



PTFlash : A deep learning framework for isothermal two-phase equilibrium calculations

Jingang Qu, Thibault Faney, Jean-Charles de Hemptinne, Soleiman Yousef,
Patrick Gallinari

► To cite this version:

Jingang Qu, Thibault Faney, Jean-Charles de Hemptinne, Soleiman Yousef, Patrick Gallinari. PT-Flash : A deep learning framework for isothermal two-phase equilibrium calculations. 2022. hal-03659647v1

HAL Id: hal-03659647

<https://hal.science/hal-03659647v1>

Preprint submitted on 5 May 2022 (v1), last revised 10 Sep 2022 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PTFlash : A deep learning framework for isothermal two-phase equilibrium calculations

Jingang Qu^{a,b}, Thibault Faney^b, Jean-Charles de Hemptinne^b, Soleiman Yousef^b, Patrick Gallinari^{a,c}

^a*Sorbonne Université, CNRS, ISIR, F-75005 Paris, France*

^b*IFPEN, France*

^c*Criteo AI Lab, Paris, France*

Abstract

Phase equilibrium calculations are an essential part of numerical simulations of multi-component multi-phase flow in porous media, accounting for the largest share of the computational time. In this work, we introduce a GPU-enabled, fast, and parallel framework, *PTFlash*, that vectorizes algorithms required for isothermal two-phase flash calculations using PyTorch, and can facilitate a wide range of downstream applications. In addition, to further accelerate *PTFlash*, we design two task-specific neural networks, one for predicting the stability of given mixtures and the other for providing estimates of the distribution coefficients, which are trained offline and help shorten computation time by sidestepping stability analysis and reducing the number of iterations to reach convergence.

The evaluation of *PTFlash* was conducted on three case studies involving hydrocarbons, CO_2 and N_2 , for which the phase equilibrium was tested over a large range of temperature, pressure and composition conditions, using the Soave-Redlich-Kwong (SRK) equation of state. We compare *PTFlash* with an in-house thermodynamic library, *Carnot*, written in C++ and performing flash calculations one by one on CPU. Results show speed-ups on large scale calculations up to two order of magnitudes, while maintaining perfect precision with the reference solution provided by *Carnot*.

Keywords: Isothermal flash calculations, two-phase equilibrium, vectorization, deep learning

1. Introduction

Numerical simulation of multi-component multi-phase flow in porous media is an essential tool for many subsurface applications, from reservoir simulation to long term CO₂ storage. A core element of the simulator for such applications is to determine the phase distribution of a given fluid mixture at equilibrium, also known as "flash" calculations. Starting with the seminal work of Michelsen [1, 2], researchers have developed robust and efficient algorithms for isothermal two-phase flash calculations. These algorithms have been implemented in the IFPEN thermodynamic C++ library *Carnot*.

Nonetheless, flash calculations still account for the majority of simulation time in a large range of subsurface applications [3, 4]. In most simulators, flash calculations are performed for each grid cell at each time step. Moreover, since modern simulators tend to require higher and higher grid resolutions up to billions of grid cells [5], the share of computing time due to flash calculations is expected to increase as well. In this context, speeding up flash calculations has drawn increasing research interest.

Some efforts have been made to accelerate flash calculations. [6–8] proposed a reduction method aiming to reduce the number of independent variables by leveraging the sparsity of the binary interaction parameter matrix, resulting in a limited speed-up [4]. [9] introduced the shadow region method using the results of previous time steps to initiate the current one, which assumes that the changes in pressure, temperature, and composition of a given block are small between two adjacent time steps in typical compositional reservoir simulation. [10] presented tie-line based methods, which approximate the results of flash calculations through linear interpolation between existing tie-lines and can be seen as a kind of look-up table. In [11–14], the authors focused on the use of machine learning, which provides a collection of techniques that can effectively discover patterns and regularities in data. They used support vector machine [15], relevance vector machine [16] and neural networks [17] to directly predict equilibrium phases and provide more accurate initial estimates for flash calculations. In [5, 18], researchers focused on developing faster parallel linear solvers, with [5] mentioning specifically that the vectorization of partial equation of state (EOS) related operations would lead to faster execution.

In this work, we introduce *PTFlash*, a framework for isothermal two-phase equilibrium calculations based on the SRK equations of state [19]. *PTFlash* is built on the deep learning framework PyTorch [20] and consists in two

main elements, namely the vectorization of algorithms and the use of neural networks. First, we perform a complete rewrite of isothermal two-phase flash calculation algorithms of *Carnot* using PyTorch. This enables the systematic vectorization of the complex iterative algorithms implemented in *Carnot*, allowing in turn to efficiently harness modern hardware with the help of, e.g., Advanced Vector Extensions AVX for Intel CPUs [21] and CUDA for Nvidia GPUs [22]. Note that vectorization of complex iterative algorithms with branching is not straightforward and needs specific care. Second, we replace repetitive and time consuming parts of the original algorithms with deep neural networks trained on the exact solution. More specifically, one neural network is used to predict the stability of given mixtures, and another one is used to provide initial estimates for the iterative algorithms. Once well trained, neural networks are seamlessly incorporated into *PTFlash*. These two elements allow *PTFlash* to provide substantial speed-ups compared to *Carnot*, especially so in the context of flow simulations where parallel executions of flash calculations for up to a billion grid cells are needed.

The rest of this article is organized as follows. In Section 2, we introduce the fundamentals of isothermal two-phase flash calculations and present three case studies. In Section 3, we explain how to efficiently vectorize flash calculations using PyTorch. In Section 4, we present two neural networks to speed up calculations. In Section 5, we demonstrate the attractive speed-up due to vectorization and the introduction of neural networks. Finally, we summarize our work and suggest future research in Section 6.

2. Isothermal two-phase flash calculations

In this section, we introduce the essential concepts of isothermal two-phase flash calculations. In the following, without loss of generality, we consider the equilibrium between the liquid and vapor phases.

2.1. Problem setting

We consider a mixture of N_c components. Given pressure (P), temperature (T) and feed composition ($\mathbf{z} = (z_1, \dots, z_{N_c})$), the objective of flash calculations is to determine the system state at equilibrium: single phase or coexistence of two phases. In the latter case, we need to additionally compute the molar fraction of vapor phase θ_V , the composition of the liquid phase \mathbf{x} and that of the vapor phase \mathbf{y} . These properties are constrained by

the following mass balance equations:

$$x_i(1 - \theta_V) + y_i\theta_V = z_i, \quad \text{for } i = 1, \dots, N_c \quad (1a)$$

$$\sum_{i=1}^{N_c} x_i = \sum_{i=1}^{N_c} y_i = 1 \quad (1b)$$

In addition, the following equilibrium condition should be satisfied:

$$\frac{\varphi_i^L(P, T, \mathbf{x})}{\varphi_i^V(P, T, \mathbf{y})} = \frac{y_i}{x_i} \quad (2)$$

where the superscripts L and V refer to the liquid and vapor phases, respectively, and φ_i is the fugacity coefficient of component i , which is a known nonlinear function of P , T and the corresponding phase composition. This function depends on an equation of state that relates pressure, temperature and volume. In this work, we use the SRK equation of state [19] and solve it using an iterative approach rather than the analytical solution of the cubic equation, e.g., the Cardano's formula, which may be subject to numerical errors in certain edge cases [23]. For more details, see Appendix A.

2.2. Numerical solver

Equations 1 and 2 form a non-linear system, which is generally solved in a two-stage procedure. First, we establish the stability of a given mixture via stability analysis (Section 2.2.1). If the mixture is stable, only one phase exists at equilibrium. Otherwise, two phases coexist. Second, we determine θ_V , \mathbf{x} and \mathbf{y} at equilibrium through phase split calculations (Section 2.2.2).

2.2.1. Stability analysis

A mixture of composition \mathbf{z} is stable at specified P and T if and only if its total Gibbs energy is at the global minimum, which can be verified through the reduced tangent plane distance [1]:

$$tpd(\mathbf{w}) = \sum_{i=1}^{N_c} w_i (\ln w_i + \ln \varphi_i(\mathbf{w}) - \ln z_i - \ln \varphi_i(\mathbf{z})) \quad (3)$$

where \mathbf{w} is a trial phase composition. If $tpd(\mathbf{w})$ is non-negative for any \mathbf{w} , the mixture is stable. This involves a constrained minimization problem, which is generally reframed as an unconstrained one:

$$tm(\mathbf{W}) = \sum_{i=1}^{N_c} W_i (\ln W_i + \ln \varphi_i(\mathbf{W}) - \ln z_i - \ln \varphi_i(\mathbf{z}) - 1) \quad (4)$$

where tm is the modified tangent plane distance and \mathbf{W} is mole numbers. To locate the minima of tm , we first use the successive substitution method accelerated by the Dominant Eigenvalue Method (DEM) [24], which iterates:

$$\ln W_i^{(k+1)} = \ln z_i + \ln \varphi_i(\mathbf{z}) - \ln \varphi_i(\mathbf{W}^{(k)}) \quad (5)$$

It is customary to initiate the minimization with two sets of estimates, that is, vapour-like estimate $W_i = K_i z_i$ and liquid-like estimate $W_i = z_i / K_i$, where K_i is the distribution coefficients, defined as y_i / x_i and initialized via the Wilson approximation [19], as follows:

$$\ln K_i = \ln \left(\frac{P_{c,i}}{P} \right) + 5.373(1 + \omega_i) \left(1 - \frac{T_{c,i}}{T} \right) \quad (6)$$

where $T_{c,i}$ and $P_{c,i}$ refer to the critical temperature and pressure of component i , respectively, and ω_i is the acentric factor.

