

Emergence of Accurate Atomic Energies from Machine Learned Noble Gas Potentials

Frank Uhlig,^{*} Samuel Tovey^{*}, and Christian Holm[†]

Institute for Computational Physics, University of Stuttgart, Stuttgart, Germany

(Dated: March 4, 2024)

The quantum theory of atoms in molecules (QTAIM) gives access to well-defined local atomic energies. Due to their locality, these energies are potentially interesting in fitting atomistic machine learning models as they inform about physically relevant properties. However, computationally, quantum-mechanically accurate local energies are notoriously difficult to obtain for large systems. Here, we show that by employing semi-empirical correlations between different components of the total energy, we can obtain well-defined local energies at a moderate cost. We employ this methodology to investigate energetics in noble liquids or argon, krypton, and their mixture. Instead of using these local energies to fit atomistic models, we show how well these local energies are reproduced by machine-learned models trained on the total energies. The results of our investigation suggest that smaller neural networks, trained only on the total energy of an atomistic system, are more likely to reproduce the underlying local energy partitioning faithfully than larger networks. Furthermore, we demonstrate that networks more capable of this energy decomposition are, in turn, capable of transferring to previously unseen systems. Our results are a step towards understanding how much physics can be learned by neural networks and where this can be applied, particularly how a better understanding of physics aids in the transferability of these neural networks.

Introduction Machine learning (ML) has impacted many aspects of modern society. From artificial intelligence, smart homes, and social media to an industrial context, e.g., predictive maintenance or medical diagnosis, ML methods have changed how we interact with computers and our environment. In recent years, ML methods have been exploited in the development of interatomic potentials for liquids [1, 2], solids [3–5] and interfaces [6, 7] with almost ab initio accuracy but at a highly reduced cost for running associated molecular dynamics (MD) simulations. Many reviews have been dedicated to ML and its various applications to potential energy surfaces [8], computational/theoretical chemistry [9, 10], and cheminformatics [11, 12]. ML can interpolate between training data without knowledge of the differential equations describing the underlying physics, thus offering considerable flexibility at the cost of interpretability in the final model. The most common approaches to developing machine-learned potentials include directly fitting the total energy of a system [8], splitting the total energy into local atomic energies [13, 14], or using ML to produce parameters for physically motivated potentials from which the energy of a system can be calculated [15, 16]. Splitting the total energy into a sum of local atomic energies has important features: it offers accuracy, flexibility, and scalability if models are fitted only for the involved atomic species. This is achieved by representing atomic environments centred on individual atom sites as feature vectors. Different variations of these feature vectors, commonly referred to as *descriptors*, have been developed with varying success and accuracy [13, 14, 17–26]. However, it is unclear whether assuming that a system’s total energy can be split into a simple sum of atomic energies

is valid. It can be argued that, due to the locality of quantum mechanics, the energy of an atom is mainly dependent on its local environment [21]. Indeed, the *quantum theory of atoms in molecules* (QTAIM) allows us to dissect the total quantum-mechanical energy of any system into well-defined atomic energies at stationary states of the atomic configuration [27, 28]. Further development has enabled the calculation of atomic energies for non-stationary states by directly calculating them from density matrices of the system. [29] This has been exploited to train machine learning models to predict these local atomic energies. [30, 31] One can then accurately predict atomic energies and effective atomic charges [15] that can be used to describe the interaction of atoms and molecules. One of the most challenging aspects of machine learning is to provide transferable models that can be trained on a small set of reference data and applied to a more extensive set with comparable accuracy [32]. In the field of computer vision, some transferability is assumed in the context of transfer learning, i.e., refining pre-trained models for a specific task. Similarly, models for mixed component systems have been systematically fitted on the pure components first and only then refined for the full systems [33]. Recent efforts regarding the prediction electron densities show that transferability can be achieved by using atom-centered symmetry functions as descriptors [34]. To achieve transferability and scalability, it is crucial that the employed descriptor can suitably accommodate the inclusion of different atomic species without growing in a computationally intractable manner. Models trained on large subsets of with vast chemical diversity offer good performance for various tasks, from predicting energetics to identifying chemical concepts. [35] Current developments focus on generating transferable descriptors that offer a reduction in scaling with respect to the number of involved chemical species, and thus, training data [16, 36–40]. Another

^{*} Contributed equally

[†] holm@icp.uni-stuttgart.de

approach to enforce transferability is to reproduce the underlying physics more accurately. This could mean ensuring that the predicted local energies correspond to a physical theory in the above context. Here, to the best of our knowledge, for the first time, we report a direct comparison between physically meaningful local energies and energies obtained with different machine-learning algorithms for bulk liquids. This work investigates under which conditions machine learning models reproduce the underlying energetics and how much data is needed to learn the corresponding physics. We argue that smaller networks are forced to learn more fundamental representations of the underlying data they fit, sometimes leading to the emergence of learned physics.

Theory The quantum mechanical virial theorem establishes a connection between the Laplacian of the electron density $\rho(\mathbf{r})$, the kinetic energy density $\mathcal{T}(\mathbf{r})$, and the total virial $\mathcal{V}(\mathbf{r})$ [27]:

$$\left(\frac{\hbar^2}{4m}\right) \nabla^2 \rho(\mathbf{r}) = 2\mathcal{T}(\mathbf{r}) + \mathcal{V}(\mathbf{r}). \quad (1)$$

If integrated over the whole system, e.g., an isolated molecule, the left-hand side of eq. 1 disappears and relates the kinetic energy to the total virial. The same holds true for zero-flux surfaces $S(\mathbf{r})$ of the electron density, defined by means of Gauss' divergence theorem using the gradient of the electron density and the normal vector, $\mathbf{n}(\mathbf{r})$ as

$$\nabla \rho(\mathbf{r}) \cdot \mathbf{n}(\mathbf{r}) = 0 \quad (2)$$

which, within the framework of QTAIM, determines atomic regions, Ω . Hence, for QTAIM, we can also define a relationship between the local virial and the kinetic energy. At a stationary state with respect to the nuclei, the virial becomes the local potential energy of the i^{th} atom:

$$-2\mathcal{T}(\Omega_i) = \mathcal{V}(\Omega_i) \stackrel{\text{stat.}}{=} \stackrel{\text{state}}{=} V(\Omega_i) \quad (3)$$

and the total energy, E can be obtained as a sum of local energies of atoms in molecules by:

$$E = \sum_i V(\Omega_i). \quad (4)$$

With eq. 4, or modified versions of it [29], one has access to well-defined local atomic energies that can be used to fit machine-learned models. [30, 31] A disadvantage of eq. 4 is that the connection between potential and kinetic energy only exists as such for stationary states. Another caveat is that the kinetic energy density is difficult to obtain at an adequate resolution for large-scale analysis. Due to the involved gradient of the wavefunction, very fine grids must be used, which quickly become impractical to handle. However, one might expect that certain correlations between energy components still exist, even for structures sampled at finite temperature. If so, it would unlock new methods for validating and even training machine learned interatomic potentials.

In typical machine-learning approaches for the fitting of machine-learned inter-atomic potentials, a model is fit to reproduce the same total total energy of the system by first computing the local energies of so-called atomic environments, or descriptors, \vec{g} , and then summing over them by

$$E = \sum_i f_\theta(\vec{g}_i), \quad (5)$$

where f_θ is a neural network parameterised by the set θ . The parameters used in constructing these descriptors will define the environment in which the local energy is computed. While many descriptors exist, this work uses the invariant 4D-Bispectrum descriptor in all model training [20]. These descriptors are fixed, meaning that they do not evolve during training, and they include a species scaling to differentiate different atom types. The network is updated via gradient descent using a loss function of the form

$$\mathcal{L} = \left(E - \sum_i f_\theta(\vec{g}_i)\right)^2. \quad (6)$$

Under this construction, there is much ambiguity about how atomic energies are assigned to specific atoms, i.e., which partitioning scheme is learned by the network. We argue that the correct partitioning of energies into their atomic contributions will result in better models as the underlying physics is better represented. This can be considered a feature learning problem, as the networks must learn how to correctly represent specific atomic environments in the latent space to be transferable to unseen local environments, something only possible if the partitioning is done physically.

