

434
435
436

Supplementary material for Variational Monte Carlo on a Budget – Fine-tuning pre-trained Neural Wavefunction

437 **A Molecule datasets**

438 **Bicyclobutane** For the Bicyclobutane to 1,3-butadiene transition we use the geometries from Kinal
439 et al. [30] and compare against the reference energies stated in Spencer et al. [5].

440 **N₂** For the N₂ potential energy surface with various bond-lengths we used the geometries including
441 reference calculations from [8].

442 **Propadiene** The global rotation of 360° degrees for propadiene is performed on the geometry
443 which is part of the test set for 3 heavy atoms. For the torsion experiment we used the equilibrium
444 geometry and rotated the torsion angle by 90° degrees in steps of 10° degrees.

445 **Zero-shot and fine-tuning dataset** The results on zero-shot and few-shot predictions for increasing
446 number of heavy atoms are performed on random subsets of molecules. For 5-7 heavy atoms we
447 sample 4 unique and distorted molecules from QM7-X [26]. For 4 heavy atoms we use all geometries
448 from the Bicyclobutane dataset. For 3 heavy atoms we use the ablation dataset.

449 **Ablation dataset** For the ablation study, we use one geometry per molecule from the out-of-
450 distribution test set from Scherbela et al. [25], leading to a set of four distinct molecules. We ensure
451 that these molecules are not part of the training set.

452 **Large scale experiment** For the large scale experiment we used a stratified random sample of 250
453 molecules from QM7 [31]. It contains all molecules with up to 4 heavy atoms, and additionally 65
454 randomly chosen molecules for 5, 6 and 7 heavy atoms each.

455 **Pre-training dataset for transferable neural wavefunctions** To train our pre-trained wavefunc-
456 tions we use two datasets, consisting of 18 and 98 disparate molecules. For part of the ablation we use
457 the dataset proposed in [25] and an extended version with 80 additional molecules. The additional
458 compounds are a combination of all valid SMILES generated with RDKit [34] with 3 heavy atoms,
459 allowing only Nitrogen, Oxygen and Carbon with single-, double- or triple-bonds, and all molecules
460 up to four heavy atoms from QM7-X [26] (excluding molecule containing Fluorine). To prevent a
461 train-test leakage, we remove Bicyclobutane (including all conformations) and the four molecules
462 from the ablations dataset. Since the normal-mode-distortions by design do not generate strongly
463 distorted geometries, we augment the 98-molecule-dataset with rotated dihedral angles. To generate a
464 subset of all possible dihedral angles for a heavy-atom bond we first generate samples with equidistant
465 angles for all possible dihedral angles and compute Hartree-Fock energies with a minimal basis-set.
466 We include the equilibrium geometry and all extrema of the potential energy surface with respect to
467 the rotation of a single dihedral angle if the energy of the extrema is significantly different to already
468 included geometries of the same molecule. Additionally, we include the transition geometry towards
469 the respective extrema and again only include energetic diverse states. Finally, to make sure that
470 certain molecules are not underrepresented in the dataset we make sure that all molecules have at
471 least 5 geometries that get distorted during pre-training by adding copies of the equilibrium geometry.
472 Overall this yields 699 initial geometries R^0 for pre-training.

473 **B Electron MCMC initialization**

474 To investigate the impact of the initial distribution of electron positions on the equilibration of the
475 Markov Chain, we run two evaluations for a glycine molecule, using a pre-trained wavefunction. We
476 perform no initial burn-in and use every 50th sample for energy evaluation. If the chain was perfectly
477 equilibrated right after initialization, all sampled energies would fluctuate around the mean energy.
478 However as Fig. 6b shows, it takes several thousand steps for the sampled energies to converge to
479 the correct mean. This is particularly pronounced with Gaussian initialization of electron positions,

480 which is the default in state-of-the-art DL-VMC codes such as FermiNet [3]. Using an exponential
 481 distribution of the initial electron positions much more closely resembles the correct electron density
 482 ψ^2 (cf. Fig. 6a) and thus reaches equilibrium substantially faster.

