

<https://doi.org/10.1038/s41524-024-01373-9>

Gas permeability, diffusivity, and solubility in polymers: Simulation-experiment data fusion and multi-task machine learning

Check for updates

Brandon K. Phan¹, Kuan-Hsuan Shen¹, Rishi Gurnani¹, Huan Tran¹, Ryan Lively² & Rampi Ramprasad¹ ✉

Machine learning (ML) models for predicting gas permeability through polymers have traditionally relied on experimental data. While these models exhibit robustness within familiar chemical domains, reliability wanes when applied to new spaces. To address this challenge, we present a multi-tiered multi-task learning framework empowered with advanced machine-crafted polymer fingerprinting algorithms and data fusion techniques. This framework combines scarce “high-fidelity” experimental data with abundant diverse “low-fidelity” simulation or synthetic data, resulting in predictive models that display a high level of generalizability across novel chemical spaces. Additionally, this multi-task scheme capitalizes on known physics and interrelated properties, such as gas diffusivity and solubility, both of which are closely tied to permeability. By amalgamating high throughput generated simulation data with available experimental data for gas permeability, diffusivity, and solubility for various gases, we construct multi-task deep learning models. These models can simultaneously predict all three properties for all gases under consideration, with markedly enhanced predictive accuracy, particularly compared to traditional models reliant solely on experimental data for a singular property. This strategy underscores the potential of coupling high-throughput classical simulations with data fusion methodologies to yield state-of-the-art property predictors, especially when experimental data for targeted properties is scarce.

Polymer-based gas and solvent separation membrane technologies have significantly impacted a diverse range of applications, including carbon capture, water purification, drug delivery, and food packaging^{1,2}. Crucial to propelling widespread adoption and advancement of this technology is the identification and design of polymer materials endowed with a desired set of properties and performance attributes. A key figure of merit in gas separations is gas permeability, which describes the movement of gas molecules into and through a polymer material. Based on the solution-diffusion model³, gas permeability (P) through a membrane is defined as the product of gas diffusivity (D) and gas solubility (S):

$$P = DS \quad (1)$$

Capabilities that can accurately and rapidly predict gas permeability across a diverse range of gases and polymer chemistries can be transformational and facilitate the discovery and development of new sustainable high-performance polymer membranes^{4,5}.

Traditionally, the measurement of gas permeability relies on the constant volume permeation technique⁶, which, though serving as the primary benchmark, is both time and resource intensive. In search of alternative approaches, classical molecular dynamics (MD) simulations have emerged as a complementary pathway to estimate gas permeability⁷. However, the fidelity of these simulations is constrained by the intrinsic limitations of the classical force fields employed and timescales that are computationally accessible. As a result, they can only achieve, at best, semi-quantitative agreement with experimental measurements, despite correctly capturing general trends.

In recent times, data-driven machine learning (ML) methods have achieved remarkable strides, fundamentally reshaping the landscape of materials property predictions and the tailored design of materials with specific target characteristics^{4,5,8–13}. ML methods have found extensive applications in the polymer gas transport domain, encompassing a diverse array of studies varying in the number of polymers investigated and the types of features used to train models. An early example of this is the work by Wessline et al. in 2006, where a neural network was used to correlate the

¹School of Materials Science and Engineering, Georgia Institute of Technology, Atlanta, GA, USA. ²School of Chemical and Biomolecular Engineering, Georgia Institute of Technology, Atlanta, GA, USA. ✉e-mail: rampi.ramprasad@mse.gatech.edu

infrared spectra of 33 polymers with their carbon dioxide permeability¹⁴. In a more current study, Yuan et al. utilized Multivariate Imputation by Chained Equations (MICE) to predict missing gas permeability values in a dataset spanning hundreds of polymers across six gases¹⁵. These examples only scratch the surface. In a comprehensive perspective paper, Ricci et al. delve deeper into the evolution of ML in modeling gas separation with polymer membranes, highlighting strategies, challenges, and future directions¹⁶.

These informatics approaches require a critical initial step: defining the feature space in which the models are trained by mapping features to the properties being learned. Early machine learning (ML) studies employed simple feature sets; for instance, in a 2006 study by Wang et al., six features related to the experimental setup, such as temperature, feed gas flux, and permeate-side pressure, were used¹⁷. These approaches have transitioned to incorporate more descriptive and comprehensive features, capturing atomistic to morphological structural details^{18,19}. In this paradigm, a polymer's chemical structure is converted into a machine readable numerical representation, commonly known as a fingerprint or feature vector. This fingerprint allows an ML algorithm in the second step (during the training phase) to discern intricate chemistry-morphology-property relationships and subsequently generate predictive models for the properties. While traditional hand-crafted fingerprints^{20,21} have conventionally represented polymer structures in machine learning models, recent endeavors have expanded the horizons of this methodology and have led to learned fingerprinting techniques, which we adopt in this study. These techniques involve machine learning key features directly from polymer repeat units, offering faster feature extraction with comparable accuracy^{22,23}. Despite these advancements, a common challenge these methods encounter is extrapolating outside of the known polymer-property space, i.e., outside of the training data space²⁴. Exploring new chemical spaces through various avenues, including experiments, simulations, and machine learning models poses unique limitations that necessitate innovative solutions.

In the present contribution, we demonstrate the power of multi-task (MT) learning, harnessing both experimental and computational data to address and bridge the shortcomings outlined above, to build a best-in-class gas transport property predictor. MT learning is a type of transfer learning in which a model is trained on more than one task, learning multiple properties and/or data sources simultaneously²⁵. In contrast, single-task (ST) learning involves the consideration of a singular property and data source. The MT architecture, which integrates various data sources and exploits underlying correlations and calibrations, has shown improved predictive performance and enhanced transferability, compared to ST methods^{26,27}. In the polymer gas transport ML space, MT learning has been commonly implemented by incorporating permeability data for various gases and utilizing datasets that encompass a broad spectrum of properties, including mechanical, thermal, and thermodynamic^{5,28,29}. We expand on these previous works by utilizing MT learning in two novel ways. The first aspect leverages data fidelity by fusing “high-fidelity” experimental data with “low-fidelity” simulation data. While experimentally measured data serves as the ground truth, it often grapples with constraints stemming from labor-intensive protocols and associated expenses. Conversely, simulation-generated data can be produced on a grander scale, but it may exhibit diminished accuracy due to necessary approximations made in the theory to make the simulations practical. MT algorithms learn to calibrate the low-fidelity (simulation) data against the high-fidelity (measured) data across the whole space of the data, thus leading to a high level of generalizability^{27,30}. Typically, gas simulations have been used to validate ML predictions. Here, we integrate the simulation data into the model itself.

The second innovative aspect of the MT learning approach extends the general gas permeability ML model to include directly correlated gas transport data; diffusivity and solubility. Gas transport experiments reported in the literature do not always include all three properties, and they tend to focus on testing specific gases of interest. Consequently, some property values may not be available for certain cases. MT learning offers a solution to this challenge by drawing on available properties to learn correlations between them and make effective generalizations²⁸. Incorporated within

these two outlined MT aspects is the integration of gas transport data spanning a variety of gases. Our MT learning strategy leverages potential correlations between the transport characteristics of multiple similar (or dissimilar) gases¹⁵. A unified model that harnesses data from (1) diverse sources (i.e., measured and simulated), (2) spanning multiple correlated properties (i.e., P , D , & S), and (3) for various gases, can lead to enhanced predictive performance and generalizability, as will be demonstrated here.

A key ingredient of our MT learning approach involves simulation data that could complement measured data for gas transport, as illustrated in Fig. 1a. To achieve this, we have designed a high-throughput molecular dynamics (MD) and Monte Carlo (MC) simulation pipeline, depicted in Fig. 1b. This pipeline generates data for gas diffusivity (D_{sim}) and solubility (S_{sim}); the subscripts explicitly indicate the source of the data. Simulated gas permeability (P_{sim}) is then derived from the product of D_{sim} and S_{sim} as prescribed by Eq. 1. Experimental data are labeled as P_{expt} , D_{expt} and S_{expt} . Data for 6 different gases (CO₂, CH₄, O₂, N₂, H₂, and He) span this study. An overview of the dataset is presented in Fig. 1c. With this fused dataset, ML models for gas transport properties are created using our newly-developed graph neural networks method – polyGNN²², thus completing the MT learning pipeline as visualized in Fig. 1a. The input for polyGNN consists of polymer “Simplified Molecular-Input Line-Entry System” (SMILES)³¹ strings. These SMILES strings are translated into graph representations and fingerprints, an essential ingredient for the property prediction model trained on the integrated dataset. The architecture of polyGNN, exhibited in Fig. 1d, illustrates this process.