Feature learning, or representation learning as it is often referred to, describes the ability of a neural network to learn a fundamental representation of a data manifold rather than simply fitting to training points [41]. For neural networks, this means that features in the data, for example certain shapes in the case of images, have been learned during propagation through layers. In the context of fitting machine-learned inter-atomic potentials, feature learning can refer to both the learning of a descriptor internally within the neural network or, as is discussed here, learning the most physical distribution of local energies, that is, the distribution closest to the physically motivated QTAIM method. It has been shown that feature learning is fundamental to the success of both pre-training and transfer learning paradigms [42, 43]. However, not all neural networks are capable of feature learning. In their 2021 report, *Yang and Hu* demonstrated that large neural networks' ability to perform feature learning strongly depends on their initialisation. Often, they are incapable of feature learning, becoming purely regressive machines; that is to say, they learn purely the data they are trained on and not features connecting the points [44].

It is not difficult to argue that learning a fundamental physical principle would amount to feature learning for a neural network or, to say it differently, a neural network with knowledge of physics has learned the correct features to describe a problem. This point is often leveraged in developing physics-informed networks, which have been shown to lead to better transferrability [45–47]. However, it is not well studied whether a neural network, constrained in its size and, therefore, in the complexity of features that can be learned, will, in turn, develop knowledge of physics. This work examines this question by determining under what conditions a neural network can learn the correct energy decomposition in an atomistic system.

Local energy computation Before identifying the role of architecture and the emergence of correct energy partitioning, we demonstrate that it is possible to apply the QTAIM method to the extraction of local energies in the non-stationary states produced in our noble-liquid DFT simulations. To do so, a correlation must be identified between the classical electrostatic contribution to the system’s total energy and the other Hamiltonian components. Such a correlation would allow us to extract local energy information on much larger grids, thereby making the computation accessible to large, periodic systems. To investigate if this correlation exists, we performed ab initio molecular dynamics simulations (AIMD) based on density functional theory (DFT) for three different bulk liquids, Ar(108), Kr(108), and Kr(54)Ar(54) at 120 K. AIMD simulations used Kohn-Sham density functional theory (KS-DFT) and were carried out with CP2K [48, 49]. The exchange-correlation energy was approximated using a van-der-Waals density functional [50, 51]. The Kohn-Sham one-particle wavefunctions were expanded into an atom-centered double- ζ basis set optimized for dense liquids[52]. The electron density was expanded in an auxiliary plane-wave basis set with a kinetic energy cutoff of 600 Ry. Initial configurations were obtained from molecular dynamics simulations using suitable noble liquid potentials. Newton’s equations of motion were integrated using a timestep of 2 fs, and the temperature of the system was kept constant at 120 K using a Nosé-Hoover chain thermostat [53] (chain length: 3, time constant: 100 fs). This work focuses on noble liquid mixtures as a proxy for more complicated liquids, specifically liquid argon, krypton and their equimolar mixture. Figure 1 outlines the results of this study for the liquid argon (see Figures 5 and 6 in SI for the mixture and pure Kr respectively.). In these systems, we observe a strong correlation between the classical electrostatic energy and the sum of all other contributions to the total energy, i.e., kinetic energy and exchange-correlation energy. The correlation between these variables can be fit linearly, and while the values for slope and y-intercept are system-specific, they can be attributed to physical interpretations. For the pure systems, the intercepts can be related to the classical electrostatic energy of the isolated atoms in vacuo. This is not a general characteristic

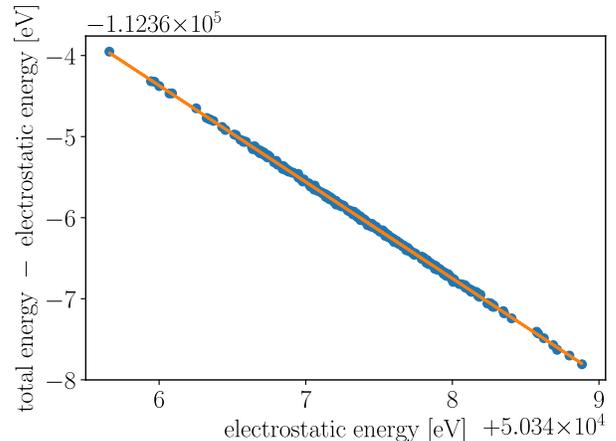


FIG. 1: Total energy minus classical electrostatic energy versus the classical electrostatic energy in pure Ar.

as it does not apply to the mixed system. Nevertheless, this study demonstrates a way to relate the electrostatic energy effectively to the remaining energy components. The fit’s root mean squared error is approximately 0.3 meV/atom, the same order as a typically achieved fit accuracy in machine learned potentials. We can largely ignore this source of uncertainty as the error in the machine learning model fit will outweigh it. Using these results, we can dissect the classical electrostatic interaction between all charged particles (electron density and protons) as it is readily available and accurately reproduced on grids of moderate resolution and can therefore be used for large, 3D-periodic systems. The dissection follows in accordance with eq. 3 and is considered a robust numerical procedure[54]. The electrostatic energy of an atom E_{Ω} is calculated using the electron density ρ and the Hartree potential V_H as $E_{\Omega} = \frac{1}{2} \int_{\Omega} \rho V_H$. We will refer to these energies as local DFT energies, E_{loc}^{DFT} .

Local energy fitting With the ability to produce reference data for the local energies, we can turn our attention to understanding under which conditions neural networks, trained purely on the total energy of ab initio simulations, can accurately reproduce the atomic energies computed using QTAIM. To do so, neural networks of various architectures were trained on pure argon and krypton systems and the ArKr mixture using the total energy of each configuration in the loss function. Due to the similarity of the results, only the results from a single architecture are shown in the manuscript, namely, a single layer network with hyperbolic tangent activation. Other architectures are displayed in the appendix. Neural networks utilised hyperbolic-tangent and linear activation functions and were trained using a mean-squared-error loss function. In all cases, the correlation score for the total energy, the quantity on which the network was trained, remains high across the architecture space (see appendix Figure 7). For small models, the initialisation

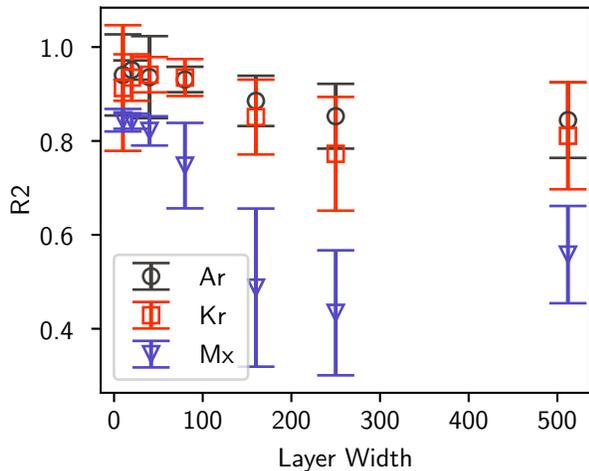


FIG. 2: Correlation score for the local energy of each system, the argon (Ar), krypton (Kr), and the 50-50 mixture (Mx).

variance results in large error bars, which are saturated for the wider and deeper models.

