
How Well Does Kohn-Sham Regularizer Work for Weakly Correlated Systems?

Bhupalee Kalita

Department of Chemistry, University of California, Irvine

Ryan Pederson and Jielun Chen (陈捷伦)

Department of Physics and Astronomy, University of California, Irvine

Li Li (李力)

Google Research

Kieron Burke

Department of Chemistry, University of California, Irvine and

Department of Physics and Astronomy, University of California, Irvine

Kohn-Sham regularizer (KSR) is a differentiable machine learning approach to finding the exchange-correlation functional in Kohn-Sham density functional theory (DFT) that works for strongly correlated systems. Here we test KSR for weak correlation. We propose spin-adapted KSR (sKSR) with trainable local, semilocal, and nonlocal approximations found by minimizing density and total energy loss. We assess the atoms-to-molecules generalizability by training on one-dimensional (1D) H, He, Li, Be, Be^{++} and testing on 1D hydrogen chains, LiH, BeH_2 , and helium hydride complexes. The generalization error from our semilocal approximation is comparable to other differentiable approaches, but our nonlocal functional outperforms any existing machine learning functionals, predicting ground-state energies of test systems with a mean absolute error of 2.7 milli-Hartrees.

Determining the ground-state properties of many-electron systems is fundamental to molecular modeling problems in chemical and material sciences. However, solving the Schrödinger equation explicitly for more than a few hundred electrons is computationally intractable. Among several methods of approximation, Kohn-Sham density functional theory (KS-DFT or simply DFT) [1, 2], a method based on the electron density distribution rather than the many-electron wave function, provides chemically useful results with $O(N^3)$ scaling for an N -electron system [3]. DFT is formally exact, but the exchange-correlation (XC) energy, resulting from the quantum-mechanical interaction between electrons, must be approximated in practice. Hundreds of XC energy functional approximations have been formulated in the past few decades [4]. Functionals can be designed non-empirically, for example using physics and chemical-based intuition and satisfying known exact constraints [5], or can involve some fitting to reference data [6]. However, in any approach, these functional approximations do not yield chemical accuracy in general, that is, with errors less than 1.6 milli-Hartrees (mH) in atomic units (or 1 kcal/mol). Improving the accuracy of XC functional approximations often incurs additional computational cost in the practical DFT calculation [7]. However, there is no systematic way in general to develop and improve XC functional approximations.

In recent years, machine learning (ML) has been used to find better DFT approximations. Attempts have been made to enhance either the speed or accuracy of DFT. Some used ML techniques to boost computational efficiency by approximating the non-interacting kinetic energy without solving the KS equations [8–11]. In an

effort to improve the accuracy of ML-DFT, a significant leap was achieved by Nagai et al. [12], who used a neural network (NN) model to approximate the XC functional and trained it with high accuracy coupled cluster (CCSD(T)) energies and densities of just three small molecules, while self-consistently solving the KS equations. This functional impressively generalized to 148 small molecules [13] to predict their energies and densities with accuracies comparable to human-designed functionals. However, the test set atomization energies were not chemically accurate. Also, they didn't have access to gradient information and were therefore limited to a gradient-free optimization scheme, which is inherently slow, often suffers poor convergence issues, and is difficult to scale to more complex NN models.

In DFT, many useful properties are extracted from the density, although an XC functional approximation need not produce accurate densities along with accurate energies [14]. In KS-DFT, we calculate the density self-consistently, and there is a nonlinear dependence of the XC functional on the density. Learning this relationship requires not only the ground truth mapping of the functional inputs to outputs but also how the functional performs in the underlying process. Hence the use of differentiable programming [15] becomes more intuitive [16]. With differentiable programming, conditioning the networks with physical insights becomes much simpler, and it can further help to ease the process of training.

Recently, Li et al. [17] made a valuable step in this direction by considering the entire DFT self-consistent calculation as a differentiable program. They implemented an end-to-end differentiable DFT code for 1-dimensional (1D) systems using JAX [18], a library

that provides differentiation, vectorization, just-in-time compilation, and other composable transformations of Python and NumPy programs [19]. They parameterized the XC functional with an NN which incorporated non-local information about the density, along with known physical constraints. The self-consistent KS calculations were embedded into the training process by backpropagating the gradients through the KS iterations. It was dubbed the Kohn-Sham regularizer (KSR). It could yield chemically accurate energies for uniformly separated 1D hydrogen chains at any separation by training on highly accurate energies and densities from only a few separations.

Following a similar approach, Kasim and Vinko [20] implemented an end-to-end differentiable DFT code in 3D for Gaussian-type orbitals and trained local and semi-local NN-based XC functional approximations, evaluating performance on small molecules. In another work, Dick et al. [21] constructed a semilocal XC functional that was carefully curated to account for several known exact conditions and pretrained to match SCAN, a popular meta-GGA functional [22]. While both of these works explore the generalizability of ML approximations for weakly correlated molecules with differentiable DFT codes, they do not incorporate global information, and their accuracy is limited to that of human-designed semilocal functionals. A slightly different approach involves introducing an ML correction term to a nonempirical or semi-empirical XC functional within a KS-DFT self-consistent framework [23, 24]. In such an approach, only a portion of the XC energy is approximated using ML and the functionals retain the characteristics of the baseline XC functional used. The recently proposed ML local hybrid functional, DM21 [25], addresses spin-symmetry breaking and delocalization error in DFT functionals. Consequently, it performs well on several main-group benchmark datasets and also correctly dissociates molecules. Unlike KSR, this functional is trained on large datasets of highly accurate reaction energies (not densities) in the loss function without explicitly supervising the self-consistent iterations.

