

Autonomous Discovery of Unknown Reaction Pathways from Data by Chemical Reaction Neural Network

WeiQi Ji and Sili Deng*



Cite This: *J. Phys. Chem. A* 2021, 125, 1082–1092



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

ABSTRACT: Chemical reactions occur in energy, environmental, biological, and many other natural systems, and the inference of the reaction networks is essential to understand and design the chemical processes in engineering and life sciences. Yet, revealing the reaction pathways for complex systems and processes is still challenging because of the lack of knowledge of the involved species and reactions. Here, we present a neural network approach that autonomously discovers reaction pathways from the time-resolved species concentration data. The proposed chemical reaction neural network (CRNN), by design, satisfies the fundamental physics laws, including the law of mass action and the Arrhenius law. Consequently, the CRNN is physically interpretable such that the reaction pathways can be interpreted, and the kinetic parameters can be quantified simultaneously from the weights of the neural network. The inference of the chemical pathways is accomplished by training the CRNN with species concentration data via stochastic gradient descent. We demonstrate the successful implementations and the robustness of the approach in elucidating the chemical reaction pathways of several chemical engineering and biochemical systems. The autonomous inference by the CRNN approach precludes the need for expert knowledge in proposing candidate networks and addresses the curse of dimensionality in complex systems. The physical interpretability also makes the CRNN capable of not only fitting the data for a given system but also developing knowledge of unknown pathways that could be generalized to similar chemical systems.



1. INTRODUCTION

Discovering the reaction network in the chemical processes in energy conversion, environmental engineering, and biology is paramount to understanding the mechanisms of pollution and disease, as well as developing mitigation strategies and drugs. The traditional approach to building chemical reaction models is based on ab initio calculations and reaction templates developed with expert knowledge.¹ In many complex reacting systems, however, ab initio calculation is computationally intractable, and limited prior knowledge of the reaction templates has been established. Instead, time-resolved species concentration data are available because of the emerging sensing and measurement technologies^{2,3} (e.g., from high-throughput experiments in biology and materials sciences). As a result, a new paradigm of data-driven modeling has been the focus of recent efforts.^{4–14} Although many statistical regression methods have successfully learned the kinetic parameters for the given reaction pathways,^{15–19} further investigation is needed to simultaneously infer the reaction pathways and quantify the kinetic parameters from the data.

Achieving both the interpretability and accuracy of the data-driven chemical models is important and challenging. Numerous recent approaches^{7,10,12,13} leverage symbolic regression and sparse regression to identify the reaction pathways from the time series data of the species concentrations. Symbolic regression and sparse regression can produce physically interpretable kinetic models with known candidate pathways. However, such candidate pathways

can only be proposed for a few relatively simple systems because of the curse of dimensionality, as the possible interactions among species increase dramatically with the number of species.²⁰ For example, the number of possible reaction pathways scales with the fourth power of the number of species if we only consider the possible reactions involving two species in the reactants and two species in the products. In addition, multiple channel reaction pathways would require duplicated reactions, such as two or more reactions with the same set of reactants and products but different activation energies. On the contrary, neural networks have shown promise in autonomously learning the features of interactions among high-dimensional inputs.²¹ Traditional neural networks can approximate unknown reaction pathways,²² but the weights are difficult to interpret physically, that is, interpreting the reaction pathways and rate constants from the neural network weights to be correlated to traditional chemical models, limiting the capability of model generalization. The booming of deep learning has greatly benefited from designing problem-specific structures for neural network models. For

Received: October 14, 2020

Revised: December 21, 2020

Published: January 20, 2021



ACS Publications

© 2021 American Chemical Society

1082

<https://dx.doi.org/10.1021/acs.jpca.0c09316>
J. Phys. Chem. A 2021, 125, 1082–1092

example, the convolutional neural network incorporates the scale and rotation invariant of the vision. By connecting the convolutional kernel with traditional image filters, the weights of the kernel and the features learned by the kernel can be partially interpreted.²¹ Long et al.²³ have connected the residual neural networks to the first-order Euler method in the time-stepping numerical method and design specific convolutional kernels to reveal the differential operators and to discover the governing partial differential equations from the data. Conceptually inspired by these studies, we aim at designing neural networks with a problem-specific structure to learn chemical kinetic models from the data.

In this work, we propose a chemical reaction neural network (CRNN) to identify the reaction pathways from the data without any prior knowledge of the chemical system. We enable the interpretability of the neural networks by encoding governing physics laws into the architecture of the neural network while leveraging the stochastic gradient descent that is widely utilized by the deep learning community to optimize high-dimensional parameters. The physically interpretable feature enables us to relate the weights of the neural network to the reaction pathways as well as kinetic parameters and provide knowledge and insights on the chemical network. The capability of the autonomous discovery of unknown pathways enables us to learn from the given system and formulate reaction templates that can be generalized to similar systems, such that knowledge transfer can be achieved. Therefore, the CRNN approach enables the autonomous inference of new chemical reaction systems without prior knowledge of the reaction templates and has the potential to transform our understanding of complex reacting systems.

2. METHODS

Without loss of generality, we first derive the CRNN to represent an elementary reaction involving four species of $[A, B, C, D]$ with the stoichiometric coefficients of $[v_A, v_B, v_C, v_D]$.