As all model architecture are capable of fitting the total energies, the interest now lies in seeing how the change in architecture will impact the ability of the models to reproduce the local QTAIM energies of the systems correctly. To do so, Pearson correlation coefficients are computed on the local energy test data produced from the DFT simulations. This is data that the neural networks have never seen during their training Figure 2 displays these correlation coefficients between the predicted and DFT-produced local energies for the pure and mixed systems (see appendix Figure 8 for all architectures). Studying these local correlation plots, it is clear that larger neural networks assign the local energies more arbitrarily and show little correlation with the underlying physics. On the other hand, smaller networks maintain a reasonable correlation between the predicted and actual partitioning of the local energy values. It is clear, however, that the models trained on the mixtures struggle more to achieve this local energy decomposition, with these networks achieving lower maximum scores and dropping off faster than their pure system counterparts.

Transferability Given the emergence of correct local energies, it is of interest to see if this emergent structure learning can be used in transfer studies. We investigate this by testing if models more capable of local energy decomposition are also capable of being transferred, without re-training, to previously unseen configurations. The models trained on the ArKr mixture are used to compute the total and local energies of the pure Ar and Kr systems whereas the models trained on the pure systems are used to compute the energies in the mixed systems. This is done on the same test data as in the previous section using the `pair style` command in the LAMMPS simu-

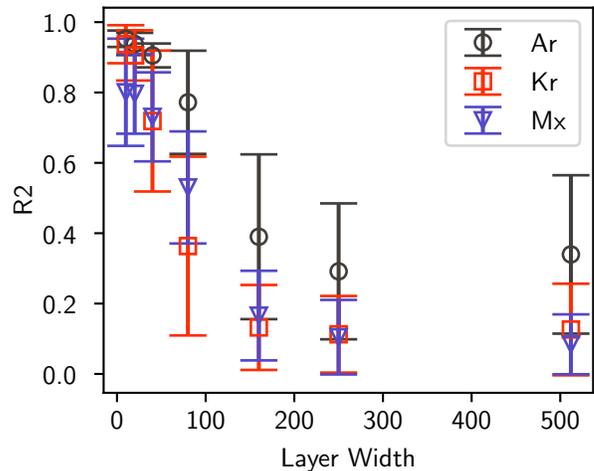


FIG. 3: Correlation score for the total energy of each system computed using the transferred models, the argon (Ar), krypton (Kr), and the 50-50 mixture (Mx).

lation engine where atomic environments centered on an Ar atom will use the pure Ar model for the force computation and the pure Kr models handle the Kr centered environments. Figure 3 outlines the results of this study on the total energy computations for all architectures and systems. As is to be expected, the models fail in their predictions early in all cases. It is, however, interesting that the smaller models can retain a high correlation score for the transferred systems. This suggests that these models represent the underlying physics of the problem better. Another point of note is the ability of the Tanh activation function models to achieve more significant correlation scores for slightly longer than the linear models. Moving away from the total energies, Figure 4 shows the results of the transfer experiment on the local energy predictions. These plots closely resemble those of the total energies in the gradual degradation of performance with increasing network width. In typical machine learning theory, the networks addressed here would not be considered purely regressive as they are not approaching an infinite width limit. However, even amongst the feature learning limit, there appears to be a discrepancy between architectures. We argue that this comes down to the number of degrees of freedom each of these neural networks has and how they can use these to distribute information. Namely, the networks that are required to find solutions with less degrees of freedom typically learn more physical representations. Such an outcome is supported by the data and resembles the principle of Occam’s razor: the simplest solution is often the correct one.

Conclusion The decomposition of total energy into atomic contributions has been significantly utilised in machine-learned inter-atomic potentials. However, whether these local energies can be considered physically meaningful requires clarification. In this investigation,

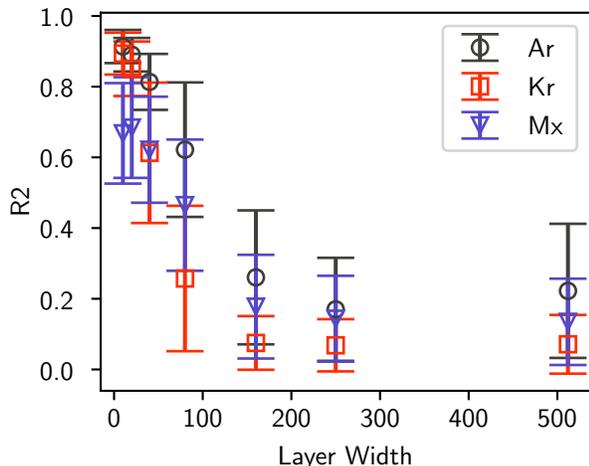


FIG. 4: Correlation score for the local energy of each system computed using the transferred models, the argon (Ar), krypton (Kr), and the 50-50 mixture (Mx).

we have restructured this problem as one of physics-aware feature learning in neural networks. By considering local energies obtained from the QTAIM as ground truth, we have argued that correct feature learning of a neural network involves learning these physically derived local energies given the system’s total energy. To this end, studies of pure liquid argon and krypton and their mixture have been performed wherein DFT-MD was performed to generate ab initio data. We have shown that QTAIM can be applied to bulk liquid systems and computed these local energies. Neural network models were then fitted to the total energies of the ab initio simulations for different architectures before the local decomposition was compared with theoretical values. We find that neural networks of smaller width (12-100 nodes) display improved decomposition over larger models. These results align with

those of feature learning in the field of learning theory, wherein it is understood that smaller neural networks are capable of so-called feature learning, i.e., extraction of features in data, while larger models are purely regressive. We further highlight that even among feature learning networks, those with fewer degrees of freedom learn more physically relevant representations, thus resulting in more physically correct energies. We conclude that smaller networks are faster to train and more easily deployed whilst retaining the accuracy of larger networks. They also learn a more physical representation of atomic decomposition and are thus favorable from a physics perspective. Finally, we have shown that the models capable of computing accurate local energies also showed the ability to accurately transfer to previously unseen configurations. While it is expected that this transferability is predominantly a geometric one, future work should identify the limits to which this can be pushed in chemical space.

ACKNOWLEDGMENTS

The authors acknowledge financial support from the German Funding Agency (Deutsche Forschungsgemeinschaft DFG) under Germany’s Excellence Strategy EXC 2075-390740016. This work was supported by SPP 2363-”Utilization and Development of Machine Learning for Molecular Applications – Molecular Machine Learning.” Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), Project-No 497249646. S.T would like to thank David Tovey for his thorough review of and comments on the manuscript. The authors would like to thank Anand Narayanan Krishnamoorthy and Matthias Bauer for their helpful input on early versions of this study involving Gaussian process-based models. All authors would like to acknowledge and thank David Beyer for his detailed reading and editing of the manuscript.

-
- [1] Samuel Tovey, Anand Narayanan Krishnamoorthy, Ganesh Sivaraman, Jicheng Guo, Chris Benmore, Andreas Heuer, and Christian Holm. Dft accurate interatomic potential for molten nacl from machine learning. *J Phys Chem C*, 124:25760–25768, 2020.
- [2] János Daru, Harald Forbert, Jörg Behler, and Dominik Marx. Coupled cluster molecular dynamics of condensed phase systems enabled by machine learning potentials: Liquid water benchmark. *Phys. Rev. Lett.*, 129:226001, Nov 2022.
- [3] Ganesh Sivaraman, Anand Narayanan Krishnamoorthy, Matthias Baur, Christian Holm, Gabor Csányi Marius Stan, Chris Benmore, and Álvaro Vázquez-Mayagoitia. Machine-learned interatomic potentials by active learning: amorphous and liquid hafnium dioxide. *npj Computational Materials*, 6:1–8, 2020.
- [4] S.S. Shishvan, G. Csányi, and V.S. Deshpande. Strain rate sensitivity of the hydrogen embrittlement of ferritic steels. *Acta Materialia*, 257:119173, 2023.
- [5] Xia Liang, Johan Klarbring, William J. Baldwin, Zhenzhu Li, Gábor Csányi, and Aron Walsh. Structural dynamics descriptors for metal halide perovskites. *The Journal of Physical Chemistry C*, 127(38):19141–19151, 2023.
- [6] Christoph Schran, Fabian L. Thiemann <https://>, Patrick Rowe, and Angelos Michaelides. Machine learning potentials for complex aqueous systems made simple. *PNAS*, 118, September 2021.
- [7] Molpeceres, G., Zaverkin, V., Furuya, K., Aikawa, Y., and Kästner, J. Reaction dynamics on amorphous solid water surfaces using interatomic machine-learned potentials - microscopic energy partition revealed from the p + h → ph reaction. *A&A*, 673:A51, 2023.