[C3] Li et al. [17] explored the generalizability of KSR for a few strongly correlated systems with stretched bonds which is a completely different domain from most chemical applications of DFT. The aim there was to generate accurate binding energy curves (all the way to the dissociation limit) using the entire density (for the nonlocal approximation called global-KSR), using inputs at only two separations, for unpolarized hydrogen chains. The generalizability was in finding the entire bond-dissociation energy curve of these chains. Moreover, only the total density was used and not the spin densities.

In the present work, we propose spin-polarized versions of local, semilocal, and nonlocal XC functional approximations within a differentiable spin-DFT implementation of KSR. We modify these approximations to predict XC energy densities using

spin-densities as feature vectors while optimizing the NN parameters using total density and energy loss. Contrary to Ref. [17], we test the KSR approach in the domain of routine DFT calculations in chemistry, namely in and around equilibrium bond lengths. We find the remarkable result that training on energies and densities of a few atoms (and ions) alone produces accurate ground-state energies for equilibrium molecules (very reminiscent of the use of appropriate norms while avoiding using any covalent bond energies). We train and test on a variety of different elements, to obtain the generalizability relevant to chemistry. Almost all previous work in the chemical domain tests various approximate functional forms employing the standard ingredients locally [12, 20, 21]. Our work achieves high accuracy using the total density and is not limited to a specific set of human-chosen features.

The practical implementation of DFT involves solving the Kohn-Sham (KS) equations to calculate the ground-state electron density,

$$\left\{ -\frac{1}{2}\nabla^2 + v_s[n](\mathbf{r}) \right\} \phi_i(\mathbf{r}) = \epsilon_i \phi_i(\mathbf{r}). \quad (1)$$

The electron density, $n(\mathbf{r})$, is the sum of the probability density over all occupied one-electron KS orbitals, $n(\mathbf{r}) = \sum_i |\phi_i(\mathbf{r})|^2$. The KS potential, $v_s[n](\mathbf{r})$, contains the external one-body potential, the Hartree potential, and the XC potentials,

$$v_s[n](\mathbf{r}) = v(\mathbf{r}) + v_H[n](\mathbf{r}) + v_{XC}[n](\mathbf{r}). \quad (2)$$

The XC potential is the functional derivative of the XC energy, $E_{XC}[n]$, with respect to the electron density [2], $v_{XC}[n](\mathbf{r}) = \delta E_{XC}[n](\mathbf{r})/\delta n(\mathbf{r})$. We can express $E_{XC}[n]$ in terms of an XC energy density per electron, $\epsilon_{XC}[n](\mathbf{r})$:

$$E_{XC}[n] = \int d^3r \epsilon_{XC}[n](\mathbf{r}) n(\mathbf{r}). \quad (3)$$

The ground-state energy is calculated from the self-consistent density by summing the non-interacting kinetic energy, T_s , the external potential energy, V , the Hartree energy, U , and the XC energy,

$$E_0 = T_s[n] + V[n] + U[n] + E_{XC}[n]. \quad (4)$$

The computational efficiency is also affected by the level of approximation used for the XC functional [26].

Density matrix renormalization group (DMRG) [27] can be used to efficiently generate highly accurate benchmark energies and densities for these 1D analog systems. We can address such systems using 1D KS-DFT calculations as well with suitable XC energy functional approximations, such as the 1D local spin-density approximation (LSDA) which was constructed in Ref. 28 from the 1D exponentially repelling uniform electron gas.

In essence, KSR is a ML-DFT regularization technique that utilizes a differentiable analog of the standard

self-consistent DFT computational flow during training to train a suitable parameterized model for $E_{\text{XC}}[n] = E_{\text{XC},\theta}[n]$, where θ are trainable parameters [17]. In this work, we consider NN-based (neural) XC models, but KSR as a regularization technique can apply more broadly to any differentiable model choice. Knowledge of physical properties and constraints in the exact XC functional can help guide the construction of a neural XC approximation. The NN that parameterizes the XC functional in KSR is carefully curated to account for a few of the expected behaviors of the exact XC functional. Nonlocality is facilitated by adding a global convolution layer in $\epsilon_{\text{XC},\theta}[n]$ to help capture long-range interactions. The sigmoid linear unit (SiLU or Swish) [29, 30] activation function is used throughout because of its infinite differentiability. The KSR network is also complemented with a self-interaction gate (SIG) that partially cancels the self-interaction error by mixing in a portion of Hartree energy density to ϵ_{XC} .

In Ref. 17 several neural XC functional models were proposed: a local functional which only depends on the density at each point (KSR-LDA), a semi-local functional that uses local and gradient information about each point (KSR-GGA), and a global functional which utilized the global convolution layer and the SIG described above (KSR-global).

A main deficiency of the KSR technique in Ref. 17 is that it does not explicitly account for spin, and so may not generalize well for spin-polarized systems. Extending this technique and associated NN models to spin DFT requires a differentiable framework that can backpropagate through resulting spin densities. Spin is often incorporated in the neural XC functional using relative polarization, ζ , as a feature [12]. For up and down spin densities, $\{n_{\uparrow}, n_{\downarrow}\}$, $\zeta = (n_{\uparrow} - n_{\downarrow})/n$. While ζ can be introduced as an additional input channel to KSR neural ϵ_{XC} , its scale can be very different relative to n in general. Instead, we use up and down spin densities as input features, which have similar scales. The usual models and concepts for KSR can be extended to obtain a spin-adapted KSR (sKSR).

