



DeepEOSNet: Capturing the dependency on thermodynamic state in property prediction tasks

Jan Pavšek^a, Alexander Mitsos^{b,a,c}, Manuel Dahmen^c, Tai Xuan Tan^a, Jan G. Rittig^a,*

^a Process Systems Engineering (AVT.SVT), RWTH Aachen University, Forckenbeckstraße 51, 52074 Aachen, Germany

^b JARA-CSD, Templergraben 55, 52056 Aachen, Germany

^c Institute of Climate and Energy Systems ICE-1: Energy Systems Engineering, Forschungszentrum Jülich GmbH, Wilhelm-Johnen-Straße, 52425 Jülich, Germany

ARTICLE INFO

Dataset link: [NIST Thermodata engine](#), [GMOlpr op](#)

Keywords:

Molecular machine learning
Physicochemical properties
Hybrid modeling
Thermodynamics
Property prediction

ABSTRACT

We propose a machine learning (ML) architecture to better capture the dependency of thermodynamic properties on independent states. When predicting state-dependent thermodynamic properties, ML models need to account for both molecular structure and the thermodynamic state, described by independent variables, typically temperature, pressure, and composition. Modern molecular ML models include state information by adding it to molecular fingerprint vectors or by embedding explicit (semi-empirical) thermodynamic relations. We propose splitting the information processing on the molecular structure and the dependency on states into two separate network channels: a graph neural network and a multilayer perceptron, whose output is combined by a dot product. We refer to our approach as DeepEOSNet, as it is based on the DeepONet architecture (Lu et al., 2021): instead of operators, we learn state dependencies, with the possibility to predict equation of states (EOS). We investigate the predictive performance of DeepEOSNet in three case studies, including the prediction of vapor pressure as a function of temperature, and mixture molar volume as a function of composition, temperature, and pressure. Our results show superior performance of DeepEOSNet for predicting vapor pressure and comparable performance for predicting mixture molar volume compared to state-of-research graph-based thermodynamic prediction models from our earlier works. In fact, we see large potential of DeepEOSNet when data is sparse in the state domain and the output function is structurally similar across different molecules. The concept of DeepEOSNet can easily be transferred to other deep learning architectures, providing a viable option for molecular property prediction.

1. Introduction

Molecular machine learning (ML) has advanced thermodynamic property prediction, promising to accelerate the identification of molecules with desired properties and thereby more efficient chemical processes (Rittig et al., 2025). Various ML methods, such as graph neural networks (GNN) (Santana et al., 2024; Rittig and Mitsos, 2024; Lansford et al., 2023; Hoffmann et al., 2025), transformers (Winter et al., 2023, 2025, 2022) and matrix completion methods (Damay et al., 2021) have shown high predictive accuracy for various thermodynamic properties of both pure components and mixtures, such as normal boiling points (Rittig et al., 2023b), solubility (Vermeire and Green, 2021), and activity coefficients (Damay et al., 2021; Rittig and Mitsos, 2024; Winter et al., 2023; Medina et al., 2022). The predictive capabilities of these methods have been further improved by utilizing thermodynamic knowledge within model training and/or as part of the model architecture (Rittig et al., 2023a; Rittig and

Mitsos, 2024; Hoffmann et al., 2025; Santana et al., 2024; Lansford et al., 2023), thereby ensuring thermodynamic consistency while preserving ML model flexibility, cf. overview in Rittig et al. (2025). Therefore, molecular ML can provide property predictions with high accuracy respecting thermodynamic principles and is highly promising for chemical engineering applications.

For state-dependent thermodynamic properties, it is critical to also account for information on the independent state variables, typically how they vary with temperature, pressure, and composition. As such, molecular ML models have been extended in various cases, including the prediction of vapor–liquid equilibria as function of composition and pressure/temperature (Sun et al., 2025; Medina and Sundmacher, 2026), vapor pressures as function of temperature (Santana et al., 2024; Lansford et al., 2023; Hoffmann et al., 2025; Lin et al., 2024; Winter et al., 2025), activity coefficients as function of composition (Rittig et al., 2023a; Rittig and Mitsos, 2024), temperature (Rittig et al., 2023c;

* Corresponding author.

E-mail address: jan.rittig@rwth-aachen.de (J.G. Rittig).

Winter et al., 2022; Damay et al., 2021; Medina et al., 2023), and both (Winter et al., 2023; Specht et al., 2024), density as function of temperature (Felton et al., 2024; Winter et al., 2025), Helmholtz free energy as function of density and temperature (Chaparro and Müller, 2024, 2023), and Gibbs free energy as function of temperature (Ham-mad and Mondal, 2025). Two different approaches to capture the state dependency are frequently used in molecular ML:

Concatenation: The state variables are concatenated to the molecular fingerprint vectors, learned by the ML models, and are then considered as an additional input to a multilayer perceptron (MLP) for property prediction, e.g., in Rittig et al. (2023c), Brozos et al. (2025), Aouichaoui et al. (2023).

Embedding of semi-empirical equations: Semi-empirical models that provide consecutive equations of states are incorporated into the ML architecture, in the form of hybrid models, e.g., in Medina et al. (2023), Winter et al. (2023), Hoffmann et al. (2025).

A challenge of simply concatenating molecular fingerprint vectors and state information, as in the *concatenation* approach, is that the MLP has to simultaneously learn the molecular structure and state dependency of the property of interest purely from the provided data, and thus in some sense distinguish between the two effects. Also physics-informed ML models that use fundamental thermodynamics have to learn the state dependency purely from data: While fundamental thermodynamics provides algebraic/differential relationships (e.g., Gibbs–Helmholtz equation and Maxwell relations) between properties and stability criteria that all state functions must satisfy in equilibrium, they do not specify the state dependency of properties for specific molecules/mixtures explicitly. The determination of state dependencies – whether $(\partial C_p/\partial P)_T$ is positive or negative, or how fugacity coefficients vary with temperature – requires molecular structure-specific information, and is typically based on statistical mechanics, molecular simulations, or experimental measurements. Notably, in physics-informed ML models using fundamental thermodynamics, learning the state dependencies is not only needed for making state dependent predictions but also for enforcing the fundamental relationships accurately. As for the *embedding of semi-empirical equations* approach, semi-empirical models typically parametrize the state dependency in few molecule- or mixture-specific parameters, thereby significantly reducing the number of parameters that need to be learned in ML model training (in comparison to an MLP). This potentially reduces data demands for training, but also relies on simplifying assumptions, such as using ideal gas law for the vapor phase in the Antoine equation (Pfennig, 2013), hence comes with the limitations of semi-empirical models.

