

From Black Boxes to Formulas: Interpretable Kolmogorov-Arnold Networks for HF-3c-to-DFT Descriptor Prediction

Maxim A. Demekhin^{1,*} · Alexander S. Novikov^{1,2}

¹ Infochemistry Scientific Center, ITMO University, Saint Petersburg 191002, Russia

² Institute of Medicine and Life Sciences, Immanuel Kant Baltic Federal University, Kaliningrad 236041, Russia

* Corresponding author: Maxim A. Demekhin, e-mail: demekhin@itmo.ru

Contents

S1 Data and evaluation protocol

S2 Kolmogorov-Arnold networks and training

S3 Symbolic regression and feature attribution

S4 Full test-set metrics against all ten baselines (Table S1)

S5 Seed stability of the selected models

S6 Interactive application (Figs. S1 and S2)

S1 Data and evaluation protocol

We used the data set of Normatov et al. [10] comprising 1031 cyclic supramolecular structures, each characterized by six HF-3c descriptors (Gibbs free energy, electronic energy, entropy, enthalpy, dipole moment, HOMO-LUMO gap) and the six corresponding B3LYP-D4/def2-TZVP descriptors. The inputs to every model in this work are the six HF-3c descriptors only; atomic coordinates were not used. Following the original protocol, the data were split into training and test sets in a 90:10 ratio at a fixed random seed (927 and 104 structures), and features and targets were standardized using statistics estimated on the training set alone. All metrics are computed on predictions returned to the original descriptor units. The data set and the fixed split are redistributed unchanged in the accompanying repository (<https://github.com/MaksimDem/kan-hf-to-dft>).

S2 Kolmogorov-Arnold networks and training

Models were built with the pykan implementation [14, 15]. Each network places learnable univariate functions, represented as cubic B-splines on a bounded grid, on its edges, and was optimized with the L-BFGS algorithm in double precision. For every target the hidden-layer width and spline-grid resolution were selected on a 90:10 validation split of the training set (mean validation RMSE over several seeds). The search favored additive [6, 1] architectures for all targets except entropy, which used one hidden layer; the selected models therefore contain 114-252 parameters. Because spline models are sensitive to initialization, ten seeds were trained per target, and the seed with the lowest validation RMSE was taken as the representative model; its held-out test metrics are reported.

To make the reported models reproducible across machines, the selected architecture and representative seed for each target are pinned in the repository configuration (selected_configs in config.yaml). Training is made deterministic on a given machine by fixing all random seeds, running single-threaded, and enabling deterministic algorithms; the trained checkpoints are provided for exact cross-machine reproduction.

S3 Symbolic regression and feature attribution

To obtain closed-form expressions, networks were trained with a small L1 penalty to encourage sparsity, after which each spline edge was replaced by the best-fitting member of a library of elementary functions (polynomial, exponential, logarithmic, trigonometric and related forms) and the affine parameters were refitted. For the three mutually collinear energy descriptors we additionally fitted single-descriptor [1, 1] networks, which remove the coefficient ambiguity of a multivariate fit and yield unique one-line formulas. For the remaining targets, inputs whose KAN 2.0 attribution score fell below a fixed threshold were pruned before symbolic conversion.

Formulas were extracted in original descriptor units by composing the network with the inverse of the training-set standardization. Feature-attribution scores were computed with the KAN 2.0 attribution method [15] on the training activations and averaged over ten seeds. The grid-refinement analysis of model complexity is presented in the main article (Fig. 4).

S4 Full test-set metrics against all ten baselines

Table S1 lists the held-out test metrics (R^2 , RMSE, MAE, in original units) for the KAN spline model and its symbolic reduction alongside the ten models reported by Normatov et al. [10]. The KAN spline model attains the lowest RMSE for every target. For the energy descriptors, only linear regression among the baselines reaches a comparable order of magnitude. The tree- and boosting-based models, although reporting R^2 near unity, carry RMSE of hundreds of eV.