- [8] Chris M. Handley and Paul L. A. Popelier. Potential energy surfaces fitted by artificial neural networks. *The Journal of Physical Chemistry A*, 114(10):3371–3383, 2010. PMID: 20131763.
- [9] Garrett B. Goh, Nathan O. Hodas, and Abhinav Vishnu. Deep learning for computational chemistry. *Journal of Computational Chemistry*, 38(16):1291–1307, 2017.
- [10] Matthias Rupp, O. Anatole von Lilienfeld, and Kieron Burke. Guest editorial: Special topic on data-enabled theoretical chemistry. *The Journal of Chemical Physics*, 148(24):241401, 2018.
- [11] John B. O. Mitchell. Machine learning methods in chemoinformatics. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 4(5):468–481, 2014.
- [12] Yu-Chen Lo, Stefano E. Rensi, Wen Torng, and Russ B. Altman. Machine learning in chemoinformatics and drug discovery. *Drug Discovery Today*, 23(8):1538 – 1546, 2018.
- [13] Jörg Behler and Michele Parrinello. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.*, 98:146401, Apr 2007.
- [14] Albert P. Bartók, Mike C. Payne, Risi Kondor, and Gábor Csányi. Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.*, 104:136403, Apr 2010.
- [15] Timothy L. Fletcher and Paul L. A. Popelier. Fflux: Transferability of polarizable machine-learned electrostatics in peptide chains. *Journal of Computational Chemistry*, 38(13):1005–1014, 2017.
- [16] Tristan Bereau, Robert A. DiStasio, Alexandre Tkatchenko, and O. Anatole von Lilienfeld. Non-covalent interactions across organic and biological subsets of chemical space: Physics-based potentials parametrized from machine learning. *The Journal of Chemical Physics*, 148(24):241706, 2018.
- [17] Jörg Behler. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *The Journal of Chemical Physics*, 134(7):074106, 2011.
- [18] Nongnuch Artrith, Tobias Morawietz, and Jörg Behler. High-dimensional neural-network potentials for multi-component systems: Applications to zinc oxide. *Phys. Rev. B*, 83:153101, Apr 2011.
- [19] Matthias Rupp, Alexandre Tkatchenko, Klaus-Robert Müller, and O. Anatole von Lilienfeld. Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.*, 108:058301, Jan 2012.
- [20] Albert P. Bartók, Risi Kondor, and Gábor Csányi. On representing chemical environments. *Physical Review B*, 87:184115, 2013.
- [21] A.P. Thompson, L.P. Swiler, C.R. Trott, S.M. Foiles, and G.J. Tucker. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *Journal of Computational Physics*, 285:316 – 330, 2015.
- [22] Katja Hansen, Franziska Biegler, Raghunathan Ramakrishnan, Wiktor Pronobis, O. Anatole von Lilienfeld, Klaus-Robert Müller, and Alexandre Tkatchenko. Machine learning predictions of molecular properties: Accurate many-body potentials and nonlocality in chemical space. *The Journal of Physical Chemistry Letters*, 6(12):2326–2331, 2015. PMID: 26113956.
- [23] O. Anatole von Lilienfeld, Raghunathan Ramakrishnan, Matthias Rupp, and Aaron Knoll. Fourier series of atomic radial distribution functions: A molecular fingerprint for machine learning models of quantum chemical properties. *International Journal of Quantum Chemistry*, 115(16):1084–1093, 2015.
- [24] Jan Jenke, Aparna P. A. Subramanyam, Marc Densow, Thomas Hammerschmidt, David G. Pettifor, and Ralf Drautz. Electronic structure based descriptor for characterizing local atomic environments. *Phys. Rev. B*, 98:144102, Oct 2018.
- [25] Behnam Parsaeifard and Stefan Goedecker. Manifolds of quasi-constant soap and acsf fingerprints and the resulting failure to machine learn four-body interactions. *Journal of Chemical Physics*, 156:034302, 2022.
- [26] V. Zaverkin and J. Kästner. Gaussian moments as physically inspired molecular descriptors for accurate and scalable machine learning potentials. *Journal of Chemical Theory and Computation*, 16(8):5410–5421, 2020. PMID: 32672968.
- [27] R. F. W. Bader. Atoms in molecules. *Accounts of Chemical Research*, 18(1):9–15, 1985.
- [28] R.F.W. Bader. *Atoms in Molecules: A Quantum Theory*. International Ser. of Monogr. on Chem. Clarendon Press, 1994.
- [29] P. L. A. Popelier and D. S. Kosov. Atom–atom partitioning of intramolecular and intermolecular coulomb energy. *The Journal of Chemical Physics*, 114(15):6539–6547, 2001.
- [30] Paul L. A. Popelier. Qctff: On the construction of a novel protein force field. *International Journal of Quantum Chemistry*, 115(16):1005–1011, 2015.
- [31] Paul L A Popelier. Molecular simulation by knowledgeable quantum atoms. *Physica Scripta*, 91(3):033007, 2016.
- [32] Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M. Elena, Dávid P. Kovács, Janosh Riebesell, Xavier R. Advincula, Mark Asta, William J. Baldwin, Noam Bernstein, Arghya Bhowmik, Samuel M. Blau, Vlad Cărare, James P. Darby, Sandip De, Flaviano Della Pia, Volker L. Deringer, Rokas Elijošius, Zakariya El-Machachi, Edwin Fako, Andrea C. Ferrari, Annalena Grenreith-Schriever, Janine George, Rhys E. A. Goodall, Clare P. Grey, Shuang Han, Will Handley, Hendrik H. Heenen, Kersti Hermansson, Christian Holm, Jad Jaafar, Stephan Hofmann, Konstantin S. Jakob, Hyunwook Jung, Venkat Kapil, Aaron D. Kaplan, Nima Karimipari, Namu Kroupa, Jolla Kullgren, Matthew C. Kuner, Domantas Kuryla, Guoda Liepuoniute, Johannes T. Margraf, Ioan-Bogdan Magdău, Angelos Michaelides, J. Harry Moore, Aakash A. Naik, Samuel P. Niblett, Sam Walton Norwood, Niamh O’Neill, Christoph Ortner, Kristin A. Persson, Karsten Reuter, Andrew S. Rosen, Lars L. Schaaf, Christoph Schran, Eric Sivonxay, Tamás K. Stenczel, Viktor Svahn, Christopher Sutton, Cas van der Oord, Eszter Varga-Umbrich, Tejs Vegge, Martin Vondrák, Yangshuai Wang, William C. Witt, Fabian Zills, and Gábor Csányi. A foundation model for atomistic materials chemistry, 2023.
- [33] Samad Hajinazar, Junping Shao, and Aleksey N. Kolmogorov. Stratified construction of neural network based interatomic models for multicomponent materials. *Phys. Rev. B*, 95:014114, Jan 2017.
- [34] Clemence Corminboeuf, Michele Certiotti, benjamin meyer, Alberto Fabrizio, Andrea Grisafi, and David M. Wilkins. A transferable machine-learning model of the electron density, Sep 2018.