We propose an alternative molecular ML approach to simultaneously learn the molecule- and the state dependency of thermodynamic properties, without introducing modeling limitations and thus preserving ML model flexibility. To this end, we do not rely on a semi-empirical thermodynamic model as in the *embedding of semi-empirical equations* approach; we rather propose a modification to the *concatenation* approach, inspired by DeepONets (deep operator networks) (Lu et al., 2021). We refer to our architecture as DeepEOSNet, as we learn state functions corresponding to equations of state (EOS) instead of operators, and therefore replace “O” by “EOS”.

Lending from the DeepONet structure, in DeepEOSNet we separate the learning of the molecular structure and state dependencies into two different channels. The molecular channel encodes structure-specific information into the molecular fingerprint vector that is relevant for the property of interest; it can thus be based on any molecular ML model, such as transformers and GNNs. The aim of the state channel is to capture the general functional form of the influence of the independent state variables on the property. The outputs of the two channels are then combined by a dot product to give the final property prediction. In this work, we apply our approach to learning individual properties,

i.e., vapor pressure and total molar volume of mixtures. We stress that the DeepEOSNet approach is also directly transferable to learning fundamental properties, e.g., Gibbs free energy, as introduced by Rittig and Mitsos (2024).

2. Fundamentals

We first describe the concepts of DeepONets for operator learning and GNNs for molecular property prediction, on which we then build our proposed DeepEOSNet architecture.

2.1. DeepONet for operator learning

DeepONets, introduced by Lu et al. (2021), are a deep neural network architecture designed to learn operators, i.e., a mapping from one function $f(\cdot)$ to another $g(\cdot)$, which they apply to ordinary differential equations (ODEs) and partial differential equations (PDEs). The DeepONet architecture is illustrated in Fig. 1: Two separate neural networks, referred to as branch and trunk net, are employed for splitting information processing. The branch net processes the information of the input function, whereas the trunk net processes information on the domain of the output function. The results of these networks are then combined by a dot product, such that the output can be described by

$$g(x, f(y)) = \sum_{i=1}^{N_{hd}} t_i(x) b_i(f(y)), \quad (1)$$

where $g(\cdot)$ denotes the output function, N_{hd} is the hidden dimension of the branch and trunk net, b_i is the output of the i -th node of the branch net, t_i is the output of the i -th node of the trunk net, $f(\cdot)$ is the input function, y is the independent variable of the input function, and x is the independent variable of the output function. In the example of PDEs, the branch net processes the PDE with specific initial/boundary conditions, whereas the trunk net processes the evaluation points. As such, DeepONets can learn to solve the PDE for different input conditions, rather than at a single condition, hence they learn solution operators.

Practically, such DeepONets approximate the function mapping by predicting a target value $g(x, f(y))$ that is dependent on the two input values $f(y)$ and x . At the same time, correlation between $f(y)$ and x might be present, in case x and y are correlating or data sampling is unequally distributed across the domain. Separating the information processing into branch and trunk net appears to mitigate the effects of potential correlations in the data points of $f(\cdot)$ and x might have on training. The separate networks ensure that the influence each of the inputs $f(y)$ and x have on the target $g(x, f(y))$ is separated from the other input. At the same time this architecture is more versatile during inference, as the evaluation of $g(x, f(y))$ at new domain points x only requires passing through the trunk net in case $f(y)$ is constant, and vice versa. Notably, separate processing of inputs has also been explored in other domains, such as image/video processing using SplitNets (Dong et al., 2022).

2.2. Graph neural networks for molecular property prediction

GNNs learn properties directly from a graph representation of molecules, as illustrated in Fig. 2. For a detailed overview of GNN architectures in the molecular context, we refer the interested reader to Schweidtmann et al. (2020), Rittig et al. (2023c), Reiser et al. (2022), Gilmer et al. (2017). The general working principle of GNN is as follows: Graph convolutional layers first encode structural information from the molecular graph, which results in a continuous vector representation of the molecule, referred to as molecular fingerprint vector. This molecular fingerprint vector is then passed through an MLP which provides a prediction for the target property $\hat{p}$.




Fig. 1. Schematic illustration of the DeepONet architecture.




Fig. 2. Architecture of a graph neural network for predicting state-dependent thermodynamic properties.




Fig. 3. Architecture of DeepEOSNet.

In case the prediction target additionally depends on at least one independent thermodynamic state variable – typically temperature, pressure, and/or composition –, these additional state variables s need to be accounted for in addition to the molecular structure. Most commonly, the state variables s are concatenated with the molecular fingerprint vector (Rittig et al., 2023c; Aouichaoui et al., 2023), hence, they serve as additional input to the MLP. The MLP then simultaneously learns the molecular structure and state dependencies.

3. DeepEOSNet for thermodynamic properties

We propose DeepEOSNet, an ML architecture for state-dependent thermodynamic property prediction. DeepEOSNet splits the information processing of the molecular structure input and the state variables into two separate network channels, as shown in Fig. 3. The branch net is a GNN, that extracts structural information from the molecular graph, resulting in a molecular fingerprint vector. The trunk net is an MLP, which takes the state variable s as input and outputs a continuous vector with the same dimension as the molecular fingerprint vector. The two output vectors are then combined by a dot product to yield the final property prediction $\hat{p}^{\text{mol}}(s)$.

By separating the information processing of the molecular structure and state variables, a stronger inductive bias is introduced to learn

the influences on the molecular property separately: The state channel (trunk net) does not have the molecular structure as an input, and thus learns a general functional form of the property as function of states. Vice versa, the molecule channel (branch net) does not have direct access to the state variables, and therefore extracts molecular structure-specific information. We stress that the molecule channel can still learn any state-dependent information. In fact, during model training, the molecule channel can learn from state-dependent differences between the predicted and actual values (i.e., the experimental or simulated property values). This allows the molecule channel to capture the molecule-specific state dependency within the general form of state dependency induced by the state channel.