Recalling the law of mass action discovered by Guldberg in 1879, the reaction rate r of eq 1 can be described by the power law expression, with the rate constant k and species concentrations of the reactants $[A]$ and $[B]$, as shown in eq 2. The expression can also be written as a cascade of exponential operation and weighted summation of species concentrations in the logarithmic scale, as shown in eq 3.

$$r = k[A]^{v_A}[B]^{v_B}[C]^0[D]^0 \quad (2)$$

$$r = \exp(\ln k + v_A \ln[A] + v_B \ln[B] + 0 \ln[C] + 0 \ln[D]) \quad (3)$$

Recall the formula of a neuron $y = \sigma(\mathbf{w}\mathbf{x} + b)$, in which $\mathbf{x}$ is the input to the neuron, y is the output, $\mathbf{w}$ are the weights, b is the bias, and $\sigma(\cdot)$ is the nonlinear activation function. Therefore, an elementary reaction can be represented as a neuron, as shown in Figure 1a. The inputs are the species concentrations in the logarithmic scale, and the outputs are the production rates of all species $\left[\frac{d[A]}{dt}, \frac{d[B]}{dt}, \frac{d[C]}{dt}, \frac{d[D]}{dt}\right]$, in short, $[[\dot{A}], [\dot{B}], [\dot{C}], [\dot{D}]]$. The weights in the input layer correspond to the reaction orders, that is, $[v_A, v_B, 0, 0]$ for $[A, B, C, D]$, respectively, as any species not presented in the reactants have the weights of zero. The bias corresponds to the rate constant in the logarithmic scale. The weights in the output layer




Figure 1. Schematic of the CRNN illustrated for a reaction system with four species and four reaction steps. (a) First, a neuron is designed based on the law of mass action and the Arrhenius law, such that the learned neural network can be translated into interpretable reactions corresponding to the traditional chemical reaction network. (b) Then, the neurons are stacked into one hidden layer to formulate a CRNN for multi-step reactions.

correspond to the stoichiometric coefficients, that is, $[-v_A, -v_B, v_C, v_D]$ for $[A, B, C, D]$, respectively.

In many chemical reaction systems, the rate constants are temperature-dependent. For example, the Arrhenius law discovered in 1889 can describe such dependence. The modified three-parameter Arrhenius formula states that

$$k = AT^b \exp\left(-\frac{E_a}{RT}\right) \quad (4)$$

where A is the prefactor or collision frequency factor, b is a fitting parameter to capture the nonexponential temperature dependence, E_a is the activation function, and R is the gas constant. Equation 4 can also be written as a linear operation of temperature and the prefactor A , as shown in eq 5, such that the Arrhenius law is also represented in the CRNN, as shown in Figure 1a.

$$\ln k = \ln A + b \ln T - \frac{E_a}{RT} \quad (5)$$

A reaction network involving multiple elementary reactions is therefore represented by forming a neural network with one hidden layer. For example, the CRNN representation of a multistep reaction network consisting of four species and four reactions is shown in Figure 1b. The number of hidden nodes is equal to the number of reactions.

As the weights and biases in the CRNN are physically interpretable, the CRNN is essentially the digital twin of the classical chemical reaction network. The inference of the reaction network can then be accomplished by training the CRNN with experimental measurements. Considering a general chemical reaction system where the vector of species concentration $\mathbf{Y}$ evolves with time, we are trying to discover a CRNN that satisfies




Figure 2. Schematic of the CRNN–ODE illustrated for a reaction system with five species $[A, B, C, D, E]$ and four reaction steps. (a) Y^{exp} corresponds to the noisy concentration time series, and Y^{CRNN} refers to the solution of ODE integrations of the CRNN model. (b) Learned CRNN weights consist of three blocks: input weights as the reaction orders (left), biases as the rate constants (middle), and output weights as the stoichiometric coefficients (right). (c) Interpreted reactions translated from the pruned CRNN weights and biases.

$$\dot{Y} = CRNN(Y) \quad (6)$$

The CRNN can be trained from the concentration and production rate data pair of $\{Y, \dot{Y}\}$. In practice, we are able to measure the concentration time series data, whereas it is challenging to measure the derivative directly because of the presence of noise in the measurements. The production rate can be approximated from noisy data, as demonstrated in a previous study;^{7,11} however, additional efforts are required to tune the hyperparameters, such as the regularization parameters, involved in the estimation.²³ Instead, we learn CRNN in the context of neural ordinary differential equations (ODE).²⁴ Specifically, an ODE system can be formulated with eq 6 and numerically solved in an ODE integrator by providing initial conditions. The solution is denoted as $Y^{CRNN}(t)$ in eq 7, in which Y_0 is a vector containing the initial conditions. We then define the loss function as the difference between the measured and predicted concentration time series data, as illustrated in Figure 2a. In the present work, the mean absolute error (MAE) is utilized as the loss metric in eq 8.

$$Y^{CRNN}(t) = ODESolve(CRNN(Y), Y_0) \quad (7)$$

$$loss = MAE(Y^{CRNN}(t), Y^{data}(t)) \quad (8)$$

The recently developed differential programming package of DifferentialEquations.jl²⁵ written in the Julia language has readily enabled the computation of the gradients of the above loss functions with respect to the CRNN parameters via backpropagation over the ODE integrators. Then, we can use the stochastic gradient descent optimization approach to learn the CRNN parameters,²⁴ such as the popular optimizer of Adam proposed by Kingma and Ba.²⁶ We denote the framework of learning CRNN using neural ODE as CRNN–ODE.

With all of the species known, the number of nodes in the hidden layer is the major hyperparameter to be determined, which corresponds to the number of reactions involved in the CRNN. We propose a grid searching approach to determine the number of hidden nodes, that is, increasing the number of proposed reactions until the performance cannot be further improved.