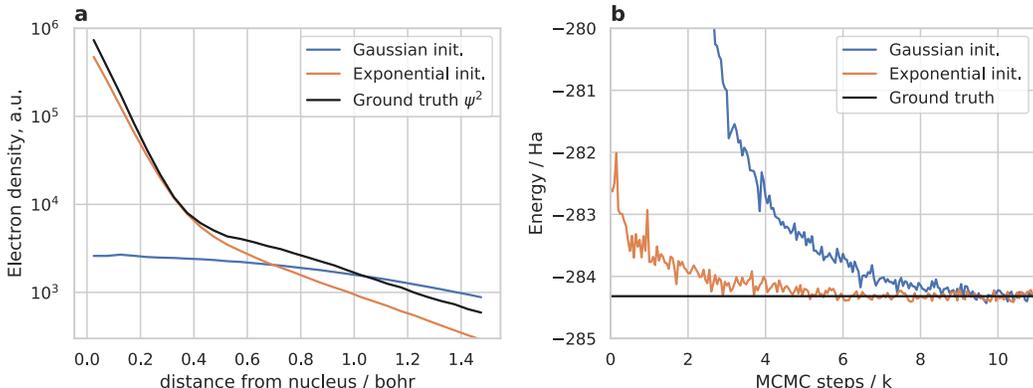


Figure 6: **Effect of electron initialization:** Initializing the electron positions using an exponential distribution instead of Gaussian, better fits the actual density (a), and thus leads to faster equilibration of observables during evaluation (b).

483 C Orbital localization

484 Our model uses orbital embeddings \mathbf{x}^{orb} as inputs to parameterize the backflows \mathbf{f}^{orb} , and exponents
 485 \mathbf{g}^{orb} of the orbitals. These orbital embeddings were introduced by Scherbela et al. [25] in the form of
 486 molecular orbital expansion coefficients, obtained from a self consistent Hartree-Fock calculation. In
 487 this setting, the coefficients \mathbf{x}^{orb} are not uniquely defined, but only up to a linear transformation U
 488 with determinant ± 1

$$\mathbf{x}_{Ik}^{\text{orb}} = \sum_{n=1}^{N_{\text{orb}}} U_{kn} \hat{\mathbf{x}}_{In}^{\text{orb}}, \quad U \in \mathbb{R}^{N_{\text{orb}} \times N_{\text{orb}}}, \quad \det U = \pm 1. \quad (15)$$

489 This stems from the fact that the corresponding Hartree-Fock wavefunction is invariant under such a
 490 transformation. Consequently there is free choice, which linear combination of embeddings \mathbf{x}^{orb} to
 491 choose from without any loss of information. We follow the approach of [25], by choosing U such
 492 that the corresponding Hartree-Fock orbitals are maximally localized according to the Foster-Boys
 493 metric, i.e. minimize the spatial variance \mathcal{L} :

$$\phi_k(\mathbf{r}, U) = \sum_{I=1}^{N_{\text{nuc}}} \sum_{\mu=1}^{N_{\text{basis}}} b_{I\mu}(\mathbf{r}) U_{kn} \hat{\mathbf{x}}_{In\mu}^{\text{orb}} \quad (16)$$

$$\mathcal{L}(U) = \sum_k \int \phi_k^2(\mathbf{r}, U) \mathbf{r}^2 d\mathbf{r} - \left(\int \phi_k^2(\mathbf{r}, U) \mathbf{r} d\mathbf{r} \right)^2 \quad (17)$$

494 Here $b_{I\mu}(\mathbf{r})$ denotes μ -th basis function of the Hartree-Fock expansion, centered on the I -th nucleus.
 495 In practice the integrals of Eq. 17 do not have to be evaluated explicitly, but can instead be computed
 496 via the overlap matrix \mathcal{S} . The minimization of \mathcal{L} is typically done iteratively, requires on the order of
 497 10 steps, and is readily implemented in many open-source quantum chemistry codes such as pySCF
 498 [35].

