

scientific data



OPEN

DATA DESCRIPTOR

Dataset of solution-based inorganic materials synthesis procedures extracted from the scientific literature

Zheren Wang^{1,2}, Olga Kononova^{1,2}, Kevin Cruse^{1,2}, Tanjin He^{1,2}, Haoyan Huo^{1,2}, Yuxing Fei^{1,2}, Yan Zeng², Yingzhi Sun^{1,2}, Zijian Cai^{1,2}, Wenhao Sun³ & Gerbrand Ceder^{1,2}✉

The development of a materials synthesis route is usually based on heuristics and experience. A possible new approach would be to apply data-driven approaches to learn the patterns of synthesis from past experience and use them to predict the syntheses of novel materials. However, this route is impeded by the lack of a large-scale database of synthesis formulations. In this work, we applied advanced machine learning and natural language processing techniques to construct a dataset of 35,675 solution-based synthesis procedures extracted from the scientific literature. Each procedure contains essential synthesis information including the precursors and target materials, their quantities, and the synthesis actions and corresponding attributes. Every procedure is also augmented with the reaction formula. Through this work, we are making freely available the first large dataset of solution-based inorganic materials synthesis procedures.

Background & Summary

Big-data-driven approaches have helped to establish a new paradigm of scientific research^{1–3}. In materials science specifically, the Materials Genome Initiative (MGI) effort has significantly facilitated and accelerated materials discovery and design by deploying large-scale *ab initio* computation and building computed databases of structure–property relationships^{4–6}. Unlike computational data, the experimentally determined properties and structures of inorganic materials are mainly available in manually curated databases^{7–11}. Such well-curated experimental databases have led to early machine learning models that address difficult problems in materials research, such as structure prediction³. The ability to efficiently design and predict the structure of novel advanced materials with the assistance of computed and experimental databases has shifted the materials innovation challenge toward understanding and determining the synthesis routes for novel materials¹². In principle, this challenge could also be mastered using data-driven approaches, since only limited predictive theories are available for materials synthesis^{13–23}. Indeed, in organic chemistry, AI-guided synthesis planning^{24,25} has already been successfully implemented in certain cases, such as in predicting retrosynthesis²⁶ and in complex natural product synthesis design²⁷. Although datasets for the synthesis of organic materials are widely available^{28,29}, there is not yet a large-scale database of inorganic synthesis routes, which is needed to train advanced deep-learning models to enable a breakthrough in AI-assisted design and optimization of inorganic materials synthesis.

Scientific publications represent the largest repository of knowledge about materials synthesis and can be used as a reliable source of data. However, human-written descriptions of syntheses require additional levels of interpretation for conversion into a codified, machine-operable format. Additionally, manual extraction of synthesis information is laborious, even for a very limited number of papers^{30–33}. Given these obstacles, an automated information extraction pipeline can accelerate data collection and assist in building structured synthesis procedures from scientific text. Natural language processing (NLP) approaches have been widely developed in the past decade, and various advanced tools for high-quality information extraction from unstructured text are available to researchers.

¹Department of Materials Science and Engineering, University of California, Berkeley, CA, 94720, USA. ²Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA. ³Department of Materials Science and Engineering, University of Michigan, Ann Arbor, MI, USA. ✉e-mail: gceder@berkeley.edu

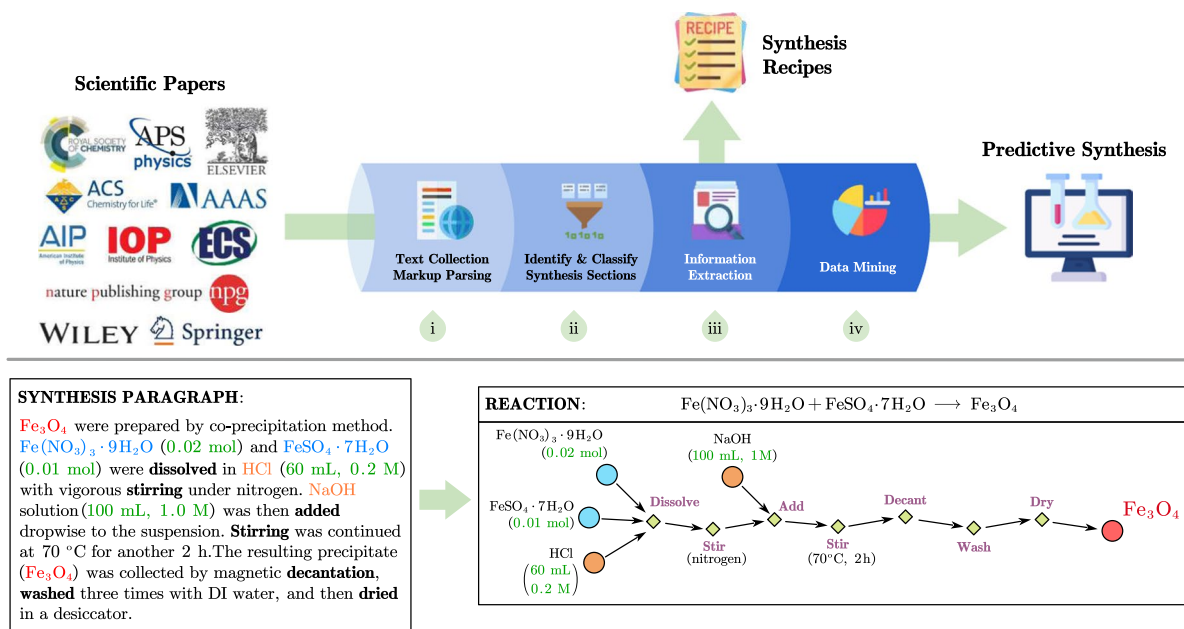


Fig. 1 Extraction pipeline and example. *Top panel:* Schematic representation of the standard text mining pipeline: (i) scrape papers in markup format from the major publishers; (ii) identify and classify synthesis sections; (iii) extract key information including materials, amounts, sequenced operations, and conditions; (iv) store synthesis procedures into the database for future data mining. *Bottom panel:* Example of a codified procedure extracted from a synthesis paragraph.

In materials science, NLP has been used to extract and analyze materials properties^{34–36}, applications^{37,38}, and synthesis conditions for some limited cases³⁹. Various NLP tools, including ChemDataExtractor⁴⁰, OSCAR⁴¹, ChemicalTagger⁴², and others^{30,43,44}, have been developed to extract information from chemical text. Recently, advanced models based on deep convolutional and recurrent neural networks^{45–48} have been proposed to improve the accuracy of chemical data extraction.