- [35] Kristof T. Schütt, Farhad Arbabzadah, Stefan Chmiela, Klaus R. Müller, and Alexandre Tkatchenko. Quantum-chemical insights from deep tensor neural networks. *Nature Communications*, 8:13890 EP –, Jan 2017. Article.
- [36] M. Gastegger, L. Schwiedrzik, M. Bittermann, F. Berzsenyi, and P. Marquetand. wacsf—weighted atom-centered symmetry functions as descriptors in machine learning potentials. *The Journal of Chemical Physics*, 148(24):241709, 2018.
- [37] Nongnuch Artrith, Alexander Urban, and Gerbrand Ceder. Efficient and accurate machine-learning interpolation of atomic energies in compositions with many species. *Phys. Rev. B*, 96:014112, Jul 2017.
- [38] Oliver T. Unke and Markus Meuwly. A reactive, scalable, and transferable model for molecular energies from a neural network approach based on local information. *The Journal of Chemical Physics*, 148(24):241708, 2018.
- [39] S. Jindal and S. S. Bulusu. A transferable artificial neural network model for atomic forces in nanoparticles. *ArXiv e-prints*, October 2018.
- [40] Ryo Tamura, Jianbo Lin, and Tsuyoshi Miyazaki. Machine learning forces trained by gaussian process in liquid states: Transferability to temperature and pressure. *arXiv preprint arXiv:1810.05586*, 2018.
- [41] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- [42] Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research*, 11(19):625–660, 2010.
- [43] Oscar Day and Taghi M. Khoshgoftaar. A survey on heterogeneous transfer learning. *Journal of Big Data*, 4(1):29, Sep 2017.
- [44] Greg Yang and Edward J. Hu. Tensor programs iv: Feature learning in infinite-width neural networks. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11727–11737. PMLR, 18–24 Jul 2021.
- [45] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what’s next. *Journal of Scientific Computing*, 92(3):88, Jul 2022.
- [46] Anshul Choudhary, John F. Lindner, Elliott G. Holliday, Scott T. Miller, Sudeshna Sinha, and William L. Ditto. Physics-enhanced neural networks learn order and chaos. *Phys. Rev. E*, 101:062207, Jun 2020.
- [47] Rahul S. Patel, Sharad Bhartiya, and Ravindra D. Gudi. Physics constrained learning in neural network based modeling. *IFAC-PapersOnLine*, 55(7):79–85, 2022. 13th IFAC Symposium on Dynamics and Control of Process Systems, including Biosystems DYCOPS 2022.
- [48] Joost VandeVondele, Matthias Krack, Fawzi Mohamed, Michele Parrinello, Thomas Chassaing, and Jürg Hutter. Quickstep: Fast and accurate density functional calculations using a mixed gaussian and plane waves approach. *Computer Physics Communications*, 167(2):103–128, April 2005.
- [49] Jürg Hutter, Marcella Iannuzzi, Florian Schiffmann, and Joost VandeVondele. CP2K: atomistic simulations of condensed matter systems. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 4(1):15–25, 2014.
- [50] Oleg A. Vydrov and Troy Van Voorhis. Nonlocal van der waals density functional: The simpler the better. *The Journal of Chemical Physics*, 133(24):244103, 2010.
- [51] Riccardo Sabatini, Tommaso Gorni, and Stefano de Gironcoli. Nonlocal van der waals density functional made simple and efficient. *Phys. Rev. B*, 87:041108, Jan 2013.
- [52] Joost VandeVondele and Jürg Hutter. Gaussian basis sets for accurate calculations on molecular systems in gas and condensed phases. *The Journal of Chemical Physics*, 127(11):114105, 2007.
- [53] Glenn J. Martyna, Michael L. Klein, and Mark Tuckerman. Nosé–hoover chains: The canonical ensemble via continuous dynamics. *Journal of Chemical Physics*, 97(4):2635–2643, 1992.
- [54] Min Yu and Dallas R. Trinkle. Accurate and efficient algorithm for bader charge integration. *The Journal of Chemical Physics*, 134(6):064111, 2011.

Appendix A: Appendix A: Energy Decomposition

Here we show the energy decomposition for the pure krypton and argon-krypton mixture systems. These results highlight that using the electrostatic component of the energy to infer local energies is valid across all studied systems.

Appendix B: Appendix B: Architectures

Here we outline the results of local energy prediction investigations on different neural network architectures.

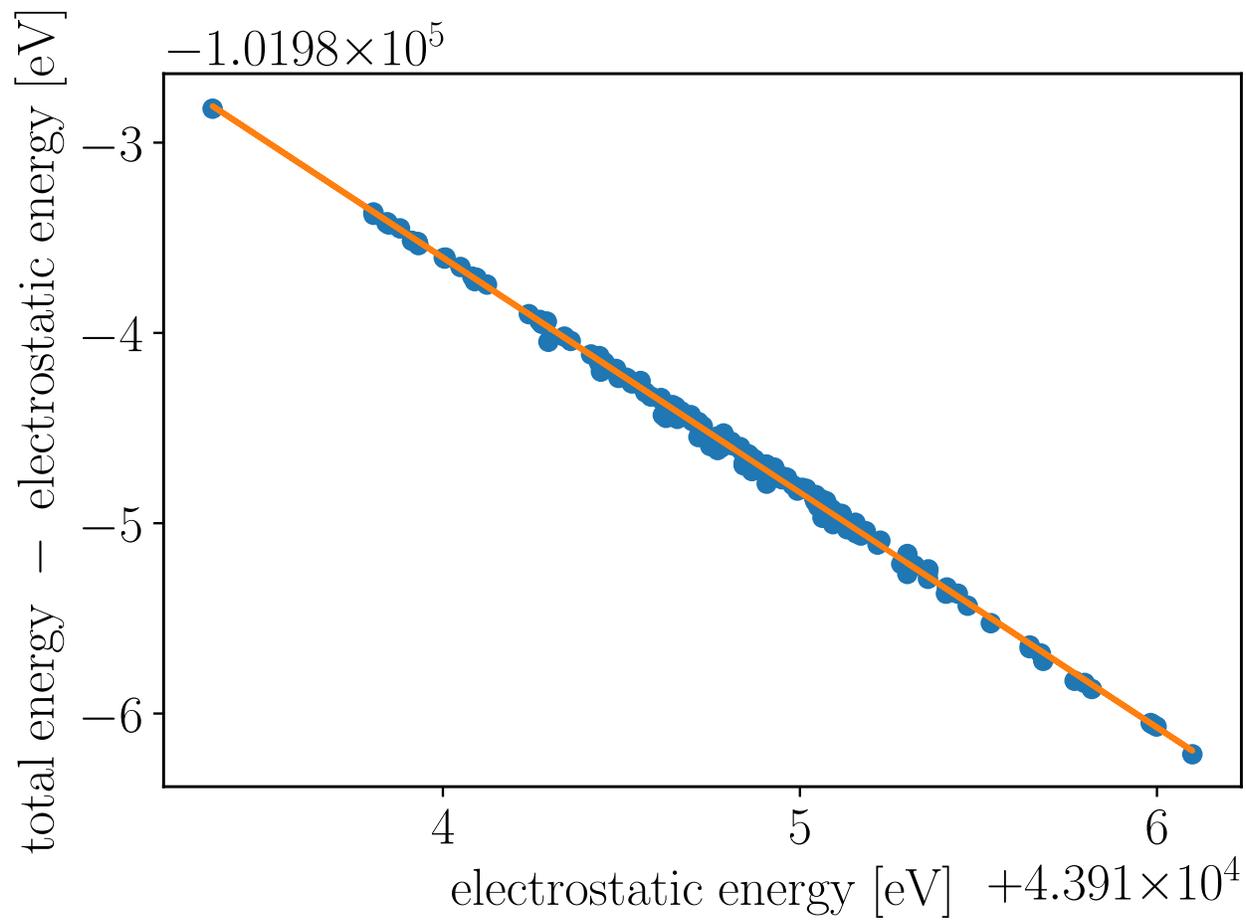


FIG. 5: Total energy minus classical electrostatic energy versus the classical electrostatic energy in the ArKr mixture.