Once converging to a stationary point (i.e., $\max(|\partial tm / \partial \mathbf{W}|) < 1.0e-6$) or a negative tm is found, successive substitution stops. If this does not happen after a fixed number of iterations (9 in our work), especially in the vicinity of critical points, we resort to a second-order optimization technique, i.e., the trust-region method [25], to minimize $tm(\mathbf{W})$, which we describe in Appendix B.1. In addition, based on the results of stability analysis, we can re-estimate K_i more accurately as z_i / W_i^L if $tm^L < tm^V$ or W_i^V / z_i otherwise, where the superscripts V and L denote the results obtained using the vapor-like and liquid-like estimates, respectively.

2.2.2. Phase split calculations

Substituting $K_i = y_i / x_i$ into Equations 1 yields the following Rachford Rice equation [26]:

$$f_{RR}(\theta_V, \mathbf{K}) = \sum_{i=1}^{N_c} \frac{(K_i - 1)z_i}{1 + (K_i - 1)\theta_V} = 0 \quad (7)$$

Given $\mathbf{K} = (K_1, \dots, K_{N_c})$, we solve the above equation using the method proposed by [27] to get θ_V , which is detailed in Appendix C.

To obtain θ_V , \mathbf{x} and \mathbf{y} at equilibrium, phase split calculations start with the accelerated successive substitution method, as illustrated in Figure 1, and the corresponding convergence criterion is $\max(|K_i^{(k+1)} / K_i^{(k)} - 1|) < 1.0e-8$. If

successive substitution fails to converge after a few iterations (9 in our work), we use the trust-region method to minimize the reduced Gibbs energy:

$$G = \sum_{i=1}^{N_c} n_i^L (\ln x_i + \ln \varphi_i^L) + \sum_{i=1}^{N_c} n_i^V (\ln y_i + \ln \varphi_i^V) \quad (8)$$

where $n_i^L = x_i(1 - \theta_V)$ and $n_i^V = y_i\theta_V$ are the mole numbers of liquid and vapor phases, respectively. The convergence criterion is $\max(|\partial G / \partial n_i^V|) < 1.0e-8$. For more details, see Appendix B.2.

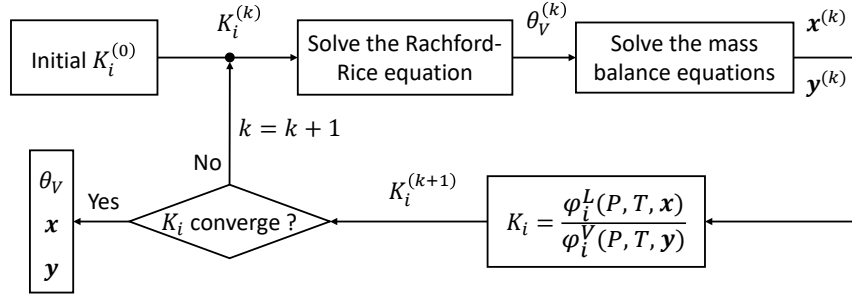


Figure 1: Successive substitution of phase split calculations

2.2.3. Strategy for isothermal two-phase flash calculations

We basically adopt the rules of thumb proposed by Michelsen in the book [28] to implement flash calculations, as shown in Figure 2. In the flowchart, we first initialize the distribution coefficients K_i using the Wilson approximation. Subsequently, in order to avoid computationally expensive stability analysis, we carry out the successive substitution of phase split calculations 3 times, which will end up with 3 possible cases: (1) θ_V is out of bounds (0, 1) during iterations. (2) None of ΔG , $tpd(\mathbf{x})$ and $tpd(\mathbf{y})$ are negative, where $tpd(\mathbf{x})$ and $tpd(\mathbf{y})$ are reduced tangent plane distances using current vapor and liquid phases as trial phases, and $\Delta G = \theta_V \times tpd(\mathbf{x}) + (1 - \theta_V) \times tpd(\mathbf{y})$. (3) Any of ΔG , $tpd(\mathbf{x})$ and $tpd(\mathbf{y})$ is negative.

For the first two cases, we cannot be sure of the stability of the given mixture, thus continuing with stability analysis. For the third case, we can conclude that the given mixture is unstable, thereby sidestepping stability analysis. Finally, if two phases coexist, we perform phase split calculations to get θ_V , \mathbf{x} and \mathbf{y} at equilibrium.

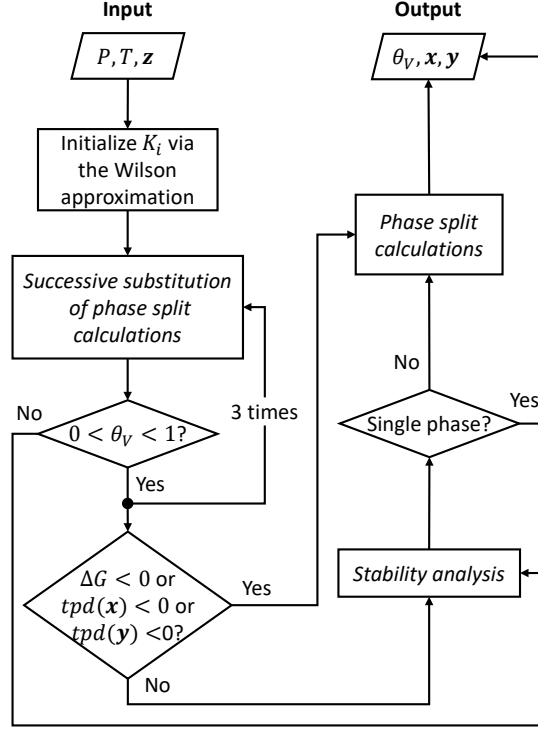


Figure 2: Flowchart of isothermal two-phase flash calculations

2.3. Case studies

Here, we introduce three case studies involving hydrocarbons, CO_2 and N_2 , whose properties are shown in Table 1. In this work, we only consider the binary interaction parameter (BIP) between CH_4 and CO_2 , which is 0.0882. The BIPs between the others are 0.

The first case study focuses on a system of two components (CH_4 and C_6H_{14}), and the second one involves four components (CH_4 , C_2H_6 , C_3H_8 and C_4H_{10}). For these two case studies, the ranges of pressure and temperature are 0.1MPa - 10MPa and 200K - 500K, respectively, and we consider the entire compositional space, i.e., $0 < z_i < 1$ for $i = 1, \dots, N_c$. The third case study includes all 9 components in Table 1. The bounds of pressure and temperature are 5MPa - 25MPa and 200K - 600K, respectively. In addition, from a practical perspective, given that some mixtures do not exist in nature, rather than considering the entire compositional space, we specify four different compositional ranges, as shown in Table 2, each of which represents

Table 1: Component properties

	P_c (MPa)	T_c (K)	w
CH_4	4.6	190.55	0.0111
C_2H_6	4.875	305.43	0.097
C_3H_8	4.268	369.82	0.1536
$n-C_4H_{10}$	3.796	425.16	0.2008
$n-C_5H_{12}$	3.3332	467.15	0.2635
C_6H_{14}	2.9688	507.4	0.296
$C_7H_{16}^+$	2.622	604.5	0.3565
CO_2	7.382	304.19	0.225
N_2	3.3944	126.25	0.039

one of the common reservoir fluid types, namely wet gas, gas condensate, volatile oil, and black oil. Figure 3(a) shows phase diagrams of four typical reservoir fluids at fixed compositions (Appendix D), and we can see that the more heavy hydrocarbons there are, the lower the pressure range of the phase envelope and the less volatile the fluid is.

Table 2: Four fluid types characterized by different compositional ranges

	CH_4	C_2H_6	$C_7H_{16}^+$	Others
Wet gas	80% - 100%	2% - 5%	$\leq 1\%$	$\leq 5\%$
Gas condensate	60% - 80%	5% - 10%	5% - 10%	$\leq 5\%$
Volatile oil	50% - 70%	6% - 10%	10% - 30%	$\leq 5\%$
Black oil	20% - 40%	3% - 6%	45% - 65%	$\leq 5\%$

2.4. Data generation

To efficiently sample input data including P , T and \mathbf{z} , we first use Latin Hypercube Sampling (LHS) technique to take space-filling samples [29]. Subsequently, for P and T , we linearly transform the uniform distribution $\mathcal{U}(0, 1)$ to the expected ranges. For \mathbf{z} subject to $\sum z_i = 1$, we transform a set of $\mathcal{U}(0, 1)$ into the Dirichlet distribution $Dir(\boldsymbol{\alpha})$ whose support is a simplex, as




Figure 3: In Figure (a), the squares on the phase envelopes represent critical points. In Figure (b), z_1 , z_2 and z_7 are the molar fractions of CH_4 , C_2H_6 and $C_7H_{16}^+$, respectively.

follows:

$$x_i \stackrel{i.i.d.}{\sim} \mathcal{U}(0, 1) \text{ using LHS} \quad (9a)$$

$$y_i = \Gamma(\alpha_i, 1).ppf(x_i) \quad (9b)$$

$$z_i = \frac{y_i}{\sum_{i=1}^{N_c} y_i} \quad (9c)$$

where $\alpha = (\alpha_1, \dots, \alpha_{N_c})$ is the concentration parameters of the Dirichlet distribution and controls its mode, $\Gamma(\alpha_i, 1)$ is the Gamma distribution, ppf represents the percent-point function, and $\mathbf{z} = (z_1, \dots, z_{N_c}) \sim Dir(\alpha)$.

For the first two case studies, the concentration parameters are $\alpha = \mathbf{1}$, i.e., all-ones vector. For the third case study, we adjust α for different fluid types to make the probability of each compositional range as large as possible, as shown in Table 3. Figure 3(b) presents the marginal distribution of z_i for black oil. In summary, we sample \mathbf{z} using Equations 9 with different α specified in Table 3, and then we single out the acceptable samples located in the compositional ranges defined in Table 2. In the following, unless otherwise specified, four fluid types are always equally represented.

Eventually, the samples of P , T and \mathbf{z} are concatenated together to form the complete input data.