Text-mining approaches in materials science have also been used to construct automated pipelines for collecting information about materials synthesis from publications and to build large-scale publicly available datasets from such collected data, including datasets of synthesis formulations for metal oxides^{39,49,50}, germanium-containing zeolites⁵¹, and perovskites⁵². In recent work, our group has developed a text-mining pipeline to construct the first large-scale dataset of solid-state ceramics synthesis “recipes”, which includes not only the starting materials and final products but also the synthesis actions, their attributes, and balanced chemical-reaction equations⁵³.

In the current work, we built a more advanced extraction pipeline (Fig. 1) which uses various advanced machine learning and natural language processing techniques to extract precise data for solution-based inorganic materials synthesis procedures from the scientific literature. Solution procedures are considerably more complex than solid state synthesis and require the precise extraction of not only the chemicals involved but also their respective amounts (since they determine concentration in solution). In addition, more complex organic and mixed organic-inorganic compounds are used to solubilize ions or to control solution conditions. By applying the extraction pipeline, we codified 35,675 solution-based inorganic materials synthesis procedures from over 4 million papers. Extracted information includes target material and precursors, their quantities, and the synthesis operations and their attributes. Information about the targets and precursors is then used to build a reaction formula for every synthesis procedure. This dataset is the first large-scale dataset of solution-based synthesis procedures, and provides a foundation to test and verify existing empirical synthesis rules, improve prediction accuracy, and even data-mine new rules to guide synthesis. Also, this codified dataset should pave the way to design optimized synthesis procedures in automated experimentation.

Methods

Content acquisition. The journal articles used in this work were downloaded with publisher consent from Wiley, Elsevier, the Royal Society of Chemistry, the Electrochemical Society, the American Chemical Society, the American Physical Society, the American Institute of Physics, and Nature Publishing Group. A customized web-scraaper, Borges (see Codes Availability section below), was used to automatically download a broad selection of materials-relevant papers published after the year 2000 from publishers’ websites in HTML/XML format. We selected 2000 as the cutoff year as parsing of materials science papers stored as image PDFs (as for most papers published before 2000) introduces a significant number of errors due to the limitations of currently available optical character recognition models on chemistry-containing text^{54,55}.

To convert the articles from HTML/XML into raw-text files, we developed the LimeSoup toolkit (see Codes Availability section below), which takes into account the specific format standards of various publishers and

journals. The full-text and metadata of the articles such as the journal name, article title, abstract, author names, etc., are stored in a MongoDB (www.mongodb.com) database collection. To date, we have accumulated 4.06 million articles, which are used for further processing down the pipeline (Fig. 1).

Paragraph classification. Paragraphs containing information about solution synthesis (referred to as “synthesis paragraphs” throughout this paper) were identified using a Bidirectional Encoder Representations from Transformers (BERT) model⁵⁶. The model was pre-trained on full-text paragraphs of 2 million papers randomly drawn from our database in a self-supervised way, i.e., by predicting masked words based on their surrounding context. After training the BERT model, we fine-tuned the paragraph classifier using 7,292 paragraphs labeled as either “solid-state synthesis”, “sol-gel precursor synthesis”, “hydrothermal synthesis”, “precipitation synthesis”, or “none of the above”. The resulting F1 score of the paragraph classification is 99.5%, an improvement over the F1 score of 94.6% in our previous work⁵⁷, when evaluated using the same labeled training dataset.

Synthesis procedure extraction. A solution-based synthesis procedure includes the precursors and target materials, their quantities, and the synthesis actions and their attributes, properly sequenced. This is the minimum essential information required to complete a synthesis route. A schematic representation of the procedure is shown in the bottom panel in Fig. 1. In the sections below, we provide a brief overview of the methods used for each step of the procedure extraction.

Materials entity recognition (MER). Materials entities in synthesis paragraphs are identified and classified as *target*, *precursor*, or *other* via a two-step sequence-to-sequence model as introduced in our previous work⁴⁶. In the current work, we replaced the original Word2Vec embedding model used previously⁵⁸ with a BERT model trained on papers from the materials science domain (see Section “Paragraph classification” above). First, each word token was transformed into a digitized BERT embedding vector. A bi-directional long-short-term memory neural network with a conditional random-field top layer (BiLSTM-CRF) was used to determine whether the token was a materials entity or a regular word, and each materials entity was replaced with the keyword <MAT> before being classified as either a *target*, *precursor*, or *other* material using a second BERT-based BiLSTM-CRF network. In addition to the 834 annotated solid-state synthesis paragraphs from 750 papers used in our previous work⁴⁶, we manually annotated 447 solution-based synthesis paragraphs from 405 papers by labeling each word token as *material*, *target*, *precursor*, or *outside*. The annotated dataset was split into training, validation, and test sets with a paper-wise ratio of 700:150:305 to train the aforementioned two neural networks.

Extraction of synthesis actions and attributes. We implemented an algorithm which combines a neural network and sentence dependency tree analysis to identify synthesis actions in the text. First, the Word2Vec model from the Gensim library⁵⁹ was re-trained on ~400,000 synthesis paragraphs of four synthesis types (see Section “Paragraph classification” above). These word embeddings were used as the input for a recurrent neural network that takes a sentence word-by-word and assigns labels to the verb tokens: *not-operation*, *mixing*, *heating*, *cooling*, *shaping*, *drying*, or *purifying*. For each obtained synthesis action, we parsed a dependency sub-tree using the SpaCy library⁶⁰ to obtain information about the corresponding temperature, time, and environment. To extract the corresponding values of these attributes, we used a rule-based regular expression approach⁶¹.

Extraction of material quantities. To extract the numerical values of material quantities and assign them to the corresponding materials obtained using the MER model (see Section “MER” above), we applied a rule-based approach to search along the syntax tree⁶¹. The NLTK library⁶² was used to build the syntax trees for each sentence in a paragraph. The words in given sentences are leaf nodes of syntax trees. We then applied an algorithm to cut the syntax tree of each sentence into the largest sub-trees for every material, with each sub-tree having only one material entity: 1. we first identified the materials on leaf nodes; 2. starting from each material, we identified the largest sub-trees, i.e., we traversed the syntax tree upwards until there was more than one material leaf node descending from the same node; 3. the largest sub-tree for a given material was defined as the sub-tree formed by the node and its descendants identified in step 2. Next, we searched for the quantities in each sub-tree given as molarity, concentration, or volume. Finally, we assigned the quantities found to the unique material entity in the sub-tree.

Building reaction formulas. For every synthesis procedure described in a paragraph, we built a chemical formula. Every material entity was converted from a text-string representation into a chemical-data structure using an in-house material parser toolkit (see Codes Availability section below). The data structure included information about the material formula, composition, and ions. We then paired the target with precursors containing at least one element in the target except for hydrogen and oxygen and defined those precursors as “precursor candidates”. Next, we computed the oxidation state change of elements from each “precursor candidate” to the target and determined whether the precursor was oxidized or reduced. If precursors were reduced or oxidized, we also included the corresponding redox agents in the reaction formula. The agents can either be another “precursor candidate” or a commonly used oxidizing or reducing agent from the remaining material entities marked as *other* or *precursor* by the MER algorithm (see Section “MER” above).

