; Yamane, N.; Igarashi, Y.; Nabaie, Y.; Hayakawa, T. Prompt Engineering of Gpt-4 for Chemical Research: What Can/Cannot Be Done? *Sci. Technol. Adv. Material, Meth.* **2023**, *3*, 2260300.



- (57) Vayansky, I.; Kumar, S. A. P. A Review of Topic Modeling Methods. *Inf. Syst.* **2020**, *94*, 101582.
- (58) Huang, E. H.; Socher, R.; Manning, C. D.; Ng, A. Y. Proceedings of the 50th annual meeting of the association for computational linguistics (Volume 1: Long papers), Jeju, Republic of Korea, 2012; p 873-882.
- (59) Mikolov, T.; Chen, K.; Corrado, G.; Dean, J. Efficient Estimation of Word Representations in Vector Space. *arXiv preprint arXiv:1301.3781* **2013**.
- (60) Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; Dean, J. Distributed Representations of Words and Phrases and Their Compositionality. *NeurIPS* **2013**, *26*.
- (61) Kusner, M.; Sun, Y.; Kolkin, N.; Weinberger, K. In *Proceedings of the 32nd International Conference on Machine Learning*; Francis, Bach; David, Blei, Eds.; PMLR: Proceedings of Machine Learning Research, 2015; Vol. 37.
- (62) Grootendorst, M. Bertopic: Neural Topic Modeling with a Class-Based Tf-Idf Procedure. *arXiv preprint arXiv:2203.05794* **2022**.
- (63) Blei, D. M.; Ng, A. Y.; Jordan, M. I. Latent Dirichlet Allocation. *JMLR* **2003**, *3*, 993-1022.
- (64) Huang, S.; Cole, J. M. Chemdatawriter: A Transformer-Based Toolkit for Auto-Generating Books That Summarise Research. *Digit. Discov.* **2023**, *2*, 1710-1720.
- (65) Jelodar, H.; Wang, Y.; Yuan, C.; Feng, X.; Jiang, X.; Li, Y.; Zhao, L. Latent Dirichlet Allocation (Lda) and Topic Modeling: Models, Applications, a Survey. *Multimed. Tools Appl.* **2019**, *78*, 15169-15211.
- (66) Kang, H. J.; Kim, C.; Kang, K. Analysis of the Trends in Biochemical Research Using Latent Dirichlet Allocation (Lda). *Processes* **2019**, *7*, 379.
- (67) Tran, B. X.; Latkin, C. A.; Sharafeldin, N.; Nguyen, K.; Vu, G. T.; Tam, W. W. S.; Cheung, N.-M.; Nguyen, H. L. T.; Ho, C. S. H.; Ho, R. C. M. Characterizing Artificial Intelligence Applications in Cancer Research: A Latent Dirichlet Allocation Analysis. *JMIR Med. Inform.* **2019**, *7*, e14401.
- (68) Karabacak, M.; Jagtiani, P.; Jain, A.; Panov, F.; Margetis, K. Tracing Topics and Trends in Drug-Resistant Epilepsy Research Using a Natural Language Processing-Based Topic Modeling Approach. *Epilepsia* **2024**, *65*, 861-872.
- (69) Schneider, N.; Fechner, N.; Landrum, G. A.; Stiefl, N. Chemical Topic Modeling: Exploring Molecular Data Sets Using a Common Text-Mining Approach. *J. Chem. Inf. Model.* **2017**, *57*, 1816-1831.
- (70) Nandy, A.; Taylor, M. G.; Kulik, H. J. Identifying Underexplored and Untapped Regions in the Chemical Space of Transition Metal Complexes. *The Journal of Physical Chemistry Letters* **2023**, *14*, 5798-5804.
- (71) Taylor, M. G.; Yang, T.; Lin, S.; Nandy, A.; Janet, J. P.; Duan, C.; Kulik, H. J. Seeing Is Believing: Experimental Spin States from Machine Learning Model Structure Predictions. *J. Phys. Chem. A* **2020**, *124*, 3286-3299.
- (72) Kim, E.; Huang, K.; Jegelka, S.; Olivetti, E. Virtual Screening of Inorganic Materials Synthesis Parameters with Deep Learning. *npj Computational Materials* **2017**, *3*, 53.



