

# Assessment of Fine-Tuned Large Language Models for Real-World Chemistry and Material Science Applications

Joren Van Herck <sup>1,\*\*</sup>, María Victoria Gil <sup>1,2,\*\*</sup>, Kevin Maik Jablonka <sup>1, 3, 4,\*\*</sup>, Alex Abrudan <sup>5</sup>, Andy S. Anker <sup>6,7</sup>, Mehrdad Asgari <sup>8</sup>, Ben Blaiszik <sup>9,10</sup>, Leander Choudhury <sup>11</sup>, Clemence Corminboeuf <sup>12</sup>, Hilal Daglar <sup>13</sup>, Ian T. Foster <sup>9,10</sup>, Susana Garcia <sup>14</sup>, Matthew Garvin <sup>14</sup>, Guillaume Godin <sup>15</sup>, Lydia L. Good <sup>5,16</sup>, Jianan Gu <sup>17</sup>, Noémie Xiao Hu <sup>1</sup>, Xin Jin <sup>1</sup>, Tanja Junkers <sup>18</sup>, Seda Keskin <sup>13</sup>, Tuomas P.J. Knowles <sup>5,19</sup>, Ruben Laplaza <sup>12</sup>, Sauradeep Majumdar <sup>1</sup>, Hossein Mashhadimoslem <sup>20</sup>, Ruairaidh D. McIntosh <sup>21</sup>, Seyed Mohamad Moosavi <sup>22</sup>, Beatriz Mouriño <sup>1</sup>, Francesca Nerli <sup>23</sup>, Covadonga Pevida <sup>2</sup>, Neda Poudineh <sup>14</sup>, Mahyar Rajabi-Kochi <sup>22</sup>, Kadi L. Saar <sup>5</sup>, Fahimeh Hooriabad Saboor <sup>24</sup>, Morteza Sagharichiha<sup>25</sup>, KJ Schmidt <sup>9</sup>, Jiale Shi <sup>26,27</sup>, Dennis Svatunek <sup>28</sup>, Marco Taddei <sup>23</sup>, Igor Tetko <sup>15, 29</sup>, Domonkos Tolnai <sup>17</sup>, Sahar Vahdatifar <sup>25</sup>, Jonathan Whitmer <sup>27,30</sup>, D.C. Florian Wieland <sup>17</sup>, Regine Willumeit-Römer <sup>17</sup>, Andreas Züttel <sup>31</sup>, and Berend Smit <sup>1,\*</sup>

<sup>1</sup>Laboratory of molecular simulation (LSMO), Institut des Sciences et Ingénierie Chimiques, École Polytechnique Fédérale de Lausanne (EPFL), Rue de l'Industrie 17, CH-1951 Sion, Switzerland

<sup>2</sup>Instituto de Ciencia y Tecnología del Carbono (INCAR), CSIC, Francisco Pintado Fe 26, 33011 Oviedo, Spain

<sup>3</sup>Laboratory of Organic and Macromolecular Chemistry (IOMC), Friedrich Schiller University Jena, Humboldtstrasse 10, 07743 Jena, Germany

<sup>4</sup>Helmholtz Institute for Polymers in Energy Applications Jena (HIPOLE Jena), Lessingstrasse 12-14, 07743 Jena, Germany

<sup>5</sup>Yusuf Hamied Department of Chemistry, University of Cambridge, Cambridge CB2 1EW, United Kingdom

<sup>6</sup>Department of Energy Conversion and Storage, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark

<sup>7</sup>Department of Chemistry, University of Oxford, Oxford OX1 3TA, United Kingdom

- <sup>8</sup>Department of Chemical Engineering & Biotechnology, University of Cambridge, Philippa Fawcett Drive, Cambridge CB3 0AS, United Kingdom
- <sup>9</sup>Department of Computer Science, University of Chicago, Chicago, IL 60637, United States
- <sup>10</sup>Data Science and Learning Division, Argonne National Laboratory, Lemont, IL 60439, United States
- <sup>11</sup>Laboratory of Catalysis and Organic Synthesis (LCSO), Institute of Chemical Sciences and Engineering (ISIC), École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
- <sup>12</sup>Laboratory for Computational Molecular Design (LCMD), Institute of Chemical Sciences and Engineering (ISIC), École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
- <sup>13</sup>Department of Chemical and Biological Engineering, Koç University, Rumelifeneri Yolu, Sariyer, 34450 Istanbul, Turkey
- <sup>14</sup>The Research Centre for Carbon Solutions (RCCS), School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, EH14 4AS, United Kingdom
- <sup>15</sup>BIGCHEM GmbH, Valerystraße 49, 85716 Unterschleißheim, Germany
- <sup>16</sup>Laboratory of Chemical Physics, National Institute of Diabetes and Digestive and Kidney Diseases, National Institutes of Health, Bethesda, Maryland 20892, United States
- <sup>17</sup>Institute of Metallic Biomaterials, Helmholtz Zentrum Hereon, Geesthacht, Germany
- <sup>18</sup>Polymer Reaction Design group, School of Chemistry, Monash University, 17 Rainforest Walk, Clayton VIC 3800, Australia
- <sup>19</sup>Cavendish Laboratory, Department of Physics, University of Cambridge, Cambridge CB3 0HE, United Kingdom
- <sup>20</sup>Department of Chemical Engineering, University of Waterloo, Waterloo, N2L3G1, Canada
- <sup>21</sup>Institute of Chemical Sciences, School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, EH14 4AS, United Kingdom
- <sup>22</sup>Chemical Engineering & Applied Chemistry, University of Toronto, Toronto, Ontario M5S 3E5, Canada
- <sup>23</sup>Dipartimento di Chimica e Chimica Industriale, Unitá di Ricerca INSTM, Università di Pisa, Via Giuseppe Moruzzi 13, 56124 Pisa, Italy
- <sup>24</sup>Chemical Engineering Department, University of Mohaghegh Ardabili, P.O. Box 179,

Ardabil, Iran

<sup>25</sup>Department of Chemical Engineering, College of Engineering, University of Tehran,  
Tehran, Iran

<sup>26</sup>Department of Chemical Engineering, Massachusetts Institute of Technology,  
Cambridge, MA 02139, United States

<sup>27</sup>Department of Chemical and Biomolecular Engineering, University of Notre Dame,  
Notre Dame, Indiana 46556, United States.

<sup>28</sup>Institute of Applied Synthetic Chemistry, TU Wien, Getreidemarkt 9, 1060, Vienna,  
Austria

<sup>29</sup>Institute of Structural Biology, Molecular Targets and Therapeutics Center, Helmholtz  
Munich - Deutsches Forschungszentrum für Gesundheit und Umwelt (GmbH),  
Ingolstädter Landstraße 1, 85764 Neuherberg, Germany

<sup>30</sup>Department of Chemistry and Biochemistry, University of Notre Dame, Notre Dame,  
Indiana 46556, United States.

<sup>31</sup>Laboratory of Materials for Renewable Energy (LMER), Institut des Sciences et  
Ingénierie Chimiques, École Polytechnique Fédérale de Lausanne (EPFL), Rue de  
l'Industrie 17, CH-1951 Sion, Switzerland

\*\*These authors contributed equally

\*Berend.Smit@epfl.ch

July 3, 2024

## Abstract

The current generation of large language models (LLMs), like ChatGPT, have limited chemical knowledge. Recently, it has been shown that these LLMs can learn and predict chemical properties through fine-tuning. In this work, we explore the potential and limitations of this approach. We studied the performance of fine-tuning GPT-J-6B, a public-domain version of the GPT family, for a range of different chemical questions. We find that in most, if not all, cases, this approach outperforms the benchmark (random guessing) for a simple classification problem. Depending on the size of the dataset and the type of questions, we can also address more sophisticated problems. The most important conclusions of this work are that, for all datasets considered, their conversion into an LLM fine-tuning training set is straightforward and that fine-tuning with even relatively small datasets leads to predictive models. These results suggest that the systematic use of LLMs to guide experiments and simulations will be a powerful technique in any research study, significantly reducing unnecessary experiments or computations.

## 1 Introduction

The traditional machine-learning workflow starts with the painstaking process of harvesting the literature for the relevant data. Some help can be obtained from text harvesting programs.<sup>1</sup> These data can be used to find correlations in the properties or synthesis of molecules or materials or correlations in any other relevant chemical question. For this, it is crucial to describe the system with features fed into a model. Ultimately, the trained model allows us to make predictions from the features of unknown materials. These models typically improve when more data becomes available.