In sKSR-global, we have a global convolution layer that takes spin densities as inputs, and the kernel takes the form:

$$G(n_{\sigma}(x), \xi_p) = \frac{1}{2}\xi_p \int dx' n_{\sigma}(x') e^{-|x-x'|/\xi_p}, \quad (5)$$

where $\sigma \in \{\uparrow, \downarrow\}$ and ξ_p is a trainable parameter that represents an interaction scale. To keep the number of parameters comparable with KSR-global, we input each spin density to a global convolution layer consisting of 8 channels. We then concatenate the output on the channel dimension and input it to the latter convolution layers. For weakly correlated systems and greater generalizability, this approximation does not include any SIG. The rest of the network architecture is kept unchanged. sKSR-LDA and sKSR-GGA approximations to XC are devoid of global information. For sKSR-LDA,

two convolution layers with filter size one and 8 channels map the spin-density to ϵ_{XC} at the same spatial point x . In sKSR-GGA, we specify the total density gradient explicitly as an additional input channel along with the spin-densities. Instead of using one convolution layer with filter size three, we use three convolution layers with filter size one and 8 channels each. The rest of the sKSR-LDA and sKSR-GGA architectures are also similar to KSR-LDA and KSR-GGA. Fig. 1(a) shows the comparative network structures for all three types of approximations. In all cases, the resulting ϵ_{XC} is symmetrized with respect to the input of the up and down densities:

$$\epsilon_{\text{XC}}^{\text{symm}}[n_{\uparrow}, n_{\downarrow}] = \frac{1}{2} \left[\epsilon_{\text{XC}}[n_{\uparrow}, n_{\downarrow}] + \epsilon_{\text{XC}}[n_{\downarrow}, n_{\uparrow}] \right]. \quad (6)$$

Our approximation replaces the ϵ_{XC} in a spin-polarized self-consistent KS-DFT framework. For spin-polarized systems we perform the above spin-unrestricted KS-DFT procedure, however for unpolarized systems we use spin-restricted KS-DFT to preserve spin-symmetry. Fig. 1(b) shows the conventional computational flow and the flow of the gradients during the self-consistent optimization. To train the neural XC functional, we use the following loss function:

$$L(\theta) = \underbrace{\mathbb{E}_{\text{train}} [(E^{\text{sKSR}} - E^{\text{DMRG}})^2 / N_e]}_{\text{energy loss } L_E} + \underbrace{\mathbb{E}_{\text{train}} \left[\int dx (n^{\text{sKSR}} - n^{\text{DMRG}})^2 / N_e \right]}_{\text{density loss } L_n}, \quad (7)$$

where E^{sKSR} and n^{sKSR} are the converged total energy and total density obtained from the neural XC functional approximations, and E^{DMRG} and n^{DMRG} are the exact ground-state electronic energy and total density for each of the test systems. The total loss is evaluated as an expectation over training examples, where N_e is the number of electrons for a given training example. All quantities are in atomic units. We only consider the converged energy in the energy loss term rather than the energy trajectory throughout KS iterations, which was explored in Ref. 17. In this work we find that the self-consistent calculations converge quickly for the small atoms and ions used in training, and incorporating energy loss from each KS iteration minimally affects the efficiency of the optimization process. The gradients are calculated based on the total loss with respect to the parameters, θ , through automatic differentiation. They are back-propagated across the self-consistent cycles and the parameters of the neural XC functionals are updated until the total loss is minimized.

As a simple consistency test, we pose the question: can KSR learn human-designed functionals from their observable results? Here we specifically investigate whether sKSR-LDA can learn the relatively simple but general human-designed 1D LSDA XC functional. Since

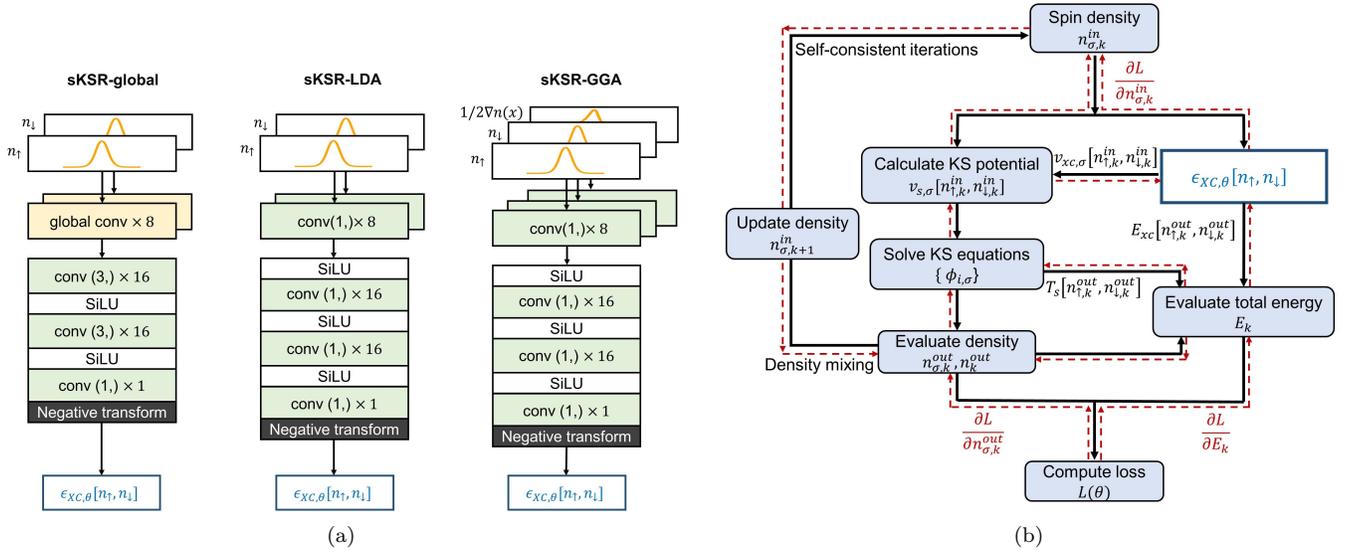


FIG. 1. (a) sKSR-global, sKSR-LDA and sKSR-GGA architectures to calculate ϵ_{xc} from spin-densities. (b) sKSR – differentiable KS-DFT with spin-polarization. Black arrows refer to the conventional computational flow. The gradients flow along red-dashed arrows to minimize the loss during training.

our sKSR-LDA model utilizes hundreds of parameters, it is unclear whether training on just a few LSDA generated DFT results will yield a neural XC model that matches LSDA. We find that by training sKSR-LDA on LSDA-generated He and Li^{++} , we recover the LSDA XC functional almost exactly for unpolarized and fully polarized systems, see Figure 2. The sKSR-LDA model deviates at the high-density limit (low r_s limit) due to the limitation that our training densities only consist of $r_s > 0.5$.

