

Efficiently mapping structure–property relationships of gas adsorption in porous materials: application to Xe adsorption†

A. R. Kaija* and C. E. Wilmer 

Received 27th January 2017, Accepted 17th March 2017

DOI: 10.1039/c7fd00038c

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 structure–property 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.

Introduction

Porous materials are used widely in many applications related to catalysis,^{1–4} separations,^{5–13} gas storage,^{14–19} and chemical sensing,^{20–23} among others. Many of these applications stand to be improved from better porous materials, and so significant effort is devoted to searching for them.^{24–28}

Department of Chemical & Petroleum Engineering, University of Pittsburgh, Pittsburgh, PA 15261, USA. E-mail: ark111@pitt.edu

† Electronic supplementary information (ESI) available. See DOI: 10.1039/c7fd00038c

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.^{1,29,30} 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.³¹ This large design space of possible MOFs makes it an ideal class to search within for better porous materials.³²

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,^{7,33} 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.⁹ In that work, some trends appeared robust, such as CO₂ 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 (CO₂ loading *vs.* isosteric heat of adsorption) at the same conditions, but used molecular simulation data from over 130 000 hypothetical MOFs,¹³ 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.^{8,10,15,31,34–37} 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 0.8.¹⁴

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 after 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 ^{85}Kr from spent nuclear fuels.^{6,38,39} 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.^{6,40} 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.

Methodology

An overview schematic of our six step computational method is presented in Fig. 1. 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

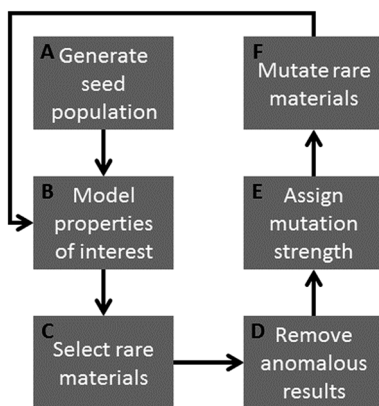


Fig. 1 Flow chart describing the method in six parts: generating a seed population (A), calculating properties of interest (B), selecting rare pseudomaterials as candidate parents (C), removing anomalous results due to statistical undersampling (D), adjusting mutation strengths (E), and mutating rare parents to create new materials (F).

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.

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 Fig. 2). 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




Fig. 2 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 σ and ϵ values respectively.

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 1.49×10^{-5} and 0.02122 atoms per Å³. 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: σ , van der Waals radius, and ϵ , the potential well depth. Values of σ and ϵ were randomly chosen in the range between 1.052 and 6.549 Å for σ and between 1.258 and 513.26 K/ k_B for ϵ . 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),⁴¹ where we allowed σ and ϵ 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.

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,⁴² with a helium probe ($\alpha = 2.96$ Å). See the ESI† for more details.

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 cm³ Xe per cm³ framework at 1 bar, 100 cm³ cm^{−3} at 5 bar, or 150 cm³ cm^{−3} at 10 bar. Pseudomaterials were chosen as candidate parents with probability inversely proportional to the number of materials in the same bin (see the ESI† for exact weighting function).

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.

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.

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.

Uniformity metric

To measure progress in exploring the structure–property space, we defined a uniformity metric (U). First each bin-count (c) was normalised: $c^* = c/c_{\max}$, where c_{\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): $U = \Sigma[c^* - (\Sigma c^*/N)]/N$. 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.

Results

In seed populations (Fig. 3) 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 (Fig. 4) 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 Fig. 4A); after fifty




Fig. 3 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 cm³ xenon per cm³ framework at 1, 5, and 10 bar respectively.

generations, the number of accessed bins has nearly doubled and the distribution continues to flatten (Fig. 4B).

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 Fig. 4




Fig. 4 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.

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 20th and 30th generations in the 1 bar run and the 20th generation in the 5 bar run (see Fig. 5A and B) 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 20th through 30th 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 Fig. 5B).

The uniformity metric (see Fig. 6) 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