Table 3: Concentration parameters α for different fluid types in Table 2

	α_1 for CH_4	α_2 for C_2H_6	α_7 for $C_7H_{16}^+$	α_i for others
Wet gas	100	5	1	1
Gas condensate	40	5	5	1
Volatile oil	55	8	20	1
Black oil	25	4	40	1

3. Vectorization of isothermal two-phase flash calculations

We vectorize the isothermal two-phase flash so that it takes as inputs $\mathbf{P} = (P_1, \dots, P_n)$, $\mathbf{T} = (T_1, \dots, T_n)$ and $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_n)$, where \mathbf{P} and \mathbf{T} are vectors, \mathbf{z} is a matrix, and n denotes the number of samples processed concurrently and is often referred to as the batch dimension.

In recent years, Automatic Vectorization (AV) has emerged and developed¹, e.g., JAX [30], which can automatically vectorize a function through the batching transformation that adds a batch dimension to its input. In this way, the vectorized function can process a batch of inputs simultaneously rather than processing them one by one in a loop. However, AV comes at the expense of performance to some extent and is slower than well-designed manual vectorization, which vectorizes a function by carefully revamping its internal operations to accommodate to a batch of inputs. For example, matrix-vector products for a batch of vectors can be directly replaced with a matrix-matrix product. In addition, flash calculations have the iterative nature and complicated control flow, which is likely to result in the failure of AV. Consequently, for finer-grained control, more flexibility, and better performance, we manually vectorize all algorithms involved in flash calculations, including the solution of the SRK equation of state and the Rachford-Rice equation, stability analysis and phase split calculations.

To achieve efficient vectorization, one difficulty is asynchronous convergence, that is, for each algorithm, the number of iterations required to reach convergence generally varies for different samples, which hinders vectorization and parallelism. To alleviate this problem, we design a general-purpose paradigm, *synchronizer*, to save converged results in time at the end of each

¹At the time of writing, PyTorch team released a fledgling library, *functorch*, which takes inspiration from JAX and supports Automatic Vectorization.

iteration and then remove the corresponding samples in order not to waste computational resources on them in the following iterations, which is achieved by leveraging a one-dimensional Boolean mask encapsulating convergence information to efficiently access data in vectors and matrices, as follows:

$$\mathbf{X}^{(k+1)} \leftarrow f(\mathbf{X}^{(k)}) \quad (10a)$$

$$\text{Save } \mathbf{X}^{(k+1)}[\text{mask}] \text{ to } \widetilde{\mathbf{X}} \quad (10b)$$

$$\mathbf{X}^{(k+1)} \leftarrow \mathbf{X}^{(k+1)}[\sim \text{mask}] \quad (10c)$$

$$k \leftarrow k + 1 \quad (10d)$$

where k is the number of iterations, $f(\mathbf{X})$ is a vectorized iterated function taking as input $\mathbf{X} \in \mathbb{R}^{n \times m}$ (n is the batch dimension, i.e., number of samples, and m is the dimension of X), $\widetilde{\mathbf{X}}$ is a placeholder matrix used to save converged results, mask is a Boolean vector where True means convergence, and \sim denotes the logical NOT operator. The number of unconverged samples gradually decreases as a result of incremental convergence. For the full version of *synchronizer*, refer to Appendix E.1. We can use *synchronizer* to wrap and vectorize any iterative algorithm. For instance, we illustrate how to perform vectorized stability analysis in Appendix E.2.

The efficiency of *synchronizer* may be questioned because previously converged samples are still waiting for unconverged ones before moving to the next step. This is true, but the situation is not as pessimistic since we try to shorten the waiting time as much as possible. For example, if successive substitution fails to converge quickly, we immediately use the trust-region method. In any case, the delay caused by waiting is insignificant compared to the acceleration due to vectorization. Furthermore, we leverage neural networks to provide more accurate initial estimate $\mathbf{X}^{(0)}$ so that all samples converge as simultaneously as possible, thereby reducing asynchrony, which we will present in Section 4.

Once all algorithms are well vectorized, another problem is how to globally coordinate different subroutines. To this end, we add barrier synchronization to the entry points of stability analysis and phase split calculations in Figure 2, which can avoid any subroutine connected to it proceeding further until all others terminate and arrive at this barrier.

We also optimized the code using TorchScript [20], allowing for more efficient execution through algebraic peephole optimizations and fusion of some operations, and more practical asynchronous parallelism without the Python

global interpreter lock [31], whereby vapor-like and liquid-like estimates are dealt with in parallel in stability analysis.

4. Acceleration of flash calculations using neural networks

To further accelerate flash calculations, we create and train two task-specific neural networks, classifier and initializer. Classifier is used to predict the probability p that a given mixture is stable, i.e., $p = \text{classifier}(P, T, \mathbf{z})$, which involves a binary classification problem. It can predict the stability of most samples, thereby bypassing stability analysis and saving time. Initializer is able to initialize K_i more accurately than the Wilson approximation, i.e., $\ln K_i = \text{initializer}(P, T, \mathbf{z})$, which relates to a regression problem. It can reduce the number of iterations required to reach convergence and alleviate the asynchronous convergence we introduced before. Note that the hyper-parameters of neural networks presented below, e.g., the number of units and layers, are dedicated to the case study containing 9 components. Nonetheless, the basic architecture of neural networks and the training methods can be generalized to any case.

4.1. Classifier

4.1.1. Architecture

Classifier has 3 hidden layers with 32 neurons and using the SiLU activation function [32–34]. The output layer has only one neuron and uses the sigmoid activation function compressing a real number to the range (0, 1). The input \mathbf{x} consists of P , T and \mathbf{z} , and the output is the probability p that a given mixture is stable. The scaling layer standardizes the inputs as $(\mathbf{x} - \mathbf{u})/\mathbf{s}$, where \mathbf{u} and \mathbf{s} are the mean and standard deviation of \mathbf{x} over the training set. To train classifier, we use the binary cross-entropy (bce), which is the de-facto loss function for binary classification problems and defined as:

$$\text{bce}(y, p) = y \ln p + (1 - y) \ln(1 - p) \quad (11)$$

where y is either 0 for unstable mixtures or 1 for stable ones.

The architecture of classifier is obtained by tuning hyper-parameters using Tree-Structured Parzen Estimator optimization algorithm [35] with Asynchronous Successive Halving algorithm [36] as an auxiliary tool to early stop less promising trials. We create a dataset containing 100 thousand samples (80% for training and 20% for validation), and then tune hyper-parameters

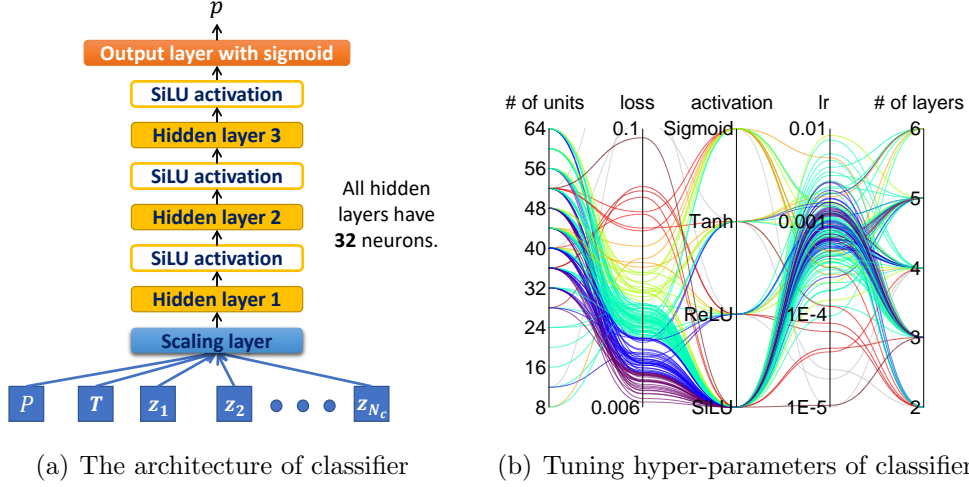


Figure 4: Figure (a) shows the architecture of classifier. Figure (b) is a parallel coordinates plot used to visualize the results of tuning hyper-parameters of classifier, where lr stands for learning rate. The colors of lines are mapped to the value of the loss.

of classifier with 150 trials to minimize the loss on the validation set (We use Adam [37] as optimizer and the batch size is 512), as shown in Figure 4(b). We can see that SiLU largely outperforms other activation functions.

4.1.2. Training

We first generate one million samples in the way described in Section 2.3, and then feed them to *PTFlash* to determine stability (no need for phase split calculations), which takes about 2 seconds. Subsequently, these samples are divided into the training (70%), validation (15%) and test (15%) sets. To train classifier, we set the batch size to 512 and use Adam with Triangular Cyclic Learning Rate (CLR) [38, 39], which periodically increases and decreases the learning rate during training, as shown in Figure 5(a). CLR helps neural networks achieve superb performance using fewer epochs and less time [40]. Early stopping is also used to avoid overfitting [41]. The total training time is about 5 minutes using Nvidia RTX 3080. The final performance of classifier on the test set is $bce = 0.002$ and $accuracy = 99.93\%$. For a more intuitive understanding of performance, Figure 5(b) shows the contours of probabilities predicted by classifier, where the blue contour of $p = 0.5$ basically coincides with the phase envelope. In the zoomed inset, the additional green and yellow contours correspond to $p=0.02$ and 0.98 ,

respectively, which are quite close to that of $p = 0.5$. This means that the prediction of classifier is very accurate.