Interestingly, employing a dot product based calculation of the final output (see Eq. (2)) is conceptually similar to the commonly employed representation using summation over simple basis functions used in thermodynamics to capture the dependency on independent states. A well-known example of this is the NASA polynomial representation (McBride, 2002) (see Eq. (3)) of thermodynamic properties, which is commonly used to capture temperature dependency:

$$\hat{p}_{\text{DeepEOSNet}}^{\text{mol}}(s) = \sum_{i=1}^{N_{\text{hd}}} t_i(s) * b_i(G^{\text{mol}}) \quad (2)$$

$$\hat{p}_{\text{NASA}}^{\text{mol}}(T) = a_1 T^{-2} + a_2 T^{-1} + a_3 + a_4 T + a_5 T^2 + a_6 T^3 + a_7 T^4 \quad (3)$$

where G^{mol} is a molecular graph, T is the temperature and a_i are exemplary molecule parameters for a seventh-order NASA polynomial. Specifically, the temperature polynomials correspond to $t_i(s)$ and the molecule-specific parameters to $b_i(G^{\text{mol}})$. Therefore, the output layer of the DeepEOSNet architecture is essentially a learnable and more general form of the NASA polynomial, as $t_i(s)$ is not constrained to polynomials but rather the MLP which is a universal function approximator.

Notably, the same functional behavior, as given by the structure of the DeepEOSNet, can also be learned by the molecular ML architectures that simply concatenate the state variable to the molecular fingerprint vector (see Fig. 2 and approach (1) in the introduction). Yet, constraining the network architecture in DeepEOSNet is expected to facilitate training, as it reduces the search space of functions that can be learned for each molecule specifically to the general functional form induced by the state channel (trunk net) (Wilson and Izmailov, 2020). We hypothesize that this is particularly helpful for properties for which the general functional form of the state dependency can be easily approximated by the state channel, which is likely if state dependencies are similar for the different molecules of interest.

We further note that Sun et al. (2023) have proposed DeepGraPhONet which modifies DeepONet by employing a GNN as branch net, and is thus similar to our architecture. However, in their application, namely the prediction of system behavior of power grids and traffic flow, the input to the branch net $f(\cdot)$ (see Fig. 1) is a graph which itself is a function of the input to the trunk net $\mathbf{x}$. That is, the graph structure itself depends on the output domain, in this case the time. In contrast, in DeepEOSNet, the input to the branch net, which is the molecular graph, does not depend on the state feature vector, cf. Fig. 3. Therefore, the input to the branch net in DeepEOSNet is not a function. We thus rather learn state *functions* instead of operators, over the discrete space of molecules/mixtures. Further, as the focus of our approach is the application of separate information processing to state dependencies in molecular prediction tasks in general, we refer to our approach as DeepEOSNet.

4. Property prediction case studies

For testing the DeepEOSNet approach, we consider two target properties of high practical relevance for chemical engineering applications, namely vapor pressure and total molar volume of binary mixtures. Note that traditionally the total molar volume of mixtures is calculated by summing up the pure component molar volumes weighted with their mole fractions and adding the excess volume of mixing (since the ideal mixing volume is zero). However, here we directly predict the total molar volume of the mixture. The reason for this is that we want to avoid the introduction of modeling assumptions that potentially limit predictive performance. Predicting first the behavior of ideal mixtures and then the excess volume might introduce an inductive bias for the model towards simplified ideal mixture behavior. To avoid this, we directly train on the total molar volume of the mixture in an end-to-end fashion. In the following, we use the terms total molar volume of the mixture and mixture molar volume interchangeably.

4.1. Dataset

We extract property data from the NIST ThermoData Engine (National Institute of Standards and Technology (NIST), 2025). The **vapor pressure** dataset contains 879 molecules with vapor pressures at varying temperatures in the range from 78 K to 875 K, amounting to 71,730 data points in total. Notably, the number of data points with different temperatures is unevenly distributed among the molecules, ranging from a few data points for some molecules to thousands for others. The **molar volume** dataset contains 265 molecules combined into 1444 unique binary mixtures with varying temperatures in the range from 201 K to 505 K, pressures in the range from 78 kPa to 291 MPa, and compositions between 0 and 1. The total number of molar volume data

points is 88,365. Again, the distribution of data points is rather uneven for the different mixtures, in particular with respect to temperature and pressure variation.

As both datasets exhibit largely varying numerical values for the property data points, the common logarithm of pressure and molar volume, respectively, is used as prediction target. The temperature and pressure input are linearly normalized to range 0–1 across the entire dataset, rather than per molecule and mixture. We thus report the error metrics in Section 5 on the logarithmic scale.

4.2. Prediction scenarios & baselines

Using the vapor pressure and mixture volume datasets, we create three thermodynamic property prediction case studies (see Table 1). These case studies assess the architectures' capabilities of learning the above mentioned state functions. Hence, the test sets contain unseen points in the state domain for each molecule and mixture rather than unseen molecules. The generalization capabilities of graph-based molecular encoders to unseen molecules have been extensively studied in Santana et al. (2024), Hoffmann et al. (2025), Rittig et al. (2023c). As we employ equivalent molecular encoders, we expect similar extrapolation behavior in the molecular domain, and hence focus rather on the extrapolation capabilities of the learned state functions in the state domain.

In the first two prediction tasks, the temperature-dependent vapor pressure $p^{\text{sat}}(T)$ is selected as target property. Specifically, we consider two scenarios:

$p_{\text{extra}}^{\text{sat}}(T)$: We select the 80% lowest temperature points for each molecule as training set and use the highest 20% as test set. For molecules where there are less than 5 temperature points in the dataset, all points are sorted to the training set. In doing so, we test the extrapolation capabilities in the temperature domain, i.e., to a temperature range not used in training at all, which serves as proxy for learning temperature dependency in general.

$p_{\text{SingleT}}^{\text{sat}}(T)$: We randomly select a single temperature point per molecule for training, while all remaining temperature points are sorted to the test set. This allows us to test the capabilities of models to distinguish between the molecular structure and state dependencies, as the variance of the two input dependencies overlaps entirely in this scenario. We note that this scenario does presumably not represent a practical scenario for experimental datasets and induces an unrealistic difficulty but is insightful to test the methodological model characteristics.