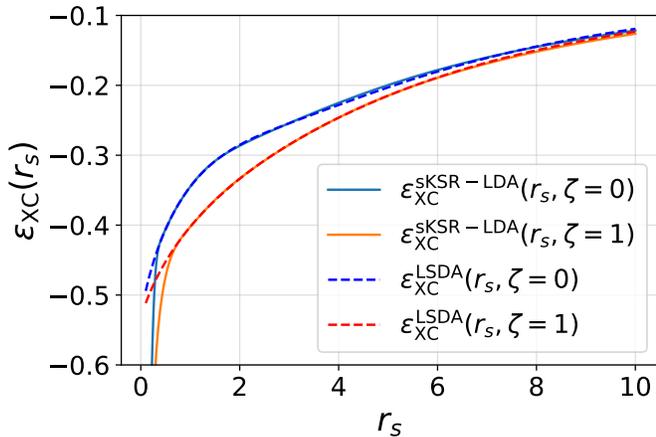


FIG. 2. sKSR-LDA trained on 1D LSDA-calculated Li^{++} and He energies and densities. Here $r_s = 1/2n$ and $\epsilon_{xc}^{\text{unif}}$ corresponds to the XC energy density of the 1D uniform electron gas [28].

Next, we assess generalizability by training sKSR models using a few 1D atomic systems and testing on

TABLE I. Training, validation and test sets for generalizability experiment. The molecules in the test set refer to the relaxed structures.

Training	Validation	Testing
H, He, Li	Be^+	$\text{H}_2, \text{H}_3, \text{H}_4, \text{H}_2^+, \text{H}_3^+$
Be, Be^{++}		$\text{LiH}, \text{BeH}_2, \text{HeH}^+$
		H-He-He-H^{2+}
		He-H-H-He^{2+}

unseen 1D molecular systems. We trained all three models on DMRG energies and densities of H, He, Li, Be, and Be^{++} and validated on Be^+ . For training and validation details, see Supporting Information. The trained model was later used to calculate the properties of several molecules in their equilibrium ground-state or relaxed form (see Table I). The errors in total energies, ionization, and atomization energies, as well as the average density losses for all three neural XC functional approximations, are reported in Table II. Compared to LSDA, the mean absolute error (MAE) in sKSR-LDA calculated energies is reduced by a factor of three. On the other hand, sKSR-global is an order of magnitude higher in accuracy and yields total energies with an MAE of 2.7 mH, not so far from the chemical accuracy limit of 1.6 mH. The cumulative MAEs for the training, validation and test datasets are reported in Supporting Information.

The importance of spin in sKSR can be seen by comparing results with the original KSR-global model from Ref. [17]. For a valid comparison, we consider KSR-global without the SIG and train it with the sets from Table. I, without adding the energy trajectory loss.

TABLE II. Total energy errors (in mH), density losses (in 10^{-4} Bohr $^{-1}$), and errors in ionization potentials for atoms and atomization energies in molecules (in mH) calculated using uniform gas LSDA [28], sKSR-LDA, sKSR-GGA, and sKSR-global respectively, for the training, validation, and test sets in Table I.

Dataset	Symbol	LSDA			sKSR-LDA			sKSR-GGA			sKSR-global		
		ΔE	L_n	ΔIP	ΔE	L_n	ΔIP	ΔE	L_n	ΔIP	ΔE	L_n	ΔIP
Training	H	26.6	5.35	-26.6	4.51	0.55	-4.50	4.49	0.31	-4.49	0.85	0.33	-0.85
	He	41.4	2.89	-8.46	20.2	0.63	-21.3	7.49	0.24	-10.2	-0.69	0.03	0.62
	Li	33.7	5.02	16.6	-11.5	0.40	37.4	-12.0	1.37	20.2	-2.37	0.12	2.79
	Be	24.5	1.18	21.4	-23.5	1.03	12.1	-2.70	0.65	-5.29	1.16	0.07	-1.23
	Be $^{++}$	55.3	0.75	-18.1	29.2	0.16	-46.1	6.55	0.49	-34.1	0.41	0.02	-1.43
	MAE	36.3	3.04	18.3	17.8	0.56	24.3	6.65	0.16	14.8	1.10	0.12	1.38
Validation	Be $^+$	46.0	1.95	9.37	-11.3	0.12	40.5	-7.99	0.61	14.5	-0.07	0.03	0.49
				ΔAE			ΔAE			ΔAE			ΔAE
Test	H $_2$	34.04	1.82	19.2	19.5	0.35	-10.5	6.83	1.99	2.14	-0.73	0.07	2.43
	H $_3$	35.6	1.93	44.3	0.45	0.21	13.1	-3.07	5.57	16.5	-3.56	3.22	6.11
	H $_4$	32.3	3.82	74.3	7.66	1.59	10.4	-9.34	4.18	27.3	2.87	1.46	0.53
	H $_2^+$	19.6	6.68	7.09	2.78	0.71	1.73	1.68	1.71	2.81	-1.94	1.04	2.79
	H $_3^+$	31.2	0.78	22.1	20.6	1.87	-11.6	15.4	11.5	-6.44	-0.40	0.47	2.09
	LiH	30.9	3.72	29.5	-8.55	2.47	1.53	-16.6	3.86	9.14	-4.38	0.66	2.86
	BeH $_2$	32.8	7.49	45.0	-27.8	5.5	13.4	-34.6	3.09	40.9	-5.07	1.29	7.93
	HeH $^+$	37.3	1.71	4.18	18.8	0.17	1.40	5.18	0.59	2.31	-1.60	0.13	0.91
	H-He-He-H $^{2+}$	36.7	14.7	46.1	5.00	6.00	35.5	-9.04	2.50	24.0	5.39	4.52	-6.77
	He-H-H-He $^{2+}$	46.1	7.40	36.7	19.9	6.48	20.6	4.35	4.75	10.6	0.79	5.47	-2.18
	MAE	33.6	5.00	32.9	13.1	2.53	12.0	10.6	3.98	14.2	2.67	1.83	3.46

