# Porous Pseudomaterials for Studying Structure-Property Relationships of Gas Adsorption

by

#### Alec Reino Kaija

Bachelor of Science, Chemical Engineering, University at Rhode Island, 2014

Bachelor of Arts, Spanish, University of Rhode Island, 2014

Submitted to the Graduate Faculty of

the Swanson School of Engineering in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2019

#### UNIVERSITY OF PITTSBURGH

#### SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

#### Alec Reino Kaija

It was defended on

May 6, 2019

#### and approved by

Christopher Wilmer, Ph.D., Assistant Professor, Department of Chemical and Petroleum Engineering

Giannis Mpourmpakis, Ph.D., Assistant Professor, Department of Chemical and Petroleum Engineering

Nathaniel Rosi, Ph.D., Professor, Department of Chemistry

Goetz Veser, Professor, Ph.D., Department of Chemical and Petroleum Engineering

Dissertation Director: Christopher E. Wilmer, Ph.D., Assistant Professor, Department of Chemical and Petroleum Engineering Copyright © by Alec Reino Kaija

2019

#### Porous Pseudomaterials for Studying Structure-Property Relationships of Gas Adsorption

Alec Kaija, PhD

University of Pittsburgh, 2019

The discovery in 1995 of metal-organic frameworks (MOFs) – with record-breaking surface areas – sparked exponential growth in research efforts dedicated to the development of new porous adsorbents, particularly for energy related gas storage applications. However, despite their promise, decades of research have yet to yield MOFs that perform well enough for many of these applications, particularly high-pressure vehicular natural gas storage and post-combustion carbon capture. To understand why, I developed a novel computational methodology for generating large (100,000+) libraries of randomly configured Lennard-Jones (LJ) crystals, or "pseudomaterials", with the intention of calculating various adsorption-related properties of interest *en masse* using grand canonical Monte Carlo (GCMC) simulations. These libraries were used to map an *n*-dimensional structure-property space, where *n* refers to the number different structure- and property-parameters.

One approach for generating these libraries of pseudomaterials is *random sampling*, where each structure is generated algorithmically at random; however, we attempt to improve the overall computational efficiency using alternative methods. These alternative methods include *mutation algorithms* and augmenting random sampling with *property prescreening*. Mutation algorithms identify pseudomaterials with unique structure-property combinations and selectively perturb their structural characteristics to sample the sparsely populated regions of the structure-property space. Property prescreening uses machine learning models to predict the properties of a pseudomaterial to justify whether it is a candidate worthy of more computationally expensive GCMC simulations; it is an attempt at reducing the computational expense associated with running simulations on pseudomaterials with redundant properties.

The overall goal of this new computational methodology was to observe structure-property relationships for porous materials *in general* (i.e., not limited to any particular sub-class). I showed that understanding these structure-property relationships provides insights into the design of better adsorbents for a wide range of gas storage and separations applications. In the future, this methodology could potentially be extended to better understanding porous materials for catalysis, sensing, and more.

### **Table of Contents**

Prefacexvii
1.0 A Brief History of Crystal Engineering: Motivation for Studying Structure-
Property Relationships1
1.1 An Introduction to Crystal Engineering1
1.2 Computational Crystal Engineering3
2.0 Pseudomaterials: Lennard-Jones Crystals for Sampling the Structure-Property
Space for Porous Adsorbents5
2.1 Algorithmically generating pseudomaterials5
2.2 HTSoHM: High-Throughput Screening of Hypothetical Materials, a Software
Package for Generating Libraries of Pseudomaterials
2.3 Atomistic models used for simulating properties of interest
2.4 Improving the computational efficiency of HTSoHM10
2.4.1 Mutation Algorithms: An Alternative Approach to Generating
Pseudomaterials10
2.4.2 OEDIPUS: A Software Package for Rapidly Testing New Methods
2.4.3 Adaptive Mutation Algorithms
2.4.4 Property Prescreening: Using Machine Learning to Improve Computational
Efficiency 32
3.0 Efficiently Mapping Structure-Property Relationships of Gas Adsorption in
Porous Materials: Application to Xe Adsorption42
3.1 Introduction

3.2 Methodology	45
3.2.1 Generating the seed population	48
3.2.2 Modeling properties of interest	50
3.2.3 Selecting rare materials	50
3.2.4 Removing anomalous results	50
3.2.5 Assigning mutation strength	51
3.2.6 Mutating parents, creating new materials	52
3.2.7 Uniformity metric	52
3.3 Results	53
3.4 Conclusions	60
4.0 High-Pressure Methane Adsorption in Porous Lennard-Jones Crystals	61
4.1 Introduction	62
4.2 Methodology	64
4.2.1 Lennard-Jones potential	64
4.2.2 Generating pseudomaterials	66
4.2.3 Simulating properties	68
4.3 Results and Discussions	68
4.4 Conclusion	81
5.0 Carbon Capture Structure-Property Relationships in Lennard-Jones + Coulomb	
Porous Crystals	82
5.1 Introduction	83
5.2 Methods	87
5.2.1 Generating pseudomaterials with partial charges	87

5.2.2 Lennard-Jones + Coulomb potential	
5.2.3 Simulating properties	
5.3 Results	
5.4 Conclusion	104
6.0 Commercial Impact of Pseudomaterials Research	105
6.1 mediPOD: Product Ideation and Business Competitions	107
6.2 Incorporation, Fundraising, and Product Development	110
Bibliography	114

### List of Tables

Table 1. Pseudomaterial generation steps and associated parameter ranges 6
Table 2. Pseudomaterial generation steps and related parameter ranges
Table 3. Adsorption evaluation criteria used by Bae and Snurr to assess sorbents for CCS
Table 4. LJ Parameters and partial charges for framework pseudoatoms and adsorbate molecules
Table 5. Aeronics' winnings via business competitions

### List of Figures

Figure 1. RDBMS schema for storing simulated data and structural information for
pseudomaterials libraries
Figure 2. Flowchart for HTSoHM software. First a worker node is launched (A) and then it enters
the worker run loop to generate a pseudomaterial structure (B), simulate properties of
interest (C), and store all these data in the database
Figure 3. Diagram of how pseudomaterials are select to parent new materials through mutation.
The least populated bins (green) contain pseudomaterials most likely to be selected as
parents whereas the most populated bins (red) contain pseudomaterials least likely to be
selected as parents
Figure 4. 2D projections of the structure-property space showing methane loading with respect to
volumetric surface area. Shown are ten generations of 1,000 children each using different
mutation strengths: 10% (A), 20% (B), 50% (C), and starting at 50% decreasing by 5%
with each generation. Children added in each generation are colored while all previously
generated pseudomaterials are represented in black14
Figure 5. Uniformity metrics across ten generations of 1,000 pseudomaterials each
Figure 6. Example of statistical outliers. 2D projection of structure-property space showing
volumetric gas loading with respect to helium void fraction. One parent is indicated by a
light green point and its children by dark green points. Another parent is indicated by a
light blue point and its children by dark blue points. The entire mapping of the structure-
property space is represented by light grey points17

Figure 7. Examples of pseudomaterials sensitive to mutation. A 2D projection of the structureproperty space (*right*) shows volumetric gas loading with respect to helium void fraction where the children of three parents are indicated by cyan, magenta, and yellow points and the remaining pseudomaterials in the library are represented by grey points. All three of these parents are in the same bin (top), yet none of their children appear in the same bin as their parents. A histogram (*left*) shows the number of children in each of the four bins they appear in, none of which are the same as the parents. This data comes from a library of 5,000 pseudomaterials; thousands of these materials were generated by mutating these Figure 8. Comparison of convergence using empty bin counts between two libraries, one generated by random sampling, indicated by dashed lines, and another by mutation at 20% strength, indicated by solid lines. Different thresholds were used to determine whether a bin was occupied; either 1 (cyan), 10 (magenta), 100 (green), or 250 (purple) pseudomaterials Figure 11. Histograms of  $\alpha$  (*left*) and  $\beta$  (*middle*) values from a population of 10,000 boxes generated using a random sampling method. A scatterplot (*right*) represents the mapping Figure 12. Convergence displayed as percentage of empty bins remaining for four 100+ generation libraries of hypothetical boxes: using random sampling (red), using a constant mutation strength of 20% (blue), using bin-specific adaptive mutation strengths (green), and using 

Flowchart for HTSoHM with propert prescreening where the property prescreening	rty prescreening
os are denoted by dahsed lines	
Comparisons of random sampling and propert prescreening (PROMETHEUS method	[HEUS method]
ng number of empty bins as a convergence criteria and a convex hull to assess whether	o assess whether

- Figure 28. 2D histograms for samples of 200 (A); 2,000 (B); 20,000 (C); and 200,000 (D) pseudomaterials from a library of 300,000 pseudomaterials. Plots show projections of the structure-property space in methane loading (at 35 bar) with respect to helium void fraction and are colored by the number of pseudomaterials in each of 40×40 equally-sized bins. 70

- Figure 38. 2D projectsions of the CCS structure-property space depicting the relationship between CO<sub>2</sub> working capacity and different structural characteristics: helium void fraction (A), volumetric surface area (B), and number density (C), all colored by regenerability....... 97

Figure 42. Preliminary data generated for oxygen storage in porous pseudomaterials at 30 bar, the
operating pressure of standard oxygen cylinders. Shown is a 3D mapping of a structure-
property space consisting of volumetric oxygen loading, helium void fraction, and
volumetric surfae area106
Figure 43. Concept art for mediPOD, cross-sectional view showing adsorbent and protective filter
to prevent inhalation of sorbent
Figure 44. Concept art for mediPOD, demonstration of flexible adaptor connecting cannister to
user's existing cannula109
Figure 45. Everyday Oxygen, recreational oxygen with three times the capacity of the leading
competitor
Figure 46. Pawprint Oxygen, emergency oxygen for the veterinary market

#### Preface

My thesis, entitled "Porous pseudomaterials for studying structure-property relationships of gas adsorption", describes a novel computational methodology which I have developed with the support and guidance of my Principal Investigator, Dr. Chris Wilmer. It has been written to fulfill the graduation requirements of the Department of Chemical & Petroleum Engineering at the University of Pittsburgh. I was engaged in the research described in this dissertation from July 2014 to April 2019.

In truth, I could not have achieved my current level of success without a strong support group. First, I would like to acknowledge, my parents, who have provided me with unwavering love and understanding. Secondly, I would like to acknowledge my committee members, who have provided patient advice and guidance through the research process. I would like to acknowledge the freedom granted to me by my Principal Investigator, Dr. Chris Wilmer, to pursue my wide and varied interests which have included – during the time of my PhD (and have not been limited to) – entrepreneurship, management consulting, and more. I would like to acknowledge the technical guidance of my colleague, Paul Boone, as well as his wealth of knowledge as a software developer. I would like to acknowledge the tutelage of Jan Steckel, my mentor at the US Department of Energy Nation Energy Technology Laboratory, where I spent the final months of PhD. And last, but not least, I would like to acknowledge my business partners, Blake Dube and Mark Spitz, with whom I co-founded Aeronics, Inc.

This work includes two peer-reviewed articles, published in *Faraday Discussions* and *Journal of Physical Chemistry Letters*, as well as one that was submitted to *Journal of Physical* 

*Chemistry C*. The work also reflects efforts put towards the publication of a review paper submitted to *Molecular Simulation*.

It would be remiss of me not to acknowledge the many struggles I have faced throughout my PhD and their contribution to developing my indomitable spirit. I am reminded of the words of Seneca the Younger: "Difficulties strengthen the mind, as labor does the body."

Thank you.

**RES SECUNDAE**.

### 1.0 A Brief History of Crystal Engineering: Motivation for Studying Structure-Property Relationships

#### 1.1 An Introduction to Crystal Engineering

The term crystal engineering was first coined in 1971 by Schmidt when he described fundamental rules which could someday be used to design crystalline structures. By the late 1980s this Holy Grail-type problem remained unsolved to the dismay of researchers like Maddox who called it "one of the continuing scandals in the physical sciences."[1] In 1989 the term received a revised definition from Desiraju: "the understanding of intermolecular interactions in the context of crystal packing and the utilization of such understanding in the design of new solids with desired physical and chemical properties."[2] Because of the many, minute intermolecular forces that govern crystal packing,[3] crystal engineering is a daunting task even with modern supercomputing resources. This has not stopped, however, a nearly exponential growth in research efforts dedicated to crystal design and crystal engineering.

Crystal engineering, particularly in the context of organic solids, has become inherently linked to concepts derived from supramolecular chemistry. As defined by Lehn in 1978, supramolecular chemistry is the "chemistry beyond the molecule," concerning complex entities which arise from the association of two or more chemical species via intermolecular forces.[4] One can consider crystals as single chemical entities and as such special examples of supramolecular assemblies; in fact, Dunitz referred to the crystal as a "supermolecule par excellence," held together by ion-ion, ion-dipole, dipole-dipole interactions, hydrogen bonding, London forces, and more.[5] Schmidt emphasized that the physical properties of crystalline solids depend upon the distribution of chemical components within the crystal lattice as properties of the individual components. Schmidt's work spawned a series of publications by Desiraju[6-8] and Etter[9-11] which used the Cambridge Structural Database (CSD) to investigate noncovalent bonding in organic solids. This work established the concept of supramolecular synthons,[6] making hydrogen bonds the most widely exploited noncovalent interaction in crystal engineering.

One approach to crystal engineering of functional solids involves constructing polymeric networks. The work of Wells, who was particularly interested in inorganic structures, [12, 13] established a system of defining structures according to their topology by reducing them to points (nodes) of a particular geometry (tetrahedral, trigonal planar, etc.). This approach wasn't reflected in laboratory work until the 1990s when Robson facilitated the development of the field of coordination polymers.[14-17] The field was further developed by Kitagawa[18] and Yaghi,[19] who showed that it was possible to apply Robson's approach to create extensive families of compounds with specific structural topologies. Yaghi coined the term reticular chemistry, referring to the systematic synthesis of families of frameworks.[20] The term metal-organic framework also first appeared in a publication by Yaghi and Li in 1995.[21] A few years later Yaghi and Li, in what would become a highly influential publication, reported the structure of MOF-5.[19] This new class of materials, because of their extremely high surface areas, has attracted substantial interest in a wide range of applications[22-29] including gas storage[19, 30-37] and separations, [38-45] catalysis, [46-51] and sensing. [52-55] Though relatively little is known about the mechanisms of self-assembly in MOFs, [56] their highly tunable nature represents a significant step towards the dream of crystal engineering.

#### **1.2 Computational Crystal Engineering**

One of the primary challenges in engineering crystals from first principles is the massiveness of the phase space. For four-element stoichiometries, ignoring geometries, there are 1012 possible structures, and when more than four elements and their geometries are considered this number exceeds 10100 structures.[57] Nevertheless, sampling this vast space, even incompletely, provides improved insight particularly for investigating structure-property relationships. In 1999 Delgado enumerated all possible networks in which each atom is connected to exactly four neighbors, representative of covalently bonded crystals, silicates, hydrates, crystalline elements, and more.[58] In the field of zeolites, Deem developed methods for generating large libraries of hypothetical structures.[59, 60] Haldoupis later conducted pore size analysis for >250,000 hypothetical zeolites in Deem's library, calculating Henry's coefficients and diffusion activation energy for  $H_2$  and  $CH_4$  in a subset of 8000 structures.[61] This was a major step towards deriving structure-property relationships in this class of materials which at the time only consisted of around 190 synthesizable structures.

For MOFs Mellot-Draznieks developed the automated assembly of secondary building units (AASBU) method, which focused on the topology of network-based structures, exploring the possible ways to assemble predefined inorganic building blocks in three-dimensional space.[62] By 2010 Yazaydin synthesized the first computationally-designed MOF, NU-100, which broke records for methane storage.[63] At the same time Wilmer sought to create an algorithm which would generate every possible MOF from a particular library of building blocks.[64] By 2013 Wilmer's methods were used to design a new record-breaking methane adsorbent whose experimental performance nearly matched computational predictions.[34] Despite the vastness of the hypothetical MOF libraries generated to date, it remains uncertain whether better adsorbents,

to name one application area, exist. More work is required to establish simple design rules or overarching laws governing the performance of these and other crystalline materials.

This work develops a novel computational methodology for generating large libraries of hypothetical structures, which we call *pseudomaterials*. Past work has focused on generating libraries of hypothetical structures by assembling modular molecular building blocks, according to design rules observed in nature. One potential shortcoming of this approach is the possibility of generating structures from molecular building blocks which have not been previously considered or assembling these molecular building blocks according to design rules which have not yet been observed in nature. This means that the regions of the structure-property space for porous crystalline materials observed to date represent an incomplete view of nature – by generating large libraries of pseudomaterials we hope to fill in the gaps.

Pseudomaterials are random configurations of matter, generated in an attempt at sampling both the previously observed regions of the structure-property space for porous crystalline materials and those disparate regions which may not yet have been explored by past work. If one were to generate an infinite number of different pseudomaterials – an admittedly impossible feat – the result would be a complete map of the structure-property space for porous crystalline materials. While it is not feasible to create an infinitely large library of pseudomaterials, we have developed strategies towards uniformly sampling the structure-property space. This work approaches two feats: a rigorous exploration of structure-property relationships and the establishment of theoretical performance limits for specific application areas of porous crystalline materials. The work described herein is the just the beginning: it is our hope that others in the scientific community will pick up where we have left off.

