Supercritical CO₂ Power Cycles

Learning Conditionally Calibrated Equations of State for Direct Fired sCO₂ Cycles with Deep Neural Networks

Luke Vilnis, Andrew McCallum
UMass Amherst CICS
Amherst, MA

David Freed, Navid Rafati
8 Rivers Capital
Durham, NC

Joe Camilo
Duke University AMLL
Durham, NC

ABSTRACT

Equation of state (EOS) form the foundation for modeling the performance of carbon dioxide, and other, power cycles. Commercial software packages such as Aspen Plus rely on EOSs to inform and predict states of matter and interactions between matter under given conditions. Direct fired critical carbon dioxide systems offer unique challenges for standard equation of states found in the literature. In particular, the Allam Cycle utilizes a high-purity carbon dioxide working fluid across a wide range of conditions, including temperatures from 30C to 1150C and pressure from atmospheric pressure to 300 bar. As with all direct fired systems, it will have impurities in the working fluid that impact key engineering variables that are derived from the chosen equation of state. Peng-Robinson (PR) is a widely used EOS to describe the thermo-physical properties of pure CO_2 and its mixtures in process modeling packages for power, oil, gas, and petrochemical industries applications. The most common methodology of calibrating the PR EOS in mixtures is by using binary interaction parameters (k_{ij}) which are typically experimentally derived in controlled volume systems and combined using the van der Waals mixing rules. However, inaccuracies in predictions from most calibrations increase when extrapolated outside of a narrow range of conditions, or when considering multi-species mixtures. This paper presents a unique and broadly applicable methodology, treating these calibration parameters not as static, or sometimes temperature-dependent, quantities, estimated *in vitro,* but instead as *learned functions* allowed access to arbitrary side information, which is termed herein as *conditional calibration.* These functions are calibrated holistically, with respect to an entire system, rather than as independent scalar parameters learned through experiment, using the differentiability of the entire process model to enable gradient-based learning. This work demonstrates that a small but deep multi-layer neural network, generating the interaction parameters of a simple PR EOS-based model of a valve, significantly outperforms fixed parameter models in a pair of synthetic experiments. The first attempts to match a set of simulation data from a more sophisticated valve model using a simple model whose parameters are not fixed but generated, allowed access to temperature, pressure, and other features of the valve. The second adds an additional physically plausible pressure perturbation on top of the simulation. In both cases, conditional calibration gives relative reductions in error on the order of 40-50%, while the fundamental physical grounding in cubic EOS models keeps the "black-box" learner reined in and easy to estimate, leveraging the best

parts of machine learning and chemical process modeling. This work aims not only to offer a powerful methodology for empirically updating existing cubic equations of state with data from empirical environments, but also to advocate for a *flexible*, *conditional*, and *differentiable* approach to computational chemical engineering and process optimization in general.

INTRODUCTION

By using oxy-combustion and high-pressure supercritical CO_2 (sCO_2) in a recuperated cycle, the sCO_2 Allam Cycle is able to generate power with near zero atmospheric emissions. This breakthrough technology has the potential to realize low cost and clean power generation.
The Allam Cycle is an oxy-fired, trans-critical CO_2 cycle with a low-pressure-ratio turbine. The turbine exhaust gases are fed into a recuperator, cooled, water is condensed and separated, and the remaining vapor phase CO_2 is compressed and pumped up to high pressure before recovering the heat in the recuperator. The hot, high-pressure CO_2 then is added as coolant to the combustor inlet to achieve a final turbine inlet temperature of approximately 1150°C. The process is currently being demonstrated at a 50MWt demonstration plant in La Porte, TX.

Figure 1. An example of sCO₂ Allam Cycle process flow.

One of the technical challenges in designing a supercritical CO_2 power cycle is the physical behavior of the impure CO_2 at relevant conditions that are typically found in the power cycle [1]. These conditions vary extensively from moderate pressure regions at around 30 bar downstream of the supercritical CO_2 turbine to high pressure at around 300 bar at the inlet of the combustor. Operating temperature also varies from ambient to about 1150°C at turbine inlet [2]. Moreover, the type and concentrations of impurities such as O_2, N_2, and Ar will have

significant impact on the predictions of physical behaviors of a CO_2 mixture such as vapor-liquid equilibrium and density at various conditions. Accurate prediction of the physical properties of CO_2 mixtures has a significant impact on the reliability of process models developed in commercial process simulator packages such as Aspen Plus and PRO/II. In a real-world application, this translates to improved accuracy in turbo-machinery efficiency and equipment sizing that can directly impact the cost and thus the economic favorability of the process.