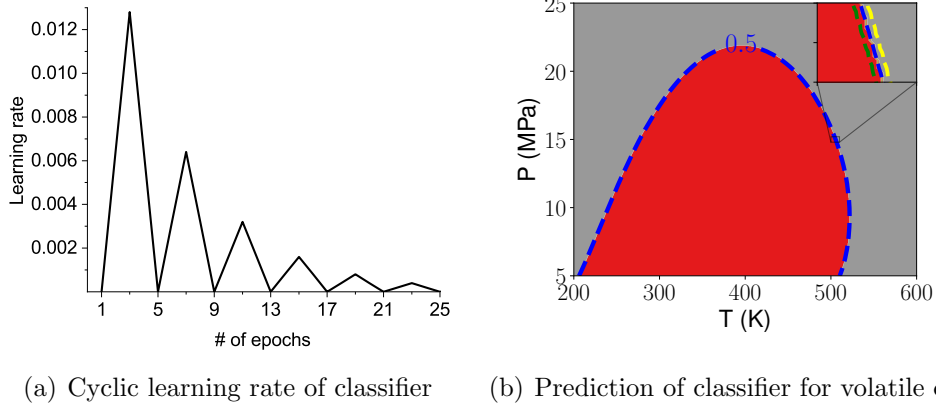


Figure 5: Figure (a) shows how the learning rate varies cyclically. Figure (b) illustrates the contours of probabilities predicted by classifier for volatile oil at fixed composition. The red and gray correspond to the two-phase and monophasic regions, respectively.

4.2. *Initializer*

4.2.1. *Architecture*

Initializer has 1 hidden layer and 3 residual blocks, as shown in Figure 6. Each residual block has 2 hidden layers and a shortcut connection adding the input of the first hidden layer to the output of the second [42]. All hidden layers have 64 neurons and use the SiLU activation function. The output layer has N_c neurons without activation function. The wide shortcut, proposed in [43], enables neural networks to directly learn simple rules via it besides deep patterns through hidden layers, which is motivated by the fact that the inputs, such as P and T , are directly involved in the calculation of K_i . The concat layer concatenates the input layer and the outputs of the last residual block (The concatenation means putting two matrices $A \in \mathbb{R}^{d_1 \times d_2}$ and $B \in \mathbb{R}^{d_1 \times d_3}$ together to form a new one $C \in \mathbb{R}^{d_1 \times (d_2 + d_3)}$). In addition, the targets of initializer are $\ln K_i$ instead of K_i , since K_i varies in different orders of magnitude, which hampers the training of initializer, whereas $\ln K_i$ does not.

We found that the convergence of phase split calculations is robuster if K_i predicted by initializer can lead to more accurate θ_V , especially around

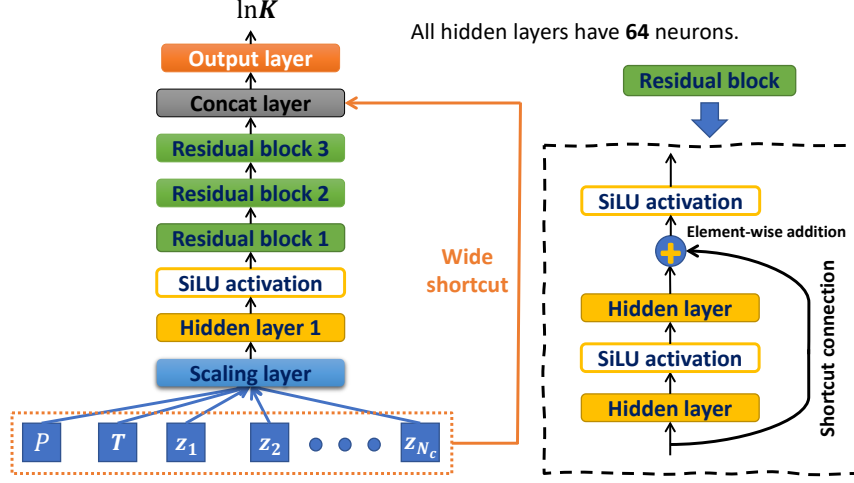


Figure 6: The architecture of initializer

critical points where calculations are quite sensitive to initial K_i and prone to degenerate into trivial solutions. As a consequence, the loss function used to train initializer consists of two parts, one is the mean absolute error (mae) in terms of K_i and the other is mae in terms of θ_V , as follows:

$$\text{mae}(\ln \mathbf{K}, \ln \hat{\mathbf{K}}) = \sum_{i=1}^{N_c} |\ln K_i - \ln \hat{K}_i| \quad (12a)$$

$$\text{mae}(\theta_V, \hat{\theta}_V) = |\theta_V - \hat{\theta}_V| \quad (12b)$$

where $\ln \mathbf{K}$ is the ground truth, $\ln \hat{\mathbf{K}}$ is the prediction of initializer, θ_V is the vapor fraction at equilibrium, and $\hat{\theta}_V$ is obtained by solving the Rice-Rachford equation given \mathbf{z} and the prediction $\hat{\mathbf{K}}$.

4.2.2. Training

We generate one million samples on the two-phase region (K_i is not available at the monophasic region), which are divided into the training (70%), validation (15%) and test (15%) sets. The training of initializer is carried out in two stages. First, we train it to minimize $\text{mae}(\ln \mathbf{K}, \ln \hat{\mathbf{K}})$, using Adam with CLR and setting the batch size to 512. Second, after the above training, we further train it to minimize $\text{mae}(\ln \mathbf{K}, \ln \hat{\mathbf{K}}) + \text{mae}(\theta_V, \hat{\theta}_V)$, using Adam with a small learning rate $1.0e-5$. Here, $\partial \hat{\theta}_V / \partial \hat{\mathbf{K}}$ is required during backpropagation and can be simply computed via PyTorch's automatic differentiation,

which, however, differentiates through all the unrolled iterations. Instead, we make use of the implicit function theorem [44], as follows:

$$\partial\hat{\theta}_V/\partial\hat{\mathbf{K}} = -[\partial_{\theta_V}f_{RR}(\hat{\theta}_V, \hat{\mathbf{K}})]^{-1}\partial_{\mathbf{K}}f_{RR}(\hat{\theta}_V, \hat{\mathbf{K}}) \quad (13)$$

In this way, we efficiently obtain $\partial\hat{\theta}_V/\partial\hat{\mathbf{K}}$ by only using the derivative information at the solution point of the Rachford-Rice equation $(\hat{\theta}_V, \hat{\mathbf{K}})$, avoiding differentiating through iterations.

Eventually, the performance of initializer on the test set is $\text{mae} = 9.66e-4$ in terms of $\ln K_i$ and $\text{mae} = 1.86e-3$ in terms of K_i .

4.3. Strategy for accelerating flash calculations using neural networks

As shown in Figure 7, given P , T and \mathbf{z} , we first use classifier to predict p . Next, based on two predefined thresholds, p_l and p_r , satisfying $p_l \leq p_r$, the given mixture is thought of as unstable if $p \leq p_l$ or stable if $p \geq p_r$. If $p_l < p < p_r$, we will use stability analysis to avoid unexpected errors. Here, we can adjust p_l and p_r to trade reliability for speed. In general, less errors occur with smaller p_l and greater p_r , but probably taking more time on stability analysis, and vice versa. A special case is $p_l = p_r = p_c$, where p_c could be a well-calibrated probability or simply set to 0.5, which means that we completely trust classifier (i.e., stable if $p \geq p_c$ or unstable otherwise), and no extra stability analysis is required. For initializer, it serves both stability analysis when $p_l < p < p_r$ and phase split calculations.

In addition, neural networks can also be used individually. If only classifier is available, initialize K_i via the Wilson approximation rather than initializer in Figure 7. If only initializer is available, use it to initialize K_i in Figure 2.

5. Results

In this section, we will compare our proposed framework for vectorized flash calculations, *PTFlash*, with *Carnot*, an in-house thermodynamic library developed by IFP Energies Nouvelles and based on C++, that performs isothermal two-phase flash calculations in the manner shown in Figure 2, but can only handle samples one at a time on CPU. Regarding the hardware, CPU is Intel 11700F, whose max turbo frequency can reach 4.9GHz, and GPU is NVIDIA RTX 3080 featuring 8704 CUDA cores and 10G memory. Note that since using multiple cores renders the frequency quite unstable due to heat




Figure 7: Acceleration of flash calculations using neural networks

accumulation, we only use one core of CPU so that the frequency can be stabilized at 4.5GHz, which allows for a consistent criterion for measuring the execution time.

PTFlash and *Carnot* gave identical results (coincidence to 9 decimal places under double-precision floating-point format) because they use exactly the same convergence criteria for all iterative algorithms. In the following, we will focus on comparing their speeds.

5.1. Vectorized flash calculations

We compare the execution time of different methods for flash calculations with respect to the workload quantified by the number of samples n , as shown in Figures 8. Due to GPU memory limitations, the maximum number of samples allowed is 10, 5, and 1 million for three case studies containing 2, 4 and 9 components, respectively. We can see that all three figures exhibit the same behavior. When the workload is relatively low, e.g., $n < 1000$, *Carnot* wins by large margins, and CPU is also preferable based on the fact that *PTFlash* runs much faster on CPU than on GPU. On the one hand, PyTorch has some fixed overhead in the setup of the working environment, e.g., the creation of tensors. On the other hand, when GPU is used, there

are some additional costs of CPU-GPU communication and synchronization. When n is small, these overheads dominate. As proof, we can see that the time of *PTFlash* on GPU hardly changes as n varies from 100 to 10^4 . On the contrary, the time of *Carnot* is almost proportional to n .

As the workload increases, the strength of *PTFlash* on GPU emerges and becomes increasingly prominent. For these three case studies, *PTFlash* on GPU is **163.4**, **106.3** and **50.5** times faster than *Carnot* at the maximum number of samples, respectively. This suggests that *PTFlash* on GPU is more suitable for large-scale computation. Interestingly, we can observe that *PTFlash* on CPU also outperforms *Carnot* when the workload is relatively heavy, e.g., $n > 10^3$. In fact, thanks to Advanced Vector Extensions, vectorization enables fuller utilization of CPU's computational power.