In chemistry and material science, however, the amount of experimental data is often, if not always, a bottleneck. Therefore, it is essential to have some leverage. One way of doing this is by expanding a dataset with computer simulations.<sup>2</sup> Alternatively, we can leverage knowledge of the system. For example, suppose we want to predict the pressure of a gas at a given density and temperature; we can focus our machine learning (ML) on predicting the deviations from the ideal gas law.<sup>3</sup> Another option is to introduce descriptors with proper inductive biases that capture our understanding of the underlying systems.<sup>4</sup>

Another way of leveraging knowledge is through transfer learning. Imagine that one has a lot of data on some particular properties of a class of materials but, as is typical in many practical applications, not enough data for the property of interest. An interesting observation is that these models can be pre-trained (without supervision) on web-scale data by simply predicting the next likely word to complete sentences. The (limited) available experimental data for the property of interest can then be used to fine-tune this pre-trained model.<sup>5</sup> The mechanism of this fine-tuning or transfer learning is that one only re-trains a small part of the transformer model (or of an added layer) and hence leverages all the pre-trained information locked in the model's part that remains unchanged. This fine-tuned model can then be used to make predictions for these properties.

In this context, a remarkable recent discovery is that one can also use fine-tuned large language models (LLMs) to answer chemistry and material science questions for which the base LLM would not know the answer.<sup>6,7</sup> For example, if we use GPT-3 (e.g., via ChatGPT) to ask a specific chemical question, say, if the high-entropy alloy  $\text{Tb}_{0.5}\text{Y}_{0.5}$  is a single phase, it will reproduce the knowledge it has. GPT-3 would not know the answer (GPT-4 knows more chemistry<sup>8</sup>), hence it will likely not get an answer on such chemical questions. However, we can fine-tune an LLM with experimental data of high-entropy alloys, of which we know whether it is a single phase or not. This gives us a new model that only aims to predict

whether a particular high entropy alloy is a single phase.

In addition, the fact that these LLMs use natural language as input instead of a descriptor is one of its most attractive features; it creates a convenient way for researchers to interact with data and tools. Numerous successful chemical applications exploiting this power of LLMs exist today, ranging from tools that summarize literature to the deployment of “chat-bots” for experimental instrumentation.<sup>9</sup>

These general-purpose LLMs are important because they do not require pre-training, can be used for any chemical question, and do not require knowledge of machine learning. In our previous work, we measured the potential of LLMs in solving chemical problems against conventional machine learning.<sup>7</sup> We showed that LLM models fine-tuned on classification, regression, and inverse design problems can be competitive with current state-of-the-art machine learning models. For this, we searched for chemical problems with a known ML solution and validated our approach against it.

In this work, we want to go a step further and attempt to address relevant chemical questions from a more practical point of view. This implies that most of the data have not been curated or selected previously for machine learning studies but are the data that researchers have at hand. The case studies we presented are guided by the questions these researchers have.

With the data, we performed simple “experiments.” First, we asked whether fine-tuned LLMs show any signs of learning. To address this question, we split the dataset in half and used a simple classification to test if the model would classify the data correctly on a holdout set that the model has not seen in training. This first step was fine-tuning the model using a standard setting without any optimization. For this step, we used an open-source model (GPT-J-6B, 6 billion parameters), which we tuned using parameter-efficient fine-tuning techniques.<sup>10-12</sup> The model showed some learning for almost all problems. The extent of learning depends on the dataset and the complexity of the question. Inevitably, the performance of such models will not be optimal for every study; therefore, we optimized the models by performing basic hyperparameter optimization in those cases.

In the next sections, we outline the methodology and then summarize the main conclusions of each case study. The corresponding section of the SI provides a detailed account of each study. The main aim of these summaries is to illustrate the range of chemical questions that can be addressed. The discussion section summarizes the lessons we have learned from these case studies.

## 2 Methods

Researchers from different disciplines presented 21 datasets. We used these datasets in case studies to better understand the potential and limitations of fine-tuning LLMs. For these case studies, we used the open-source LLM GPT-J-6B.<sup>13</sup> Rapid advancements in this field lead to fast improvements and continuously updated models. While we recognize that there are arguably more suitable models available today, we continued to perform experiments with GPT-J-6B to maintain consistency and obtain comparable results in all case studies. GPT-J-6B is a smaller model than GPT-3 (and GPT-4). Our previous results show that GPT-3 typically performs better, i.e., it requires less training data to get similar performance.<sup>7</sup> However, this increase in accuracy does not compensate for the fact that with GPT-J-6B, anybody can reproduce our results. To structure the presentation of the results, we have organized the case studies into three categories: materials and properties, reactions and synthesis, and systems and applications.

Each case study was approached similarly. The first step requires converting the dataset into a set of questions and answers that we can use for fine-tuning. To do this correctly, we obtained some general knowledge of the systems' scientific background. This background is given in more detail in the SI and summarized at the beginning of each paragraph describing the case studies. In the SI, one can also find more details on how the dataset was obtained.

The first test we carried out was a standard test to determine if our fine-tuned LLM learned anything. This test was a simple classification problem in which we split the dataset into two equally populated categories. Depending on the case study, these categories were high/low, good/bad, optimal/non-optimal, etc. This simple classification allows for a simple benchmark: random guessing. The minimum criterion the LLM should outperform on the test set is to do better than random guessing. This random guess corresponds to the situation where we have zero knowledge of the system. We will refer to this experiment as the "base case," in which random guessing was the benchmark. Hence, any model that does better might be of practical use.

If our base case model outperformed the benchmark, the next step is to make the model more useful. In most practical applications, one has more data on poor-performing materials than optimal materials. However, for fine-tuning LLMs, one also needs a reasonable number of materials above the performance threshold, which distinguishes poor from top-performing materials. This may require more data than we have available for a specific case study and may require us to optimize the model further. This part will be specific to each

case study.

We follow the same fine-tuning method as in our previous work (the reader is referred to the original work<sup>7</sup> for details on the fine-tuning) except that we only used GPT-J-6B. In this procedure, the chemical context is formatted in a single representation as a question (Table 1). The binary class is given as a numeric value, i.e., 0 or 1, representing the respective chemical property.

**Table 1. Example of a prompt used to fine-tune an LLM.** This example is a typical classification problem where we have split the dataset into two groups. In the example, we have a <Property> of a <Material>, which can either be high or low. What is High or Low is determined by the threshold we want to use. In practice, we want the model to predict 0 or 1, so in all training sets, we see that good/bad, single phase/two phases will be translated to 0 or 1.

Representation	Completion	Real
What is <Property> of <Material>?	0 or 1	Low or High

In the first iteration, default fine-tuning hyperparameter values were used (see Supplementary Note 1). This allowed us to gain some insights into whether such an approach can be used as a black box without expertise in using LLMs or if some tweaking is needed to get sufficiently accurate results. After analyzing the first result, in some case studies, increasing the number of epochs, i.e., the times the model sees the training data, significantly increased the model's performance. This gives us insights into the fine-tuning procedure. This second step typically requires some more experience with these LLMs.



## 3 Results and Discussion - Case Studies

### 3.1 Materials and Properties

Following a bottom-up approach, most chemical applications start from fundamental research of the structure-property relationships. Therefore, it is no surprise that a thorough understanding of the structures at hand is necessary before proceeding to field-specific applications. However, the chemical space is vast and complex, and finding optimal solutions often requires extensive screening and analysis. Alternatively, chemical properties can be predicted from the structural features of a molecule.<sup>14</sup> We demonstrate how our approach can help predict chemical properties in various case studies and more effectively guide computational and experimental research.

**Adhesive Energy of Polymers** As a first example, we experimented with a computational dataset of polymers and their respective adhesive free-energy on a polymer surface (see Supplementary Note 2.1).<sup>15</sup> Here, the question arises if we can predict the adhesive free-energy from any hypothetical copolymer sequence. The polymers in the dataset were chains of 20 monomers, described as either an “A” or a “B” unit. We used the sequence string of the polymers to predict the balanced binary classification of the adhesive free-energy (i.e., is this free energy high or low). The fine-tuned LLM had an accuracy of 88%, which is 38% above the random baseline (50%).

The remarkable aspect of these results is that we have a hypothetical model polymer for which simulations compute the free energies. Yet, the LLM can correlate a sequence of 20 (arbitrarily chosen) characters of the type “A” and “B” to the free energy, suggesting no potential data leakage.