An Equation of state (EOS) is a widely-used tool to calculate thermodynamic behaviors of pure substances and mixtures. An EOS relates three state variables (temperature, pressure and volume), and can be used to derive a variety of thermo-physical properties such as enthalpy, entropy, and the vapor-liquid equilibrium behavior of mixtures. Peng-Robinson (PR) is one the most widely used EOS to model variety of processes in oil, gas, petrochemical and power industries. Although PR provides satisfactory prediction of variety of systems including CO_2 mixtures, it does not accurately cover a wide range of conditions [3].

PR is a cubic Equation of State that uses Van der Waals equations to expand its application from pure components to a multi-component mixture. While general EOSs with complex structures (GECS) such as Benedict-Webb-Rubin may give better results, they contain more parameters, which may not be available for all substance and are more difficult to integrate into commercial simulation tools [4] or calibrate with non-laboratory experimental data.

EOS Ex.	Function Form	Year	Mixing Rule
RK	$$P = \frac{RT}{V-b} - \frac{a/T^{0.5}}{V(V+b)}$$	1949	$a = \sum_i \sum_j x_i x_j a_i^{\frac{1}{2}} a_j^{\frac{1}{2}} (1 - k_{ij})$ $b = \sum_i x_i b_i \ \left(k_{ij} = k_{ji} \right)$
SRK	$$P = \frac{RT}{V-b} - \frac{a(T)}{V(V+b)}$$	1972	$a = \sum_i \sum_j x_i x_j a_i^{\frac{1}{2}} a_j^{\frac{1}{2}} (1 - k_{ij})$ $b = \sum_i x_i b_i \ \left(k_{ij} = k_{ji} \right)$
PR	$$P = \frac{RT}{V-b} - \frac{a(T)}{V(V+b) + b(V-b)}$$	1976	$a = \sum_i \sum_j x_i x_j a_i^{\frac{1}{2}} a_j^{\frac{1}{2}} (1 - k_{ij})$ $b = \sum_i x_i b_i \ \left(k_{ij} = k_{ji} \right)$
PT	$$P = \frac{RT}{V-b} - \frac{a(T)}{V(V+b) + c(V-b)}$$	1982	$a = \sum_i \sum_j x_i x_j a_i^{\frac{1}{2}} a_j^{\frac{1}{2}} (1 - k_{ij})$ $b = \sum_i x_i b_i \ \left(k_{ij} = k_{ji} \right)$ $c = \sum_i x_i c_i \ \left(k_{ij} = k_{ji} \right)$

Table 1. Several canonical cubic equations of state, including Peng-Robinson, denoted PR*.*

Binary interaction parameters, k_{ij}, are used in general models with simple structures (GESS) in the mixing rule noted in the table above. They are usually determined by fitting experimental data into the EOS equation. The binary interaction parameters account for the attraction forces between pairs of non-similar molecules. In the Li paper, as of 2008, the number of experimental points in the literature for various interactions with CO_2 were over several thousand with temperatures ranging from ambient to >1000K and pressures ranging from atmospheric to 600MPa [5]. Attempts have been made to improve the PR prediction accuracy for variety of supercritical CO_2 systems by tailoring the binary interaction parameters for a particular operation regime using experimental data, thereby correcting for the differences in mixing ideal and non-ideal gases [1, 6]. However, there still exist inaccuracies in modeling multi-component mixtures.

Recent studies have demonstrated the usefulness of incorporating machine learning (ML) in the modeling of certain chemical processes [6-9]. Modeling a chemical process with the flexibility of a non-parametric machine-learning algorithm can produce predictions that better reflect experimental data. A popular, non-parametric method for chemical process modeling is the artificial neural network (NN) [10]. NNs can have a high degree of flexibility and accuracy for predication, but require large amounts of training data to ensure a robust algorithm and prevent over-fitting. In general, as the complexity of the NN model increases, there is a corresponding increase in the number of example system states that need to be presented to the algorithm during the training process.