Table S1 Per-target test metrics for all models (held-out test set, original units). Values for the ten baseline models are taken from ref. [10]. LR, linear regression; DT, decision tree; RF, random forest; SLP and MLP, single- and multi-layer perceptron; GCN, graph convolutional network. Entries of this work are in bold

Gibbs free energy (eV)

Model	R^2	RMSE	MAE
LR	1.0000	3.5806	1.4947
LASSO	0.9996	116.95	53.63
Ridge	0.9980	259.57	190.21
ElasticNet	0.9996	116.58	53.46
DT	0.9833	742.68	501.19
RF	0.9905	560.61	299.61
XGBoost	0.9941	440.23	106.85
SLP	0.9969	320.51	243.12
MLP	0.9970	314.74	210.30
GCN	0.9733	938.40	785.07
KAN, spline (this work)	1.0000	1.1581	0.6747
KAN, symbolic (this work)	1.0000	2.0004	1.5788

Electronic energy (eV)

Model	R^2	RMSE	MAE
LR	1.0000	3.5558	1.4958
LASSO	0.9996	116.89	53.32
Ridge	0.9980	259.63	190.23
ElasticNet	0.9996	116.52	53.14
DT	0.9833	742.58	501.17
RF	0.9905	560.47	299.48
XGBoost	0.9956	382.69	95.58
SLP	0.9955	386.12	266.58
MLP	0.9967	330.41	213.06
GCN	0.9731	941.56	785.22
KAN, spline (this work)	1.0000	1.2894	0.8035
KAN, symbolic (this work)	1.0000	2.0031	1.5822

Entropy (eV)

Model	R ²	RMSE	MAE
LR	0.9502	0.0399	0.0260
LASSO	0.9797	0.0255	0.0201
Ridge	0.9509	0.0396	0.0276
ElasticNet	0.9797	0.0255	0.0201
DT	0.9008	0.0563	0.0430
RF	0.9498	0.0400	0.0320
XGBoost	0.9790	0.0259	0.0201
SLP	0.9601	0.0357	0.0276
MLP	0.9655	0.0332	0.0255
GCN	0.9309	0.0470	0.0364
KAN, spline (this work)	0.9788	0.0260	0.0209
KAN, symbolic (this work)	0.9813	0.0245	0.0193

Enthalpy (eV)

Model	R ²	RMSE	MAE
LR	1.0000	3.5517	1.4877
LASSO	0.9996	116.96	53.66
Ridge	0.9980	259.57	190.21
ElasticNet	0.9996	116.58	53.48
DT	0.9833	742.67	501.18
RF	0.9905	560.60	299.61
XGBoost	0.9947	417.62	106.41
SLP	0.9953	392.99	274.88
MLP	0.9966	333.47	228.55
GCN	0.9733	938.51	784.90
KAN, spline (this work)	1.0000	1.2053	0.6962
KAN, symbolic (this work)	1.0000	1.9912	1.5735

Dipole moment (D)

Model	R ²	RMSE	MAE
LR	0.3330	1.9216	1.3337
LASSO	0.4625	1.7251	1.2070
Ridge	0.4361	1.7669	1.2268
ElasticNet	0.4624	1.7251	1.2070
DT	0.3123	1.9511	1.3412
RF	0.5030	1.6587	1.1124
XGBoost	0.4870	1.6852	1.1571
SLP	0.4097	1.8077	1.2376
MLP	0.2978	1.9717	1.3636
GCN	0.1924	2.1145	1.4325
KAN, spline (this work)	0.5395	1.5967	1.0803
KAN, symbolic (this work)	0.4886	1.6826	1.1820

Band gap (eV)

Model	R ²	RMSE	MAE
LR	0.6855	0.5072	0.4081
LASSO	0.7058	0.4906	0.3813
Ridge	0.6807	0.5111	0.3956
ElasticNet	0.7059	0.4906	0.3813
DT	0.7396	0.4616	0.3571
RF	0.8290	0.3740	0.2687
XGBoost	0.8285	0.3746	0.2595
SLP	0.6996	0.4958	0.3552
MLP	0.6936	0.5007	0.3594
GCN	0.7815	0.4228	0.3129
KAN, spline (this work)	0.8577	0.3413	0.2552
KAN, symbolic (this work)	0.6920	0.5020	0.3896

S5 Seed stability of the selected models

Training different seeds of the selected architecture gives tightly clustered test performance for the additive models. For the band gap, the ten seeds of the pinned [6, 1] model span R² 0.855-0.861 (mean 0.858, standard deviation 0.0015), and the representative seed (0.858) coincides with the mean. The energy targets vary only in the last digits of an R² that rounds to 1.0000, and entropy and the dipole moment show comparable stability. This robustness is why the additive architecture is preferred and pinned for the reported models.