Next, we focus on the mixture of 9 components and analyze the performance of *PTFlash* for this case study. Table 4 is a performance profiler of *PTFlash* on GPU at $n = 10^6$, which records the running time of each subroutine of flash calculations. As a complement, Figures 9 dissect phase split calculations by tracking the total elapsed time and the convergence percentage up to each iteration, as well as the mean of critical distances d_c of converged samples at each iteration, where d_c is defined as:

$$d_c = \sqrt{\sum_{i=1}^{N_c} \ln K_i^2} \quad (14)$$

The closer to critical points, the smaller d_c . In other words, d_c indicates the closeness to critical points.

The observations of Figures 9 are summarized as follows: (1) In Figure 9(a), the slope of time with respect to the number of iterations is decreasing because the workload is reduced due to incremental convergence. (2) In Figure 9(b), for the samples that do not converge after successive substitution, the majority of them (92.67%) converge after 3 iterations of the trust-region method. (3) In Figure 9(c), d_c decreases during iterations, which means that samples close to critical points converge last and also confirms that convergence is slow around critical points.

The above analysis gives us a general understanding of *PTFlash*, but in fact it is not easy to analyze *PTFlash* comprehensively because each subroutine also contains iterative algorithms, such as solving the SRK equation of state and the Rice-Rachford equation. Nevertheless, given the information already obtained, we know that we need to shorten the time of stability

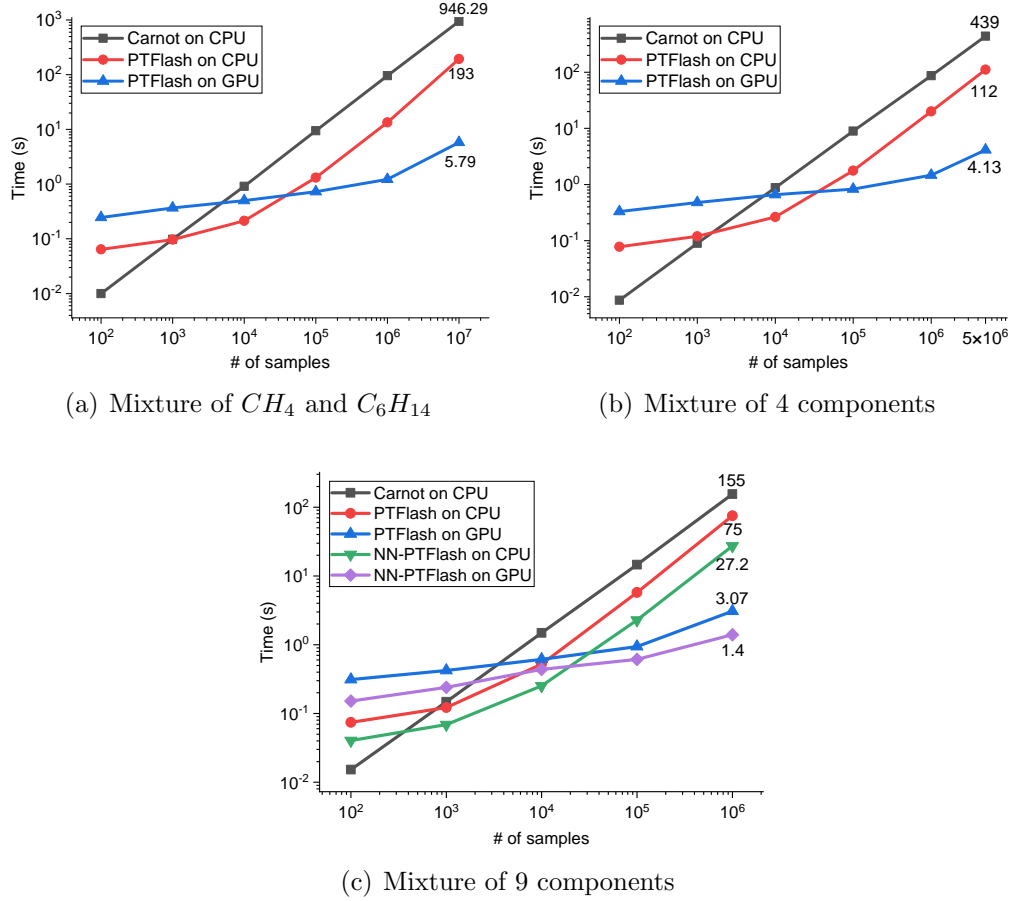


Figure 8: Comparison between *PTFlash* and *Carnot* in terms of speed. *NN-PTFlash* is *PTFlash* accelerated by neural networks, as presented in Section 4.

analysis and reduce the number of iterations in order to accelerate *PTFlash*, which is exactly the role of classifier and initializer.

5.2. Deep-learning-powered vectorized flash calculations

We trained neural networks following Section 4 for the mixture of 9 components. Here, we will explore the effect of neural networks. First of all, we set $p_l = 0.02$ and $p_r = 0.98$ as the thresholds of stability and instability, which are carefully chosen so that no misclassification occurs. In Figure 8(c), we can see that *NN-PTFlash* outpaces *PTFlash* on both CPU (2.7x speed-up) and GPU (2.2x speed-up). In addition, *NN-PTFlash* on GPU runs almost

Table 4: Performance profiler of *PTFlash* on GPU at $n = 10^6$ in Figure 8(c)

	ss of	Stability analysis				Phase split	
	phase split	vapor-like estimate		liquid-like estimate		calculations	
	calculations	ss	tr	ss	tr	ss	tr
# of samples	10 ⁶	625645	130715	625645	90179	413442	223741
Convergence	37.44% ¹	79.11%	100%	85.59%	100%	45.88%	100%
Max number of iterations	3	9	18	9	16	9	13
Total time	0.4565s	0.4136s	0.3417s	0.4044s	0.2706s	0.7412s	0.5132s
		1.3237s ²				1.2544s	
ss: successive substitution				tr: trust-region method			

¹ 37.44% is the percentage of samples for which any of ΔG , tpd_x and tpd_y is negative after 3 attempts of successive substitution, as described in Section 2.2.3.

² The total time of stability analysis is less than the sum of the times of all subroutines because vapor-like and liquid-like estimates are handled concurrently.

110.7 times faster than *Carnot* at $n = 10^6$.

Table 5 is the performance profiler of *NN-PTFlash* on GPU. We can see that classifier is able to precisely determine the stability of the vast majority of samples (99.42%), which significantly relieves the burden of stability analysis and saves time. In addition, compared to phase split calculations of *PTFlash*, the convergence percentage of successive substitution increases from 45.88% to 67.40%, and the overall time is also greatly reduced, which is contributed to better initial K_i provided by initializer.

We also performed ablation studies to compare the contributions of classifier and initializer by using them individually. For instance, when handling 1 million samples for the case study containing 9 components, *NN-PTFlash* with only classifier on GPU takes 1.88s. However, the attempt to use initializer alone fails because we find that initializer outputs considerable values (e.g., $1.0e15$) for stable mixtures far away from the boundary between the single-phase and two-phase regions, which leads to numerical overflow. From machine learning terminology, this is the out-of-distribution generalization problem, since initializer is trained on the two-phase region and may suffer from large predictive errors when dealing with the single-phase region. Nonetheless, there is no problem when initializer works in tandem with classifier because remaining samples located in the single-phase region are fairly

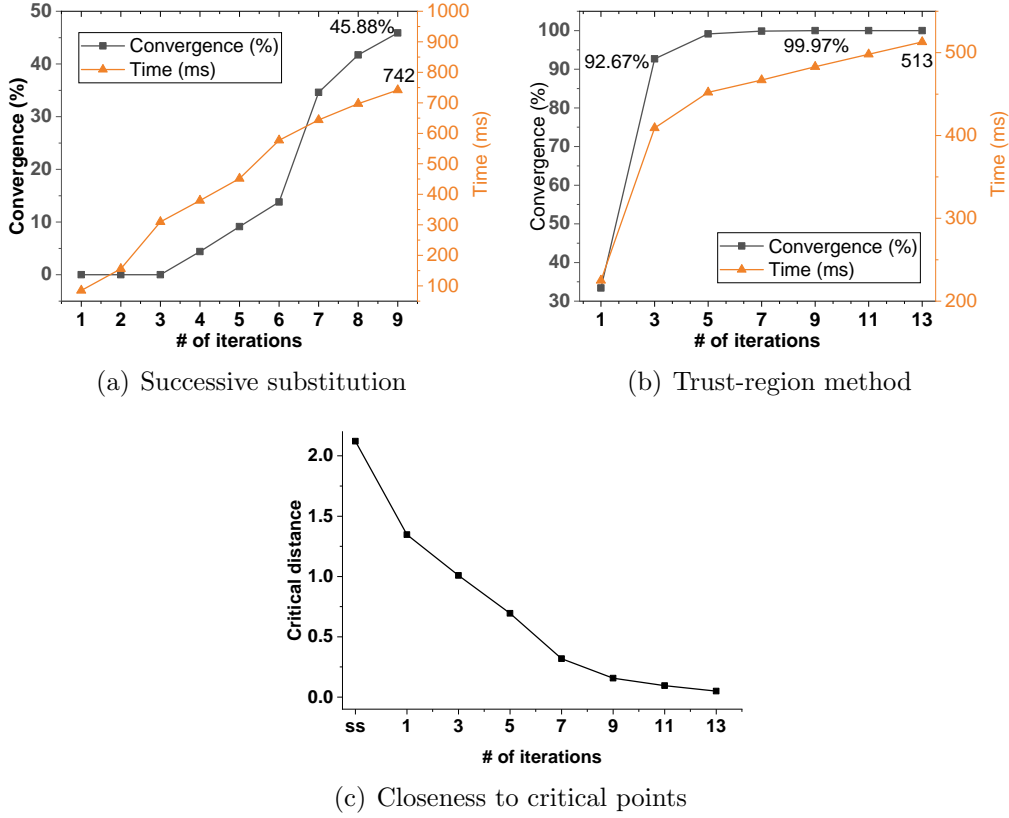


Figure 9: Dissection of phase split calculations of *PTFlash* on GPU. In Figure (c), on the x-axis, ss corresponds to the end of successive substitution and other integers are the number of iterations of the trust-region method.

close to the boundary after the filtering through classifier, as shown in Figure 5(b). In any case, based on the fact that *NN-PTFlash* using only classifier always lags behind that using both, we can conclude that both classifier and initializer play an important role in speeding up flash calculations.