There are multiple studies that show the comparison of EOSs with NNs for predicting various quantities in sCO_2 systems. For example, the capability of a NN was demonstrated in modeling black pepper essential oil extraction process using sCO_2 [10]. This study used a large amount of training data to create a reliable NN model that outperformed in oil yield predictions as compared to the previous hand-crafted mathematical model. Another large survey on the suitability of neural networks as EOSs analyzed the prediction of solid solubility in sCO_2 [7], again finding them to outperform many standard EOSs. The improvement shown in these works significantly increases the usefulness of a simulation, leading to better informed design decisions.

Rather than replace centuries of physical and chemical intuition outright with NNs, the goal of this work is to explore the best of both worlds: low-bias, assumption-free nonlinear models, but without losing physical grounding. We refer to this approach, generating input-dependent calibration parameters of physically motivated equations of state, as *conditional calibration*. Rather than requiring enormous training data to learn physical facts, this provides an inductive bias that allows the use of powerful nonlinear learners with less data and less risk of bizarre, unphysical predictions. By viewing a segment of a process model as a differentiable computation graph, any and all parameters can be directly optimized by gradient methods, holistically within their natural and un-ideal environment. The experiments first demonstrate that conditional calibration with deep neural nets allows us to better model a more complex valve equation by a simpler one when compared to a fixed-parameter model, while learning directly from simulation data. Additionally, using the power of DNNs to condition on arbitrary side information, this work shows that conditional calibration beats fixed parameter baselines in modeling a quasi-physical

simulated perturbation, differing from EOS expected values. In both cases the improvements over fixed parameter models, not only with textbook k_{ij} coefficients, but with fixed parameters fit directly to minimize the same error, are on the order of a 40-50% relative reduction in error. The approach, while demonstrated on simulated data, can be used to tune an EOS in the real-world conditions of direct-fired sCO$_2$ power cycles, and more broadly.

METHODOLOGY

Valves provide for a confined test case for which to evaluate the performance of an EOS against a real system. For certain valves, upstream pressure, temperature and mass flow are known. In simulation, the EOS calculates molar volume where molar volume is the mixture volume divided by the number of moles within the mixture.

The volumetric flow rate upstream of the valve is calculated as:
$$Q = \frac{v * \dot{m}}{MW_{mix}}$$
Where v is the molar volume and MW_{mix} is the mixture molecular weight which is calculated as follows:
$$MW_{mix} = \sum_i x_i MW_i$$
Where i represents the various constituents of the mixture and x_i their mole fraction. For an oxy-combustion system, example constituents could include oxygen (O_2), carbon dioxide (CO_2), water (H_2O), nitrogen (N_2), Argon (Ar) and others.

The outlet pressure, in the simple liquid case, is found using the valve equation:
$$P_{out} = P_{in} - SG \left(\frac{Q}{Cv}\right)^2$$
Where SG is specific gravity, Cv is the flow coefficient of the valve and Q is volumetric flow rate.

Specific gravity is defined as:
$$SG = \frac{MW_{mix}}{v * \rho_{water}}$$

In addition to the simplified valve model, in an empirical situation, it is likely that P_{out} will be mispredicted given the EOS's inability to accurately predict density across all temperature and pressure conditions expected to be found in a direct-fired facility like the Allam Cycle. In this paper, to simulate this inaccuracy in the modeling, a P_{out}^{test} was created by perturbing the calculation of the outlet pressure in a non-random manner. The perturbation, for the purposes of modeling, was based on the relative opening of the valve. Assuming the error to be a bounded by 10psi on either side of the estimated pressure, as the valve opens, the pressure is more likely to be higher than the estimation. Similarly, as the valve closes, the pressure is more likely to be lower. The function for the perturbation is based on the valve characteristic curve. In this work, an equal percentage valve was used with the following equation from Aspen's standard valves:

$$V = \frac{0.01P^2}{\sqrt{2 - 10^{-8}P^4}}$$

A neural net is trained to take as input representation the inlet temperature, inlet pressure, and valve position, and to output the binary interaction parameters. The training method is described in the sequel. The range of inputs are chosen for a valve that would experience very wide pressure and temperature ranges within the dynamic simulation model.

MODEL ARCHITECTURE

This work relies crucially on the view of the equation of state, valve equation, and neural network as a single *differentiable computation graph*. The differentiable computation graph abstraction conceptualizes a system as a group of connected computational modules, each of which is capable of computing a derivative with respect to its input(s) and parameters. In this work, the first piece of the graph is a neural network, the definition of which is briefly reviewed below.