Finally, we employ hard threshold pruning to further encourage sparsity in the learned CRNN weights, especially the reaction orders and stoichiometric coefficients. The pruning proceeds by clipping the input and output weights below a certain threshold. The threshold is also determined by grid searching, as will be illustrated in Section 3. Then, we can

translate the pruned CRNN model into the classical form of reaction equations. The effect of weight pruning is illustrated in Figure 2b,c. Figure 2b shows the learned CRNN weights without pruning, and there exist many relatively small weights, such as the reaction order of 0.002 for species C in reaction R1. These small weights are then pruned, and the interpreted reactions are presented in Figure 2c.

3. RESULTS

We demonstrate the training of CRNN–ODE to simultaneously discover reaction pathways and learn kinetic parameters from synthetic noisy data in canonical chemical engineering and biochemistry systems, including both complete measurements and incomplete measurements with missing species. The ODE solver and the optimization are implemented with Julia, and the code is open source. The code is available on GitHub at <https://github.com/DENG-MIT/CRNN>.

3.1. Case I: An Elementary Reaction Network without Temperature Dependence. The first representative reaction network, taken from Searson et al.,²⁷ comprises five chemical species labeled as [A, B, C, D, E] that are involved in four reactions, as shown in eq 9 and Table 1. As the rate constants

Table 1. Learned Pathways Interpreted from the CRNN for Case I Trained with Noisy Data, with the Standard Deviation of the Noise Being 5% of the Concentrations

ground truth		learned CRNN	
equation	rate	equation	rate
$B + D \rightarrow E$	0.3	$B + 1.006D \rightarrow 1.006E$	0.307
$2A \rightarrow B$	0.1	$2.093A \rightarrow 1.107B$	0.101
$A \rightarrow C$	0.2	$1.004A \rightarrow 0.965C$	0.206
$C \rightarrow D$	0.13	$0.999C \rightarrow 1.011D$	0.13

are not temperature-dependent, only the law of mass action needs to be satisfied. The case is chosen to assess the capability of CRNN in discovering the chemical reaction pathways from noisy measurements and to demonstrate the CRNN–ODE algorithm.



A total of 30 synthetic experimental datasets are simulated, with the initial conditions randomly sampled between [0.2, 0.2, 0, 0, 0] and [1.2, 1.2, 0, 0, 0]. Each dataset comprises 100 data points evenly distributed in the temporal space with a time step of 0.4 time units. The simulation proceeds by solving the governing ODEs (shown in Supporting Information eq S1) using the solver of the Tsitouras 5/4 Runge–Kutta method.²⁵ Gaussian noise is added to the species concentration profiles, with the standard deviation of noise being 5% of the concentrations.

The 30 datasets are randomly split into the training datasets and validation datasets by the ratio of 2:1, that is, 20 of them as training datasets and 10 as validation datasets. The cross-validation helps assessing the level of overfitting, and we adopt

early stopping based on the loss in the validation dataset to prevent overfitting.²⁸ The training of CRNN proceeds in a fashion of mini-batch. Specifically, for each parameter update, only 1 of the 20 training datasets is used for computing the loss function and the gradients to the CRNN model parameters. The mini-batch approach accelerates the training and implicitly regularizes the CRNN model.²⁸ The input weights share parameters with the output weights to accelerate optimization, that is, the reaction orders are set to be equal to the stoichiometric coefficients for the reactants. The optimizer of Adam²⁶ is adopted with a learning rate of 0.001.

We design a CRNN with five inputs to model case I. The first step of the CRNN–ODE modeling pipeline is to determine the number of hidden nodes. We then train the CRNN with different numbers of hidden nodes to study the dependence of the model performance on the number of hidden nodes. The model performance can be measured by the average MAE loss function across all experimental datasets. Figure 3 shows the typical evolution of loss functions with the




Figure 3. Typical evolution of loss functions with the number of epochs. Results shown correspond to the CRNN with four hidden nodes.

number of epochs. Normally, the training converges at around 5000 epochs. Figure 4a then shows the evolution of the minimum loss functions for both the training and validation datasets with the number of hidden nodes. The training and validation loss decrease when the number of proposed reactions is less than four and reaches a plateau after that. It then can be inferred that the kinetics could be well described using four reactions.

We then learn a CRNN with four hidden nodes and present the learned weights and biases in Figure 5a. The dimensions of inputs and hidden nodes are physically interpretable. Each row of the matrix corresponds to a reaction. The weights in the input layer reveal the reaction orders of each reaction, and they are always positive. The weights in the output layer correspond to the stoichiometric coefficients, and positive/negative values indicate that the corresponding species are produced/consumed, whereas zero values indicate that the corresponding species do not participate in the reaction. The biases reflect the rate constants. Consequently, we can infer the corresponding four reaction pathways from the learned CRNN weights.

There are relatively small (compared to unity) nonzero input and output weights in the learned CRNN model. To further encourage sparsity, we apply a hard threshold pruning to the input and output weights. Specifically, all the weights




Figure 4. Dependence of minimum loss functions for the training and validation datasets with the (a) number of hidden nodes and the (b) pruning threshold.

(a) Learned CRNN weights before pruning.

	A	B	C	D	E	k	A	B	C	D	E
R1	0.	1	0.003	1.006	0.	0.307	0.002	-1	-0.003	-1.006	1.006
R2	2.093	0.	0.	0.002	0.008	0.101	-2.093	1.107	0.003	-0.002	-0.008
R3	1.004	0.	0.	0.	0.	0.206	-1.004	0.005	0.965	0.004	0.003
R4	0.	0.	0.999	0.	0.004	0.13	0.	0.002	-0.999	1.011	-0.004

(b) Learned CRNN weights after pruning with threshold = 0.01.