5.3. Discussion

The results show that the systematic and exhaustive vectorization of isothermal two-phase flash calculations does result in attractive speed-up when large-scale computation is involved, e.g., the number of samples to process is on the order of millions. Importantly, this speed-up does not come at the cost of accuracy and stability like [11–14] which are subject to the

Table 5: Performance profiler of *NN-PTFlash* on GPU at $n = 10^6$ in Figure 8(c)

	classifier	Stability analysis				Phase split	
		vapor-like estimate		liquid-like estimate		calculations	
		ss	tr	ss	tr	ss	tr
# of samples	10 ⁶	5818	1073	5818	1704	413442	134786
Convergence	99.42% ¹	81.56%	100%	70.71%	100%	67.40%	100%
Max number of iterations		9	13	9	12	9	13
Total time	0.0005s	0.1365s	0.128s	0.0514s	0.12s	0.7043s	0.3388s
		0.34s				1.0431s	
ss: successive substitution				tr: trust-region method			

¹ 99.42% includes 58.38% predicted as stable (i.e., $p > p_r$) and 41.04% predicted as unstable (i.e., $p < p_l$).

unreliability of machine learning models. In addition, we can see that neural networks, such as classifier and initializer, really make a big difference.

Due to GPU memory limitations, the number of samples n is limited in Figures 8. Nonetheless, we can see that the slopes of time with respect to n differ significantly between different methods. The time of *Carnot* is proportional to n , in contrast, the time of *PTFlash* on GPU is increasing slowly. Therefore, it is reasonable to believe that the speed advantage of *PTFlash* on GPU will become increasingly prominent if n continues to grow.

Using PyTorch has several benefits in addition to its simplicity and flexibility. First, we can seamlessly incorporate neural networks into *PTFlash*. Second, any subroutine of *PTFlash* is fully differentiable through automatic differentiation, and we can also leverage the implicit function theorem for efficient differentiation, as we did in Section 4.2.2. Third, PyTorch’s highly optimized and ready-to-use multi-GPU parallelization largely circumvents the painstaking hand-crafted effort.

PTFlash also has several limitations. First, *PTFlash* is based on the SRK equation of state, which is relatively simple and sufficient for mixtures containing hydrocarbons and non-polar components, but does not take into account the effect of hydrogen bonding and falls short of adequacy for cross-associating mixtures having polar components, such as water and alcohol [45]. In this case, more advanced but also more complicated equations of state should be employed, such as the SAFT equation of state [46–51] or the

CPA equation of state [52, 53]. Second, *PTFlash* consumes a large amount of GPU memory, badly limiting its use on much larger batches of data. We need to optimize *PTFlash* to reduce the consumption of GPU memory, e.g., by leveraging the sparsity and symmetry of matrices. Third, *PTFlash* does not support multi-phase equilibrium. Last but not least, neural networks are subject to the out-of-distribution generalization problem. If pressure and temperature are out of predefined ranges used to train neural networks, predictive performance will deteriorate dramatically. Furthermore, once the components of the mixture change, we need to create new neural networks and train them from scratch.

6. Conclusion

In this work, we presented a fast and parallel framework, *PTFlash*, for isothermal two-phase flash calculations based on PyTorch and powered by GPU, which efficiently vectorizes algorithms and gains attractive speed-up at large-scale calculations. Two neural networks were used to predict the stability of given mixtures and to initialize the distribution coefficients more accurately than the Wilson approximation, which greatly accelerate *PTFlash*. In addition, *PTFlash* has much broader utility compared to the aforementioned methods which are mainly tailored to compositional reservoir simulation.

We compared *PTFlash* with *Carnot*, an in-house thermodynamic library, and we investigated three case studies containing 2, 4 and 9 components with maximum number of samples of 10, 5 and 1 million, respectively. The results showed that *PTFlash* on GPU is **163.4**, **106.3** and **50.5** times faster than *Carnot* at the maximum number of samples for these three cases, respectively.

In the future, we will optimize *PTFlash* to reduce the consumption of GPU memory and extend our work to support multi-phase equilibrium and more advanced equations of state. In addition, we will also apply our work to downstream applications, e.g., compositional reservoir simulation.

Appendix A. SRK equation of state and its solution

The SRK equation of state describes the relationship between pressure (P), temperature (T) and volume (V) in the following mathematical form:

$$P = \frac{RT}{V - b} - \frac{a\alpha}{V(V + b)} \quad (\text{A.1})$$

where R is the gas constant, $a\alpha$ refers to the temperature-dependent energy parameter, and b denotes the co-volume parameter. We employ the van der Waals mixing rules and the classical combining rules to calculate $a\alpha$ and b , as follows:

$$a\alpha = \sum_{i=1}^{N_c} \sum_{j=1}^{N_c} c_i c_j (a\alpha)_{ij} \quad (\text{A.2a})$$

$$(a\alpha)_{ij} = (1 - k_{ij}) \sqrt{(a\alpha)_i (a\alpha)_j} \quad (\text{A.2b})$$

$$b = \sum_{i=1}^{N_c} c_i b_i \quad (\text{A.2c})$$

$$a_i = \frac{0.42748 \cdot R^2 (T_{c,i})^2}{P_{c,i}} \quad (\text{A.2d})$$

$$b_i = \frac{0.08664 \cdot R T_{c,i}}{P_{c,i}} \quad (\text{A.2e})$$

$$\alpha_i = \left[1 + m_i \left(1 - \sqrt{\frac{T}{T_{c,i}}} \right) \right]^2 \quad (\text{A.2f})$$

$$m_i = 0.480 + 1.574 \omega_i - 0.176 \omega_i^2 \quad (\text{A.2g})$$

where the subscripts i and j refer to the components i and j , respectively, c_i denotes the mole fraction of the component i in the phase considered, k_{ij} is the binary interaction parameter between the components i and j , a_i and b_i are two substance-specific constants related to the critical temperature $T_{c,i}$ and critical pressure $P_{c,i}$, and ω_i is the acentric factor. We reformulate Equation A.1 as a cubic equation in terms of the compressibility factor Z :

$$f_{srk}(Z) = Z^3 - Z^2 + \rho_1 Z - \rho_0 = 0 \quad (\text{A.3})$$

where $\rho_0 = AB$ and $\rho_1 = A - B(1 + B)$, in which $A = a\alpha P/(R^2 T^2)$ and $B = bP/(RT)$. To find the roots of $f_{srk}(Z)$, we utilize an iterative approach

based on Halley’s method as follows:

$$Z^{(k+1)} = Z^{(k)} - \frac{f_{srk}(Z^{(k)})}{f'_{srk}(Z^{(k)})} \left[1 - \frac{f_{srk}(Z^{(k)})}{f'_{srk}(Z^{(k)})} \cdot \frac{f''_{srk}(Z^{(k)})}{2f'_{srk}(Z^{(k)})} \right]^{-1} \quad (\text{A.4})$$

The above iteration starts with a liquid-like guess and converges to a real root Z_0 (The convergence criterion is $|Z^{(k+1)}/Z^{(k)} - 1| < 1.0e-8$), and then we deflate the cubic equation as:

$$f_{srk}(Z) = (Z - Z_0)(Z^2 + pZ + q) = 0 \quad (\text{A.5})$$

where $p = Z_0 - 1$ and $q = pZ_0 + \rho_1$. If $p^2 < 4q$, only one real root Z_0 exists, otherwise, there are three real roots and the other two are $-p/2 \pm \sqrt{p^2 - 4q}/2$. In the latter case, we assign the smallest root to the liquid phase and the biggest one to the vapour phase. Subsequently, the root corresponding to the lowest Gibbs energy will be chosen. When Z is known, the fugacity coefficients φ_i are calculated as follows:

$$\begin{aligned} \ln \varphi_i(P, T, \mathbf{c}) = & \frac{b_i}{b}(Z - 1) - \ln(Z - B) \\ & + \frac{A}{B} \left(\frac{b_i}{b} - \frac{2}{a\alpha} \sum_{j=1}^{N_c} (a\alpha)_{ij} c_j \right) \ln\left(1 + \frac{B}{Z}\right) \end{aligned} \quad (\text{A.6})$$

where \mathbf{c} is the composition of the phase considered. In addition, the derivatives of the fugacity coefficients with respect to mole numbers, which are necessary for the trust-region methods of stability analysis and phase split calculations, are calculated explicitly rather than through PyTorch’s automatic differentiation, which requires retaining intermediate results and consumes prohibitive memory at large-scale computation.

Appendix B. Trust-region method

When successive substitution fails to converge quickly, particularly around critical points for which liquid and vapor phases are almost indistinguishable, we will switch to the trust-region method with restricted steps, which is a second-order optimization technique, to achieve faster convergence.