To test our MT learning approach, we constructed four distinct models to examine and benchmark the impact of incorporating multiple data streams. These models were designed to emulate real-world usage scenarios for the prediction model's application and to assess the improvements in prediction capabilities. To evaluate the efficacy of the MT learning, a comparison with ST learning is employed. Through these case studies, we demonstrate that MT learning surpasses conventional learning models by integrating diverse data sources and extracting meaningful correlations, particularly in data-scarce scenarios. Furthermore, the inclusion of diverse property data in this approach substantially broadens the coverage of the chemical space and effectively addresses the ML extrapolation problem. This is an ongoing process though, one that can lead to continuous improvement as more data becomes available. We then performed a head-to-head comparison of our new MT model against our previous, then state-of-the-art gas permeability predictor, deployed at Polymer Genome (<https://www.polymergenome.org>)¹¹, making predictions across 13 polymer classes and demonstrating the superiority of the present model.

Finally, we highlight the power of the present development in the realm of materials discovery. Robeson-type trade-off plots are created for gas permeability, diffusivity, and solubility (by pairing each with selectivity), for over 13,000 known (i.e., previously synthesized) polymers. These trade-off plots reveal interesting candidates, as well as the true property limits across the known polymer chemical space. Most importantly the limitations of the present model (in terms of recognizing chemical spaces where the model is uncertain) are also revealed.

By integrating high-throughput simulation data with available measured data and employing data fusion techniques, one can progressively enhance the accuracy and generalizability of predictions. This philosophy and strategy holds the potential to advance polymer discovery not only for membrane technology but also for other applications.

Results

Experimental data acquisition

Measured gas transport properties (permeability, diffusivity, and solubility) for six different gases (CO₂, CH₄, O₂, N₂, H₂, and He) were obtained from 84 publications listed in the Polymer Handbook³². The experimental testing temperatures ranged from 25 °C to 35 °C, and testing pressures varied between 1 and 30 atm. The dataset comprised a total of 820 polymers and included 3748, 709, and 550 P_{expt} , D_{expt} and S_{expt} values, respectively, amounting to a

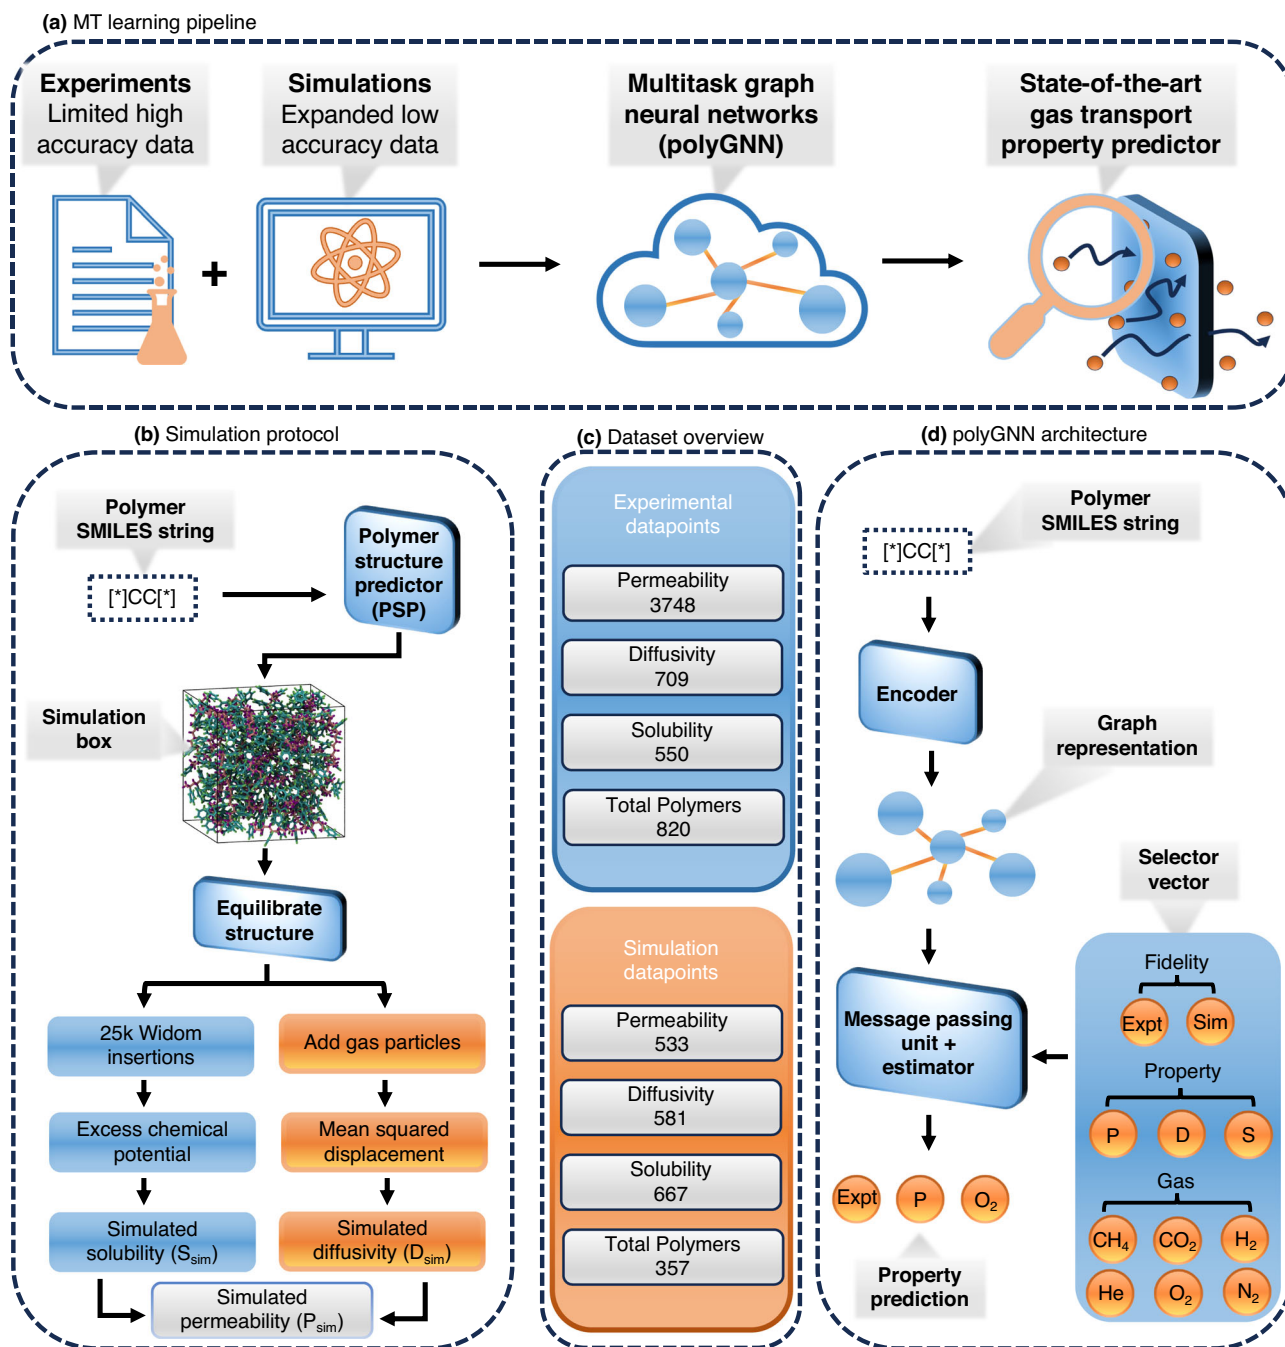


Fig. 1 | Multi-task learning approach for predicting gas transport properties in polymers. **a** MT learning pipeline. Our innovative multi-task learning approach employs the fusion of experimental and simulation data, harnessed through the power of polyGNN, a graph neural network architecture, to construct a state-of-the-art predictor for gas transport properties **b** Simulation protocol. The process begins with a polymer SMILES string³¹, from which the Polymer Structure Predictor (PSP) package³⁶ constructs a simulation box. This box undergoes a 21-step equilibration procedure³⁷. Subsequently, the equilibrated structures serve as the starting point for gas diffusivity and solubility calculations, accomplished through molecular dynamics and Monte Carlo simulations, respectively. Gas permeability is determined by the product of the simulated gas diffusivity and solubility. **c** Dataset

overview. Curated experimental and simulation data used for training the multi-task ML models. **d** polyGNN²² architecture. A method based on graph neural networks is initiated with a polymer SMILES string. The encoder converts the repeat unit SMILES string into a periodic graph along with fingerprints, followed by the computation of initial atomic and bond fingerprint vectors. Subsequently, the message passing unit generates the learned polymer fingerprint. Introducing a selector vector to convey data fidelity (experimental or simulation) and specific properties (permeability, diffusivity, solubility) for six gases, the approach then combines this fingerprint and selector vector before passing it to the estimator, resulting in the prediction of the desired property.

total of 5007 data points. Factors such as polymer process history and testing method were not directly included as parameters. Instead, the measured P_{exp} , D_{exp} and S_{exp} values are treated as samples from the distribution of possible values for a given polymer. As such, it is important to consider the uncertainty in the predictions, and not just the mean value of the prediction.

