

# Capturing Molecular Interactions in Graph Neural Networks: A Case Study in Multi-Component Phase Equilibrium

Shiyi Qin<sup>1</sup>, Shengli Jiang<sup>1</sup>, Jianping Li<sup>1</sup>  
Prasanna Balaprakash<sup>2</sup>, Reid C. Van Lehn<sup>1</sup>, and Victor M. Zavala<sup>1\*</sup>

<sup>1</sup>Department of Chemical and Biological Engineering  
University of Wisconsin - Madison, 1415 Engineering Dr, Madison, WI 53706, USA  
<sup>2</sup>Mathematics and Computer Science Division & Leadership Computing Facility  
Argonne National Laboratory, 9700 S Cass Ave, Lemont, IL 60439, USA

## Abstract

Graph neural networks (GNNs) have been widely used for predicting molecular properties, especially for single molecules. However, when treating multi-component systems, GNNs have mostly used simple data representations (concatenation, averaging, or self-attention on features of individual components) that might fail to capture molecular interactions and potentially limit prediction accuracy. In this work, we propose a GNN architecture that captures molecular interactions in an explicit manner by combining atomic-level (local) graph convolution and molecular-level (global) message passing through a molecular interaction network. We tested the architecture (which we call SolvGNN) on a comprehensive phase equilibrium case study that aims to predict activity coefficients for a wide range of binary and ternary mixtures; we built this large dataset using the CONductor-like Screening MOdel for Real Solvation (COSMO-RS). We show that SolvGNN can predict composition-dependent activity coefficients with high accuracy and show that it outperforms a previously developed GNN used for predicting infinite-dilution activity coefficients. We performed counterfactual analysis on the SolvGNN model that allowed us to explore the impact of functional groups and composition on equilibrium behavior. We also used the SolvGNN model for the development of a computational framework that automatically creates phase diagrams for a diverse set of complex mixtures. All scripts needed to reproduce the results are shared as open-source code.

---

\*Corresponding Author: victor.zavala@wisc.edu.

## 1 Introduction

Predicting the physicochemical properties of molecules is crucial for applications such as product and process design. In the past decade, machine learning (ML) techniques have been used as data-driven approaches that help accelerate molecule screening and to reduce experimentation cost, especially when a large chemical space is involved. These models have also shown to be versatile and to predict diverse molecular properties such as water solubility [1–3], toxicity [4–6], and lipophilicity [7,8]. A fundamental step in the use of ML models is the pre-definition or pre-calculation of molecular descriptors [9–12]; such descriptors are used as input data to develop quantitative structure-property relationship models [13]. Recently, there has been growing interest in applying ML models to study more complex chemical systems that might contain multiple components such as chemical reactions [14, 15], alloys [16, 17], copolymers [18, 19], and gas/liquid mixtures [20–28]. Among the ML techniques explored, graph neural networks (GNNs) [29, 30] have gained special popularity because they can directly incorporate molecular representations (in the form of graphs), which enable the capturing of key structural information while potentially avoiding the need to pre-calculate/pre-define descriptors using more advanced but computationally-intensive tools such density functional theory (DFT) or molecular dynamics (MD) models.

In a typical GNN architecture for prediction of molecular properties [31], the characteristics of the atoms and of the bonds are propagated based on the chemical structure of a single input molecule, followed by featurizing embedding via nonlinear transformation. The embedded features are then fed to fully connected layers to construct predictive models. GNNs have achieved better performance than conventional descriptor-based approaches in various benchmark datasets [32, 33]. When dealing with multiple components, several approaches have been devised; a typical way to encode multi-molecule information is to simply average or concatenate the features of individual molecules and to use these as system-level features for property inference with fully connected or attentive layers [14, 15, 19]. Previous studies have also incorporated weighted sums or concatenation to take into account composition information when needed [19]. However, these approaches do not capture molecular interactions in an explicit manner, which may limit the predictive power of GNNs for systems in which intermolecular interactions play an important role.

In this work, we present a GNN architecture that explicitly incorporates molecular interactions via the combination of atomic-level (local) graph convolution and molecular-level (global) message passing for property prediction of multi-component chemical systems. To connect local features with global features, we construct a molecular interaction network as an intermediate step. The molecular interaction network is a complete graph in which each node represents a molecule and each edge represents a hypothetical intermolecular interaction (e.g., hydrogen bonding). This representation serves as a physics-informed topological prior that aid feature extraction

from multi-component systems. The composition information is also encoded in the architecture as node-weights for the molecular interaction network. We hypothesize that, with this type of data representation and feature propagation guided by physical intuition, the proposed architecture may better model mixture properties while taking composition information into consideration.

We evaluate the proposed GNN architecture through a comprehensive case study on miscibility calculations for multi-component systems. We choose activity coefficients as the target thermodynamic property of interest, which measures the deviation of a liquid mixture from ideal solution behavior. Activity coefficients are one of the fundamental properties of a mixture and therefore can lead to the derivation of equilibrium conditions (*e.g.*, phase diagrams), which are important in physical chemistry and engineering for understanding and optimizing chemical separations [34]. Previous studies have developed ML-based methods to predict infinite-dilution activity coefficients for binary mixtures, including matrix completion on the activity coefficient matrix [27, 35] and multilayer perceptrons on the system descriptors [36]. However, these methods did not account for molecular structural information directly, and the latter is limited to systems of water in ionic liquids. A more recent study by Medina *et. al.* [28] used GNN models to tackle this problem; this approach, however, used a data representation that involves a simple concatenation of individual graph features after local embedding (*i.e.*, the GNN architecture does not explicitly capture intermolecular interactions). Furthermore, all these previous studies have focused on predicting infinite-dilution coefficients, which do not take into consideration composition information (this limits their use in more sophisticated thermodynamic predictions such as phase diagrams). To the best of our knowledge, GNNs have not been explored as a method to predict composition-dependent activity coefficients nor have they been extended to predict activity coefficients for systems with more than two components. The proposed GNN architecture is generalizable to multiple component systems and captures composition.

Through our case study, we demonstrate that the proposed GNN (which we call SolvGNN) outperforms prior architectures (that lack an explicit graph representation of molecular interactions) in terms of prediction accuracy. Our study leverages a large dataset that was developed using the COnductor-like Screening MOdel for Real Solvation (COSMO-RS). SolvGNN also enables better modeling of mixture compositions due to the incorporation of global message passing on the molecular interaction network with hydrogen bonding information. We also show that SolvGNN can be applied to both binary and ternary liquid-phase mixtures to predict composition-dependent activity coefficients. To interpret our SolvGNN predictions, we perform counterfactual analysis [37] to identify the impact of functional groups on activity coefficients. To demonstrate the applicability of SolvGNN, we developed a framework that can automatically generate phase diagrams (P-x-y) for complex binary and ternary mixtures. The generated phase diagrams can be used to study solvent miscibility and to help identify azeotrope compositions to guide the design of targeted mixtures and chemical separations. We share scripts and datasets as open-source code to enable the reproduction of the results and to conduct benchmarks.

## 2 Materials and Methods

### 2.1 Data Set Summary

We assembled a list of 700 common solvents [38], covering a wide spectrum of small molecules such as water, alcohols, esters, and ethers. We then used random sampling over the solvent space to generate 40,000 binary mixtures and 40,000 ternary mixtures. For each binary mixture, we explored five molar composition ratios—10%/90%, 30%/70%, 50%/50%, 70%/30%, 90%/10%; for each ternary mixture, we explored four molar composition ratios—15%/15%/70%, 15%/70%/15%, 70%/15%/15%, and 33.3%/33.3%/33.4%. Overall, we assembled a large database with 200,000 entries for binary mixtures and 160,000 entries for ternary mixtures.

To visualize the coverage of the chemical space, the solvents were grouped into 22 categories based on a predefined list of functional groups (details are provided in the SI). The visualization is provided in Figure 1a. The sampled binary mixtures are represented by connections between nodes. The number of solvents in each category and the number of sampled pairs are reflected by node size and edge thickness; this illustrates that our random mixture sampling covers a wide range of solvent pairs in different categories. We also visualized the chemical space of the solvents by performing a t-distributed stochastic neighbor embedding (t-SNE) [39] dimensionality reduction technique on the Morgan fingerprints [9], also known as extended connectivity fingerprints [10] in Figure 1b. The 2D map from t-SNE shows separation between some solvent categories, such as nitriles and aromatics. However, because some solvents contain more than one identifiable functional group, they may potentially be grouped into another category. As a result, the clustering in a few other categories is less clear, but in general the scattered distribution here suggests the inclusion of diverse and complex chemical structures.

We categorized the sampled binary and ternary mixtures based on whether each component in the mixture is polar or nonpolar, as summarized in Table 1. We computed the percentage of each mixture type; this information was used for stratified sampling, which creates training/validation folds by preserving the percentage of samples for each mixture type (this ensures that the model learns different types of molecular interactions). Overall, most mixtures contain at least one polar component, indicating the presence of strong intermolecular interactions (*e.g.*, dipole-dipole forces).