4

## 2.0 Pseudomaterials: Lennard-Jones Crystals for Sampling the Structure-Property Space for Porous Adsorbents

#### 2.1 Algorithmically generating pseudomaterials

A pseudomaterial, in the simplest sense, is a configuration of Lennard-Jones (LJ) spheres within a unit cell. This can be any configuration of LJ spheres, regardless of whether it represents a real or hypothetically synthesizable structure. In fact, most pseudomaterials could never be synthesized. Pseudomaterials are an abstract representation of a configuration of matter for the sake of sampling a structure-property space. A library of pseudomaterials therefore represents a sampling, or mapping, of the structure-property space for porous materials.

Pseudomaterials are generated by randomly positioning LJ spheres, which we refer to as pseudoatoms, within a randomly-sized unit cell. The unit cell dimensions were bounded between 25.6 and 51.2 Å in each of the crystallographic directions (with the lower bound twice the cutoff length used for gas-gas and gas-crystal interactions: 12.8 Å). The number of LJ spheres within a unit cell is bounded between  $1.49 * 10^{-5}$  and 0.02122 atoms / Å<sup>3</sup>, where the lower boundary ensures that each unit cell contains at least two atom sites and the upper boundary corresponds to 10% the number density of iron. Pseudoatom types are defined by their LJ parameters:  $\sigma$ , the van der Waals radius, and  $\varepsilon$ , the potential well depth. Values for  $\sigma$  are bounded between 1.052 and 6.549 Å and values for  $\varepsilon$  between 1.258 and 513.264 K. This range of LJ parameters was based on the Universal Force Field (UFF),[65] where  $\sigma/\varepsilon$ -values were allowed to be 50% lower or greater, respectively, than the minimum and maximum values present in the UFF. While a pseudomaterial's unit cell

may contain hundreds of pseudoatoms depending upon the randomly selected number density, the number of pseudoatom *types* is limited to four in all pseudomaterials.

	Pseudomaterial generation step	Parameter range
1	Select lattice constants	25.6 – 51.2 Å
2	Select number density	$1.49 \times 10^{5} - 0.02122$ atoms/Å <sup>3</sup>
3	Position pseudoatom sites	N/A
4	Select four sets of LJ parameters (one for each pseudoatom type)	σ: 1.052 – 6.549 Å ε: 1.258 – 513.264 K

Table 1. Pseudomaterial generation steps and associated parameter ranges

## 2.2 HTSoHM: *High-Throughput Screening of Hypothetical Materials*, a Software Package for Generating Libraries of Pseudomaterials

To sample a substantial region of the structure-property space, a large library (>100,000) of pseudomaterials is needed. Because of scale, the process of generating such a library and simulating properties for each pseudomaterial must be automated. This task can be divided into two smaller tasks: (1) create a data structure for storing simulations results and pseudomaterial structural data and (2) devise a system for distributing work on a computing cluster. These two

problems may seem disparate; however, the use of a computing cluster dictates the use of a database. A relational database management system (RDBMS) is used because it allows concurrent access from multiple computing nodes and provides a convenient method for storing structural and simulation data simultaneously (see Figure 1).



Figure 1. RDBMS schema for storing simulated data and structural information for pseudomaterials libraries.

The database is in PostgreSQL using SQLAlchemy as an API to interface with the main program – called HTSoHM, short for *High-Throughput Screening of Hypothetical Materials* – which is written in Python. The program consists of two main pieces of code: a PBS shell script containing instructions to start a worker node using qsub (Figure 2a) and a Python script for the worker run loop (Figure 2b, Figure 2c, Figure 2d) in which the worker creates a pseudomaterial structure (Figure 2b), simulates properties of interest (Figure 2c), and stores the structure in the database (Figure 2d), repeating the process until enough pseudomaterials have been added to the database.



Figure 2. Flowchart for HTSoHM software. First a worker node is launched (A) and then it enters the worker run loop to generate a pseudomaterial structure (B), simulate properties of interest (C), and store all these data in the database.

The methods for generating pseudomaterial structures have been covered previously and properties of interest that are simulated include volumetric gas loading, volumetric surface area, and helium void fraction. This overall approach to generating libraries of pseudomaterials at random is referred to as *random sampling*, attempts at improving the efficiency of this overall computational methodology will be described later.

The HTSoHM software package is free and open source and available online at: https://github.com/WilmerLab/HTSOHM-dev

#### 2.3 Atomistic models used for simulating properties of interest

A simulation package called RASPA was used to simulate all properties of interest. Grand canonical Monte Carlo (GCMC) simulations are used to calculate volumetric gas loading according to previously established methods.[35, 66] Each calculation uses a 1000 cycle equilibration period followed by a 1000 cycle production run. A cycle consists of n Monte Carlo steps, where n equals the number of molecules in the system (which fluctuates). Simulations include insertion, deletion, and translation moves of molecules with equal probabilities. Fugacities are calculated using the Peng-Robinson equation of state.

GCMC simulations are also used to calculate void fraction in each pseudo material according to established methods.[67] Each simulation consists of a 1000 cycle production run. Simulations use the Widom particle insertion method with a helium atom probe ( $\sigma$ =2.96Å). Average Rosenbluth weight is recorded for the void fraction of each pseudomaterial.

Interaction energies from non-bonded atoms are computed through the Lennard-Jones potential:

$$V_{ij} = 4\epsilon_{ij} \left( \left( \frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left( \frac{\sigma_{ij}}{r_{ij}} \right)^6 \right)$$
(2-1)

where *i* and *j* are interacting atoms,  $r_{ij}$  is the distance between atoms *i* and *j*,  $\sigma_{ij}$  is the LJ diameter, and  $\varepsilon_{ij}$  is the LJ well depth. The LJ cutoff was 12.8 Å. LJ parameters for framework pseudoatoms were selected at random from a range of values 50% larger than the UFF.

#### 2.4 Improving the computational efficiency of HTSoHM

We have explored various approaches improving the efficiency with which HTSoHM samples the structure-property space for porous adsorbents. Two of these approaches are *mutation algorithms* and *property prescreening*. Mutation algorithms introduce an alternative method of generating pseudomaterials via mutation or perturbing the structural characteristics of existing pseudomaterials to more rigorously sample specific regions of the structure-property space. Property prescreening augments random sampling by using machine learning models to predict the property-combinations of a pseudomaterial to determine whether it is worth spending computing resources on simulating properties of interest for that structure. We have also developed a software package, OEDIPUS, for testing the effectiveness of new methods more quickly than by generating pseudomaterials libraries and calculating properties of interest using GCMC simulations.

#### 2.4.1 Mutation Algorithms: An Alternative Approach to Generating Pseudomaterials

The approach described above, simply randomly generating structures, represents an attempt at randomly sampling the structure-property space. While this approach is effective, it stands to be improved by developing a different method for generating materials, one which

doesn't rely solely on randomness. To accomplish this, we proposed (after having randomly generated a seed population) selecting those pseudomaterials with unique structure-property combinations and using them to create new pseudomaterials via *mutation*. This approach requires the addition of two algorithms: one to select these unique, *parent* pseudomaterials and another to mutate them. There is more than one way to identify a pseudomaterial's uniqueness in this context; however, we chose to use the number of pseudomaterials in each of a congruent three-dimensional bin. This requires us to establish arbitrary maxima in each dimension, then subdivide the structure-property space into equally-sized bins, counting the number of pseudomaterials in each bin. Those bins with the fewest pseudomaterials contain the most unique pseudomaterials, which are most likely to be selected to parent new pseudomaterials as shown in Figure 3.



Figure 3. Diagram of how pseudomaterials are select to parent new materials through mutation. The least populated bins (*green*) contain pseudomaterials most likely to be selected as parents whereas the most populated bins (*red*) contain pseudomaterials least likely to be selected as parents.

Mutating pseudomaterials consists of perturbing each of the defining parameters: atom-site positions, Lennard-Jones parameters, lattice constants, and number density. Generating a mutated child can be concisely described as linearly interpolating each defining value between the parent and a completely random pseudomaterial (unrelated to the parent) to a degree dictated by the mutation strength. At 0%, the child is a clone of its parent, and at 100% the child is a completely new randomly generated pseudomaterial. We can observe the effect of using different mutation strengths in Figure 4.



Surface area (m<sup>2</sup>/cm<sup>3</sup>)

Figure 4. 2D projections of the structure-property space showing methane loading with respect to volumetric surface area. Shown are ten generations of 1,000 children each using different mutation strengths: 10% (A), 20% (B), 50% (C), and starting at 50% decreasing by 5% with each generation. Children added in each generation are colored while all previously generated pseudomaterials are represented in black.

To compare the effectiveness of each mutation strength, we defined a uniformity metric, U. First each bin-count, c, was normalised:

$$c^* = \frac{c}{c_{max}} \tag{2-2}$$

Here  $c_{max}$  is the highest bin-count observed. The uniformity metric is then calculated as the variance of the normalized bin-counts for the number of non-zero bins in the dataset, *N*:

$$U = \frac{\sum \left(c^* - \frac{\sum c^*}{N}\right)}{N} \tag{2-3}$$

The lower the value of U, the more uniformly the structure–property space has been explored. A high value of U indicates that some structure–property combinations are over-represented relative to others.

In Figure 5 we use this uniformity metric to compare the effectiveness of different mutation strengths.



Figure 5. Uniformity metrics across ten generations of 1,000 pseudomaterials each.

Figure 5 shows a 20% mutation strength as being the most effective for uniformly sampling the structure-property space; however, for all mutation strengths evaluated the uniformity remains relatively constant after a few generations. Visual inspection of the mappings in Figure 4 show that after a few generations very few of the unexplored regions of the structure-property space are accessed. One cause for this is statistical outliers; because GCMC simulations are used to calculate the properties of interest it is possible that a simulation may not converge in the number of cycles used. These statistical outliers falsely demonstrate structure-property combinations that make them unique as shown in Figure 6, causing them to be selected as parents while their children never appear in the same bin as them, increasing the uniformity metric. One way to address this issue is to recalculate all the properties of interest and compare them to the original value; if the recalculated value is substantially different from the originally calculated value the material can be flagged to prevent it from being selected as a parent.



Void fraction (dim.)

Figure 6. Example of statistical outliers. 2D projection of structure-property space showing volumetric gas loading with respect to helium void fraction. One parent is indicated by a light green point and its children by dark green points. Another parent is indicated by a light blue point and its children by dark blue points. The entire mapping of the structure-property space is represented by light grey points. Another limiting factor is the presence of pseudomaterials which are particularly sensitive to mutation, making it highly unlikely (at a particular mutation strength) that they will produce children in the same bin as the parent. This leads to a sort of infinite loop in which a unique pseudomaterial is repeatedly selected to parent children while none of the children appear in the same bin as the parent as demonstrated in Figure 7.


Figure 7. Examples of pseudomaterials sensitive to mutation. A 2D projection of the structure-property space (*right*) shows volumetric gas loading with respect to helium void fraction where the children of three parents are indicated by cyan, magenta, and yellow points and the remaining pseudomaterials in the library are represented by grey points. All three of these parents are in the same bin (*top*), yet none of their children appear in the same bin as their parents. A histogram (*left*) shows the number of children in each of the four bins they appear in, none of which are the same as the parents. This data comes from a library of 5,000 pseudomaterials; thousands of these materials were generated by mutating these three parents alone, indicative of the infite loop problem.

This infinite loop problem manifests itself when comparing libraries generated using the mutation approach with a 20% mutation strength and those generated by random sampling. In addition to the uniformity metric we may also compare convergence by counting the number of empty bins. Our search algorithm assumes arbitrary minima and maxima as the boundaries to be used when binning pseudomaterials. For volumetric methane loading these boundaries are 0 and 350 v/v STP, for volumetric surface area 0 and 4500 m<sup>2</sup>/cm<sup>3</sup>, and for helium void fraction 0 and 1. The area between these boundaries is then subdivided so it is possible to subtract the number of accessed bins by the total number of bins to get the number of empty bins, which can be used as a convergence metric. Figure 8 shows a comparison of empty bin counts for two libraries of 300,000 materials, one generated using random sampling and the other the mutation method with a strength of 20%; different thresholds were used to define occupied bins where either 1, 10, 100, or 250 pseudomaterials had to be present to consider a bin occupied.



Figure 8. Comparison of convergence using empty bin counts between two libraries, one generated by random sampling, indicated by dashed lines, and another by mutation at 20% strength, indicated by solid lines. Different thresholds were used to determine whether a bin was occupied; either 1 (*cyan*), 10 (*magenta*),

100 (green), or 250 (purple) pseudomaterials needed to be present to consider a bin occupied.

Regardless of the threshold used, random sampling was more effective than mutation at exploring the structure-property space. A likely culprit for this is the infinite loop problem, but unlike the statistical outlier problem a more sophisticated solution is required.

## 2.4.2 OEDIPUS: A Software Package for Rapidly Testing New Methods

To provide a platform for testing methods more rapidly, primarily to combat the infinite loop problem, we started studying abstract representations of pseudomaterials: rectangular prisms, or boxes. Where a pseudomaterial has a complex input space: a number density, lattice constants, many atom-site positions, many LJ parameters, many partial charges, a box has only three input values: length, width, and breadth. Just as we use GCMC simulations to calculate properties of interest from a pseudomaterial structure, we can use the inputs from a box to calculate our own properties of interest. We can carefully define these properties of interest such that we understand with certainty the resulting probability distributions and behavior of this idealized structure-property space. Suppose we are generating boxes by randomly selecting values between 0 and 1 for each input parameter: length (x), breadth (y), and width (z). Each variable, x, y, or z is represented by a uniform distribution. If we then define a property of interest,  $\alpha$  as the sum of x and y, the convolution of these two uniform distributions results in the density shown in Figure 9.



Figure 9. Convolution of two uniform distributions.

More importantly we know that all values of  $\alpha$  are between 0 and 2. If we then define a second parameter,  $\beta$ , as  $z^{12}$  we'd expect a probability density with a sharp peak near 0 and limits of 0 and 1. Our hypothetical parameter space is now a box with a width of 2 and height of 1. We can redefine  $\alpha$  as (x + y) / 2 to make our parameter space a unit square as shown in Figure 10.



Figure 10. Parameter space for boxes where  $\alpha = (x + y) / 2$  and  $\beta = z^{12}$ .

Whereas the pseudomaterials parameter space is of an unknown size and shape, we know in advance the exact size and shape of our parameter space for boxes. This allows us to know precisely how well we've sampled the parameter space using our sampling algorithms. Furthermore, because we have insight into the probability density of the parameters of interest we have some intuition as to how the space will be sampled when using a random sampling algorithm. Because the probability densities for both  $\alpha$  and  $\beta$  are non-uniform certain regions are sampled more rigorously than others via random sampling as show in Figure 11.



Figure 11. Histograms of  $\alpha$  (*left*) and  $\beta$  (*middle*) values from a population of 10,000 boxes generated using a random sampling method. A scatterplot (*right*) represents the mapping of the overall 2D parameter space using this method.

 $\beta$  was intentionally defined to make it difficult to sample the entire parameter space via random sampling to mimic challenges in sampling the pseudomaterials space. We can then compare different methods by the number of empty bins remaining after each generation using the random sampling approach as a benchmark. This idealized input/output space captures the essence of the problem while reducing the amount of time for output from weeks to just minutes.

The OEDIPUS software package is free and open source and available at:

https://github.com/akaija/OEDIPUS

## 2.4.3 Adaptive Mutation Algorithms

To address the infinite loop problem, we assume that different pseudomaterials can respond differently despite being mutated with the same mutation strength; that is, the percentage of the pseudomaterial's children which land in the same bin as the parent will vary. One way to address this is to assign mutation strengths to each bin, adjusting based on the percentage of children that land in the parent bin. Each bin is initialized with the same mutation strength value, which may be adjusted later. Starting in the second generation, mutation strengths are calculated for each bin containing parents for children in the previous generation. Mutation strengths are calculated by assessing the children from all parents within a bin and determining the fraction of those children which landed in the same bin as their parents. If less than 10% of the children occupy the same bin as their parents, then the mutation strength is decreased by approximately 10% (divided by 1.1) and if more than 50% of the children occupy the same bin as their parents the mutation strength is increased by 10% (multiplied by 1.1). The objective is to have between 10% and 50% of all children occupying the same bin as their parents. If too few children occupy their parents' bin, we get stuck in an infinite loop, constantly selecting parents from a sparsely populated bin without ever increasing the bin's relative population, causing the same parents to be selected over and over. If too many children occupy their parents' bin, no new regions of the structure-property space will ever be explored. The adaptive mutation strength is an attempt at ensuring that two things occur: new regions of the structure-property space are explored and that the overall sampling of the structure-property space is uniform.

Recall that using the methods described above, a mutation strength is assigned to each bin. Alternatively, mutation strengths can be assigned to each material, though as the number of materials in the library reaches tens of thousands this approach may become computationally expensive. Testing these two methods using the hypothetical box platform we see that bin-specific adaptive mutation strengths outperform material-specific adaptive mutation strengths, as shown in Figure 12.

25



Figure 12. Convergence displayed as percentage of empty bins remaining for four 100+ generation libraries of hypothetical boxes: using random sampling (*red*), using a constant mutation strength of 20% (*blue*), using bin-specific adaptive mutation strengths (*green*), and using material/box-specific adaptive mutation strengths (*magenta*).