499 D Tables of energies

500 D.1 Conformers of Bicyclobutane

501 In Tab. 1 we list the relative energies of our method and all reference energies corresponding to Fig.
 502 3.

Table 1: Energies relative to the energy of bicyclobutane in mHa, including the zero-point vibrational energy correction from Kinal et al. [30]

structure	CCSD(T) [30]	DMC [30]	FermiNet 200k [5]	FermiNet 10k [5]	Our work zero-shot	Our work 700 per geom
con_TS	64.4	64.4	64.1	63.9	94.0	66.6
dis_TS	34.7	93.4	92.0	87.1	183.8	94.5
g-but	-40.0	-40.2	-40.3	-44.9	-48.5	-40.4
gt-TS	-35.5	-35.4	-35.9	-42.9	-46.9	-36.7
t-but	-44.6	-44.5	-45.3	-47.5	-51.8	-43.2

503 E Reference energies

504 **CCSD(T)** All CCSD(T) energies – except explicitly stated otherwise – were obtained using ORCA
 505 [36] starting from a restricted Hartree-Fock calculation. We use correlation consistent basis sets of the
 506 cc-pCVXZ family, with X in {2, 3, 4}. To extrapolate to the complete basis set limit (CBS), we use
 507 the approach outlined in [3] and fit the following functions with free parameters $E_{\text{CBS}}^{\text{HF}}$, $E_{\text{CBS}}^{\text{corr}}$, a , b , c :

$$E_X^{\text{HF}} = E_{\text{CBS}}^{\text{HF}} + ae^{-bX}$$

$$E_X^{\text{corr}} := E_X^{\text{HF}} - E_X^{\text{CCSD(T)}} = E_{\text{CBS}}^{\text{corr}} + cX^{-3}$$

$$E_{\text{CBS}}^{\text{CCSD(T)}} = E_{\text{CBS}}^{\text{HF}} + E_{\text{CBS}}^{\text{corr}}$$

508 We stress that although CCSD(T)-energies are often considered as "gold-standard", they do not
 509 necessarily represent the actual ground-state energy. There are many cases, where CCSD(T) either
 510 overestimates the true ground-state energy, or even underestimates it, because CCSD(T) does not
 511 yield upper bounds to the true ground-state energy.

512 **PsiFormer** For Fig. 2 we used the open-source FermiNet codebase [37]. The codebase didn't allow
 513 for inference calculation, therefore a slight fix was applied. All calculations were performed with the
 514 small settings as proposed in von Glehn et al. [6].

515 F Adaption of PhisNet

516 We heavily rely on PhisNet by Unke et al. [20] to obtain orbital descriptors without the need for a
 517 separate SCF calculation. Compared to their original work, we made several simplifications, which
 518 are motivated by the fact that we do not predict final high-accuracy orbitals in a large basis set, but
 519 only use PhisNet as a feature extractor by predicting orbitals in a minimal basis-set:

- 520 • **Layer Norm** We found deep variants of PhisNet to be unstable to train and mitigated the
 521 issue by adding an (equivariant) layer norm after each PhisNet module.
- 522 • **Simplified Fock matrix prediction** The original PhisNet implementation uses a final
 523 interaction between the node embeddings, before predicting the elements of the Fock matrix.
 524 We found this interaction to be superfluous for our purposes and left it out for simplicity.
- 525 • **Separate energy head** The original PhisNet computes energies via the eigenvalues obtained
 526 by diagonalization of the Fock matrix. We instead predict energies using a separate head on
 527 top of the scalar features of the node embeddings.
- 528 • **Smaller network** We changed the hyperparameters to obtain a smaller and faster version of
 529 PhisNet which obtained sufficient accuracy for our purposes. We used 2 layers (instead of 5)
 530 and $L_{\text{max}} = 2$ (instead of 4). This reduces the number of parameters from 17M to 3M.
- 531 • **Diverse training set** While the original work optimized separate models for each molecule
 532 (e.g. by training on different geometries of a molecular dynamics simulation), we optimize a
 533 single model to predict F , S , E , and ∇E across a dataset of 47k geometries sampled from
 534 QM7-X [26].
- 535 • **JAX re-implementation** We re-implemented PhisNet in JAX, using the e3nn library [38] to
 536 construct the SE(3)-equivariant operations.