Molecular Dynamics and Monte Carlo simulations

Gas diffusivity and solubility data were generated using classical molecular dynamics (MD) and Monte Carlo (MC) simulations, respectively. These simulations were conducted using the open-source large atomic molecular massively parallel simulator (LAMMPS) package³³. The

atomic potential parameters for polymers were adopted from the general AMBER force field 2 (GAFF2)³⁴. In the simulations, the gas molecules (i.e., CO₂, CH₄, O₂, and N₂) were treated as rigid molecules, and thus were modeled with non-bonded potentials described by the TraPPE (transferable potentials for phase equilibria) models³⁵. To perform the simulations, 27 polymer chains were inserted into the simulation box, with each chain comprising of approximately 150 atoms, and their ends were capped with a methyl group. The initial polymer configurations were generated using the Polymer Structure Predictor (PSP) package³⁶, and a representative snapshot is shown in Fig. 1b. To achieve equilibrated structures, all systems underwent a 21-step relaxation procedure as recommended by Abbott et al.³⁷. The mean-squared displacement (MSD) of the polymers was then computed, and movement beyond a few times the distance of the radius of gyration on average was assessed. This step ensured that the polymers explored various conformations and reached an equilibrium conformational state and density. Once the equilibrated structures were obtained, D_{sim} and S_{sim} were calculated. The simulation protocol is outlined in Fig. 1b.

For the D_{sim} calculations, a total of 27 gas molecules were randomly added to the simulation box. This specific number of molecules was chosen to be small enough to maintain the system in the dilute Fickian regime such that the gas molecules do not significantly influence each other, and yet large enough to obtain meaningful statistics. Subsequently, all systems underwent an additional equilibration of 10 ns in the NPT ensemble, followed by a 100–200 ns production run in the NVT ensemble. The choice of a 100–200 ns production run duration was made to ensure the convergence of gas diffusivity and gas MSD slope across a broad spectrum of polymer types. While shorter time frames are adequate for certain instances, there are cases where the extended range of 100–200 ns is necessary to achieve the desired level of convergence. To illustrate this behavior, we present an analysis of simulation time versus methane diffusivity for polyethylene, polyimide, polystyrene, and polymethyl methacrylate, with the results detailed in Supplementary Figure S1. The box size in the NVT run was fixed using the average spacing and density obtained from the last 1 ns of the NPT run. Nosé-Hoover thermostat and barostat were employed with a damping parameter of 100 time steps for each, and a time step of 1 fs was used in all MD simulations. The barostat coupled the three dimensions of the box to maintain a cubic box for all systems. Simulation outputs were saved every 1000 fs and block averaging from one polymer configuration was used to calculate an average D_{sim} and standard deviation from the gas MSD. Block averaging allows for the reduction of random noise and more reliable statistical measures³⁸.

For the S_{sim} calculations, a 5 ns production run was performed on equilibrated structures in an NVT ensemble. During this 5 ns run, a snapshot of the structure was captured every 100 ps, resulting in a total of 50 snapshots. Employing an ensemble of snapshots allows for improved sampling and a standard error, which is crucial for accurate estimation of S_{sim} ³⁹. Using a built-in LAMMPS function (https://docs.lammps.org/fix_widom.html), 25,000 gas particles were inserted per snapshot, at random positions, following the Widom insertion method⁴⁰. This method involves determining the excess chemical potential resulting from the insertion of gas molecules into the polymer, which allows for the estimation of Henry's constant. Henry's constant indicates how easily a particular gas dissolves in the polymer. Henry's Law is then used to obtain gas solubility from Henry's constant, with an assumption of a partial pressure equal to 1 atm, which is the IUPAC standard testing condition⁴¹. This derivation is detailed in the methods section. No relaxation was performed to adjust the positions of the polymer atoms or the gas particles during the insertion process. Langevin thermostat was used with a time step of 1 fs for all MC simulations. 25 polymer configurations were used to calculate the S_{sim} standard deviation, and the standard error from the excess chemical potential.

Figure 1b provides an overview of the simulation protocol used, and details of D_{sim} estimation from gas MSD and S_{sim} from the excess chemical potential are described in the Methods section.

Validation of MD and MC simulations

As an essential step of this investigation, we aimed to validate and calibrate the accuracy of the MD and MC predictions and assess the extent to which the simulations capture trends in gas transport properties. Performing classical simulations with a specific force field for polymer-gas systems across extensive chemical spaces to estimate gas diffusivity and solubility is a relatively rare endeavor. While generic force fields like GAFF2 are designed for a wide variety of materials, they often require fine-tuning of potential parameters for each unique material to attain better accuracy.

A total of 584 polymer-gas systems were simulated, out of which 342 systems had corresponding experimental measurements. The additional simulated systems were intended to expand the chemical coverage of the model. A comparison of P_{sim} , D_{sim} , and S_{sim} against their respective experimental values, P_{expt} , D_{expt} , and S_{expt} is illustrated in Fig. 2. Overall, the simulations tend to overestimate the measured values, but they effectively capture the general trends across the polymer-gas chemical space considered. More specifically, the simulated polymer systems often exhibit lower densities compared to experimental systems, as modeled systems are approximations of the real polymeric materials and may include lower molecular weights and limited equilibration times. In our methodology, we employ a 21-step polymer equilibration relaxation procedure, which results in consistent density trends compared to experimental systems. However, a slight underestimation of density remains, as also observed by Abbott et al. in their study employing the same procedure³⁷. This increased free volume allows gas molecules to move more easily and quickly through the polymer system, resulting in higher diffusivity.

Similarly, the discrepancies of S_{sim} relative to S_{expt} may be due to the approximations inherent to the Widom insertion approach and the quality of the classical force fields across chemical spaces. Nonetheless, the favorable trends that the force fields can capture provide optimism for the usage of such simulation-derived datasets, albeit with lower fidelity, in multi-task learning frameworks. Another essential aspect of the validation is the derivation of P_{sim} from the product of D_{sim} and S_{sim} using Eq. 1. While non-equilibrium MD can be used to simulate P_{sim} , it requires a more complex setup and can be computationally intensive.

Multi-task model benchmark

To elucidate the effect of data fusion, we train and compare both ST and MT polyGNN models, using a subset of the experimental data collected and simulation data generated. These models were evaluated based on the predictive accuracy of P_{expt} using various holdout train and test splits of 293 systems (comprised of 80 unique polymers with varying available gas data). For instance, in a 20/80 split, 20% of the P_{expt} data is set aside as testing data, while 80% is used to train the model. To ensure representative data sampling, stratified sampling based on polymer SMILES was used when splitting the data into train and test sets. In this type of sampling, when a polymer is selected for the test set, all gas data for P_{expt} associated with that polymer are withheld from the training set. This also provides insight into how well the model extrapolates to new unknown polymers. The polyGNN model training parameters used are detailed in Supplementary Table S1.

In Fig. 3, we illustrate the two model types, ST and MT, along with the details of the train and test splits. The performance of the models was evaluated using two key metrics: the coefficient of determination (R^2) and the order of magnitude error (OME)—units in Barrer. R^2 assesses how well a model predicts an outcome, while OME quantifies the prediction error in terms of orders of magnitude (taken as the logarithm of the mean absolute error). We conducted four random seed selections of the training and test sets to compute the statistics of the model performance.