A simple type of feedforward neural network is a series of affine transformations followed by pointwise nonlinear functions, stacked in layers where the output of each is input to the next layer. The dimensionality of the affine transformation is called the number of *hidden units* in the layer. Often an affine transformation of the final layer's output produces the function of interest.

$$Layer_i(x) = f(W_i x + b_i)$$

$$h_{i+1} = Layer_{i+i}(h_i)$$

Where f is a pointwise nonlinearity. In this work, all experiments use the nonlinearity $f(x) = \max(0, x)$, known as the *rectified linear unit* or *ReLU* [11].

A NN with several layers is known as a deep neural network (DNN). These affine transformations can be fit to minimize a loss function, making DNNs a powerful class of nonlinear function approximators.

$$DNN(x) = Wh_n + b$$

$$x = h_0$$

Because each element of the weight matrix W_i represents an interaction parameter between one of the input units with one of the output units, these scalars are often called *connections* and a DNN can be visualized as a graph of these connections, as can be seen in Figure 2.

The model uses three input variables: inlet pressure, inlet temperature, and valve position. These three input variables are expanded as a 4th degree polynomial, giving a total of 35 input variables for the DNN. These variables are passed into a 4-layer DNN of 64 hidden units per layer, with a final affine transformation producing the necessary $n(n-1)$ k_{ij}'s, which are then reshaped and symmetrized into the interaction matrix of the Peng-Robinson EOS.
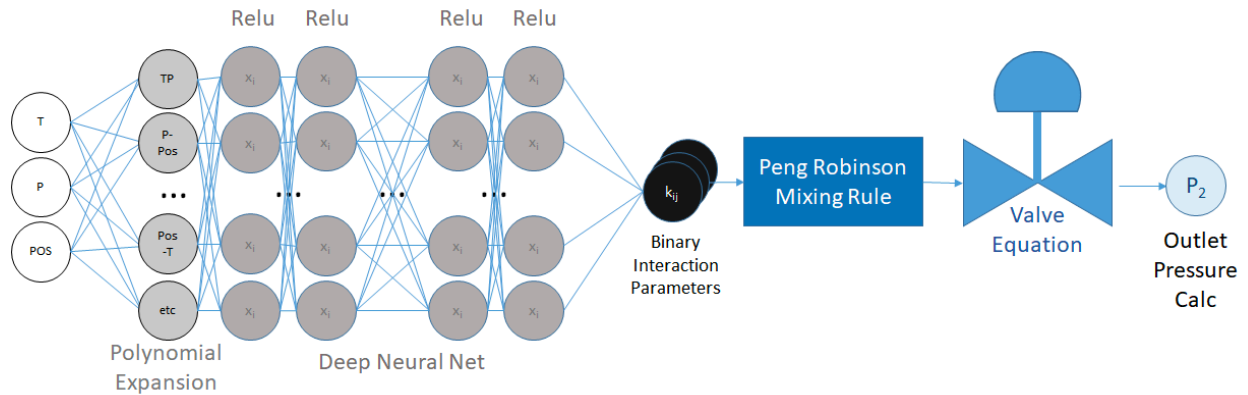
Figure 2. The conditionally calibrated EOS predicting valve outlet pressure, viewed as an end-to-end learnable computation graph. Note also the connections between nodes in the neural network, representing entries of the matrices W_i.

The PR EOS can itself be viewed as a component in the differentiable computation graph: a function of k_{ij} that outputs molar volume (given temperature, pressure, mixing proportions, critical properties, etc.) Since the need is only for the largest real root of the cubic, the solver is implemented with trigonometric and hyperbolic methods, sidestepping the necessity of complex numbers for simplicity. Omitting the equations for brevity, the solution for the largest real root of a cubic can be described by a piecewise combination of 3 nested series of radicals and hyperbolic trigonometric functions, each one a differentiable component of the graph.

The valve equation is obviously differentiable (at the points of interest), as is the squared error function between predicted and desired outlet pressure. The gradient of this error with respect to the k_{ij}'s and the parameters of the DNN can be computed in the same amount of time it takes to compute the function itself, using the backpropagation algorithm [12], also known as reverse-mode automatic differentiation [13], and the adjoint state method in PDE-constrained optimization [14].

The model is implemented in TensorFlow [15], a software package for differentiable computation graphs. The parameters are optimized with *minibatch stochastic gradient descent* (sometimes called the Robbins-Monro algorithm [16]), selecting 64 input-output pairs at a time at random and updating the parameters with the average error gradient and training until convergence.