For the third prediction task, we consider the total molar volume of binary mixtures as a function of temperature, pressure, and composition. This scenario adds difficulty to the prediction in two ways: First, accounting for three state dependencies results in a three dimensional state variable vector $\mathbf{s}$ instead of a one-dimensional one as in the vapor pressure scenarios. Secondly, considering binary mixtures requires considering two molecular graphs as model input instead of a single one. For this, we use a GNN architecture from our previous works (Rittig et al., 2023a; Rittig and Mitsos, 2024): The graphs are passed separately through graph encoders. The two resulting molecular fingerprints are then used to construct a mixture graph with nodes corresponding to molecules and edges to intermolecular interactions. The mixture graph is passed through another graph encoder that yields a mixture fingerprint that enters an MLP for the molar volume prediction. This third prediction task is denoted as follows:

$v_{\text{SingleX}}^{\text{mix}}(T, p, \mathbf{x})$: We perform a train-test split such that each binary mixture is only available at a single composition (but at varying temperatures and pressures) during training, whereas all data points corresponding to the remaining compositions of this binary mixtures are used for testing. Notably, this scenario represents a more realistic case compared to the $p_{\text{SingleT}}^{\text{sat}}(T)$, as density measurements for mixtures might be more sparse in the composition domain.

Table 1
Overview case studies.

No.	Abbreviation	s	$\hat{p}^{\text{mol}}(\text{s})$	Train set
1	$p_{\text{extra}}^{\text{sat}}(T)$	T	$p^{\text{sat}}(T)$	Lowest 80% temperature Points per molecule
2	$p_{\text{SingleT}}^{\text{sat}}(T)$	T	$p^{\text{sat}}(T)$	One temperature Point per molecule
3	$V_{\text{SingleX}}^{\text{mix}}(T, p, x)$	$[T, p, x]^T$	$V^{\text{mix}}(T, p, x)$	One composition point Per mixture, T and p

In all scenarios, we employ a GNN using concatenation of the state variables to the molecular fingerprint vectors as a baseline, denoted as $\text{GNN}_{\text{concat}}$. For the two vapor pressure prediction scenarios, we also consider as a baseline a GNN that embeds the Antoine equation, which is a hybrid model that predicts the parameters of the Antoine equation in the output head, similar to [Lansford et al. \(2023\)](#), [Santana et al. \(2024\)](#), [Hoffmann et al. \(2025\)](#), [Lin et al. \(2024\)](#). We refer to this as $\text{GNN}_{\text{Antoine}}$.

4.3. Implementation & hyperparameter

We implement DeepEOSNet and the GNN baselines in Python using PyTorch Geometric ([Fey and Lenssen, 2019](#)), based on our graph ML framework for molecular and mixture property prediction `\emph{GMolProp}`. All GNN model parts share the same basic architecture. The GNN baseline that embeds the Antoine equation includes adaptations following the approach by [Hoffmann et al. \(2025\)](#): A tanh activation function is applied in the last layer before entering the Antoine output head, and training follows a two step procedure, that is, first training with mean squared error (MSE) loss and then switching to the Huber loss ([Hoffmann et al., 2025](#)).

We optimize model hyperparameters in a two-step grid search approach. First, the hyperparameters concerning the model architecture are optimized for all three models. Then, with fixed optimized architectural hyperparameters, the training parameters are determined. The optimized hyperparameters are then used to generate the results for all three case studies. In the first step, the optimized parameters are the dimension of the molecular fingerprint $\in \{64, 128\}$, the number of MLP layers $\in \{1, 2, 3\}$, the activation function in the branch net $\in \{\text{LeakyReLU}, \text{ReLU}\}$ and in the trunk net $\in \{\text{tanh}, \text{sigmoid}\}$, and the latent feature dimension $N_{\text{hd}} \in \{20, 30, 40\}$. During the first step, the learning rate is fixed to 0.001, the batch size to 64, and the drop out rate to 0. For the DeepEOSNet, the hyperparameter optimization results in a fingerprint dimension of 128, 3 MLP layers in the branch net, LeakyReLU as activation in the GNN/branch net and tanh in the trunk net, and a latent feature dimension of 40. The GNN with state variable concatenations exhibits a fingerprint dimension of 128, 3 MLP layers, and the LeakyReLU activation function. For the $\text{GNN}_{\text{Antoine}}$, the fingerprint dimension is 128, the number of MLP layers is 3, and LeakyReLU is chosen as the activation function. In the second step, the initial learning rate $\in \{0.01, 0.001, 0.0001\}$, the batch size $\in \{32, 64, 128\}$ and the dropout rate $\in \{0.1, 0.05, 0\}$ are optimized. This yields the same training procedures for all three models, with a learning rate of 0.001, a batch size of 64, and a dropout rate of 0, which correspond to the hyperparameters chosen in the first step. Furthermore, we train all models for 300 epochs with early stopping, and repeat training 10 times with different seeds.

5. Results and discussion

We first present and discuss the results from the vapor pressure scenarios (Section 5.1) and then consider the more difficult mixture molar volume scenario (Section 5.2). For all numeric performance metrics and training curves reported in the subsequent sections, the averages and standard deviations across the 10 different seeds are reported. Parity plots and individual molecule examples are shown for the best run, in terms of performance on the test set, of each model across the 10 different seeds.

5.1. Vapor pressure prediction

Table 2 provides an overview of the prediction test performance of DeepEOSNet and the baseline GNN models, $\text{GNN}_{\text{concat}}$ and $\text{GNN}_{\text{Antoine}}$, for the two vapor pressure case studies. In general, both the $\text{GNN}_{\text{concat}}$ and the DeepEOSNet achieve high prediction accuracy across both vapor pressure prediction scenarios with R^2 values above 0.9. The $\text{GNN}_{\text{Antoine}}$ exhibits higher RMSE and MAE values, and an average R^2 below 0. Thus, in both case studies, both the $\text{GNN}_{\text{concat}}$ and the DeepEOSNet significantly outperform the $\text{GNN}_{\text{Antoine}}$.

Notably, the performance metrics reported here for the $\text{GNN}_{\text{Antoine}}$ are larger than those reported by [Hoffmann et al. \(2025\)](#), with an MAE of 0.41 compared to 0.2. However, Hoffmann et al. use a different train-test split, sorting unseen molecules to the test set, whereas we test for extrapolation capabilities in the temperature domain. As extrapolation to unseen molecules is an inherently different task than extrapolation to unseen temperature regions, the metrics cannot be directly compared.