Our MT learning methodology comprises two primary components: the integration of simulation data and the inclusion of correlated experimental data. To establish a baseline for comparison, we employ a ST model. Shown in Fig. 3a and represented by the “ST” row in Fig. 3b, the ST model is exclusively trained using P_{expt} data. Due to its reliance on limited data and the absence of diverse property inputs, the ST model's coverage of the chemical space is inherently constrained. As the test set percentage

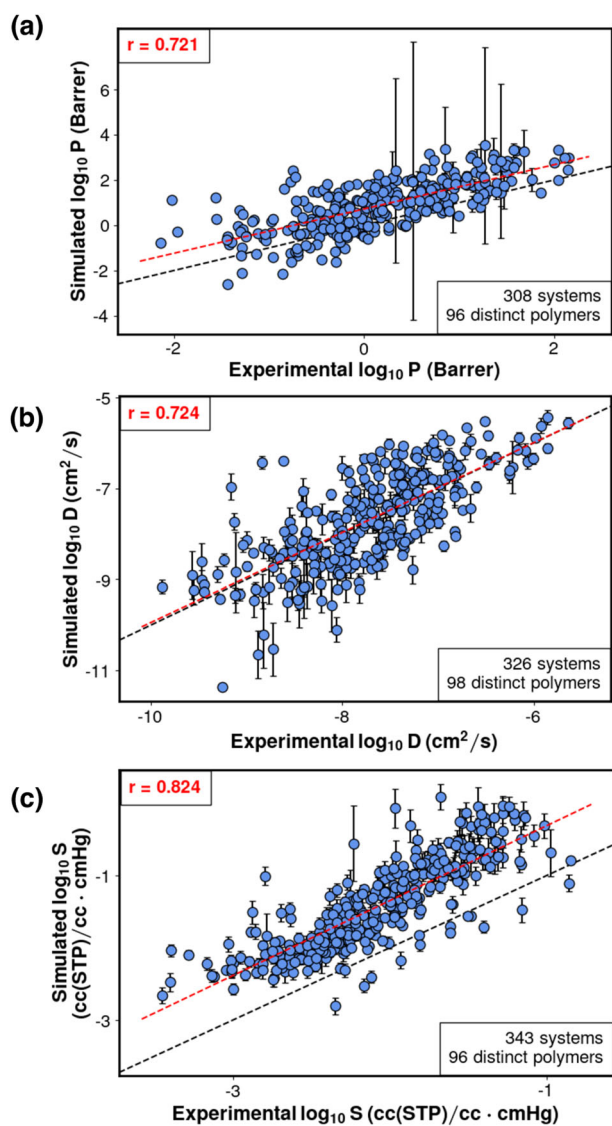
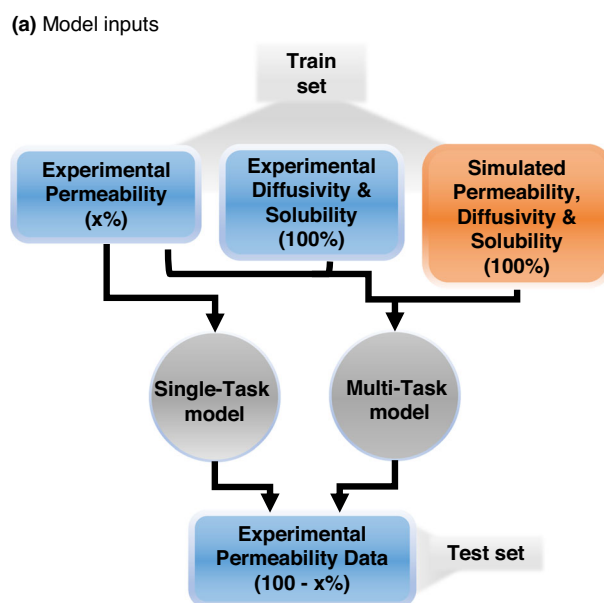


Fig. 2 | Validation of gas transport simulations. **a** Gas permeability parity plot, **b** Gas diffusivity parity plot, and **c** Gas solubility parity plot. Parity plots comparing the results from simulations against experiment data. Simulated gas permeability was derived using Eq. 1, using simulated gas diffusivity and solubility as inputs. The red lines represent trends in predicted data, while the black lines depict the parity lines of optimal fit. The error bars for all plots are represented in standard deviations. Error propagation techniques were employed to calculate the error bars for gas permeability. While some overestimation is expected across all cases, a qualitative correlation is demonstrated.

increases, this model is trained on progressively reduced amounts of data, leading to an anticipated decrease in predictive performance. This trend is evident in Fig. 4 where the R^2 decreases and the OME increases as the ST model is trained on diminishing data portions. In the most challenging scenario (80% test set size), the R^2 dropped to less than 0.50 and the OME increased to ≈ 0.44 Barrer.

Now let's consider the first element of our MT learning approach, specifically the augmentation of P_{expt} training data with P_{sim} , represented by the "MT-1" row in Fig. 3. The MT-1 model is enriched with simulation data spanning the test set space. Its primary purpose is to exploit the correlations between measured and simulated data learned from the training set. This scenario mirrors situations where experimental data is unavailable, and simulation data is introduced to guide the model's predictions. Upon examining the MT-1 model, its performance noticeably surpasses that of the baseline ST model. The MT-1 model achieves an average R^2 and OME



(b) Benchmark models

Model \ Data	P_{expt}	D_{expt}	S_{expt}	P_{sim}	D_{sim}	S_{sim}
ST	✓					
MT-1	✓			✓		
MT-2	✓	✓	✓			
MT-3	✓	✓	✓	✓	✓	✓

Fig. 3 | Comparison of single task and multi-task learning models. **a** Model inputs. This schema illustrates the train and test splits for two model variants: Regular Single Task (ST) and data-fused Multi-Task (MT) models. In the ST models, solely experimental gas permeability data is incorporated. Conversely, the MT model encompasses a possible amalgamation of experimental gas diffusivity and solubility, along with simulated gas permeability, diffusivity, and solubility data. **b** Benchmark models. Four distinct models were developed to assess the impact of MT learning. The first model (ST) exclusively incorporated experimental gas permeability data. In contrast, the subsequent MT models progressively integrated additional data. The presence of each data type in the model is indicated by a black checkmark. Here, P , D , and S represent gas permeability, diffusivity, and solubility, respectively. The abbreviation *expt* corresponds to experimental data, while *sim* signifies simulation data.

of ≈ 0.77 and ≈ 0.30 , respectively, as shown in Fig. 4 (MT-1). This improvement is particularly pronounced when the test set size reaches 80%, where the coverage of experimental data within the chemical space is most limited. This accentuates the ability of data fusion models, reinforced with simulation data, to effectively mitigate the challenges of extrapolation that conventional models (trained solely on a single experimental property) would inevitably confront. Furthermore, as another demonstration, this analysis was extended to experimental and simulation data for gas diffusivity, resulting in a similar strengthening in performance, as illustrated in Supplementary Figure S2. This observation underlines the value of bolstering experimental data with simulation data, indicating its potential extension to other properties of interest as well.

Moving on to the second component of our MT learning methodology, we focus on augmenting the P_{expt} training data with D_{expt} and S_{expt} , represented by the "MT-2" row in Fig. 3b. The inclusion of this supplementary data serves the purpose of empowering the model to leverage knowledge from other available pertinent properties and established physics and make predictions for the P_{expt} values. In this scenario, a remarkable enhancement is observed, leading to a significant boost in predictive performance. Specifically, the average R^2 and OME is




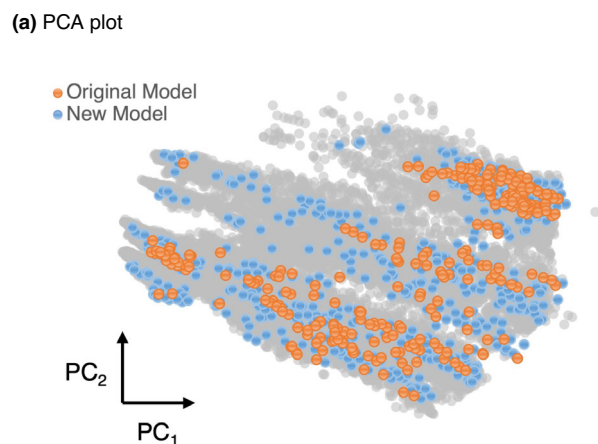
Fig. 4 | Predictions of Permeability at various train test splits. **a** Coefficient of determination (R^2). **b** Order of magnitude error (OME). R^2 evaluates the predictive performance of a model, whereas OME measures the prediction error by considering orders of magnitude, represented as the logarithm of the mean absolute error. The ST and MT models are compared based on varying percentages of the unseen test set. The different test set sizes illustrate the impact of reducing training data. At 80%, the model is trained on only 20% of the dataset and tested on the remaining 80%, reflecting a data-scarce region with limited chemical coverage. Comparatively, the MT models show significant improvement over the ST model, particularly at higher percentages of the unseen test set.

