

A FUNCTION-CENTRIC GRAPH NEURAL NETWORK APPROACH FOR PREDICTING ELECTRON DENSITIES

005 **Anonymous authors**

006 Paper under double-blind review

ABSTRACT

011 Electronic structure predictions are relevant for a wide range of applications, from
 012 drug discovery to materials science. Since the cost of purely quantum mechanical
 013 methods can be prohibitive, machine learning surrogates are used to predict the
 014 result of these calculations. This work introduces the Basis Overlap Architecture
 015 (BOA), an equivariant graph neural network architecture based on a novel mes-
 016 sage passing scheme that utilizes the overlap matrix of the basis functions used
 017 to represent the ground state electron density. BOA is evaluated on QM9 and
 018 MD density datasets, surpassing the previous state-of-the-art in predicting accu-
 019 rate electron densities. Excellent generalization to larger molecules of up to nearly
 020 200 atoms is demonstrated using a model trained only on QM9 molecules of up
 021 to 9 heavy atoms.

1 INTRODUCTION

025 Accurate electronic structure predictions are crucial for the development of new catalysts, improved
 026 batteries or more specific drugs. Today’s gold standard for reasonably sized systems is Kohn-Sham
 027 density functional theory (KS-DFT). It accounts for a significant fraction of worldwide supercom-
 028 puting time, and three of its cornerstone methods are amongst the ten most cited publications of all
 029 times and fields (Van Noorden, 2025). Still, its computational cost prohibits routine use on large sys-
 030 tems, or in very high throughput scenarios. In response, machine learning surrogates are developed
 031 to either circumvent or speed up KS-DFT calculations.

032 These methods range from property prediction (directly predicting observables from molecular ge-
 033 ometry (Batatia et al., 2022; Ko et al., 2023; Kozinsky et al., 2023; Simeon & Fabritiis, 2023; Batatia
 034 et al., 2025; Liao et al., 2023; Anstine et al., 2025; Wood et al., 2025)) to incorporating more physi-
 035 cal knowledge into the model. Examples for the latter include Hamiltonian prediction (Zhang et al.,
 036 2024a; Yuan et al., 2024; Wang et al., 2024; Qian et al., 2025; Luo et al., 2025), mimicking the
 037 self-consistent iterations of KS-DFT (Song & Feng, 2024), or predicting the electron density using
 038 orbital-free schemes (Zhang et al., 2024b; Remme et al., 2025). Directly predicting the electron den-
 039 sity, as done here, lies between these extremes. The ground state electron density is an observable of
 040 central interest as it uniquely determines all ground state molecular properties in theory (Hohenberg
 041 & Kohn, 1964), and many properties of interest can be derived from the density in practice. Another
 042 use for electron densities is to reduce the number of self-consistent field iterations required in KS-
 043 DFT (Koker et al., 2024; Sunshine et al., 2023). Even a single KS-DFT diagonalization step can be
 044 enough to reach chemically accurate energies relative to a KS self-consistent solution (Jørgensen &
 045 Bhowmik, 2022; Li et al., 2025).

046 Most previous work on directly predicting ground state electron densities can be broadly grouped
 047 into two categories based on the representation used: The first class is based on a representation of
 048 the density in a typically atom-centered basis (Cuevas-Zuviría & Pacios, 2021; Rackers et al., 2023;
 049 del Rio et al., 2023; Elsborg et al., 2025; Qiao et al., 2022; 2020; Fu et al., 2024; Cheng & Peng,
 050 2023; Kim & Ahn, 2024; Mitnikov & Jacobson, 2024; Febrer et al., 2025). A model would then
 051 predict the coefficients of an expansion of the ground state electron density in the given basis. This
 052 approach has the advantage of being relatively scalable, as the density only needs to be evaluated on
 053 a volumetric grid at the end of the model, if at all. The choice of basis is however crucial, and to
 achieve high accuracy a very large number of basis functions is typically required. This problem was
 partially addressed by Fu et al. (2024), where virtual nodes are employed, and the basis functions