537 We train the PhisNet-model on a dataset of 47k molecules from QM7X [26], using the Adam optimizer
 538 [39] on the following loss

$$\mathcal{L} = \sum_n (E^{\text{phis}}(\mathbf{R}^n, \mathbf{Z}^n) - E^{\text{ref},n})^2 + \quad (18)$$

$$+ \sum_{nI\zeta} \left(\frac{\partial}{\partial R_{I\zeta}^n} E^{\text{phis}}(\mathbf{R}^n, \mathbf{Z}^n) - G_{I\zeta}^{\text{ref},n} \right)^2 + \quad (19)$$

$$+ \sum_{nIJ\mu\nu} \left(F_{IJ\mu\nu}^{\text{phis}}(\mathbf{R}^n, \mathbf{Z}^n) - F_{IJ\mu\nu}^{\text{ref},n} \right)^2 + \quad (20)$$

$$+ \sum_{nIJ\mu\nu} \left(S_{IJ\mu\nu}^{\text{phis}}(\mathbf{R}^n, \mathbf{Z}^n) - S_{IJ\mu\nu}^{\text{ref},n} \right)^2. \quad (21)$$

539 Here E denotes energies, G denotes gradients of energies, F Fock matrices, and S overlap matrices.
 540 The indices I, J run over nuclei, the indices μ, ν over basis functions, and the index n over samples
 541 in a batch.

542 G Hyperparameters

543 A detailed description of the hyperparameter used in this work can be found below (cf. Tab. 2). For
 544 the mapping of the orbital descriptors to the electron embeddings to build the orbitals we rely on the
 545 hyperparameter from [25]. For optimization we rely on the second-order method KFAC [29] and use
 546 their Python implementation [40]. During the continuous sampling of the geometries we allow each
 547 geometry to perform a maximum of 20 steps of normal-mode distortion from the initial geometry and
 548 reset to the original one once the threshold is reached.

549 H Computational resources

550 We used ≈ 5 k GPUhs (A100) for development and training of our base models, and another 5k GPUhs
 551 (A40) on evaluations and fine-tuning. Additionally we required ≈ 20 k CPUhs for CCSD(T) reference
 552 calculations.

553 I Code and data availability

554 All code, configuration files, geometries, datasets and obtained energies are available in the supple-
 555 mentary information.

Table 2: Hyperparameter settings used in this work

Electron Embedding	Hidden dimension N_{emb}	256
	N ^o iterations	4
Nuclear Embedding	Hidden dimension \tilde{x}^{nuc}	64
	N ^o layer MLP	1
Message passing	Activation function	SiLU
	N ^o layer edge embedding	3
	Dimension edge embedding	64
	Dimension linear layer	32
Markov Chain Monte Carlo	N ^o walkers	2048
	N ^o decorrelation steps	50
	Target acceptance prob.	50%
PhisNet [20]	Pre-trained against basis set	STO-6G
	N ^o iterations	2
	Harmonic degree L	2
	N ^o radial basis functions	128
	Hidden dimension of \mathbf{x}^{nuc}	128
	Distance cutoff (bohr)	30
Transferable atomic orbitals [25]	N ^o determinants N_{det}	8
	N ^o hidden layers \mathbf{f}^{orb}	2
	Hidden dimension of \mathbf{f}^{orb}	256
	N ^o hidden layers \mathbf{g}^{orb}	2
	Hidden dimension \mathbf{g}^{orb}	128
	N ^o iterations MPNN	2
	N ^o radial basis functions	16
	Hidden edge embedding dimension	32
	Hidden node embedding dimension	16
	Activation function	SiLU
Variational pre-training	Optimizer	KFAC
	Batch size	2048
	Norm constraint	3×10^{-3}
	Initial damping d_0	1
	Minimal damping d_{min}	0.001
	Damping rate decay	$d(t) = d_0 \exp(-t/20000)$
	Initial learning rate lr_0	0.1
	Learning rate decay	$lr(t) = lr_0(1 + t/6000)^{-1}$
	Optimization steps	128,000 - 256,000
Changes for fine-tuning	Learning rate decay	$lr(t) = lr_0(7 + t/6000)^{-1}$
	Optimization steps	0 - 32,000
Sampling geometries	Distortion energy β	0.005 Ha
	Max age	20
	Bias towards original geometry α	0.2