The MAE in KSR-global predictions for total energies of the test molecules is 10.02 mH, comparable to sKSR-GGA, but much worse than sKSR-global (see Table. A2 in Supporting Information). sKSR-global also converges more quickly than KSR-global, reaching lower training losses with fewer training steps (see Fig. A7 in Supporting Information).

The size of our dataset is practically limited by the chemical space provided by 1D and the associated exponential interaction. Even though we are dealing with a much smaller dataset, we trained the sKSR models on the ground-state energies and densities of 5 atomic systems only and did not include any molecules, contrary to results in Ref. 12 and Ref. 20 which train on derived quantities, such as atomization and ionization energies, and include molecules in training.

Using sKSR-global, the predicted densities of each molecule have little noticeable error, see Fig. 3(a). The corresponding XC potentials are shown in Fig. 3(b). For all unpolarized systems, we run restricted KS calculations, and the up and down XC potentials match, while for polarized systems (Li, Be $^+$, H $_2^+$, and H $_3$ only) we run unrestricted KS calculations. The sKSR-LDA and sKSR-GGA total densities and XC potentials for the test set are included in the Supporting Information. The comparison to exact XC potentials is not expected to be as precise as potentials are extremely sensitive to densities. However, for each of these examples, we see

that the sKSR-global XC potential closely mimics the exact XC potential, even though we did not include XC potentials in the training. Furthermore, seemingly large deviations in the XC potentials can result in similar resulting densities. For example, this can be seen in the case of BeH $_2$ where the XC potentials are noticeably different but the resulting densities are very similar. The KS potentials are reasonably accurate for the test set (see Supporting Information). Note that similar to the exact XC potentials, the sKSR-global XC potentials are smooth, due to the use of a smooth activation function.

We can use these potentials to validate the known theoretical properties of the exact XC potentials for different test systems, compare with other XC approximations, and utilize them to introduce corrections to existing local and semilocal approximations. Similarly, sKSR-global can also produce quite accurate spin-densities even though we did not incorporate spin-densities in the loss function while training the XC functionals (see Fig. A1 in Supporting Information).

A very interesting question is: how does our weakly-correlated sKSR behave for strongly-correlated systems? We answer this by studying the paradigm case of the H $_2$ binding curve in Fig. 4, where the sKSR-global curve remains highly accurate up to at least 3 Bohr. Just as with all single-particle methods, the restricted calculation yields energy that is far too high in the