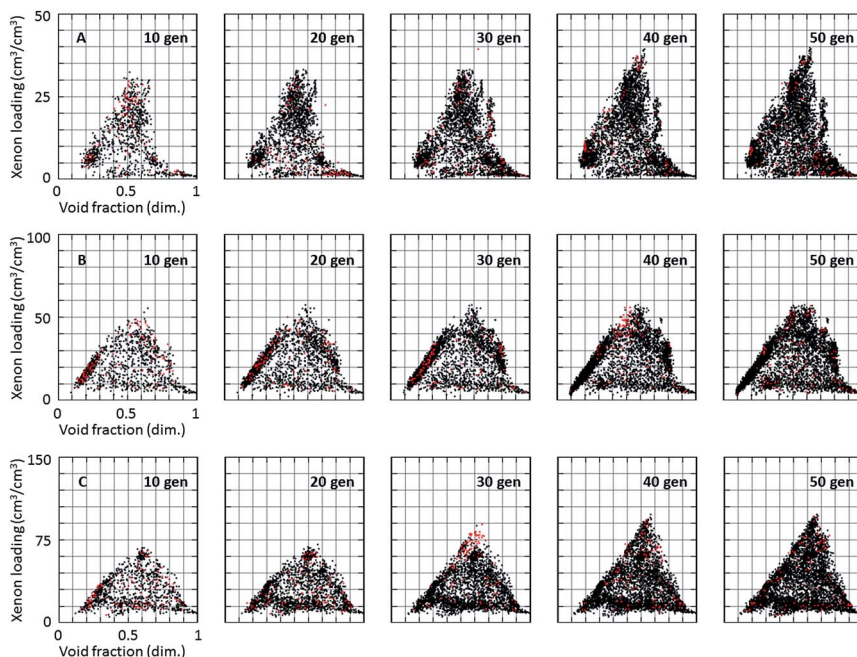


Fig. 5 Scatterplots after 10, 20, 30, 40, and 50 generations (from left to right) of children had been added for runs at 1 bar (A), 5 bar (B), and 10 bar (C). Children added in the last generation are highlighted in red.

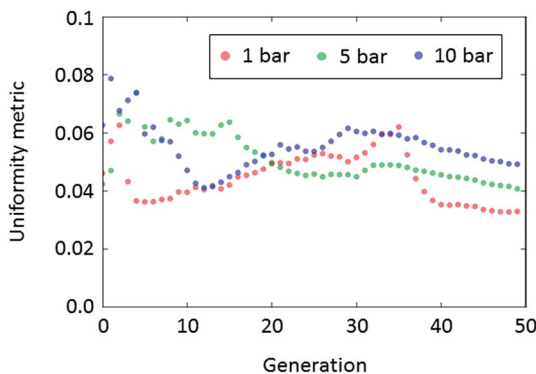


Fig. 6 The uniformity metric after each generation for each of the 1 (red), 5 (blue), and 10 bar (green) runs.

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 (Fig. 6), 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 Fig. 5A, 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. Fig. 6 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.

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 Fig. 5) closely resembles what was seen for Xe adsorption in a high throughput screening study on a database of MOFs.⁸ In a similarly 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.

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.

Acknowledgements

Computational work was supported through the resources provided by the Center for Simulation and Modeling (SAM) at the University of Pittsburgh. A. R. K. and C. E. W. gratefully acknowledge support from University of Pittsburgh, Swanson School of Engineering.

Notes and references

- 1 J. Lee, O. K. Farha, J. Roberts, K. A. Scheidt, S. T. Nguyen and J. T. Hupp, *Chem. Soc. Rev.*, 2009, **38**, 1450.
- 2 A. Taguchi and F. Schüth, *Microporous Mesoporous Mater.*, 2005, **77**, 1–45.
- 3 M. E. Davis, *Nature*, 2002, **417**, 813–821.
- 4 H. Li, M. Eddaoudi, M. O’Keeffe and O. M. Yaghi, *Nature*, 1999, **402**, 276–279.
- 5 H. K. Chae, D. Y. Siberio-Pérez, J. Kim, Y. Go, M. Eddaoudi, A. J. Matzger, M. O’Keeffe and O. M. Yaghi, *Nature*, 2004, **427**, 523–527.
- 6 J. Liu, P. K. Thallapally and D. Strachan, *Langmuir*, 2012, **28**, 11584–11589.
- 7 J. Gascon and F. Kapteijn, *Angew. Chem., Int. Ed.*, 2010, **49**, 1530–1532.
- 8 B. J. Sikora, C. E. Wilmer, M. L. Greenfield and R. Q. Snurr, *Chem. Sci.*, 2012, **3**, 2217–2223.
- 9 Y.-S. Bae and R. Q. Snurr, *Angew. Chem., Int. Ed.*, 2011, **50**, 11586–11596.
- 10 B. C. Yeo, D. Kim, H. Kim and S. S. Han, *J. Phys. Chem. C*, 2016, **120**, 24224–24230.
- 11 F. G. Kerry, *Industrial gas handbook: gas separation and purification*, CRC Press, Boca Raton, FL, 2007.
- 12 J.-R. Li, R. J. Kuppler and H.-C. Zhou, *Chem. Soc. Rev.*, 2009, **38**, 1477.