**Properties of Monomers** Focusing on a more standardized and widespread descriptor of molecular structures, we investigated the Simplified Molecular-Input Line-Entry System (SMILES) notation. These textual strings capture the elemental composition, bonds, branches, and stereochemistry of chemical compounds. The monomer database, computationally generated by Schneider et al.<sup>16</sup>, served as an ideal test case to validate the synergy of SMILES and LLMs (see Supplementary Note 2.2). Schneider et al.<sup>16</sup> obtained from simulation many different properties, including the glass transition temperature, cohesive energy density, squared radius of gyration, and density of a wide range of monomers. We obtained four unique binary classification case studies by taking the median for every property. We fine-

tuned LLMs to predict the specific property from the monomer's SMILES. For all cases, an accuracy above 80% (average accuracy of 86% over the four properties) with non-optimized hyperparameters was obtained.

**Melting Point of Molecules** The following case study concerned the prediction of the melting point of small molecules (see Supplementary Note 2.3). The 274983 structures were all represented by their SMILES notation and IUPAC name. Therefore, in this particular example, we further explored the chemical representation and how it affects the quality of the predictions.

As the melting point of many chemicals is reported, we first studied how well ChatGPT (OpenAI's GPT-3.5) can classify the melting point as high or low. Using the frontend interface, we prompted "What is the melting point of ;name of molecule;?" and saw that only 50% of the time, it predicted this correctly, which is no better than random guessing. In contrast, models trained on the IUPAC name reached an accuracy of 60%. Interestingly, our fine-tuned models trained on the SMILES of the molecules could predict the melting point with an accuracy of 69%.

**Dynamic Viscosity of Molecules** In this case study, we also used the SMILES notation of some chemicals. Our objective was to predict their dynamic viscosity with fine-tuned LLMs (see Supplementary Note 2.4).

As the dynamic viscosity value of many chemicals is also reported, we evaluated (via ChatGPT) how well OpenAI's GPT-3.5 model can classify the viscosity as high or low. Our prompt was "What is the viscosity of ;name of molecule;?" Our results showed that viscosity was not better predicted than random guessing, with an accuracy of 55% when the chemical name was provided as input to the model.

In contrast, for a median split balanced dataset, with a training set size of 80 examples and 30 fine-tuning epochs, our fine-tuned LLM model reached an accuracy of 80% for binary classification. We also trained a model to predict whether a chemical had a dynamic viscosity in the top 28% of the values in the dataset. After reducing the dataset size to obtain a balanced dataset, we also obtained a predictive accuracy of 80% using a training set of 50 data points by increasing the number of fine-tuning epochs to 140.

**Microstructural Properties of Magnesium Alloys** Due to their lightweight, Mg alloys gain popularity in structural applications where weight saving is of importance.<sup>17</sup> Besides a small

portion produced with powder-metallurgical processing,<sup>18–20</sup> the majority of these alloys are cast and subsequently subjected to thermo-mechanical treatment to obtain a microstructure corresponding to a suitable property profile.<sup>21,22</sup> To understand the connection between process routes, their specific processing parameters, and microstructure evolution, we can use the LLM's modular and versatile way of featurizing to include all relevant parameters, independent of the production route (see Supplementary Note 2.5).

We found that the fine-tuned models can deal with incomplete and multivariable inputs, reaching accuracies of 84% to predict the classification of the material either belonging to the class high or low amount of second phases. Interestingly, we acquired similar accuracies when only using the production route to represent the model irrespective of the individual process parameters. Despite the small dataset, the LLM is able to catch the material science properties and classify them accordingly.

**Phase Separation Propensity of Proteins** Phase separation of proteins and other biomolecules is recognized as an important intracellular process that affects cellular compartmentalization and regulation.<sup>23</sup> However, the mechanisms that drive biomolecular phase behavior are still under active investigation. Saar et al.<sup>24</sup> performed an *in silico* study to understand the link between protein sequence and its liquid-liquid phase separation (LLPS). This is an interesting challenge from a data-science perspective, as the protein sequences are long strings of letters, each representing a single amino acid. Such non-numeric input often requires additional data pre-processing steps and/or dedicated statistical techniques. Saar et al.<sup>24</sup> developed binary classifiers based on extracted physical features of the protein sequence and a word embedding of the sequence made using a word2vec model. From these sequence-based embeddings, the model was able to classify proteins based on their propensity to undergo LLPS. Identifying proteins capable of undergoing LLPS into protein-rich biomolecular condensates is important for understanding cellular function and pathology.

This is an interesting case for an LLM model, as a protein represented by a string, like RRGDGRRRG . . . GGGRGQGRGR, can be inputted directly in the prompt (see Supplementary Note 2.6). We obtained accuracies reaching 95% for models that distinguish proteins on their phase separation propensity, which is similar to the accuracy obtained by Saar et al.<sup>24</sup>. We want to stress that no extra data manipulation was needed. The protein sequence as received was used as input for the prompt, again demonstrating the versatility of LLMs.

In addition, we carried out some experiments in which we changed the original sequences (e.g., making them shorter or creating randomized sequences of the same letter/aminoacid

composition). The most interesting observation was that a model trained on randomized sequences of the same letters resulted in a relatively small drop in accuracy from 95% to 86%, which shows that a significant part (but not all) of the predictive ability can be obtained from the protein's sequence composition without any positional information of sequence order. Interestingly, the addition of positional information was also not found to increase the performance in predicting the apparent and shear modulus of materials.<sup>25</sup>

**Structure of Nanoparticles** Developing new nanomaterials for energy technologies requires a deep understanding of the intricate relation between material properties and atomic structure. Solving the atomic structure of nanomaterials from their x-ray total scattering data is challenging. Generative models such as Conditional Variational Autoencoder (CVAE) have been proposed to obtain valid chemical structures from the scattering pattern.<sup>26,27</sup> In this case study, we predict the structure type of nanomaterials and the number of atoms in the material from their scattering pattern. Predicting these values accurately is an easier task than solving the structure, and the LLMs are more convenient to use for researchers than CVAEs as they are purely based on natural language.

This is also an interesting case for an LLM model since the scattering pattern consists of a very long series of numbers represented as a string. With our approach, we obtained an accuracy of 94% to predict the structure type of nanoparticles from scattering patterns simulated from 7 highly unbalanced structure types with between 5 and 100 atoms (30 epochs, 1800 data points). We found that, for complex input variables, where the information is embedded along long sequences, using a relatively large training set size, the fine-tuned model can predict an unbalanced dataset with 7 classes. However, if the number of training data points is very low (200), the fine-tuned model is not even predictive on a balanced dataset.

For the prediction of the number of atoms in the nanomaterial, we obtained accuracies of 91% and 85% for datasets with 4 and 10 balanced classes. However, from a practical point of view, given the interest in predicting the number of atoms with very high accuracy, we also developed a regression model. This model predicted the number of atoms with an  $R^2$  of 98% and a maximum absolute error (MAE) between 1 and 2 atoms (XGBoost baseline:  $R^2 = 92\%$ , MAE = 5.6), i.e., with excellent performance.

## 3.2 Reactions and Synthesis

Going beyond chemical properties, we explored the potential of our methodology in predicting reaction outcomes. Predicting reaction outcomes is a field in which conventional theory has made little progress. The complexity and specificity of chemical reactions make developing a general theoretical framework extremely challenging.<sup>28</sup> Practically, indicating the success rate or yield could prioritize the synthesis of top candidates or omit protocols that lead to certain failures, thereby saving resources, time, and money.

**Activation Energy of Cycloadditions** A first example in this regard is based on a previous study that screened the influence of substituents of one of the reactants on the activation energy of bioorthogonal tetrazine-alkene cycloadditions. From a total of 966 different reactants, the free energy barrier was computed using density functional theory (DFT, see Supplementary Note 3.1). We took the median free energy as the threshold for ‘good’ reactants for the specific reaction and the SMILES notation to represent the molecule in question. With a training set size of 250 examples, our LLM model reached an accuracy of 94%, significantly higher than a random guess, i.e., 50% for a balanced dataset.

This case study is a great example of how machine learning, specifically LLMs, can significantly impact expensive computational studies. This model can be used as a first screening to filter out poor-performing structures. More expensive calculations can then be used for a more detailed analysis of well-performing structures.