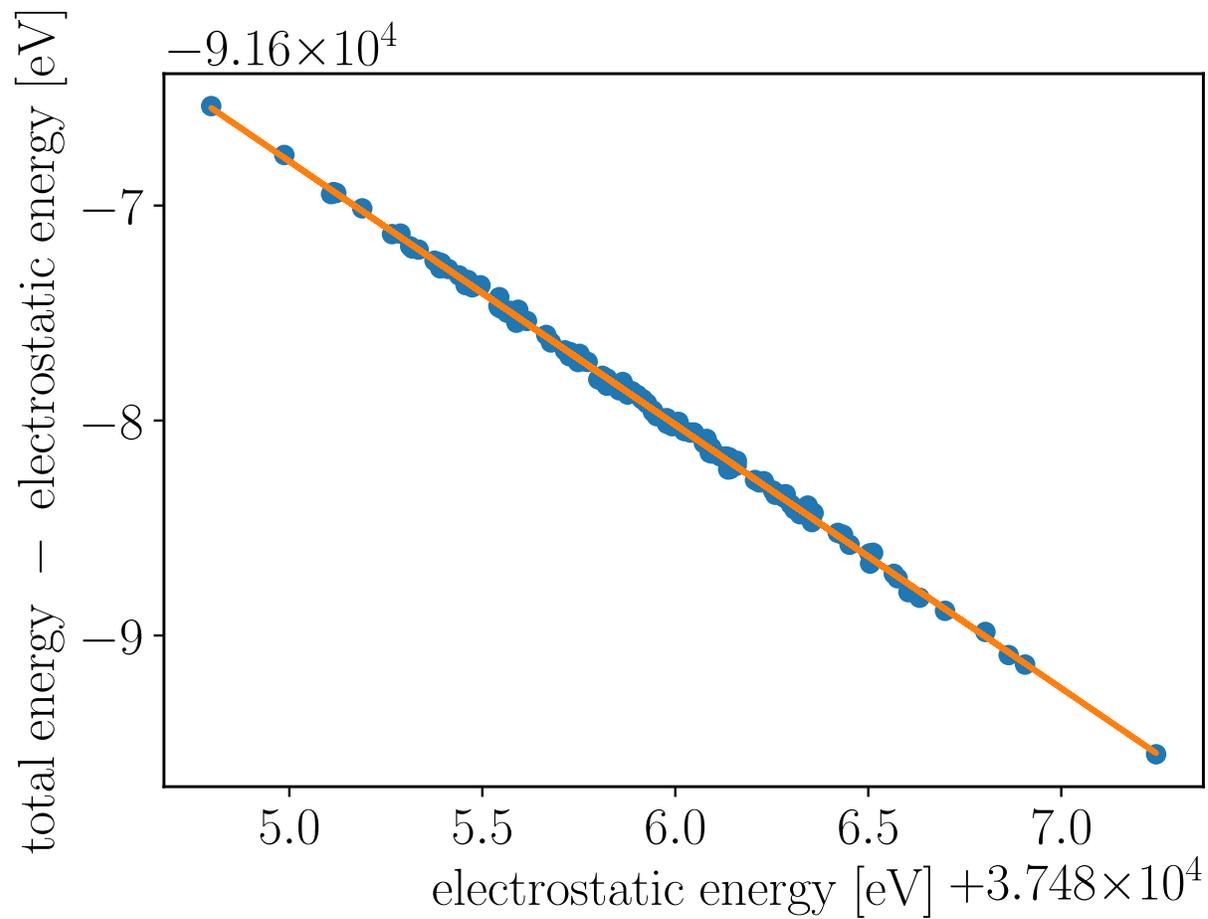


FIG. 6: Total energy minus classical electrostatic energy versus the classical electrostatic energy in the pure Kr system.

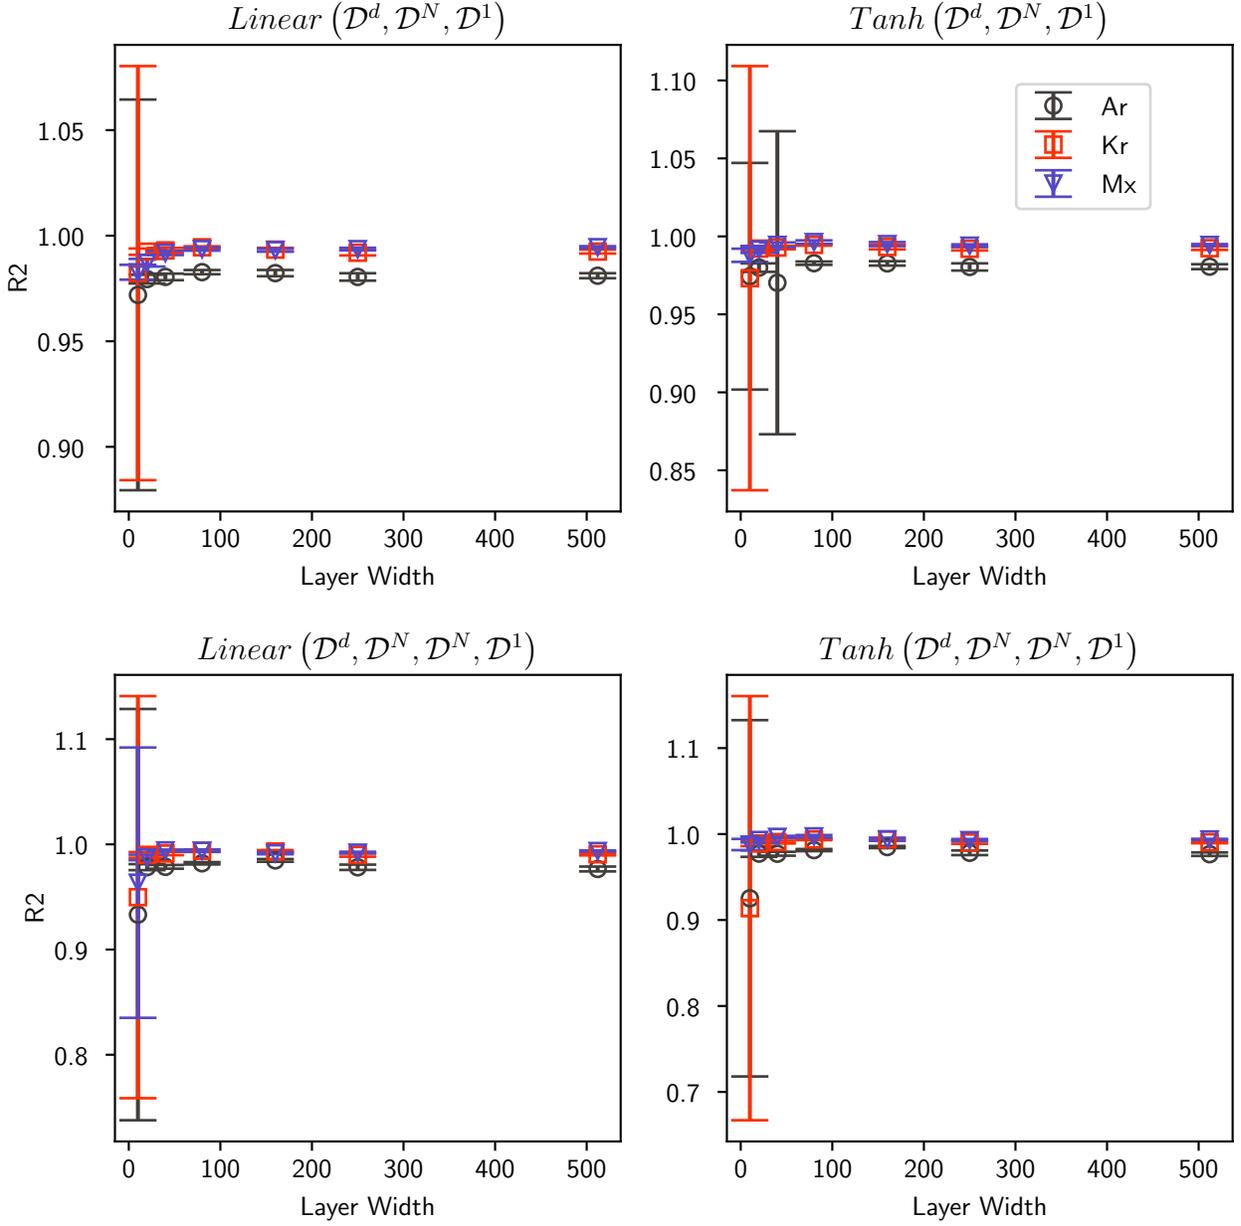


FIG. 7: Correlation score for the total energy of each system. As the networks are trained on the total energy, we expect the score to remain high over the full space of architecture.