Bin-specific mutation strengths even perform slightly better than a constant mutation strength, whereas material-specific mutation strengths do not. Unfortunately, none of these mutation methods perform better than random sampling when comparing convergence as number of empty bins. However, observing histograms of the number of materials in each bin after 100 generations shows that these three mutation methods (using constant, bin-specific adaptive, and material-specific adaptive mutation strengths) sample the structure-property space more uniformly, as shown in Figure 21.



Figure 13. Histograms showing the number of boxes generated in each bin after 100 generations of 100 materials using random sampling (*red*), constant 20% mutation strength (*blue*), bin-specific adaptive mutation strengths (*green*), and material/box-specific adaptive mutation strengths (*magenta*). Bin-count variances are displayed in each subplot.

While the three mutation methods don't perform as well as random sampling when measured by the total area sampled, they are more effective at sampling a smaller area more uniformly. One simple approach to improve performance is to combine one of the mutation methods with random sampling. An implementation of this approach is to randomly choose whether to create a new material via mutating an existing structure or generating one randomly each time a new material is added to the library. We see this "hybrid" approach outperforms the mutation methods and random sampling when compared using empty bin counts, as shown in Figure 14.



Figure 14. Convergence displayed as percentage of empty bins remaining for four 100+ generation libraries of hypothetical boxes: using random sampling (*red*), using a constant mutation strength of 20% (*blue*), using bin-specific adaptive mutation strengths (*green*), using material/box-specific adaptive mutation strengths (*magenta*), and using a hybrid approach where each material is created by randomly choosing between mutating an existing box or generating one randomly (*cyan*).

The hybrid approach is somewhat less promising when comparing how uniformly the space is sampled. While it performs substantially better than random sampling, the hybrid approach samples the structure-property space far less uniformly than any of the mutation methods, as shown in Figure 15.



Figure 15. Histograms showing the number of boxes generated in each bin after 100 generations of 100 materials using random sampling (*red*), constant 20% mutation strength (*blue*), and a hybrid approach where each material is created by randomly choosing between mutating an existing box or generating one randomly (*cyan*). Bin-count variances are displayed in each subplot.

This hybrid approach could be further tuned by weighting the random selection between creating a new material via mutation or creating one randomly. One might expect that a selection more heavily weighted toward mutation would yield a more uniform sampling of the structure-property space.

# 2.4.4 Property Prescreening: Using Machine Learning to Improve Computational Efficiency

Property prescreening augments the random sampling approach by using machine learning models to determine which pseudomaterial structures should be considered for GCMC simulations. This approach adds a few steps to the random sampling process. First the overall random sampling routine is used to create some number of pseudomaterials. Then, prior to generating additional pseudomaterials, a machine learning model is fit to this set of data. We found by testing all of the regressors available in the Scikit Learn Python package that Gradient Boosting Regressors were the most effective for predicting values outside of the training data used to fit the model. To arrive at this conclusion, we built models using the average  $\sigma$ -values, average  $\epsilon$ -values (where we averaged the values for each atom-site across the pseudomaterial unit cell), number density, and unit cell dimensions as inputs with methane loading at 35 bar and helium void fraction as outputs. We then used simulated data from HTSoHM to build training and testing data sets. The training data consisted of all points within an arbitrarily-sized "box", while testing data consisted of all point outside of this box. This accomplished two tasks: (1) it determined which regressors were capable of predicting values "outside" of an arbitrary boundary (as certain regressors, such as those which use decision trees, are not capable of predicting values outside the convex hull formed by the training set) and (2) it determined which of these regressors was the most accurate at predicting values outside of this box. The results of this test are shown in Figure 16.



Figure 16. Scatterplot of testing data (*black*) and predicted results from various regressors with non-negative R-scores. Each regressor is represented by a different color with its R-score displayed in the legend. Training data came from various points within the *red* box, none of the testing data came from within this box, so a perfect regressor would never predict values within it.

To provide a more realistic estimate of the accuracy of a Gradient Boosting Regressor, we then preformed a second test. This time training data came from simulation of a library of 1000 pseudomaterials. A convex hull was then generated from this data and an additional 10,000 pseudomaterials were used to create a testing set. The accuracy of the Gradient Boosting Regressor was determined by its ability to predict whether values within this testing set were inside or outside the convex hull formed by the training data. The results of this test are shown in Figure 17.



Figure 17. Predictions of where 10,000 pseudomaterials would exhibit properties inside or outside of a convex hull formed by a different set of 1,000 pseudomaterials using a Gradient Boosting Regressor.

We observed that the Gradient Boosting Regressor was accurate at predicting whether pseudomaterials would exhibit unobserved properties (that is, properties that fall outside of the convex hull formed by the training set) more than 85% of the time. Assuming this accuracy was sufficient for our purposes, we used this regressor to prescreen pseudomaterials. After generating some number of pseudomaterials using random sampling, we would train a Gradient Boosting Regressor to predict if newly-generated structures exhibited unobserved properties, if so the pseudomaterial structure would be investigated using GCMC simulations (if not, a new structure would be generated in loop that exited when the regressor predicts that the newly-generated psuedomaterial would exhibit unobserved properties). The overall schematic appears in Figure 18, where the property prescreening steps are denoted by dashed lines.



Figure 18. Flowchart for HTSoHM with propert prescreening where the property prescreening steps are denoted by dahsed lines.

Finally, to determine whether it would be worth moving forward with this approach, which we have dubbed PROMETHEUS, we used a library of 300,000 pseudomaterials with data for volumetric methane loadings at 35 bar and helium void fractions to conduct a final test. This test would compare random-sampling to random-sampling with property prescreening. Instead of generating new data, which would incur substantial computational expense, we would use a set of 1,000 pseudomaterials for training data and then use a Gradient Boosting Regressor to iterate over the remaining 299,000 pseudomaterials to predict if the exhibited unobserved properties. If a pseudomaterial was predicted to have unobserved properties, instead of simulating its properties (this was already done when the original 300,000 pseudomaterials library was generated) we would add the previously-simulated data to the training set, retrain the model, and continue iterating over the remaining pseudomaterials. If property prescreening worked, we would be able to observe a comparable region of the structure-property space with far fewer pseudomaterials. However, one question remains: how do we determine if a material has unobserved properties? We attempted two methods: one using a convex hull and another using binning. With the convex hull, a pseudomaterial with properties predicted to be outside of the convex hull would be subject to simulation. With binning, we would bin the structure-property space and a pseudomaterial with properties predicted to access a new, unoccupied bin would be subject to simulation. Results from both PROMETHEUS tests are shown in Figure 19 and Figure 20.



Figure 19. Comparisons of random sampling and propert prescreening (PROMETHEUS method) using number of empty bins as a convergence criteria and a convex hull to assess whether predicted properties were unobserved.



Figure 20. Comparisons of random sampling and propert prescreening (PROMETHEUS method) using number of empty bins as a convergence criteria and binning to assess whether predicted properties were

unobserved.

We observed that in both cases, the PROMETHEUS method would improve computational efficiency by reducing the number pseudomaterials needed to rigorously sample the structure-property space. Between the convex hull and binning approaches, binning proved to be more effective at sampling a large region of the structure-property space. Using binning, we were able to access ~550 bins with ~3000 pseudomaterials, while random-sampling alone access ~1100 bins with 300,000 pseudomaterials. This means the PROMETHEUS method was able to sample ~50% of the same space at 1% of the computational expense. This work represents a promising new direction for pseudomaterials research and has been implemented into the HTSoHM codebase for future testing.

# 3.0 Efficiently Mapping Structure-Property Relationships of Gas Adsorption in Porous Materials: Application to Xe Adsorption

Alec R. Kaija, Christopher E. Wilmer

Faraday Discuss., vol. 201, pgs. 221-232, (2017).

Designing better porous materials for gas storage or separations applications frequently leverages known structure-property relationships. Reliable structure-property relationships, however, only reveal themselves when adsorption data on many porous materials are aggregated and compared. Gathering enough data experimentally is prohibitively time consuming, and even approaches based on large-scale computer simulations face challenges. Brute force computational screening approaches that do not efficiently sample the space of porous materials may be ineffective when the number of possible materials is too large. Here we describe a general and efficient computational method for mapping structure-property spaces of porous materials that can be useful for adsorption related applications. We describe an algorithm that generates random porous "pseudomaterials", for which we calculate structural characteristics (e.g., surface area, pore size and void fraction) and also gas adsorption properties via molecular simulations. Here we chose to focus on void fraction and Xe adsorption at 1 bar, 5 bar, and 10 bar. The algorithm then identifies pseudomaterials with rare combinations of void fraction and Xe adsorption and mutates them to generate new pseudomaterials, thereby selectively adding data only to those parts of the structureproperty map that are the least explored. Use of this method can help guide the design of new porous materials for gas storage and separations applications in the future.

## **3.1 Introduction**

Porous materials are used widely in many applications related to catalysis, separations, [38, <u>41, 44, 45, 68-71</u>] gas storage, [27, 34-36, 72, 73] and chemical sensing, [23, 52, 53, 55] among others. Many of these applications stand to be improved from better porous materials, and so signi cant effort is devoted to searching for them. [28, 29, 46, 74, 75] In particular, enormous effort has been devoted over the past decade to designing better metal-organic frameworks (MOFs), which are porous crystals with extremely high surface areas that are synthesized by the self-assembly of modular chemical building blocks. [20, 49, 76] Due to the modularity of the building blocks used in their synthesis, thousands of new MOFs have been synthesized over the past decade, and there are undoubtedly millions of possible MOFs that have not yet been created.[64] This large design space of possible MOFs makes it an ideal class to search within for better porous materials.[77] However, the availability of a large design space does not itself provide any guidance on how to design a better porous material. Although simple trial-and error synthesis combined with chemical intuition have yielded numerous successes in the field of MOFs, [69, 78] rational design of materials is generally considered preferable. One way to improve on the trial-and-error approach is to search for design rules that can be extracted from observations of structure-property relationships. For example, Bae et al. aggregated experimental data of CO2 adsorption in zeolites and MOFs from over 40 sources in the literature and examined trends with respect to material properties. [45] In that work, some trends appeared robust, such as  $CO_2$  loading at 0.5 bar varying linearly with the isosteric heat of adsorption (up to a point). However, at 2.5 bar, the authors wrote that no trend could be inferred from the data. In a follow up work by us, which considered the same relationships (CO<sub>2</sub> loading vs. isosteric heat of adsorption) at the same conditions, but used molecular simulation data from over 130 000 hypothetical MOFs, [71] we

found the same trends as Bae et al., but also many well-defined relationships that were not visible when only 40 data points were available. There are many recent examples of other researchers using molecular simulations and large datasets (numbering anywhere from 5000 to over 600 000 materials) to obtain structure-property relationships for porous materials related to gas adsorption applications. [35, 44, 61, 64, 66, 79, 80] It is worth noting that these computational studies often do lead to better synthesized materials, usually because the observed structure-property relationships led to design rules that were then followed by experimentalists. For example, a promising MOF for natural gas storage, NU-125, was found almost immediately after the discovery that methane storage at high pressure peaks sharply at void fractions of 0.8.[34] Their demonstrated utility notwithstanding, the large-scale computational studies described above were nevertheless very inefficient with regards to mapping structure-property relationships. This is because many of the materials considered in those large-scale studies had very similar structures, and hence similar adsorption properties. In other words, the structure-property spaces were not explored uniformly; some properties (e.g., small pores) were encountered much more frequently than others (e.g., large pores). Undoubtedly to the frustration of many, it has usually been the most sought a er properties (e.g., high concentrations of adsorbed H2 at 298 K and 100 bar) that have been the most poorly represented in the datasets, dwarfed in number by materials whose properties do not lend them any obvious application value. Here we describe a method whose aim is to uniformly, and thus efficiently, explore structure-property spaces related to gas adsorption. Briefly, the method revolves around the use of crystalline porous "pseudomaterials", which are collections of randomly arranged Lennard-Jones spheres in a periodic unit cell. Despite their randomness within the unit cell, they have long-range order and are expected to behave like porous crystals rather than like amorphous materials. Furthermore, a subset of pseudomaterials have

highly symmetric arrangements of spheres even within their unit cells, thus exactly resembling typical zeolites and MOFs. Using the same computational methodology that is used for zeolites and MOFs, we are able to calculate surface areas, void fractions, and gas adsorption in these pseudomaterials. We then identify pseudomaterials in regions of the structure-property space that are underexplored and mutate them to generate additional (but not identical) child materials. Pseudomaterials in well-explored regions of the structure-property space are ignored. The details are described in the Methods section. We focused in this initial study on Xe adsorption at various pressures (1 bar, 5 bar, and 10 bar) as a function of the void fraction of the porous pseudomaterials. In addition to the relative simplicity of modelling Xe adsorption, it is a gas of industrial importance as it is a component in fluorescent lights and must be removed along with radioactive isotopes of 85Kr from spent nuclear fuels.[81-83] Both of these applications typically employ cryogenic distillation to obtain pure Xe, which is very energy intensive and could potentially be replaced by the use of the right porous adsorbent in process that operates under ambient conditions.[81, 82, 84] A primary focus of this study was to explore how quickly the space of Xe adsorption vs. void fraction is explored using this mutation strategy, and to see whether the structure-property space of porous pseudomaterials resembles that of real materials.

## 3.2 Methodology

An overview schematic of our six step computational method is presented in Figure 21. The details of each step are given in the subsections that follow, but a concise overview is provided here. First, a seed population of pseudomaterials is generated. Second, various structural characteristics (*e.g.*, surface area, pore size and void fraction) and adsorption properties (*e.g.*, Xe loading at 298 K at 5 bar) are calculated for each newly generated pseudomaterial. Third, pseudomaterials with rare structure–property combinations are identified as candidate "parents" that can spawn "child" materials with similar properties. Fourth, parent candidates are tested to ensure that the rareness of their properties is not due to calculation inaccuracies stemming from statistical undersampling. Fifth, a mutation strength parameter is adjusted to ensure that child pseudomaterials are not too similar, nor too different, from their parents (this only affects the efficiency with which the structure–property space is explored). Sixth, parent pseudomaterials are mutated to generate a new population of child pseudomaterials. At that point, we return to step two where the process repeats for as many generations as are needed to explore the structure–property space completely.



Figure 21. Flow chart describing the method in six parts: generating a seed population (A), calculating properties of interest (B), selecting rare pseudomaterials as candidate parents (C), removing anomalous results due to statistical undersampling (D), adjusting mutation strengths (E), and mutating rare parents to create new materials (F).

## **3.2.1** Generating the seed population

Pseudomaterials are generated by randomly positioning Lennard-Jones spheres, which we refer to as pseudoatoms, within a randomly sized unit cell (see Figure 22). The crystal lattice constants (i.e., unit cell dimensions) were bounded between 25.6 and 51.2 Å for each of the crystallographic directions (the lower bound was set as twice the cutoff length used for the interactions between non-bonded atoms: 12.8 Å). The number density of pseudoatoms was bounded between  $1.49 \times 10^{-5}$  and 0.02122 atoms per Å<sup>3</sup>. The density minimum was chosen to ensure that each pseudomaterial would have at least two pseudoatoms, and the maximum corresponds to 10% the number density of iron. Pseudoatom types were defined by Lennard-Jones (LJ) interaction parameters:  $\sigma$ , van der Waals radius, and  $\varepsilon$ , the potential well depth. Values of  $\sigma$  and  $\varepsilon$  were randomly chosen in the range between 1.052 and 6.549 Å for  $\sigma$  and between 1.258 and 513.26K/k<sub>B</sub> for  $\varepsilon$ . Although a pseudomaterial might have hundreds of pseudoatoms, each material had only four pseudoatom types. No two materials shared the same pseudoatom types. This range of LJ parameters was based on the Universal Force Field (UFF),[65] where we allowed  $\sigma$  and  $\varepsilon$  values to be 50% lower or greater, respectively, than the minimum and maximum values present in that force field. In this study, the seed population, and each subsequent child population, contained 100 pseudomaterials.



Figure 22. Orthogonal and perspective views of two randomly generated pseudomaterials, A and B. The black wireframes represent the unit cells. Pseudoatoms are shown as spheres, whose radii and colour indicate  $\sigma$  and  $\varepsilon$  values respectively.

### **3.2.2** Modeling properties of interest

After a generation of pseudomaterials was created, we used grand canonical Monte Carlo (GCMC) simulations to determine xenon adsorption at 298 K and 1 bar, 5 bar, and 10 bar. Pseudomaterials were treated as rigid structures, where the positions of the pseudoatoms were held fixed in space throughout the simulation. The helium void fraction of each pseudomaterial was calculated using a Widom insertion method, [67] with a helium probe ( $\alpha = 2.96$  Å).

#### **3.2.3** Selecting rare materials

To explore new regions of the structure–property space, pseudomaterials with rare structure–property combinations were preferentially selected as candidate "parents" in the process of creating new materials. First, the structure–property space is subdivided into bins. In our case, the Xe-adsorption-void-fraction space was divided into 100 bins (10 bins along the Xe loading axis times 10 bins along the void fraction axis). Void fraction ranged from 0 to 100%, and Xe loading from 0 to either 50 cm<sup>3</sup>Xe per cm<sup>3</sup> framework at 1 bar, 100 cm<sup>3</sup> cm<sup>-3</sup> at 5 bar, or 150 cm<sup>3</sup> cm<sup>-3</sup> at 10 bar. Pseudomaterials were chosen as candidate parents with probability inversely proportional to the number of materials in the same bin.