**Free Energy of Catalyzed Cleavage Reaction** In the same category of organic reaction, a reaction dataset on a Ni-catalyzed aryl ether cleavage reaction was explored.<sup>29</sup> (see Supplementary Note 3.2). Similarly, a large computational study was performed to gain knowledge about the efficiency of a set of catalysts. Our LLM-based approach has the potential to minimize expensive calculations and predict whether new molecules might be a suitable candidate for this reaction. As we know the chemical structure of the catalysts, we can base our predictions on the SMILES notation as inputs. For a median split base case, an accuracy of 90% was reached.

While excellent in performance, the scientific relevance of evaluating these systems based on a single threshold is limited. Rather, a small range of continuous descriptor values is often considered ‘good’. All values above and below this range are then considered ‘poor’. From the dataset of catalysts, only 3.8% was labeled a ‘good’ catalyst. As a result, we were

forced to reduce significantly our training set to get a balanced dataset. Nevertheless, even with a training size of 100 data points, the model was able to correctly classify 79% of the test data.

An interesting future strategy would be to use an LLM model combined with more expensive quantum calculations. Initially, one would aim for a band of ‘good’ structures that is broader than one would like from a catalysis point of view but more balanced to get a more accurate model. Then, if we get more ‘good’ materials, we would retrain the model with a narrower band.

**Yield of Catalytic Isomerization** A catalytic isomerization was examined in this case study (see Supplementary Note 3.3). The dataset is part of a scoping study where the catalytic activity of PtBr<sub>2</sub> was assessed. Apart from the starting material, all reaction conditions were kept constant. The amount of data is limited, with 16 experimental entries. In addition, the noise in the data is expected to be rather high as both the starting material and the product are not very stable. We were interested in predicting the success of the isomerization based on the yield (>50%). Even after optimizing hyperparameters, no valuable models could be created.

Despite the low predictive power of the models, this case study exemplifies an untouched application of LLMs in chemical research, namely scoping studies. Such studies are traditionally performed to gain a thorough understanding of whether a reaction is successful and how efficient the reaction is. While the amount of different substrates tested can be large, they often still screen a small fraction of the chemical space. Predicting the outcome of similar structures based on a set of experimentally assessed reactions could accelerate material discovery.

**Kinetics of Polymerization** This case study focused on an experimental kinetic polymer screening (see Supplementary Note 3.4). The dataset contained 23 entries with different monomer-monomer concentration combinations, each with their experimental measured kinetic profile. Our objective was to predict the polymerization rate from the reaction conditions, i.e., we used the monomer representation and the monomer concentration as input variables. Given the smaller size of the dataset, hyperparameters were modified slightly to obtain acceptable performances. When comparing different representations of the monomer, we see that fine-tuned models based on the SMILES of the monomer (accuracy of 76%) do better than models based on the IUPAC name of the monomers (accuracy of 57%). We also

noticed that reactions with a polymerization rate close to the binary split threshold, i.e., the median of the dataset, remained difficult to predict.

Not only are high-throughput kinetic screenings an excellent way to gain in-depth insights into reaction mechanisms, but they also produce datasets that can be used to train ML models and guide further research and development. Here, only two reaction parameters were varied. Combining screenings in a multi-parameter landscape with predictive models could accelerate polymer synthesis optimizations.

**Photocatalytic Water Splitting Activity of MOFs** Following with reactions, in this case, we explored a dataset containing 95 MOF structures with different properties related to photocatalytic water splitting, as obtained from DFT calculations (see Supplementary Note 3.5). Such DFT calculations require significant computational resources.

In this study, we predicted various photocatalytic properties of MOFs, thereby assessing whether a given material has the right band alignments for water splitting and absorbs visible light. We used the elemental composition of the MOF's linker and metal node to represent the material. The fine-tuned models could successfully predict the various properties with accuracies of  $\geq 85\%$ .

**Photocatalytic CO<sub>2</sub> Conversion Activity of MOFs** We explored another example of the use of MOFs, in this case study for photocatalytic CO<sub>2</sub> conversion (see Supplementary Note 3.6). The dataset ( $n = 77$ ) contained the catalyst system (metal source, linker, phase, sacrificial agent, and cocatalyst), the band gaps of the MOF and cocatalyst, and the photocatalytic activity studied. This data allowed us to investigate the use of different parameters in the prompt as predictors of the photocatalytic activity of MOFs. When we used the catalytic system, we obtained an accuracy of 65%. When we combined the SMILES notation and the catalytic system parameters in the prompt, for a total of six features as predictors, the accuracy increased to 68%. However, the predictions of this model were not better than random guessing for samples with higher values of photocatalytic activity. We probably need more data to predict the outcome of complex processes such as photocatalysis. On the other hand, adding the values of band gaps, conducting band, and valence band to the prompt, i.e., seven extra features, did not increase accuracy (58%).

**Success of MOF Synthesis** In this case related to the synthesis of materials, an interesting dataset of MOFs was investigated (see Supplementary Note 3.7). The objective here was to



predict the success of the synthesis of a MOF given experimental parameters extracted from reaction protocols. Interestingly, the majority of reported protocols often lead to the desired product, hindering the creation of a balanced dataset and subsequently making unbiased predictions hard.

This quantity issue is reflected in the provided dataset, which has only 25 different reaction conditions. Taking a yield of 20% as the threshold for success could create a fairly balanced dataset. After training for just 10 epochs on a training set size of 20 examples, the fine-tuned models could not recognize the prompt/completion structure and thus failed to output a meaningful prediction. Using a training set size of 20 and increasing the number of epochs to 50 leads to the expected binary responses, i.e., 0 or 1, with an average accuracy of 89%.

### 3.3 Systems and Applications

Here, we explore the potential of LLMs to predict the outcomes of different systems and applications.<sup>30</sup> Knowing the effect of processing parameters on process performance can help to optimize a system and increase its efficiency. Gaining knowledge of the influence of the experimental conditions on the results can allow us to tailor the process according to our specific objectives.

**Gas Uptake and Diffusion of MOFs** In this case, an extended version of the SMILES notation of MOFs was explored.<sup>31</sup> Here, the MOFid, a combined string of the SMILES of the individual building blocks that construct a MOF, served as the chemical descriptor in predicting its gas uptake and diffusion (see Supplementary Note 4.1).<sup>32</sup> In an extensive computational study, Daglar and Keskin<sup>31</sup> simulated hydrogen, nitrogen, helium, and methane gas uptake and diffusion in more than 5000 structures, which served as our dataset. These eight individual case studies yielded an average accuracy of 68% for predicting the different properties. Notably, these results are in line with the models in the original work.

Some MOFs (e.g., ZIFs) have different isomers with the same chemical building blocks. Therefore, it is interesting to investigate whether adding further details on the structure in the prompt will improve learning. Apart from the MOFid, and the uptake and diffusion values, Daglar and Keskin<sup>31</sup> included 20 additional simulated features of the MOF structures, all of which are numeric and grouped based on their chemical and physical relevance. We used a combined feature vector (per group) to create a prompt for predicting the binary



class, i.e., above or below the median, for helium diffusion. In the first experiment, we included the largest cavity diameter, pore limiting diameter, and the pore size ratio, i.e., group A. Secondly, a prompt with the density, pore volume, porosity, and surface area, i.e., group B, was created. For these “single group” experiments, we obtained an accuracy of 68% and 62%, respectively, for groups A and B. Only when we combine groups A and B, thus creating a prompt with seven features, we obtained a performance improvement (73% accuracy). Adding eight extra features related to the elemental composition increased the accuracy to 77%.

**Hydrogen Storage Capacity of Metal Hydrides** A potential replacement for fossil fuels is hydrogen. One of the disadvantages of hydrogen compared to fossil fuels is its low energy density. Finding ways to store hydrogen is, therefore, an important research theme. Metal hydrides are a promising class of materials for their capacity to store hydrogen.<sup>33</sup>

The heat of formation of a metal hydride is often used as an indicator for their potential hydrogen storage use. Theoretically, this value is related to the equilibrium pressure. We, therefore, started our experiment by validating if fine-tuned LLMs could capture this relation. The ML-HydPARK dataset created by Witman et al.<sup>34</sup> contains 430 metal hydrides with their respective heat of formation and equilibrium pressure (see Supplementary Note 4.2). In these initial experiments, we used the median heat of formation as the threshold for the binary classification. We fine-tuned an LLM that could answer the question, “Is the heat of formation, and thus its potential for hydrogen storage, of a metal hydride with an equilibrium pressure of  $\mu$ value high or low?”. Such models indeed predicted the heat of formation from a material’s equilibrium pressure with an accuracy of 76%.