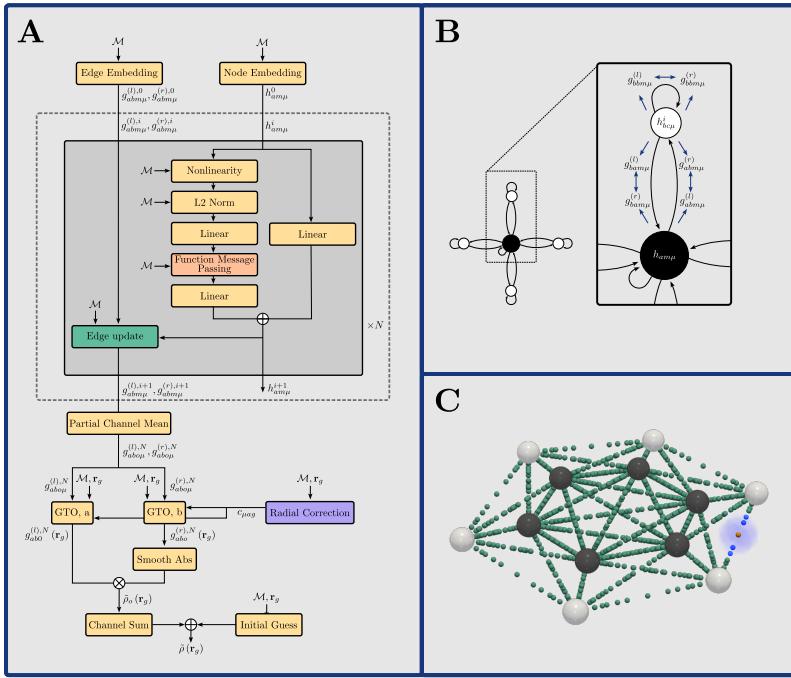


Figure 1: **The Basis Overlap Architecture (BOA).** An overview of BOA is given in (A). The input molecular geometry is fed into the edge and node embedding layers. The node embeddings are updated using BOA blocks, which contain a function message passing step to facilitate communication between nodes. The edge features are modified using the edge update block, which uses the current edge and node features to calculate new edge features. The output of the BOA backbone consists of coefficients used to expand the density in atom-centered Gaussian-type basis functions as described in Section 2.1. In the partial channel mean the number of channels is reduced by taking the mean of groups of channels. M denotes the molecular geometry and \mathbf{r}_g are the grid positions with $g \in \{1, \dots, N^g\}$ and N^g being the number of grid points. (B) shows the flow of information in the BOA message passing block. While the edge features are updated using the node features, there is no flow of information from the edge features to the node features. Every edge in the graph has two edge features with superscript (l) and (r) respectively, corresponding to the two nodes the edge connects. Self-loops are included in the graph. (C) shows a benzene molecule from the MD dataset with carbon atoms (black) and hydrogen atoms (white). Shown as green spheres are the centers of the Gaussian parts of the products of the atom-centered basis functions. The smaller def2-SVP (Weigend & Ahlrichs, 2005) basis set is used here for better visualization. A product of $l = 0$ basis functions from two different hydrogen atoms is shown in blue. The center of the product of these two basis functions is shown in yellow, lying in the middle between the two atoms on which the original basis functions are centered. Since the products of basis functions are well distributed in space, a highly accurate representation of the density can be achieved without the need for floating basis functions or virtual nodes.

adapted in a fine-tuning step to mitigate the impact of a suboptimal basis. Virtual nodes are placed in the molecule, e.g. at the midpoint of bonds, and additional localized basis functions are centered at these nodes. Another approach to address this problem was taken by Elsborg et al. (2025), where floating basis functions are employed. These floating orbitals are no longer centered at the atoms of the molecule and instead the basis function positions are predicted per molecule individually, enabling a much more flexible representation of the density.

The second class of methods works directly with a representation of the electron density on a volumetric grid (Jørgensen & Bhowmik, 2022; Li et al., 2025; Koker et al., 2024; Gong et al., 2019; Sunshine et al., 2023). While this approach avoids basis set related inaccuracies, it is typically much more memory intensive, due to the large number of grid points required to accurately represent the density.

108 **Contributions** This work falls into the first category of methods, i.e. we represent the density
 109 in an atom-centered basis and the model predicts coefficients in this basis. In contrast to previous
 110 work, we, however, choose to represent the density in a quadratic expansion of the basis functions,
 111 inspired by the internal representation of the density in KS-DFT calculations using a density matrix.
 112 A product of two atom-centered Gaussian-type basis functions, as employed in this work, will be
 113 centered between the two atoms, avoiding the need for floating orbitals or virtual nodes. For a
 114 benzene molecule, this is shown in Figure 1C. This work avoids explicitly predicting coefficients
 115 for each pair of basis functions, effectively employing a low-rank representation of blocks of the
 116 density matrix, without constructing the full coefficient matrix at any point.

117 To effectively utilize this representation, we introduce the Basis Overlap Architecture (BOA), a
 118 novel equivariant message passing neural network architecture. The full BOA architecture is shown
 119 in Figure 1A. The fundamental idea of BOA is to imbue the model with information about the un-
 120 derlying basis by *interpreting*, at appropriate points in the model, the internal features *as functions*
 121 represented in the given basis. Most notably the message passing formulation utilizes the overlap
 122 between functions represented in atom-centered basis functions at different atoms to facilitate com-
 123 munication between nodes. During message passing each message is transformed from the basis of
 124 the sending node to the basis of the receiving node, essentially giving the best fit of the message in
 125 the basis of the receiving node. This instills the model with information not only about the under-
 126 lying basis, but also the geometry of the molecule, as the overlap between basis functions centered at
 127 different atoms depends on their relative position. This novel approach to message passing is shown
 128 to be very effective for the task of predicting the electron density.

129 The internal features of BOA are separated into node and edge features, with the bulk of the compu-
 130 tation happening in the node features. This separation avoids the high computational costs of basing
 131 the full procedure solely on edge features, while still allowing for a rich representation needed to
 132 accurately predict the density. The flow of information between node and edge features is shown in
 133 Figure 1B.

134 BOA is evaluated on electron densities generated from the widely used QM9 dataset (Jørgensen &
 135 Bhowmik, 2022; Li et al., 2025; Ruddigkeit et al., 2012; Ramakrishnan et al., 2014) and a molec-
 136 ular dynamics (MD) dataset of small organic molecules (Cheng & Peng, 2023; Bogojeski et al.,
 137 2020; Brockherde et al., 2017). On all evaluated datasets BOA outperforms previous state-of-the-art
 138 methods by a significant margin.