EXPERIMENTAL DETAILS

This paper presents two sets of experiments, treating both unperturbed and perturbed simulation results from a commercial package incorporating more sophisticated vapor-fluid mixture modeling as "ground truth," and testing the ability of two different models to calibrate their EOS to the valve simulations, end-to-end. The first proposed model uses a fixed (but still learned by gradient descent in the same manner) parameter matrix of k_{ij}'s, the second is the DNN-based model which conditionally calibrates the EOS based on the current inputs. In order to evaluate the capability of a model that includes a powerful function approximator like a DNN,

the data is partitioned into a train/development/test split, holding out approximately 20% of the dataset of 4906 points at random as testing points, and another 20% as a development set. The parameters of the models are fit by gradient descent on the remaining 60% of the data, while architectural and optimization decisions (called *hyperparameters)* are evaluated by their performance on the development set. Final evaluation is performed by finding the model that performs best on the development set and applying it to the test data. This is most useful when using iterative methods like gradient descent and low-bias models like DNNs, as the rule of thumb to stop optimizing (or optimize slower) once your performance on the development data stops improving. The goal is not optimization *per se* but statistical learning; a low train error is undesirable if it fails to usefully generalize to the population distribution.

RESULTS AND DISCUSSION

As documented in Table 2, the conditionally calibrated model significantly outperforms the statically parameterized model, and both outperform the standard textbook EOS parameters since they are learned jointly in a >2 species mixture.

On the perturbed task, our approach achieves a ~44% relative reduction in error for dynamic calibration over static parameters fit through the same technique. On the unperturbed task, the best model achieves a remarkably similar ~45% relative reduction in error. Both static and conditional jointly calibrated parameters dominate the standard Peng Robinson k_{ij} for these synthetic tasks.

Many interesting phenomena are apparent when fitting both the static and conditionally calibrated models to empirical valve data. k_{ij}'s are considered "semi-physical," representing attraction between different species and calculated from VLE data in laborious pairwise experiments, and considered to be independent of temperature range, additional composition, and volume. There is evidence [16-18] that temperature and composition play a role, motivating the need for conditional calibration.

Model	Perturbed	Test AAD %
Standard k_{ij}	Yes	0.7052
Learned k_{ij} (fixed)		0.6625
Conditional Calibrated k_{ij}		
Depth		
2		0.5213
4		0.4112
6		0.3752
8		**0.3703**
Standard k_{ij}	No	0.4417
Learned k_{ij} (fixed)		0.4205
Conditional Calibrated k_{ij}		
Depth		
2		0.3659
4		0.3574
6		**0.2334**
8		0.2569

Table 2. Results of traditional fixed k_{ij}, fixed learned, and conditionally calibrated parameters on both tasks. Test AAD % refers to the metric of average absolute deviation on test data set, using the model with best development set performance. Bold numbers mark best performance.

Further, when directly fit to the conditions of a simulated system, quite different, often larger magnitude k_{ij} arise. Even in the unperturbed case, the computation graph model does not possess all of the liquid-vapor correction terms of the ground truth simulation and is able to improve its error versus the standard k_{ij}'s by gradient descent, and even further by conditional calibration using the same inputs of pressure, temperature, and valve position.

Similarly, in the perturbed experiments, explaining the positive or negative pressure perturbations require the conditionally calibrated model to posit large repulsive or attractive forces between species, learned from the input.

This stands to reason, as the calibration is attempting to explain significant modeling differences between the basic EOS and the gold-standard simulation, especially around the critical point, while being only able to affect the outlet pressure by modifying the k_{ij} parameters. While this is somewhat of an artificial constraint, it is meant to demonstrate several things:

(1)   Laboriously calculated binary EOS parameters can be calibrated on-the-fly to custom, non-idealized systems with heterogeneous conditions and multi-species mixtures in a way that is both less time consuming and more accurate within its area of data input.
(2)   Powerful conditional machine learning models like DNNs, allowing arbitrary covariates and side information to inform their predictions, can be safely "regularized", trained on little data, and prevented from overfitting by allowing them to only affect certain parameters of a semi-empirical model derived from reasonable physical principles, like cubic EOS, rather than having to learn a whole DNN-based EOS from scratch with the attendant burdens of massive training data and apprehensions regarding non-physicality.
(3)   Tunable semi-physical degrees of freedom in a process model, rather than posing estimation challenges *in vitro*, should be looked at as opportunities to fit directly *in vivo* to complex systems through automatic differentiation, with modern machine learning methods providing powerful tools to condition on sensors and other side information.