≈ 0.93 and ≈ 0.12 , respectively, as displayed in Fig. 4 (MT-2). Comparing the MT learning component in the previous passage with this second component reveals a notable difference in performance. While both approaches expand the coverage of the chemical space, MT-2 stands out due to the incorporation of high-fidelity experimental data. Unlike MT-1, where all augmented data comes from simulation, the new information in MT-2 originates from additional experimental sources, contributing to superior predictive capabilities. The MT-2 model can be likened to an ideal scenario where complementary or correlated high fidelity data is readily available. In scenarios where such ideal conditions are not met, the MT-1 approach excels by effectively integrating simulation data to achieve a respectable level of prediction accuracy.

In our final model, we combine the strategies embedded in both the MT-1 and MT-2 models, creating a unified model represented by row “MT-3” in Fig. 3b. This comprehensive model encompasses all available experimental and simulation data points. The performance of the MT-3 model slightly outperforms that of the MT-2 model, exhibiting an elevated average R^2 of ≈ 0.96 and a comparable average OME of ≈ 0.10 , as depicted in Fig. 4 (MT-3). Overall, this model achieves superior performance compared to the base ST model, which had an average R^2 and OME is ≈ 0.57 and ≈ 0.38 , respectively. These results establish the efficacy of integrating simulation and correlated experimental data in successfully addressing the challenges posed by ML extrapolation.

Production model benchmark

In the first iteration of our gas permeability prediction work, deployed at Polymer Genome (<https://www.polymergenome.org>), a Gaussian process



(b) Dataset comparison

Dataset of Original Polymer Genome model and new model		
Data Type	Original Model	New Model
Experimental Gas Permeability	1501	3748
Experimental Gas Diffusivity		709
Experimental Gas Solubility		550
Simulated Gas Permeability		533
Simulated Gas Diffusivity		581
Simulated Gas Solubility		667
Total datapoints	1501	6788
Polymers	315	1050

Fig. 5 | Comparison of chemical space and data coverage by the original and production models. **a** Principal Component Analysis (PCA) plot. The PCA plot demonstrates an expanded coverage of chemical space by both the original and production models. The orange and blue dots correspond to the coverage of the original and production model, respectively, while the grey dots represent the 13,000 known polymers in our database. **b** Dataset comparison. A comparison between the original and production models reveals an incorporation of diverse data types. The production model integrates experimental and simulation data for permeability, diffusivity, and solubility properties.

regression algorithm was employed alongside a hierarchical polymer fingerprinting scheme to train a ST model¹¹. In the present work, a transition is made to polyGNN (a recently published Graph Neural Network model that automatically generates fingerprints from SMILES strings), data augmentation, and invariant transformations to train a MT model. The models presented in the preceding section were trained using a subset of our dataset, a deliberate choice made to clearly illustrate the impact of incorporating diverse data types on prediction capabilities in a multi-task setting. Our final production model adopts the MT-3 model scheme and now incorporates all the available experimental and simulation data for gas permeability, diffusivity, and solubility. With this latest model iteration, our objective is to achieve substantial improvements over the previous version and to push the boundaries of transport predictions through polymers. The principal component analysis (PCA) plot in Fig. 5a, created using Polymer Genome fingerprints, displays the chemical space of the present study against 13,000 known polymers in our database. This plot visually demonstrates our production model's expansion to include additional chemical compositions, increasing the range of polymers for which the model can make accurate predictions. We also present a PCA plot in Supplementary Figure S3, illustrating the chemical coverage of our simulation data in comparison to experimental gas transport data and 13,000 known polymers in our database. Figure 5b highlights the data fusion aspect of the model, showcasing the contrast between the datasets employed in the original and current models. Particularly noteworthy is the considerable enlargement of our dataset, expanding from 315 to 1050 polymers, accompanied by a significant increase in the total number of data points. With this amplified dataset, our

Table 1 | Benchmarking model performances

Polymer Class	Original Model R ²	Production Model R ²
Conjugated Polymers	0.99	0.99
High Temperature Polymers	0.94	0.99
Parylenes	0.89	0.98
Poly(aryl ethers) & Poly(aryl ether ketones)	0.92	0.98
Polyamides & Poly(amide-imides)	0.96	0.99
Polyarylates	0.95	0.97
Polycarbonates	0.75	0.98
Polyimides & Polypyrrolones	0.97	0.92
Polynorbornenes	0.51	0.96
Polyphosphazenes	0.49	0.90
Polypropynes, polyacetylenes, polypentynes	0.56	0.87
Polysulfones	0.80	0.98
Vinyl & Vinylidene Polymers	0.77	0.96
Overall R ²	0.93	0.95

A comparative analysis of the predictions of the original and production models across 153 systems, encompassing 31 polymers and 13 polymer classes. The production model integrates data fusion, integrating multiple data sources and multiple properties, whereas the original model relies solely on experimentally measured gas permeability. Bold values signify the model with a superior R².

model gains the capability to not only predict gas permeability but also include gas diffusivity and solubility. This broader scope of predictions reflects the power of our MT learning approach and its ability to leverage diverse data sources for a more comprehensive understanding of gas transport properties through polymers.

To highlight the superior performance of the present model, P_{expt} predictions were made on a holdout test set of 153 systems, consisting of 31 polymers across 13 polymer classes, following a similar approach as the original model¹¹. The summarized outcomes are presented in Table 1, and the specific polymers selected for this assessment are listed in Supplementary Table S2. The overall R² has increased from 0.93 to 0.95 in the updated model compared to the original. Upon a more detailed examination of individual polymer classes, it becomes evident that the R² metric exceeds 90% for all classes for the production model, with a particularly significant enhancement observed for polyphosphazenes, where the R² value has risen from 0.49 to 0.90. Additionally, substantial advancements have been achieved in prediction accuracy for polymers such as polynorbornenes, polypropynes, substituted polyacetylenes, and polypentynes. The diminished performance of the original model in these cases could be attributed to either limited data availability for certain polymer classes or inherent uncertainties within the experimental data. Importantly, it should be noted that the test data points for these specific polymer classes vary widely, ranging from 4 to 53 data points. This variability in data availability across diverse classes could potentially contribute to lower individual R² values for specific classes while concurrently contributing to a higher overall model R². Nonetheless, the updated model effectively overcomes these performance variations, highlighting its robustness and versatility. Further insights into the model's performance are depicted in parity plots showcasing train and test set predictions for the 31 evaluated polymers, shown in Supplementary Figures S4 and S5.

Forward-looking design

The ideal performance of gas separation membranes is related to two intrinsic material properties: the gas permeability and the permselectivity between specific target gas pairs. Ideally, a membrane would provide high permeability and permselectivity to maximize throughput and minimize costs. In 1991, Robeson⁴² documented a trade-off relationship between these

two characteristics for polymers, often referred to as “the upper bound”. This principle asserts that polymers with high permeability typically exhibit diminished selectivity, and vice versa. These upper bounds illustrate the trade-off relationship for pairs of common gases (CO₂, CH₄, O₂, N₂, H₂, and He), highlighting the best possible combination of permeability and permselectivity. This upper bound establishes a comparative benchmark for evaluating the performance metrics when designing novel membranes. As such, data driven methods that establish a relationship between polymer structure and polymer membrane performance hold immense potential in accelerating the design of tailor-made polymers for specific separation tasks.

To this end, we demonstrate our model's capability to make these assessments. We constructed permeability trade-off plots for ≈13,000 known polymers (i.e., previously synthesized) for the gas pairs; CO₂/CH₄, CO₂/N₂, H₂/CH₄, H₂/CO₂, O₂/N₂, and N₂/CH₄. Figure 6a shows the permeability trade-off plot for CO₂/CH₄, while the other gas pairs are shown in Supplementary Figure S6. The ML predicted gas pair permeability and selectivity closely align with the available experimental data and the bounds, while simulation data over-predicts as expected. Both experimental and simulation data are also shown in Fig. 6a. By predicting property values for the ≈13,000 known polymers, we can gain a clearer understanding of the overall trade-off behaviors. Robeson's upper bound, initially established in 1991, is presented alongside updated bounds introduced in 2008 and 2019^{43,44}. PIM-DM-BTrip, a polymer with superior performance, is highlighted as a part of the set of polymers that helped define the 2019 bound.