139 2 METHODS

140 To achieve the desired inductive bias, internal features are designed to be interpreted as functions
 141 represented in an atom-centered basis. At each point the node features are given by $h_{am\mu}$, where
 142 $a \in \mathcal{N}$ is the node with atom type Z_a , \mathcal{N} is the set of all nodes, $m \in \{1, \dots, N^c\}$ is the feature
 143 channel, and $\mu \in \{1, \dots, N_{Z_a}^B\}$ is the basis index. We use Gaussian-type ("GTO") basis functions,
 144 specifically a fully uncontracted version of def2-QZVPPD (Weigend & Ahlrichs, 2005). A correc-
 145 tion to the radial part is predicted by the model; see Section 2.9 for details. The internal features
 146 $h_{am\mu}$ are understood as expansion coefficients of N^c functions of space

$$147 h_m(\mathbf{r}) = \sum_{a \in \mathcal{N}} h_{am}(\mathbf{r}) = \sum_{a \in \mathcal{N}} \sum_{\mu} h_{am\mu} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a), \quad (1)$$

148 consisting of a sum over functions $h_{am}(\mathbf{r})$ localized at the node positions \mathbf{r}_a . Each of these functions
 149 in turn is represented in the atom type specific basis $\omega_{\mu}^{Z_a}(\mathbf{r})$. The same principle holds for the edge
 150 features $g_{abm\mu}^{(l)}$ and $g_{abm\mu}^{(r)}$, where $(a, b) \in \mathcal{E}_e$ is a directed edge in the graph, with \mathcal{E}_e being the set
 151 of all edges. Superscripts l and r denote left and right, respectively, for reasons that are clear from
 152 Equation 2. The l edge features are interpreted as functions localized at node a , while the r edge
 153 features are interpreted as functions localized at node b .

154 To guarantee rotational equivariance the choice of basis functions is essential. Choosing basis func-
 155 tions where the angular part is given by spherical harmonics ensures that the basis function coeffi-
 156 cients transform according to irreducible representations of $SO(3)$, making it possible to construct
 157 equivariant operations. BOA is fully equivariant under rotations and translations by construction,
 158 details are shown in Appendix Section A.6.

162 2.1 DENSITY REPRESENTATION
163

164 Previous work explored different ways to represent the density, either in a basis or directly on the
165 grid. We choose a representation in an atom-centered basis. In contrast to previous work, we
166 however do not expand the density (Cuevas-Zuviría & Pachos, 2021; Rackers et al., 2023; del Rio
167 et al., 2023; Qiao et al., 2022; Fu et al., 2024; Cheng & Peng, 2023) or its square root (Mitnikov &
168 Jacobson, 2024) directly as a linear expansion of the basis functions $\rho(\mathbf{r}) = \sum_{a \in \mathcal{N}} \sum_{\mu} p_{a\mu} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a)$. Instead, we choose a quadratic expansion, inspired by the natural expansion of the density in
169 squared orbital functions in KS-DFT:
170

$$171 \rho(\mathbf{r}) = \sum_{a \in \mathcal{N}} \hat{g}_a^{(l)}(\mathbf{r}) \hat{g}_a^{(r)}(\mathbf{r}) + \sum_{(a,b) \in \mathcal{E}_e} \sum_{o}^{N^o} g_{abo}^{(l)}(\mathbf{r}) g_{abo}^{(r)}(\mathbf{r}), \quad (2)$$

174 where N^o is the number of function pairs per edge and $g_{abo}^{(l)}(\mathbf{r})$, $g_{abo}^{(r)}(\mathbf{r})$, $\hat{g}_a^{(l)}(\mathbf{r})$ and $\hat{g}_a^{(r)}(\mathbf{r})$ are
175 expanded in the localized basis functions $\omega_{\mu}^{Z_a}(\mathbf{r})$ and $\omega_{\mu}^{Z_b}(\mathbf{r})$:

$$176 \hat{g}_{abo}^{(l)}(\mathbf{r}) = \sum_{\mu} g_{ab\mu}^{(l)} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a), \quad g_{abo}^{(r)}(\mathbf{r}) = \sum_{\mu} g_{ab\mu}^{(r)} \omega_{\mu}^{Z_b}(\mathbf{r} - \mathbf{r}_b), \quad (3)$$

$$179 \hat{g}_a^{(l)}(\mathbf{r}) = \sum_{\mu} \hat{g}_{a\mu}^{(l)} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a), \quad \hat{g}_a^{(r)}(\mathbf{r}) = \sum_{\mu} \hat{g}_{a\mu}^{(r)} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a). \quad (4)$$

181 This representation can be rewritten by defining $\Gamma_{ab\mu\nu} = \sum_o^{N^o} g_{ab\mu\nu}^{(l)} g_{ab\mu\nu}^{(r)} + \delta_{ab} \hat{g}_{a\mu}^{(l)} \hat{g}_{a\nu}^{(r)}$, with $\delta_{ab} =$
182 1 if $a = b$ and 0 otherwise, resulting in $\rho(\mathbf{r}) = \sum_{(a,b) \in \mathcal{E}_e} \Gamma_{ab\mu\nu} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) \omega_{\nu}^{Z_b}(\mathbf{r} - \mathbf{r}_b)$. Each $\Gamma_{ab\mu\nu}$
183 can be interpreted as a block of the full density matrix $\Gamma_{\mu\nu}$, where the indices μ and ν are understood
184 to run over all basis functions of all atoms. Using the full density matrix, the density is represented
185 by $\rho(\mathbf{r}) = \Gamma_{\mu\nu} \bar{\omega}_{\mu}(\mathbf{r}) \bar{\omega}_{\nu}(\mathbf{r})$, where $\bar{\omega}(\mathbf{r})$ is the concatenation of the basis functions of all atoms. This
186 is exactly the representation of the density used in KS-DFT calculations. The full density matrix $\Gamma_{\mu\nu}$
187 is never explicitly constructed in BOA, instead the functions $g_{abo}^{(l)}(\mathbf{r})$ and $g_{abo}^{(r)}(\mathbf{r})$ are evaluated on
188 the grid and the density is obtained using Equation 2, avoiding the costly evaluation of all pairwise
189 products of basis functions on the grid. Equation 2 amounts to evaluating a density given by a
190 low-rank representation of each block $\Gamma_{ab\mu\nu}$ of the density matrix, where the rank is given by N^o .
191