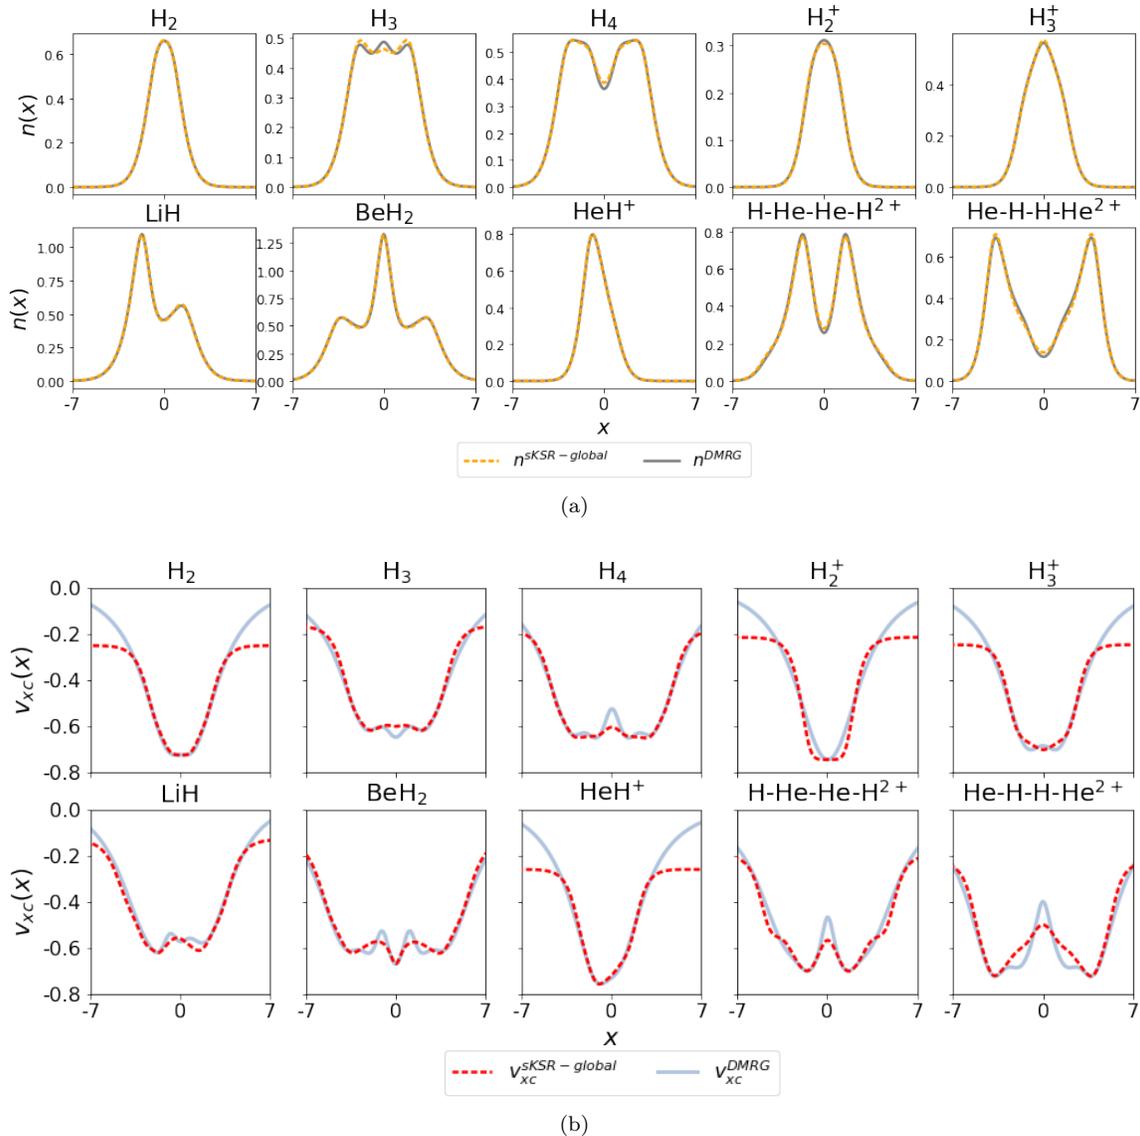


FIG. 3. (a) The densities obtained using sKSR-global (orange dashes) and the exact ground-state densities (gray), (b) average XC potentials calculated from sKSR-global approximation (red dashes) to ϵ_{xc} and their exact counterparts calculated with DMRG (light blue) for the test molecules in Table. I at equilibrium separations. The sKSR potentials are shifted by a constant for a better comparison with the exact XC potentials. sKSR-global was trained on H, He, Li, Be, and Be^{++} and validated on Be^+ . Note that, in general, these 1D densities and XC potentials can differ even qualitatively from their 3D analogs.

dissociated limit. On the other hand, an unrestricted calculation, which breaks spin-symmetry beyond about 4 Bohr, does dissociate correctly, but at the price of poor spin densities and a kink in the binding energy curve. Fig. A6 in Supporting Information shows analogous features for sKSR-LDA and sKSR-GGA, and also shows the accuracy of the total density of the unrestricted solutions at large separations. Fig. 4 also shows the result of a KSR-global calculation (i.e., total density only), but trained just on atoms. While it naturally dissociates correctly, it is much less accurate. Of course, the KSR-global of Ref. 17 is chemically accurate for the entire curve because its training included a stretched bond.

In many cases, the predictability of sKSR can extend well beyond the equilibrium bond distance. Fig. 5 shows the complete dissociation energy curve of LiH obtained from restricted calculation. Near equilibrium, sKSR-LDA and sKSR-GGA underestimate the binding energy but perform better than LSDA. As the bond is stretched, sKSR-GGA and sKSR-LDA quickly deviate from the expected trajectory. However, sKSR-global performs well throughout, extending its predictive accuracy well beyond the equilibrium bond distance. We show the total density and the XC potential of stretched LiH at 5.92 Bohr in Fig. 6. LSDA largely overestimates the total energy of the stretched molecule, but its

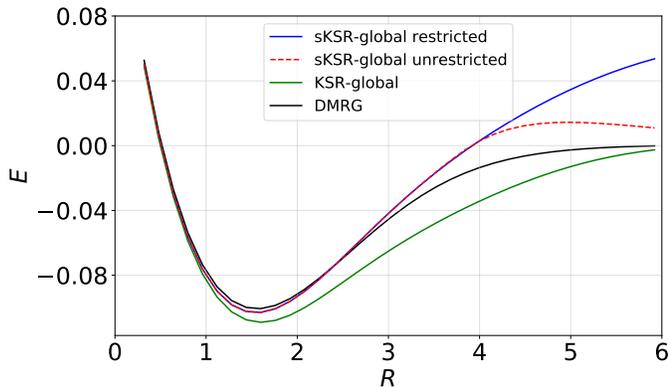


FIG. 4. The binding energy curve of H_2 molecule calculated based on the total energy prediction for H_2 molecule and the energy of the individual H atoms. sKSR-global was evaluated using restricted KS (blue) and unrestricted KS (red dashes) scheme. The DMRG (black) and KSR-global (green) results are also shown. All the neural approximations, with and without spin, are trained on the dataset given in Table. I.

density remains reasonably accurate. The XC potentials calculated from neural XC functional approximations are comparable, with sKSR-global closely approximating the exact behavior. A comparison of the sKSR-global and the exact total density and XC potential of stretched LiH with respect to the atomic contributions from Lithium and Hydrogen is included in the Supporting Information.