Dataset generation. The dataset generation followed the protocol displayed in Fig. 1. We downloaded a total of 4,061,814 papers using web scraping and identified the experimental sections by keyword matching in section headings, with keywords including “experiment”, “synthesis”, “preparation”, and their morphological derivations. ChemDataExtractor⁴⁰ was used to split the plain-text paragraphs into sentences and words. After classification (see Section “Paragraph classification” above), 364,076 paragraphs describing solid-state, hydrothermal, sol-gel, and precipitation syntheses were obtained. Among them, 189,553 paragraphs described hydrothermal or

precipitation syntheses, which we categorize as solution-based synthesis methods. These paragraphs were further processed to extract the precursors, targets, quantities, and operations with corresponding conditions and to build the reaction formula (Fig. 1).

Data Records

The solution-based synthesis dataset is provided as a single JSON file, available at <https://doi.org/10.6084/m9.figshare.16583387.v4>⁶³. There are 20,037 hydrothermal synthesis reactions and 15,638 precipitation synthesis reactions. Each record corresponds to a synthesis procedure extracted from a paragraph and is represented as an individual JSON object. If a paragraph reported the synthesis of several materials, the corresponding reactions were split into separate data records. In addition to the chemical formula, the metadata for each reaction returns the data structure used in our previous work⁵³, which includes: DOI of the paper, a snippet of the corresponding synthesis paragraph (50 first and 50 last characters to facilitate its lookup), chemical information about the target and precursor materials used in the reaction, and operations with their corresponding attributes. We also included the materials with their corresponding quantities in the metadata. The details of the data format are given in Table 1

The chemical formula for the reaction is stored as a string (`reaction_string`) as well as in a dictionary containing lists of precursors (`left_side`) and target materials (`right_side`) in the reaction.

The metadata for target materials and precursors used to construct the chemical formula are represented by the following data structure:

- `material_string`: string of material as given in the original paragraph before being parsed into a chemical composition.
- `material_formula`: chemical formula associated with the material (given originally or constructed empirically by parser).
- `composition`: chemical composition of the material derived from its formula. Aside from single-compound materials, we found that a large portion of the materials (predominantly target materials) are composites, mixtures, solid solutions, or alloys written as a sequence of compound-fraction pairs. Therefore, a chemical-composition entity is represented by a list of dictionary entries, where each item is associated with a compound found in the materials formula. The fraction of each compound in the material is given in `amount`, and its chemical composition (i.e., the elements and stoichiometry) is given in `elements`. If a material is one compound, the list has only one item and `amount` = 1.0. If a material is a hydrate, water is added to the `composition` list with its amount corresponding to the amount of water molecules (if specified).
- `additives`: list of additive elements (i.e., elements used for doping, stabilization, or substitution) resolved from the material string.
- `elements_vars`: lists all variable elements and their corresponding values found in the materials.
- `amounts_vars`: lists all variable element ratios and their corresponding values found in the material formula. The values of each variable are given as a structure with `values` listing the values of each specific variable and `max_value/min_value` values if a range is given in the paragraph.
- `oxygen_deficiency`: yes/no attribute that reflects whether a material was synthesized with unspecified oxygen stoichiometry.
- `mp_id`: ID of the lowest-energy polymorph entry in the Materials Project database (<https://materialsproject.org/materialsproject.org>) if the material is found there.

To facilitate querying of the dataset, the `targets_string` field contains the target material formulas, and the `solvents` field contains all solvent(s) from matching material entities marked as *other* by the MER model with a table of common solvents adopted from Common Solvents Used in Organic Chemistry (<https://organicchemistrydata.org/solvents/organicchemistrydata.org/solvents>).

Technical Validation

Extraction completeness and accuracy. To ensure high accuracy of the dataset, we included only those data that produced complete reaction formulas at the final step of the pipeline. This strategy reduced potential errors in the dataset that may have been caused by composition-parsing failure, incomplete extraction, or incomplete information provided by the text. We applied the extraction pipeline to 189,553 solution-based synthesis paragraphs, 28,749 of which generated a reaction formula, giving an extraction yield of ~15%. To evaluate the source of the loss, we randomly selected and manually checked 100 solution-based synthesis paragraphs that did not produce any reactions. Among those 100 paragraphs, 36 were written with an incomplete list of precursors or targets in the text, such that human experts would not be able to reconstruct the reaction based solely on the information provided in the paragraph. For the remaining 64 paragraphs, the loss was due to: 1. the use of organic precursors with complex groups or complicated notation (e.g., acronyms) that could not be parsed into a chemical composition by our parser or 2. MER misidentification resulting in an incomplete or incorrect list of precursors and (or) target entities such that the reactions could not be built.

To evaluate the quality of the dataset, we had a human expert test 100 randomly pulled entries. The human expert manually extracted the information presented in the procedure, and the results were compared with those extracted by the pipeline. Table 2 presents the accuracy statistics, which include the precision, recall, and F1 scores calculated from the tested entries. For the fields that included reaction, targets, precursors, operations, operation temperatures, time, and atmosphere, the F1 scores were over 90%. The relatively low recall, and hence F1 score, for the extraction of materials quantities can be mainly explained by the MER algorithm missing the corresponding material entity and, thus, the quantities not being assigned. The accuracy of the obtained dataset is comparable to that in our previous work⁵³ and of other text-mined datasets⁴⁹.

Data description	Data Key Label	Data Type
DOI of the original paper	doi	string
Snippet of the raw text	paragraph_string	string
Chemical formula	reaction	Object (dict):
		- left_side: list of strings
		- right_side: list of strings
Chemical formula in string format	reaction_string	string
Target material data	target	Object (dict):
		- material_string: string,
		- material_formula: string,
		- composition: list of Objects ¹ ,
		- additives: list of strings
		- elements_vars: {var: list of strings}
		- amounts_vars: {var: list of Objects ² }
		- oxygen_deficiency: boolean
- mp_id: string		
List of target formulas obtained after variables substitution	targets_string	list of strings
Precursor materials data	precursors	list of Objects (See target)
List of solvent formulas	solvents_string	list of strings
Sequence of synthesis steps and corresponding conditions	operations	list of Objects (dict):
		- token: string,
		- type: string
		- conditions: Object
		- temperature: list of Objects ³
		- time: list of Objects ³ ,
		- atmosphere: list of strings
		- mixing_device: list of strings
- mixing_media: list of strings		
Materials with corresponding quantities	quantities	list of Objects (dict):
		- material: string,
		- quantity: list of Objects ⁴
Synthesis type	type	string

Table 1. Format of each data record: description, key label, data type. ¹{formula: string, elements: {elements: amount of element}, amount: string}. ²{max_value: float, min_value: float, values: list of floats}. ³{max_value: float, min_value: float, values: list of floats, units: string}. ⁴{number: float, unit: string}.