### 3.2.4 Removing anomalous results

Because the Monte Carlo methods used to model Xe loading and void fraction are stochastic, there is always a finite probability that a candidate parent is selected on the basis of anomalous simulation results (*e.g.*, zero Xe loading at 10 bar in a material with over 90% void

fraction, which contradicts the ideal gas law). Such anomalous results would get the highest weighting based on the selection criteria of rarity, but would have no children with similar properties. To prevent these anomalous results from creating such inefficiencies, the Xe loading and void fraction of each candidate parent were re-simulated five times. If the average value from these retests varied more than one bin-width from the original, the selected pseudomaterial was disqualified as a parent and ignored subsequently.

## 3.2.5 Assigning mutation strength

When a parent is selected due to a low bin count, the expected outcome is that many children (though not all) are generated in the same bin. However, if the mutation strength is too high, most of the children may land far from the parent with respect to their structure–property combinations. This is particularly common near edges or cusps of the structure–property space. Conversely, if the mutation strength is too low, all of the child materials can end up in the same bin as the parent, thus not exploring the structure–property space outwards.

To address this, we apply an adaptive scheme that adjusts the mutation strength every generation for each bin. In this study, each bin had an initial mutation strength of 20% and was increased or decreased as follows. If in any generation, a bin produced 90% of its children in other bins (not including itself) the mutation strength was halved. If more than 50% of its children ended up in the same bin as the parent, the mutation strength was doubled (up to a maximum of 40%). In the in-between case, the mutation strength was not adjusted.

## **3.2.6** Mutating parents, creating new materials

Once the mutation strengths are adjusted for each bin, a new generation of child pseudomaterials is created. Here mutation refers to the process of randomly perturbing each value that defines the parent material's structure (coordinates of each pseudoatom, number density, LJ values of pseudoatom types, and unit cell dimensions). Generating a mutant child can be concisely described as linearly interpolating each defining value between the parent and a completely random pseudomaterial (unrelated to the parent) to a degree dictated by the mutation strength. At 0%, the child is a clone of its parent, and at 100% the child is a completely new randomly generated pseudomaterial.

## 3.2.7 Uniformity metric

To measure progress in exploring the structure–property space, we defined a uniformity metric (U). First each bin-count (c) was normalised:

$$c^* = \frac{c}{c_{max}} \tag{3-1}$$

Where  $c_{\text{max}}$  is the highest bin-count up to the current generation. Then *U* is the variance of the normalised bin-counts for the number of non-zero bins in the dataset (*N*):

$$U = \sum \frac{c^* - \frac{\sum c^*}{N}}{N} \tag{3-2}$$

The lower the value of U, the more uniformly the structure–property space has been explored. A high value of U indicates that some structure–property combinations are over-represented relative to others.

#### **3.3 Results**

In seed populations (Figure 23) in all three runs (1, 5, and 10 bar) more materials were concentrated in the 0.9 to 1 void fraction domain than others, with very few to no materials in the 0 to 0.1 range. In the seed population of the 1 bar run (Figure 24) the most populated bin contained twice as many materials as the next most populated bin. While the density limits could be adjusted to produce a more even distribution of structure–property combinations right from the start, the method should adjust for initial unevenness because we then selectively mutate rare pseudomaterials. Already after one generation we see a more even distribution of structure property-combinations as well as six new bins (see Figure 24a); after fifty generations, the number of accessed bins has nearly doubled and the distribution continues to flatten (Figure 24b).



Figure 23. Scatterplots for seed-populations of 100 pseudomaterials from three separate runs at 1 bar (A), 5 bar (B), and 10 bar (C). The entire parameter-space is plotted as searched in each case, each data point represents a different material. Search limits were set at 50, 100, and 150 cm<sup>3</sup> xenon per cm<sup>3</sup> framework at 1, 5, and 10 bar respectively.


Figure 24. Bar chart of all bin-counts after seed population (black) and addition of first generation (red, A) and first fifty generations (red, B). Bin-counts were normalised with the maximum bin-count.

Because new pseudomaterials are created by mutating existing ones, new bins are accessed by few children while the majority of new pseudomaterials end up in bins that have already been populated. This contributes to the histogram in Figure 24 having a "sloped" shape, where bins accessed in earlier generations tend to have higher bin-counts than bins accessed later. Even after 50 generations, we observed that the most populated bin had nearly twice as many materials as the next most populated bin, and that the 0 to 0.1 void fraction bin was still not accessed (*i.e.*, empty). Because this bin clearly corresponds to a physically feasible region of structure–property space (*i.e.*, that of solid materials), the emptiness of the bin indicated that the method has not yet explored the entire space after 50 generations. However, the significantly flattening of the distribution clearly shows that the method is functioning as intended.

We can observe clear examples of new bins being accessed after the 20<sup>th</sup> and 30<sup>th</sup> generations in the 1 bar run and the 20<sup>th</sup> generation in the 5 bar run (see Figure 25a, Figure 25b) as indicated by a single red data point in its own bin. The method then successfully fills these bins within the next 10 to 20 generations. In the 20<sup>th</sup> through 30<sup>th</sup> generations in the 1 bar run, for example, the process of filling the newly-accessed bins results in a dense cluster of data points with a very slight gap between it and the adjacent bin. The gap is slight because the binning routine does not notice the lack of data points as it occurs in the same bin as the dense cluster. The method does not smooth the distribution of pseudomaterials within bins, only between them. This results in some clustered data points, most clearly visible in the 5 bar data (see Figure 25b).



Figure 25. Scatterplots after 10, 20, 30, 40, and 50 generations (from left to right) of children had been added for runs at 1 bar (A), 5 bar (B), and 10 bar (C). Children added in the last generation are highlighted in red.

The uniformity metric (see Figure 26) gives insight into how new bins are accessed and then filled with materials. Immediately after the seed population was generated there is a slight increase in the uniformity metric, as new bins that are adjacent to sparsely populated regions, are accessed. These new bins are then filled relatively quickly (10 generations or less) as indicated by a decrease in the uniformity metric. The uniformity metric then steadily rose in the 1 and 10 bar runs before starting to decrease (Figure 26), as new bins were accessed and then filled with pseudomaterials. In the 1 bar run, for example, there was a slight decrease followed by an increase in uniformity (or a slight increase followed by a decrease in uniformity metric) from generation 30 to 50 corresponding to the sudden migration of child pseudomaterials into the three bins in the upper Xe loading domain (see Figure 25a, 30 vs. 40 generations). This trend is also seen in the 5 bar run between generations 30 and 50, as well as in the 10 bar run between generations 15 and 30. This process of discovering new bins can be observed as peaks in the uniformity metric, where the downward slope represents the filling of those bins. Figure 26 indicates that the uniformity metric is decreasing but has not yet reached a steady minimum value. We expect that allowing this method to proceed with more generations, beyond the 50 shown here, would eventually yield a more complete mapping of the structure-property space, where the uniformity metric would approach zero and every bin would have an equal number of materials.



Figure 26. The uniformity metric after each generation for each of the 1 (red), 5 (blue), and 10 bar (green) runs.

Because pseudomaterials are random configurations of Lennard-Jones spheres and thus are not synthesizable, it is reasonable to ask how representative they are of real materials. The general shape of the structure–property maps of pseudomaterials we observed here (as shown in Figure 25) closely resembles what was seen for Xe adsorption in a high throughput screening study on a database of MOFs.[44] In a similar study on MOFs, but looking at methane adsorption at 35 bar, there were observations of a sharp peak in loading at 0.8 void fraction, not unlike the sharp peaks we observe here for Xe loading in pseudomaterials. These similarities in structure–property relationships between MOFs and pseudomaterials are encouraging and support the possibility that the latter can have utility in helping understand real materials.

#### **3.4 Conclusions**

Despite significant strides in discovering better porous materials for a wide range of applications, particularly in the case of MOFs, it remains challenging to find the right design rules, especially as they vary from application to application. To extract design rules for a new application, one requires large datasets, for which experimental screening is too costly and time consuming, and for which high throughput computational screening can be very inefficient if the parameter space is non-uniformly sampled.

We have developed a method for efficiently exploring structure–property maps that relate to gas adsorption in sorbents through the creation of multiple generations of porous pseudomaterials. In each generation we simulated Xe loading and void fraction using classical GCMC techniques. Then we selected pseudomaterials with rare structure–property combinations and mutated them to generate child materials in the lesser explored areas of the structure–property space. We considered 50 generations for Xe loading at 1 bar, 5 bar, and 10 bar. We observed the method systematically accessing new bins then filling them with new pseudomaterials, which gradually makes the distribution of properties in the dataset more uniform. To measure this uniformity, we introduced a quantitative metric for tracking this process called the uniformity metric.

Encouragingly, structure–property maps from simulations of real materials (MOFs) closely resembled those found for our pseudomaterials. The method presented here represents a computationally efficient means by which to rapidly map a structure–property space that can then be used to extract potentially useful design rules. We hope this methodology may someday aid in the design of better porous adsorbents.

60

#### 4.0 High-Pressure Methane Adsorption in Porous Lennard-Jones Crystals

Alec R. Kaija, Christopher E. Wilmer

J. Phys. Chem. Lett., vol. 9, pgs. 4275-4281, (2018)

Decades of research have yet to yield porous adsorbents that meet the US Department of Energy's methane storage targets. To better understand why this might be, we calculated high pressure methane adsorption in 600,000 randomly generated porous crystals, or "pseudomaterials," using atomistic grand canonical Monte Carlo (GCMC) simulations. These pseudomaterials are periodic configurations of Lennard-Jones (LJ) spheres whose coordinates in space, along with corresponding well depths ( $\varepsilon$ ) and radii ( $\sigma$ ), are all chosen at random. GCMC simulations were performed for pressures of 35 and 65 bar at a temperature of 298 K. Methane adsorption was compared for all materials against a range of other properties: average  $\varepsilon$  and  $\sigma$  value, number density, helium void fraction, and volumetric surface area. The results reveal structure-property relationships that resemble those observed previously for MOFs and other porous materials. A common characteristic among top performers was a combination of atoms with small radii and strong binding (i.e., low  $\sigma$  and high  $\varepsilon$  values). We find that our computational methodology can be useful for discovering structure-property relationships related to gas adsorption without requiring detailed structural data of real (or even realistic) materials.

#### 4.1 Introduction

The discovery of tunable, high-surface area, porous materials known as metal-organic frameworks (MOFs)[21] has sparked considerable interest over the past two decades in developing adsorbents for various industrial applications[22, 23, 25-29, 85] including gas storage[19, 30-32, 34-36, 73] and separations,[38-45] catalysis,[46-51] and sensing.[52-55] In particular, the application of high pressure methane storage has driven significant exploratory efforts in the MOF field. Although tens of thousands of MOFs have been synthesized, there are still potentially many millions more that remain undiscovered.[49, 55, 64, 75] In addition to experimental efforts, large-scale computational screening has been used extensively on hundreds of thousands of real and hypothetical MOFs and related materials to find promising targets for methane storage.[66, 80, 86-88] In addition to identifying useful targets for synthesis, these large-scale screening studies have contributed significant insights via the discovery of clear structure-property relationships.[64, 71, 89, 90] Many researchers have utilized large libraries of hypothetical porous materials (5,000 to 600,000 materials) to observe structure-property relationships related to gas adsorption.[35, 61, 64, 66, 79, 80]

Despite these efforts, MOFs discovered to date (including the hypothetical ones) have not met the high methane storage targets set by the US Department of Energy (DOE): 315 cm<sup>3</sup>(STP)/cm<sup>3</sup> at 35 or 65 bar and ambient temperature.[91-93] However, given the vast space of *possible* MOFs, it is not clear whether a performance ceiling has been reached or whether higher performing MOFs wait to be discovered. A higher performing porous material might also be discovered in a different material class altogether (e.g., zeolites).

In addition to libraries of MOFs, there are libraries of hypothetical zeolites,[<u>60</u>, <u>94</u>] porous polymer networks (PPNs),[<u>87</u>] and zeolitic imidazolate frameworks (ZIFs).[<u>95</u>] To date, virtually

all attempts at creating large libraries of hypothetical porous materials have relied on modular chemical building blocks and design rules inherent to a particular material class. But what if new materials exist which do not obey previously observed design rules, or cannot be constructed from previously-studied building blocks? Using these design rules alone to create a comprehensive library of porous materials might result in holes where new materials might exist. These holes might limit our understanding of the full range of structure-property relationships for porous materials.

In this study, in an attempt to generalize structure-property observations of methane adsorption across disparate material classes, we have generated and studied porous "pseudomaterials": periodic configurations of Lennard-Jones (LJ) spheres meant to represent arbitrary porous crystals. This pseudomaterials approach makes it possible to sample regions of the structure-property space that may have been missed by previous studies (at the expense of potentially sampling structure-property combinations that are not physically realizable).[96]

Briefly, our approach involved generating a library of 600,000 porous crystals, represented by configurations of LJ spheres within a unit cell. Six properties were then evaluated for each pseudomaterial: methane loading (at both 35 and 65 bar), void fraction, volumetric surface area, average  $\varepsilon$  value, average  $\sigma$  value, and number density. An earlier report on Xe and Kr adsorption in pseudomaterials, which was more preliminary in nature, used a much smaller library of materials and did not consider as many material properties.[96] We then generated 2D-projections of the property space, where the data were grouped into bins that were colored by some third property; the result is a heatmap showing various structure-property relationships in three dimensions.

#### 4.2 Methodology

## 4.2.1 Lennard-Jones potential

Our pseudomaterials are constructed using LJ spheres, meant to represent different chemical species or moieties, to create structures for use in grand canonical Monte Carlo (GCMC) simulations of methane physisorption. The LJ potential is commonly used in computational studies of physisorption in porous materials and has been experimentally validated numerous times.[64, 71, 79, 97] There are well-documented limitations of the LJ potential in the context of adsorption, such as chemisorption or physisorption where the binding is very strong (as is the case for certain gases and open metal sites). It may stand that even higher methane capacities may be attainable in systems that rely on chemisorption or take advantage of other physical phenomena not captured by LJ interactions.



Figure 27. Renderings of a synthesizable MOF, NU-125 (A), and an algorithmically generated configuration of LJ spheres, or pseudomaterial (B).

#### 4.2.2 Generating pseudomaterials

Pseudomaterials were generated by randomly positioning LJ spheres, which we refer to as pseudoatoms, within a randomly-sized unit cell. The unit cell dimensions were bounded between 25.6 and 51.2 Å in each of the crystallographic directions (with the lower bound twice the cutoff length used for gas-gas and gas-crystal interactions: 12.8 Å). The number of LJ spheres within a unit cell was bounded between 1.49 \* 10<sup>-5</sup> and 0.02122 atoms / Å<sup>3</sup>, where the lower boundary ensures that each unit cell contains at least two atom sites and the upper boundary corresponds to 10% the number density of iron. Pseudoatom types were defined by their LJ parameters:  $\sigma$ , the van der Waals radius, and  $\varepsilon$ , the potential well depth. Values for  $\sigma$  were bounded between 1.052 and 6.549 Å and values for ε between 1.258 and 513.264 K. This range of LJ parameters was based on the Universal Force Field (UFF), [65] where  $\sigma/\epsilon$ -values were allowed to be 50% lower or greater, respectively, than the minimum and maximum values present in the UFF. While a particular pseudomaterial's unit cell may contain hundreds of pseudoatoms depending upon the randomly selected number density, the number of pseudoatom types was limited to four in all pseudomaterials. In this study we generated a library of 600,000 pseudomaterials: two independent sets of 300,000 materials for each methane loading pressure.

# Table 2. Pseudomaterial generation steps and related parameter ranges.

	Pseudomaterial generation step	Parameter range
1	Select lattice constants	25.6 – 51.2 Å
2	Select number density	$1.49 \times 10^{5} - 0.02122$ atoms/Å <sup>3</sup>
3	Position pseudoatom sites	N/A
4	Select four sets of LJ parameters (one for each pseudoatom type)	σ: 1.052 – 6.549 Å ε: 1.258 – 513.264 K
5	Assign each pseudoatom site to one of the four pseudoatom types	N/A

## 4.2.3 Simulating properties

After each pseudomaterial was created, grand canonical Monte Carlo (GCMC) simulations were used to calculate methane adsorption at 298 K and either 35 bar or 65 bar. These pressures were chosen primarily because the community has converged on these two pressures as benchmarks for high pressure methane storage. It may interest the reader to know that 35 bar is the typical pressure of US interstate natural gas pipelines, and 65 bar is the upper limit achievable with inexpensive two-stage compressors.[92] Pseudomaterials were treated as rigid structures, where pseudoatom site positions were held constant throughout the simulation. Void fractions were calculated using a Widom insertion method[67] using a helium probe ( $\sigma = 2.96$  Å). Volumetric surface areas were calculated in a Monte Carlo search, rolling a nitrogen probe ( $\sigma =$ 3.31 Å) over the surface of the unit cell. All of these properties were calculated using a simulation software package for adsorption in nanoporous materials called RASPA.[67]

#### 4.3 Results and Discussions

As expected, the more pseudomaterials we generated, the larger the volume of the structure-property space we sampled (see Figure 28). However, we also expected that at a certain size the library would be sufficiently large such that no new features would be observed with the addition of new pseudomaterials. Figure 28 demonstrates that, when generated at random using our approach discussed in the Methodology section, several hundred thousand pseudomaterials were needed before the addition of new materials became redundant. It was not until tens of thousands of materials had been sampled (Figure 28c) that the upper limits of methane capacity

were observed. Also, certain combinations of properties were much more likely than others. For example, we found that the randomly generated pseudomaterials most commonly had void fractions between 0.5 and 0.9, as can be seen most clearly in Figure 28d.