Research endeavors commonly focus on permeability trade-off plots, however as permeability can be broken down into diffusivity and solubility components, we also created CO₂/CH₄ trade-off plots for these properties, as shown in Fig. 6b, c. Diffusivity and solubility trade-off plots for CO₂/N₂, O₂/N₂, and N₂/CH₄ are illustrated in Supplementary Fig. S7 and S8. When using these models, the sensibility of predictions can be evaluated by observing common trends for the properties. For example, gas diffusivity tends to follow the relationship of $D_{O_2} > D_{CO_2} > D_{N_2} > D_{CH_4}$, a pattern primarily driven by the molecular diameter effects. However, Fig. 6b, illustrates instances where the CO₂/CH₄ diffusivity selectivity falls below 1 (i.e., below 0 in the log scale). This contradicts the intuition that CO₂ diffusivity should almost always be greater than CH₄. Although there are cases where $D_{CO_2}/D_{CH_4} < 1$, it is a rare occurrence. A closer examination of these suspicious predictions reveals that most of them fall in the lower diffusivity regime. In this regime, prediction uncertainty, calculated using Monte Carlo dropout, tends to be inflated, revealing lower confidence for predictions. This heightened uncertainty can be directly attributed to the scarcity of data in this specific range, a challenge that is particularly pronounced in both simulations and experimental measurements. Indeed, as can be seen from Fig. 6b, there are no measured data and only sparse simulation data points in this property range; hence, the ML predictions must be viewed with extreme caution and suspicion. This underscores the importance of recognizing that these models are valuable tools, but they must be used in conjunction with chemical intuition and an understanding of prediction uncertainties, especially for predictions in regions far away from the chemical space of the training set. This becomes especially critical when assessing areas with limited data or when venturing into new domains. These considerations thus mandate either experiments or simulations in such unexplored chemical spaces to better inform the ML models.

Trade-off plots are typically employed in designing amorphous polymers for gas separation, but when considering other applications such as packaging, the degree of crystallinity of the polymer must be considered. Gas transport behavior in semi-crystalline polymers varies due to the crystalline regions acting as impermeable barriers against gas penetration⁴⁵. Michaels et al. originally described this behavior using a two-phase model, which comprises a crystalline phase and an amorphous phase, where impedance is directly proportional to crystallinity⁴⁶. Weinkauff et al. extended this model into a three-phase model that incorporates the ratio between the rigid amorphous phase fraction and mobile amorphous phase fraction (RAF/MAF)⁴⁵. These models provide critical insights into the behavior of semi-crystalline polymers and offer guidelines for tailoring their gas transport

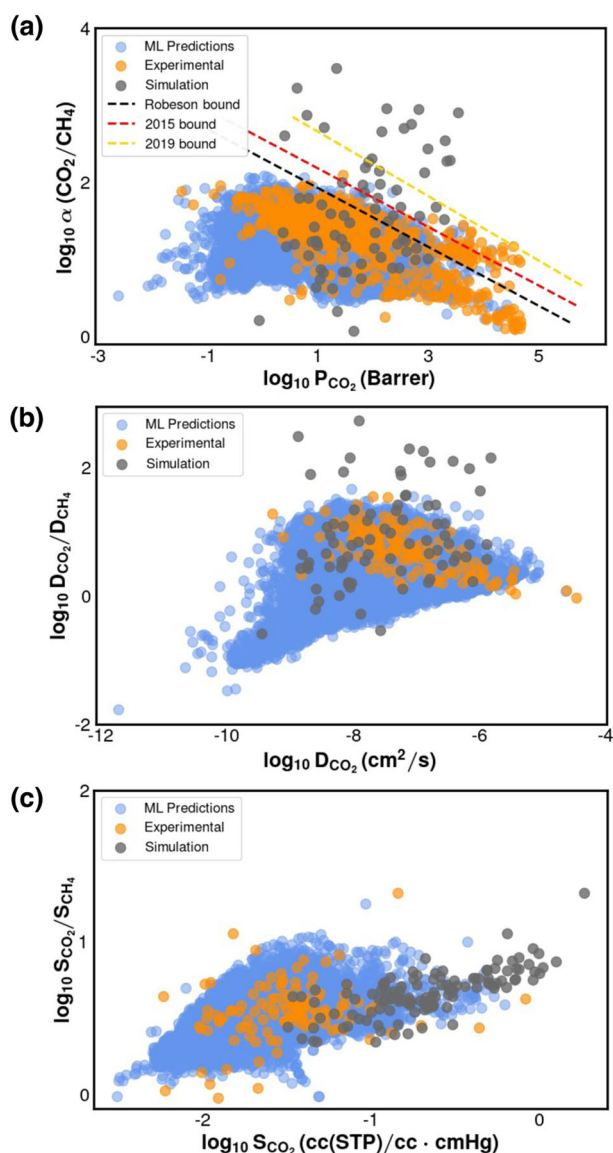


Fig. 6 | Gas transport trade-off plots. **a** Gas permeability, **b** Gas diffusivity, and **c** Gas solubility. Using our model to predict $\approx 13,000$ known polymers, we compare the results to experimental data. The orange and blue dots correspond to the experimental data and machine learning predictions, respectively, while the grey dots represent simulated values. The original Robeson upper bound (1991) is depicted as the dashed black line, with reevaluated bounds from 2008 and 2019 shown as red and yellow dashed lines, respectively, for gas permeability. There are no established bounds for gas diffusivity or solubility, but the model predictions closely align with the experimental data values. In the case of CO_2/CH_4 diffusivity selectivity, the low diffusivity regime has high prediction uncertainty and should be taken with caution.

properties to specific applications. The present work may be extended to address such practical situations by planning simulations of gas diffusivity and solubility through amorphous, crystalline, and amorphous-crystalline interfaces.

Discussion

In this study, we introduce a novel multi-task (MT) learning approach that leverages a combination of measured and simulation data, along with correlated properties to create a state-of-the-art predictor for gas transport properties. To thoroughly evaluate the effectiveness of this approach, we performed a benchmark study in which we compared the individual

impacts of each of these tasks and their collective effect when considered together. It was revealed that the addition of interrelated measured data led to a bigger benefit in enhancing the predictive capabilities of the ML model. The situation indicates that multiple correlated ground truth (i.e., measured) property data are most desirable to generate accurate property forecasts. However, in instances where rich measured data is unavailable, easily producible simulation data, when combined with measured data, demonstrates its potential by offering informed predictions. In any case, both scenarios of MT learning were able to learn underlying physical correlations and are superior to single-task (ST) models that have a less robust basis for predictions.

These ideas have been unified to create a model representing a major advancement in predicting gas transport properties through polymers. This model in a comparative analysis with the prior work displayed concrete improvements, across 13 different polymer classes. Using our new ML model we also generate selectivity trade-off plots for gas permeability, diffusivity, and solubility for $\approx 13,000$ known polymers (i.e., previously synthesized). These plots provide insights into the strengths and limitations of the models, but more importantly, the need for data across diverse chemical spaces, e.g., via simulations if measured data proves to be laborious to generate. The prospect of continual expansion of the accessible polymer universe will push the frontiers of what is achievable in terms of properties and performance.

Methods

Gas diffusivity calculation in MD simulations

From our MD simulations, the diffusivity (D_{sim}) of gas molecules was obtained by:

$$D_{sim} = \frac{1}{6N_{gas}} \lim_{t \rightarrow \infty} \frac{d}{dt} \sum_{i=1}^{N_{gas}} \langle \Delta r_i(t)^2 \rangle \quad (2)$$

where N_{gas} is the number of gas molecules in the simulation cell, t is the simulation time, $r_i(t)$ is the position of the gas molecule i at time t , $\Delta r_i(t) = r_i(t) - r_i(0)$, is the displacement of gas i between time 0 and time t , and $\langle \Delta r_i(t)^2 \rangle$ is the mean square displacement (MSD) of gas molecule i at time t . The gas MSD was block averaged over 2-5 non-overlapping trajectories depending on the system dynamics (depending on whether breaking the trajectory into shorter blocks still allows its MSD curve to reach the Fickian regime). This block averaging was also used to calculate the standard deviation. The diffusivity was obtained using the slope from a least-squares linear fit of the final decade of the MSD data. The log-log slopes of all the systems' MSD curves are in the range of 0.95–1.05, which is a common range for classical MD simulation diffusivity studies^{47–49}.