Exploratory data analysis. To test the diversity of the dataset and its coverage of the materials space, we analyzed unique materials (targets and precursors) and reactions. The dataset contains 11,603 unique reactions that include 2,870 unique precursors and 5,416 unique targets. The ten most frequent targets in the dataset and their corresponding precursors are listed in Table 3. The target list captures materials that have drawn substantial attention in the past two decades: catalysts (ZnO, Fe₂O₃, TiO₂, Fe₃O₄, SnO₂, ZrO₂, CuO), adsorbents (SiO₂), various materials for sensors (ZnO, Fe₂O₃, WO₃), quantum dots (CdS), and semiconductors (ZnO, TiO₂, SnO₂, CdS). Unsurprisingly, these most frequent target materials usually appear in multiple applications, as they possess desirable physical and chemical properties in many scientific and engineering fields.

We use the periodic table representation (Fig. 2) to visualize the chemical space covered by the dataset. For each element, the fraction of synthesis procedures containing this element in the target formula is shown with the yellow-to-navy blue gradient framed at the top of each element box. The most data-rich elements are transition metals in the third period, such as Zn, Fe, Ti, Ni, and Co, in accordance with the compounds listed in Table 3. The next-most prevalent targets are materials with Bi, Sn, Al, W, Mo, Cu, Zr, or Li. The least common elements are rare elements such as Ru, Rh, Hf, Ta, Re, and Ir. The elements Fr, Ra, Tc, and Pm are not present as target materials in the dataset, likely due to their radioactivity. Additionally, we calculated the frequency of co-occurrence of chemical elements and common ions in precursor materials to understand how different ions are brought into solution. In Fig. 2, the frequencies for each ion are displayed as colored bars. The length of the bar is the fraction of one specific ion paired with the element normalized over all precursors for this element.

The commonly used precursors are mainly those that are widely available from companies such as Sigma-Aldrich and Fisher Scientific. For example, Li₂CO₃ or LiOH for Li and sulfate or chloride for Fe. Inorganic salts, such as nitrates, sulfates, and chlorides, are often used because of their high solubility⁶⁴. Halides are not used for Lanthanide metals because Lanthanide halides are highly hygroscopic, and thus their molar weight is not well defined. We observed that precursors with neighboring elements in the periodic table tend to have

Data attribute	Precision	Recall	F1 score
Balanced reactions	0.94	/	/
- targets	0.97	/	/
- precursors	0.98	0.99	0.98
Operations	0.96	0.85	0.90
Conditions			
- temperature	0.97	0.92	0.94
- time	0.98	0.89	0.93
- atmosphere	0.97	0.92	0.94
Quantities	0.90	0.85	0.87

Table 2. Performance of data extraction for dataset entries.

Targets	Common Precursors
ZnO	Zn(NO ₃) ₂ , Zn(Ac) ₂ , ZnCl ₂
TiO ₂	Ti(OCH(CH ₃)) ₄ , Ti(OC ₄ H ₉) ₄ , TiCl ₄
Fe ₃ O ₄	FeCl ₃ , FeCl ₂
Fe ₂ O ₃	FeCl ₃ , Fe(NO ₃) ₃
SnO ₂	SnCl ₄
ZrO ₂	ZrOCl ₂ , ZrO(NO ₃) ₂
CuO	Cu(NO ₃) ₂ , Cu(Ac) ₂ , CuCl ₂ , CuSO ₄
SiO ₂	Si(OC ₂ H ₅) ₄
WO ₃	Na ₂ WO ₄ , WCl ₆ , H ₂ WO ₄
CdS	Na ₂ S, CH ₄ N ₂ S, CdCl ₂ , Cd(NO ₃) ₂

Table 3. Ten most common targets in the dataset with their corresponding precursors.

similar ions paired. For instance, nitrates, sulfates, and chlorides are commonly used anions for 3rd-period transition metals, whereas the precursors for lanthanides are mostly oxides and nitrates.

We used information about the extracted materials and sequences of synthesis actions to classify the solution-based synthesis procedures into four categories of synthesis protocols (table in Fig. 3) according to the following definitions:

- *solution-mixing with heat treatment step* has a final heat treatment step after the precipitate is obtained from the solution;
- *aqueous solution synthesis* has no final heat treatment step after precipitating the compound from the solution and the solvent is water;
- *non-aqueous solution synthesis* has no final heat treatment step after precipitating the compound from solution and the solvent(s) is (are) organic;
- *aqueous-non-aqueous mixed solution synthesis* has no final heat treatment step after precipitating the compound from solution and the solvents are a mixture of water and organic solvent(s).

The resulting distributions of synthesis protocols over the aforementioned categories are shown in the two pie charts in the top-right corner of Fig. 3. Note that as solution-based synthesis includes both hydrothermal and precipitation synthesis according to our definition (see Section “Paragraph classification”), we analyzed these synthesis types separately. As observed in the pie charts, only 20% of the procedures in the hydrothermal synthesis subset have a heat treatment step after solution mixing. Among those that do not have heat treatment step, 63% use only water as a solvent, 8% use only organic solvents, and 9% use both water and organic solvents. In contrast, the fractions in the precipitation synthesis subset are 43%, 46%, 5%, and 6%, respectively.

A heat treatment step after solution mixing can be used to dehydrate the targets, decompose the intermediates to produce the final products, change the oxidation state, change the morphology, or improve crystallization^{65–67}. To explore this in more detail, we split the targets according to their anion type (oxide, sulfide, etc.) and different oxidation states of several data-rich transition-metal elements. We then computed the distribution of synthesis categories for each of the split subsets. Figure 3 presents the results for the most prevalent subsets of oxides, sulfides, and elements Fe²⁺, Fe³⁺, Co³⁺, Ni²⁺, Cu²⁺, and Zn²⁺. The fraction of procedures with a heat treatment step in precipitation synthesis is larger than that in hydrothermal synthesis. This observation holds for all targets, all oxides, all sulfides, and individual oxides and sulfides with queried oxidation states of transition metals. This finding can also be interpreted as hydrothermal synthesis often being used to obtain final products in a “one-shot” process, without subsequent heat treatment after solution mixing, likely because many compounds can be crystallized as anhydrous powders with controlled size and morphology directly from hydrothermal synthesis. In a standard hydrothermal synthesis procedure, the reaction is performed in an autoclave with autogenic pressure so that it can operate in a wider temperature window, including temperatures above the atmospheric boiling point of the solvent. In contrast, precipitation synthesis is performed under normal pressure.