Figure 28. 2D histograms for samples of 200 (A); 2,000 (B); 20,000 (C); and 200,000 (D) pseudomaterials from a library of 300,000 pseudomaterials. Plots show projections of the structure-property space in methane loading (at 35 bar) with respect to helium void fraction and are colored by the number of pseudomaterials in each of 40×40 equally-sized bins.

Once our library was sufficiently large, we were able to observe the distribution of materials across various 2D projections (see Figure 29). We observed that the highest methane loadings (regardless of pressure) occurred at void fractions between 0.7 and 0.9 and at volumetric surface areas exceeding 3,150 m<sup>2</sup>/cm<sup>3</sup>, which is similar to what has been reported previously by others in the case of MOFs.[64] We found that our pseudomaterials most commonly had volumetric surface areas between 1,350 and 3,600 m<sup>2</sup>/cm<sup>3</sup>, as shown in Figure 29c and Figure 29f (note that Figure 29c and Figure 29f are nearly identical, as expected, since the properties displayed do not depend on pressure). We also found the highest surface areas occurred in a void fraction range of 0.7 to 0.95. Not surprisingly, a larger distribution of pseudomaterials occupied the high methane loading domain in the 65-bar library. Most notably, we observed the highest volumetric surface area pseudomaterials associated with the highest methane loadings.



Figure 29. 2D histograms for the full library of 300,000 pseudomaterials. Plots show different projections of the high-pressure methane storage structure-property space colored by number of pseudomaterials in each of the 40×40 equally-sized bins. Projections shown here: methane capacity with respect to void fraction (A, D), methane capacity with respect to surface area (B, E), and surface area with respect to void fraction (C, F). The top row of plots corresponds to an operating pressure (for methane adsorption simulations) of 35 bar (A,

B, C) with results at 65 bar below (D, E, F).

Two important parameters to explore in our LJ-based model are, of course, the influence of the  $\varepsilon$  and  $\sigma$  values on methane loading. Since each pseudomaterial contained a range of both, we used average values in plotting structure-property relationships. The average  $\varepsilon$ -value in a pseudomaterial provided a measure of the availability of strong (or weak) binding sites in a pseudomaterial. In Figure 30, we show the relationship between this average  $\varepsilon$ -value and void fraction, surface area, and methane loading at both pressures. Figure 30a and Figure 30c show that, at void fraction values above ~0.3, the presence of strong binding sites (here represented by high average  $\varepsilon$ -values) was necessary to maximize methane loading. Also, median average  $\varepsilon$ -values typically corresponded to median methane loading; however, sometimes high surface area pseudomaterials were able to store more methane than ones with higher epsilon values but lower surface areas (see Figure 30b and Figure 30d). At 65 bar, we see higher methane loadings out of lower average  $\varepsilon$ -value pseudomaterials; this shows, as one might expect, that increasing the pressure decreases the need for strong binding sites given the same porosity.



Figure 30. 2D projections of the high-pressure methane storage structure-property space colored by average ε-value for all pseudomaterials within each 40×40 bin. Projections shown here: methane capacity with respect to void fraction (A, C) and methane capacity with respect to surface area (B, D). The top row of plots corresponds to an operating pressure (for methane adsorption simulations) of 35 bar (A, B) with results at 65 bar below (C, D).

We also looked at similar relationships considering the average  $\sigma$ -values (see Figure 31). Interestingly, we observed that both the highest methane loadings and highest surface areas occurred when the average atom/moiety sizes (represented here by average  $\sigma$ -values) were the smallest. Conversely, large atom sites/moieties corresponded with below average surface areas and methane loadings. The very highest methane loadings were not achievable above a certain average  $\sigma$ -value; however, below this domain, average  $\sigma$ -values were constant at constant surface areas. These low to medium-high methane loadings were much more strongly influenced by the availability of strong binding sites than by the average pseudoatom size.



Figure 31. 2D projections of the high-pressure methane storage structure-property space colored by average σ-value for all pseudomaterials within each 40×40 bin. Projections shown here: methane capacity with respect to void fraction (A, C) and methane capacity with respect to surface area (B, D). The top row of plots corresponds to an operating pressure (for methane adsorption simulations) of 35 bar (A, B) with results at 65 bar below (C, D).

Finally, we considered the effect of number density on methane loadings. First, we divided our two 300,000 material libraries into low, medium, and high number density groups. Then we colored by methane loading and plotted average  $\sigma$ -value with respect to average  $\varepsilon$ -values (see Figure 32). We found that the highest methane loadings were observed among medium-to-high number density pseudomaterials. Lower average  $\sigma$ -values were observed in the high number density materials. The highest methane loadings were observed in those pseudomaterials with higher-than-average  $\varepsilon$ -values and lower-than-average  $\sigma$ -values. With high number densities, methane loading clearly depended on both  $\sigma$  and  $\varepsilon$  (see Figure 32d and Figure 32g). However, at lower number densities, methane loading appeared independent of  $\sigma$  (see Figure 32a and Figure 32e).



Figure 32. 2D projections of the high-pressure methane storage structure-property space in average σ-value with respect to average ε-value colored by methane capacity. The plots represent subsets of the 300,000 pseudomaterial library containing low (A, D), medium (B, E), and high (C, F) number density pseudomaterials. The top row of plots corresponds to an operating pressure (for methane adsorption simulations) of 35 bar (A, B, C) with results at 65 bar below (D, E, F).

Figure 33 shows pseudomaterial structures from disparate regions of the structure-property space. In low methane loading pseudomaterials (I, III, and V), we see fewer strong binding sites than in higher methane loading pseudomaterials (II and IV). Pseudomaterial IV had a lower average  $\sigma$ -value than pseudomaterial II, which contributed to its higher methane loading. Pseudomaterial IV is substantially less porous than pseudomaterial V, which also lacks a dense network of strong binging sites. From examining these structures, we continue to see that strong binding sites (high average  $\varepsilon$ -values), as well as relatively low  $\sigma$ -values and an appropriately high level of porosity, contribute to the best methane adsorbents.



Figure 33. Renderings of pseudomaterials from different regions of the high-pressure methane storage structure-property space. The plot (top left) shows where each pseudomaterial is located with respect to void fraction and methane loading: (I) low void fraction and low methane loading, (II) medium low void fraction and methane loading, (IV) medium high void

fraction and high methane loading, and (V) high void fraction and low methane loading.

#### **4.4 Conclusion**

Here we used a novel approach to examine structure-property relationships of physisorption in porous materials, namely by generating and screening libraries of "pseudomaterials," periodic configurations of Lennard-Jones (LJ) spheres, each representing a particular chemical species or moiety. The use of these more abstract structures allowed us to flexibly explore the space of porous materials and avoid potential limitations and challenges in generating structures using chemical building blocks and design rules.

We generated two libraries of 300,000 pseudomaterials and for each material simulated methane capacity at 35 and 65 bar. We then compared the methane capacities for each material with their helium void fraction, volumetric surface area, average  $\sigma$ -values, average  $\epsilon$ -values, and number density. We found the randomly generated pseudomaterials most commonly had void fractions between 0.5 and 0.9 and volumetric surface areas between 1,350 and 3,600 m<sup>2</sup>/cm<sup>3</sup>. The highest methane loadings occurred at void fractions between 0.5 and 0.9 and at surface areas exceeding 3,150 m<sup>2</sup>/cm<sup>3</sup>. Strong binding sites were necessary for higher methane loadings as were smaller atoms/moieties and medium-to-high number densities.

Finally, we looked at several individual structures from different regions of our structureproperty space. We believe this pseudomaterials approach will be useful information for efficiently creating structure-property maps for gas adsorption that experimental researchers who seek to design new adsorbents can use.

81

# 5.0 Carbon Capture Structure-Property Relationships in Lennard-Jones + Coulomb Porous Crystals

Alec R. Kaija, Christopher E. Wilmer

Submitted to J. Phys. Chem. C

Global economic growth is correlated with increased energy demand and – because of our reliance on fossil fuels – increased greenhouse gas emissions. Carbon capture and sequestration (CCS) is therefore not only a potential solution for environmental concerns but economic ones as well. Here we focus on vacuum swing adsorption (VSA) as a process for post-combustion CO<sub>2</sub> capture. VSA relies on porous adsorbents with high CO<sub>2</sub>/N<sub>2</sub> selectivity and CO<sub>2</sub> working capacity. The discovery of such adsorbents, in turn, can be accelerated by developing a better understanding the structure-property relationships for VSA. Here, using a library of 20,000 randomly generated porous crystals that interact with adsorbed CO<sub>2</sub> and N<sub>2</sub> molecules via a Lennard-Jones + Coulomb interaction, we have evaluated their CO<sub>2</sub>/N<sub>2</sub> selectivity, CO<sub>2</sub> working capacity, and other VSA performance metrics, in relation to several structural parameters: atom-site number density, average well depths and van der Waals radii, helium void fraction, and volumetric surface area. The resulting data exhibit sharply defined structure-property relationships not previously observed across different classes of porous materials (classes which we have abstracted away in our generalized approach). The best performing sorbents had small atom sites/moieties and a limited number of strong binding sites. Performance was dictated more by Coulombic interactions than by van der Waals interactions. It may be worthwhile to sacrifice CO<sub>2</sub> working capacity and CO<sub>2</sub> selectivity to ensure high regenerability. These reported structure-property relationships may

provide valuable insight into the structural characteristics necessary to design better adsorbents of VSA CCS processes.

### **5.1 Introduction**

Because of the dependence on fossil fuels to meet energy demands, global economic growth is strongly correlated with increased CO<sub>2</sub> emissions. In 2017, the energy demand growth rate was more than double the previous year; 72% of this rise was met by fossil fuels.[98] Fossil fuels currently account for approximately 85% of energy produced worldwide.[99] This dependence on fossil fuels has negative impacts on the environment which are already becoming apparent. Atmospheric CO<sub>2</sub> levels have risen from 278 ppm at the start of the industrial revolution to over 400 ppm today; 75% of this increase has taken place over the past fifty years.[100] This rise in greenhouse gas emissions[101] and increased demand for natural gas[98, 102] have sparked considerable interest in separating CO<sub>2</sub> from flue gases. Power plants are one of the major sources of greenhouse gas emissions; a 500 MW coal-fired power plant, for example, can produce 3 million tons of CO<sub>2</sub> annually.[103] CCS from these large point sources is crucial to reducing CO<sub>2</sub> emissions.

Various methods have been proposed for separating CO<sub>2</sub> from flue gases: absorption, adsorption, gas separation membranes, and cryogenic distillation.[104] Adsorption processes are promising because they have lower energy requirements than absorption-based processes.[105] Solid sorbents can be used over a wider range of temperatures than liquid sorbents, yield less waste during cycling, and are more environmentally-friendly at end of life.[106] Adsorption processes using solid sorbents fall into two categories: pressure swing adsorption (PSA) and vacuum swing

adsorption (VSA). Both PSA and VSA are technologies considered feasible for industrial scale CCS,[107] where in PSA the adsorption step is done at elevated pressure and in VSA adsorption is performed at or below atmospheric pressure.[108]

In VSA processes sorbent selection precedes process design. Solid sorbents can be classified according to their sorption and desorption temperatures: (1) low-temperature ( $< 200 \,^{\circ}$ C), (2) intermediate-temperature (200 - 400 °C), and (3) high-temperature (>400 °C).[106] Lowtemperature sorbents include amine impregnated adsorbents (e.g. PEI/MCM-41),[109] polymeric amines (e.g. ion-exchange resins (IER) with amine functionality),[110, 111] carbon-based adsorbents (e.g. carbon aerogels),[112] graphite/graphene-based adsorbents (e.g. graphene/chitosan),[113] zeolites (e.g. ZSM),[114, 115] metal-organic frameworks (MOFs) (e.g. alkylamine tethered MIL-101),[81, 116-119] silica-based adsorbents (e.g. 3-(2-aminoethylamino)propyl-dimethoxymethylsilane (APMS)),[120] porous polymers (e.g. PIMs),[121-125], and more. Intermediate-temperature sorbents include LDH-base sorbents (e.g. nano-sized spherical Mg<sub>3</sub>Al- $CO_3$  LDHs)[126] and MgO-based sorbents (e.g.  $Cs_2CO_3$ -doped MgO).[127] High-temperature sorbents include CaO-based sorbents (e.g. CaO-Ca<sub>12</sub>Al<sub>14</sub>O<sub>33</sub>)[128] and alkali silicate-based sorbents (e.g. Li4SiO4).[129] This work focuses on low-temperature sorbents, by simulating  $CO_2$ adsorption at ambient temperatures. Instead of evaluating materials in a full-scale VSA process, approximate quantitative performance metrics can be calculated from the results of simulated adsorption of pure gases (instead of mixtures). Bae and Snurr have established four "adsorption evaluation criteria" for evaluating the effectiveness of sorbents using pure CO<sub>2</sub> and N<sub>2</sub> adsorption measurements. [45] See Table 3; here, N is the component absolute volumetric loading, q is the component absolute loading, and p is the partial pressure. The subscripts 1 and 2 refer to  $CO_2$  and  $N_{2}$ , respectively. The subscripts *ads* and *des* refer to adsorption and desorption conditions, respectively.

Table 3. Adsorption evaluation criteria used by Bae and Snurr to assess sorbents for CCS

Working CO <sub>2</sub> capacity (v/v STP), $N_1^{ads} - N_1^{des}$	$\Delta N_1$
<b>Regenerability</b> (%), $\Delta N_1 / N_1^{ads} \times 100$	R
Selectivity under adsorption conditions, $(q_1^{ads}/q_2^{ads})/(p_1^{ads}/p_2^{ads})$	$\alpha_{12}^{ads}$
Sorbent selection parameter, $(\alpha_{12}^{ads})^2 / \alpha_{12}^{des} \times (\Delta N_1 / \Delta N_2)$	<i>S</i>

These criteria were established because prior work often only considered CO<sub>2</sub> uptake when evaluating a sorbent's potential for use in CCS processes. CO<sub>2</sub> working capacity is defined as the difference between the amounts of adsorbed gas between the adsorption pressure and desorption pressure. The adsorption selectivity is defined as the ratio of equilibrium uptakes in the pore to the ratio of the molar fractions in the bulk phase. CO<sub>2</sub> working capacity and selectivity dictate the amount of gas that can be recovered in single pass, or cycle, through an adsorption column. For environmental and economic feasibility, however, a sorbent must be capable of operating for many cycles. Regnerability – defined as the ratio of the working capacity and amount adsorbed at the adsorption pressure – is a measure of a sorbent's ability to be used for multiple cycles. A sorbent selection parameter was proposed by Rege and Yang and combines the selectivity with the ratio of the working capacities of the two components, as opposed to considering the working capacity of the more strongly adsorbed component alone.[130] Bae and Snurr used these criteria to assess over forty sorbents, identifying the most promising materials for further investigation. In this work

the operating conditions considered represent VSA separation of CO<sub>2</sub> from flue gas assuming a mixture composition of CO<sub>2</sub>-N<sub>2</sub> = 10 : 90, an adsorption pressure of 1bar, and a desorption pressure of 0.1 bar. This work complements past large-scale computational screening efforts focused on CCS.[68, 71, 95, 131, 132] Lin *et al.* and Haldoupis *et al.* computationally screened hundreds of thousands of zeolites/zeolitic imidazolate frameworks (ZIFs) and ~500 MOFs, respectively, for their ability to separate CO<sub>2</sub> from N<sub>2</sub>, but did not describe correlations of performance with the structural features considered in this work (e.g. void fraction, surface area) or other adsorption properties (e.g. working CO<sub>2</sub> capacity, CO<sub>2</sub> selectivity).[68, 95, 132] Wilmer *et al.* expanded the scope of Bae and Snurr's survey, simulating pure component CO<sub>2</sub> and N<sub>2</sub> adsorption data required to calculate the four adsorption evaluation criteria (as shown in Table 3) in over 130,000 hypothetical MOFs for VSA under the same operating conditions considered in this work.[71] Unlike these past studies, this work takes an abstract approach towards sampling the structure-property space for VSA sorbents such that it is inclusive of various different classes of porous materials.

In this work, to generalize structure-property observations of VSA sorbents across disparate material classes, we have generated a library of porous "pseudomaterials": periodic configurations of spheres, each of which being defined by two Lennard-Jones (LJ) interaction parameters and a partial charge. This pseudomaterial approach samples regions of the structure-property space that may have been missed in prior works that focused on specific material classes (at the expense of potentially sampling structure-property combinations that are non-physical).[96, 133] We generated a library of 20,000 pseudomaterials, simulating pure component CO<sub>2</sub> and N<sub>2</sub> adsorption at the operating conditions previously described to simulate a VSA process for CCS, as well as calculating void fraction and volumetric surface area. We used this data to consider the

four adsorption evaluation criteria (see Table 3) as dimensions to the property space. We then generated 2D-projections of this space, where data were grouped into bins that were colored by a third property, resulting in a heatmap representing various structure-property relationships in three dimensions.