Appendix B.1. Trust-region method for stability analysis

The objective function to be minimized is the modified tangent plane distance:

$$tm(\mathbf{W}) = \sum_{i=1}^{N_c} W_i (\ln W_i + \ln \varphi_i(\mathbf{W}) - \ln z_i - \ln \varphi_i(\mathbf{z}) - 1)$$

The minimization is accomplished by iterating the following equations:

$$\boldsymbol{\beta}^{(k)} = 2\sqrt{\mathbf{W}^{(k)}} \quad (\text{B.1a})$$

$$(\mathbf{H}^{(k)} + \eta^{(k)}\mathbf{I}) \cdot \Delta\boldsymbol{\beta} + \mathbf{g}^{(k)} = \mathbf{0} \quad \text{s.t.} \quad \|\Delta\boldsymbol{\beta}\| \leq \Delta_{max}^{(k)} \quad (\text{B.1b})$$

$$\boldsymbol{\beta}^{(k+1)} = \boldsymbol{\beta}^{(k)} + \Delta\boldsymbol{\beta} \quad (\text{B.1c})$$

$$\mathbf{W}^{(k+1)} = \left(\frac{\boldsymbol{\beta}^{(k+1)}}{2} \right)^2 \quad (\text{B.1d})$$

where \mathbf{I} is the identity matrix, \mathbf{g} and \mathbf{H} are the gradient and Hessian matrix of tm with respect to $\boldsymbol{\beta}$, respectively, and are calculated as follows:

$$g_i = \sqrt{W_i} (\ln W_i + \ln \varphi_i(\mathbf{W}) - \ln z_i - \ln \varphi_i(\mathbf{z})) \quad (\text{B.2a})$$

$$H_{ij} = \sqrt{W_i W_j} \frac{\partial \ln \varphi_i}{\partial W_i} + \sigma_{ij} \left(1 + \frac{g_i}{\beta_i} \right) \quad \text{where } \sigma_{ij} = 1 \Leftrightarrow i = j \quad (\text{B.2b})$$

In addition, η is the trust-region size used to guarantee the positive definiteness of $\mathbf{H} + \eta\mathbf{I}$ and to tailor the step size to meet $\|\Delta\boldsymbol{\beta}\| \leq \Delta_{max}$, where Δ_{max} is adjusted during iterations depending on the match between the actual reduction $\delta_{tm} = tm^{(k+1)} - tm^{(k)}$ and the predicted reduction based on the quadratic approximation $\hat{\delta}_{tm} = \Delta\boldsymbol{\beta}^T \mathbf{g} + \frac{1}{2} \Delta\boldsymbol{\beta}^T \mathbf{H} \Delta\boldsymbol{\beta}$, using the following heuristic rules:

$$\Delta_{max}^{(k+1)} = \begin{cases} \frac{\Delta_{max}^{(k)}}{2}, & \text{if } \left| \delta_{tm} / \hat{\delta}_{tm} \right| \leq 0.25 \\ 2\Delta_{max}^{(k)}, & \text{if } \left| \delta_{tm} / \hat{\delta}_{tm} \right| \geq 0.75 \\ \Delta_{max}^{(k)}, & \text{otherwise} \end{cases} \quad (\text{B.3})$$

The convergence criterion of Equations B.1 is $\max(|\mathbf{g}|) < 1.0e-6$.

Appendix B.2. Trust-region method for phase split calculations

The objective function to be minimized is the reduced Gibbs energy:

$$G = \sum_{i=1}^{N_c} n_i^L (\ln x_i + \ln \varphi_i^L) + \sum_{i=1}^{N_c} n_i^V (\ln y_i + \ln \varphi_i^V)$$

where $n_i^L = x_i(1 - \theta_V)$ and $n_i^V = y_i\theta_V$ are the mole numbers of liquid and vapor phases, respectively. We choose n_i^V as the independent variable and perform the following iteration:

$$\left(\tilde{\mathbf{H}}^{(k)} + \tilde{\eta}^{(k)} \cdot \mathbf{D} \left(\frac{\mathbf{z}}{\mathbf{xy}} \right) \right) \cdot \Delta \mathbf{n}^V + \tilde{\mathbf{g}}^{(k)} = \mathbf{0} \quad s.t. \quad \|\Delta \mathbf{n}^V\| \leq \tilde{\Delta}_{max}^{(k)} \quad (\text{B.4a})$$

$$\mathbf{n}^{V,k+1} = \mathbf{n}^{V,k} + \Delta \mathbf{n}^V \quad (\text{B.4b})$$

where $\tilde{\mathbf{H}}^{(k)}$ and $\tilde{\mathbf{g}}^{(k)}$ are the gradient and hessian matrix of G with respect to n_i^V , respectively, and are calculated as follows:

$$\tilde{g}_i = \ln y_i + \ln \varphi_i^V - \ln x_i - \ln \varphi_i^L \quad (\text{B.5a})$$

$$\tilde{H}_{ij} = \frac{1}{\theta_V(1 - \theta_V)} \left(\frac{z_i}{x_i y_i} \sigma_{ij} - 1 + \theta_V \frac{\partial \ln \varphi_i^L}{\partial n_j^L} + (1 - \theta_V) \frac{\partial \ln \varphi_i^V}{\partial n_j^V} \right) \quad (\text{B.5b})$$

In addition, $\mathbf{D}(\cdot)$ is a diagonal matrix with diagonal entries in parentheses. The above iteration stops if $\max(|\tilde{\mathbf{g}}|) < 1.0e-8$. Here, the trust-region method is implemented in the same way as in stability analysis.

Appendix C. Solution of the Rachford-Rice equation

The Rachford-Rice equation is as follows:

$$f_{RR}(\theta_V, \mathbf{K}) = \sum_{i=1}^{N_c} \frac{(K_i - 1)z_i}{1 + (K_i - 1)\theta_V} = 0$$

Given \mathbf{K} , the solution of the above equation amounts to finding an appropriate zero yielding all non-negative phase compositions. Concretely, we adopt the method proposed by [27], which transforms f_{RR} into a helper function h_{RR} which is more linear in the vicinity of the zero:

$$h_{RR}(\theta_V, \mathbf{K}) = (\theta_V - \alpha_l) \cdot (\alpha_r - \theta_V) \cdot f_{RR}(\theta_V) = 0 \quad (\text{C.1})$$

where $\alpha_l = 1/(1 - \max(K_i))$ and $\alpha_r = 1/(1 - \min(K_i))$. The above equation is solved by alternating between the Newton method and the bisection method used when the Newton step renders θ_V out of the bounds which contain the zero and become narrower during iterations. When the Newton step size is smaller than $1.0e-8$, the iteration stops.

Appendix D. Some typical reservoir fluid compositions

Table D.6: Some typical reservoir fluid compositions

	Wet gas	Gas condensate	Volatile oil	Black oil
CH_4	92.46%	73.19%	57.6%	33.6%
C_2H_6	3.18%	7.8%	7.35%	4.01%
C_3H_8	1.01%	3.55%	4.21%	1.01%
$n-C_4H_{10}$	0.52%	2.16%	2.81%	1.15%
$n-C_5H_{12}$	0.21%	1.32%	1.48%	0.65%
C_6H_{14}	0.14%	1.09%	1.92%	1.8%
$C_7H_{16}^+$	0.82%	8.21%	22.57%	57.4%
CO_2	1.41%	2.37%	1.82%	0.07%
N_2	0.25%	0.31%	0.24%	0.31%

Appendix E. Vectorized algorithms

Appendix E.1. Synchronizer

Algorithm 1: PyTorch pseudo-code of *synchronizer* to save converged results after iteration and remove the corresponding samples

Input: Vectorized iterated function $f(\mathbf{X}, \mathbb{O})$, initial estimate $\mathbf{X}^{(0)}$, other f -related inputs \mathbb{O} , convergence criterion C , maximum number of iterations K

1 Initialization

2 Set the number of iterations $k \leftarrow 1$
 3 Generate a vector \mathbf{i} containing indices from 0 to $n - 1$
 /* n is the number of samples and indexing starts from 0. */
 4 Create a placeholder matrix $\widetilde{\mathbf{X}}$ of the same shape as $\mathbf{X}^{(0)}$

5 while $k \leq K$ do

6 $\mathbf{X}^{(k+1)} \leftarrow f(\mathbf{X}^{(k)}, \mathbb{O})$
 7 $\text{mask} \leftarrow C(\dots)$
 /* C returns a Boolean vector and True means convergence. */

8 Saving

9 indices $\leftarrow \mathbf{i}[\text{mask}]$
 10 $\widetilde{\mathbf{X}}[\text{indices}] \leftarrow \mathbf{X}^{(k+1)}[\text{mask}]$

11 Removing

12 $\mathbf{i} \leftarrow \mathbf{i}[\sim \text{mask}]$
 13 $\mathbb{O} \leftarrow \mathbb{O}[\sim \text{mask}]$
 /* Apply this operation to every element in \mathbb{O} */
 14 $\mathbf{X}^{(k+1)} \leftarrow \mathbf{X}^{(k+1)}[\sim \text{mask}]$

15 $k \leftarrow k + 1$

16 if $\text{len}(\mathbf{i}) \neq 0$ then

17 $\widetilde{\mathbf{X}}[\mathbf{i}] \leftarrow \mathbf{X}$
 /* Also save unconverged results for further utilization. */

Output: Converged results $\widetilde{\mathbf{X}}$ and unconverged indices \mathbf{i}

Appendix E.2. Vectorized stability analysis

Algorithm 2: PyTorch pseudo-code of vectorized stability analysis

Input: Pressure \mathbf{P} , temperature \mathbf{T} , feed composition \mathbf{z} , component properties $(\mathbf{P}_c, \mathbf{T}_c, \boldsymbol{\omega}, \text{BIPs})$, initial estimate $\mathbf{W}^{(0)}$, convergence criteria C_{ss} and C_{tr} , maximum numbers of iterations $K_{ss} = 9$ and $K_{tr} = 20$