In an alternative approach, we hypothesized that the metal in the material is an indicator of success. We substituted the equilibrium pressure with the elemental formula of the material and repeated the training. Instead of a numeric feature, we now describe the material with a simple textual string, fully exploiting the potential of LLMs. These binary classification models performed significantly better, with an accuracy of 84%. A possible explanation for this increase in performance might be rooted in the augmented information present in the chemical composition of the material. When we combined both the pressure and the chemical formula in the feature vector, we saw a slightly higher accuracy than with the model trained on only pressure values, suggesting that the additional chemical information had extra predictive power.

From a practical point of view, a realistic threshold for defining promising materials

would be useful. As suggested by the literature, this range (heat of formation values between  $-40 \text{ kJ mol}^{-1}$  to  $-20 \text{ kJ mol}^{-1}$ ) created a slightly unbalanced dataset. Nevertheless, acceptable performances were still achieved with accuracies of 75%.

**CO<sub>2</sub> Adsorption of Biomass-derived Adsorbents** In this example, we investigated a dataset on the synthesis of activated carbons from different biomass precursors for CO<sub>2</sub> capture by adsorption processes<sup>35</sup> (see Supplementary Note 4.3). From these data, we analyzed two binary classification case studies to predict whether the BET surface area and CO<sub>2</sub> adsorption capacity of biomass-based adsorbents are high or low.

The dataset contains data on 33 biomass precursors and ten activating agents. We fine-tuned LLMs to predict the BET surface area and CO<sub>2</sub> adsorption capacity from the biomass precursor, activation conditions, adsorbent textural properties, and adsorption conditions. An interesting aspect of this example is that, unlike in conventional machine learning models that require conversion to numerical values, the biomass precursor's name and the activating agent's chemical formula were entered as textual strings into the model.

By taking the median as the threshold to classify adsorbent materials, an accuracy of 84% was obtained for the prediction of CO<sub>2</sub> adsorption capacity from precursor activation conditions, activated carbon textural properties, and adsorption conditions, with non-optimized hyperparameters. Since the dataset is smaller (training set of 65 data points), we had to increase the number of fine-tuning epochs from 30 to 140 to predicting the BET surface area from precursor activation conditions with an accuracy of 72%. We also found that models trained without the precursor name performed slightly worse than models trained on the full feature vector, indicating that the model could also learn some trends associated with the biomass name.

Under a more practical classification, we also evaluated a threshold value that would allow us to predict which materials are really 'good,' which forced us to reduce the training set to obtain a balanced dataset. Nevertheless, CO<sub>2</sub> adsorption capacity was predicted with an accuracy of 82% by increasing the number of fine-tuning epochs to 100. Likewise, using a smaller dataset, BET surface area was predicted with 75% accuracy by increasing the number of fine-tuning epochs to 200.

**Thermal Desalination of Water** In this case study, we focused on thermal desalination of saline or brackish water sources (see Supplementary Note 4.4). Knowing the behavior of thermal desalination units is crucial to optimize their design.<sup>36</sup> Here, we evaluated two

case studies to predict the Gain Output Ratio (GOR), which is a measure of the thermal energy utilization efficiency, and the specific heat transfer surface, which will determine the size of the plant of a solar desalination system from the number of effects and the steam temperature.

The dataset was split to obtain a balanced binary classification problem. The relatively small size of the dataset forced us to use a maximum training set size of 25 example prompts. The first model with non-optimized hyperparameters showed no predictive power. By increasing the number of epochs from 30 to 100, we obtained an accuracy of 83% for the specific heat transfer surface and 92% for GOR.

We also trained binary classification models using unbalanced datasets to simulate a more realistic case for finding top-performing conditions. The model did not perform better than the random guessing baseline of 80% to predict the specific heat transfer surface, given an accuracy of 80%, but showed acceptable performance in predicting GOR, with 93% accuracy.

**Detection Response of Gas Sensors** Following with applications, in this example, we analyzed a dataset on the sensing behavior of gas sensors (see Supplementary Note 4.5). Here, the objective was to predict the detection response of gas sensors by training a model using data from 56 different experimentally analyzed conditions, varying three design parameters: sensor type (i.e., core-shell and composite sensors), zinc oxide concentration, and operating temperature.

A binary classification model trained on 45 example prompts (100 epochs) could predict whether a given sensor was in the top half performing conditions with 89% accuracy.

**Stability of Gas Sensors** In the field of gas sensor applications, we also explored a dataset on the long-term stability of gas sensors (Supplementary Note 4.6). In this case, the objective was to predict whether a SnO<sub>2</sub>-based gas sensor is stable or not as a function of the type of dopant material, its dosage, and the calcination temperature during synthesis. By accurately predicting stability, a more efficient search for ideal sensors could accelerate the field of gas sensing.

With a rather small dataset of 19 data points, this case study tested the limits on the size of the training set. With slightly optimized hyperparameters by increasing the number of epochs to 120, a predictive accuracy of 71% was achieved for a binary classification model that was able to predict whether a sensor was stable, i.e., had a response loss between days 5

and 15 of less than 12%.

**Gasification of Biomass** A dataset on the biomass gasification process was also used to validate our predictive framework in applications,<sup>37</sup> such as thermochemical conversion of biomass to produce energy carriers. The data described the H<sub>2</sub>/CO ratio in the syngas obtained from the gasification of solid biomasses (see Supplementary Note 4.7). Here, we are interested in predicting whether the gasification process of a given biomass gives a H<sub>2</sub>/CO ratio higher than 1.8, meaning that it is suitable for fuel and chemical synthesis. Since data points with such values only represent 30% of the overall dataset, we were forced to reduce its size to obtain a balanced dataset. With a relatively small training set of 25 data points, we increased the number of epochs to 140 to obtain an accuracy of 70%, i.e., higher than the benchmark value of 50% (random guess).

## 4 Discussion

In chemistry, we mostly have to deal with a limited amount of data. Hence, it is essential to use leverage in any machine-learning approach applied to chemistry. In an LLM, one leverages the linguistic nuances, patterns, and knowledge captured in correlations harvested from large quantities of internet text. This new source of data, combined with using natural language to interface with the model, makes this approach, potentially, more powerful than machine learning models trained only on conventional data sources.

In this work, we tried to obtain some insights into the performance of such LLMs by looking at 21 case studies describing many different systems, ranging from predicting simple thermodynamic properties to device performance. The obvious question is whether it works. In this section, we try to answer these questions in parts.

We must remember that the original corpus of text used to train these models was not specifically curated for chemical questions. It is remarkable that we can create specific solutions for a range of chemical subfields, spanning from a molecular level to reaction kinetics to high-end applications.

In most of these case studies, the LLMs demonstrated their ability to predict basic structure-property relationships. Various cases concerning reactions showcased that LLMs can predict reaction outcomes and yield determination, thereby facilitating reaction optimizations, scoping studies, or catalyst designs. In our applied chemistry cases, the versatility of our

approach was further underscored by predicting system parameters, thereby assisting the optimization of real-life chemical processes.