CONCLUSION AND FUTURE WORK

This work demonstrates through synthetic experiments that standard EOS can achieve large decreases in error by conditionally calibrating their parameters holistically with respect to a process model. We advocate not only for nonlinear conditional calibration of EOS, but also increased use of differentiable computation graph methods for computational chemical and process engineering in general.

Future work will continue to explore this exciting intersection between chemical process modeling and machine learning, past the proof-of-concept and towards the development of a rich "modeling language." Idealized component models (e.g. valves, turbomachinery and heat exchangers) and parsimonious, physically-motivated correction parameters, increase

tremendously in power when those parameters are dynamically conditioned on side-information through deep learning models, while simultaneously encouraging these over-parameterized models to generalize from far fewer training samples. Machine learning methods will help to ease the burden of parameter estimation and conditional modeling, while centuries of physical intuition provide a powerful inductive bias for the learner.

REFERENCES

[1] A new equation of state for CCS pipeline transport: Calibration of mixing rules for binary mixtures of CO_2 with N2, O2 and H2: https://doi.org/10.1016/j.jct.2015.07.045

[2] Demonstration of the Allam Cycle: An Update on the Development Status of a High Efficiency Supercritical Carbon Dioxide Power Process Employing Full Carbon Capture: https://doi.org/10.1016/j.egypro.2017.03.1731

[3] Selection of a Proper Equation of State for the Modeling of a Supercritical CO_2 Brayton Cycle: Consequences on the Process Design. DOI: 10.1021/acs.iecr.7b00917

[4] H. Li, J. Y. (2009). Evaluating cubic equations of state for calculation of vapor–liquid equilibrium of CO_2 and CO_2-mixtures for CO_2 capture and storage processes. Applied Energy 86, 826-836.

[5] Li, H. (2008). "Thermodynamic Properties of CO_2 Mixtures and Their Applications in Advanced Power Cycles with CO_2 Capture Processes . TRITA-CHE Report, 58

[6] Comparison the capability of artificial neural network (ANN) and EOS for prediction of solid solubilities in supercritical carbon dioxide. doi:10.1016/j.fluid.2011.06.002

[7] A. Tarjomannejad, "Prediction of the liquid vapor pressure using the artificial neural network-group contribution method," Iran. J. Chem. Chem. Eng., vol. 34, no. 4, pp. 97–111, 2015.

[8] M. J. Kamali and M. Mousavi, "Analytic, neural network, and hybrid modeling of supercritical extraction of α-pinene," J. Supercrit. Fluids, vol. 47, no. 2, pp. 168–173, 2008.

[9] M. Izadifar and F. Abdolahi, "Comparison between neural network and mathematical modeling of supercritical CO_2 extraction of black pepper essential oil," J. Supercrit. Fluids, vol. 38, no. 1, pp. 37–43, 2006.

[10] C. M. Bishop, "Machine learning and pattern recognition," Inf. Sci. Stat. Springer, Heidelb., 2006.

[11] Xavier Glorot, Antoine Bordes and Yoshua Bengio (2011). Deep sparse rectifier neural networks (PDF). AISTATS.

[12] Rumelhart, David E.; Hinton, Geoffrey E.; Williams, Ronald J. (8 October 1986). "Learning representations by back-propagating errors". Nature. 323 (6088): 533–536. doi:10.1038/323533a0.

[13] Rall, Louis B. (1981). Automatic Differentiation: Techniques and Applications. Lecture Notes in Computer Science. 120. Springer. ISBN 3-540-10861-0.

[14] R.-E. Plessix; A review of the adjoint-state method for computing the gradient of a functional with geophysical applications, *Geophysical Journal International*, Volume 167, Issue 2, 1 November 2006, Pages 495–503, https://doi.org/10.1111/j.1365-246X.2006.02978.x

[15] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. TensorFlow: a system for large-scale machine learning. In Proceedings of the 12th USENIX conference on Operating Systems Design and Implementation (OSDI'16). USENIX Association, Berkeley, CA, USA, 265-283.

[16] Robbins, Herbert, and Sutton Monro. "A stochastic approximation method." *The annals of mathematical statistics* (1951): 400-407.