192 The expansion coefficients $g_{ab\mu\nu}^{(l)}$ and $g_{ab\mu\nu}^{(r)}$ are predicted from the molecule geometry \mathcal{M} by BOA
193 as described in the following sections. The self-loop coefficients $\hat{g}_{a\mu}^{(l)}$ and $\hat{g}_{a\mu}^{(r)}$ are used to represent
194 an initial guess for the density, which is added to the density predicted by BOA. Taking the smooth
195 absolute value of one of the functions in each pair, as described in Appendix A.8, proved beneficial.
196

197 2.2 INITIAL GUESS OF THE ELECTRON DENSITY
198

199 BOA uses an initial guess of the ground state electron density based on the atom types of the nodes.
200 After pre-training the initial guess, only an offset to that guess is learned. The learned guess amounts
201 to one pair of edge functions as described in Section 2.1 per atom. Consequently, the initial guess for
202 the density is given by $\sum_{a \in \mathcal{N}} \hat{g}_a^{(l)}(\mathbf{r}) \hat{g}_a^{(r)}(\mathbf{r})$ and is part of the predicted electron density as shown
203 in Equation 2. The initial guess coefficients are concatenated to the output of the BOA model in
204 the channel dimension before evaluating the density, resulting in the initial guess being added to the
205 density predicted by the BOA model. The initial guess coefficients are pre-trained for 1000 steps,
206 but are not fixed after that, i.e. the initial guess is refined during the full training process.
207

208 2.3 NODE AND EDGE EMBEDDINGS
209

210 The node and edge embeddings are used to set the initial node and edge features respectively. Both
211 are lookup tables, which map the atom type Z_a of the node a and the atom types of the nodes
212 connected by the edge (a, b) to a learned set of features. To keep equivariance under rotations, only
213 the coefficients corresponding to $l = 0$ basis functions are set in the embeddings, while the rest is
214 initialized to zero. The node features are therefore initialized as
215

$$h_{am\mu} = \begin{cases} W_{Z_a m \mu}^{(n)}, & \text{if } \omega_{\mu}^{Z_a}(\mathbf{r}) \text{ is an } l_{Z_a \mu} = 0 \text{ basis function} \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

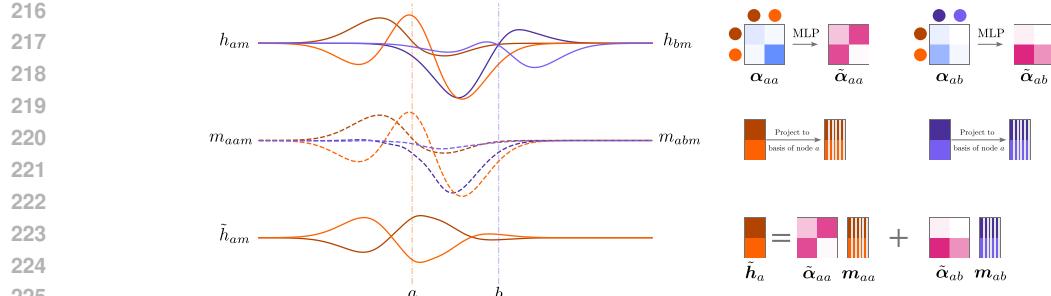


Figure 2: **1D illustration of basis overlap message passing.** Two nodes a and b are shown, each with two feature functions (blue and orange) centered at their respective positions. These functions are expanded in a local basis build from monomials multiplied with Gaussians. The first line shows the original feature functions h_{am} and h_{bm} . The second line shows the messages send from node a and node b to node a . These are the projection of the respective feature functions to the basis of node a . The third line shows the new features \tilde{h}_{am} after message passing which are a weighted sum of all incoming messages.

where $W_{Z_a m \mu}^{(n)}$ is learned during training and $l_{Z_a \mu}$ is labeling the irreducible representation of the basis function $\omega_{\mu}^{Z_a}(\mathbf{r})$. The edge features are initialized similarly as

$$g_{abm\mu}^{(l)} = \begin{cases} W_{Z_a Z_b m \mu}^{(e,l)}, & \text{if } \omega_{\mu}^{Z_a}(\mathbf{r}) \text{ is an } l_{Z_a \mu} = 0 \text{ basis function} \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

$$g_{abm\mu}^{(r)} = \begin{cases} W_{Z_a Z_b m \mu}^{(e,r)}, & \text{if } \omega_{\mu}^{Z_b}(\mathbf{r}) \text{ is an } l_{Z_b \mu} = 0 \text{ basis function} \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

where $W_{Z_a Z_b m \mu}^{(e,l)}$ and $W_{Z_a Z_b m \mu}^{(e,r)}$ are again learned.

2.4 BASIS OVERLAP MESSAGE PASSING

The fundamental principle behind the message passing mechanism introduced in this section is that, each feature channel should be interpretable as a function represented in a given basis. To respect this interpretation, a basis change is employed to transform each message from the basis of the sending node to the basis of the receiving node, a concept similar to frame-to-frame transitions employed in local canonicalization schemes (Lippmann et al., 2025). Inspired by the attention mechanism (Vaswani et al., 2017), the messages are weighted by attention weights, which are calculated from the overlap integrals between the features of the sending and receiving node. The full message passing mechanism is shown in Figure 3B and a one dimensional illustration is given in Figure 2.

BOA operates using two different cutoff radii $r_e, r_{mp} \in \mathbb{R}$, with r_e being a smaller cutoff defining the edges \mathcal{E}_e used for the edge features and in the edge update module defined in 2.8, and r_{mp} being a larger cutoff defining the edges \mathcal{E}_{mp} used in the message passing step.