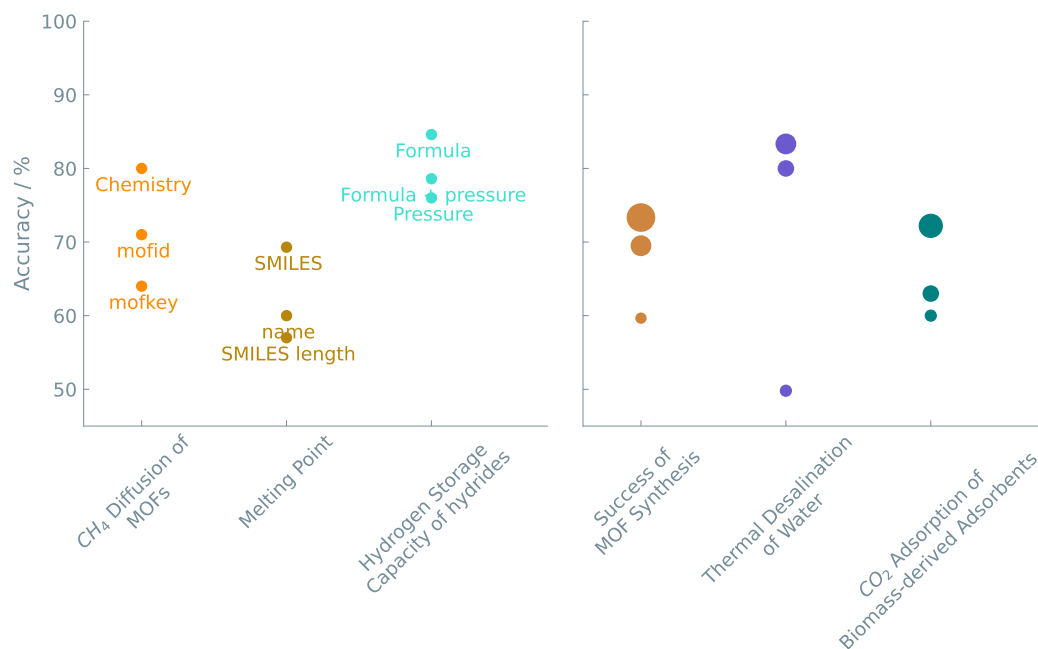
Our results also make clear that the LLM approach works best with a reasonably balanced dataset. However, in practice, one is often interested in the (small) subset of top-performing materials, and we observe that the training set quickly becomes too unbalanced to make sufficiently accurate predictions. The solution to this problem is to start with a model trained on a less narrow window. We typically observe that depending on the size of the dataset, this approach is better than random guessing. So, for problems that are too complex or for which we do not have any intuition, we already gain. If we then collect more data in the region of interest, we can narrow the window, making the model increasingly useful.

Our study also provided useful insights into more specific issues related to featurization, feature importance and the size of the dataset.

**Featurization** Can we translate the prediction of the properties of a material or a chemical reaction into a set of simple questions and answers that can be used to fine-tune an LLM? Conventional machine learning requires featurization, i.e., converting the chemical system into a feature vector that quantifies the similarity between systems one wants to compare. Especially in experimental datasets, feature extraction is not as obvious as it might seem. Indeed, many tools and molecular fingerprints exist, but these often pose a burden for non-experts aiming to integrate machine learning into their workflows.

As the case studies show, translating a chemical question into a prompt for fine-tuning is straightforward. The primary challenge is choosing how to represent a material or chemical. One can try one of a number of different representations or even use a combination of such representations (see fig. 1). Standardized notations like SMILES can be exploited to represent chemical structures in LLMs. The readability of SMILES strings makes them convenient for researchers and chemical toolkits to interpret. We show that text-based descriptors like SMILES (see ‘Melting Point of Molecules’ study), MOFid (see ‘Gas Uptake and Diffusion of MOFs’ study), or even non-standardized strings (see ‘Adhesion Energy of Polymers’ and ‘Structure of Nanoparticles’ studies) perform well in connecting structural information with physical/chemical properties or reaction outcomes.

Machine learning approaches become even more powerful when dealing with multiple variables. Thus, we extended our prompts with additional data to allow for multi-variable predictions. For instance, in the ‘Hydrogen Storage Capacity of Metal Hydrides’ study, we



**Figure 1. Accuracy of the different representations.** The color indicates the particular case study. The same color coding is used in fig. 2 to compare between all case studies. In the left figure, the annotations are the representations used in the prompting. In the right figure, the size of the circles is related to the number of epochs that are used.

combined molecular information and equilibrium pressure in one prompt to predict a material's heat of formation. Interestingly, we noticed that this longer prompt outperformed the model that was only trained on the pressure data. Similar trends are also seen in the 'Gas Uptake and Diffusion of MOFs' and 'Photocatalytic Water Splitting Activity of MOFs' studies. This methodology becomes particularly interesting for predicting the experimental success of a synthetic reaction. Reported reaction protocols are generally described as a combination of textual (e.g., reagents, solvent system) and numeric (e.g., reaction time, temperature) data. In the 'CO<sub>2</sub> Adsorption of Biomass-derived Adsorbents' study, the dataset consisted of 8 variables used to fine-tune the models. Again, models trained without one of the textual variables performed slightly worse, highlighting the synergy between text-based data and LLMs.

**Feature Importance** Feature importance is often an interesting analysis of a trained ML model to know which features carry the most weight in the model. Synthetic chemists also do this daily by asking, for example, “Which parameter do I most likely need to vary to get the desired result?”. Combining natural language and machine learning, i.e., using LLMs, could facilitate this process and create an understandable approach to optimizing chemical systems. By iteratively removing one particular feature in the multi-variable prompt and assessing the accuracy of the resulting fine-tuned model, the influence of the respective parameter on the final objective can be evaluated. We used this approach in the ‘Gas Uptake and Diffusion of MOFs’ dataset. We saw high accuracies when structures were represented as 20 individual features. By removing descriptors, we noticed a drop in accuracies, hinting at the importance of the omitted descriptors on the final predictions.

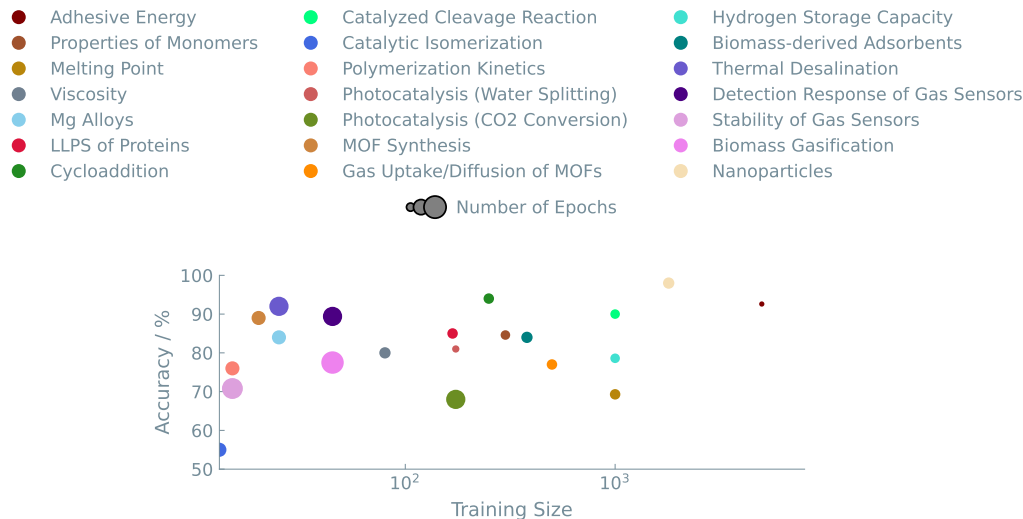
**Size of the dataset** Data quantity, as well as quality, plays an important role in fine-tuning LLMs. We explored datasets ranging from as few as 20 to as many as 5000 entries. Fig. 2 summarises the accuracies that we obtained for the different datasets. We see that our approach works well for datasets in the low-data regime. A consistent trend over all experiments suggests that models trained on larger datasets have excellent predictive performance. We typically get accuracies above 80%, allowing us to make balanced training sets if one wants to identify a (much smaller) sub-set of interesting materials.

On the other end of the plot, i.e., really small datasets ( $\leq 20$  data points), the LLMs initially have difficulties predicting any meaningful output. For experiments with small datasets, we slightly optimized the hyperparameters to elevate the robustness of the models. By increasing the number of times that the models see the training data, i.e., the number of epochs during training, the performance of these models came close to large dataset models in terms of accuracy (see fig. 1 (right)).

## 5 Conclusions

For most, if not all, of these 20 case studies, the LLMs performed (much) better than random guessing. This is a remarkable result, given that the LLMs were not pre-trained on the bulk of scientific literature. In addition, our effort in making a predictive model using an LLM is modest. We also have to note that GPT-J-6B is by now an old and small model and that we did not carry out a full fine-tuning.

We focused on binary classifications that provide a simple ‘yes’ or ‘no’ answer. We see



**Figure 2. Accuracy as a function of the size of the dataset.** The color indicates the case study and the size of the circles is related to the number of epochs used.

such tasks as a first step; if the LLM did not outperform random guessing, we would conclude that there has been no learning. If we can sufficiently accurately predict such a simple classification, one can proceed to develop a regression model as the next step. Yet, even binary classifications can be useful, especially in experimental settings where a continuous value is often unnecessary to streamline decision-making. An accurate binary classifier can already facilitate various aspects of today’s research. For example, ML-based screenings of a particular chemical system can significantly reduce computational resources or experimental work, e.g., “Is it worth doing this experiment?”.