Here, we generally find a high variance in prediction performance of the $\text{GNN}_{\text{Antoine}}$. In fact, despite the low average R^2 values (see **Table 2**), it achieves R^2 values above 0.9 for some of the ten runs, which is reflected in the high standard deviation and indicates training instabilities. This highlights that, in general, the Antoine equation is capable of achieving acceptable extrapolative performance in the temperature domain. However, embedding it into neural network architectures appears to significantly destabilize the training process, as the variance in the performance metrics is significantly higher in the $\text{GNN}_{\text{Antoine}}$ than in the other two models, which do not embed any semi-empirical equations and show a reasonably low standard deviation in the performance metrics. In **Fig. 4**, we show the training and validation loss for both vapor pressure scenarios, where the line represents the average and the shaded area the standard deviation between the 10 runs. We find that both the $\text{GNN}_{\text{concat}}$ and the DeepEOSNet have significantly faster convergence and, indeed, more stable training across different seeds than the $\text{GNN}_{\text{Antoine}}$, so the ease of training is an advantage of both the $\text{GNN}_{\text{concat}}$ and DeepEOSNet.

Notably, in the $p_{\text{extra}}^{\text{sat}}(T)$ scenario, the performance of the DeepEOSNet and the $\text{GNN}_{\text{concat}}$ is on par (see **Fig. 4(a)**). This highlights that given sufficiently many data points (on average 80 temperature points per molecule for this scenario), the $\text{GNN}_{\text{concat}}$ is also capable of disentangling the influence of the two dependencies using a single MLP to learn both molecule and state dependency. In the data scarce $p_{\text{SingleT}}^{\text{sat}}(T)$ scenario, an advantage of the DeepEOSNet becomes apparent. Here, the inductive bias that is introduced by the separate information processing on state and molecule dependency is leveraged to extract both dependencies in a data-efficient way.

To further compare the model capabilities, we analyze the two vapor pressure prediction scenarios individually in the following.


5.1.1. $p_{\text{extra}}^{\text{sat}}(T)$ scenario

In **Fig. 5**, we show parity plots for the $p_{\text{extra}}^{\text{sat}}(T)$ scenario, i.e., excluding all vapor pressure data points in a higher temperature range (cf. Section 4.2). We stress that the results in the parity plots corresponds to the best out of ten different training runs for each model (based on test error), allowing to compare the models for the case of stable training of the $\text{GNN}_{\text{Antoine}}$ model.

Table 2

Performance metrics, root mean squared error (RMSE), mean absolute error (MAE), and coefficient of determination (R^2) evaluated on the logarithmic scale, of different prediction models on test set for the two vapor pressure prediction scenarios.

Scenario		GNN _{concat}			DeepEOSNet			GNN _{Antoine}		
		RMSE	MAE	R^2	RMSE	MAE	R^2	RMSE	MAE	R^2
$p_{extra}^{sat}(T)$	avg	0.12	0.045	0.98	0.13	0.060	0.98	0.75	0.41	< 0
	σ	0.008	0.006	0.002	0.011	0.010	0.003	0.73	0.34	1.8
$p_{SingleT}^{sat}(T)$	avg	0.34	0.21	0.95	0.31	0.20	0.96	8.9	0.64	< 0
	σ	0.032	0.026	0.010	0.049	0.047	0.014	16	0.34	420

**Fig. 4.** Training curves.

Both the GNN_{concat} as well as the DeepEOSNet exhibit a tight distribution along the diagonal, indicated by the similar RMSE of 0.115 for the DeepEOSNet and 0.109 for the GNN_{concat}. For the GNN_{Antoine} the best performing run achieves comparative performance to the other two models, with an R^2 of 0.98, which is significantly above its average (see Table 2). Only in the high vapor pressure region, the distribution is slightly wider than in the other two models with the general tendency to overestimate (see Fig. 5(c)). In the high vapor pressure region, also the DeepEOSNet slightly deviates from the diagonal with a tendency to underestimate. The GNN_{concat} remains in its trend closer to the diagonal in the high vapor pressure region. Furthermore, for all three models, numerous outliers are visible; however, these do not indicate general patterns in mispredictions. As they seem to be present in equal amounts for all models, we rather hypothesize that these outliers can be partly explained by noise and erroneous points in the experimental data used for training. In particular, the vertical line that becomes apparent in all parity plots close to a real value of two can be explained by the fact that it corresponds to ambient pressure. As the data quantity is overall above-average at ambient pressure, it is also likely that the number of erroneous points is larger, which then results in the observed vertical line instead of more evenly distributed noise scatter. Overall, the results of the $p_{extra}^{sat}(T)$ case suggest that the DeepEOSNet as the more restrictive architecture poses neither a benefit nor a disadvantage for predictive performance. In fact, all models show the capability to predict vapor pressures at temperatures higher than those used in training with reasonable accuracy.

5.1.2. $p_{SingleT}^{sat}(T)$ scenario

Next, in Fig. 6, we consider the parity plots for the $p_{SingleT}^{sat}(T)$ case, where only a single, randomly selected temperature point per molecule is used for training (cf. Section 4.2); again showing the results for the best out of ten runs.

We see that the DeepEOSNet predictions are closely aligned along the diagonal, with a slight tendency to overestimate in the low vapor pressure region. The GNN_{concat} overestimates the vapor pressures in both the low and high pressure region. This reflects the difference in average RMSE between the DeepEOSNet with 0.31 and the GNN_{concat} with 0.34 (see Table 2). In general, both have a more widened distribution along the diagonal than in the $p_{extra}^{sat}(T)$ scenario, which can be explained by the fact that much less data is used for training, in fact, only a single temperature data point per molecule compared to 80% of available data. The best GNN_{Antoine} performs significantly worse in the $p_{SingleT}^{sat}(T)$ scenario, with a strong tendency to underestimate in the lower vapor pressure regions and overall numerous structural outliers. This is also reflected in the high average RMSE, which is an order of magnitude larger compared to the $p_{extra}^{sat}(T)$ scenario. We thus hypothesize that the limited amount of data available for training in the $p_{SingleT}^{sat}(T)$ scenario does not suffice to learn suitable Antoine parameters. To a certain extent this is expected as in traditional parameter estimation the Antoine parameters would not be identifiable with a single temperature point per molecule. Thus, we focus on the DeepEOSNet and GNN_{concat} in the following.