# **5.2 Methods**

# 5.2.1 Generating pseudomaterials with partial charges

Pseudomaterials were generated by randomly positioning LJ spheres, which we refer to as pseudoatoms, within a randomly-sized unit cell according to previously established methods.[96, 133] Past work using pseudomaterials has considered the adsorption of relatively non-polar gases (Xe, Kr, and CH<sub>4</sub>). In order to effectively model the Coulombic interactions associated with the adsorption of CO<sub>2</sub> it was necessary to assign partial charges to pseudoatoms within the pseudomaterial structure. After an uncharged pseudomaterial was generated, we iterated over every pseudoatom and added a partial charge to one while subtracting an equal amount from another randomly selected pseudoatom. Finally, the partial charges were scaled to ensure that all resulting partial charges were bounded between -1 and 1. In this work a library 20,000 pseudomaterials was generated, a number selected for being substantially large enough to observe structure-property relationships without incurring excessive computational expense.

# 5.2.2 Lennard-Jones + Coulomb potential

Here we define pseudomaterials as hypothetical porous materials composed of random configurations of spheres (see Figure 34b), where each sphere is defined by a coordinate, a pair of LJ parameters,  $\varepsilon$  (the LJ well depth) and  $\sigma$  (the LJ diameter), and a partial charge, q, which define its interaction energy with atoms of adsorbing gases as follows:

$$V_{ij} = 4\varepsilon_{ij} \left( \left( \frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left( \frac{\sigma_{ij}}{r_{ij}} \right)^6 \right) + \frac{q_i q_j}{4\pi\epsilon_0 r_{ij}}$$
(5 - 1)

In Equation (5-1), the subscripts *i* and *j* refer to the *i*<sup>th</sup> pseudomaterial sphere interacting with the *j*<sup>th</sup> atom of a gaseous adsorbate,  $r_{ij}$  is the distance between *i* and *j*, *q* is the partial charge, and the parameters  $\sigma_{ij}$  and  $\varepsilon_{ij}$  are obtained using Lorenz-Bertholot mixing rules:

$$\sigma_{ij} = \frac{\sigma_i + \sigma_j}{2} \tag{5-2}$$

$$\varepsilon_{ij} = \sqrt{\varepsilon_i \varepsilon_j} \tag{5-3}$$

For  $N_2$  we used a model with three partial charges: a positive charge at the center of the linear molecule (N\_com) for which no LJ interactions were assigned and two negative charges on the two nitrogen (N\_n2) pseudoatoms. For CO<sub>2</sub> we assigned positive charges to carbon

pseudoatoms and a negative charge to the oxygen pseudoatom. See Table 4 for partial charges used in this work.

Atom type	σ (Å)	ε/k <sub>B</sub> (K)	q (C/particle)
pseudoatoms	1.052 - 6.549	1,258 - 513.264	-1 - (+)1
C_co2	2.80	27.0	+0.7
O_co2	3.05	79.0	-0.35
N_n2	3.31	36.0	-0.4048
N_com	-	-	+0.8096

Table 4. LJ Parameters and partial charges for framework pseudoatoms and adsorbate molecules

This LJ + Coulomb potential is commonly used to model physisorption in porous media and has been extensively validated experimentally.[64, 71, 79, 97] However, the LJ potential has well-documented limitations with regard to its inability to capture the behavior of chemisorption or physisorption when binding is very strong.



Figure 34. Renderings of a MOF, NU-125 and a randomly generated configuration of LJ spheres, which we refer to here as a pseudomaterial.
## 5.2.3 Simulating properties

After each pseudomaterial was created, grand canonical Monte Carlo (GCMC) simulations were used to simulate pure component physisorption. These pure component simulations are representative of a mixture of composition  $CO_2-N_2 = 10$ : 90, with an adsorption pressure of 1 bar and a desorption pressure of 0.1 bar. We simulated  $CO_2/N_2$  adsorption at 298K to make straightforward comparisons to experimental data in literature, while flue gas at point of emission is typically closer to 310K. Pseudomaterials were treated as rigid structures, where pseudoatom site positions were held constant throughout the simulation. Void fractions were calculated using a Widom insertion method<sup>46</sup> using a helium probe ( $\sigma = 2.96$  Å). Volumetric surface areas were calculated in a Monte Carlo search, rolling a nitrogen probe ( $\sigma = 3.31$  Å) over the surface of the unit cell. All of these properties were calculated using a simulation software package for adsorption in nanoporous materials called RASPA.[67]

#### **5.3 Results**

With our large dataset, we were able to observe sharply defined structure-property relationships between the four adsorption evaluation criteria (see Table 3) and other material properties such as surface area, void fraction, average  $\sigma/\epsilon$ -values, atom-site number density, and heats of adsorption (separated by their Coulombic and van der Waals contributions). First, we evaluated the relationships between the four adsorption evaluation criteria and the average van der Waals radii (also referred to as average  $\sigma$ -values), which are calculated by averaging the  $\sigma$ -values for all pseudoatoms within the pseudomaterial unit cell. We observed that low-to-average  $\sigma$ -values

were correlated with the highest CO<sub>2</sub> selectivities (see Figure 35a), that low  $\sigma$ -values were correlated with the highest CO<sub>2</sub> working capacities (see Figure 35b), that outside of the very lowest  $\sigma$ -value domain (where regenerability tapered off) regenerability was independent of  $\sigma$ -values (see Figure 35c), and that average  $\sigma$ -values were correlated with the highest sorbent selection parameters (see Figure 35d). Maximizing surface area optimizes the CO<sub>2</sub> working capacity (see Figure 35b) and approaches the highest CO<sub>2</sub> selectivity (see Figure 35a) and sorbent selection parameter (see Figure 35d), but leads to poor regenerability (see Figure 35c).



Figure 35. 2D projections of the CCS sturcture-property space depicting the relationship between average van der Waals radii of interaction (σ-value) and the four adsorption evaluation criteria: CO<sub>2</sub> selectivity (A), CO<sub>2</sub> working capacity (B), regenerability (C), and sorbent selection parameters (D), all colored by volumetric

surface area.

Next, we evaluated the relationships between the four adsorption evaluation criteria and the average potential well depths (also referred to as average  $\varepsilon$ -values), which are calculated by averaging the  $\varepsilon$ -values for all pseudoatoms within the pseudomaterial unit cell. We observed in this work that CO<sub>2</sub> working capacity was relatively independent of average  $\varepsilon$ -values except in the absolute lowest average  $\varepsilon$ -value domain (see Figure 36b). Regenerabilities were also relatively independent of average  $\varepsilon$ -values, except that the lower boundary was higher for lower  $\varepsilon$ -values than in higher average  $\varepsilon$ -value domains (see Figure 36c). The lowest average  $\varepsilon$ -values were correlated with the highest CO<sub>2</sub> selectivities (see Figure 36a) as well as the highest sorbent selection parameters (see Figure 36d). Once again, it is apparent that maximizing surface area improves CO<sub>2</sub> working capacity (see Figure 36b) and to a lesser extent CO<sub>2</sub> selectivity (see Figure 36a) and sorbent selection parameter (see Figure 36d), but negatively impacts regenerability (see Figure 36c).



Figure 36. 2D projections of the CCS sturcture-property space depicting the relationship between average potential well depths (ε-value) and the four adsorption evaluation criteria: CO<sub>2</sub> selectivity (A), CO<sub>2</sub> working capacity (B), regenerability (C), and sorbent selection parameters (D), all colored by volumetric surface area.

Beyond the aggregate characteristics of the pseudoatoms themselves we have also considered the role of other structural features such as helium void fraction, volumetric surface area, and number density. First, we compared these three features to the CO<sub>2</sub> selectivity. One might naively expect that the most porous materials (those with the highest void fraction and surface area) would exhibit the best performance across all four adsorption evaluation criteria, however, we found that in general there were no strong correlations between any of these three structural features and CO<sub>2</sub> selectivity, but that the very highest CO<sub>2</sub> selectivity was observed in those pseudomaterials with below-average helium void fractions (see Figure 37a) and volumetric surface areas (see Figure 37b) as well as those pseudomaterials with relatively average number densities (see Figure 37c). Unfortunately, the highest CO<sub>2</sub> selectivities we observed among pseudomaterials with well-below-average regenerabilities, which means there is a substantial trade-off between these two adsorption evaluation criteria.



Figure 37. 2D projections of the CCS structure-property space depicting the relationship between CO<sub>2</sub> selectivity and different structural characteristics: helium void fraction (A), volumetric surface area (B), and number density (C), all colored by regenerability.

When comparing the same three structural features to  $CO_2$  working capacities, we observed – not surprisingly – that highly porous materials performed better (see Figure 38a, Figure 38b). As in the case of  $CO_2$  selectivities, the highest  $CO_2$  working capacity was observed among pseudomaterials of average number densities (see Figure 38c). Again, there is a substantial tradeoff between regenerability and  $CO_2$  working capacity.



Figure 38. 2D projectsions of the CCS structure-property space depicting the relationship between CO<sub>2</sub> working capacity and different structural characteristics: helium void fraction (A), volumetric surface area (B), and number density (C), all colored by regenerability.

Finally, because this model differs from our previous studies using pseudomaterials[96, 133] in that partial charges have been assigned to each pseudoatom, it is worthwhile to compare the contributions of van der Waals and Coulombic interactions to the performance of different pseudomaterials. First, we will consider the Coulombic interactions.  $CO_2$  selectivity and  $CO_2$  working capacity – the two most tightly correlated adsorption evaluation criteria – are both maximized at relatively average Coulombic host-adsorbate energies (U<sub>c</sub>) (see Figure 39a, Figure

39b), while regenerability is maximized when the magnitude of  $U_c$  is the lowest, that is when the Coulombic interactions are the weakest (see Figure 39c). The sorbent selection parameter is maximized when Coulombic interactions are weak (see Figure 39d), suggesting that weak Coulombic interactions give rise to the best sorbents.



Figure 39. 2D projections of the CCS structure-property space depicting the relationship between Coulombic heat of adsorption and the four adsorption evaluation criteria: CO<sub>2</sub> selectivity (A), CO<sub>2</sub> working capacity (B), regenerability (C), and sorbent selection parameter (D), all colored by helium void fraction.

Finally, we have evaluated the contributions of van der Waals interactions in a same manner as Coulombic interactions as well as by visualizing pseudomaterials from different regions of the structure-property space. We observe that the magnitude of van der Waals host-adsorbate energies  $(U_{vdw})$  were about half those their Coulombic counterparts. CO<sub>2</sub> selectivity, regenerability, and sorbent selection parameters were all maximized when the magnitudes of  $U_{vdw}$  were the lowest, that is when the van der Waals interactions were the weakest (see Figure 40a, Figure 40c, Figure 40d). CO2 working capacity was maximized with relatively weak van der Waals interactions (see Figure 40b).  $U_{vdw}$  is a function of the number density and average  $\varepsilon$ -value. Because we observe different upper boundaries for CO<sub>2</sub> working capacity (Figure 36b, Figure 40b) and regenerability (Figure 36c, Figure 40c) when considering  $U_{vdw}$  and average  $\varepsilon$ -values, we can assume that number density has a substantial effect upon these two adsorption evaluation criteria.



Figure 40. 2D projections of the CCS structure-property space depicting the relationship between van der Waals heat of adsorption and the four adsorption evaluation criteria: CO<sub>2</sub> selectivity (A), CO<sub>2</sub> working capacity (B), regenerability (C), and sorbent selection parameter (D), all colored by helium void fraction.

We have visualized different pseudomaterials from different regions of the structureproperty space, coloring pseudoatoms by their  $\varepsilon$ -value. We observe that the poorest performer has the most void space, so much so that the accessible surface area is diminished (see Figure 41a); it has more strong binding sites than the next poorest performer, whose performance is improved by an increased accessible surface area (see Figure 41b). As CO<sub>2</sub> working capacity increases (see Figure 41c, Figure 41d) we observe an increase in the number of strong binding sites as well as a decrease in average  $\sigma$ -values. As CO<sub>2</sub> selectivity increases (see Figure 41e) we observe fewer strong binding sites, as selectivity decreases with increasing average  $\varepsilon$ -values (see Figure 36a), and large pseudoatoms, as the smallest average  $\sigma$ -values are not correlated with the highest selectivities (Figure 35a).



Figure 41. Renderings of five pseudomaterials from different regions of the CCS structure-property space, indicated the lettering overlaid upoon the top-right heat map depicting CO<sub>2</sub> working capacity with respect to CO<sub>2</sub> selectivity. Here pseudoatoms are represented by spheres where the radius corresponds to the van der Waals radius (σ-value) and color corresponds to the potential well depth (ε-value).

## **5.4 Conclusion**

Our approach considers a large number and diversity of structures, not limited to any particular class of materials, resulting in a continuous spectrum of structural characteristics and observed properties. Our results are derived from an abstract model which is approximate and best suited for high-level, high-throughput analysis. We have considered four adsorption evaluation criteria: CO<sub>2</sub> working capacity, CO<sub>2</sub> selectivity, regenerability, and a previously-established sorbent selection parameter. We found the relatively small atom sites/moieties are ideal when considering the impact of their size upon the four adsorption evaluation criteria and that the presence of strong binding sites was not conducive to producing the best performing sorbents. Maximizing surface area or void fraction does not appear to optimize sorbent performance. We also found substantial tradeoffs between regenerability, CO2 working capacity, and CO2 selectivity where it may be worthwhile to sacrifice CO<sub>2</sub> working capacity and CO<sub>2</sub> selectivity to ensure high regenerability. We found that Coulombic interactions dictated a sorbent's performance more than van der Waals interactions. While these insights are valuable, experimental researchers would benefit from more rigorous characterization of these pseudomaterials. Future work may investigate the geometric and topological "fingerprints" of these abstract structures to bridge this theoretical approach to one that results in synthesizable structures.

We have simulated  $CO_2/N_2$  adsorption in 20,000 hypothetical pseudomaterials assessing the overall structure-property relationships governing VSA for CCS. The resulting structureproperty relationships with include surface area, void fraction, heats of adsorption, and more, are sharply defined and provide high-level insights into the design of adsorbents in VSA CCS processes.

# 6.0 Commercial Impact of Pseudomaterials Research

My pseudomaterials research lead to the successful spin-out of a startup company – Aeronics, Inc. – which specializes in lightweight, portable gas delivery systems. Pseudomaterials were used to investigate the potential of porous adsorbents for oxygen storage (see Figure 42).



Figure 42. Preliminary data generated for oxygen storage in porous pseudomaterials at 30 bar, the operating pressure of standard oxygen cylinders. Shown is a 3D mapping of a structure-property space consisting of volumetric oxygen loading, helium void fraction, and volumetric surfae area.

These results suggested a theoretical performance limit of over 600  $v_{STP}/v$  for oxygen at the same operating pressure as standard gas cylinders, which compress oxygen to ~26 v/v.

## 6.1 mediPOD: Product Ideation and Business Competitions

Considering the possibility of improving the storage capacity of standard oxygen cylinder by nearly thirty times, we proposed storing an hour's supply of oxygen in a soda can (a product *called mediPOD*). This would provide a lightweight, portable supply of medical grade oxygen for people suffering from chronic lung conditions, such as Chronic Obstructive Pulmonary Disorder (COPD). This design was intended to be a drop-in replacement for soft drink filling lines, where a supply of liquid oxygen would replace the soft drink and a porous adsorbent would ensure that the internal pressure would remain below the burst pressure for a standard soda can as the liquid oxygen entered the gas phase. Concept art for this design is shown in Figure 43 and Figure 44.



Figure 43. Concept art for mediPOD, cross-sectional view showing adsorbent and protective filter to prevent

inhalation of sorbent.



Figure 44. Concept art for mediPOD, demonstration of flexible adaptor connecting cannister to user's existing cannula.

Using this concept (mediPOD), Aeronics entered numerous business plan competitions. A complete listing of all prize money earned by Aeronics through business competitions appears in Table 5.

Competition Name	Year	Winnings
Randall Family Big Idea Competition, University of Pittsburgh	2016	\$25,000
Michael G. Wells Competition, University of Pittsburgh	2016	\$10,000
Kuzneski Innovation Cup, University of Pittsburgh	2016	\$5,000
TigerLaunch Startup Challenge, Princeton University	2017	\$10,000
xTechSearch, US Army ASA(ALT)	2018	\$130,000
•		

#### Table 5. Aeronics' winnings via business competitions.

#### 6.2 Incorporation, Fundraising, and Product Development

In 2017, Aeronics was incorporated as a C Corporation in the state of Delaware. It received pre-seed investment from Quake Capital, in New York City, NY, where it participated in the second cohort of its in-house accelerator. During this time Aeronics developed its first product, Everyday Oxygen (shown in Figure 45), a recreational oxygen (95%) product that contains three times the capacity of the leading competitor at the same price-point.



Figure 45. Everyday Oxygen, recreational oxygen with three times the capacity of the leading competitor.

While the release of Everyday Oxygen was delayed on account of shelf-life issues, Aeronics entered the veterinary market with an adsorbent-less product, called Pawprint Oxygen (shown in Figure 46).



Figure 46. Pawprint Oxygen, emergency oxygen for the veterinary market.

Since its launch, Pawprint Oxygen has generated \$18,000 in sales in Q1 2019; sales are projected to more than double to \$40,000 in Q2 2019. Aeronics has also raised over a \$1,000,000 in both dilutive and non-dilutive funding to date. Aeronics is current engaged key decision makes and subject matter experts within the Department of Defense to develop products for the US Army and US Special Operations Command (SOCOM).