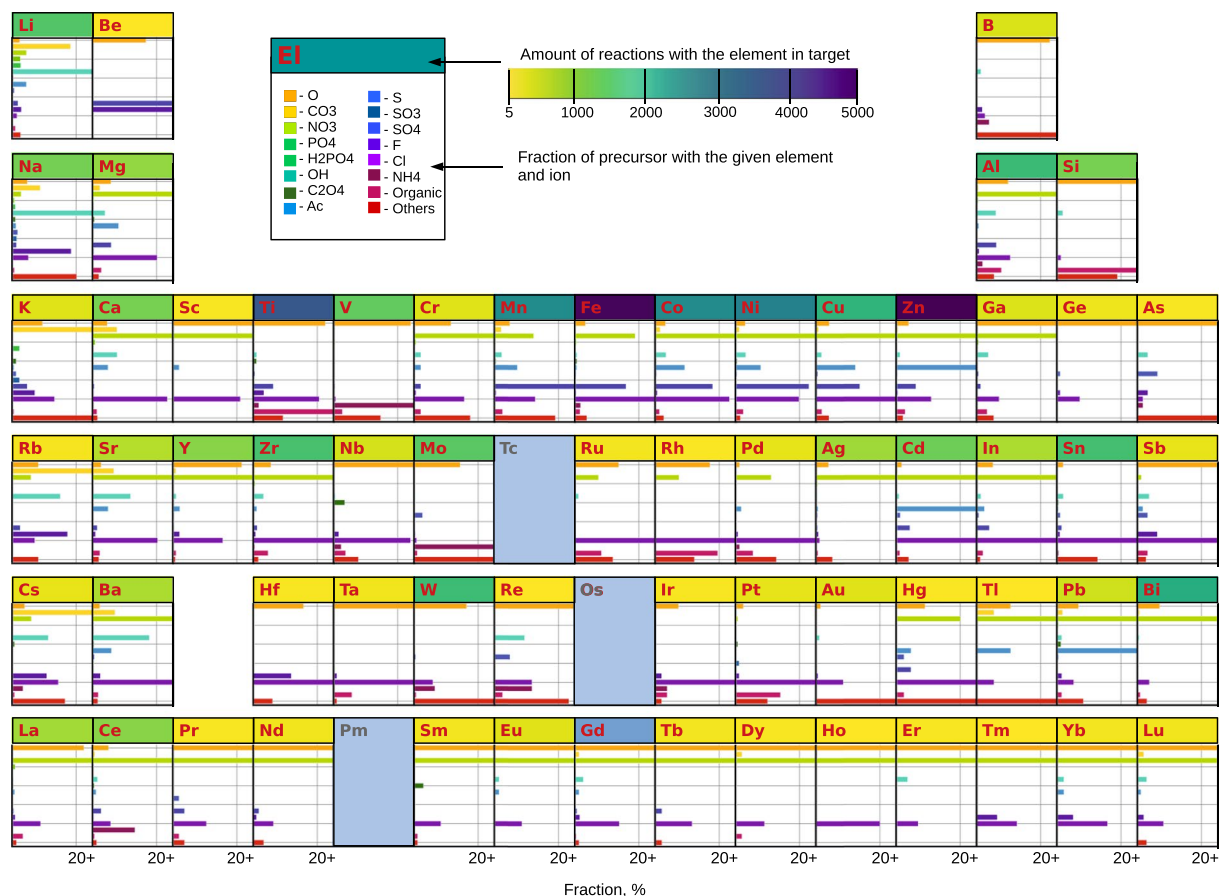


Fig. 2 The chemical space covered by the dataset. For each element, the box containing the element name is colored in a yellow-to-navy blue gradient representing the total amount of reactions that produce a target compound containing the element. The bar graph below each element shows the list of ions paired with the element in precursor compounds. The fractions of the precursors (i.e. element + ion) used are shown by the length of the bars. Boxes with no bar graph represent elements occurring in five and fewer targets. “Ac” stands for acetate radical CH_3COO^- in the compound formula.

The higher temperature possible in hydrothermal synthesis is associated with enhanced kinetics in chemical transport, nucleation, and crystal growth and thus with a more effective dissolution–recrystallization process, which can help remove defects and improve crystallinity. Furthermore, the physico-chemical properties, such as the viscosity and dielectric constant of water or other solvents, change pronouncedly under conditions of hydrothermal synthesis, affecting the solubility and mobility of species in the solution and eventually facilitating crystallization⁶⁸. Therefore, hydrothermal synthesis does not need a post-synthesis heat treatment as often as precipitation synthesis.

Solution-based synthesis is an important area of materials synthesis⁵⁷ and this dataset can help with advancing the science and model building for solution synthesis. Nevertheless, challenges remain in the mining of scientific literature and construction of robust and accurate large-scale datasets. First, the organic precursors with complex radicals commonly used in solution-based synthesis pose a challenge for parsing and extracting chemical information. Constructing reaction formulas becomes problematic when the precursor information is lost. Therefore, these entries are mostly dropped out later in the pipeline. To address this issue, a universal parser that can parse chemical tokens needs to be developed.

Second, our data was extracted from the experimental section in the main body of each paper and does not include any information about the actual synthesis results, e.g., whether the material was synthesized using the reported procedure or which structure was obtained. This problem could be overcome by introducing a model that can parse characterization data (e.g., X-ray diffraction patterns or electron microscopy images) and relate them to the corresponding synthesis conditions, something which, to the best of our knowledge, has not yet been performed. Even though the actual results of a synthesis can be extracted from a paper, there remains the challenge of data interpretation and usage, as the authors usually report only successful and “cherry-picked” experimental results. This introduces significant anthropogenic bias toward “positive” data with little “negative” content in the dataset, thus limiting the tasks for future machine-learning applications^{69,70}. A promising approach to solve this issue is to incorporate results obtained by autonomous robotic synthesis platforms that can provide a vast amount of “negative” data in a reasonable time frame^{71,72}.