329 **References**

- 330 [1] M. Born and R. Oppenheimer. “Zur Quantentheorie der Molekeln”. In: *Annalen der Physik*
 331 389.20 (1927), pp. 457–484. DOI: <https://doi.org/10.1002/andp.19273892002>.
- 332 [2] László Gyevi-Nagy, Mihály Kállay, and Péter R. Nagy. “Accurate Reduced-Cost CCSD(T)
 333 Energies: Parallel Implementation, Benchmarks, and Large-Scale Applications”. In: *Journal of*
 334 *Chemical Theory and Computation* 17.2 (Feb. 2021), pp. 860–878. DOI: [10.1021/acs.jctc](https://doi.org/10.1021/acs.jctc.0c01077)
 335 [0c01077](https://doi.org/10.1021/acs.jctc.0c01077).
- 336 [3] David Pfau et al. “Ab Initio Solution of the Many-Electron Schrödinger Equation with
 337 Deep Neural Networks”. In: *Phys. Rev. Res.* 2.3 (Sept. 2020), p. 033429. DOI: [10.1103/](https://doi.org/10.1103/PhysRevResearch.2.033429)
 338 [PhysRevResearch.2.033429](https://doi.org/10.1103/PhysRevResearch.2.033429).
- 339 [4] W. K. Hastings. “Monte Carlo Sampling Methods Using Markov Chains and Their Applica-
 340 tions”. In: *Biometrika* 57.1 (Apr. 1970), pp. 97–109. DOI: [10.1093/biomet/57.1.97](https://doi.org/10.1093/biomet/57.1.97).
- 341 [5] James S. Spencer et al. “Better, Faster Fermionic Neural Networks”. In: (2020).
- 342 [6] Ingrid von Glehn, James S. Spencer, and David Pfau. *A Self-Attention Ansatz for Ab-initio*
 343 *Quantum Chemistry*. Nov. 2022. DOI: [10.48550/arXiv.2211.13672](https://doi.org/10.48550/arXiv.2211.13672).
- 344 [7] Weiluo Ren et al. “Towards the ground state of molecules via diffusion Monte Carlo on neural
 345 networks”. In: *Nature Communications* 14.1 (Apr. 2023), p. 1860. DOI: [10.1038/s41467-](https://doi.org/10.1038/s41467-023-37609-3)
 346 [023-37609-3](https://doi.org/10.1038/s41467-023-37609-3).
- 347 [8] Leon Gerard et al. “Gold-Standard Solutions to the Schrödinger Equation Using Deep Learning:
 348 How Much Physics Do We Need?” In: *Advances in Neural Information Processing Systems*.
 349 Oct. 2022.
- 350 [9] M. T. Entwistle et al. “Electronic excited states in deep variational Monte Carlo”. In: *Nature*
 351 *Communications* 14.1 (Jan. 2023), p. 274. DOI: [10.1038/s41467-022-35534-5](https://doi.org/10.1038/s41467-022-35534-5).
- 352 [10] Michael Scherbela et al. “Solving the Electronic Schrödinger Equation for Multiple Nuclear
 353 Geometries with Weight-Sharing Deep Neural Networks”. In: *Nature Computational Science*
 354 2.5 (May 2022), pp. 331–341. DOI: [10.1038/s43588-022-00228-x](https://doi.org/10.1038/s43588-022-00228-x).
- 355 [11] Yubing Qian et al. “Interatomic Force from Neural Network Based Variational Quantum
 356 Monte Carlo”. In: *The Journal of Chemical Physics* 157.16 (Oct. 2022), p. 164104. DOI:
 357 [10.1063/5.0112344](https://doi.org/10.1063/5.0112344).
- 358 [12] Gino Cassella et al. “Discovering Quantum Phase Transitions with Fermionic Neural
 359 Networks”. In: *Physical Review Letters* 130.3 (Jan. 2023), p. 036401. DOI: [10.1103/](https://doi.org/10.1103/PhysRevLett.130.036401)
 360 [PhysRevLett.130.036401](https://doi.org/10.1103/PhysRevLett.130.036401).
- 361 [13] Xiang Li, Zhe Li, and Ji Chen. “Ab initio calculation of real solids via neural network ansatz”.
 362 In: *Nature Communications* 13.1 (Dec. 2022), p. 7895. DOI: [10.1038/s41467-022-35627-](https://doi.org/10.1038/s41467-022-35627-1)
 363 [1](https://doi.org/10.1038/s41467-022-35627-1).
- 364 [14] Giuseppe Carleo and Matthias Troyer. “Solving the Quantum Many-Body Problem with
 365 Artificial Neural Networks”. In: *Science* 355.6325 (Feb. 2017), pp. 602–606. DOI: [10.1126/](https://doi.org/10.1126/science.aag2302)
 366 [science.aag2302](https://doi.org/10.1126/science.aag2302).
- 367 [15] Giuseppe Carleo et al. “NetKet: A Machine Learning Toolkit for Many-Body Quantum
 368 Systems”. In: *SoftwareX* 10 (July 2019), p. 100311. DOI: [10.1016/j.softx.2019.100311](https://doi.org/10.1016/j.softx.2019.100311).
- 369 [16] Yuan-Hang Zhang and Massimiliano Di Ventra. “Transformer Quantum State: A Multipur-
 370 pose Model for Quantum Many-Body Problems”. In: *Physical Review B* 107.7 (Feb. 2023),
 371 p. 075147. DOI: [10.1103/PhysRevB.107.075147](https://doi.org/10.1103/PhysRevB.107.075147).
- 372 [17] Nicholas Gao and Stephan Günemann. *Generalizing Neural Wave Functions*. Feb. 2023. DOI:
 373 [10.48550/arXiv.2302.04168](https://doi.org/10.48550/arXiv.2302.04168).
- 374 [18] Simon Batzner et al. “E(3)-Equivariant Graph Neural Networks for Data-Efficient and Accurate
 375 Interatomic Potentials”. In: *Nature Communications* 13.1 (May 2022), p. 2453. DOI: [10.1038/](https://doi.org/10.1038/s41467-022-29939-5)
 376 [s41467-022-29939-5](https://doi.org/10.1038/s41467-022-29939-5).
- 377 [19] Ilyes Batatia et al. “MACE: Higher Order Equivariant Message Passing Neural Networks for
 378 Fast and Accurate Force Fields”. In: *Advances in Neural Information Processing Systems*.
 379 2022.
- 380 [20] Oliver Unke et al. “SE(3)-Equivariant Prediction of Molecular Wavefunctions and Electronic
 381 Densities”. In: *Advances in Neural Information Processing Systems*. 2021.
- 382 [21] M. Gastegger et al. “A Deep Neural Network for Molecular Wave Functions in Quasi-Atomic
 383 Minimal Basis Representation”. In: *The Journal of Chemical Physics* 153.4 (July 2020),
 384 p. 044123. DOI: [10.1063/5.0012911](https://doi.org/10.1063/5.0012911).