	A	B	C	D	E	k	A	B	C	D	E
R1	0.	1	0.	1.006	0.	0.307	0.	-1	0.	-1.006	1.006
R2	2.093	0.	0.	0.	0.	0.101	-2.093	1.107	0.	0.	0.
R3	1.004	0.	0.	0.	0.	0.206	-1.004	0.	0.965	0.	0.
R4	0.	0.	0.999	0.	0.	0.13	0.	0.	-0.999	1.011	0.

Figure 5. Learned CRNN weights and biases for case I with five inputs and four hidden nodes. (a) Learned CRNN weights before pruning. (b) Learned CRNN weights after pruning with threshold = 0.01.




Figure 6. Noisy species concentration profiles and the predictions using the learned CRNN model presented in Table 1.

that have absolute values below the threshold are clipped to zero. The threshold is also determined via grid searching. Figure 4b shows the dependence of the loss function with the threshold value. It is found that the loss function, that is, the model performance, is insensitive to the value of the threshold

when the threshold is much smaller than unity. For example, the threshold values of 0.01 and 0.5 make no difference in the model performance. Figure 5b shows the learned CRNN weights after pruning with the threshold of 0.01. The weights and bias are further translated to reaction equations and rate

constants, and they are shown in Table 1. The learned pathways are close to the ground truth, as the reactants and products are correctly learned. Furthermore, the learned stoichiometric coefficients and rate constants are also close to the ground truth values, with the maximum relative error within 10%. Finally, the noisy species concentration profiles and the predictions using the learned CRNN model presented in Table 1 are shown in Figure 6, and they agree very well.

It is worth mentioning that the learned stoichiometric coefficients and reaction orders are very close to integers, although we have not applied any regularization to force them to be close to integers. This makes the approach suitable for learning elementary reactions, in which the stoichiometric coefficients and reaction orders are usually assumed to be integers. In addition, the stoichiometric coefficients and reaction orders for species not participating are learned to be very close to zero.

It is not surprising that the learned CRNN model is not exactly the same as the ground truth. The uncertainties in the chemical kinetic model parameters depend on the noise level in the data, the amount of data, and the design of initial conditions. The parameter uncertainties can be estimated using Bayesian inference,¹⁷ and the computational framework developed for traditional chemical reaction networks can be extended to the CRNN framework straightforwardly. As a rule of thumb, the model parameter uncertainties will be reduced as the experimental data uncertainties reduce. We then present the learned CRNN model under 1% noise in Table 2. As can

under a moderate level of noise, for example, 5% noise, and reduce the number of experimental datasets required.


3.2. Case II: Biodiesel Production with Temperature Dependence. The second case is to demonstrate the capability of learning the temperature dependence of the rate constants, that is, the Arrhenius parameters in eq 4. The reaction system, studied by Burnham et al.,¹² is for biodiesel production via the transesterification of large, branched triglyceride (TG) molecules into smaller, straight-chain molecules of methyl esters. Darnoko and Cheryan³⁰ described three consecutive reactions, shown in eq 10, which produce three by-products, diglyceride (DG), monoglyceride (MG), and glycerol (GL). The governing equations are detailed in the Supporting Information eq S2. The prefactor A_0 in the logarithmic scale for the three reactions are [18.60, 19.13, 7.93], and the activation energy is [14.54, 14.42, 6.47] kcal/mol.

A total of 30 experiments are simulated, and the temperature is randomly drawn from [323, 343 K], according to the experimental studies by Darnoko and Cheryan.³⁰ The initial reactants are TG and ROH, with the initial conditions randomly sampled between [0.2, 0.2, 0, 0, 0] and [2.2, 2.2, 0, 0, 0]. Each dataset comprises 50 data points evenly spaced in time with a time step of 1. The 30 datasets are also randomly split into training and validation datasets with a ratio of 2:1. Same as the study in case I, we also add 5% Gaussian noise to the simulated species profiles. The training algorithms are the same as those for case I.

Table 2. Learned Pathways Interpreted from the CRNN for Case I Trained with Noisy Data, with the Standard Deviation of the Noise Being 1% of the Concentrations; Weights are Pruned with a Threshold of 0.0045

ground truth		learned CRNN	
equation	rate	equation	rate
$B + D \rightarrow E$	0.3	$B + 1.002D \rightarrow E$	0.303
$2A \rightarrow B$	0.1	$2.017A \rightarrow 1.013B$	0.099
$A \rightarrow C$	0.2	$0.999A \rightarrow 0.991C$	0.201
$C \rightarrow D$	0.13	$0.999C \rightarrow 1.004D$	0.13

be seen, the learned stoichiometric coefficients are closer to the ground truth compared to the results under 5% noise. In future works, we shall employ the design of experiments²⁹ into the modeling pipeline to reduce the model parameter uncertainties



We also first study the dependence of the minimum loss functions with the number of hidden nodes. Typical loss curves are presented in Figure S1. The training converges at around 6000 epochs. There is no obvious overfitting observed as the training and validation loss functions are close to each other. Figure 7a then shows that the loss functions decrease when the number of hidden nodes is less than three and almost unchanged when further increasing beyond three. Note that the model performances not necessarily strictly decrease with the number of hidden nodes because of the randomness in the training. Instead, the general trend is that the model




Figure 7. Dependence of minimum loss functions for the training and validation datasets with the (a) number of hidden nodes and the (b) pruning threshold.

(a) Learned CRNN weights before pruning.