Gas solubility calculation in MC simulations

The Widom Insertion method was used to calculate the Henry's constant (k) of gas molecules within a Monte Carlo simulation. In this method, we insert N gas molecules into the simulation box (one at a time, at various random locations), and the excess chemical potential (μ_{ex}) of the gas in the membrane is obtained. The estimation uses an ensemble average of the N separate, random, insertions, the i th of which will change the internal energy of the system by ΔE_i .

$$\mu_{ex} = -k_b T \ln \left\langle \exp \left(\frac{-\Delta E_i}{k_b T} \right) \right\rangle_N \quad (3)$$

where k_b is Boltzmann constant, T is temperature and $N = 25,000$ insertions. This estimation is defined for the dilute limit, where there are no gas-gas interactions. When an insertion overlaps with the polymer in the system, the μ_{ex} results in minimal contribution towards gas solubility, and thus these insertions with energies greater than $5k_b T$ are discarded. With μ_{ex} , we can obtain Henry's constant, k , of gas

molecules by:

$$k = \exp\left(\frac{\mu_{ex}}{k_b T}\right) \quad (4)$$

With k , we then use Henry's law which states that the solubility (S) of a gas is directly proportional to the partial pressure of the gas (P_{gas}), which takes the form:

$$S = kP_{gas} \quad (5)$$

We assume a standard testing condition of partial pressure equal to 1 atm, and thus:

$$S = k \quad (6)$$

In this work, for each polymer configuration, we take $m = 50$ dynamic snapshots from a 5 ns production run on the equilibrated structure in an NVT ensemble. An average solubility, S_i of the configuration is obtained by:

$$S_i = \frac{1}{m} \sum_{j=1}^m S_{ij} \quad (7)$$

For proper convergence and as a measure of standard error, we use $n = 25$ polymer configurations and S_{sim} is obtained by:

$$S_{sim} = \frac{1}{n} \sum_{i=1}^n S_i \quad (8)$$

The standard deviation, σ , and standard error, SE, are then calculated as:

$$\sigma^2 = \frac{1}{n} \sum_{i=1}^n (S_i - S_{sim})^2 \quad (9)$$

$$SE = \frac{\sigma}{\sqrt{n}} \quad (10)$$

Final gas solubility values were screened for a SE of less than 5%.

polyGNN

The predictive model we used was polyGNN, a multitask graph neural network method that has shown promising results when dealing with large-scale multi-property datasets²². Briefly, polyGNN contains three modules: the Encoder, Message Passing Block, and the Estimator. The inputs to polyGNN are a polymer repeat unit and a property of interest (or, equivalently, the property's associated selector vector). The two outputs of a polyGNN model are the repeat unit's fingerprint and the value of the property of interest. In the Encoder, the repeat unit is first converted to a periodic graph, with each atom as a node and each bond as an edge. Then, each node and edge in the graph are given an initial fingerprint. After the graph elements have been assigned their initial features, the graph is passed to the Message

Passing Block. Messages between neighboring atoms are iteratively passed along chemical bonds. After each iteration, every node fingerprint is updated using the messages, while each bond fingerprint remains the same. The message passed from atom j to atom i at time step k is calculated according to Eq. 11.

$$m_{i,j}^{(k)} = \phi^{(k)}\left(x_i^{(k)}, x_j^{(k)}, e_{i,j}\right) \quad (11)$$

where each $\phi^{(k)}$ is a parameterized function, $x_i^{(k)}$ and $x_j^{(k)}$ are the encodings of neighboring ij atoms after time step k , and $e_{i,j}$ is the fingerprint of the bond that joins atoms i and j . $m_{i,j}^{(k)} = 0$ if i, j do not share a chemical bond. After initialization, each node receives messages from all of its neighbors. These

messages are aggregated by some permutation-invariant function f (e.g., sum, mean, max). We use the sum in this work. The aggregated message, along with the current node encoding, is used to update the node encoding. The node update process is defined in Eq. 12.

$$x_i^{(k)} = \chi^{(k)}\left(x_i^{(k-1)}, f\left(\left\{m_{i,j} \forall j \in [1, N_p]\right\}\right)\right) + x_i^{(k-2)} \quad (12)$$

where each $\chi^{(k)}$ is a parameterized function, p is a polymer, $[1, N_p]$ is the set of integers between 1 and N_p , N_p is the number of atoms in the repeat unit of p , and $x_i^{(k)} = 0, \forall k < 0$. Messages are passed for τ time steps, where τ is also the capacity in this work. The fingerprint of the entire polymer, x_p , is calculated by the graph aggregation function A_g , as shown in Eq. 13.

$$x_p = A_g\left(x_i^{(\tau)}, x_i^{(0)}\right) = \frac{1}{N_p} \sum_{i=1}^{N_p} x_i^{(\tau)} + x_i^{(0)} \quad (13)$$

Finally, x_p and the selector s can be passed to the Estimator. Here, these inputs are mapped to a polymer property prediction, y_p , via a parameterized function ψ , which represents the multilayer perceptron (MLP) depth, as shown in Eq. 14.

$$y_p = \psi\left(x_p, s\right) \quad (14)$$

ψ specifies the number of hidden layers between the input and output layers, with this depth parameterized to range from 2 to 14 layers. During training, the parameters of all $\phi(k)$, $\chi(k)$, ψ are learned simultaneously. As shown in Eq. 12, our update step leverages skip connections, which have been shown to improve the optimization of shallow layers in deep neural networks.

All neural network architectures used dropout layers, fully connected layers, and Leaky ReLU activations (with a negative slope equal to 0.01). MC dropout was implemented by performing 10 forward passes through the network, each time applying dropout to different subsets of nodes. All architectures were created using PyTorch and PyTorch Geometric. The weights of all models were optimized using the Adam optimizer and the mean squared error loss function.

Training procedure

The training procedure used is similar to that in the polyGNN work, where the models are ensemble models, composed of several submodels²². The output of the ensemble is computed by the average of each submodel's output. The data used for training was grouped based on gas transport type (P , D , & S), gas type, and data source (experiment or simulation). Once grouped, each data subset was then min-max scaled between 0 and 1. The polyGNN model training parameters used are detailed below and also compiled in Supplementary Table S1.

Next, the entire data set was stratified and split into training and test sets (percentages of test sets were 20%, 40%, 60%, or 80%) based on polymer SMILES strings three times. Using the NNDebugger package⁵⁰, the optimal capacity was found by attempting to overfit ($R^2 > 0.97$) the entire training data set. If the data was not overfit, then the capacity corresponding to the highest R^2 value was used. The capacity range considered was between two and fourteen. The training data set was then divided into an 80% hyperparameter (HP) training set and a 20% HP validation set. The remaining HPs (batch size, learning rate, dropout percentage) were optimized using the package scikit-optimize⁵¹. The set of HPs corresponding to the lowest RMSE on the HP validation set was considered optimal.

Finally, the training data set was split into five folds using cross-validation (CV), producing one CV train data set and one CV validation data set per fold. For each fold, the model's HPs were fixed as the optimal HPs and the model's learnable parameters were fit to the CV train data set

for 1000 epochs. At the end of 1000 epochs, the model parameters corresponding to the epoch with the lowest RMSE in the CV validation data set were chosen. After all five models were trained on their respective CV splits, the models were placed in an ensemble. The ensemble was used to make predictions of the test set, that were completely unseen by the ensemble during HP optimization or model training with CV.

Data availability

The experimental sources of data used are reported in the paper. All data, experimental and simulation, are available free of charge at https://github.com/Ramprasad-Group/polyVERSE/tree/main/Other/Gas_permeability_solubility_diffusivity.

Code availability

The Polymer Structure Predictor (PSP) package to create simulation polymer structures is available free of charge at <https://github.com/Ramprasad-Group/PSP>. The code used to perform molecular dynamics (MD) and Monte Carlo (MC) simulations is available free of charge at https://github.com/Ramprasad-Group/polyVERSE/tree/main/Other/Gas_permeability_solubility_diffusivity. The code used to train our polyGNN models is available at <https://github.com/Ramprasad-Group/polyggn> for academic use.

Received: 26 November 2023; Accepted: 4 August 2024;