The first step in transforming the features of node b into the basis of node a is calculating the overlap integrals between these features and the basis functions of node a . This overlap $o_{abm\mu}$ is given by

$$o_{abm\mu} = \int d\mathbf{r} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) h_{bm}(\mathbf{r}) = \int d\mathbf{r} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) \sum_{\nu} \omega_{\nu}^{Z_b}(\mathbf{r} - \mathbf{r}_b) h_{bm\nu} \quad (8)$$

$$= \sum_{\nu} W_{\mu\nu}^{ab} h_{bm\nu}, \quad (9)$$

where $W_{\mu\nu}^{ab} = \int d\mathbf{r} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) \omega_{\nu}^{Z_b}(\mathbf{r} - \mathbf{r}_b)$ is the overlap matrix between the basis functions of node a and node b . From these overlap integrals, the representation of the features of node b in the basis of node a can be obtained by multiplying with the inverse of the overlap matrix $(W^{aa})_{\mu\nu}^{-1}$ of

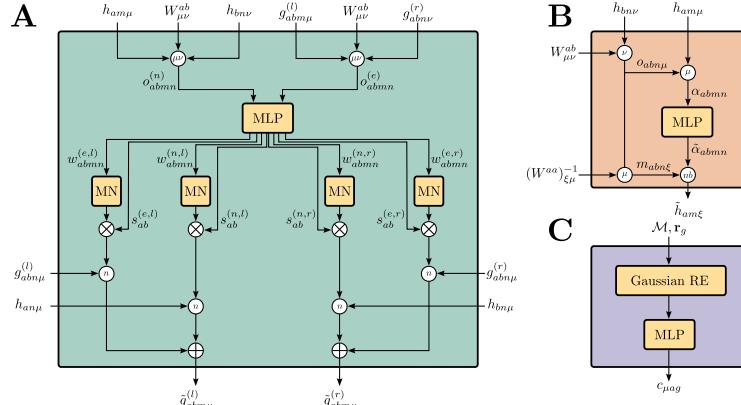
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285

Figure 3: **Subblocks of the BOA architecture.** An encircled index in the graph indicates a contraction over the corresponding dimension, i.e. multiplication of the inputs and summation over the given index. **(A)** The new edge features $\tilde{g}_{abmu}^{(l)}$ and $\tilde{g}_{abmu}^{(r)}$ are generated as a superposition of the old edge features $g_{abmu}^{(l)}$ and $g_{abmu}^{(r)}$ and the node features $h_{am\mu}$ and $h_{bn\mu}$. The weights of this superposition are generated by passing the overlap integrals of the edge and node features through an MLP. After the MLP the weights are normalized using the matrix normalization (MN) block and scaled by the predicted factors $s_{ab}^{(\cdot, \cdot)}$. The resulting weight matrices are then used to linearly mix the edge and node features. **(B)** Incoming messages are calculated by transforming the feature functions from the basis of the sending node to the basis of the receiving node. Attention weights $\tilde{\alpha}_{abmn}$ are calculated from the overlap between the feature functions of the two nodes and used to weight the messages. **(C)** A correction to the radial part of the basis functions is learned for each atom type separately by a small MLP. The radius is passed through a Gaussian radial embedding (Passaro & Zitnick, 2023) before being fed into the MLP. The output c_μ of the MLP is then used to scale the radial part of the basis functions by $1 + c_\mu$.

300
301 the basis functions of node a , resulting in the message m_{abmu} from node b to node a

$$m_{abmu} = \sum_{\nu} (W^{aa})_{\mu\nu}^{-1} o_{abmu\nu}. \quad (10)$$

302
303
304 These messages m_{abmu} are the coefficients in the basis of node a that best represent the features of
305 node b in the least squares sense. A derivation of this result can be found in Appendix A.5.

306
307 To weight the messages, an attention matrix $\tilde{\alpha}_{abmn}$ is calculated for each edge (a, b) . Each entry in
308 this attention matrix describes the weight of the message from channel n of node b to channel m of
309 node a . The attention matrix is calculated from the overlap between the feature functions of the two
310 nodes

$$\alpha_{abmn} = \int d\mathbf{r} h_{am}(\mathbf{r}) h_{bn}(\mathbf{r}) = \int d\mathbf{r} \sum_{\mu\nu} h_{am\mu} \omega_{\mu}^a(\mathbf{r} - \mathbf{r}_a) h_{bn\nu} \omega_{\nu}^b(\mathbf{r} - \mathbf{r}_b) \quad (11)$$

$$= \sum_{\mu\nu} h_{am\mu} W_{\mu\nu}^{ab} h_{bn\nu} = \sum_{\mu} h_{am\mu} o_{abn\mu}. \quad (12)$$

311
312
313 This feature overlap matrix α_{abmn} is processed by a multi-layer perceptron to produce the attention
314 weights $\tilde{\alpha}_{abmn}$, which determine the contribution of messages from neighboring nodes. These
315 weighted messages are then aggregated to update the features of node a

$$\tilde{h}_{am\mu} = \sum_{b \in \mathcal{N}_{mp}(a)} \sum_{n} \tilde{\alpha}_{abmn} m_{abn\mu}, \quad (13)$$

316
317 where $\mathcal{N}_{mp}(a) = \{b \mid (a, b) \in \mathcal{E}_{mp}\}$ is the set of neighboring nodes of node a defined by the
318 message passing edges \mathcal{E}_{mp} . Since the edges \mathcal{E}_{mp} include self-loops, the original features of each
319 node are included in the message passing step, which amounts to a residual connection weighed by
320 the attention weights $\tilde{\alpha}_{aamn}$.

324
325

2.5 NONLINEARITY

326
327
328
329
330

Searching for a suitable nonlinearity for features that should be interpretable as functions, one faces similar problems as for nonlinearities applied in equivariant neural networks. Here features are grouped together to form tensors with known transformation behavior under rotations, so that the nonlinearity can not be applied independently to each feature. In our case, the features are not only grouped to form tensors, but also form larger groups that each represent a function.