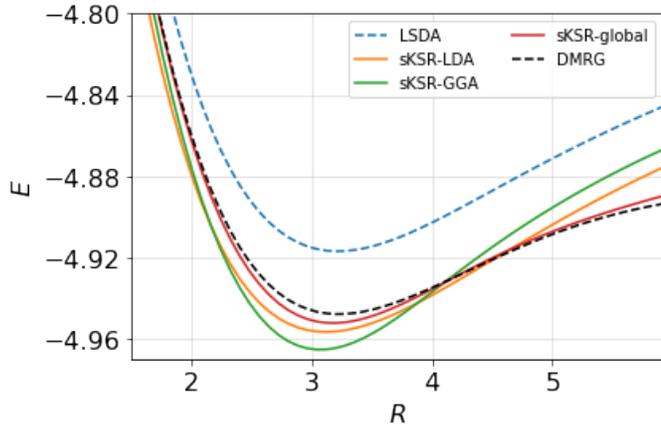
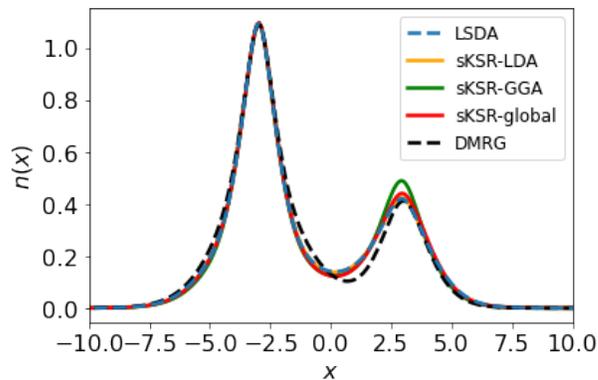
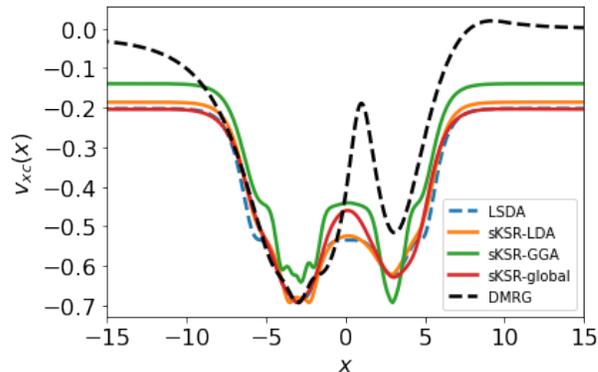


FIG. 5. The complete dissociation energy curve of LiH molecule generated with sKSR-LDA (orange), sKSR-GGA (green) and sKSR-global (red). The DMRG (black dashes) and the uniform gas LSDA (blue dashes) results are also shown. The neural XC functional approximations were trained and validated on atoms and ions given in Table. I.

The approximate total energy of a molecule can have two types of error contributions: the error due to the approximate functional and the error arising from the self-consistent density [31]. For most XC functionals, the total density calculated from the self-consistent solution



(a)



(b)

FIG. 6. (a) The total density and (b) the average XC potentials of LiH at a bond-distance of 5.92 Bohr calculated with the three neural XC functionals as well as uniform-gas LSDA. The exact (DMRG) average XC potentials are included for comparison.

of the KS equations works as an excellent approximation to the exact density for most systems. Hence, the density-driven error is often negligible. However, some approximations can have significant density-driven errors [32]. For our test molecules, the errors in the self-consistent densities were trivial and consequently had minimal impacts on the atomization energy errors. The functional and density-driven errors in our neural XC functional approximations are reported for the hydrogen molecule in the supplementary information section.

We found that sKSR-global achieves remarkable accuracy and generalization in a very data-efficient manner by including the self-consistent KS equations into the training. sKSR-global predicts the ground-state energy of ten unseen molecules in equilibrium with a mean absolute error of 2.7 mH (~ 1.7 kcal/mol) when trained with just five atoms and ions. Hence, a nonlocal XC functional approximation trained on atomic energies and densities can generate predictions for weakly-correlated molecules with near chemical accuracy. An extension of this work can lead to an ML functional that is applicable across a broad chemical spectrum

without using an exceedingly large training set. The end-to-end differentiable implementation also ensures smooth and reasonable XC potentials. In addition, sKSR-global trained on atoms can adequately describe a molecule with a stretched bond. Combining differentiable programming with inherent physical intuition thus takes us one step closer to a generalizable, chemically accurate ML XC functional.

The application of the current sKSR algorithm is limited to 1D systems and our test set does not include real 3D molecules. However, the methods presented are transferable to 3D and we anticipate that the characteristic performance is not unique to 1D systems, as these systems tend to mimic their 3D analogs [33]. The low-dimensional examples are useful for quick and rigorous assessment of the quality of an approximation. Besides, the predictions from the local and semilocal approximation explored in our study are consistent with the 3D differentiable formulations in Ref 20 and Ref. 21.

ACKNOWLEDGMENTS

This work is supported by National Science Foundation, grant no. DGE-1633631 (B. K.), CHE-1856165 (B. K., K. B.), and Department of Energy, grant no. DE-SC0008696 (R. P.).

1. SUPPORTING INFORMATION

The Supporting Information is available at <https://pubs.acs.org/doi/10.1021/acs.jpcllett.2c00371>.

2. DATA AVAILABILITY STATEMENT

The training and testing data and the one-dimensional density functional theory solver used for the uniform electron-gas LDA calculations are available at https://github.com/pedersor/DFT_1d. The ML models and the JAX version of the DFT code are available from the corresponding author upon reasonable request.