In Fig. 7, we depict the $p^{sat}(T)$ relation for six sample molecules for the $p_{SingleT}^{sat}(T)$ case, which are randomly sampled from the dataset,




Fig. 5. Parity plots of test set $p_{extra}^{sat}(T)$.

together with the predictions by the DeepEOSNet and GNN_{concat} . For (a) benzo(k)fluoranthene and (d) benzo(g,h,i)perylene, the GNN_{concat} predicts near-linear behavior in the logarithmic vapor pressure scale across the entire temperature domain; for (f) eicosanoic acid, the prediction is nearly linear. As such, the GNN_{concat} fails to capture the nonlinearity in the temperature dependency for these molecules, resulting in systematic over- and underestimations. The DeepEOSNet model correctly identifies nonlinearities across the whole vapor pressure domain, without any systematic deviations. In the other three molecule examples, namely (b) acetonitrile, (c) chloroform and (e) 1,2-dichloro-4-(trifluoromethyl)benzene both models capture the curvature of the experimental data well. For (b) acetonitrile, the DeepEOSNet exhibits deviations from the experimental values in the high temperature region, whereas the GNN_{concat} deviates in the low temperature region. For (c) chloroform, both models deviate in the low temperature region, with the GNN_{concat} also underestimating slightly at moderate temperatures. The widened distribution along the diagonal in the parity plot (see Fig. 6) is likely to be attributed to such deviations. Given that the $p_{SingleT}^{sat}(T)$ case has a roughly 99% smaller training set size than the $p_{extra}^{sat}(T)$ case, the prediction accuracy achieved by both models – in particular of the DeepEOSNet – is overall highly promising.

Notably, the individual vapor pressure curves (see Fig. 7) also highlight problems with experimental data quality. Particularly for chloroform (see Fig. 7(c)), the outlier at 300 K is potentially a conversion error, as an order of magnitude smaller would better align with the remainder of the dataset. This highlights that even curated experimental databases, e.g., as here provided by NIST, can include erroneous data points, which might explain at least some of the outliers observed in the parity plots (see Fig. 5). Such erroneous data points occur frequently in molecular property datasets, are difficult to verify and therefore represent a practical scenario. Additionally, we are rather interested in comparing different methods, including how they are

affected in such cases, hence we decided to keep the dataset as provided by NIST.

5.2. Mixture molar volume prediction

Lastly, we consider the $V_{SingleX}^{mix}(T,p,x)$ case, i.e., prediction of the molar volume of binary mixtures as a function temperature, pressure, and composition (cf. Section 4.2). Here, we compare the DeepEOSNet to the GNN_{concat} architecture, i.e., to simply concatenating the three state variables to the binary mixture fingerprint vector.

In Table 3, we show the prediction metrics for both models. Both the DeepEOSNet and the GNN_{concat} achieve good overall accuracies with average R^2 values at 0.99, and low RMSEs and MAEs in a similar range. The relatively low standard deviation obtained from ten different training runs indicates stable training, as in the vapor pressure case studies; notably, we also find similar convergence rates for both models.

We further show the parity plots for the $V_{SingleX}^{mix}(T,p,x)$ scenario in Fig. 8. For both models, a large fraction of the predictions matches the diagonal closely, emphasizing the high prediction accuracy. Given that the $V_{SingleX}^{mix}(T,p,x)$ is a more high-dimensional prediction task, which is inherently more challenging, both models show good overall prediction performance on a similar level. However, for both models, systematic misalignment for some groups of points is visible, almost in a horizontal alignment. Such prediction behavior can indicate that one or multiple state dependencies are not learned correctly but systematically over- or underestimated, which we investigate by means of individual molecules in the following.

Fig. 9 shows the molar volume over the composition of four individual binary mixtures, where (a) is selected randomly and the other three are selected to show different types of prediction difficulties. We choose the composition as independent variable, as we only include one composition point per binary mixture in the training (but at varying


Fig. 6. Parity plots of test set $p_{SingleT}^{sat}(T)$.

Table 3

Performance metrics, root mean squared error (RMSE), mean absolute error (MAE), and coefficient of determination (R^2) evaluated on the logarithmic scale, of different prediction models on test set for the mixture molar volume prediction scenario.

Scenario		GNN _{concat}			DeepEOSNet		
		RMSE	MAE	R^2	RMSE	MAE	R^2
$V_{SingleX}^{mix}(T,p,x)$	avg	0.013	0.0050	0.99	0.015	0.0064	0.99
	σ	0.0048	0.0036	0.0073	0.0076	0.0058	0.015

temperature and pressure) in the $V_{SingleX}^{mix}(T,p,x)$ case, hence learning the composition dependency is presumably most challenging.

For the mixture of (a) bicyclohexyl & methylcyclohexane, which behaves close to an ideal mixture with near-linear composition dependency, both the GNN_{concat} and the DeepEOSNet capture the trends well. For the mixture of (b) 1,2-dibromoethane & butyl propanoate both models fail to capture the composition dependency correctly, despite a rather simple near-linear relation. This is also reflected in the horizontal lines in the parity plots (see Fig. 8) as described above. This highlights that even for some mixtures with near-ideal behavior, a single data point in the composition domain is not sufficient for capturing the composition dependency accurately.

For the mixture of (c) cyclohexene & n-methyl-2-pyrrolidone, which exhibits highly non-ideal behavior, both models are not capable to approximate the composition dependency and its local maximum accurately. Hence for this mixture, a single data point in the composition domain is not sufficient for accurately predicting non-idealities for neither of the two models. We note that less than 2% of mixtures in the dataset exhibit this level of non-ideal behavior, which explains the tendency of both models to predict ideal mixture behavior. Contrary, for the mixture of (d) bromoform & 2-methoxyethanol example, the GNN_{concat} falsely predicts a local minimum, whereas the DeepEOSNet more accurately captures the near-linear curvature. Examples (b), (c)

and (d) explain the misalignment for specific points that we observe in the parity plots (Fig. 8). Specifically, the last two mixture examples stress that highly non-ideal behavior poses a significant challenge to both models in a data scarce scenario, and it remains unclear which model is more suitable for predicting this. As the accurate prediction of non-idealities has high implications on chemical engineering applications, these limitations should be considered when using the models in practice.

In general, the mixture examples in Fig. 9 highlight the added difficulty that the $V_{SingleX}^{mix}(T,p,x)$ scenario entails: The composition dependency of the molar volume differs for the individual mixtures across the dataset. This is structurally different to the first two case studies focusing on the vapor pressure, as the temperature dependency of the vapor pressure typically has a similar form for all molecules (Pfennig, 2013). The advantage of the DeepEOSNet observed in the $p_{SingleT}^{sat}(T)$ case seems to appear in scenarios where data is sparse in the state domain and the state dependency is structurally similar across the entire dataset. Overall, we therefore find the DeepEOSNet architecture to result in at least similar prediction accuracy compared to current approaches in capturing state dependencies, additionally providing benefits for properties for which the state dependency is similar between molecules. DeepEOSNet is thus a promising alternative for predicting the state dependencies of molecular and mixture properties.