331
332
333
334

One solution in the case of equivariant networks are gated nonlinearities (Weiler et al., 2018), where first scalar features are computed to which some nonlinear function is applied. The resulting scalar features are then used to scale the tensors, i.e. each group of features forming a tensor is scaled with the same scalar. In a similar approach we first calculate scalar features

335
336
337

$$l_{amn} = \int d\mathbf{r} \int d\mathbf{r}' \frac{h_{am}(\mathbf{r})h_{an}(\mathbf{r}')}{\|\mathbf{r} - \mathbf{r}'\|} = \sum_{\mu\nu} h_{am\mu} C_{\mu\nu}^{aa} h_{an\nu}, \quad (14)$$

338
339
340
341

where $C_{\mu\nu}^{aa} = \int d\mathbf{r} \int d\mathbf{r}' \frac{\omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a)\omega_{\nu}^{Z_a}(\mathbf{r}' - \mathbf{r}_a)}{\|\mathbf{r} - \mathbf{r}'\|}$ is the Coulomb matrix, which can be generated for Gaussian-type basis functions using the PySCF package (Sun et al., 2020). The resulting scalar features l_{amn} are flattened and passed through an MLP. The result

342
343
344

$$\mathbf{w}_a = \text{MLP}_{Z_a}(\mathbf{l}_a) \quad (15)$$

is reshaped to the original shape of l_{amn} and then used to linearly transform the features of node a

345
346

$$\tilde{h}_{am\mu} = \sum_n w_{amn} h_{an\mu}. \quad (16)$$

347
348

Z_a denotes the atom type of node a , i.e. separate MLPs are learned for each atom type.

349
350

2.6 L2 NORMALIZATION

351
352
353
354
355

The normalization also aims to respect the function nature of the features, so the L2 norm of each of the per-atom channel functions is normalized. The norm is calculated as $n_{am} = \sqrt{\int d\mathbf{r} (h_{am}(\mathbf{r}))^2} = \sqrt{\sum_{\mu\nu} h_{am\mu} W_{\mu\nu}^{aa} h_{am\nu}}$ and the features are normalized by $\tilde{h}_{am\mu} = \frac{h_{am\mu}}{n_{am} + \epsilon}$ where ϵ is chosen as 10^{-6} to avoid division by zero.

356
357
358

2.7 LINEAR LAYERS

359
360
361
362
363

Standard equivariant linear layers, as implemented in the e3nn package (Geiger & Smidt, 2022), are applied to the features of each channel and node. This means that only tensors of the same type are mixed and biases are only applied to scalars to preserve equivariance. The linear layers used depend on the atom type of each node, i.e. separate parameters are learned for each atom type. After the application of these equivariant linear layers, the channels are mixed by a weight matrix, which again depends on the atom type of the node. This results in

364
365
366

$$\tilde{h}_{am\mu}^{Z_a} = \sum_{\mu} \sum_{\nu} W_{mn}^{Z_a} W_{\mu\nu}^{Z_a, (\text{eq})} h_{an\nu} \quad (17)$$

367
368
369
370
371
372
373
374

where $W_{mn}^{Z_a}$ is the weight matrix for the mixing of the channels and $W_{\mu\nu}^{Z_a, (\text{eq})}$ is the weight matrix for the equivariant linear layer. While the linear layer used here is fully equivariant, it arguably does not fully respect the function interpretation of the features. Initial experiments with linear layers that respect the function interpretation, i.e. by applying an integration $\int d\mathbf{r} w_{mn}(\mathbf{r})h_n(\mathbf{r})$ where $w_{mn}(\mathbf{r})$ are learned functions represented in some basis, showed strong instabilities in training. Since the linear layers shown here worked well in practice, we decided to use them in the final architecture, opting for stability over additional inductive bias.

375
376

2.8 EDGE UPDATE

377

The edge functions $g_{abm}^{(l)}(\mathbf{r}), g_{abm}^{(r)}(\mathbf{r})$ are updated after each BOA block. The flow of information is illustrated in Figure 1B. The node features $h_{am}(\mathbf{r})$ are used together with the old edge features

$g_{abm}^{(l)}(\mathbf{r}), g_{abm}^{(r)}(\mathbf{r})$ to generate new edge features $\tilde{g}_{abm}^{(l)}(\mathbf{r}), \tilde{g}_{abm}^{(r)}(\mathbf{r})$. There is however no flow of information from the edge features to the node features. The full edge update procedure is shown in Figure 3A. Intermediate features are generated using the overlap integrals of edge and node features

$$o_{abmn}^{(n)} = \int d\mathbf{r} h_{am}(\mathbf{r}) h_{bn}(\mathbf{r}) = \sum_{\mu\nu} h_{am\mu} W_{\mu\nu}^{ab} h_{bn\nu}, \quad (18)$$

$$o_{abmn}^{(e)} = \int d\mathbf{r} g_{abm}^{(l)}(\mathbf{r}) g_{abn}^{(r)}(\mathbf{r}) = \sum_{\mu\nu} g_{abm\mu}^{(l)} W_{\mu\nu}^{ab} g_{abn\nu}^{(r)}, \quad (19)$$