### 2.2 Activity Coefficient Calculations

To overcome the challenge of limited experimental data availability, we used the COnductor-like Screening MOdel for Real Solvation (COSMO-RS) to generate ground-truth labels for supervised ML. COSMO-RS calculations are based on surface charge densities ( $\sigma$ -profiles) of mixture components, which are obtained from DFT calculations coupled with the COSMO continuum solvation model [40], and it can be used to calculate the activity coefficients for any mixture as long as the

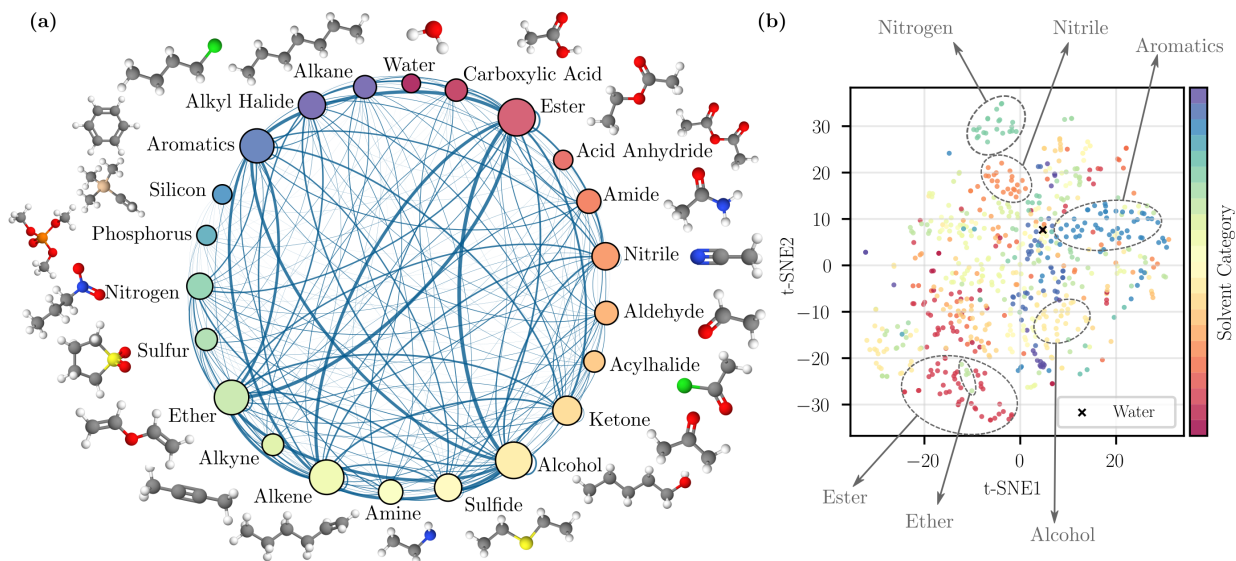


Figure 1: Dataset visualization. (a) Solvent categories based on the primary functional group of individual solvents with examples. The size of a node reflects the number of solvents in that category, and the thickness of an edge reflects the number of sampled pairs between two categories. (b) 2D map obtained by t-SNE dimensionality reduction [39] applied to molecular fingerprints.

Table 1: Mixture types based on polarity of individual components.

	Mixture Type	Count	Percentage
Binary (200,000)	polar-polar (p-p)	116,105	58%
	polar-nonpolar (p-n)	72,795	36%
	nonpolar-nonpolar (n-n)	11,100	6%
Ternary (160,000)	polar-polar-polar (p-p-p)	71,136	45%
	polar-polar-nonpolar (p-p-n)	66,040	41%
	polar-nonpolar-nonpolar (p-n-n)	20,848	13%
	nonpolar-nonpolar-nonpolar (n-n-n)	1,976	1%

chemical structures are provided and optimized. For each solvent mixture, we obtained activity coefficients  $\gamma_i$  for individual components  $i$  from COSMO-RS and constructed large and structured data sets for model training and evaluation.

COSMOtherm [41] (version 2019), a software that implements COSMO-RS, was used to obtain composition-dependent activity coefficients for the individual components of each sampled mix-

ture. Prior to COSMO-RS calculations, chemical structures were generated from CirPy (version 1.0.2), a Python library that serves as the interface for the Chemical Identifier Resolver (CIR) [42]; this searches the National Institutes of Health (NIH) database for the chemical structures and provides the optimized coordinates for the atoms. We next conducted DFT calculations using TURBOMOLE [43] (version 7.5) at the BP-TZVP theory level with the Becke-Perdew (BP) functional and the resolution of identity approximation under ideal screening condition ( $\epsilon^\infty$ , COSMO continuum solvation model). A single-point calculation was then conducted with the def2-TZVPD basis set and fine cavity parameter to create the  $\sigma$ -profiles. Activity coefficients were then calculated given the  $\sigma$ -profiles of individual components, the mixture compositions in the liquid phase, and temperature (298K).

### 2.3 GNN Model Architecture

As shown in Figure 2, each input mixture is represented as molecular graphs  $\mathcal{G} = (V, E, H)$  of individual components with nodes  $v \in V$ , edges  $e \in E$ , and node feature matrix

$$H = \begin{bmatrix} - & h_{v_1}^T & - \\ - & h_{v_2}^T & - \\ & \vdots & \end{bmatrix} \quad (1)$$

that encodes atom and bond information such as atom types and degrees [44]. A local graph convolution [45] was applied to each of the input molecular graphs, and the node features were updated through

$$H^{(t+1)} = \text{ReLU} \left( \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H^{(t)} W^{(t)} \right). \quad (2)$$

Here,  $\tilde{A}$  is the adjacency matrix of graph  $\mathcal{G}$  with self-loops,  $\tilde{D} = \sum_j \tilde{A}_{ij}$  is the degree matrix and  $W^{(t)}$  is the learnable weight matrix at time step  $t$ .  $\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}}$  is derived from normalized graph Laplacian that accounts for graph topology and implicitly imitates molecular interactions. The  $W^{(t)}$  values are kept the same for each component in the mixture. After local graph convolution, node-level features are averaged to generate the graph-level feature  $u_{\mathcal{G}} = \frac{1}{|V|} \sum_v h_v$ .

We compared several approaches to capture molecular interactions. The first approach is illustrated in Figure 2a and referred to as SolvCAT. In this approach,  $u_{\mathcal{G}}$ 's undergo a composition  $x$ -weighted concatenation to form a fixed-length latent feature vector. For a ternary system, for example,  $u_{mix} = x_1 u_{\mathcal{G}_1} | x_2 u_{\mathcal{G}_2} | x_3 u_{\mathcal{G}_3}$ . The system-level feature vector is then sent to fully connected neural network layers for activity coefficient predictions. The second approach is illustrated in Figure 2b and referred to as SolvGCN. In this approach, a molecular interaction network was constructed to explicitly simulate molecular interactions between the components in a system. The molecular interaction network  $\mathcal{G}_{int} = (V_{mol}, E_{int}, H_{mol})$  is a complete graph where each node  $v_{mol} \in V_{mol}$  denotes a molecule, each edge  $e_{int} \in E_{int}$  denotes the existence of certain intermolec-

ular interaction, and molecular-level node feature matrix:

$$H_{mol} = \begin{bmatrix} -h_{v_{mol1}}^T = x_1 u_{\mathcal{G}_1}^T & - \\ -h_{v_{mol2}}^T = x_2 u_{\mathcal{G}_2}^T & - \\ \vdots & \end{bmatrix}. \quad (3)$$

A global graph convolution is applied using the same updating rules as Eq.2; in this case,  $u_{mix}$  is obtained by concatenating the latent node features  $h_{mol}$ 's after global graph convolution.

The third approach is illustrated in Figure 2c and referred to as SolvGNN. Building on SolvGCN, for this approach we developed a more informative representation of the molecular interaction network; we encoded hydrogen-bond (H-bond) information, one of the strongest form of dipole-dipole interactions, as the edge feature. For a ternary system, this feature is formulated as:

$$e_{int(v_{mol_i} v_{mol_j})} = \min(\#H\text{-bond acceptor}_{v_{mol_i}}, \#H\text{-bond donor}_{v_{mol_j}}) \quad (4)$$

where H-bond information between like molecules ( $i = j$ ) and unlike molecules ( $i \neq j$ ) are both captured. In this case, the global graph convolution integrates edge features and is achieved via message passing [30] expressed by

$$m_{v_{mol}}^{(t+1)} = \sum_{v_{mol_i} \in N(v_{mol})} M_t(h_{v_{mol}}^{(t)}, h_{v_{mol_i}}^{(t)}, e_{int(v_{mol} v_{mol_i})}) \quad (5)$$

and

$$h_{v_{mol}}^{(t+1)} = U_t(h_{v_{mol}}^{(t)}, m_{v_{mol}}^{(t+1)}), \quad (6)$$

Here, we used the original message passing formulation [30], where the message function  $M_t$  is a fully-connected edge network to compute a message matrix based on graph topology as well as edge features, and the node update function  $U_t$  is a gated recurrent unit (GRU) [46] to aggregate "message" and the original node feature, which can be viewed as a generalization of the plain GCN.