Fig. 7. Vapor pressure curve for sample molecules $p_{SingleT}^{sat}(T)$.

6. Conclusion

We propose an architectural adaptation to GNNs for predicting the state-dependent thermodynamic properties of molecules. The proposed architecture, DeepEOSNet, takes inspiration from DeepONets and SplitNets, i.e., the prediction is separated into two network channels, capturing molecular structure and general state dependencies, respectively. We apply DeepEOSNet to three property prediction case studies and compare its performance to simply concatenating state information to molecular fingerprints and embedding semi-empirical thermodynamic equations. In the first case study, where vapor pressure is to be predicted and data is abundant along the temperature domain, both the DeepEOSNet and GNN_{concat} perform equally well, outperforming the GNN with embedded semi-empirical equations in terms of

training stability. In case of data sparsity along the temperature domain of vapor pressures, which is tested in the second case study, where only a single temperature point per molecule is used for training, the DeepEOSNet shows the highest prediction performance. This highlights the capabilities of the DeepEOSNet to learn functional relations, here state dependencies, in a data efficient way. The results from the third case study, that is, predicting the molar volume of binary mixtures as a function of temperature, pressure, and composition, show that the DeepEOSNet's capabilities can be limited to prediction tasks where the dependency on the additional input feature is structurally similar across the dataset.

In summary, DeepEOSNet is a promising architecture for capturing the state dependencies of thermodynamic properties when training data are sparse in the state domain and the state dependency of the property




Fig. 8. Parity plots of test set $V_{SingleX}^{mix}(T,p,x)$.




Fig. 9. Molar volume over mole fractions for four exemplary mixtures in the $V_{SingleX}^{mix}(T,p,x)$ case.

of interest is structurally similar across the dataset. In scenarios where training data are abundant, concatenation-based approaches for the combination of molecule and state-dependency are at least equally suitable. Correctly predicting the state dependency with varying functional form across the dataset remains a challenge for both DeepEOSNet and GNN_{concat}. Therefore, further research is needed to establish the most suitable approach in scenarios where the state dependency varies structurally across the dataset, and compare different approaches with respect to prediction uncertainty, cf. [Hirschfeld et al. \(2020\)](#), [Heid et al. \(2023\)](#), [Jiang et al. \(2024\)](#), [Gao et al. \(2025\)](#).

Future work can easily transfer DeepEOSNet to predicting further state-dependent molecular and mixture properties. Here, a promising extension is the combination of DeepEOSNet with thermodynamics-consistent architectures ([Rittig and Mitsos, 2024](#)).

CRedit authorship contribution statement

Jan Pavšek: Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Alexander Mitsos:** Writing – review & editing, Supervision, Funding

acquisition, Conceptualization. **Manuel Dahmen**: Writing – review & editing, Conceptualization. **Tai Xuan Tan**: Writing – review & editing, Methodology. **Jan G. Rittig**: Writing – original draft, Supervision, Software, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This project was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 466417970 – within the Priority Programme “SPP 2331: Machine Learning in Chemical Engineering”.

Alexander Mitsos and Manuel Dahmen received funding from the Helmholtz Association of German Research Centers.

Model training was performed with computing resources granted by RWTH Aachen University.

The authors thank René Görgen for his software engineering support.

Data availability

The data used for training of the presented models are confidential and can be accessed through the [NIST Thermodata engine](#). The code is available in our [GitLab](#) repository [GMolprop](#).