- [1] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* **136**, B864 (1964).
- [2] W. Kohn and L. J. Sham, Self-consistent equations including exchange and correlation effects, *Phys. Rev.* **140**, A1133 (1965).
- [3] R. M. Dreizler and E. K. U. Gross, *Density Functional Theory: An Approach to the Quantum Many-Body Problem* (Springer Berlin Heidelberg, Berlin, Heidelberg, 1990).
- [4] N. Mardirossian and M. Head-Gordon, Thirty years of density functional theory in computational chemistry: an overview and extensive assessment of 200 density functionals, *Molecular Physics* **115**, 2315 (2017).
- [5] J. P. Perdew, K. Burke, and M. Ernzerhof, Generalized gradient approximation made simple, *Phys. Rev. Lett.* **77**, 3865 (1996).
- [6] Y. Zhao and D. G. Truhlar, A new local density functional for main-group thermochemistry, transition metal bonding, thermochemical kinetics, and noncovalent interactions, *The Journal of Chemical Physics* **125**, 194101 (2006), <https://doi.org/10.1063/1.2370993>.
- [7] K. Burke, Perspective on density functional theory, *J. Chem. Phys.* **136**, 150901 (2012).
- [8] J. C. Snyder, M. Rupp, K. Hansen, K.-R. Müller, and K. Burke, Finding density functionals with machine learning, *Phys. Rev. Lett.* **108**, 253002 (2012).
- [9] L. Li, J. C. Snyder, I. M. Pelaschier, J. Huang, U.-N. Niranjan, P. Duncan, M. Rupp, K.-R. Müller, and K. Burke, Understanding machine-learned density functionals, *International Journal of Quantum Chemistry* **116**, 819 (2016), <https://onlinelibrary.wiley.com/doi/pdf/10.1002/qua.25040>.
- [10] F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K.-R. Müller, Bypassing the Kohn-Sham equations with machine learning, *Nature Communications* **8**, 872 (2017).
- [11] B. Kalita, L. Li, R. J. McCarty, and K. Burke, Learning to Approximate Density Functionals, *Accounts of Chemical Research* **54**, 818 (2021).
- [12] R. Nagai, R. Akashi, and O. Sugino, Completing density functional theory by machine learning hidden messages from molecules, *npj Computational Materials* **6**, 43 (2020).
- [13] L. A. Curtiss, K. Raghavachari, G. W. Trucks, and J. A. Pople, Gaussian-2 theory for molecular energies of first- and second-row compounds, *The Journal of Chemical Physics* **94**, 7221 (1991), <https://doi.org/10.1063/1.460205>.
- [14] M.-C. Kim, E. Sim, and K. Burke, Understanding and reducing errors in density functional calculations, *Phys. Rev. Lett.* **111**, 073003 (2013).
- [15] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, Automatic differentiation in machine learning: A survey, *J. Mach. Learn. Res.* **18**, 5595–5637 (2017).
- [16] M. Innes, A. Edelman, K. Fischer, C. Rackauckas, E. Saba, V. B. Shah, and W. Tebbutt, A differentiable programming system to bridge machine learning and scientific computing (2019), [arXiv:1907.07587](https://arxiv.org/abs/1907.07587).
- [17] L. Li, S. Hoyer, R. Pederson, R. Sun, E. D. Cubuk, P. Riley, and K. Burke, Kohn-sham equations as regularizer: Building prior knowledge into machine-learned physics, *Phys. Rev. Lett.* **126**, 036401 (2021).
- [18] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, *JAX: composable transformations of Python+NumPy programs* (2018).
- [19] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, Array programming with NumPy, *Nature* **585**, 357 (2020).
- [20] M. F. Kasim and S. M. Vinko, Learning the exchange-correlation functional from nature with fully differentiable density functional theory, *Phys. Rev. Lett.* **127**, 126403 (2021).
- [21] S. Dick and M. Fernandez-Serra, Using differentiable programming to obtain an energy and density-optimized exchange-correlation functional (2021), [arXiv:2106.04481](https://arxiv.org/abs/2106.04481).
- [22] J. Sun, A. Ruzsinszky, and J. P. Perdew, Strongly constrained and appropriately normed semilocal density functional, *Phys. Rev. Lett.* **115**, 036402 (2015).
- [23] S. Dick and M. Fernandez-Serra, Machine learning accurate exchange and correlation functionals of the electronic density, *Nature Communications* **11**, 3509 (2020).
- [24] Y. Chen, L. Zhang, H. Wang, and W. E, DeePKS: A Comprehensive Data-Driven Approach toward Chemically Accurate Density Functional Theory, *Journal of Chemical Theory and Computation* **17**, 170 (2021).
- [25] J. Kirkpatrick, B. McMorrow, D. H. P. Turban, A. L. Gaunt, J. S. Spencer, A. G. D. G. Matthews, A. Obika, L. Thiry, M. Fortunato, D. Pfau, L. R. Castellanos, S. Petersen, A. W. R. Nelson, P. Kohli, P. Mori-Sánchez, D. Hassabis, and A. J. Cohen, Pushing the frontiers of density functionals by solving the fractional electron problem, *Science* **374**, 1385 (2021), <https://www.science.org/doi/pdf/10.1126/science.abj6511>.
- [26] J. P. Perdew, A. Ruzsinszky, J. Tao, V. N. Staroverov, G. E. Scuseria, and G. I. Csonka, Prescription for the design and selection of density functional approximations: More constraint satisfaction with fewer fits, *The Journal of Chemical Physics* **123**, 062201 (2005), <https://doi.org/10.1063/1.1904565>.
- [27] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [28] T. E. Baker, E. M. Stoudenmire, L. O. Wagner, K. Burke, and S. R. White, One-dimensional mimicking of electronic structure: The case for exponentials, *Phys. Rev. B* **91**, 235141 (2015).
- [29] S. Elfving, E. Uchibe, and K. Doya, Sigmoid-weighted linear units for neural network function approximation in reinforcement learning, *Neural Networks* **107**, 3 (2018), special issue on deep reinforcement learning.
- [30] P. Ramachandran, B. Zoph, and Q. V. Le, Searching for activation functions (2018).

-
- [31] S. Vuckovic, S. Song, J. Kozłowski, E. Sim, and K. Burke, Density functional analysis: The theory of density-corrected dft, *Journal of Chemical Theory and Computation* **15**, 6636 (2019), pMID: 31682433.
- [32] Y. Kim, S. Song, E. Sim, and K. Burke, Halogen and chalcogen binding dominated by density-driven errors, *The Journal of Physical Chemistry Letters* **10**, 295 (2019).
- [33] L. O. Wagner, E. Stoudenmire, K. Burke, and S. R. White, Reference electronic structure calculations in one dimension, *Physical Chemistry Chemical Physics* **14**, 8581 (2012).