	TG	ROH	DG	MG	GL	R'CO ₂ R	E _a	lnA	TG	ROH	DG	MG	GL	R'CO ₂ R
R1	1.	0.99	0.	0.	0.	0.	14.44	18.42	-1.	-0.99	1.01	0.	0.	1.
R2	0.	0.99	0.	0.99	0.	0.	6.42	7.81	0.01	-0.99	0.01	-0.99	1.01	1.09
R3	0.01	1.	1.	0.	0.	0.	14.43	19.18	-0.01	-1.	-1.	0.97	0.	0.95

(b) Learned CRNN weights after pruning with threshold = 0.015.

	TG	ROH	DG	MG	GL	R'CO ₂ R	E _a	lnA	TG	ROH	DG	MG	GL	R'CO ₂ R
R1	1.	0.99	0.	0.	0.	0.	14.44	18.42	-1.	-0.99	1.01	0.	0.	1.
R2	0.	0.99	0.	0.99	0.	0.	6.42	7.81	0.	-0.99	0.	-0.99	1.01	1.09
R3	0.	1.	1.	0.	0.	0.	14.43	19.18	0.	-1.	-1.	0.97	0.	0.95

Figure 8. Learned CRNN weights and biases for case II. (a) Learned CRNN weights before pruning. (b) Learned CRNN weights after pruning with a threshold of 0.015.**Table 3.** Learned Pathways Interpreted from the CRNN for Case II Trained with Noisy Data, with the Standard Deviation of the Noise Being 5% of the Concentrations

ground truth			learned CRNN		
equation	E _a	ln A	equation	E _a	ln A
TG + ROH → DG + R'CO ₂ R	14.54	18.60	TG + 0.99ROH → 1.01DG + R'CO ₂ R	14.44	18.42
MG + ROH → GL + R'CO ₂ R	6.47	7.93	0.99MG + 0.99ROH → 1.01GL + 1.09R'CO ₂ R	6.42	7.81
DG + ROH → MG + R'CO ₂ R	14.42	19.13	DG + ROH → 0.97MG + 0.95R'CO ₂ R	14.43	19.18

**Figure 9.** Noisy species concentration profiles and the predictions using the learned CRNN models. M1 corresponds to the CRNN learned from the complete dataset, M2 corresponds to the CRNN learned from the incomplete dataset, with DG included in the CRNN inputs and outputs, and M3 corresponds to the CRNN learned from the incomplete dataset but excluding DG from the CRNN inputs and outputs.

performance becomes insensitive to the number of hidden nodes beyond three. It is then suggested that the system can be well modeled with three reactions. The dependence of the loss functions with the threshold value is further shown in Figure 7b. Similar to the results in case I, the model performance is insensitive to the threshold value below unity. Figure 8 shows the weights of the learned CRNN model before pruning and after pruning with a threshold of 0.015, and the interpreted CRNN is compared with the ground truth in Table 3. Overall, the reaction pathways are accurately learned, with the maximum error in the stoichiometric coefficients within 10%.

Finally, the noisy species concentration profiles and the predictions using the learned CRNN model are shown in Figure 9, and they agree very well. The associated model is denoted as CRNN M1 in Figure 9.

Although the above demonstrations require training the CRNN with complete datasets, assuming that the concentrations of all species are measured, such complete datasets might not always exist in real applications. We then explore the feasibility to train CRNN with incomplete datasets. Without loss of generality, we assume that the concentration history of the intermediate species DG is missing. Then, we use the same

	TG	ROH	DG	MG	GL	R'CO ₂ R	E _a	lnA	TG	ROH	DG	MG	GL	R'CO ₂ R
R1	0.	1.	1.	0.	0.	0.	13.73	18.07	0.	-1.	-1.	0.97	0.	1.07
R2	0.	1.	0.	1.01	0.	0.	6.36	7.77	0.	-1.	0.	-1.01	1.01	1.01
R3	1.	1.01	0.	0.	0.	0.	14.85	19.02	-1.	-1.01	1.01	0.	0.	0.97

Figure 10. Learned CRNN weights and biases for case II trained with incomplete datasets. The weights are pruned with a threshold of 0.01.





Figure 11. Dependence of minimum loss functions for the training and validation datasets with the (a) number of hidden nodes and the (b) pruning threshold.

data generation and training algorithms as previously discussed for training with complete datasets, except that the species DG is not included in the loss functions. Figure S2 shows the dependence of the model performance with the number of hidden nodes, and it is suggested that the system can be modeled with three reactions. We then present the learned CRNN weights after pruning in Figure 10, and the predictions with the associated model denoted as CRNN M2 are shown in Figure 9. Satisfactorily, the learned CRNN model is also close to the ground truth, and the species profiles are well predicted. Moreover, with the learned CRNN model, we are also able to infer the species profiles of unmeasured DG.

While in the above demonstration we assume that we know the existence of the intermediate species DG in the system although we cannot measure DG, there are often cases where we do not know the existence of many intermediate species in prior. For instance, chemical kinetic modeling involves discovering not only reaction pathways but also unknown species. The CRNN approach offers a flexible framework for proposing new species by treating the number of inputs and outputs as an additional hyperparameter and using the grid searching to determine the minimum number of required species. For example, if we exclude the species DG from the CRNN inputs and outputs, the species profiles of the rest of the five species cannot be well predicted, such as the CRNN M3 shown in Figure 9. Therefore, the CRNN shall include at least six species. We shall further explore the potentials of CRNN in discovering unknown species in future studies.