S6 Interactive application

An interactive application built with Streamlit [16] applies the exported KAN models to new HF-3c inputs and returns the six predicted B3LYP-D4/def2-TZVP descriptors. It has two modes. In the single-compound mode (Fig. S1) the user enters the six HF-3c descriptors and reads off the six predictions together with each model’s test R² and RMSE; the closed-form formulas are shown alongside. In the batch mode (Fig. S2) the user uploads a CSV file with one row per compound. The columns are matched to the six descriptors automatically (a downloadable template is provided), predictions are computed for all rows, and a CSV file with six added pred_<target> columns can be downloaded.

The application is part of the repository and can be launched locally with a single command. It is also publicly available at <https://kan-hf-to-dft.streamlit.app/>.

HF-3c to DFT descriptor prediction

Kolmogorov-Arnold Networks predicting six B3LYP-D4/def2-TZVP descriptors from the six corresponding HF-3c descriptors.

Single compound **Batch (CSV)**

HF-3c descriptors

Enter the input descriptors (defaults are the dataset medians).

HF-3c Gibbs free energy (eV)	-18808,4435	-	+
HF-3c electronic energy (eV)	-18819,7443	-	+
HF-3c entropy (eV)	1,7446	-	+
HF-3c enthalpy (eV)	-18806,5946	-	+
HF-3c dipole moment (D)	3,1655	-	+
HF-3c HOMO-LUMO gap (eV)	18,7608	-	+

Predicted DFT descriptors

DFT descriptor	Predicted value	Test R2	Test RMSE
Gibbs free energy (eV)	-19054.6581	1	1.1581
Electronic energy (eV)	-19063.7075	1	1.2894
Entropy (eV)	1.814	0.9788	0.026
Enthalpy (eV)	-19052.7297	1	1.2053
Dipole moment (D)	2.9732	0.5395	1.5967
Band gap (eV)	7.041	0.8577	0.3413

Fig. S1 Single-compound mode of the interactive application

HF-3c to DFT descriptor prediction

Kolmogorov-Arnold Networks predicting six B3LYP-D4/def2-TZVP descriptors from the six corresponding HF-3c descriptors.

Single compound **Batch (CSV)**

Batch prediction from a CSV file

Upload a CSV with one row per compound and the six HF-3c descriptors as columns. Columns are matched automatically by name; any extra columns (e.g. an identifier) are passed through to the output.

Download CSV template

HF-3c CSV

Upload 200MB per file • CSV

> Closed-form symbolic expressions

Representative single-seed models behind Figure 1 (validation-selected seed). Their held-out test metrics are the ones reported in comparison_main.csv. Headline metrics in the paper are ten-seed means.

Fig. S2 Batch (CSV) mode of the interactive application

References

Reference numbering follows the main article.

10. Normatov S, Nesterov PV, Aliev TA, Timralieva AA, Novikov AS, Skorb EV (2025) Search for correlations between the results of the density functional theory and Hartree–Fock calculations using neural networks and classical machine learning algorithms. *ACS Omega* 10:5919–5933. <https://doi.org/10.1021/acsomega.4c09861>
14. Liu Z, Wang Y, Vaidya S, Ruehle F, Halverson J, Soljačić M, Hou TY, Tegmark M (2024) KAN: Kolmogorov-Arnold networks. *arXiv:2404.19756*. <https://doi.org/10.48550/arXiv.2404.19756>
15. Liu Z, Ma P, Wang Y, Matusik W, Tegmark M (2024) KAN 2.0: Kolmogorov-Arnold networks meet science. *arXiv:2408.10205*. <https://doi.org/10.48550/arXiv.2408.10205>
16. Streamlit Inc (2024) Streamlit: an open-source Python framework for building interactive data applications. <https://streamlit.io>. Accessed 24 September 2026