Even a modest accuracy can be helpful if the alternative is random guessing. Moreover, we also show that these models improve significantly if more data is collected. In this context, we must mention the importance of balanced datasets. In most practical cases, there are many more failures than successful experiments. Hence, in our training, we had to reduce the training set to have a reasonably balanced dataset. If we were to use literature data, we would have the opposite problem. In most, if not all, studies, only successful results are published. Machine learning, like human learning, learns even more from its failures.<sup>38</sup> Thus, if we want to take full advantage of the tools explored here, we need to rethink how data are reported.<sup>39</sup>

In addition to the remarkable performances of the trained models, we also want to stress



that natural language in ML models facilitates various aspects of the case studies. By obviating the need to featurize the chemical system, this use of textual descriptors of molecules points to an attractive alternative interface to chemical knowledge suitable for non-experts. Moreover, we noticed that natural language greatly improves scientific interpretation, effective discussions, and communication between different research fields.

## Data and code availability

The datasets and Jupyter Notebooks used in this work are available at <https://github.com/JorenBE/GPT-Challenge>.

## Acknowledgements

The research of J.V.H., and B.S. is supported by the Swiss Science Foundation through a Project Funding (214872) and Advanced Grant (216165).

M.V.G. and C.P. gratefully acknowledge financial support from the Spanish Agencia Estatal de Investigación (AEI) through Grants TED2021-131693B-I00 (M.V.G. and C.P.) and CNS2022-135474 (M.V.G.), funded by MICIU/AEI/ 10.13039/501100011033 and by the European Union NextGenerationEU/PRTR. M.V.G. acknowledges support from the Spanish National Research Council (CSIC) through Programme for internationalization i-LINK 2023 (Project ILINK23047). M.V.G. acknowledges the access granted by the Galician Supercomputing Center (CESGA) to the FinisTerae III supercomputer, funded by the Spanish Ministry of Science and Innovation, the Galician Government, and the European Regional Development Fund (ERDF), and the access granted by the CSIC to the Drago supercomputer.

Parts of the work of K.M.J. were supported by the Carl Zeiss Foundation.

S.G., M.G., N.P., B.S., and J.V.H. are partly supported by the USorb-DAC Project through a grant from The Grantham Foundation for the Protection of the Environment to RMI's climate tech accelerator program, Third Derivative.

The work of A.S.A. is supported by Novo Nordisk Foundation grant NNF23OC0081359.

S.M.M. and M.R.K. work is partly supported by grant number DSI-CGY3R1P16 from the Data Sciences Institute at the University of Toronto.

M.A. expresses gratitude to the European Research Council (ERC) for evaluating the project with the reference number 101106377 titled "CLARIFIER" and accepting it for funding under the HORIZON TMA MSCA Postdoctoral Fellowships - European Fellowships.

Furthermore, M.A. acknowledges the funding by UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee (EP/Y023447/1; organization reference:101106377).

L.L.G., A.A., and T.P.J.K. gratefully acknowledge funding from the European Research Council under the European Union's Horizon 2020 research and innovation program through the ERC grant DiProPhys (agreement ID 101001615) (L.L.G., A.A., T.P.J.K.). the National Institutes of Health Oxford-Cambridge Scholars Program (L.L.G.), the Cambridge Trust's Cambridge International Scholarship (L.L.G.), the Intramural Research Program of the National Institute of Diabetes and Digestive and Kidney Diseases at the National Institutes of Health (L.L.G.). the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013; T.P.J.K.) and the Frances and Augustus Newman Foundation (T.P.J.K.).

K.L.S. acknowledges funding from the Schmidt Science Fellowship in partnership with the Rhodes Trust and from St. John's College Research Fellowship programme.

F.N. and M.T. thank the Italian MUR for provision of funding through the PRIN 2020 Project doMino (ref 2020P9KKBKZ).

The research of N.X.H. was supported by the NCCR MARVEL, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (grant number 205602).

B.M. acknowledges the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 945363.

## Author contributions

Author	Conceptualization	Data curation	Formal analysis	Funding acquisition	Investigation	Methodology	Project administration	Resources	Software	Supervision	Validation	Visualization	Writing – original draft	Writing – review & editing
Joren van Herk														
Maria Victoria Gil														
Kevin Maik Jablonka														
Alex Abrudan														
Andy S. Anker														
Mehrdad Asgari														
Ben Blaszczak														
Leander Choudhury														
Clemence Corninboeuf														
Hilal Daglar														
Ian T. Foster														
Susana Garcia														
Matthew Garvin														
Guillaume Godin														
Lydia L. Good														
Jianan Gu														
Noemie Xiao Hu														
Xin Jin														
Tanja Junkers														
Seda Keskin														
Tuomas P.J. Knowles														
Ruben Laplaza														
Sauradeep Majumdar														
Hossein Mashhadimoslem														
Ruairadh D. McIntosh														
Seyed Mohammad Moosavi														
Beatriz Mourino														
Francesca Neri														
Covadonga Pevida														
Neda Poudineh														
Mahyar Rajabi-Kochi														
Kadi L. Saar														
Fahimeh H. Saboor														
Morteza Sagharichiha														
KJ Schmidt														
Jiale Shi														
Dennis Svatoněk														
Marco Taddei														
Igor Tetko														
Domonkos Tolnai														
Sahar Vahdatfar														
Jonathan Whitmer														
Florian Wieland														
Regine Willumeit-Romer														
Andreas Zuttel														
Berend Smit														

## Competing interests

K.M.J. and A.S.A. have been contractors for OpenAI (as part of the red teaming network). K.L.S. is a consultant for Transition Bio.

## References

- [1] Kononova, O.; He, T. J.; Huo, H. Y.; Trewartha, A.; Olivetti, E. A.; Ceder, G. Opportunities and challenges of text mining in materials research. *Iscience* **2021**, *24*, DOI: 10.1016/j.isci.2021.102155.
- [2] Han, H.; Choi, S. Transfer Learning from Simulation to Experimental Data: NMR Chemical Shift Predictions. *J Phys Chem Lett* **2021**, *12*, 3662–3668, DOI: 10.1021/acs.jpcllett.1c00578.
- [3] Ramakrishnan, R.; Dral, P. O.; Rupp, M.; von Lilienfeld, O. A. Big Data Meets Quantum Chemistry Approximations: The  $\Delta$ -Machine Learning Approach. *J Chem Theory Comput* **2015**, *11*, 2087–2096, DOI: 10.1021/acs.jctc.5b00099.
- [4] Margraf, J. T. Science-Driven Atomistic Machine Learning. *Angew Chem Int Edit* **2023**, *62*, DOI: 10.1002/anie.202219170.
- [5] Howard, J.; Ruder, S. Universal Language Model Fine-tuning for Text Classification. 2018.
- [6] Castro Nascimento, C. M.; Pimentel, A. S. Do Large Language Models Understand Chemistry? A Conversation with ChatGPT. *Journal of Chemical Information and Modeling* **2023**, *63*, 1649–1655, DOI: 10.1021/acs.jcim.3c00285, PMID: 36926868.
- [7] Jablonka, K. M.; Schwaller, P.; Ortega-Guerrero, A.; Smit, B. Leveraging large language models for predictive chemistry. *Nature Machine Intelligence* **2024**, DOI: 10.1038/s42256-023-00788-1.
- [8] Mirza, A. et al. Are large language models superhuman chemists? 2024.
- [9] Jablonka, K. M. et al. 14 examples of how LLMs can transform materials science and chemistry: a reflection on a large language model hackathon. *Digital Discovery* **2023**, *2*, 1233–1250, DOI: 10.1039/D3DD00113J.