3.3. Case III: An Enzyme Reaction Network. The third case is to demonstrate the learning of catalytic (enzyme) reactions in which some species are present in both reactants and products. The mitogen-activated protein kinase (MAPK) pathway, taken from the work of Hoffmann et al.,⁷ is an important regulatory mechanism for biological cells to respond to stimuli and is widely involved in proliferation, differentiation, inflammation, and apoptosis. The MAPK pathway consists of multiple stages of kinases that are either inactive or active (denoted by “*”). The activation occurs because of phosphorylation catalyzed by the kinase from the previous

stage, and the dephosphorylation is catalyzed by phosphatases. When the kinase is active, it can activate subsequent kinases for the next stage. Following Hoffmann et al.,⁷ the MAPK pathway is modeled with three stages of kinases, MAPK, MAP2K, and MAP3K. The initial stimulus is S, and the final substrate to be activated is a transcription factor TF. The ground truth reaction network consists of activation/phosphorylation and deactivation/dephosphorylation reactions, as shown in eq 11. The governing equations are shown in Supporting Information eq S3. All of the reaction rate constants are assigned to be 1.0.



A total of 100 experiments are simulated, with the initial conditions randomly sampled from the concentration of all species within [0.001, 1]. Each dataset comprises 100 data points evenly spaced in time, with a time step of 0.1. The 100 datasets are randomly split into training and validation datasets with a ratio of 70:30. Same as the study in cases I and II, we also add 5% Gaussian noise to the simulated species profiles. The training algorithms are the same as that for case I, except that the sharing parameters between the input weights and output weights are relaxed as the stoichiometric coefficients (output weights) for the catalysis could be zero, whereas the reaction orders (input weights) are nonzero.

We also first study the dependence of minimum loss functions with the number of hidden nodes. Typical loss curves are presented in Figure S3. The training converges at around 1000 epochs. Figure 11a then shows that the loss functions decrease when the number of hidden nodes is less than eight

	S1	S2	S3	S4	S5	S6	S7	S8	S9	k	S1	S2	S3	S4	S5	S6	S7	S8	S9
R1	1.	1.	0.	0.	0.	0.	0.	0.	0.	1.	0.	-1.	1.	0.	0.	0.	0.	0.	0.
R2	0.	0.	0.	0.	0.	0.	1.	1.01	0.	1.03	0.	0.	0.	0.	0.	0.	0.	-1.	1.
R3	0.	0.	1.	0.	0.	0.	0.	0.	0.	1.	0.	1.	-1.	0.	0.	0.	0.	0.	0.
R4	0.	0.	0.	0.	1.	0.	0.	0.	0.	1.	0.	0.	0.	1.	-1.	0.	0.	0.	0.
R5	0.	0.	0.	0.	0.	0.	1.	0.	0.	1.	0.	0.	0.	0.	0.	1.	-1.	0.	0.
R6	0.	0.	1.	1.	0.	0.	0.	0.	0.	0.99	0.	0.	0.	-1.	1.	0.	0.	0.	0.
R7	0.	0.	0.	0.	1.	1.	0.	0.	0.	1.	0.	0.	0.	0.	0.	-1.	1.	0.	0.
R8	0.	0.	0.	0.	0.	0.	0.	0.	1.	1.01	0.	0.	0.	0.	0.	0.	0.	1.	-0.99

Figure 12. Learned CRNN weights and biases for case III. The weights are pruned with a threshold of 0.01. Indices of species: “S”(1), “MAP3K”(2), “MAP3K”(3), “MAP2K”(4), “MAP2K”(5), “MAPK”(6), “MAPK”(7), “TF”(8), and “TF”(9).

Table 4. Learned Pathways Interpreted from the CRNN for Case III Trained with Noisy Data, with the Standard Deviation of the Noise Being 5% of the Concentrations

ground truth		learned CRNN	
equation	rate	equation	rate
$S + \text{MAP3K} \rightarrow S + \text{MAP3K}^*$	1	$S + \text{MAP3K} \rightarrow S + \text{MAP3K}^*$	1
$\text{MAPK}^* + \text{TF} \rightarrow \text{MAPK}^* + \text{TF}^*$	1	$\text{MAPK}^* + \text{TF} \rightarrow \text{MAPK}^* + \text{TF}^*$	1.03
$\text{MAP3K}^* \rightarrow \text{MAP3K}$	1	$\text{MAP3K}^* \rightarrow \text{MAP3K}$	1
$\text{MAP2K}^* \rightarrow \text{MAP2K}$	1	$\text{MAP2K}^* \rightarrow \text{MAP2K}$	1
$\text{MAPK}^* \rightarrow \text{MAPK}$	1	$\text{MAPK}^* \rightarrow \text{MAPK}$	1
$\text{MAP3K}^* + \text{MAP2K} \rightarrow \text{MAP3K}^* + \text{MAP2K}^*$	1	$\text{MAP3K}^* + \text{MAP2K} \rightarrow \text{MAP3K}^* + \text{MAP2K}^*$	0.99
$\text{MAP2K}^* + \text{MAPK} \rightarrow \text{MAP2K}^* + \text{MAPK}^*$	1	$\text{MAP2K}^* + \text{MAPK} \rightarrow \text{MAP2K}^* + \text{MAPK}^*$	1
$\text{TF}^* \rightarrow \text{TF}$	1	$\text{TF}^* \rightarrow 0.99\text{TF}$	1.01




Figure 13. Noisy species concentration profiles and the predictions using the learned CRNN model presented in Table 4.

and keep almost unchanged when further increasing the number of hidden nodes beyond eight. It is then suggested that the system can be well modeled with eight reactions. The dependence of the loss functions with the threshold value is further shown in Figure 11b. Similarly, the model performance is insensitive to the threshold value below unity. We then show the pruned weights in Figure 12 with a pruning threshold of 0.01 and compared the interpreted CRNN with the ground truth in Table 4. Overall, the reaction pathways are accurately learned, with the maximum error in the stoichiometric coefficients within 3%. Finally, the noisy species concentration profiles and the predictions using the learned CRNN model are shown in Figure 13, and they agree very well.