References

- Aouichaoui, A.R., Cogliati, A., Abildskov, J., Sin, G., 2023. S-GNN: State-dependent graph neural networks for functional molecular properties. In: *Computer Aided Chemical Engineering*. vol. 52, Elsevier, pp. 575–581.
- Brozos, C., Rittig, J.G., Akanny, E., Bhattacharya, S., Kohlmann, C., Mitsos, A., 2025. Predicting the temperature-dependent CMC of surfactant mixtures with graph neural networks. *Comput. Chem. Eng.* 198, 109085.
- Chaparro, G., Müller, E.A., 2023. Development of thermodynamically consistent machine-learning equations of state: Application to the mie fluid. *J. Chem. Phys.* 158 (18).
- Chaparro, G., Müller, E.A., 2024. Development of a Helmholtz free energy equation of state for fluid and solid phases via artificial neural networks. *Commun. Phys.* 7 (1), 406.
- Damay, J., Jirasek, F., Kloft, M., Bortz, M., Hasse, H., 2021. Predicting activity coefficients at infinite dilution for varying temperatures by matrix completion. *Ind. Eng. Chem. Res.* 60 (40), 14564–14578.
- Dong, X., De Salvo, B., Li, M., Liu, C., Qu, Z., Kung, H.-T., Li, Z., 2022. Splitnets: Designing neural architectures for efficient distributed computing on head-mounted systems. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12559–12569.
- Felton, K.C., Raßpe-Lange, L., Rittig, J.G., Leonhard, K., Mitsos, A., Meyer-Kirschner, J., Knösche, C., Lapkin, A.A., 2024. ML-SAFT: a machine learning framework for PCP-SAFT parameter prediction. *Chem. Eng. J.* 492, 151999.
- Fey, M., Lenssen, J.E., 2019. Fast graph representation learning with PyTorch Geometric. In: *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- Gao, Q., Miedema, D.C., Zhao, Y., Weber, J.M., Tao, Q., Schweidtmann, A.M., 2025. Bayesian uncertainty quantification of graph neural networks using stochastic gradient Hamiltonian Monte Carlo. *Syst. Control. Trans.* 1360–1364.
- Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E., 2017. Neural message passing for quantum chemistry. *Int. Conf. Mach. Learn.* 1263–1272, URL <https://proceedings.mlr.press/v70/gilmer17a>.
- Hammad, R., Mondal, S., 2025. Advancements in thermochemical predictions: a multi-output thermodynamics-informed neural network approach. *J. Cheminformatics* 17 (1), 95.
- Heid, E., McGill, C.J., Vermeire, F.H., Green, W.H., 2023. Characterizing uncertainty in machine learning for chemistry. *J. Chem. Inf. Model.* 63 (13), 4012–4029.
- Hirschfeld, L., Swanson, K., Yang, K., Barzilay, R., Coley, C.W., 2020. Uncertainty quantification using neural networks for molecular property prediction. *J. Chem. Inf. Model.* 60 (8), 3770–3780.
- Hoffmann, M., Hasse, H., Jirasek, F., 2025. GRAPPA—A hybrid graph neural network for predicting pure component vapor pressures. *Chem. Eng. J. Adv.* 100750.
- Jiang, S., Qin, S., Van Lehn, R.C., Balaprakash, P., Zavala, V.M., 2024. Uncertainty quantification for molecular property predictions with graph neural architecture search. *Digit. Discov.* 3 (8), 1534–1553.
- Lansford, J.L., Jensen, K.F., Barnes, B.C., 2023. Physics-informed transfer learning for out-of-sample vapor pressure predictions. *Propellants, Explos. Pyrotech.* 48 (4), e202200265.
- Lin, Y.-H., Liang, H.-H., Lin, S.-T., Li, Y.-P., 2024. Advancing vapor pressure prediction: A machine learning approach with directed message passing neural networks. *J. the Taiwan Inst. Chem. Eng.* 105926.
- Lu, L., Jin, P., Pang, G., Zhang, Z., Karniadakis, G.E., 2021. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nat. Mach. Intell.* 3 (3), 218–229.
- McBride, B.J., 2002. NASA Glenn coefficients for calculating thermodynamic properties of individual species. National Aeronautics and Space Administration, John H. Glenn Research Center.
- Medina, E.I.S., Linke, S., Stoll, M., Sundmacher, K., 2022. Graph neural networks for the prediction of infinite dilution activity coefficients. *Digit. Discov.* 1 (3), 216–225.
- Medina, E.I.S., Linke, S., Stoll, M., Sundmacher, K., 2023. Gibbs–Helmholtz graph neural network: capturing the temperature dependency of activity coefficients at infinite dilution. *Digit. Discov.* 2 (3), 781–798.
- Medina, E.I.S., Sundmacher, K., 2026. Graph neural networks embedded into margules model for vapor–liquid equilibria prediction. *Fluid Phase Equilib.* 599, 114520.
- National Institute of Standards and Technology (NIST), 2025. ThermoData engine, NIST standard reference database 103b, version 10.4.5. critically evaluated thermodynamic and transport property data. URL <https://www.nist.gov/mml/acmd/trc/thermodata-engine/srd-nist-tde-103b>. (Accessed 20 June 2025).
- Pfennig, A., 2013. *Thermodynamik der Gemische*. Springer-Verlag.
- Reiser, P., Neubert, M., Eberhard, A., Torresi, L., Zhou, C., Shao, C., Metni, H., van Hoesel, C., Schopmans, H., Sommer, T., Friederich, P., 2022. Graph neural networks for materials science and chemistry. *Commun. Mater.* 3 (1), 93. <http://dx.doi.org/10.1038/s43246-022-00315-6>.
- Rittig, J.G., Dahmen, M., Grohe, M., Schwaller, P., Mitsos, A., 2025. Molecular machine learning in chemical process design. *arXiv preprint arXiv:2508.20527*.
- Rittig, J.G., Felton, K.C., Lapkin, A.A., Mitsos, A., 2023a. Gibbs–duhem-informed neural networks for binary activity coefficient prediction. *Digit. Discov.* 2 (6), 1752–1767.
- Rittig, J.G., Gao, Q., Dahmen, M., Mitsos, A., Schweidtmann, A.M., 2023b. Graph neural networks for the prediction of molecular structure–property relationships. In: Zhang, D., Del Rio Chanona, E.A. (Eds.), *Machine Learning and Hybrid Modelling for Reaction Engineering*. Royal Society of Chemistry, pp. 159–181.
- Rittig, J.G., Hicham, K.B., Schweidtmann, A.M., Dahmen, M., Mitsos, A., 2023c. Graph neural networks for temperature-dependent activity coefficient prediction of solutes in ionic liquids. *Comput. Chem. Eng.* 171, 108153.
- Rittig, J.G., Mitsos, A., 2024. Thermodynamics-consistent graph neural networks. *Chem. Sci.* 15 (44), 18504–18512.
- Santana, V.V., Rebello, C.M., Queiroz, L.P., Ribeiro, A.M., Shardt, N., Nogueira, I.B., 2024. PUFFIN: A path-unifying feed-forward interfaced network for vapor pressure prediction. *Chem. Eng. Sci.* 286, 119623.
- Schweidtmann, A.M., Rittig, J.G., Konig, A., Grohe, M., Mitsos, A., Dahmen, M., 2020. Graph neural networks for prediction of fuel ignition quality. *Energy & Fuels* 34 (9), 11395–11407.
- Specht, T., Nagda, M., Fellenz, S., Mandt, S., Hasse, H., Jirasek, F., 2024. HANNA: hard-constraint neural network for consistent activity coefficient prediction. *Chem. Sci.* 15 (47), 19777–19786.
- Sun, Y., Moya, C., Lin, G., Yue, M., 2023. Deepgraphonnet: A deep graph operator network to learn and zero-shot transfer the dynamic response of networked systems. *IEEE Syst. J.* 17 (3), 4360–4370.
- Sun, J., Xue, J., Yang, G., Li, J., Zhang, W., 2025. Vapor–liquid phase equilibrium prediction for mixtures of binary systems using graph neural networks. *AIChE J.* 71 (2), e18637.
- Vermeire, F.H., Green, W.H., 2021. Transfer learning for solvation free energies: From quantum chemistry to experiments. *Chem. Eng. J.* 418, 129307.
- Wilson, A.G., Izmailov, P., 2020. Bayesian deep learning and a probabilistic perspective of generalization. *Adv. Neural Inf. Process. Syst.* 33, 4697–4708.
- Winter, B., Rehner, P., Esper, T., Schilling, J., Bardow, A., 2025. Understanding the language of molecules: predicting pure component parameters for the PC-SAFT equation of state from SMILES. *Digit. Discov.* 4 (5), 1142–1157.
- Winter, B., Winter, C., Esper, T., Schilling, J., Bardow, A., 2023. SPT-NRTL: A physics-guided machine learning model to predict thermodynamically consistent activity coefficients. *Fluid Phase Equilib.* 568, 113731.
- Winter, B., Winter, C., Schilling, J., Bardow, A., 2022. A smile is all you need: predicting limiting activity coefficients from SMILES with natural language processing. *Digit. Discov.* 1 (6), 859–869.