- (73) Richardson, L. Beautiful Soup Documentation. *April* **2007**.
- (74) Bird, S.; Klein, E.; Loper, E. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*, " O'Reilly Media, Inc.", 2009.
- (75) Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V. Scikit-Learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825-2830.
- (76) Reimers, N.; Gurevych, I. Sentence-Bert: Sentence Embeddings Using Siamese Bert-Networks. *arXiv preprint arXiv:1908.10084* **2019**.
- (77) McInnes, L.; Healy, J.; Melville, J. Umap: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv preprint arXiv:1802.03426* **2018**.
- (78) Blei, D. M.; Ng, A. Y.; Jordan, M. I. Latent Dirichlet Allocation. *Journal of machine Learning research* **2003**, *3*, 993-1022.
- (79) Zenodo Repository for Leveraging Natural Language Processing to Curate the Tmcat, Tmphoto, Tmbio, and Tmsco Datasets of Functional Transition Metal Complexes. <https://zenodo.org/records/11404217>. (Accessed May 31, 2024).
- (80) Campello, R. J. G. B.; Moulavi, D.; Sander, J. Advances in Knowledge Discovery and Data Mining, Berlin, Heidelberg, 2013; p 160-172.
- (81) Balcells, D.; Skjelstad, B. B. Tmqm Dataset—Quantum Geometries and Properties of 86k Transition Metal Complexes. *Journal of Chemical Information and Modeling* **2020**, *60*, 6135-6146.
- (82) Lan, M.; Tan, C.-L.; Low, H.-B.; Sung, S.-Y. In *Special interest tracks and posters of the 14th international conference on World Wide Web*, Association for Computing Machinery: Chiba, Japan, 2005, DOI:10.1145/1062745.1062854 10.1145/1062745.1062854.
- (83) Zeimentz, P. M.; Arndt, S.; Elvidge, B. R.; Okuda, J. Cationic Organometallic Complexes of Scandium, Yttrium, and the Lanthanoids. *Chemical Reviews* **2006**, *106*, 2404-2433.
- (84) Wang, X.; Brosmer, J. L.; Thevenon, A.; Diaconescu, P. L. Highly Active Yttrium Catalysts for the Ring-Opening Polymerization of ϵ -Caprolactone and δ -Valerolactone. *Organometallics* **2015**, *34*, 4700-4706.
- (85) Biffis, A.; Centomo, P.; Del Zotto, A.; Zecca, M. Pd Metal Catalysts for Cross-Couplings and Related Reactions in the 21st Century: A Critical Review. *Chemical Reviews* **2018**, *118*, 2249-2295.
- (86) Nolan, S. P.; Clavier, H. Chemoselective Olefin Metathesis Transformations Mediated by Ruthenium Complexes. *Chem. Soc. Rev* **2010**, *39*, 3305-3316.
- (87) Iglesias, M.; Oro, L. A. A Leap Forward in Iridium–Nhc Catalysis: New Horizons and Mechanistic Insights. *Chemical Society Reviews* **2018**, *47*, 2772-2808.
- (88) Zi, W.; Dean, T. F. Recent Advances in Enantioselective Gold Catalysis. *Chem. Soc. Rev* **2016**, *45*, 4567-4589.



- (89) Chen, J.; Lu, Z. Asymmetric Hydrofunctionalization of Minimally Functionalized Alkenes Via Earth Abundant Transition Metal Catalysis. *Organic Chemistry Frontiers* **2018**, *5*, 260-272.
- (90) Liao, P.; Getman, R. B.; Snurr, R. Q. Optimizing Open Iron Sites in Metal–Organic Frameworks for Ethane Oxidation: A First-Principles Study. *ACS Applied Materials & Interfaces* **2017**, *9*, 33484-33492.
- (91) Andrikopoulos, P. C.; Michel, C.; Chouzier, S.; Sautet, P. In Silico Screening of Iron-Oxo Catalysts for C–H Bond Cleavage. *ACS Catalysis* **2015**, *5*, 2490-2499.
- (92) Nandy, A.; Zhu, J.; Janet, J. P.; Duan, C.; Getman, R. B.; Kulik, H. J. Machine Learning Accelerates the Discovery of Design Rules and Exceptions in Stable Metal–Oxo Intermediate Formation. *ACS Catal.* **2019**, *9*, 8243-8255.
- (93) Wu, K.; Doyle, A. G. Parameterization of Phosphine Ligands Demonstrates Enhancement of Nickel Catalysis Via Remote Steric Effects. *Nat. Chem.* **2017**, *9*, 779-784.
- (94) Vayansky, I.; Kumar, S. A. P. A Review of Topic Modeling Methods. *Information Systems* **2020**, *94*, 101582.
- (95) Van der Maaten, L.; Hinton, G. Visualizing Data Using T-Sne. *J. Mach. Learn. Res.* **2008**, *9*.
- (96) Marion, N.; Nolan, S. P. N-Heterocyclic Carbenes in Gold Catalysis. *Chem. Soc. Rev.* **2008**, *37*, 1776-1782.
- (97) Nolan, S. P. The Development and Catalytic Uses of N-Heterocyclic Carbene Gold Complexes. *Acc. Chem. Res.* **2011**, *44*, 91-100.
- (98) Ioannidis, E. I.; Gani, T. Z. H.; Kulik, H. J. Molsimplify: A Toolkit for Automating Discovery in Inorganic Chemistry. *J. Comput. Chem.* **2016**, *37*, 2106-2117.
- (99) Hazari, N.; Melvin, P. R.; Beromi, M. M. Well-Defined Nickel and Palladium Precatalysts for Cross-Coupling. *Nature Reviews Chemistry* **2017**, *1*, 0025.
- (100) Campeau, D.; León Rayo, D. F.; Mansour, A.; Muratov, K.; Gagosz, F. Gold-Catalyzed Reactions of Specially Activated Alkynes, Allenes, and Alkenes. *Chemical Reviews* **2021**, *121*, 8756-8867.
- (101) Gessler, S.; Randl, S.; Blechert, S. Synthesis and Metathesis Reactions of a Phosphine-Free Dihydroimidazole Carbene Ruthenium Complex. *Tetrahedron Letters* **2000**, *41*, 9973-9976.
- (102) Vougioukalakis, G. C.; Grubbs, R. H. Ruthenium-Based Heterocyclic Carbene-Coordinated Olefin Metathesis Catalysts. *Chem. Rev.* **2010**, *110*, 1746-1787.
- (103) Bie, Z.; Li, B.-G.; Hu, J.; Yao, Z. Studies on Alternating Copolymerization of Ethylene and Carbon Monoxide Using Nickel-Based Catalyst: Cocatalyst and the Polarity of Solvent. *Macromolecular Reaction Engineering* **2022**, *16*, 2100047.
- (104) Gibson, V. C.; Redshaw, C.; Solan, G. A. Bis(Imino)Pyridines: Surprisingly Reactive Ligands and a Gateway to New Families of Catalysts. *Chemical Reviews* **2007**, *107*, 1745-1776.