- 385 [22] Albert Musaelian et al. “Scaling the Leading Accuracy of Deep Equivariant Models to
386 Biomolecular Simulations of Realistic Size”. In: (Apr. 2023). DOI: [10.48550/arXiv.2304.
387 10061](https://doi.org/10.48550/arXiv.2304.10061), arXiv: [2304.10061](https://arxiv.org/abs/2304.10061)
- 388 [23] Nicholas Gao and Stephan Günnemann. “Ab-Initio Potential Energy Surfaces by Pairing
389 GNNs with Neural Wave Functions”. In: *arXiv:2110.05064 [physics]* (Nov. 2021). arXiv:
390 [2110.05064 \[physics\]](https://arxiv.org/abs/2110.05064).
- 391 [24] Nicholas Gao and Stephan Günnemann. “Sampling-free Inference for Ab-Initio Potential
392 Energy Surface Networks”. In: (2022). DOI: [10.48550/ARXIV.2205.14962](https://doi.org/10.48550/ARXIV.2205.14962).
- 393 [25] Michael Scherbela, Leon Gerard, and Philipp Grohs. “Towards a Foundation Model for Neural
394 Network Wavefunctions”. In: *arXiv.org* (Mar. 2023).
- 395 [26] Johannes Hoja et al. “QM7-X, a comprehensive dataset of quantum-mechanical properties
396 spanning the chemical space of small organic molecules”. In: *Scientific Data* 8.1 (Feb. 2021),
397 p. 43. DOI: [10.1038/s41597-021-00812-2](https://doi.org/10.1038/s41597-021-00812-2).
- 398 [27] János Pipek and Paul G. Mezey. “A Fast Intrinsic Localization Procedure Applicable for Ab
399 Initio and Semiempirical Linear Combination of Atomic Orbital Wave Functions”. In: *The
400 Journal of Chemical Physics* 90.9 (May 1989), pp. 4916–4926. DOI: [10.1063/1.456588](https://doi.org/10.1063/1.456588).
- 401 [28] Claudia Filippi and C. J. Umrigar. “Correlated Sampling in Quantum Monte Carlo: A Route
402 to Forces”. In: *Physical Review B* 61.24 (June 2000), R16291–R16294. DOI: [10.1103/
403 PhysRevB.61.R16291](https://doi.org/10.1103/PhysRevB.61.R16291).
- 404 [29] James Martens and Roger Grosse. “Optimizing neural networks with kronecker-factored
405 approximate curvature”. In: *International conference on machine learning*. PMLR. 2015,
406 pp. 2408–2417.
- 407 [30] Armağan Kinal and Piotr Piecuch. “Computational Investigation of the Conrotatory and
408 Disrotatory Isomerization Channels of Bicyclo[1.1.0]butane to Buta-1,3-diene: A Completely
409 Renormalized Coupled-Cluster Study”. In: *The Journal of Physical Chemistry A* 111.4 (Feb.
410 2007). Publisher: American Chemical Society, pp. 734–742. ISSN: 1089-5639. DOI: [10.1021/
411 jp065721k](https://doi.org/10.1021/jp065721k). URL: <https://doi.org/10.1021/jp065721k>.
- 412 [31] Matthias Rupp et al. “Fast and Accurate Modeling of Molecular Atomization Energies with
413 Machine Learning”. In: *Phys. Rev. Lett.* 108 (5 Jan. 2012), p. 058301. DOI: [10.1103/
414 PhysRevLett.108.058301](https://doi.org/10.1103/PhysRevLett.108.058301).
- 415 [32] Kristof T. Schütt et al. “SchNet: A Continuous-Filter Convolutional Neural Network for
416 Modeling Quantum Interactions”. In: (June 2017). DOI: [10.48550/arXiv.1706.08566](https://doi.org/10.48550/arXiv.1706.08566).
- 417 [33] Gabriel Pescia et al. *Message-Passing Neural Quantum States for the Homogeneous Electron
418 Gas*. May 2023. DOI: [10.48550/arXiv.2305.07240](https://doi.org/10.48550/arXiv.2305.07240).
- 419 [34] Greg Landrum. *RDKit: Open-source cheminformatics*. [https://github.com/rdkit/
420 rdkit](https://github.com/rdkit/rdkit). GitHub repository. 2009.
- 421 [35] Qiming Sun et al. “Recent Developments in the PySCF Program Package”. In: *The Journal of
422 Chemical Physics* 153.2 (July 2020), p. 024109. DOI: [10.1063/5.0006074](https://doi.org/10.1063/5.0006074).
- 423 [36] Frank Neese et al. “The ORCA Quantum Chemistry Program Package”. In: *The Journal of
424 Chemical Physics* 152.22 (June 2020), p. 224108. DOI: [10.1063/5.0004608](https://doi.org/10.1063/5.0004608).
- 425 [37] David Pfau James S. Spencer and FermiNet Contributors. *FermiNet*. 2020. URL: [http:
426 //github.com/deepmind/ferminet](http://github.com/deepmind/ferminet).
- 427 [38] Mario Geiger and Tess Smidt. *E3nn: Euclidean Neural Networks*. July 2022. arXiv: [2207.
428 09453 \[cs\]](https://arxiv.org/abs/2207.09453).
- 429 [39] Diederik P. Kingma and Jimmy Ba. “Adam: A Method for Stochastic Optimization”. In:
430 arXiv:1412.6980 (Jan. 2017). DOI: [10.48550/arXiv.1412.6980](https://doi.org/10.48550/arXiv.1412.6980), arXiv: [1412.6980 \[cs\]](https://arxiv.org/abs/1412.6980).
431 (Visited on 05/17/2023).
- 432 [40] Aleksandar Botev and James Martens. *KFAC-JAX*. Version 0.0.1. 2022. URL: [http://github.
433 com/deepmind/kfac-jax](http://github.com/deepmind/kfac-jax).