In all three cases mentioned above, the embedded features after "intermolecular interactions" are flattened and sent to the fully connected readout layers for the final activity coefficient ( $\gamma_i$ ) prediction.

## 2.4 GNN Training and Hyperparameter Tuning

SMILES strings were used as molecule identifiers and processed by RDKit (version 2019.03.2) [47] to generate molecular graphs. The GNN models were constructed using PyTorch (version 1.2.0) [48] and Deep Graph Library (version 0.4.3) [49]. The major hyperparameters we varied include the number of graph convolution layers (1,2), the number of fully connected readout layers (1,2,3), the number of hidden neurons (128,256), and the learning rate (0.0005,0.001). The model was trained with the average mean-squared-error (MSE) loss for the  $\ln \gamma_i$  values, the Adam optimizer, a learning rate of 0.001, and a batch size of 100 for 100 epochs. We performed 5-fold cross-validation

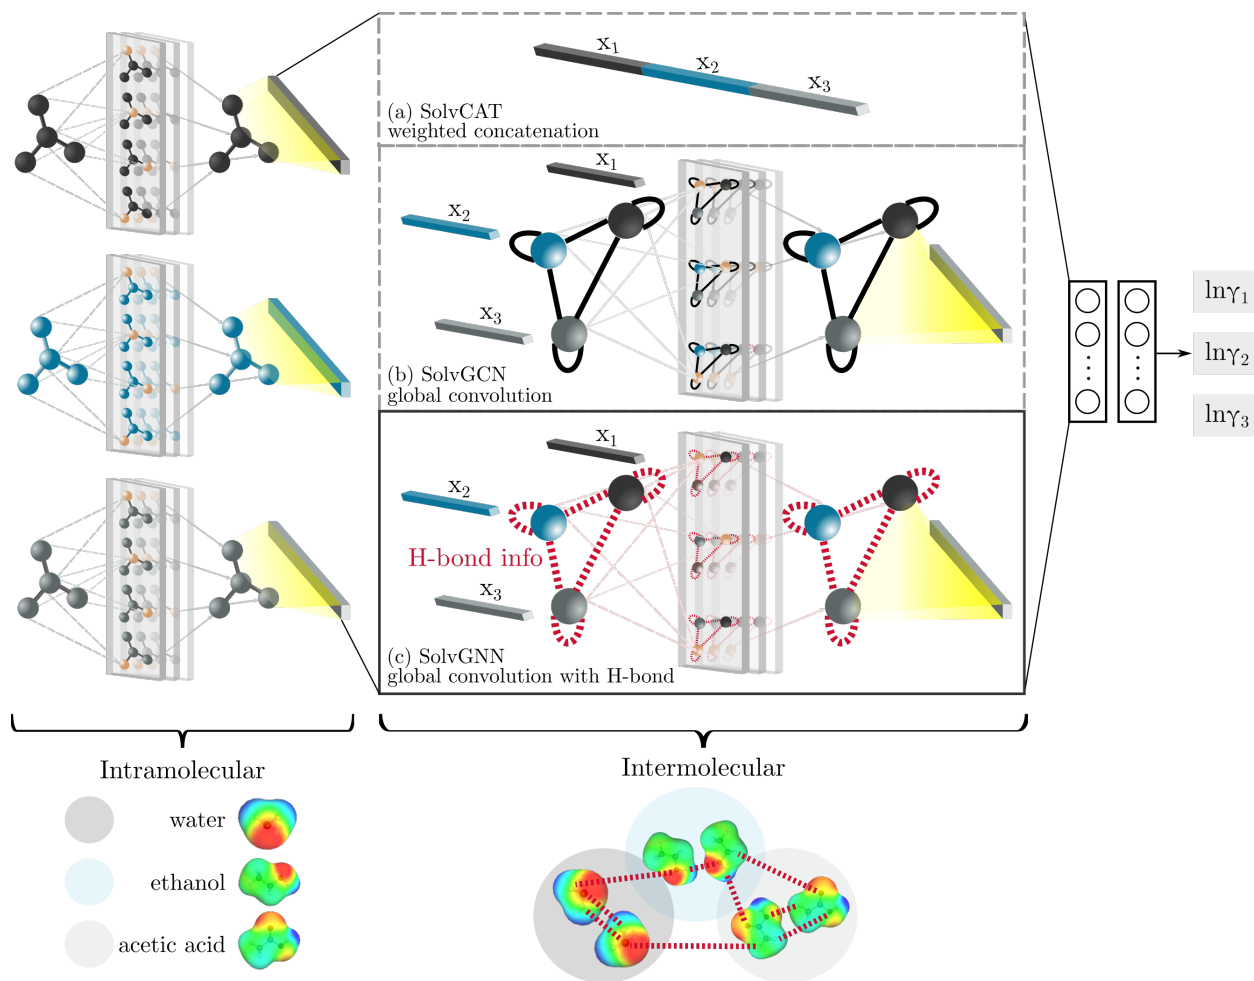


Figure 2: GNN architectures studied. All three GNN architectures undergo the same graph convolution for feature embedding of individual components at a local level. They differ in their way of capturing intermolecular interactions. SolvCAT (a) conducts a simple weighted concatenation of the locally embedded features; SolvGCN (b) constructs an intermediate molecular interaction network followed by global convolution without explicit edge information; SolvGNN (c) explicitly incorporates H-bond information as the edge feature in the interaction network, which undergoes global message passing for "intermolecular"-level feature embedding. In each case, the globally embedded features are flattened and used for activity coefficient ( $\gamma_i$ ) predictions through fully connected readout layers. Images at the bottom illustrate screening charge densities computed from COSMO-RS and representative interactions.

(CV) with stratified random sampling for hyperparameter tuning and model evaluation. Stratification is based on the type of mixture (*e.g.*, polar-polar or polar-nonpolar). All evaluation metrics are computed using the compilation of the validation data in each fold to obtain a realistic estimation of the model performance. More implementation details can be found in the SI.



Because the data sets contain a large number of binary or ternary mixtures at different compositions, it is computationally expensive to generate the corresponding molecular graphs for every training/validation instance. As a result, we designed our data loading and model training algorithm to lower the training time. Upon data set initiation, we generated and stored all 700 molecular graphs at once in a dictionary format. When a training/validation instance was passed to the algorithm, the molecular graphs were obtained from the dictionary using the index and only require simple manipulation (*e.g.*, calculation for intermolecular H-bond) to form the desired mixture data. Doing so largely reduced redundant calculations and saved time (from days to a couple of hours).

## 2.5 Counterfactual Analysis

To interpret the trained model, we adapted the counterfactual framework proposed in [37] to understand which chemical structures and functional groups lead to certain activity coefficient predictions. Here, we generated two types of counterfactuals for our dataset. Counterfactual Type I (Eq.7) focused on searching for mixture samples with minimal input differences but maximal output deviations from a base mixture. Counterfactual Type II (Eq.8) focused on the mixture samples with the maximal input differences and minimal output deviations. The distance between mixtures  $d(\text{mixture}, \text{mixture}')$  was obtained via the mean Tanimoto similarity [50] of the pair, and the difference between predicted activity coefficient predictions was computed with the MAE of the  $\ln \gamma_i$  values using the trained SolvGNN denoted as  $\hat{f}$ . The parameter  $\lambda$  is a trade-off parameter that controls the relative importance of mixture (input) similarity and prediction (outcome) difference. The parameter  $\lambda$  was set to 0.9 to generate Type I counterfactuals with a similarity value of at least 0.6. The search space was limited to the 700 solvents in our data set to keep the computational cost tractable, especially for Type II counterfactuals.