These intermediate features are flattened and passed through an MLP to obtain weights $w_{abmn}^{(n,l)}, w_{abmn}^{(n,r)}, w_{abmn}^{(e,l)}, w_{abmn}^{(e,r)}, s_{ab}^{(n,l)}, s_{ab}^{(n,r)}, s_{ab}^{(e,l)}, s_{ab}^{(e,r)}$. The weight matrices $\mathbf{w}_{ab}^{(n,l)}, \mathbf{w}_{ab}^{(n,r)}, \mathbf{w}_{ab}^{(e,l)}, \mathbf{w}_{ab}^{(e,r)}$ are normalized by

$$\tilde{\mathbf{w}}_{ab}^{(\cdot,\cdot)} = \frac{\mathbf{w}_{ab}^{(\cdot,\cdot)}}{\|\mathbf{w}_{ab}^{(\cdot,\cdot)}\|_f + \epsilon} \gamma^{(\cdot,\cdot)}, \quad (20)$$

where (\cdot, \cdot) denotes either (n, l) , (n, r) , (e, l) or (e, r) and $\gamma^{(\cdot,\cdot)}$ is a learned scalar factor. The normalization is computed in the Frobenius norm $\|\cdot\|_f$. Using these weights the new edge features

$$\tilde{g}_{abm\mu}^{(\cdot)} = \sum_n s_{ab}^{(e,\cdot)} \tilde{w}_{abmn}^{(e,\cdot)} g_{abn\mu}^{(\cdot)} + \sum_n s_{ab}^{(n,\cdot)} \tilde{w}_{abmn}^{(n,l)} h_{\star n\mu}^{(n,l)} \quad (21)$$

are generated, where $\star = a$ for the l features and $\star = b$ for the r features.

2.9 RADIAL CORRECTION FACTOR

While internally all features are represented in a basis consisting of Gaussian-type basis functions, for the final density representation an additional correction factor to the radial part of the basis functions is learned. For each atom type Z separately a small MLP is used to predict a correction factor c_μ based on the radius $r = \|\mathbf{r}\|$. As shown in Figure 3C the radius is passed through a Gaussian radial embedding (Passaro & Zitnick, 2023) before being fed into the MLP. The corrected basis functions $\tilde{\omega}_\mu^Z(\mathbf{r})$ are then given by $\tilde{\omega}_\mu^Z(\mathbf{r}) = \omega_\mu^Z(\mathbf{r}) (1 + c_\mu(r))$. The impact of this radial correction factor is evaluated in Appendix Section A.1.1.

3 EXPERIMENTS

The performance of BOA is evaluated on two electron density datasets based on QM9 (Jørgensen & Bhowmik, 2022; Li et al., 2025; Ruddigkeit et al., 2012; Ramakrishnan et al., 2014) and on a dataset based on MD trajectories (Cheng & Peng, 2023; Bogojeski et al., 2020; Brockherde et al., 2017). These datasets provide ground state electron densities for given geometries on a volumetric grid. Two versions of BOA are evaluated, a small version and a large version, differing only in the number of grid points used to evaluate the loss during training, and the batch size. The small version uses 5000 grid points and a batch size of 12 while the large version uses 6000 grid points and a batch size of 24. Using these settings, the small version fits on a single A100 GPU with \sim 40GB of memory while the large version requires an H100 GPU with \sim 94GB of memory or some setup with similar memory. All hyperparameters and training details are listed in Appendix A.7.

Additionally, the generalization capabilities of BOA are evaluated on larger molecules with up to almost 200 atoms from the QMugs dataset (Isert et al., 2022) not seen during training. Ground-truth ground state electron densities are generated using the same method as used for the QM9 PySCF dataset (Li et al., 2025), the dataset on which the evaluated models were trained.

3.1 EVALUATION ON SMALL MOLECULES

As in previous works, the predicted electron density $\tilde{\rho}(\mathbf{r})$ is compared to the reference electron density $\rho(\mathbf{r})$ using the normalized mean absolute error $\text{NMAE}(\tilde{\rho}, \rho) = \frac{\int d\mathbf{r} |\tilde{\rho}(\mathbf{r}) - \rho(\mathbf{r})|}{\int d\mathbf{r} |\rho(\mathbf{r})|}$. The integration is approximated on the full grid by summing over all grid points.

The two QM9 based datasets differ in the way the reference electron density was calculated, either using VASP (Jørgensen & Bhowmik, 2022) or PySCF (Li et al., 2025). BOA surpasses the previous state-of-the-art on both QM9 datasets, as shown in Table 1. The split between training, validation and test set is taken from Fu et al. (2024) for the QM9 VASP dataset and from Li et al. (2025) for the QM9 PySCF dataset.

Table 1: **Comparison of BOA with previous best methods on the QM9 charge density datasets.** Two datasets based on QM9 are evaluated, differing in the reference electron density calculation method. Errors are reported as NMAE [%]. For BOA the mean and standard error over three runs are reported for the small models. For the large models the mean and standard error over five runs are reported. Errors of eqDeepDFT, InfGCN, ChargE3Net, and SCDP are reproduced from Fu et al. (2024). The ResNet results are taken from Li et al. (2025).

NMAE [%]	VASP ground truth	PySCF ground truth
eqDeepDFT (Jørgensen & Bhowmik, 2022)	0.284	n/a
InfGCN (Cheng & Peng, 2023)	0.869	n/a
ChargE3Net (Koker et al., 2024)	0.196	n/a
SCDP (Fu et al., 2024)	0.178	n/a
ELECTRA (Elsborg et al., 2025)	0.177	n/a
ResNet (Li et al., 2025)	n/a	0.14
BOA small	0.1381 ± 0.0003	0.13 ± 0.01
BOA large	0.1339 ± 0.0005	0.116 ± 0.006

Additionally, the performance of BOA is evaluated on the MD dataset. As described in Cheng & Peng (2023), the MD dataset is curated from two sources. The ethanol, benzene, phenol and resorcinol data is taken from Bogojeski et al. (2020), while the ethane and malonaldehyde data is taken from Brockherde et al. (2017). BOA outperforms all previous methods on all molecules but one, where it matches the state-of-the-art, as shown in Table 2. Especially promising is the fact that no additional hyperparameter tuning was needed to achieve these results. The model is exactly the same as in the QM9 experiments, with the only change being a reduction in training steps to 200,000. Since only a training and test set is provided for the MD dataset, 10% of the training set is randomly sampled and used as a validation set to choose the best model during training.