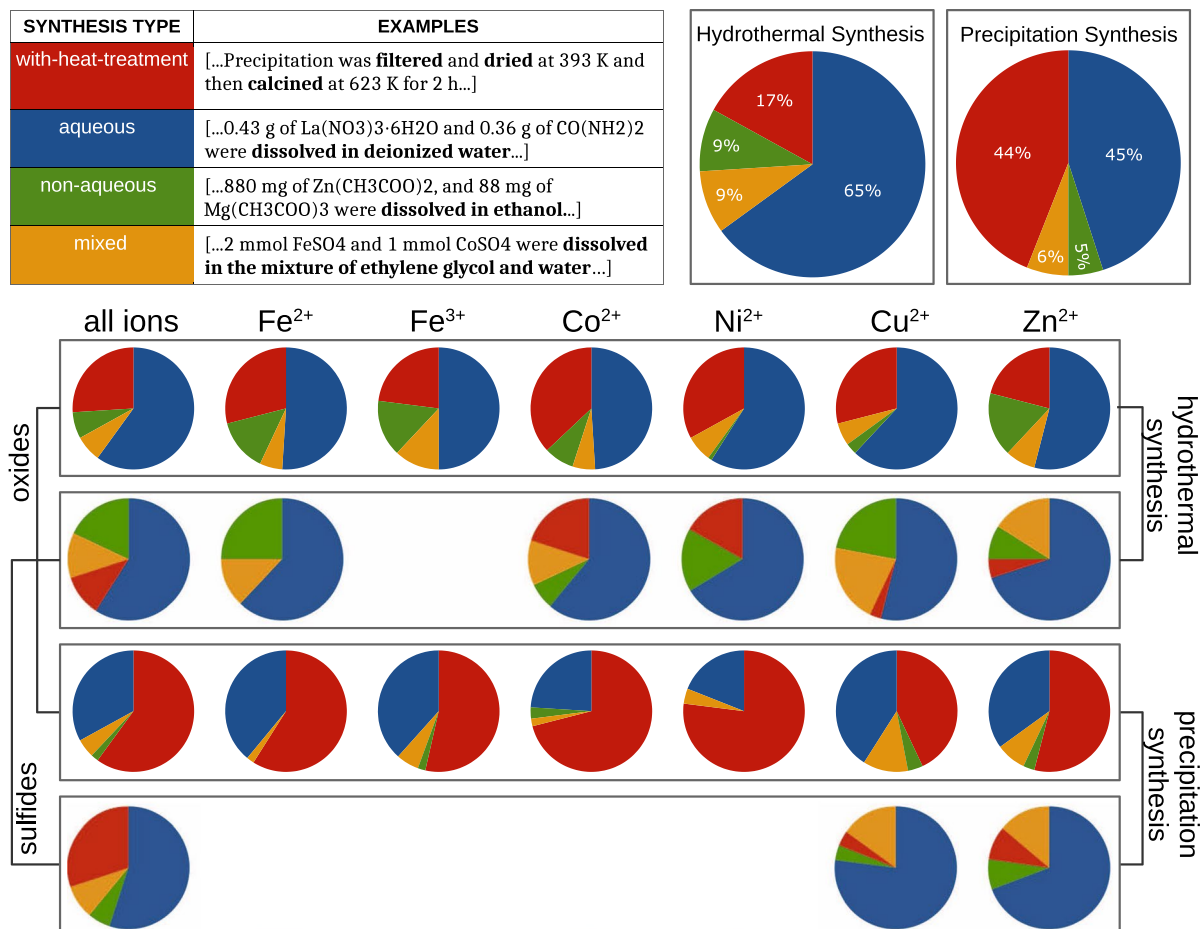


Fig. 3 Correspondence between choice of synthesis route and selected types of targets. The top table gives an example of the four synthesis categories defined: with heat treatment step, aqueous, non-aqueous, and mixed. The two pie-charts on the top-right show the fractions of synthesis routes in the hydrothermal and precipitation datasets separately. The four rows of pie charts in the lower half of the figure represent the fractions of the four synthesis routes (given in the table) for all oxides, all sulfides, and individual oxides and sulfides with different oxidation states of data-rich transition metals separately. The first and second rows are results from the hydrothermal dataset. The third and fourth rows are results from the precipitation dataset. Each blank space means that there is not enough data to form a statistic for the corresponding type of target.

Finally, solution-based synthesis is advantageous when the control of specimen morphology is required, e.g., when synthesizing noble-metal nanoparticles. However, this dataset does not provide information about the morphology of the synthesized materials, though such information is often contained in characterization or results paragraphs instead of the experimental section. The extraction of morphology and other solution synthesis outcomes is another text-mining challenge in materials science research that requires the development of advanced algorithms and models³⁰, which is beyond the scope of the current study.

Usage Notes

The dataset is provided in JSON format as a single file. All major programming languages, such as Python, Matlab, R, and Wolfram Mathematica, can be used to read it. No particular dependency is required.

Because the dataset contains detailed information about chemical formulas as well as the compositions of the target materials and precursors for each procedure, it can be easily used to conduct a literature review by querying desired precursors and (or) targets in different chemical spaces. For example, selecting all TiO_2 synthesized from TiCl_4 allows an exploration of how other synthesis formulations, such as synthesis actions, attributes, and quantities, affect the results. Furthermore, the materials entries in the dataset are supplied with the Materials Project⁵ identifiers, thus facilitating the integration of the procedures with the thermochemical data available in the Materials Project^{73,74}.

In addition, this solution-based synthesis dataset keeps the same data structure as that in the solid-state dataset generated in our previous work⁵³. Therefore, it is easy to analyze the procedures from the two datasets.

Despite the dataset being provided as a static snapshot⁶³, we intend to update it on a regular basis.

Code availability

The scripts used to classify paragraphs and extract procedures as well as to perform the data analysis are home-written codes which are publicly available at the GitHub repository https://github.com/CederGroupHub/text-mined-solution-synthesis_public with acknowledgement of the current paper.

The underlying libraries used in this project are all open-source:

Tensorflow (www.tensorflow.org)
Keras (keras.io/keras.io)
SpaCy (spacy.io/spacy.io)⁶⁰
NLTK (<https://www.nltk.org/>)⁶²
gensim (radimrehurek.com/radimrehurek.com)⁵⁹
scikit-learn (scikit-learn.org)⁷⁵
ChemDataExtractor (chemdataextractor.org)⁴⁰
Material Parser (github.com/CederGroupHub/MaterialParser)
Borges (github.com/CederGroupHub/Borges)
LimeSoup (github.com/CederGroupHub/LimeSoup).