## 2.6 Phase Diagram Calculations

For an illustration of real-world applications, we set up a computational framework that can intake the chemical structures from diverse binary or ternary mixtures and make activity coefficient predictions with uncertainties by averaging the predicted values from individually trained SolvGNNs in each CV fold. P-x-y phase diagrams were then generated from the predicted activity coefficients  $\gamma_i$  using modified Raoult’s Law  $P = \sum_i y_i P = \sum_i x_i \gamma_i P_i^{sat}$ . The saturation pressure  $P_i^{sat}$  for each component was obtained using the Antoine Equation  $\log_{10} P_i^{sat} = A_i - \frac{B_i}{C_i + T}$  with coefficients collected from the National Institute of Standards and Technology (NIST) via web scraping [51]. We sampled the liquid-phase compositions  $x_i$  and calculated the equilibrium pressures  $P$  with the specified compositions at 298 K. For ternary systems, we computed phase diagrams following the same method for the binary systems by sampling the mixture compositions followed by equilibrium pressure calculations. The isobaric liquid-liquid-equilibrium (LLE) lines were then obtained by filtering the sampled points around a selected equilibrium pressure.

$$\max_{mixture'} d(mixture, mixture') + \lambda(MAE(\hat{f}(mixture), \hat{f}(mixture'))) \quad (7)$$

and

$$\min_{mixture'} d(mixture, mixture') + \lambda(MAE(\hat{f}(mixture), \hat{f}(mixture'))) \quad (8)$$

## 3 Results and Discussion

### 3.1 Model Performance on Binary Mixtures

We compared the three GNN architectures (SolvCAT, SolvGCN, and SolvGNN) introduced in the previous section in terms of their ability to predict the composition-dependent activity coefficients. SolvCAT takes the weighted concatenation of embedded features after local graph convolutions on individual components; SolvGCN constructs a complete interaction network after local convolution without any assumptions on the edge weights for another layer of graph convolution at the global level. SolvGNN takes this SolvGCN one step further by introducing H-bond information as an example prior knowledge on intermolecular interactions for further message passing. The performance of the three GNN architectures was evaluated on the binary mixture data set by the cumulative frequency plot, as shown in Figure 3a. In the cumulative frequency plot, the absolute errors of the natural logarithms of the activity coefficients,  $\ln \gamma_1$  and  $\ln \gamma_2$ , (between true and predicted values from CV) for each data point were first averaged, and the cumulative frequencies for the averaged error values were then plotted in the ascending order. Among the three GNN architectures, SolvGNN exhibits the best performance; specifically, it shows that almost 96% of the data points are predicted with an error of less than 0.1. SolvCAT performs slightly worse, with 88% of the data points falling within the 0.1 error range. SolvGCN shows the worst performance, with only 46% of the data points are predicted with an error less than 0.1. These observations are also supported by the mean absolute errors (MAEs), which are 0.03, 0.05, and 0.20 for SolvGNN, SolvCAT, and SolvGCN (respectively).

The above results indicate that the inclusion of the global interaction network with H-bond information in SolvGNN provides an effective method for improving the prediction accuracy for activity coefficients. However, when H-bond information is excluded, the pure global graph convolution worsens the model performance, possibly due to the unbiased "averaging" without any physics-informed resemblance to intermolecular interactions. On the other hand, SolvCAT, despite the lack of explicit global graph convolution that depicts intermolecular interactions, still exhibits satisfactory predictive power. This is consistent with an earlier study [19], which has found that aggregation over composition-weighted latent features provides an effective approach to handle information of mixture composition.

Figure 3b and 3c show individual parity plots from SolvGNN for  $\ln \gamma_1$  and  $\ln \gamma_2$ . All the predictions shown are from the validation process yet still exhibit high accuracy, with average  $\ln \gamma_i$  MAE

being 0.03 and average  $\ln \gamma_i$  RMSE being 0.08. The data points are colored by the mixture type defined earlier. In general, the values of  $\ln \gamma$  for nonpolar-nonpolar interactions are close to 0 (ideal behavior) and have smaller MAE, while the values of  $\ln \gamma$  for mixtures with polar components spread across the entire data range and have larger MAE. With respect to composition, mixtures that are rich in one of the components (10%/90%) exhibit a higher MAE ( $\sim 0.04$ ), whereas the mixtures with equimolar components exhibit a lower MAE ( $\sim 0.025$ ). We also identified a couple of outliers in the plots; these mixtures contain amines with hydrogen-bonding solutes or solvents. The extreme  $\ln \gamma$  values of these mixtures can be the result of limitations of COSMO-RS, which has been specifically noted to incorrectly simulate the interactions of secondary and tertiary amines when hydrogen-bond donors or acceptors are present in the system [52].

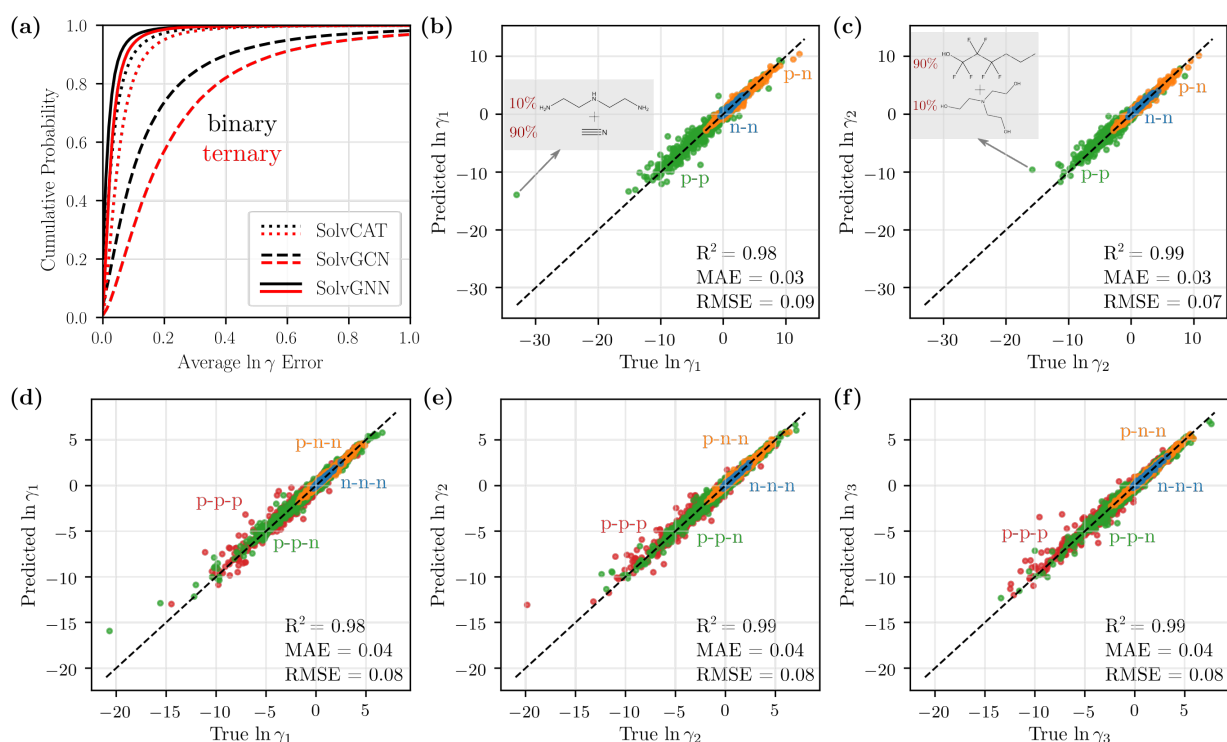


Figure 3: Model comparison and parity plots for binary and ternary mixtures. (a) Cumulative frequency plots for the average  $\ln \gamma_i$  errors for binary (black) and ternary (red) mixtures to compare SolvCAT, SolvGCN, and SolvGNN. Additionally, the parity plots for individual  $\ln \gamma_i$ 's between the true (COSMO-RS) and predicted (SolvGNN) values from CV are displayed for binary (b-c) and ternary (d-f) mixtures. The points are colored by the type of mixtures defined in Table 1 based on polarity.

Besides the regular CV using stratified sampling that relies on the type of mixture, we also tested the robustness of the SolvGNN using an alternative CV method. Here, for each CV fold, we trained the model on only two of the three mixture types (polar-polar, polar-nonpolar, or nonpolar-nonpolar; see Table 1) and validated the rest. Results have shown that, although the model could

achieve similar training losses (MAE=0.04) to the regular CV (MAE=0.03), the validation accuracy was reduced accordingly. For the case where we trained the model with polar-polar and polar-nonpolar mixtures (94% of the data set) while validating on nonpolar-nonpolar mixtures, the model demonstrated suitable transferability by comparable validation losses (MAE=0.04). However, when we trained the model on only polar-polar and nonpolar-nonpolar samples (64%) while validating on polar-nonpolar samples, validation MAE increased by 0.09, which indicates the distinct nature of polar-nonpolar interactions and suggests that it is non-trivial and therefore cannot be omitted during model training. Additionally, we examined the condition when the model was trained on only polar-nonpolar and nonpolar-nonpolar mixtures (42%) and validated on polar-polar mixtures. The convergence plot (included in the SI) indicates over-training with high validation losses; such behavior was expected because the training size is less than half of the data set, and the majority of the training samples lack H-bond acceptors or donors, which are present in most of the validation set. This result again suggests that polar-polar and polar-nonpolar mixtures, despite possessing strong intermolecular interactions such as H-bonding in both cases, are intrinsically different and therefore are both required in the training process. These results are in general agreement with chemical intuition.