- 13 C. E. Wilmer, O. K. Farha, Y.-S. Bae, J. T. Hupp and R. Q. Snurr, *Energy Environ. Sci.*, 2012, **5**, 9849–9856.
- 14 C. E. Wilmer, O. K. Farha, T. Yildirim, I. Eryazici, V. Krungleviciute, A. A. Sarjeant, R. Q. Snurr and J. T. Hupp, *Energy Environ. Sci.*, 2013, **6**, 1158–1163.
- 15 N. S. Bobbitt, J. Chen and R. Q. Snurr, *J. Phys. Chem. C*, 2016, **120**, 27328–27341.
- 16 N. L. Rosi, *Science*, 2003, **300**, 1127–1129.
- 17 M. Eddaoudi, *Science*, 2002, **295**, 469–472.
- 18 C. Sanchez, B. Julián, P. Belleville and M. Popall, *J. Mater. Chem.*, 2005, **15**, 3559.
- 19 H. Furukawa, K. E. Cordova, M. O’Keeffe and O. M. Yaghi, *Science*, 2013, **341**, 1230444.
- 20 Z. Hu, B. J. Deibert and J. Li, *Chem. Soc. Rev.*, 2014, **43**, 5815–5840.
- 21 B. Chen, S. Xiang and G. Qian, *Acc. Chem. Res.*, 2010, **43**, 1115–1124.
- 22 H.-L. Jiang, Y. Tatsu, Z.-H. Lu and Q. Xu, *J. Am. Chem. Soc.*, 2010, **132**, 5586–5587.
- 23 S. T. Meek, J. A. Greathouse and M. D. Allendorf, *Adv. Mater.*, 2011, **23**, 249–267.
- 24 J. Liu, L. Chen, H. Cui, J. Zhang, L. Zhang and C.-Y. Su, *Chem. Soc. Rev.*, 2014, **43**, 6011–6061.
- 25 D. Zhao, D. J. Timmons, D. Yuan and H.-C. Zhou, *Acc. Chem. Res.*, 2011, **44**, 123–133.
- 26 Y. Bai, Y. Dou, L.-H. Xie, W. Rutledge, J.-R. Li and H.-C. Zhou, *Chem. Soc. Rev.*, 2016, **45**, 2327–2367.
- 27 B. Li, M. Chrzanowski, Y. Zhang and S. Ma, *Coord. Chem. Rev.*, 2016, **307**, 106–129.
- 28 L. Wang, Y. Han, X. Feng, J. Zhou, P. Qi and B. Wang, *Coord. Chem. Rev.*, 2016, **307**, 361–381.
- 29 N. W. Ockwig, O. Delgado-Friedrichs, M. O’Keeffe and O. M. Yaghi, *Acc. Chem. Res.*, 2005, **38**, 176–182.
- 30 N. L. Rosi, J. Kim, M. Eddaoudi, B. L. Chen, M. O’Keeffe and O. M. Yaghi, *J. Am. Chem. Soc.*, 2005, **127**, 1504–1518.
- 31 C. E. Wilmer, M. Leaf, C. Y. Lee, O. K. Farha, B. G. Hauser, J. T. Hupp and R. Q. Snurr, *Nat. Chem.*, 2012, **4**, 83–89.
- 32 S. Keskin and S. Kizilel, *Ind. Eng. Chem. Res.*, 2011, **50**, 1799–1812.
- 33 H.-C. Joe Zhou and S. Kitagawa, *Chem. Soc. Rev.*, 2014, **43**, 5415–5418.
- 34 Y. J. Colón and R. Q. Snurr, *Chem. Soc. Rev.*, 2014, **43**, 5735–5749.
- 35 S. Li, Y. G. Chung and R. Q. Snurr, *Langmuir*, 2016, **32**, 10368–10376.
- 36 E. Haldoupis, S. Nair and D. S. Sholl, *Phys. Chem. Chem. Phys.*, 2011, **13**, 5053–5060.
- 37 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, *Chem. Mater.*, 2014, **26**, 6185–6192.
- 38 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, *Nat. Commun.*, 2016, **7**, ncomms11831.
- 39 P. Ryan, O. K. Farha, L. J. Broadbelt and R. Q. Snurr, *AIChE J.*, 2011, **57**, 1759–1766.

- 40 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, *Acc. Chem. Res.*, 2015, **48**, 211–219.
- 41 A. K. Rappe, C. J. Casewit, K. S. Colwell, W. A. Goddard and W. M. Skiff, *J. Am. Chem. Soc.*, 1992, **114**, 10024–10035.
- 42 D. Dubbeldam, S. Calero, D. E. Ellis and R. Q. Snurr, *Mol. Simul.*, 2016, **42**, 81–101.