As catalytic species are present in both reactants and products, their stoichiometric coefficients shown in the output layer are zero. Together with the reaction orders inferred from the input weights, the participation of the catalyst and the catalytic pathways can be revealed, that is, the catalyst has a positive reaction order in the input weights. The successful demonstration of using CRNN to infer catalytic reaction systems shall open the possibility of future applications in biology and chemical engineering.

4. DISCUSSION

Machine learning techniques, especially deep neural networks, have revolutionized the fields of computer vision³¹ and natural

language processing. One of the key driving forces is the stochastic gradient descent with backpropagation for high-dimensional nonlinear and nonconvex optimization problems, which enables the autonomous learning of features and precludes the expert knowledge in designing grammar rules. However, a common challenge for applying machine learning techniques to physical problems is that it is difficult to impose physical constraints on the model. Specifically, the neural network trained solely based on fitting training data may fail to generalize to regimes beyond the training set. In the present work, we have demonstrated an approach to encode fundamental physics laws into the neural network structure, such that the learned model satisfies the physical constraints while maintaining the accuracy in fitting the data. Therefore, the physically interpretable model is expected to generalize across a wide range of regimes and enable knowledge transfer among similar chemical systems.

Despite the successful demonstrations of the physically interpretable CRNN, further studies are needed to improve the robustness and generality of the approach. For example, the current approach assumes that the chemical system is not very stiff, and all of the species concentrations are within similar orders of magnitude, such that we can use eq 8 to adequately represent the fitness of the CRNN model. Although this could be a challenge when the reacting systems involve a wide range of timescales and concentration levels, approaches such as adaptive weights for the loss components³² of different species could potentially tackle this problem by rescaling the loss functions based on the gradients of the CRNN model parameters.

Besides the encoded law of mass action and Arrhenius law, we acknowledge that the chemical systems are governed by many other physics laws as well. Encoding all of them into the structure of the neural network could be challenging. For example, the stoichiometric coefficients for elementary reactions should be integers; however, such nondifferentiable constraints could be difficult to implement with stochastic gradient descent. Although we do not enforce these constraints in the CRNN, the demonstrations show that the trained models automatically satisfy these constraints. Additional physical constraints such as element conservation and collision limit could potentially be included in defining the loss function to reduce the required training data and improve the generality. Identifying how “physical” the neural network needs to be is of interest for future study.

5. CONCLUSIONS

We have presented a CRNN approach for the autonomous discovery of reaction pathways and kinetic parameters from the concentration time series data. The CRNN is the digital twin of the classical chemical reaction network, and it is formulated based on the fundamental physics laws of the law of mass action and Arrhenius law. The reaction pathways and rate constants can be interpreted from the weights and biases of the CRNN. Stochastic gradient descent is adopted to optimize the large-scale nonlinear CRNN models. The approach is demonstrated in three representative chemical systems in chemical engineering and biochemistry. Both the reaction pathways and the kinetic parameters can be accurately learned. These demonstrations shall open the possibility of discovering a large number of hidden reaction pathways in life sciences, environmental sciences, and engineering, including but not

limited to gene expressions, disease progression, virus spreading, material synthesis, and energy conversion.

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jpca.0c09316>.

Governing ODEs, loss curves, and dependence of model performance on the number of proposed reactions (PDF)

■ AUTHOR INFORMATION

Corresponding Author

Sili Deng – Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, United States; orcid.org/0000-0002-3421-7414; Email: silideng@mit.edu

Author

Weiqi Ji – Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, United States; orcid.org/0000-0002-7097-0219

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acs.jpca.0c09316>

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

W.J. would like to acknowledge the help from Christopher Rackauckas on the usage of DifferentialEquations.jl. S.D. would like to acknowledge the support from Karl Chang (1965) Innovation Fund at Massachusetts Institute of Technology.

■ REFERENCES

- (1) Gao, C. W.; Allen, J. W.; Green, W. H.; West, R. H. Reaction Mechanism Generator: Automatic Construction of Chemical Kinetic Mechanisms. *Comput. Phys. Commun.* **2016**, *203*, 212–225.
- (2) Coley, C. W.; Thomas, D. A.; Lummiss, J. A. M.; Jaworski, J. N.; Breen, C. P.; Schultz, V.; Hart, T.; Fishman, J. S.; Rogers, L.; Gao, H.; Hicklin, R. W.; Plehiers, P. P.; Byington, J.; Piotti, J. S.; Green, W. H.; Hart, A. J.; Jamison, T. F.; Jensen, K. F. A Robotic Platform for Flow Synthesis of Organic Compounds Informed by AI Planning. *Science* **2019**, *365*, No. eaax1566.
- (3) Hanson, R. K.; Davidson, D. F. Recent Advances in Laser Absorption and Shock Tube Methods for Studies of Combustion Chemistry. *Prog. Energy Combust. Sci.* **2014**, *44*, 103–114.
- (4) Champion, K.; Lusch, B.; Kutz, J. N.; Brunton, S. L. Data-Driven Discovery of Coordinates and Governing Equations. *Proc. Natl. Acad. Sci. U.S.A.* **2019**, *116*, 22445–22451.
- (5) Costello, Z.; Martin, H. G. A Machine Learning Approach to Predict Metabolic Pathway Dynamics from Time-Series Multiomics Data. *npj Syst. Biol. Appl.* **2018**, *4*, 19.
- (6) Schmidt, M.; Lipson, H. Distilling Free-Form Natural Laws from Experimental Data. *Science* **2009**, *324*, 81–85.
- (7) Hoffmann, M.; Fröhner, C.; Noé, F. Reactive SINDy: Discovering Governing Reactions from Concentration Data. *J. Chem. Phys.* **2019**, *150*, 025101.
- (8) Ranade, R.; Alqahtani, S.; Farooq, A.; Echekki, T. An ANN Based Hybrid Chemistry Framework for Complex Fuels. *Fuel* **2019**, *241*, 625–636.
- (9) Rudy, S. H.; Brunton, S. L.; Proctor, J. L.; Kutz, J. N. Data-Driven Discovery of Partial Differential Equations. *Sci. Adv.* **2017**, *3*, No. e1602614.