- 1 **Initialization**
- 2 Instantiate $pteos = PTEOS(\mathbf{P}_c, \mathbf{T}_c, \boldsymbol{\omega}, \text{BIPs})$
 /* *PTEOS* is a PyTorch-based class to efficiently calculate the
 fugacity coefficients and their partial derivatives. */
- 3 **Successive substitution**
- 4 Iterated function f_{ss} specified by Equation 5
- 5 Other inputs $\mathbb{O}_{ss} \leftarrow \{\mathbf{P}, \mathbf{T}, \mathbf{z}\}$
- 6 $\mathbf{W}, \mathbf{i}_{ss} \leftarrow \text{synchronizer}(f_{ss}, \mathbf{W}^{(0)}, \mathbb{O}_{ss}, \mathbb{C}_{ss}, K_{ss})$
- 7 **Trust-region method**
- 8 Iterated function f_{tr} specified by Equations B.1
- 9 $\mathbf{W}_{tr}^{(0)} \leftarrow \mathbf{W}[\mathbf{i}_{ss}]$
- 10 Other inputs $\mathbb{O}_{tr} \leftarrow \{\mathbf{P}[\mathbf{i}_{ss}], \mathbf{T}[\mathbf{i}_{ss}], \mathbf{z}[\mathbf{i}_{ss}]\}$
- 11 $\mathbf{W}_{tr}, \mathbf{i}_{tr} \leftarrow \text{synchronizer}(f_{tr}, \mathbf{W}_{tr}^{(0)}, \mathbb{O}_{tr}, \mathbb{C}_{tr}, K_{tr})$
- 12 $\mathbf{W}[\mathbf{i}_{ss}] \leftarrow \mathbf{W}_{tr}$ and $\mathbf{i} \leftarrow \mathbf{i}_{ss}[\mathbf{i}_{tr}]$

Output: Converged results \mathbf{W} and unconverged indices \mathbf{i}

References

- [1] M. L. Michelsen, The isothermal flash problem. part II. phase-split calculation 9 (1) 21–40, publisher: Elsevier.
- [2] M. L. Michelsen, The isothermal flash problem. part i. stability 9 (1) 1–19, publisher: Elsevier.
- [3] P. Wang, E. H. Stenby, Non-iterative flash calculation algorithm in compositional reservoir simulation 95 93–108, publisher: Elsevier.
- [4] A. Belkadi, W. Yan, M. L. Michelsen, E. H. Stenby, Comparison of two methods for speeding up flash calculations in compositional simulations, in: SPE Reservoir Simulation Symposium, OnePetro.
- [5] A. H. Dogru, L. S. K. Fung, U. Middy, T. Al-Shaalan, J. A. Pita, A next-generation parallel reservoir simulator for giant reservoirs, in: SPE Reservoir Simulation Symposium, OnePetro.
- [6] M. L. Michelsen, Simplified flash calculations for cubic equations of state 25 (1) 184–188, publisher: ACS Publications.
- [7] E. M. Hendriks, Reduction theorem for phase equilibrium problems 27 (9) 1728–1732, publisher: ACS Publications.
- [8] E. M. Hendriks, A. Van Bergen, Application of a reduction method to phase equilibria calculations 74 17–34, publisher: Elsevier.
- [9] C. P. Rasmussen, K. Krejbjerg, M. L. Michelsen, K. E. Bjurstrøm, Increasing the computational speed of flash calculations with applications for compositional, transient simulations 9 (1) 32–38, publisher: OnePetro.
- [10] D. Voskov, H. A. Tchelepi, Compositional space parameterization for flow simulation, in: SPE Reservoir Simulation Symposium, OnePetro.
- [11] V. Gaganis, N. Varotsis, An integrated approach for rapid phase behavior calculations in compositional modeling 118 74–87, publisher: Elsevier.

- [12] V. Gaganis, N. Varotsis, Machine learning methods to speed up compositional reservoir simulation, in: SPE Europec/EAGE annual conference, OnePetro.
- [13] A. Kashinath, M. Szulczewski, A. Dogru, A fast algorithm for calculating isothermal phase behavior using machine learning 465 73–82. doi:10.1016/j.fluid.2018.02.004.
- [14] S. Wang, N. Sobecki, D. Ding, L. Zhu, Y.-S. Wu, Accelerating and stabilizing the vapor-liquid equilibrium (VLE) calculation in compositional simulation of unconventional reservoirs using deep learning based flash calculation 253 209–219. doi:10.1016/j.fuel.2019.05.023.
- [15] C. Cortes, V. Vapnik, Support-vector networks 20 (3) 273–297, publisher: Springer.
- [16] M. E. Tipping, Sparse bayesian learning and the relevance vector machine 1 211–244.
- [17] I. Goodfellow, Y. Bengio, A. Courville, Deep learning, MIT press.
- [18] Z. Chen, H. Liu, S. Yu, B. Hsieh, L. Shao, GPU-based parallel reservoir simulators, in: Domain Decomposition Methods in Science and Engineering XXI, Springer, pp. 199–206.
- [19] G. Soave, Equilibrium constants from a modified redlich-kwong equation of state 27 (6) 1197–1203, publisher: Elsevier.
- [20] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshein, L. Antiga, others, Pytorch: An imperative style, high-performance deep learning library 32 8026–8037.
- [21] C. Lomont, Introduction to intel advanced vector extensions 23.
- [22] J. Sanders, E. Kandrot, CUDA by example: an introduction to general-purpose GPU programming, Addison-Wesley Professional.
- [23] Y. Zhi, H. Lee, Fallibility of analytic roots of cubic equations of state in low temperature region 201 (2) 287–294, publisher: Elsevier.

- [24] O. Orbach, C. Crowe, Convergence promotion in the simulation of chemical processes with recycle-the dominant eigenvalue method 49 (4) 509–513, publisher: Wiley Online Library.
- [25] M. Hebden, An algorithm for minimization using exact second derivativesPublisher: Citeseer.
- [26] H. H. Rachford, J. Rice, Procedure for use of electronic digital computers in calculating flash vaporization hydrocarbon equilibrium 4 (10) 19–3, publisher: OnePetro.
- [27] C. Leibovici, J. Neoschil, A new look at the rachford-rice equation 74 303–308. doi:10.1016/0378-3812(92)85069-K.
- [28] M. L. Michelsen, J. Mollerup, Thermodynamic modelling: fundamentals and computational aspects, Tie-Line Publications.
- [29] M. D. McKay, R. J. Beckman, W. J. Conover, A comparison of three methods for selecting values of input variables in the analysis of output from a computer code 42 (1) 55–61, publisher: Taylor & Francis.
- [30] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, Q. Zhang, JAX: composable transformations of python+NumPy programs.
URL <http://github.com/google/jax>
- [31] G. Van Rossum, F. L. Drake, The python language reference manual, Network Theory Ltd.
- [32] D. Hendrycks, K. Gimpel, Gaussian error linear units (gelus).
- [33] S. Elfving, E. Uchibe, K. Doya, Sigmoid-weighted linear units for neural network function approximation in reinforcement learning 107 3–11, publisher: Elsevier.
- [34] P. Ramachandran, B. Zoph, Q. V. Le, Swish: a self-gated activation function 7 1, publisher: Technical report.
- [35] J. Bergstra, R. Bardenet, Y. Bengio, B. Kégl, Algorithms for hyperparameter optimization 24.

- [36] L. Li, K. Jamieson, A. Rostamizadeh, E. Gonina, J. Ben-Tzur, M. Hardt, B. Recht, A. Talwalkar, A system for massively parallel hyperparameter tuning 2 230–246.
- [37] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization.
- [38] L. N. Smith, No more pesky learning rate guessing games 5.
- [39] L. N. Smith, Cyclical learning rates for training neural networks, in: 2017 IEEE winter conference on applications of computer vision (WACV), IEEE, pp. 464–472.
- [40] L. N. Smith, N. Topin, Super-convergence: Very fast training of neural networks using large learning rates, in: Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications, Vol. 11006, International Society for Optics and Photonics, p. 1100612.
- [41] L. Prechelt, Early stopping-but when?, in: Neural Networks: Tricks of the trade, Springer, pp. 55–69.
- [42] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778.
- [43] H.-T. Cheng, L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhye, G. Anderson, G. Corrado, W. Chai, M. Ispir, others, Wide & deep learning for recommender systems, in: Proceedings of the 1st workshop on deep learning for recommender systems, pp. 7–10.
- [44] S. G. Krantz, H. R. Parks, The implicit function theorem: history, theory, and applications, Springer Science & Business Media.
- [45] G. M. Kontogeorgis, G. K. Folas, Thermodynamic models for industrial applications: from classical and advanced mixing rules to association theories, John Wiley & Sons.
- [46] M. S. Wertheim, Fluids with highly directional attractive forces. II. thermodynamic perturbation theory and integral equations 35 (1) 35–47, publisher: Springer.
- [47] M. Wertheim, Fluids with highly directional attractive forces. i. statistical thermodynamics 35 (1) 19–34, publisher: Springer.

- [48] M. Wertheim, Fluids with highly directional attractive forces. IV. equilibrium polymerization 42 (3) 477–492, publisher: Springer.
- [49] M. Wertheim, Fluids with highly directional attractive forces. III. multiple attraction sites 42 (3) 459–476, publisher: Springer.
- [50] W. G. Chapman, K. E. Gubbins, G. Jackson, M. Radosz, New reference equation of state for associating liquids 29 (8) 1709–1721, publisher: ACS Publications.
- [51] S. H. Huang, M. Radosz, Equation of state for small, large, polydisperse, and associating molecules 29 (11) 2284–2294, publisher: ACS Publications.
- [52] G. M. Kontogeorgis, E. C. Voutsas, I. V. Yakoumis, D. P. Tassios, An equation of state for associating fluids 35 (11) 4310–4318, publisher: ACS Publications.
- [53] G. M. Kontogeorgis, I. V. Yakoumis, H. Meijer, E. Hendriks, T. Moorwood, Multicomponent phase equilibrium calculations for water–methanol–alkane mixtures 158 201–209, publisher: Elsevier.