- [10] Dettmers, T.; Lewis, M.; Belkada, Y.; Zettlemoyer, L. GPT3.int8(): 8-bit Matrix Multiplication for Transformers at Scale. *Advances in Neural Information Processing Systems*. 2022.
- [11] Dettmers, T.; Lewis, M.; Shleifer, S.; Zettlemoyer, L. 8-bit Optimizers via Block-wise Quantization. *The Tenth International Conference on Learning Representations, ICLR*. 2022.
- [12] Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Chen, W. LoRA: Low-Rank Adaptation of Large Language Models. *International Conference On Learning Representations*. 2021.
- [13] Wang, B.; Komatsuzaki, A. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>, 2021.
- [14] Heid, E.; Greenman, K. P.; Chung, Y.; Li, S.-C.; Graff, D. E.; Vermeire, F. H.; Wu, H.; Green, W. H.; McGill, C. J. Chemprop: A Machine Learning Package for Chemical Property Prediction. *Journal of Chemical Information and Modeling* **2024**, *64*, 9–17, DOI: 10.1021/acs.jcim.3c01250, PMID: 38147829.
- [15] Shi, J.; Quevillon, M. J.; Valença, P. H. A.; Whitmer, J. K. Predicting Adhesive Free Energies of Polymer–Surface Interactions with Machine Learning. *ACS Appl. Mater. Interfaces* **2022**, *32*, 37161–37169.
- [16] Schneider, L.; Schwarting, M.; Mysona, J.; Liang, H.; Han, M.; Rauscher, P. M.; Ting, J. M.; Venkatram, S.; Ross, R. B.; Schmidt, K. J.; Blaiszik, B.; Foster, I.; de Pablo, J. J. *In silico* active learning for small molecule properties. *Mol. Syst. Des. Eng.* **2022**, *7*, 1611–1621, DOI: 10.1039/d2me00137c.
- [17] Dieringa, H.; Kainer, K. U. In *Springer Handbook of Materials Data*; Warlimont, H., Martienssen, W., Eds.; Springer International Publishing: Cham, 2018; pp 151–159.
- [18] Wolff, M.; Ebel, T.; Dahms, M. Sintering of Magnesium. *Advanced Engineering Materials* **2010**, *12*, 829–836, DOI: <https://doi.org/10.1002/adem.201000038>.
- [19] Wolff, M.; Schaper, J. G.; Suckert, M. R.; Dahms, M.; Feyerabend, F.; Ebel, T.; Willumeit-Römer, R.; Klassen, T. Metal Injection Molding (MIM) of Magnesium and Its Alloys. *Metals* **2016**, *6*.

- [20] Wolff, M.; Schaper, J. G.; Dahms, M.; Ebel, T.; Kainer, K. U.; Klassen, T. Magnesium powder injection moulding for biomedical application. *Powder Metallurgy* **2014**, *57*, 331–340, DOI: 10.1179/1743290114Y.0000000111.
- [21] Pekguleryuz, M. O., Kainer, K. U., Arslan Kaya, A., Eds. *Fundamentals of Magnesium Alloy Metallurgy*; Woodhead Publishing Series in Metals and Surface Engineering; Woodhead Publishing, 2013; p iv, DOI: <https://doi.org/10.1016/B978-0-85709-088-1.50012-4>.
- [22] Humphreys, F., Hatherly, M., Eds. *Recrystallization and Related Annealing Phenomena (Second Edition)*, second edition ed.; Elsevier: Oxford, 2004; pp 527–540, DOI: <https://doi.org/10.1016/B978-008044164-1/50021-9>.
- [23] Boeynaems, S.; Alberti, S.; Fawzi, N. L.; Mittag, T.; Polymenidou, M.; Rousseau, F.; Schymkowitz, J.; Shorter, J.; Wolozin, B.; Van Den Bosch, L.; Tompa, P.; Fuxreiter, M. Protein Phase Separation: A New Phase in Cell Biology. *Trends in Cell Biology* **2018**, *28*, 420–435, DOI: <https://doi.org/10.1016/j.tcb.2018.02.004>.
- [24] Saar, K. L.; Morgunov, A. S.; Qi, R.; Arter, W. E.; Krainer, G.; Lee, A. A.; Knowles, T. P. Learning the molecular grammar of protein condensates from sequence determinants and embeddings. *Proceedings of the National Academy of Sciences* **2021**, *118*, e2019053118.
- [25] Alampara, N.; Miret, S.; Jablonka, K. M. MatText: Do Language Models Need More than Text & Scale for Materials Modeling? 2024; <https://arxiv.org/abs/2406.17295>.
- [26] Anker, A. S.; T. S. Kjær, E.; Dam, E. B.; J. L. Billinge, S.; Jensen, K. M. O.; Selvan, R. Characterising the Atomic Structure of Mono-Metallic Nanoparticles from X-Ray Scattering Data Using Conditional Generative Models. **2020**, DOI: 10.26434/chemrxiv.12662222.v1.
- [27] Kjær, E. T. S.; Anker, A. S.; Weng, M. N.; Billinge, S. J. L.; Selvan, R.; Jensen, K. M. O. DeepStruc: towards structure solution from pair distribution function data using deep generative models. *Digital Discovery* **2023**, *2*, 69–80, DOI: 10.1039/d2dd00086e.
- [28] Tu, Z.; Stuyver, T.; Coley, C. W. Predictive chemistry: machine learning for reaction

- deployment, reaction development, and reaction discovery. *Chem. Sci.* **2023**, *14*, 226–244, DOI: 10.1039/D2SC05089G.
- [29] Cordova, M.; Wodrich, M. D.; Meyer, B.; Sawatlon, B.; Corminboeuf, C. Data-Driven Advancement of Homogeneous Nickel Catalyst Activity for Aryl Ether Cleavage. *ACS Catal.* **2020**, *10*, 7021–7031, DOI: 10.1021/acscatal.0c00774.
- [30] Dobbelaere, M. R.; Plehiers, P. P.; Van de Vijver, R.; Stevens, C. V.; Van Geem, K. M. Machine Learning in Chemical Engineering: Strengths, Weaknesses, Opportunities, and Threats. *Engineering* **2021**, *7*, 1201–1211, DOI: <https://doi.org/10.1016/j.eng.2021.03.019>.
- [31] Daglar, H.; Keskin, S. Combining Machine Learning and Molecular Simulations to Unlock Gas Separation Potentials of MOF Membranes and MOF/Polymer MMMs. *ACS Applied Materials & Interfaces* **2022**, *14*, 32134–32148.
- [32] Bucior, B.; Rosen, A.; Haranczyk, M.; Yao, Z.; Ziebel, M.; Farha, O.; Hupp, J.; Siepmann, J.; Aspuru-Guzik, A.; Snurr, R. Identification Schemes for Metal–Organic Frameworks To Enable Rapid Search and Cheminformatics Analysis. *Crystal Growth & Design* **2019**, *19*, DOI: 10.1021/acs.cgd.9b01050.
- [33] Klopčič, N.; Grimmer, I.; Winkler, F.; Sartory, M.; Trattner, A. A review on metal hydride materials for hydrogen storage. *Journal of Energy Storage* **2023**, *72*, 108456, DOI: <https://doi.org/10.1016/j.est.2023.108456>.
- [34] Witman, M.; Ling, S.; Grant, D. M.; Walker, G. S.; Agarwal, S.; Stavila, V.; Allendorf, M. D. Extracting an Empirical Intermetallic Hydride Design Principle from Limited Data via Interpretable Machine Learning. *The Journal of Physical Chemistry Letters* **2019**, *11*, 40–47, DOI: <https://doi.org/10.1021/acs.jpcllett.9b02971>.
- [35] Mashhadimoslem, H.; Vafaeinia, M.; Safarzadeh, M.; Ghaemi, A.; Fathalian, F.; Maleki, A. Development of Predictive Models for Activated Carbon Synthesis from Different Biomass for CO<sub>2</sub> Adsorption Using Artificial Neural Networks. *Industrial and Engineering Chemistry Research* **2021**, *60*, 13950–13966, DOI: 10.1021/acs.iecr.1c02754.
- [36] Darre, N.; Toor, G. Desalination of Water: a Review. *Current Pollution Reports* **2018**, *4*, 1–8, DOI: 10.1007/s40726-018-0085-9.

- [37] Gil, M. V.; Jablonka, K. M.; Garcia, S.; Pevida, C.; Smit, B. Biomass to energy: a machine learning model for optimum gasification pathways. *Digital Discovery* **2023**, *2*, 929–940, DOI: 10.1039/d3dd00079f.
- [38] Moosavi, S. M.; Chidambaram, A.; Talirz, L.; Haranczyk, M.; Stylianou, K. C.; Smit, B. Capturing chemical intuition in synthesis of metal-organic frameworks. *Nat. Commun.* **2019**, *10*, 539, DOI: 10.1038/s41467-019-08483-9.
- [39] Jablonka, K. M.; Patiny, L.; Smit, B. Making the collective knowledge of chemistry open and machine actionable. *Nat Chem* **2022**, *14*, 365–376, DOI: 10.1038/s41557-022-00910-7.