Received: 7 October 2021; Accepted: 5 April 2022;

Published online: 25 May 2022

References

- Pankratius, V. *et al.* Computer-aided discovery: Toward scientific insight generation with machine support. *IEEE Intelligent Systems* **31**, 3–10 (2016).
- Tolle, K., Tansley, D. & Hey, A. The fourth paradigm: Data-intensive scientific discovery. *Proceedings of the IEEE* **99**, 1334–1337 (2011).
- Fischer, C. C., Tibbetts, K. J., Morgan, D. & Ceder, G. Predicting crystal structure by merging data mining with quantum mechanics. *Nature Materials* **5**, 641–646 (2006).
- Holden, J. Materials Genome Initiative for global competitiveness. Tech. Rep., National Science and Technology Council (2011).
- Ong, S. P. *et al.* Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science* **68**, 314–319 (2013).
- Jain, A. *et al.* Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL Materials* **1**, 011002 (2013).
- Bergerhoff, G. & Brown, I. D. *Crystallographic Databases* (F. H. Allen *et al.* (Hrsg.) Chester, International Union of Crystallography, 1987).
- Belsky, A., Hellenbrandt, M., Karen, V. L. & Luksch, P. New developments in the Inorganic Crystal Structure Database (ICSD): accessibility in support of materials research and design. *Acta Crystallogr. B* **58**, 364–369 (2002).
- Linstrom, P. & Mallard, W. (eds.) *NIST Chemistry WebBook, NIST Standard Reference Database Number 69* (National Institute of Standards and Technology, Gaithersburg MD, 20899, 2019).
- Blokhin, E. & Villars, P. *The PAULING FILE Project and Materials Platform for Data Science: From Big Data Toward Materials Genome*, 1–26. Handbook of Materials Modeling (Springer, Cham, 2018).
- Villars, P. & Cenzual, K. Pearson's Crystal Data: Crystal Structure Database for Inorganic Compounds (on DVD) (Release 2018/19). ASM International®, Materials Park, Ohio, USA.
- Sumpter, B. G., Vasudevan, R. K., Potok, T. & Kalinin, S. V. A bridge for accelerating materials by design. *npj Computational Materials* **1**, 15008 (2015).
- Bianchini, M. *et al.* The interplay between thermodynamics and kinetics in the solid-state synthesis of layered oxides. *Nature Materials* **19**, 1088–1095 (2020).
- Kovnir, K. Predictive synthesis. *Chemistry of Materials* **33**, 4835–4841 (2021).
- Kohlmann, H. Looking into the black box of solid-state synthesis. *European Journal of Inorganic Chemistry* **2019**, 4174–4180 (2019).
- Chamorro, J. R. & McQueen, T. M. Progress toward solid state synthesis by design. *Accounts of Chemical Research* **51**, 2918–2925 (2018).
- Sun, W. *et al.* The thermodynamic scale of inorganic crystalline metastability. *Science Advances* **2**, e1600225 (2016).
- Sun, W. *et al.* A map of the inorganic ternary metal nitrides. *Nature Materials* **18**, 732–739 (2019).
- Aykol, M., Dwaraknath, S. S., Sun, W. & Persson, K. A. Thermodynamic limit for synthesis of metastable inorganic materials. *Science Advances* **4**, eaaq0148 (2018).
- Aykol, M. *et al.* Network analysis of synthesizable materials discovery. *Nature Communications* **10**, 2018 (2019).
- Aykol, M., Montoya, J. H. & Hummelshøj, J. Rational solid-state synthesis routes for inorganic materials. *Journal of the American Chemical Society* **143**, 9244–9259 (2021).
- Bartel, C. J. *et al.* Physical descriptor for the gibbs energy of inorganic crystalline solids and temperature-dependent materials chemistry. *Nature Communications* **9**, 4168 (2018).
- McDermott, M. J., Dwaraknath, S. S. & Persson, K. A. A graph-based network for predicting chemical reaction pathways in solid-state materials synthesis. *Nature Communications* **12**, 3097 (2021).
- Duvenaud, D. K. *et al.* Convolutional networks on graphs for learning molecular fingerprints. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M. & Garnett, R. (eds.) *Advances in Neural Information Processing Systems* **28**, 2224–2232 (Curran Associates, Inc., 2015).
- Ley, S. V., Fitzpatrick, D. E., Ingham, R. J. & Myers, R. M. Organic synthesis: March of the machines. *Angew. Chem. Int. Edit.* **54**, 3449–3464 (2015).
- Segler, M. H. S., Preuss, M. & Waller, M. P. Planning chemical syntheses with deep neural networks and symbolic ai. *Nature* **555**, 604–610 (2018).
- Mikulak-Klucznik, B. *et al.* Computational planning of the synthesis of complex natural products. *Nature* **588**, 83–88 (2020).
- Goodman, J. Computer software review: Reaxys. *J. Chem. Inf. Model.* **49**, 2897–2898 (2009).
- Kim, S. *et al.* PubChem 2019 update: improved access to chemical data. *Nucleic Acids Res.* **47**, D1102–D1109 (2018).
- Kononova, O. *et al.* Opportunities and challenges of text mining in materials research. *iScience* **24** (2021).
- Gaultois, M. *et al.* Data-driven review of thermoelectric materials: Performance and resource considerations. *Chem. Mater.* **25**, 2911–2920 (2013).
- Ghadbeigi, L., Harada, J. K., Lettiere, B. R. & Sparks, T. D. Performance and resource considerations of Li-ion battery electrode materials. *Energy Environ. Sci.* **8**, 1640–1650 (2015).
- Oh, E. *et al.* Meta-analysis of cellular toxicity for cadmium-containing quantum dots. *Nature Nanotech.* **11**, 479 (2016).