Published online: 15 August 2024

References

- Ferreira, A., Alves, V. & Coelho, I. Polysaccharide-based membranes in food packaging applications. *Membranes* **6**, 22 (2016).
- Baker, R. W. Membrane technology. *Encyclopedia of Polymer Science and Technology* 2001,
- Wijmans, J. & Baker, R. The solution-diffusion model: a review. *J. Membr. Sci.* **107**, 1–21 (1995).
- Tran, H., Shen, K.-H., Shukla, S., Kwon, H.-K. & Ramprasad, R. Informatics-driven selection of polymers for fuel-cell applications. *J. Phys. Chem. C* **127**, 977–986 (2023).
- Barnett, J. W. et al. Designing exceptional gas-separation polymer membranes using machine learning. *Sci. Adv.* **6**, eaaz4301 (2020).
- Moore, T. T., Damle, S., Williams, P. J. & Koros, W. J. Characterization of low permeability gas separation membranes and barrier materials; design and operation considerations. *J. Membr. Sci.* **245**, 227–231 (2004).
- Müller-Plathe, F. Permeation of polymers—a computational approach. *Acta Polymerica* **45**, 259–293 (1994).
- Audus, D. J. & de Pablo, J. J. Polymer informatics: Opportunities and challenges. *ACS Macro Lett.* **6**, 1078–1082 (2017).
- Batra, R., Song, L. & Ramprasad, R. Emerging materials intelligence ecosystems propelled by machine learning. *Nat. Rev. Mater.* **6**, 655–678 (2021).
- Chen, L. et al. Polymer informatics: Current status and critical next steps. *Mater. Sci. Eng.: R: Rep.* **144**, 100595 (2021).
- Zhu, G. et al. Polymer genome-based prediction of gas permeabilities in polymers. *J. Polym. Eng.* **40**, 451–457 (2020).
- Wu, C. et al. Rational design of all-organic flexible high-temperature polymer dielectrics. *Matter* **5**, 2615–2623 (2022).
- Chen, L. et al. Frequency-dependent dielectric constant prediction of polymers using machine learning. *npj Computational Mater.* **6**, 61 (2020).
- Wessling, M. et al. Modelling the permeability of polymers: a neural network approach. *J. Membr. Sci.* **86**, 193–198 (1994).
- Yuan, Q. et al. Imputation of missing gas permeability data for polymer membranes using machine learning. *J. Membr. Sci.* **627**, 119207 (2021).
- Ricci, E. & De Angelis, M. G. A perspective on data-driven screening and discovery of polymer membranes for gas separation, from the molecular structure to the industrial performance. *Reviews in Chemical Engineering* 2023,
- Wang, L., Shao, C., Wang, H. & Wu, H. Radial basis function neural networks-based modeling of the membrane separation process: hydrogen recovery from refinery gases. *J. Nat. Gas. Chem.* **15**, 230–234 (2006).
- Rogers, D. & Hahn, M. Extended-connectivity fingerprints. *J. Chem. Inf. modeling* **50**, 742–754 (2010).
- Landrum, G. Rdkit documentation. *Release* **1**, 4 (2013).
- Huan, T. D., Mannodi-Kanakithodi, A. & Ramprasad, R. Accelerated materials property predictions and design using motif-based fingerprints. *Phys. Rev. B* **92**, 014106 (2015).
- Le, T., Epa, V. C., Burden, F. R. & Winkler, D. A. Quantitative structure–property relationship modeling of diverse materials properties. *Chem. Rev.* **112**, 2889–2919 (2012).
- Gurnani, R., Kuenneth, C., Toland, A. & Ramprasad, R. Polymer informatics at scale with multitask graph neural networks. *Chem. Mater.* **35**, 1560–1567 (2023).
- Kuenneth, C. & Ramprasad, R. polyBERT: a chemical language model to enable fully machine-driven ultrafast polymer informatics. *Nat. Commun.* **14**, 4099 (2023).
- Ramprasad, R., Batra, R., Pilania, G., Mannodi-Kanakithodi, A. & Kim, C. Machine learning in materials informatics: recent applications and prospects. *npj Computational Mater.* **3**, 54 (2017).
- Hutchinson, M. L. et al. Overcoming data scarcity with transfer learning. *arXiv e-prints*. arXiv–1711 (2017).
- Caruana, R. *Learning to Learn*; Springer, 1998.
- Patra, A. et al. A multi-fidelity information-fusion approach to machine learn and predict polymer bandgap. *Computational Mater. Sci.* **172**, 109286 (2020).
- Kuenneth, C. et al. Polymer informatics with multi-task learning. *Patterns* **2**, 100238 (2021).
- Yang, J., Tao, L., He, J., McCutcheon, J. R. & Li, Y. Machine learning enables interpretable discovery of innovative polymers for gas separation membranes. *Sci. Adv.* **8**, eabn9545 (2022).
- Venkatram, S. et al. Predicting crystallization tendency of polymers using multifidelity information fusion and machine learning. *J. Phys. Chem. B* **124**, 6046–6054 (2020).
- Weininger, D. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.* **28**, 31–36 (1988).
- Brandrup, J., Immergut, E. H., Grulke, E. A., Abe, A. & Bloch, D. R. *Polymer handbook*, Vol. 89 (Wiley New York, 1999).
- Thompson, A. P. et al. LAMMPS—a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Comput. Phys. Commun.* **271**, 108171 (2022).
- Wang, J., Wolf, R. M., Caldwell, J. W., Kollman, P. A. & Case, D. A. Development and testing of a general amber force field. *J. Comput. Chem.* **25**, 1157–1174 (2004).
- Potoff, J. J. & Siepmann, J. I. Vapor–liquid equilibria of mixtures containing alkanes, carbon dioxide, and nitrogen. *AIChE J.* **47**, 1676–1682 (2001).
- Sahu, H., Shen, K.-H., Montoya, J. H., Tran, H. & Ramprasad, R. Polymer structure predictor (psp): a python toolkit for predicting atomic-level structural models for a range of polymer geometries. *J. Chem. Theory Comput.* **18**, 2737–2748 (2022).
- Abbott, L. J., Hart, K. E. & Colina, C. M. Polymatic: a generalized simulated polymerization algorithm for amorphous polymers. *Theor. Chem. Acc.* **132**, 1–19 (2013).
- Frenkel, D.; Smit, B. *Understanding molecular simulation: from algorithms to applications*; Elsevier, 2023.
- Khawaja, M., Sutton, A. & Mostofi, A. Molecular simulation of gas solubility in nitrile butadiene rubber. *J. Phys. Chem. B* **121**, 287–297 (2017).
- Longuet-Higgins, H. & Widom, B. A rigid sphere model for the melting of argon. *Mol. Phys.* **8**, 549–556 (1964).

41. Mocak, J., Bond, A. M., Mitchell, S. & Scollary, G. A statistical overview of standard (IUPAC and ACS) and new procedures for determining the limits of detection and quantification: application to voltammetric and stripping techniques (technical report). *Pure Appl. Chem.* **69**, 297–328 (1997).
42. Robeson, L. M. Correlation of separation factor versus permeability for polymeric membranes. *J. Membr. Sci.* **62**, 165–185 (1991).
43. Robeson, L. M. The upper bound revisited. *J. Membr. Sci.* **320**, 390–400 (2008).
44. Comesaña-Gándara, B. et al. Redefining the Robeson upper bounds for CO_2/CH_4 and CO_2/N_2 separations using a series of ultrapermeable benzotriptycene-based polymers of intrinsic microporosity. *Energy Environ. Sci.* **12**, 2733–2740 (2019).
45. Weinkauff, D. & Paul, D. *Effects of structural order on barrier properties*; ACS Publications, 1990.
46. Michaels, A. S. & Bixler, H. J. Solubility of gases in polyethylene. *J. Polym. Sci.* **50**, 393–412 (1961).
47. Shen, K.-H., Brown, J. R. & Hall, L. M. Diffusion in lamellae, cylinders, and double gyroid block copolymer nanostructures. *ACS Macro Lett.* **7**, 1092–1098 (2018).
48. Shen, K.-H. & Hall, L. M. Effects of ion size and dielectric constant on ion transport and transference number in polymer electrolytes. *Macromolecules* **53**, 10086–10096 (2020).
49. Shen, K.-H. & Hall, L. M. Ion conductivity and correlations in model salt-doped polymers: Effects of interaction strength and concentration. *Macromolecules* **53**, 3655–3668 (2020).
50. Gurnani, R. P. Debugging Neural Networks. <https://nanohub.org/resources/netdebugger> (2021).
51. Head, T. et al. scikit-optimize/scikit-optimize: v0. 5.2. *Version v0* 2018, 5.
52. PACE Partnership for an Advanced Computing Environment (PACE). (2017).

Acknowledgements

This work is financially supported by Toyota Research Institute through the Accelerated Materials Design and Discovery program and the Office of Naval Research through a multidisciplinary university research initiative (MURI) grant N00014-20-1-2586. This research is supported in part through research cyber-infrastructure resources and services provided by the Partnership for an Advanced Computing Environment (PACE) at the Georgia Institute of Technology, Atlanta, Georgia, USA⁵². The authors thank XSEDE/ACCESS for computational support through Grant No. TG-DMR080058N.

Author contributions

The work was conceived and guided by R.R. B.K.P. designed, trained, and evaluated the machine learning models. B.K.P. and K.H.S. developed the simulation pipelines. B.K.P., K.H.S., R.G., H.T., R.L., and R.R. discussed the results and commented on the manuscript.

Competing interests

R.R. is a founder of Matmerize, Inc., a company specializing in materials informatics software and services. The other authors have no conflicts of interest to declare.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41524-024-01373-9>.

Correspondence and requests for materials should be addressed to Rampi Ramprasad.

Reprints and permissions information is available at <http://www.nature.com/reprints>

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

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

© The Author(s) 2024, corrected publication 2025