### 3.2 Scale up to Ternary Mixtures

We next scaled up the proposed SolvGNN architecture to ternary mixtures. As shown in Figure 3a (red), the cumulative frequency of the average  $\ln \gamma$  errors demonstrates similar trends as in binary mixtures. SolvGNN provides the best model performance, followed by SolvCAT and SolvGCN. Here, we observed that the gaps between the curves appear wider, suggesting a more significant advantage of SolvGNN over SolvCAT and SolvGCN. For SolvGNN, more than 93% of the data points were predicted with an error less than 0.1. SolvCAT was the second-most accurate model, with around 80% of the data points falling within the 0.1 error range, showing a more notable performance drop (13%) than the results for binary mixtures (8%). SolvGCN continues to exhibit the worst performance, and only around 30% of the data points are predicted with an error less than 0.1. These observations are supported by MAE values, which are 0.04, 0.07, and 0.27 for SolvGNN, SolvCAT, and SolvGCN.

For SolvGNN, comparable model accuracy was obtained even though the number of training/validation samples was reduced compared to the binary mixture data set, as shown in Figure 3d-f. The CV  $R^2$ , MAE, and RMSE are similar to the results from binary systems, with corresponding values around 0.99, 0.04, and 0.08. When breaking down the predicted values based on the mixture type, we found that samples containing only nonpolar components tend to have smaller errors and systems containing only polar components have larger errors. Mixtures with both polar and non-polar components have MAEs and RMSEs lying somewhere in between the extremes. When grouping by composition, mixtures that are rich in one of the components tend to have higher prediction errors than equimolar mixtures. This observation is consistent with previous

model performance on binary mixtures and could be caused by the fact that the majority of the training data are not equimolar systems.

Overall, SolvGNN exhibited satisfactory performance in making predictions for activity coefficients of binary and ternary systems, given the advantage of explicitly including H-bond information (as a representative and primary intermolecule force) via global message passing on the molecular interaction network. To the best of our knowledge, this is the first time that such a GNN architecture is used to make predictions for general activity coefficients that depend on composition (compared to infinite-dilution activity coefficients) and for ternary systems (compared to binary systems).

### 3.3 Comparison to Previous GNN for Infinite-Dilution Activity Coefficients

To compare SolvGNN with a recently developed GNN for infinite dilution activity coefficient ( $\ln \gamma^\infty$ ) prediction by Medina et al. [28], we conducted a benchmark of our model on the same experimental data set used in their study, which contains 2,810 binary mixtures (with specific solute/solvent assignment) and values of  $\ln \gamma^\infty$  for the solute. Because infinite-dilution activity coefficients do not rely on composition, we removed the composition information  $x_i$ , but kept the same interaction network and the H-bond method. To conduct a fair comparison, we applied the same training / validation / testing method described in their research [28] through ensemble learning (bagging), which splits the training/validation data randomly 30 times and averages the predictions. We also used the same batch size (32) and epoch number (200), and calculated the evaluation metrics on the unscaled  $\gamma^\infty$  values. The results are summarized in Table 2 for a comparison of the test data. We observed that the performance of SolvGNN (without composition) is better than the previous GNN model [28] for  $\ln \gamma^\infty$  prediction. The proposed SolvGNN shows improvements in all metrics used to evaluate the model in the original paper, including a significant decrease in the mean absolute percent error (MAPE) by 44%.

Table 2: Simplified SolvGNN without composition  $x_i$  for prediction of infinite dilution activity coefficients ( $\ln \gamma^\infty$ ) and a comparison with the previously developed GNN [28] on the test data.

Model	MAE	SDEP	MSE	RMSE	MAPE	$R^2$
Previous GNN [28]	3.91	26.73	729.69	27.01	22.66	0.82
SolvGNN w/o $x_i$	3.73	21.25	465.66	21.58	12.72	0.87
% difference	-5%	-21%	-36%	-20%	-44%	+6%

In general, our results provide evidence that the simplified version of SolvGNN (without composition information) can be used to predict infinite-dilution activity coefficients in a satisfactory manner, thus illustrating that the architecture is versatile. Comparison of the evaluation metrics indicates that there is a benefit in including intermolecular interactions in the GNN architecture.

These results also provide evidence that the SolvGNN architecture can be used to learn not only from simulation data (*e.g.*, COSMO-RS) but also from experimental data. To further confirm comparable performance on both COSMO-RS and experimental data, we also generated  $\ln \gamma^\infty$  (on our binary data set) using COSMO-RS and trained the simplified SolvGNN. The results are summarized in the SI and again demonstrate accurate predictions.

### 3.4 Counterfactual Analysis

We derived counterfactuals [37] as a way to provide some interpretability to SolvGNN predictions. Here, we investigated two types of counterfactuals: mixtures with the highest similarity yet the most different predictions (Type I), and mixtures with the lowest similarity yet the most similar predictions (Type II). As illustrated in Figure 4, we started with a base mixture (50% benzene and 50% toluene) that exhibits nearly ideal behavior ( $\gamma_i = 1$  for both components). Since input chemical ratios are also a contributing factor to activity coefficients, we first identified the composition with the same two molecular species that leads to the farthest deviation in activity coefficients, as illustrated by Counterfactual 1. We found that increasing the composition in toluene to the extreme has the most significant impact on activity coefficients, although the deviation from ideal behavior is still small. Next, we fixed one of the components and varied the other to find the mixture with the highest structural resemblance yet the most dissimilar activity coefficients, illustrated by Counterfactual 2 and Counterfactual 3. When fixing benzene, Counterfactual 2 shows that replacing the methyl group with a hydroxyl group, coupled with a change in composition, largely influences activity coefficients. This can be explained by the fact that removing the methyl group converts one of the components from nonpolar to polar, thus resulting in strong deviations from ideal behavior. Counterfactual 3 shows a similar tendency. When fixing toluene, the other component in the counterfactual tends to converge to a more polar chemical, such as pyridine which converts one of the carbons on the benzene ring to nitrogen.

On the other hand, Type II counterfactuals also reveal interesting trends to identify mixtures with dissimilar chemical structures but similar  $\gamma_i$ 's. When fixing one of the components, counterfactuals 4 and 5 both acquire an alternative component that is highly symmetric, such as divinyl sulfide and diethyl disulfide. In both cases, one of the nonpolar components is replaced by a polar molecule as the result of the effort to minimize similarity, but the mixtures exhibit near-ideal behavior because the nonpolar content dominates as reflected by the activity coefficients and the compositions. Lastly, when we relaxed the constraint and allowed both components to vary, counterfactual 6 picks out the mixture from the data set that shows two polar components with near-ideal behavior.

In general, the counterfactual analysis has shown coherent physical insights regarding how compositions and structural features may lead to variations in activity coefficient, and these findings in turn agree with our chemical understanding of mixture behavior. Such interpretation,

especially Type II counterfactuals, can be used to apply SolvGNN to procedures such as the selection of a candidate good solvent for a desired solute. For example, counterfactuals could be used to identify an antisolvent given a known good solvent for a specific polymer for polymer recycling applications [53,54]. The antisolvent is expected to be miscible with the solvent while immiscible with the polymer, and therefore counterfactual Type I may be identified as the candidate antisolvent.

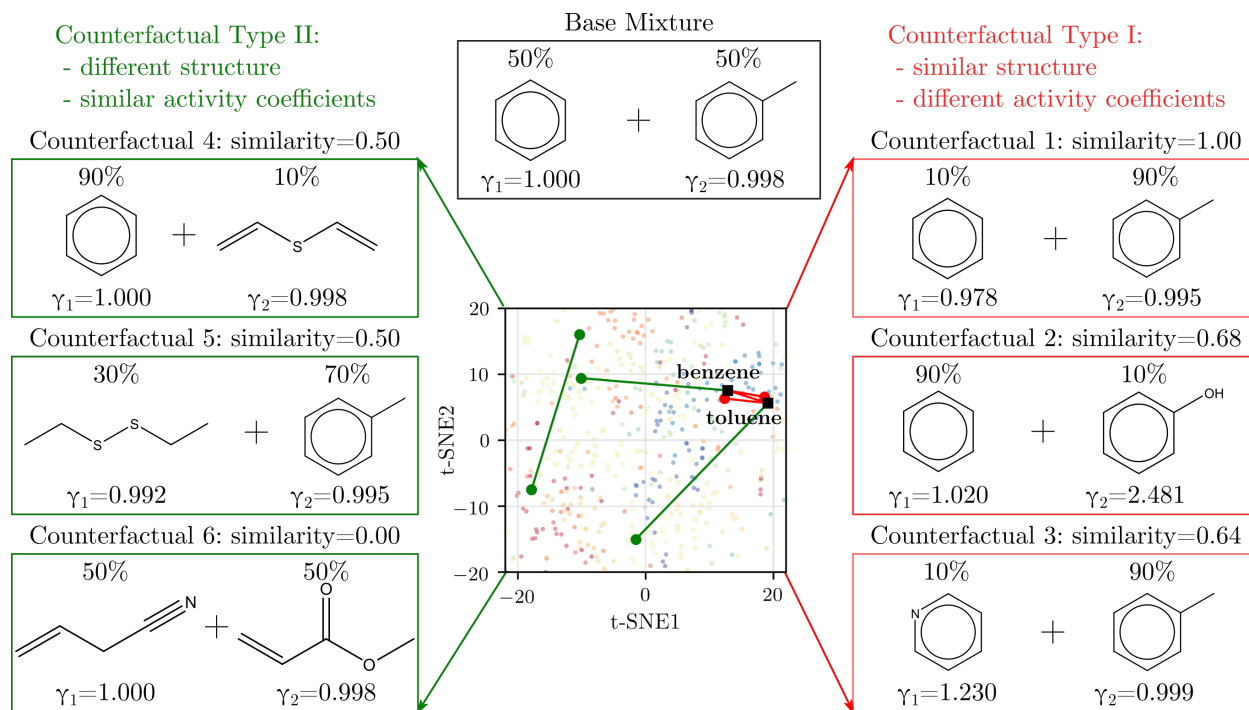


Figure 4: Counterfactual analysis. Type I (red) shows mixtures with the most similar structures but the most different activity coefficients from the base mixture whereas Type II (green) shows the opposite. The corresponding solvents are labeled on the 2D t-SNE map introduced in Figure 1 to help illustrate similarity.