## **Bibliography**

- [1] F. C. Hawthorne, Crystals from First Principles, Nature 345 (1990) 297-297. 10.1038/345297a0
- [2] G. R. Desiraju, Crystal Engineering: The Design of Organic Solids, Elsevier Scientific Publishers (1989)
- [3] A. Gavezzotti, Are Crystal Structures Predictable?, Accounts of Chemical Research 27 (1994) 309-314. 10.1021/ar00046a004
- [4] J. M. Lehn, Cryptates: Inclusion Complexes of Macropolycyclic Receptor Molecules, Pure and Applied Chemistry 50 (1978) 871. 10.1351/pac197850090871
- [5] J. D. Dunitz, Phase Transitions in Molecular Crystals from a Chemical Viewpoint, Pure and Applied Chemistry 63 (1991) 177. 10.1351/pac199163020177
- [6] G. R. Desiraju, Supramolecular Synthons in Crystal Engineering—a New Organic Synthesis, Angewandte Chemie International Edition in English 34 (1995) 2311-2327. 10.1002/anie.199523111
- [7] G. R. Desiraju, The C-H.Cntdot..Cntdot..Cntdot.O Hydrogen Bond in Crystals: What Is It?, Accounts of Chemical Research 24 (1991) 290-296. 10.1021/ar00010a002
- [8] G. R. Desiraju and A. Gavezzotti, From Molecular to Crystal Structure; Polynuclear Aromatic Hydrocarbons, Journal of the Chemical Society, Chemical Communications (1989) 621-623. 10.1039/C39890000621
- [9] M. C. Etter, Hydrogen Bonds as Design Elements in Organic Chemistry, The Journal of Physical Chemistry 95 (1991) 4601-4610. 10.1021/j100165a007
- [10] M. C. Etter, Encoding and Decoding Hydrogen-Bond Patterns of Organic Compounds, Accounts of Chemical Research 23 (1990) 120-126. 10.1021/ar00172a005
- [11] M. C. Etter, J. Macdonald and J. Bernstein, Graph-Set Analysis of Hydrogen Bond Patterns, (1990)
- [12] A. F. Wells, Three-Dimensional Nets and Polyhedra, Wiley (1977)
- [13] A. F. Wells, Structural Inorganic Chemistry, Oxford University Press (1984)
- [14] B. F. Abrahams, B. F. Hoskins and R. Robson, A New Type of Infinite 3d Polymeric Network Containing 4-Connected, Peripherally-Linked Metalloporphyrin Building Blocks, Journal of the American Chemical Society 113 (1991) 3606-3607. 10.1021/ja00009a065

- [15] S. Batten, R. Robson, P. Jensen, B. Moubaraki and K. Murray, Structure and Molecular Magnetism of the Rutile-Related Compounds M(Dca)2, M = Coii, Niii, Cuii, Dca = Dicyanamide, N(Cn)2-, (1998)
- [16] R. W. Gable, B. F. Hoskins and R. Robson, Synthesis and Structure of [Nme4][Cupt(Cn)4]: An Infinite Three-Dimensional Framework Related to Pts Which Generates Intersecting Hexagonal Channels of Large Cross Section, Journal of the Chemical Society, Chemical Communications (1990) 762-763. 10.1039/C39900000762
- [17] B. F. Hoskins and R. Robson, Design and Construction of a New Class of Scaffolding-Like Materials Comprising Infinite Polymeric Frameworks of 3d-Linked Molecular Rods. A Reappraisal of the Zinc Cyanide and Cadmium Cyanide Structures and the Synthesis and Structure of the Diamond-Related Frameworks [N(Ch3)4][Cuiznii(Cn)4] and Cui[4,4',4",4"'-Tetracyanotetraphenylmethane]Bf4.Xc6h5no2, Journal of the American Chemical Society 112 (1990) 1546-1554. 10.1021/ja00160a038
- [18] M. Kondo, T. Yoshitomi, H. Matsuzaka, S. Kitagawa and K. Seki, Three-Dimensional Framework with Channeling Cavities for Small Molecules: {[M2(4, 4'-Bpy)3(No3)4]·Xh2o}N (M □ Co, Ni, Zn), Angewandte Chemie International Edition in English 36 (1997) 1725-1727. 10.1002/anie.199717251
- [19] H. Li, M. Eddaoudi, M. O'Keeffe and O. M. Yaghi, Design and Synthesis of an Exceptionally Stable and Highly Porous Metal-Organic Framework, Nature 402 (1999) 276-279. 10.1038/46248
- [20] N. W. Ockwig, O. Delgado-Friedrichs, M. O'Keeffe and O. M. Yaghi, Reticular Chemistry: Occurrence and Taxonomy of Nets and Grammar for the Design of Frameworks, Accounts of Chemical Research 38 (2005) 176-182. 10.1021/ar0200221
- [21] O. M. Yaghi and H. Li, Hydrothermal Synthesis of a Metal-Organic Framework Containing Large Rectangular Channels, Journal of the American Chemical Society 117 (1995) 10401-10402. 10.1021/ja00146a033
- [22] C. Janiak, Engineering Coordination Polymers Towards Applications, Dalton Transactions (2003) 2781-2804. 10.1039/B305705B
- [23] S. T. Meek, J. A. Greathouse and M. D. Allendorf, Metal-Organic Frameworks: A Rapidly Growing Class of Versatile Nanoporous Materials, Advanced Materials 23 (2011) 249-267. 10.1002/adma.201002854
- [24] J. L.C. Rowsell and O. M. Yaghi, Metal-Organic Frameworks: A New Class of Porous Materials, (2004)
- [25] U. Mueller, M. Schubert, F. Teich, H. Puetter, K. Schierle-Arndt and J. Pastré, Metal-Organic Frameworks—Prospective Industrial Applications, Journal of Materials Chemistry 16 (2006) 626-636. 10.1039/B511962F

- [26] S. L. James, Metal-Organic Frameworks, Chemical Society Reviews 32 (2003) 276-288. 10.1039/B200393G
- [27] C. Sanchez, B. Julián, P. Belleville and M. Popall, Applications of Hybrid Organic– Inorganic Nanocomposites, Journal of Materials Chemistry 15 (2005) 3559-3592. 10.1039/B509097K
- [28] D. Zhao, D. J. Timmons, D. Yuan and H.-C. Zhou, Tuning the Topology and Functionality of Metal–Organic Frameworks by Ligand Design, Accounts of Chemical Research 44 (2011) 123-133. 10.1021/ar100112y
- [29] Y. Bai, Y. Dou, L.-H. Xie, W. Rutledge, J.-R. Li and H.-C. Zhou, Zr-Based Metal–Organic Frameworks: Design, Synthesis, Structure, and Applications, Chemical Society Reviews 45 (2016) 2327-2367. 10.1039/C5CS00837A
- [30] J. L. C. Rowsell and O. M. Yaghi, Strategies for Hydrogen Storage in Metal–Organic Frameworks, Angewandte Chemie International Edition 44 (2005) 4670-4679. 10.1002/anie.200462786
- [31] R. B. Getman, Y.-S. Bae, C. E. Wilmer and R. Q. Snurr, Review and Analysis of Molecular Simulations of Methane, Hydrogen, and Acetylene Storage in Metal–Organic Frameworks, Chemical Reviews 112 (2012) 703-723. 10.1021/cr200217c
- [32] C. Liu, F. Li, L. P. Ma and H. M. Cheng, Advanced Materials for Energy Storage, Advanced materials (Deerfield Beach, Fla.) 22 (2010) E28-62. 10.1002/adma.200903328
- [33] S. Ma and H.-C. Zhou, Gas Storage in Porous Metal–Organic Frameworks for Clean Energy Applications, Chemical Communications 46 (2010) 44-53. 10.1039/B916295J
- [34] C. E. Wilmer, O. K. Farha, T. Yildirim, I. Eryazici, V. Krungleviciute, A. A. Sarjeant, R. Q. Snurr and J. T. Hupp, Gram-Scale, High-Yield Synthesis of a Robust Metal–Organic Framework for Storing Methane and Other Gases, Energy & Environmental Science 6 (2013) 1158-1163. 10.1039/C3EE24506C
- [35] N. S. Bobbitt, J. Chen and R. Q. Snurr, High-Throughput Screening of Metal–Organic Frameworks for Hydrogen Storage at Cryogenic Temperature, The Journal of Physical Chemistry C 120 (2016) 27328-27341. 10.1021/acs.jpcc.6b08729
- [36] N. L. Rosi, J. Eckert, M. Eddaoudi, D. T. Vodak, J. Kim, M. Keeffe and O. M. Yaghi, Hydrogen Storage in Microporous Metal-Organic Frameworks, Science 300 (2003) 1127. 10.1126/science.1083440
- [37] H. Furukawa, K. E. Cordova, M. O'Keeffe and O. M. Yaghi, The Chemistry and Applications of Metal-Organic Frameworks, Science 341 (2013) 1230444. 10.1126/science.1230444

- [38] J.-R. Li, R. J. Kuppler and H.-C. Zhou, Selective Gas Adsorption and Separation in Metal– Organic Frameworks, Chemical Society Reviews 38 (2009) 1477-1504. 10.1039/B802426J
- [39] J. S. Seo, D. Whang, H. Lee, S. I. Jun, J. Oh, Y. J. Jeon and K. Kim, A Homochiral Metal– Organic Porous Material for Enantioselective Separation and Catalysis, Nature 404 (2000) 982-986. 10.1038/35010088
- [40] J.-R. Li, J. Sculley and H.-C. Zhou, Metal–Organic Frameworks for Separations, Chemical Reviews 112 (2012) 869-932. 10.1021/cr200190s
- [41] H. K. Chae, D. Y. Siberio-Pérez, J. Kim, Y. Go, M. Eddaoudi, A. J. Matzger, M. O'Keeffe and O. M. Yaghi, A Route to High Surface Area, Porosity and Inclusion of Large Molecules in Crystals, Nature 427 (2004) 523-527. 10.1038/nature02311
- [42] X. Chen, A. M. Plonka, D. Banerjee, R. Krishna, H. T. Schaef, S. Ghose, P. K. Thallapally and J. B. Parise, Direct Observation of Xe and Kr Adsorption in a Xe-Selective Microporous Metal–Organic Framework, Journal of the American Chemical Society 137 (2015) 7007-7010. 10.1021/jacs.5b02556
- [43] T. Rodenas, I. Luz, G. Prieto, B. Seoane, H. Miro, A. Corma, F. Kapteijn, F. X. Llabrés i Xamena and J. Gascon, Metal–Organic Framework Nanosheets in Polymer Composite Materials for Gas Separation, Nature Materials 14 (2015) 48-55. 10.1038/nmat4113
- [44] B. J. Sikora, C. E. Wilmer, M. L. Greenfield and R. Q. Snurr, Thermodynamic Analysis of Xe/Kr Selectivity in over 137 000 Hypothetical Metal–Organic Frameworks, Chemical Science 3 (2012) 2217-2223. 10.1039/C2SC01097F
- [45] Y.-S. Bae and R. Q. Snurr, Development and Evaluation of Porous Materials for Carbon Dioxide Separation and Capture, Angewandte Chemie International Edition 50 (2011) 11586-11596. 10.1002/anie.201101891
- [46] J. Liu, L. Chen, H. Cui, J. Zhang, L. Zhang and C.-Y. Su, Applications of Metal–Organic Frameworks in Heterogeneous Supramolecular Catalysis, Chemical Society Reviews 43 (2014) 6011-6061. 10.1039/C4CS00094C
- [47] L. Ma, C. Abney and W. Lin, Enantioselective Catalysis with Homochiral Metal–Organic Frameworks, Chemical Society Reviews 38 (2009) 1248-1256. 10.1039/B807083K
- [48] A. Corma, H. García and F. X. Llabrés i Xamena, Engineering Metal Organic Frameworks for Heterogeneous Catalysis, Chemical Reviews 110 (2010) 4606-4655. 10.1021/cr9003924
- [49] J. Lee, O. K. Farha, J. Roberts, K. A. Scheidt, S. T. Nguyen and J. T. Hupp, Metal–Organic Framework Materials as Catalysts, Chemical Society Reviews 38 (2009) 1450-1459. 10.1039/B807080F

- [50] A. Taguchi and F. Schüth, Ordered Mesoporous Materials in Catalysis, Microporous and Mesoporous Materials 77 (2005) 1-45. 10.1016/j.micromeso.2004.06.030
- [51] M. E. Davis, Ordered Porous Materials for Emerging Applications, Nature 417 (2002) 813-821. 10.1038/nature00785
- [52] B. Chen, S. Xiang and G. Qian, Metal–Organic Frameworks with Functional Pores for Recognition of Small Molecules, Accounts of Chemical Research 43 (2010) 1115-1124. 10.1021/ar100023y
- [53] Z. Hu, B. J. Deibert and J. Li, Luminescent Metal–Organic Frameworks for Chemical Sensing and Explosive Detection, Chemical Society Reviews 43 (2014) 5815-5840. 10.1039/C4CS00010B
- [54] L. E. Kreno, K. Leong, O. K. Farha, M. Allendorf, R. P. Van Duyne and J. T. Hupp, Metal– Organic Framework Materials as Chemical Sensors, Chemical Reviews 112 (2012) 1105-1125. 10.1021/cr200324t
- [55] H.-L. Jiang, Y. Tatsu, Z.-H. Lu and Q. Xu, Non-, Micro-, and Mesoporous Metal–Organic Framework Isomers: Reversible Transformation, Fluorescence Sensing, and Large Molecule Separation, Journal of the American Chemical Society 132 (2010) 5586-5587. 10.1021/ja101541s
- [56] D. Biswal and P. G. Kusalik, Probing Molecular Mechanisms of Self-Assembly in Metal– Organic Frameworks, ACS Nano 11 (2017) 258-268. 10.1021/acsnano.6b05444
- [57] Daniel W. Davies, Keith T. Butler, Adam J. Jackson, A. Morris, Jarvist M. Frost, Jonathan M. Skelton and A. Walsh, Computational Screening of All Stoichiometric Inorganic Materials, Chem 1 (2016) 617-627. 10.1016/j.chempr.2016.09.010
- [58] O. D. Friedrichs, A. W. M. Dress, D. H. Huson, J. Klinowski and A. L. Mackay, Systematic Enumeration of Crystalline Networks, Nature 400 (1999) 644-647. 10.1038/23210
- [59] D. J. Earl and M. W. Deem, Toward a Database of Hypothetical Zeolite Structures, Industrial & Engineering Chemistry Research 45 (2006) 5449-5454. 10.1021/ie0510728
- [60] M. W. Deem, R. Pophale, P. A. Cheeseman and D. J. Earl, Computational Discovery of New Zeolite-Like Materials, The Journal of Physical Chemistry C 113 (2009) 21353-21360. 10.1021/jp906984z
- [61] E. Haldoupis, S. Nair and D. S. Sholl, Pore Size Analysis of > 250 000 Hypothetical Zeolites, Physical Chemistry Chemical Physics 13 (2011) 5053-5060. 10.1039/c0cp02766a
- [62] C. Mellot-Draznieks, S. Girard, G. Férey, J. C. Schön, Z. Cancarevic and M. Jansen, Computational Design and Prediction of Interesting Not-yet-Synthesized Structures of Inorganic Materials by Using Building Unit Concepts, Chemistry – A European Journal 8

(2002) 4102-4113. CHEM4102>3.0.CO;2-3

- [63] O. K. Farha, A. Ö. Yazaydın, I. Eryazici, C. D. Malliakas, B. G. Hauser, M. G. Kanatzidis, S. T. Nguyen, R. Q. Snurr and J. T. Hupp, *De Novo* Synthesis of a Metal–Organic Framework Material Featuring Ultrahigh Surface Area and Gas Storage Capacities, Nature Chemistry 2 (2010) 944-948. 10.1038/nchem.834
- [64] C. E. Wilmer, M. Leaf, C. Y. Lee, O. K. Farha, B. G. Hauser, J. T. Hupp and R. Q. Snurr, Large-Scale Screening of Hypothetical Metal–Organic Frameworks, Nature Chemistry 4 (2012) 83-89. 10.1038/nchem.1192
- [65] A. K. Rappe, C. J. Casewit, K. S. Colwell, W. A. Goddard and W. M. Skiff, Uff, a Full Periodic Table Force Field for Molecular Mechanics and Molecular Dynamics Simulations, Journal of the American Chemical Society 114 (1992) 10024-10035. 10.1021/ja00051a040
- [66] Y. J. Colón and R. Q. Snurr, High-Throughput Computational Screening of Metal–Organic Frameworks, Chemical Society Reviews 43 (2014) 5735-5749. 10.1039/C4CS00070F
- [67] D. Dubbeldam, S. Calero, D. E. Ellis and R. Q. Snurr, Raspa: Molecular Simulation Software for Adsorption and Diffusion in Flexible Nanoporous Materials, Molecular Simulation 42 (2016) 81-101. 10.1080/08927022.2015.1010082
- [68] D. Wu, Q. Yang, C. Zhong, D. Liu, H. Huang, W. Zhang and G. Maurin, Revealing the Structure–Property Relationships of Metal–Organic Frameworks for Co2 Capture from Flue Gas, Langmuir 28 (2012) 12094-12099. 10.1021/la302223m
- [69] J. Gascon and F. Kapteijn, Metal-Organic Framework Membranes—High Potential, Bright Future?, Angewandte Chemie International Edition 49 (2010) 1530-1532. 10.1002/anie.200906491
- [70] F. G. Kerry, Industrial Gas Handbok: Gas Separation and Purification, CRC Press (2007)
- [71] C. E. Wilmer, O. K. Farha, Y.-S. Bae, J. T. Hupp and R. Q. Snurr, Structure–Property Relationships of Porous Materials for Carbon Dioxide Separation and Capture, Energy & Environmental Science 5 (2012) 9849-9856. 10.1039/C2EE23201D
- [72] M. Eddaoudi, J. Kim, N. Rosi, D. Vodak, J. Wachter, M. Keeffe and O. M. Yaghi, Systematic Design of Pore Size and Functionality in Isoreticular Mofs and Their Application in Methane Storage, Science 295 (2002) 469. 10.1126/science.1067208
- [73] H. Furukawa, K. E. Cordova, M. O'Keeffe and O. M. Yaghi, The Chemistry and Applications of Metal-Organic Frameworks, Science 341 (2013) 1230444. 10.1126/science.1230444