- (105) Bielawski, C. W.; Grubbs, R. H. Living Ring-Opening Metathesis Polymerization. *Progress in Polymer Science* **2007**, *32*, 1-29.
- (106) Yang, C.; Mehmood, F.; Lam, T. L.; Chan, S. L.-F.; Wu, Y.; Yeung, C.-S.; Guan, X.; Li, K.; Chung, C. Y.-S.; Zhou, C.-Y.; Zou, T.; Che, C.-M. Stable Luminescent Iridium(III) Complexes with Bis(N-Heterocyclic Carbene) Ligands: Photo-Stability, Excited State Properties, Visible-Light-Driven Radical Cyclization and CO₂ Reduction, and Cellular Imaging. *Chem. Sci.* **2016**, *7*, 3123-3136.
- (107) Kalinowski, J.; Fattori, V.; Cocchi, M.; Williams, J. A. G. Light-Emitting Devices Based on Organometallic Platinum Complexes as Emitters. *Coord. Chem. Rev.* **2011**, *255*, 2401-2425.
- (108) Lazorski, M. S.; Castellano, F. N. Advances in the Light Conversion Properties of Cu(I)-Based Photosensitizers. *Polyhedron* **2014**, *82*, 57-70.
- (109) Prestayko, A. W.; D'Aoust, J. C.; Issell, B. F.; Crooke, S. T. Cisplatin (Cis-Diamminedichloroplatinum II). *Cancer Treat. Rev.* **1979**, *6*, 17-39.
- (110) Kartalou, M.; Essigmann, J. M. Mechanisms of Resistance to Cisplatin. *Mutat. Res.* **2001**, *478*, 23-43.
- (111) Weiss, R. B.; Christian, M. C. New Cisplatin Analogues in Development. *Drugs* **1993**, *46*, 360-377.
- (112) Murray, B. S.; Babak, M. V.; Hartinger, C. G.; Dyson, P. J. The Development of Raptam Compounds for the Treatment of Tumors. *Coord. Chem. Rev.* **2016**, *306*, 86-114.
- (113) Swaminathan, S.; Haribabu, J.; Balakrishnan, N.; Vasanthakumar, P.; Karvembu, R. Piano Stool Ru(II)-Arene Complexes Having Three Monodentate Legs: A Comprehensive Review on Their Development as Anticancer Therapeutics over the Past Decade. *Coord. Chem. Rev.* **2022**, *459*, 214403.
- (114) de Castro, B.; Ferreira, R.; Freire, C.; García, H.; Palomares, E. J.; Sabater, M. J. Photochemistry of Nickel Salen Based Complexes and Relevance to Catalysis. *New Journal of Chemistry* **2002**, *26*, 405-410.
- (115) Teegardin, K.; Day, J. I.; Chan, J.; Weaver, J. Advances in Photocatalysis: A Microreview of Visible Light Mediated Ruthenium and Iridium Catalyzed Organic Transformations. *Org. Process Res. Dev.* **2016**, *20*, 1156-1163.
- (116) Bevernaegie, R.; Doix, B.; Bastien, E.; Diman, A.; Decottignies, A.; Feron, O.; Elias, B. Exploring the Phototoxicity of Hypoxic Active Iridium(III)-Based Sensitizers in 3d Tumor Spheroids. *Journal of the American Chemical Society* **2019**, *141*, 18486-18491.



Data for this article, including the analysis of coordination geometries and substructure analysis, the corpus curated for identifying catalysts, the resulting datasets of all complexes, and the machine learning models for the curation are available in a Zenodo repository at <https://zenodo.org/records/11404217>.

The data supporting this article, including relevant figures and tables, have also been included as part of the Supplementary Information.