### 3.5 From Activity Coefficients to Phase Diagrams

Although activity coefficients suggest the miscibility of different components in a mixture, they are not a direct measurement of equilibrium behavior (which might exhibit complex behavior such as the presence of azeotropes). Therefore, we further developed a framework to generate phase diagrams directly from chemical structures using the trained SolvGNN. These results aim to show the potential use of SolvGNN in industrially-relevant applications or experimental studies (e.g., miscibility or separation of target components). The framework uniformly samples the compositions of the input mixtures and predicts the corresponding activity coefficients using SolvGNN, which are then used for calculating equilibrium bubble and dew pressures via modified Raoult's Law. Figure 5 showcases several phase diagrams generated from the framework for both binary

and ternary mixtures. We would like to point out that most of the shown mixtures (all except for water-methanol) are not in our training or validation data, so they can be viewed as additional test instances, in spite of the fact that they are commonly used as mixture examples with contrasting equilibrium behavior.

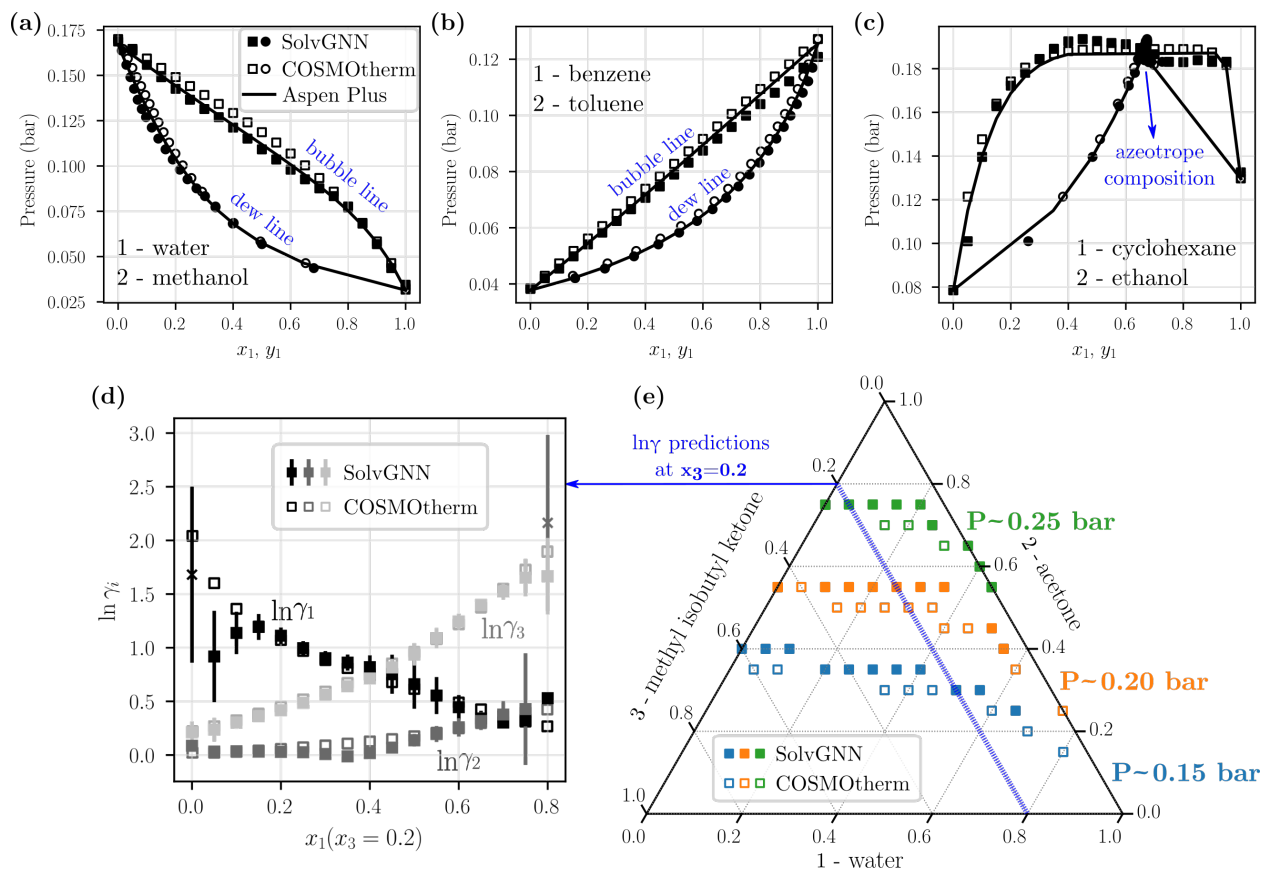


Figure 5: Example phase diagrams generated from SolvGNN. (a-c) P-x-y diagrams of three binary mixtures, each representing a different type of mixture (polar-polar, nonpolar-nonpolar, and polar-nonpolar). The equilibrium pressure is computed with the modified Raoult’s Law that uses the predicted activity coefficients from SolvGNN. The phase diagrams are compared with those generated from two other state-of-the-art tools, including COSMOtherm that implements COSMORS [52] and Aspen Plus that implements UNIFAC [55] (as well as other activity models [56–60]). The vapor compositions are represented as circles and liquid compositions are represented as squares. The phase diagram for an example ternary mixture (e) is calculated in the similar way to indicate isobaric curves for liquid-liquid equilibrium. Here, Aspen Plus (UNIFAC) fails to converge and therefore only COSMOtherm is used for comparison. Additionally, (d) displays the estimated activity coefficients for individual components at different  $x_1$  and  $x_2$  compositions while fixing  $x_3 = 0.2$  from SolvGNN and COSMOtherm. In all of phase diagram calculations, the  $\ln \gamma_i$ ’s are obtained by averaging the predictions of SolvGNN trained from each CV fold, and the standard deviations are visualized as the error bars.



Figure 5a-c includes representative examples of polar-polar, nonpolar-nonpolar, and polar-nonpolar binary mixtures. At 298K, a water-methanol mixture deviates positively from ideal solution behavior and shows higher equilibrium bubble pressure as a result of unfavorable unlike-molecule interactions. By contrast, a benzene-toluene mixture exhibits near-ideal behavior, as indicated by a bubble line that is almost linear, which suggests a homogeneous solution where molecular interactions between like and unlike components are viewed the same. Additionally, we showcase a cyclohexane-ethanol mixture that forms an azeotrope, which was successfully identified by SolvGNN, and the predicted azeotrope composition ( $x_{cyclohexane} \sim 0.65$ ) is consistent with the estimates from COSMOtherm and Aspen Plus. In all three cases, the predicted phase diagrams obtained by SolvGNN are consistent with the phase diagrams generated using COSMOtherm (COSMO-RS) or Aspen Plus (UNIFAC), and the MAE values in the equilibrium pressure range from 0.002 to 0.005 bar. We also observed that, compared to Aspen Plus (UNIFAC), SolvGNN tends to underestimate equilibrium pressure values, whereas COSMOtherm tends to overestimate these values.