- [74] B. Li, M. Chrzanowski, Y. Zhang and S. Ma, Applications of Metal-Organic Frameworks Featuring Multi-Functional Sites, Coordination Chemistry Reviews 307 (2016) 106-129. 10.1016/j.ccr.2015.05.005
- [75] L. Wang, Y. Han, X. Feng, J. Zhou, P. Qi and B. Wang, Metal–Organic Frameworks for Energy Storage: Batteries and Supercapacitors, Coordination Chemistry Reviews 307 (2016) 361-381. 10.1016/j.ccr.2015.09.002
- [76] N. L. Rosi, J. Kim, M. Eddaoudi, B. Chen, M. O'Keeffe and O. M. Yaghi, Rod Packings and Metal–Organic Frameworks Constructed from Rod-Shaped Secondary Building Units, Journal of the American Chemical Society 127 (2005) 1504-1518. 10.1021/ja0451230
- [77] S. Keskin and S. Kızılel, Biomedical Applications of Metal Organic Frameworks, Industrial & Engineering Chemistry Research 50 (2011) 1799-1812. 10.1021/ie101312k
- [78] H.-C. J. Zhou and S. Kitagawa, Metal–Organic Frameworks (Mofs), Chemical Society Reviews 43 (2014) 5415-5418. 10.1039/C4CS90059F
- [79] S. Li, Y. G. Chung and R. Q. Snurr, High-Throughput Screening of Metal–Organic Frameworks for Co2 Capture in the Presence of Water, Langmuir 32 (2016) 10368-10376. 10.1021/acs.langmuir.6b02803
- [80] Y. G. Chung, J. Camp, M. Haranczyk, B. J. Sikora, W. Bury, V. Krungleviciute, T. Yildirim, O. K. Farha, D. S. Sholl and R. Q. Snurr, Computation-Ready, Experimental Metal–Organic Frameworks: A Tool to Enable High-Throughput Screening of Nanoporous Crystals, Chem. Mater. 26 (2014) 6185-6192. 10.1021/cm502594j
- [81] Y. Liu, Z. U. Wang and H.-C. Zhou, Recent Advances in Carbon Dioxide Capture with Metal-Organic Frameworks, Greenhouse Gases: Science and Technology 2 (2012) 239-259. 10.1002/ghg.1296
- [82] D. Banerjee, A. J. Cairns, J. Liu, R. K. Motkuri, S. K. Nune, C. A. Fernandez, R. Krishna, D. M. Strachan and P. K. Thallapally, Potential of Metal–Organic Frameworks for Separation of Xenon and Krypton, Accounts of Chemical Research 48 (2015) 211-219. 10.1021/ar5003126
- [83] P. Ryan, O. K. Farha, L. J. Broadbelt and R. Q. Snurr, Computational Screening of Metal-Organic Frameworks for Xenon/Krypton Separation, AIChE Journal 57 (2011) 1759-1766. 10.1002/aic.12397
- [84] D. Banerjee, C. M. Simon, A. M. Plonka, R. K. Motkuri, J. Liu, X. Chen, B. Smit, J. B. Parise, M. Haranczyk and P. K. Thallapally, Metal–Organic Framework with Optimally Selective Xenon Adsorption and Separation, Nature Communications 7 (2016) ncomms11831. 10.1038/ncomms11831
- [85] J. L. C. Rowsell and O. M. Yaghi, Metal–Organic Frameworks: A New Class of Porous Materials, Microporous and Mesoporous Materials 73 (2004) 3-14. 10.1016/j.micromeso.2004.03.034

- [86] S. Curtarolo, G. L. W. Hart, M. B. Nardelli, N. Mingo, S. Sanvito and O. Levy, The High-Throughput Highway to Computational Materials Design, Nature Materials 12 (2013) 191-201. 10.1038/nmat3568
- [87] R. L. Martin, C. M. Simon, B. Smit and M. Haranczyk, In Silico Design of Porous Polymer Networks: High-Throughput Screening for Methane Storage Materials, Journal of the American Chemical Society 136 (2014) 5006-5022. 10.1021/ja4123939
- [88] C. M. Simon, J. Kim, D. A. Gomez-Gualdron, J. S. Camp, Y. G. Chung, R. L. Martin, R. Mercado, M. W. Deem, D. Gunter, M. Haranczyk, D. S. Sholl, R. Q. Snurr and B. Smit, The Materials Genome in Action: Identifying the Performance Limits for Methane Storage, Energy & Environmental Science 8 (2015) 1190-1199. 10.1039/C4EE03515A
- [89] S. Bonakala and S. Balasubramanian, Structure–Property Relationships in Amorphous Microporous Polymers, J. Phys. Chem. B 120 (2016) 557-565. 10.1021/acs.jpcb.5b08842
- [90] M. A. B. Meador, L. A. Capadona, L. McCorkle, D. S. Papadopoulos and N. Leventis, Structure–Property Relationships in Porous 3d Nanostructures as a Function of Preparation Conditions: Isocyanate Cross-Linked Silica Aerogels, Chem. Mater. 19 (2007) 2247-2260. 10.1021/cm070102p
- [91] J. A. Mason, M. Veenstra and J. R. Long, Evaluating Metal–Organic Frameworks for Natural Gas Storage, Chemical Science 5 (2013) 32-51. 10.1039/C3SC52633J
- [92] Y. Peng, V. Krungleviciute, I. Eryazici, J. T. Hupp, O. K. Farha and T. Yildirim, Methane Storage in Metal–Organic Frameworks: Current Records, Surprise Findings, and Challenges, Journal of the American Chemical Society 135 (2013) 11887-11894. 10.1021/ja4045289
- [93] F. Gándara, H. Furukawa, S. Lee and O. M. Yaghi, High Methane Storage Capacity in Aluminum Metal–Organic Frameworks, Journal of the American Chemical Society 136 (2014) 5271-5274. 10.1021/ja501606h
- [94] R. Pophale, P. A. Cheeseman and M. W. Deem, A Database of New Zeolite-Like Materials, Physical Chemistry Chemical Physics 13 (2011) 12407-12412. 10.1039/C0CP02255A
- [95] L.-C. Lin, A. H. Berger, R. L. Martin, J. Kim, J. A. Swisher, K. Jariwala, C. H. Rycroft, A. S. Bhown, M. W. Deem, M. Haranczyk and B. Smit, <I>in Silico</I> Screening of Carbon-Capture Materials, Nature Materials 11 (2012) 633-641. 10.1038/nmat3336
- [96] A. R. Kaija and C. E. Wilmer, Efficiently Mapping Structure–Property Relationships of Gas Adsorption in Porous Materials: Application to Xe Adsorption, Faraday Discuss. 201 (2017) 221-232. 10.1039/C7FD00038C
- [97] H. Frost and R. Q. Snurr, Design Requirements for Metal-Organic Frameworks as Hydrogen Storage Materials, The Journal of Physical Chemistry C 111 (2007) 18794-18803. 10.1021/jp076657p

- [98] International Energy Agency, Global Energy & Co<sub>2</sub> Status Report 2017 (2018)
- [99] S. Chu, Y. Cui and N. Liu, The Path Towards Sustainable Energy, Nature Materials 16 (2017) 16-22. 10.1038/nmat4834
- [100] J. H. Butler and S. A. Montzka, NOAA Earth System Research Laboratory, The Noaa Annual Greenhouse Gas Index (Aggi) (2018)
- [101] S. Chu, Carbon Capture and Sequestration, Science 325 (2009) 1599-1599. 10.1126/science.1181637
- [102] M. Tagliabue, D. Farrusseng, S. Valencia, S. Aguado, U. Ravon, C. Rizzo, A. Corma and C. Mirodatos, Natural Gas Treating by Selective Adsorption: Material Science and Chemical Engineering Interplay, Chemical Engineering Journal 155 (2009) 553-566. 10.1016/j.cej.2009.09.010
- [103] W. Liu, D. King, J. Liu, B. Johnson, Y. Wang and Z. Yang, Critical Material and Process Issues for Co2 Separation from Coal-Powered Plants, JOM 61 (2009) 36-44. 10.1007/s11837-009-0050-6
- [104] J. C. M. Pires, F. G. Martins, M. C. M. Alvim-Ferraz and M. Simões, Recent Developments on Carbon Capture and Storage: An Overview, Chemical Engineering Research and Design 89 (2011) 1446-1460. 10.1016/j.cherd.2011.01.028
- [105] T. C. Drage, K. M. Smith, C. Pevida, A. Arenillas and C. E. Snape, Development of Adsorbent Technologies for Post-Combustion Co2 Capture, Energy Procedia 1 (2009) 881-884. 10.1016/j.egypro.2009.01.117
- [106] J. Wang, L. Huang, R. Yang, Z. Zhang, J. Wu, Y. Gao, Q. Wang, D. O'Hare and Z. Zhong, Recent Advances in Solid Sorbents for Co2 Capture and New Development Trends, Energy & Environmental Science 7 (2014) 3478-3518. 10.1039/C4EE01647E
- [107] B. Metz, O. Davidson, H. de Coninck, M. Loos and L. Meyer, Intergovernmental Panel on Climate Change, Ipcc Special Report on Carbon Dioxide Capture and Storage (2005)
- [108] R. Ben-Mansour, M. A. Habib, O. E. Bamidele, M. Basha, N. A. A. Qasem, A. Peedikakkal, T. Laoui and M. Ali, Carbon Capture by Physical Adsorption: Materials, Experimental Investigations and Numerical Modeling and Simulations a Review, Applied Energy 161 (2016) 225-255. 10.1016/j.apenergy.2015.10.011
- [109] X. Xu, C. Song, J. M. Andresen, B. G. Miller and A. W. Scaroni, Novel Polyethylenimine-Modified Mesoporous Molecular Sieve of Mcm-41 Type as High-Capacity Adsorbent for Co<Sub>2</Sub> Capture, Energy & Fuels 16 (2002) 1463-1469. 10.1021/ef020058u
- [110] K. S. Lackner, Capture of Carbon Dioxide from Ambient Air, The European Physical Journal Special Topics 176 (2009) 93-106. 10.1140/epjst/e2009-01150-3

- [111] T. Wang, K. S. Lackner and A. Wright, Moisture Swing Sorbent for Carbon Dioxide Capture from Ambient Air, Environ. Sci. Technol. 45 (2011) 6670-6675. 10.1021/es201180v
- [112] C. Robertson and R. Mokaya, Microporous Activated Carbon Aerogels Via a Simple Subcritical Drying Route for Co2 Capture and Hydrogen Storage, Microporous and Mesoporous Materials 179 (2013) 151-156. 10.1016/j.micromeso.2013.05.025
- [113] A. A. Alhwaige, T. Agag, H. Ishida and S. Qutubuddin, Biobased Chitosan Hybrid Aerogels with Superior Adsorption: Role of Graphene Oxide in Co2 Capture, RSC Advances 3 (2013) 16011-16020. 10.1039/C3RA42022A
- [114] B. Bonelli, B. Onida, B. Fubini, C. O. Areán and E. Garrone, Vibrational and Thermodynamic Study of the Adsorption of Carbon Dioxide on the Zeolite Na–Zsm-5, Langmuir 16 (2000) 4976-4983. 10.1021/la991363j
- [115] S. K. Wirawan and D. Creaser, Co2 Adsorption on Silicalite-1 and Cation Exchanged Zsm-5 Zeolites Using a Step Change Response Method, Microporous and Mesoporous Materials 91 (2006) 196-205. 10.1016/j.micromeso.2005.11.047
- [116] Y. Hu, W. M. Verdegaal, S.-H. Yu and H.-L. Jiang, Alkylamine-Tethered Stable Metal-Organic Framework for Co2 Capture from Flue Gas, ChemSusChem 7 (2014) 734-737. 10.1002/cssc.201301163
- [117] Z. Xiang, X. Peng, X. Cheng, X. Li and D. Cao, Cnt@Cu3(Btc)2 and Metal–Organic Frameworks for Separation of Co2/Ch4 Mixture, The Journal of Physical Chemistry C 115 (2011) 19864-19871. 10.1021/jp206959k
- [118] Z. Xiang, D. Cao, J. Lan, W. Wang and D. P. Broom, Multiscale Simulation and Modelling of Adsorptive Processes for Energy Gas Storage and Carbon Dioxide Capture in Porous Coordination Frameworks, Energy & Environmental Science 3 (2010) 1469-1487. 10.1039/C0EE00049C
- [119] X. Peng, X. Cheng and D. Cao, Computer Simulations for the Adsorption and Separation of Co 2 /Ch 4 /H 2 /N 2 Gases by Umcm-1 and Umcm-2 Metal Organic Frameworks, Journal of Materials Chemistry 21 (2011) 11259-11270. 10.1039/C1JM10264H
- [120] M. Yao, Y. Dong, X. Feng, X. Hu, A. Jia, G. Xie, G. Hu, J. Lu, M. Luo and M. Fan, The Effect of Post-Processing Conditions on Aminosilane Functionalization of Mesocellular Silica Foam for Post-Combustion Co2 Capture, Fuel 123 (2014) 66-72. 10.1016/j.fuel.2014.01.059
- [121] R. Dawson, D. J. Adams and A. I. Cooper, Chemical Tuning of Co 2 Sorption in Robust Nanoporous Organic Polymers, Chemical Science 2 (2011) 1173-1177. 10.1039/C1SC00100K

- [122] N. Du, H. B. Park, G. P. Robertson, M. M. Dal-Cin, T. Visser, L. Scoles and M. D. Guiver, Polymer Nanosieve Membranes for Co<Sub>2</Sub>-Capture Applications, Nature Materials 10 (2011) 372-375. 10.1038/nmat2989
- [123] M. R. Liebl and J. Senker, Microporous Functionalized Triazine-Based Polyimides with High Co2 Capture Capacity, Chem. Mater. 25 (2013) 970-980. 10.1021/cm4000894
- [124] Y. Luo, B. Li, L. Liang and B. Tan, Synthesis of Cost-Effective Porous Polyimides and Their Gas Storage Properties, Chemical Communications 47 (2011) 7704-7706. 10.1039/C1CC11466B
- [125] N. Ritter, I. Senkovska, S. Kaskel and J. Weber, Intrinsically Microporous Poly(Imide)S: Structure–Porosity Relationship Studied by Gas Sorption and X-Ray Scattering, Macromolecules 44 (2011) 2025-2033. 10.1021/ma102448h
- [126] Q. Wang, Y. Gao, J. Luo, Z. Zhong, A. Borgna, Z. Guo and D. O'Hare, Synthesis of Nano-Sized Spherical Mg 3 Al–Co 3 Layered Double Hydroxide as a High-Temperature Co 2 Adsorbent, RSC Advances 3 (2013) 3414-3420. 10.1039/C2RA22607C
- [127] M. Liu, C. Vogt, A. L. Chaffee and S. L. Y. Chang, Nanoscale Structural Investigation of Cs2co3-Doped Mgo Sorbent for Co2 Capture at Moderate Temperature, The Journal of Physical Chemistry C 117 (2013) 17514-17520. 10.1021/jp4024316
- [128] C. S. Martavaltzi and A. A. Lemonidou, Development of New Cao Based Sorbent Materials for Co2 Removal at High Temperature, Microporous and Mesoporous Materials 110 (2008) 119-127. 10.1016/j.micromeso.2007.10.006
- [129] M. Y. Veliz-Enriquez, G. Gonzalez and H. Pfeiffer, Synthesis and Co2 Capture Evaluation of Li2–Xkxzro3 Solid Solutions and Crystal Structure of a New Lithium–Potassium Zirconate Phase, Journal of Solid State Chemistry 180 (2007) 2485-2492. 10.1016/j.jssc.2007.06.023
- [130] S. U. Rege and R. T. Yang, A Simple Parameter for Selecting an Adsorbent for Gas Separation by Pressure Swing Adsorption, Separation Science and Technology 36 (2001) 3355-3365. 10.1081/SS-100107907
- [131] R. Krishna and J. M. v. Baten, In Silico Screening of Metal–Organic Frameworks in Separation Applications, Physical Chemistry Chemical Physics 13 (2011) 10593-10616. 10.1039/C1CP20282K
- [132] E. Haldoupis, S. Nair and D. S. Sholl, Finding Mofs for Highly Selective Co2/N2 Adsorption Using Materials Screening Based on Efficient Assignment of Atomic Point Charges, Journal of the American Chemical Society 134 (2012) 4313-4323. 10.1021/ja2108239
- [133] A. R. Kaija and C. E. Wilmer, High-Pressure Methane Adsorption in Porous Lennard-Jones Crystals, J. Phys. Chem. Lett. 9 (2018) 4275-4281. 10.1021/acs.jpclett.8b01421