- (10) Mangan, N. M.; Brunton, S. L.; Proctor, J. L.; Kutz, J. N. Inferring Biological Networks by Sparse Identification of Nonlinear Dynamics. *IEEE Trans. Mol. Biol. Multi-Scale Commun.* **2016**, *2*, 52–63.
- (11) Brunton, S. L.; Proctor, J. L.; Kutz, J. N. Discovering Governing Equations from Data by Sparse Identification of Nonlinear Dynamical Systems. *Proc. Natl. Acad. Sci. U.S.A.* **2016**, *113*, 3932–3937.
- (12) Burnham, S. C.; Searson, D. P.; Willis, M. J.; Wright, A. R. Inference of Chemical Reaction Networks. *Chem. Eng. Sci.* **2008**, *63*, 862–873.
- (13) Langary, D.; Nikoloski, Z. Inference of Chemical Reaction Networks Based on Concentration Profiles Using an Optimization Framework. *Chaos* **2019**, *29*, 113121.
- (14) Bongard, J.; Lipson, H. Automated Reverse Engineering of Nonlinear Dynamical Systems. *Proc. Natl. Acad. Sci. U.S.A.* **2007**, *104*, 9943–9948.
- (15) Wang, H.; Sheen, D. A. Combustion Kinetic Model Uncertainty Quantification, Propagation and Minimization. *Prog. Energy Combust. Sci.* **2015**, *47*, 1–31.
- (16) Frenklach, M. Systematic Optimization of a Detailed Kinetic Model Using a Methane Ignition Example. *Combust. Flame* **1984**, *58*, 69–72.
- (17) Najm, H. N.; Debusschere, B. J.; Marzouk, Y. M.; Widmer, S.; le Maître, O. P. Uncertainty Quantification in Chemical Systems. *Int. J. Numer. Methods Eng.* **2009**, *80*, 789–814.
- (18) Ji, W.; Wang, J.; Zahm, O.; Marzouk, Y. M.; Yang, B.; Ren, Z.; Law, C. K. Shared Low-Dimensional Subspaces for Propagating Kinetic Uncertainty to Multiple Outputs. *Combust. Flame* **2018**, *190*, 146–157.
- (19) Fröhlich, F.; Kaltenbacher, B.; Theis, F. J.; Hasenauer, J. Scalable Parameter Estimation for Genome-Scale Biochemical Reaction Networks. *PLoS Comput. Biol.* **2017**, *13*, No. e1005331.
- (20) Nagy, T.; Tóth, J.; Ladics, T. Automatic Kinetic Model Generation and Selection Based on Concentration versus Time Curves. *Int. J. Chem. Kinet.* **2020**, *52*, 109–123.
- (21) LeCun, Y.; Bengio, Y.; Hinton, G. Deep Learning. *Nature* **2015**, *521*, 436–444.
- (22) Ranade, R.; Alqahtani, S.; Farooq, A.; Echekki, T. An Extended Hybrid Chemistry Framework for Complex Hydrocarbon Fuels. *Fuel* **2019**, *251*, 276–284.
- (23) Long, Z.; Lu, Y.; Ma, X.; Dong, B. PDE-Net: Learning PDEs from Data. **2017**, arXiv Prepr. arXiv1710.09668.
- (24) Chen, R. T. Q.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D. Neural Ordinary Differential Equations. *Advances in Neural Information Processing Systems*; 2018; Vol. 2018, pp 6571–6583.
- (25) Rackauckas, C.; Nie, Q. DifferentialEquations.jl – A Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia. *J. Open Res. Software* **2017**, *5*, 15.
- (26) Kingma, D. P.; Ba, J. Adam: A Method for Stochastic Optimization. **2014**, arXiv Prepr. arXiv1412.6980.
- (27) Searson, D. P.; Willis, M. J.; Wright, A. Reverse Engineering Chemical Reaction Networks from Time Series Data. *Statistical Modelling of Molecular Descriptors in QSAR/QSPR*; Wiley-VCH Verlag GmbH & Co. KGaA: Weinheim, Germany, 2012; Vol. 2, pp 327–348.
- (28) Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT press, 2016.
- (29) Huan, X.; Marzouk, Y. M. Simulation-Based Optimal Bayesian Experimental Design for Nonlinear Systems. *J. Comput. Phys.* **2013**, *232*, 288.
- (30) Darnoko, D.; Cheryan, M. Kinetics of Palm Oil Transesterification in a Batch Reactor. *J. Am. Oil Chem. Soc.* **2000**, *77*, 1263–1267.
- (31) Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis.* **2015**, *115*, 211.
- (32) Wang, S.; Teng, Y.; Perdikaris, P. Understanding and Mitigating Gradient Pathologies in Physics-Informed Neural Networks. **2020**, arXiv Prepr. arXiv2001.04536.