[17] Moysan J., Huron M., Paradowski H., Vidal J. Prediction of the solubility of hydrogen in hydrocarbon solvents through cubic equation of state for fluids and fluids mixtures Chem. Eng. Sci. 77(3), 463-473, 1983

[18] Voros N.G., Tassios D.P. Vapor-liquid equilibria in non-polar/weakly polar systems with different types of mixing rules. Fluid Phase Equilibria 91:1-29, 1985.

[19] Bjorlykke O.P., Firoozabadi A. Measurement and computation of retrograde condensation and near critical phase behavior, In:

BIOGRAPHIES

Luke Vilnis is a Ph.D candidate in Computer Science at the University of Massachusetts Amherst, advised by Andrew McCallum. He received his M.S. in Computer Science from UMass in 2015 after passing his candidacy exam with distinction. Prior to that he received his B.S. in Mathematics and Economics from Duke University in 2010 and spent a couple of years as a software engineering consultant for the financial sector. He now researches machine learning. Favorite topics include knowledge representation, reasoning, and structured prediction, all of which go better with deep learning. Favorite applications include common sense, language (both generation and consumption), medicine, and energy. He has co-authored multiple publications awarded Best or Outstanding Paper, since two is a multiple of one.

David Freed is a Principal at 8 Rivers Capital, where he is the company's Project Manager for NET Power, LLC's 50MWt demonstration plant. Mr. Freed also leads the Intellectual Capital Committee at 8 Rivers, focused on company-wide technology development and deployment. As part of the initial and current NET Power team, Mr. Freed has supported the technical and commercial sides of the business, helping raise over $140M in equity capital for the company and additional grant capital. Mr. Freed also helped oversee the engineering, procurement, construction and now commissioning and testing of NET Power's demonstration plant in La Porte, TX. Prior to joining 8

Rivers, Mr. Freed worked at Boston Consulting Group in New York. He attended Duke University where he earned a Bachelor of Science in Engineering with distinction and a Certificate in Markets and Management. Mr. Freed was recently named by Forbes as the feature for energy in its 30 Under 30 awards and as a guest judge for the 2018 awards. For more information: david.freed@8rivers.com
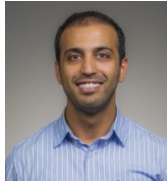
Navid Rafati is a senior process development engineer at 8 Rivers Capital. As a part of NET Power team, he has supported various design and engineering aspects 50MWth demonstration plant and also the commercial scale system. He is also involved in the development of several technologies for the production of energy, alternative fuels and chemicals. He received his bachelors and masters degrees in chemical engineering from University of Tehran and a PhD degree in energy and environmental studies from NC A&T. Prior to doctoral studies, he worked with a major EPC contractor and contributed to the design of South Pars Gas Field phases 15&16. For more information: navid.rafati@8rivers.com

Joseph A. Camilo received his B.S. from Clarkson University, Potsdam, NY, in 2013. He completed his M.S. and Ph.D. in electrical engineering at Duke University, Durham, NC, in 2015 and 2017, respectively. He is currently an associate in research in the Applied Machine Learning Lab (AMLL). His current research interests are machine learning, radar imaging, and computer vision applications.

Andrew McCallum is a Professor and Director of the Information Extraction and Synthesis Laboratory, as well as Director of Center for Data Science in the College of Information and Computer Science at University of Massachusetts Amherst. He has published over 250 papers in many areas of AI, including natural language processing, machine learning and reinforcement learning; his work has received over 50,000 citations. He obtained his PhD from University of Rochester in 1995 with Dana Ballard and a postdoctoral fellowship from CMU with Tom Mitchell and Sebastian Thrun. In the early 2000's he was Vice President of Research and Development at at WhizBang Labs, a 170-person start-up company that used machine learning for information extraction from the Web. He is a AAAI Fellow, the recipient of the UMass Chancellor's Award for Research and Creative Activity, the UMass NSM Distinguished Research Award, the UMass Lilly Teaching Fellowship, and research awards from Google, IBM, Microsoft, and Yahoo. He was the General Chair for the International Conference on Machine Learning (ICML) 2012, and is the current President of the International Machine Learning Society, as well as member of the editorial board of the Journal of Machine Learning Research. For the past ten years, McCallum has been active in research on statistical machine learning applied to text, especially information extraction, entity resolution, social network analysis, structured prediction, semi-supervised learning, and deep neural networks for knowledge representation. His work on open peer review can be found at http://openreview.net. McCallum's web page is http://www.cs.umass.edu/~mccallum.