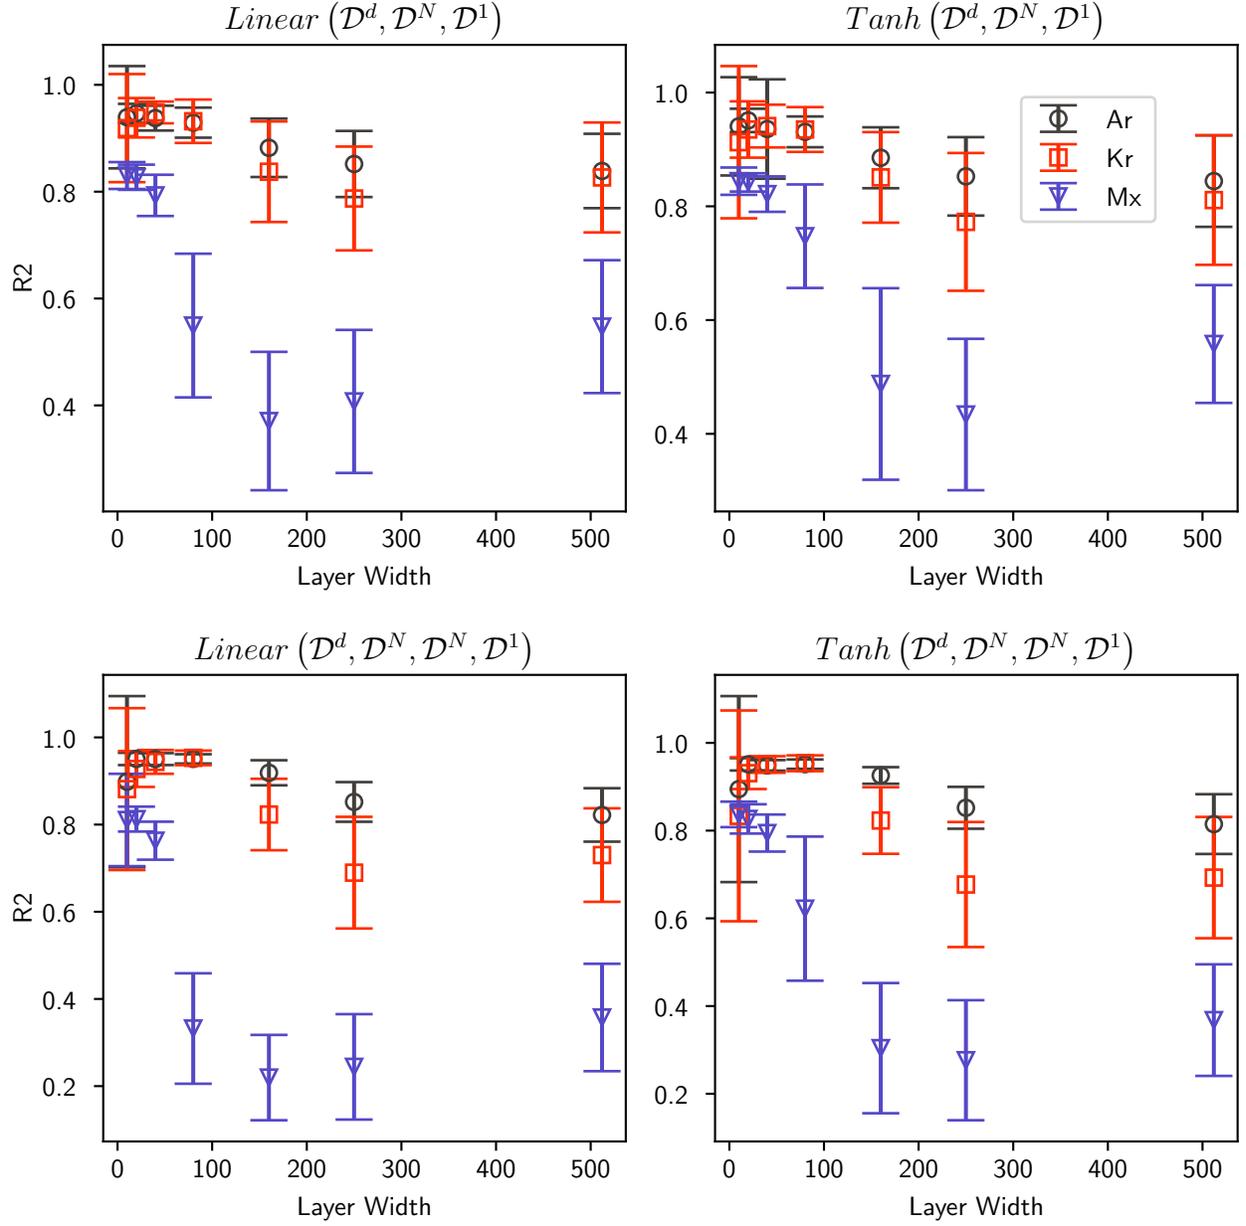


FIG. 8: Correlation score for the local energy of each system.