34. Court, C. J. & Cole, J. M. Auto-generated materials database of Curie and Néel temperatures via semi-supervised relationship extraction. *Sci. Data* **5**, 180111 (2018).
35. Court, C. & Cole, J. Magnetic and superconducting phase diagrams and transition temperatures predicted using text mining and machine learning. *npj Comput. Mater* **6**, 1–9 (2020).
36. Huang, S. & Cole, J. M. A database of battery materials auto-generated using chemdataextractor. *Scientific Data* **7**, 260 (2020).
37. Onishi, T., Kadohira, T. & Watanabe, I. Relation extraction with weakly supervised learning based on process-structure-property-performance reciprocity. *Sci. Technol. Adv. Mater.* **19**, 649–659 (2018).
38. Tshitoyan, V. *et al.* Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature* **571**, 95–98 (2019).
39. Kim, E. *et al.* Materials synthesis insights from scientific literature via text extraction and machine learning. *Chem. Mater* **29**, 9436–9444 (2017).
40. Swain, M. C. & Cole, J. M. ChemDataExtractor: A toolkit for automated extraction of chemical information from the scientific literature. *J. Chem. Inf. Model.* **56**, 1894–1904 (2016).
41. Jessop, D. M., Adams, S. E., Willighagen, E. L., Hawizy, L. & Murray-Rust, P. OSCAR4: a flexible architecture for chemical text-mining. *J. Cheminformatics* **3**, 41 (2011).
42. Hawizy, L., Jessop, D. M., Adams, N. & Murray-Rust, P. ChemicalTagger: A tool for semantic text-mining in chemistry. *J. Cheminformatics* **3**, 17 (2011).
43. Krallinger, M., Rabal, O., Lourenço, A., Oyarzabal, J. & Valencia, A. Information retrieval and text mining technologies for chemistry. *Chem. Rev.* **117**, 7673–7761 (2017).
44. Eltyeb, S. & Salim, N. Chemical named entities recognition: A review on approaches and applications. *J. Cheminformatics* **6**, 17 (2014).
45. Korvigo, I., Holmatov, M., Zaikovskii, A. & Skoblov, M. Putting hands to rest: efficient deep cnn-rnn architecture for chemical named entity recognition with no hand-crafted rules. *J. Cheminform.* **10**, 28 (2018).
46. He, T. *et al.* Similarity of precursors in solid-state synthesis as text-mined from scientific literature. *Chemistry of Materials* **32**, 7861–7873 (2020).
47. Weston, L. *et al.* Named entity recognition and normalization applied to large-scale information extraction from the materials science literature. *J. Chem. Inf. Model.* **59**, 3692–3702 (2019).
48. Kuniyoshi, F., Makino, K., Ozawa, J. & Miwa, M. Annotating and extracting synthesis process of all-solid-state batteries from scientific literature. In *Proceedings of The 12th Language Resources and Evaluation Conference, 1941–1950* (European Language Resources Association, Marseille, France, 2020).
49. Kim, E. *et al.* Machine-learned and codified synthesis parameters of oxide materials. *Sci. Data* **4**, 170127 (2017).
50. Kim, E., Huang, K., Jegelka, S. & Olivetti, E. Virtual screening of inorganic materials synthesis parameters with deep learning. *npj Comput. Mater* **3**, 53 (2017).
51. Jensen, Z. *et al.* A machine learning approach to zeolite synthesis enabled by automatic literature data extraction. *ACS Cent. Sci.* **5**, 892–899 (2019).
52. Kim, E. *et al.* Inorganic materials synthesis planning with literature-trained neural networks. *J. Chem. Inf. Model.* **60**, 1194–1201 (2020).
53. Kononova, O. *et al.* Text-mined dataset of inorganic materials synthesis recipes. *Scientific Data* **6**, 203 (2019).
54. Mouchère, H., Zanibbi, R., Garain, U. & Viard-Gaudin, C. Advancing the state of the art for handwritten math recognition: the crohme competitions, 2011–2014. *IJDAR* **19**, 173–189 (2016).
55. Mahdavi, M., Zanibbi, R., Mouchère, H., Viard-Gaudin, C. & Garain, U. Icdar 2019 crohme + tfd: Competition on recognition of handwritten mathematical expressions and typeset formula detection. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, 1533–1538 (IEEE, 2019).
56. Devlin, J., Chang, M.-W., Lee, K. & Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186 (Association for Computational Linguistics, Minneapolis, Minnesota, 2019).
57. Huo, H. *et al.* Semi-supervised machine-learning classification of materials synthesis procedures. *npj Computational Materials* **5**, 62 (2019).
58. Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S. & Dean, J. Distributed representations of words and phrases and their compositionality. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z. & Weinberger, K. Q. (eds.) *Advances in Neural Information Processing Systems* **26**, 3111–3119 (Curran Associates, Inc., 2013).
59. Řehůřek, R. & Sojka, P. Software framework for topic modelling with large corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, 45–50 (ELRA, Valletta, Malta, 2010).
60. Honnibal, M. & Johnson, M. An improved non-monotonic transition system for dependency parsing. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 1373–1378 (Association for Computational Linguistics, Lisbon, Portugal, 2015).
61. Jurafsky, D. & Martin, J. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall Series in Artificial Intelligence (Pearson Prentice Hall, 2009).
62. Bird, S., Edward, L. & Ewan, K. *Natural Language Processing with Python*. (O'Reilly Media Inc, 2009).
63. Wang, Z. *et al.* Dataset of solution-based inorganic materials synthesis procedures extracted from the scientific literature. *figshare* <https://doi.org/10.6084/m9.figshare.16583387.v4> (2021).
64. Wang, D., Xie, T. & Li, Y. Nanocrystals: Solution-based synthesis and applications as nanocatalysts. *Nano Research* **2**, 30–46 (2009).
65. Rakhî, R. B., Chen, W., Cha, D. & Alshareef, H. N. Influence of calcination temperature on the morphology and energy storage properties of cobalt oxide nanostructures directly grown over carbon cloth substrates. *Materials for Renewable and Sustainable Energy* **2**, 17 (2013).
66. Augustin, M. *et al.* Manganese oxide phases and morphologies: A study on calcination temperature and atmospheric dependence. *Beilstein journal of nanotechnology* **6**, 47–59 (2015).
67. Ellingham, H. Transactions and communications. *Journal of the Society of Chemical Industry* **63**, 125–160 (1944).
68. Byrappa, K. & Yoshimura, M. *Handbook of Hydrothermal Technology* (William Andrew Publishing, Oxford, 2013).
69. Raccuglia, P. *et al.* Machine-learning-assisted materials discovery using failed experiments. *Nature* **533**, 73–76 (2016).
70. Jia, X. *et al.* Anthropogenic biases in chemical reaction data hinder exploratory inorganic synthesis. *Nature* **573**, 251–255 (2019).
71. Burger, B. *et al.* A mobile robotic chemist. *Nature* **583**, 237–241 (2020).
72. Szymanski, N. J. *et al.* Toward autonomous design and synthesis of novel inorganic materials. *Mater. Horiz.* **8**, 2169–2198 (2021).
73. Persson, K. A., Waldwick, B., Lazic, P. & Ceder, G. Prediction of solid-aqueous equilibria: Scheme to combine first-principles calculations of solids with experimental aqueous states. *Phys. Rev. B* **85**, 235438 (2012).
74. Sun, W., Kitchaev, D. A., Kramer, D. & Ceder, G. Non-equilibrium crystallization pathways of manganese oxides in aqueous solution. *Nature Communications* **10**, 573 (2019).
75. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. DMR-1922372. We also thank Chris Bartel, and Amalie Trewartha for valuable discussions.

Author contributions

Z.W. developed the material quantities extraction algorithm, reaction formula algorithm, Materials Parser, operations and conditions extraction algorithms, analyzed the data, and wrote the manuscript. O.K. developed the Materials Parser, operations and conditions extraction algorithms, analyzed the data, and wrote the manuscript. K.C. developed operations and conditions extraction algorithms, manually inspected the dataset, and wrote the manuscript. T.H. developed materials entity recognition algorithm, and manually inspected the dataset. H.H. curated paragraphs databases, and developed the paragraphs classifier. Y.F. developed building reaction formula algorithm and manually inspected the dataset. Y.Z. analyzed the data and manually inspected the dataset. Y.S. analyzed the data and manually inspected the dataset. Z.C. analyzed the data and manually inspected the dataset. W.S. analyzed the data and developed material quantities extraction algorithm. G.C. developed approach, supervised the project, and wrote the manuscript. All authors discussed the results and contributed to the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to G.C.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022