Figure 5d demonstrates conditions for liquid-liquid equilibrium (LLE) for a ternary mixture of water-acetone-methyl isobutyl ketone (MIBK). Similar to the phase diagram calculations for binary systems, we sampled different liquid compositions and calculated the equilibrium pressures using modified Raoult’s Law. The conventional way to report phase equilibrium for ternary systems is by specifying a temperature and pressure and solving for the equilibrium compositions. However, we instead computed the equilibrium pressure from sampled compositions to match the procedure used in the binary phase diagrams. For visual clarity, we picked three pressures and showed the corresponding liquid compositions on each isobaric curve. The generated phase diagram was again compared with the one generated by COSMOtherm; in general, the shapes of the curves match, yet the LLE points from SolvGNN appear above the LLE points from COSMOtherm, suggesting that SolvGNN tends to converge to a lower pressure value than COSMOtherm. Upon inspecting the activity coefficients at  $x_{MIBK} = 0.2$ , we found that, although the activity coefficients were trained only on four compositions, SolvGNN was able to make relatively accurate predictions for compositions in a continuous space, with a few exceptions at the extreme points. For example, when  $x_{water}$  is close to 0, the corresponding  $\ln \gamma_{water}$  deviates from that of COSMO-RS. This could be a result of the lack of extreme points in the training samples, since we only sampled compositions within 15% and 70%. A similar pattern is present at  $x_{water} = 0.8(x_{acetone} = 0)$  for  $\ln \gamma_{acetone}$  (dark gray). Nevertheless, these extreme points are multiplied by the corresponding liquid compositions (likely zero) in modified Raoult’s Law, thus the effect of such errors is minimal in phase equilibrium calculations.

In summary, we were able to create phase diagrams (at 298K) with a full range of compositions using SolvGNN that was only trained on a few sampled input ratios. The provided framework has shown great potential for high-throughput screening of mixtures for use cases including azeotrope identification and non-ideal behavior investigation for liquid mixtures. Incorporating SolvGNN

into such phase equilibrium calculations bypasses the need to identify functional groups with human expertise and obtain interaction parameters (as needed in UNIFAC) or to conduct DFT calculations (as needed in COSMO-RS), especially when the chemicals in a mixture are relatively uncommon. Moreover, this framework could be used in conjunction with open-source process models (*e.g.*, BioSteam [61]) as an addition to the existing computational models (*e.g.*, UNIFAC) for generating thermodynamic data.

## 4 Conclusions and Future Outlook

We developed a GNN architecture (SolvGNN) that incorporates both local (intramolecular) and global (intermolecular) convolutions on graph representations and used this for predicting activity coefficients of solvent mixtures. SolvGNN explicitly integrates intermolecular interactions through the construction of the molecular interaction network that encodes H-bonding information. We found that with such feature embedding, SolvGNN can successfully estimate the general activity coefficients that vary with chemical compositions for binary as well as ternary mixtures, which has not been explored much under the hood of ML, especially in the context of activity coefficients.

Compared to the current state-of-the-art approach for general activity coefficient estimations (*e.g.*, UNIFAC and COSMO-RS), SolvGNN achieves comparable model performance and is easy to use without any additional calculations for missing parameters or DFT. We also benchmarked SolvGNN on the same experimental dataset that was used in an earlier study for developing a GNN that predicts infinite-dilution activity coefficients (with no composition dependence) of binary mixtures [28]; SolvGNN (without composition information) outperforms the previously developed GNN in almost all evaluation metrics, proving the importance to use prior knowledge (in this case explicit topological prior pertinent to intermolecular interactions) when designing GNN architectures. These findings demonstrate the ability of SolvGNN to learn from simulation (*e.g.*, COSMO-RS) and experimental data.

Moreover, we provided an open-source computational tool for creating phase diagrams (P-x-y) using SolvGNN as an example to show its potential for real-world applications. The generated phase diagrams were consistent with those obtained from COSMOtherm and Aspen Plus (with the selection of UNIFAC as the thermodynamic method), which further illustrated the generalization ability of SolvGNN that was only trained on a minimal subset of composition cases. Besides phase diagrams, we provided algorithms to obtain counterfactuals to aid model interpretation, which may help extract physical insights that are less known and help design solvent mixtures.

The architecture and study can be expanded in a number of ways. For example, so far we have only obtained activity coefficients at room temperature, and thus SolvGNN does not have temperature dependence. However, obtaining temperature-dependent activity coefficients from

COSMO-RS and re-training SolvGNN with an additional temperature variable would be a relatively trivial, given that the computational framework is in place. Another limitation for phase equilibrium predictions is related to the availability of Antoine coefficients; in circumstances where Antoine coefficients are missing, we cannot compute the corresponding phase diagrams. A potential solution could be to develop another GNN architecture for Antoine coefficient predictions or expand the output dimension of our current SolvGNN to make these predictions. Furthermore, the presented counterfactual analysis only searches the chemical space within the data set, and therefore to obtain more meaningful results, we will adapt some of the more established chemical search algorithms [37, 62] that have been designed for single chemicals to the case of mixtures. Future studies will also explore the use of SolvGNN for other mixture properties and investigate different possible representations of intermolecular interactions (*e.g.*, Lennard-Jones potentials). We are also interested in using these types of architectures to design solvents that can selectively solubilize target molecules.

## 5 Acknowledgements

All authors of this manuscript are committed to supporting and advancing women in chemistry. R. Van Lehn and V.M. Zavala are co-supervisors of the lead author(S. Qin) in her PhD studies and are committed to provide a healthy environment that allows her to achieve her professional career and personal goals. This is achieved via careful mentoring and gathering of frequent feedback.

This material is based on work supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, under Contract DE-AC02-06CH11357.

## 6 Data Availability

All the data and scripts are available as open-source code at <https://github.com/zavalab/ML/tree/SolvGNN>.

## References

- [1] Jarmo Huuskonen, Marja Salo, and Jyrki Taskinen. Aqueous solubility prediction of drugs based on molecular topology and neural network modeling. Journal of chemical information and computer sciences, 38(3):450–456, 1998.
- [2] Alessandro Lusci, Gianluca Pollastri, and Pierre Baldi. Deep architectures and deep learning in chemoinformatics: the prediction of aqueous solubility for drug-like molecules. Journal of chemical information and modeling, 53(7):1563–1575, 2013.
- [3] Samuel Boobier, David RJ Hose, A John Blacker, and Bao N Nguyen. Machine learning with physicochemical relationships: solubility prediction in organic solvents and water. Nature communications, 11(1):1–10, 2020.
- [4] Andreas Mayr, Günter Klambauer, Thomas Unterthiner, and Sepp Hochreiter. Deeptox: toxicity prediction using deep learning. Frontiers in Environmental Science, 3:80, 2016.
- [5] Priyanka Banerjee, Andreas O Eckert, Anna K Schrey, and Robert Preissner. Protox-ii: a webserver for the prediction of toxicity of chemicals. Nucleic acids research, 46(W1):W257–W263, 2018.
- [6] Jian Jiang, Rui Wang, and Guo-Wei Wei. Ggl-tox: Geometric graph learning for toxicity prediction. Journal of chemical information and modeling, 61(4):1691–1700, 2021.
- [7] Timon Schroeter, Anton Schwaighofer, Sebastian Mika, Antonius Ter Laak, Detlev Suelzle, Ursula Ganzer, Nikolaus Heinrich, and Klaus-Robert Müller. Machine learning models for lipophilicity and their domain of applicability. Molecular pharmaceutics, 4(4):524–538, 2007.
- [8] Bowen Tang, Skyler T Kramer, Meijuan Fang, Yingkun Qiu, Zhen Wu, and Dong Xu. A self-attention based message passing neural network for predicting molecular lipophilicity and aqueous solubility. Journal of cheminformatics, 12(1):1–9, 2020.
- [9] Harry L Morgan. The generation of a unique machine description for chemical structures—a technique developed at chemical abstracts service. Journal of chemical documentation, 5(2):107–113, 1965.
- [10] David Rogers and Mathew Hahn. Extended-connectivity fingerprints. Journal of chemical information and modeling, 50(5):742–754, 2010.
- [11] Mati Karelson, Victor S Lobanov, and Alan R Katritzky. Quantum-chemical descriptors in qsar/qspr studies. Chemical reviews, 96(3):1027–1044, 1996.
- [12] Hiroto Moriawaki, Yu-Shi Tian, Norihito Kawashita, and Tatsuya Takagi. Mordred: a molecular descriptor calculator. Journal of cheminformatics, 10(1):1–14, 2018.