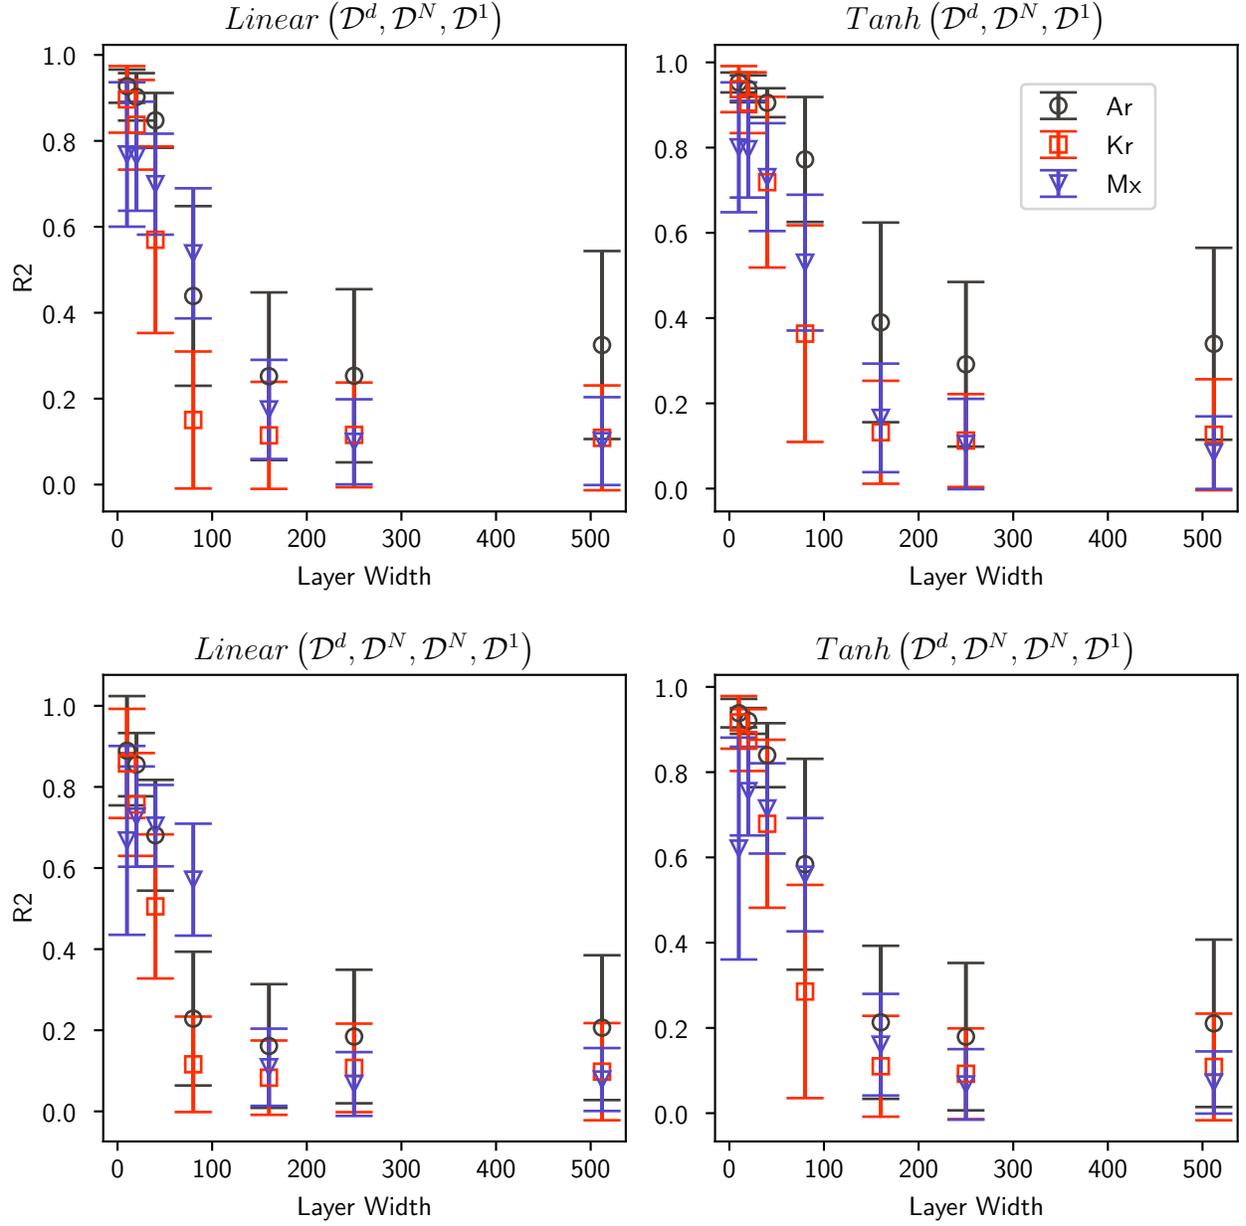


FIG. 9: Correlation score for the total energy of each system computed using the transferred models.

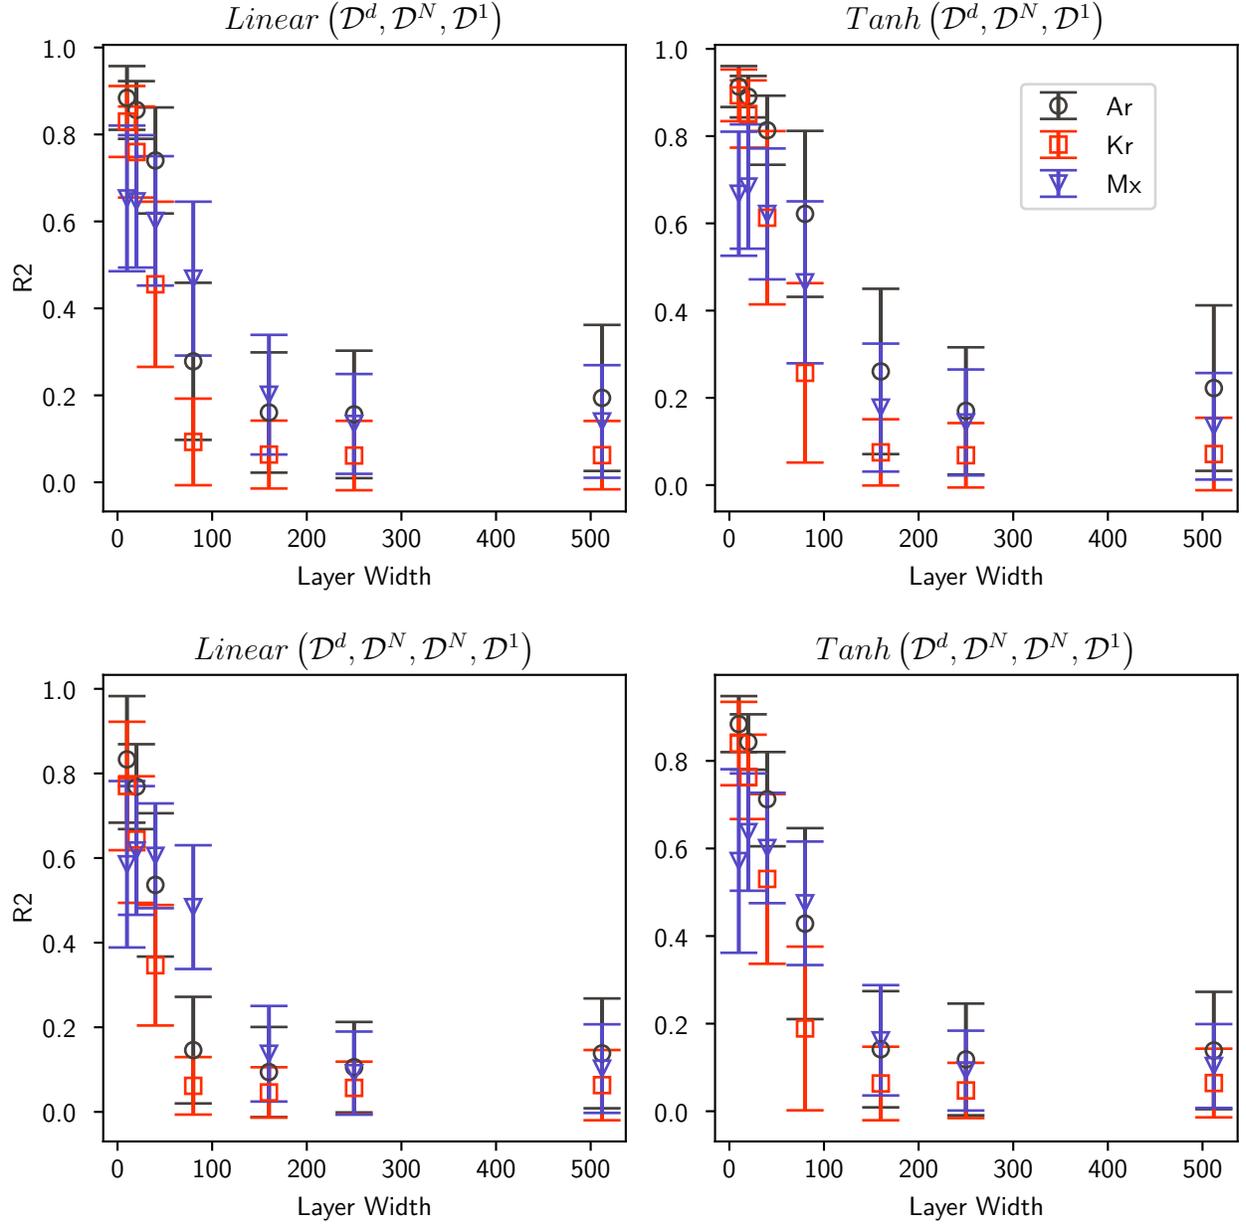


FIG. 10: Correlation score for the local energy of each system computed using the transferred models.