- [13] Yu-Chen Lo, Stefano E Rensi, Wen Torng, and Russ B Altman. Machine learning in chemoinformatics and drug discovery. *Drug discovery today*, 23(8):1538–1546, 2018.
- [14] Jennifer N Wei, David Duvenaud, and Alán Aspuru-Guzik. Neural networks for the prediction of organic chemistry reactions. *ACS central science*, 2(10):725–732, 2016.
- [15] Connor W Coley, Wengong Jin, Luke Rogers, Timothy F Jamison, Tommi S Jaakkola, William H Green, Regina Barzilay, and Klavs F Jensen. A graph-convolutional neural network model for the prediction of chemical reactivity. *Chemical science*, 10(2):370–377, 2019.
- [16] Anirudh Raju Natarajan and Anton Van der Ven. Machine-learning the configurational energy of multicomponent crystalline solids. *npj Computational Materials*, 4(1):1–7, 2018.
- [17] Xiaodi Liu, Xin Li, Quanfeng He, Dandan Liang, Ziqing Zhou, Jiang Ma, Yong Yang, and Jun Shen. Machine learning-based glass formation prediction in multicomponent alloys. *Acta Materialia*, 201:182–190, 2020.
- [18] Liam Wilbraham, Reiner Sebastian Sprick, Kim E Jelfs, and Martijn A Zwijnenburg. Mapping binary copolymer property space with neural networks. *Chemical science*, 10(19):4973–4984, 2019.
- [19] Kyohei Hanaoka. Deep neural networks for multicomponent molecular systems. *ACS omega*, 5(33):21042–21053, 2020.
- [20] Yong Pan, Xianke Ji, Li Ding, and Juncheng Jiang. Prediction of lower flammability limits for binary hydrocarbon gases by quantitative structure—property relationship approach. *Molecules*, 24(4):748, 2019.
- [21] Subhash Ajmani, Stephen C Rogers, Mark H Barley, and David J Livingstone. Application of qspr to mixtures. *Journal of chemical information and modeling*, 46(5):2043–2055, 2006.
- [22] Alan R Katritzky, Iva B Stoyanova-Slavova, Kaido Tamm, Tarmo Tamm, and Mati Karelson. Application of the qspr approach to the boiling points of azeotropes. *The Journal of Physical Chemistry A*, 115(15):3475–3479, 2011.
- [23] Alexander A Oliferenko, Polina V Oliferenko, José S Torrecilla, and Alan R Katritzky. Boiling points of ternary azeotropic mixtures modeled with the use of the universal solvation equation and neural networks. *Industrial & engineering chemistry research*, 51(26):9123–9128, 2012.
- [24] Ting Wang, Lili Tang, Feng Luan, and MNDS Cordeiro. Prediction of the toxicity of binary mixtures by qsar approach using the hypothetical descriptors. *International journal of molecular sciences*, 19(11):3423, 2018.
- [25] Guillaume Fayet and Patricia Rotureau. New qspr models to predict the flammability of binary liquid mixtures. *Molecular informatics*, 38(8-9):1800122, 2019.

- [26] Sivadurgaprasad Chinta and Raghunathan Rengaswamy. Machine learning derived quantitative structure property relationship (qspr) to predict drug solubility in binary solvent systems. Industrial & Engineering Chemistry Research, 58(8):3082–3092, 2019.
- [27] Fabian Jirasek, Rodrigo AS Alves, Julie Damay, Robert A Vandermeulen, Robert Bamler, Michael Bortz, Stephan Mandt, Marius Kloft, and Hans Hasse. Machine learning in thermodynamics: Prediction of activity coefficients by matrix completion. The journal of physical chemistry letters, 11(3):981–985, 2020.
- [28] Edgar Ivan Sanchez Medina, Steffen Linke, Martin Stoll, and Kai Sundmacher. Graph neural networks for the prediction of infinite dilution activity coefficients. Digital Discovery, 2022.
- [29] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. Advances in neural information processing systems, 28, 2015.
- [30] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In International Conference on Machine Learning, pages 1263–1272. PMLR, 2017.
- [31] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. AI Open, 1:57–81, 2020.
- [32] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. Moleculenet: a benchmark for molecular machine learning. Chemical science, 9(2):513–530, 2018.
- [33] Shiyi Qin, Tianyi Jin, Reid C Van Lehn, and Victor M Zavala. Predicting critical micelle concentrations for surfactants using graph convolutional neural networks. The Journal of Physical Chemistry B, 125(37):10610–10620, 2021.
- [34] Harrison C Carlson and Allan P Colburn. Vapor-liquid equilibria of nonideal solutions. Industrial & Engineering Chemistry, 34(5):581–589, 1942.
- [35] Julie Damay, Fabian Jirasek, Marius Kloft, Michael Bortz, and Hans Hasse. Predicting activity coefficients at infinite dilution for varying temperatures by matrix completion. Industrial & Engineering Chemistry Research, 60(40):14564–14578, 2021.
- [36] Hania Benimam, Cherif Si-Moussa, Maamar Laidi, and Salah Hanini. Modeling the activity coefficient at infinite dilution of water in ionic liquids using artificial neural networks and support vector machines. Neural Computing and Applications, 32(12):8635–8653, 2020.
- [37] Geemi P Wellawatte, Aditi Seshadri, and Andrew D White. Model agnostic generation of counterfactual explanations for molecules. Chemical Science, 2022.

- [38] Charles M Hansen. Hansen solubility parameters: a user's handbook. CRC press, 2007.
- [39] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. Journal of machine learning research, 9(11), 2008.
- [40] Andreas Klamt. Conductor-like screening model for real solvents: a new approach to the quantitative calculation of solvation phenomena. The Journal of Physical Chemistry, 99(7):2224–2235, 1995.
- [41] Release COSMOtherm. 19; cosmologic gmbh & co. kg, 2019.
- [42] Nci/cadd chemical identifier resolver.
- [43] Sree Ganesh Balasubramani, Guo P Chen, Sonia Coriani, Michael Diedenhofen, Marius S Frank, Yannick J Franzke, Philipp Furche, Robin Grotjahn, Michael E Harding, Christof Hättig, et al. Turbomole: Modular program suite for ab initio quantum-chemical and condensed-matter simulations. The Journal of chemical physics, 152(18):184107, 2020.
- [44] Mufei Li, Jinjing Zhou, Jiajing Hu, Wenxuan Fan, Yangkang Zhang, Yaxin Gu, and George Karypis. Dgl-lifesci: An open-source toolkit for deep learning on graphs in life science. ACS omega, 6(41):27233–27238, 2021.
- [45] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907, 2016.
- [46] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. arXiv preprint arXiv:1409.1259, 2014.
- [47] Greg Landrum. Rdkit documentation. Release, 1(1-79):4, 2013.
- [48] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems, 32, 2019.
- [49] Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma, Lingfan Yu, Yu Gai, et al. Deep graph library: A graph-centric, highly-performant package for graph neural networks. arXiv preprint arXiv:1909.01315, 2019.
- [50] Dávid Bajusz, Anita Rácz, and Károly Héberger. Why is tanimoto index an appropriate choice for fingerprint-based similarity calculations? Journal of cheminformatics, 7(1):1–13, 2015.
- [51] Oscar Contreras. Nist-web-book-scraping. <https://github.com/oscarcontrerasnavas/NIST-web-book-scraping>, 2019.

- [52] Andreas Klamt, Frank Eckert, and Wolfgang Arlt. Cosmo-rs: an alternative to simulation for calculating thermodynamic properties of liquid mixtures. Annual review of chemical and biomolecular engineering, 1:101–122, 2010.
- [53] Theodore W Walker, Nathan Frelka, Zhizhang Shen, Alex K Chew, Jesse Banick, Steven Grey, Min Soo Kim, James A Dumesic, Reid C Van Lehn, and George W Huber. Recycling of multilayer plastic packaging materials by solvent-targeted recovery and precipitation. Science advances, 6(47):eaba7599, 2020.
- [54] Panzheng Zhou, Kevin L Sánchez-Rivera, George W Huber, and Reid C Van Lehn. Computational approach for rapidly predicting temperature-dependent polymer solubilities using molecular-scale models. ChemSusChem, 14(19):4307–4316, 2021.
- [55] Aage Fredenslund, Russell L Jones, and John M Prausnitz. Group-contribution estimation of activity coefficients in nonideal liquid mixtures. AIChE Journal, 21(6):1086–1099, 1975.
- [56] MS Margules. On the composition of saturated vapors of mixtures. Akademie der Wissenschaften in Wien, Mathematisch-Naturwissenschaftliche Klasse Abteilung II, 104:1234–1239, 1895.
- [57] JJ Van Laar. Über dampfspannungen von binären gemischen. Zeitschrift für Physikalische Chemie, 72(1):723–751, 1910.
- [58] Grant M Wilson. Vapor-liquid equilibrium. xi. a new expression for the excess free energy of mixing. Journal of the American Chemical Society, 86(2):127–130, 1964.
- [59] Henri Renon and JM Prausnitz. Estimation of parameters for the nrtl equation for excess gibbs energies of strongly nonideal liquid mixtures. Industrial & Engineering Chemistry Process Design and Development, 8(3):413–419, 1969.
- [60] Denis S Abrams and John M Prausnitz. Statistical thermodynamics of liquid mixtures: a new expression for the excess gibbs energy of partly or completely miscible systems. AIChE Journal, 21(1):116–128, 1975.
- [61] Yoel Cortes-Peña, Deepak Kumar, Vijay Singh, and Jeremy S Guest. Biosteam: a fast and flexible platform for the design, simulation, and techno-economic analysis of biorefineries under uncertainty. ACS Sustainable Chemistry & Engineering, 8(8):3302–3310, 2020.
- [62] AkshatKumar Nigam, Robert Pollice, Mario Krenn, Gabriel dos Passos Gomes, and Alan Aspuru-Guzik. Beyond generative models: superfast traversal, optimization, novelty, exploration and discovery (stoned) algorithm for molecules using selfies. Chemical science, 12(20):7079–7090, 2021.