Table 2: **Comparison of BOA with other methods on the MD charge density dataset.** Errors are reported as NMAE [%]. For BOA the mean and standard error of the mean over three runs are reported. Errors of the other models (InfGCN (Cheng & Peng, 2023), GPWNO (Kim & Ahn, 2024), SCDP (Fu et al., 2024), ELECTRA (Elsborg et al., 2025)) are reproduced from Elsborg et al. (2025).

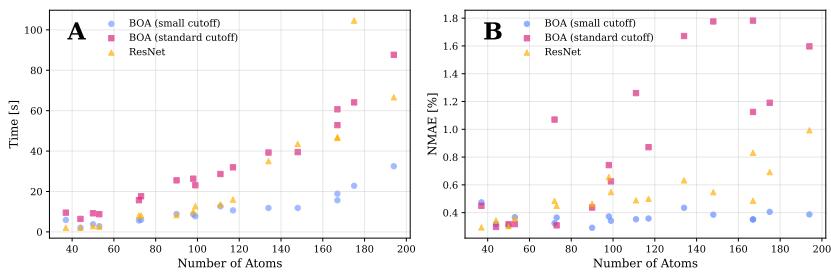
NMAE [%]	ethanol	benzene	phenol	resorcinol	ethane	malonaldehyde
InfGCN	8.43	5.11	5.51	5.95	7.01	10.34
GPWNO	4.00	2.45	2.68	2.73	3.67	5.32
SCDP	2.34	1.13	1.29	1.35	2.05	2.71
ELECTRA	1.02	0.45	0.56	0.62	0.91	0.80
BOA small	0.710 ± 0.004	0.361 ± 0.003	0.56 ± 0.03	0.371 ± 0.004	0.772 ± 0.002	0.61 ± 0.01

3.2 GENERALIZATION TO LARGER MOLECULES

To evaluate the generalization performance of BOA to larger molecules, the QMugs dataset (Isert et al., 2022) is used. Ground-truth ground state electron densities for molecules with up to almost 200 atoms are calculated using PySCF (Sun et al., 2020) with the same settings as used for the QM9 PySCF dataset (Li et al., 2025). Two BOA models trained on only the QM9 PySCF dataset are evaluated on this new test set. One of the models uses a smaller message passing cutoff of $r_{mp} = 3\text{\AA}$ instead of $r_{mp} = 6\text{\AA}$ and a smaller edge feature cutoff of $r_e = 2\text{\AA}$ instead of $r_e = 3\text{\AA}$. Additionally, we evaluate the ResNet model (Li et al., 2025) on this dataset for comparison. The time-scaling

486 behavior and the NMAE of all three models are shown in Figure 4. While the BOA model with the
 487 standard cutoffs performs worse than the ResNet model, the error of the BOA model with smaller
 488 cutoffs stays roughly constant over the evaluated molecule sizes and outperforms the ResNet model.
 489 This shows that to achieve good generalization from small to large molecules, it is beneficial to limit
 490 the field of view of the model. This may be because the contents of a large field of view differ
 491 significantly between small and large molecules, introducing a significant distribution shift, making
 492 it difficult for the model to generalize. Using a smaller field of view enables BOA to generalize well
 493 to large molecules despite being trained on only small molecules of up to 9 heavy atoms.

494 The time scaling behavior shown in Figure 4A also shows a second advantage of using smaller
 495 cutoffs. The BOA model with smaller cutoffs is significantly faster than both the ResNet model and
 496 the BOA model with larger cutoffs when evaluating larger molecules.
 497



508 **Figure 4: Generalization of BOA to larger molecules.** Two BOA models trained on the QM9
 509 PySCF data are evaluated on larger molecules from the QMugs dataset (Isert et al., 2022). One of
 510 the models uses a smaller message passing cutoff of $r_{mp} = 3\text{\AA}$ instead of $r_{mp} = 6\text{\AA}$ and a smaller
 511 edge feature cutoff of $r_e = 2\text{\AA}$ instead of $r_e = 3\text{\AA}$. (A) shows the time scaling behavior of both
 512 models compared to the ResNet model (Li et al., 2025). (B) shows the NMAE [%] of both models on
 513 the QMugs dataset over the number of atoms in the molecules. While the BOA model with the larger
 514 cutoffs performs worse than the ResNet model, the BOA model with smaller cutoffs outperforms
 515 the ResNet model with accuracy staying roughly constant over the evaluated molecule sizes.
 516

517 4 DISCUSSION

518 Electronic structure calculations plays a fundamental role in computational chemistry, with a myriad
 519 of practical applications. Accelerating these calculations has the potential of enabling new applica-
 520 tions which are currently not possible, and of cutting costs for what is already feasible. BOA takes a
 521 step towards this goal by introducing a novel architecture, permitting predictions of the electron den-
 522 sity with unprecedented accuracy. These results are enabled by a novel message passing mechanism
 523 treating the internal features as functions additional to a representation of the density in a quadratic
 524 basis expansion. The excellent generalization capabilities of BOA to larger molecules enables its
 525 application to problems were full electronic structure calculations are especially costly.
 526

527 While BOA already achieves state-of-the-art on the evaluated datasets, there are several avenues for
 528 future work. Both the QM9 and the MD dataset contain only relatively small organic molecules.
 529 Training on larger more diverse datasets will be needed to enable applicability to a wider range of
 530 practical applications. Generalization over a large part of the periodic table would be desirable,
 531 which might require changes to the architecture. Currently, BOA uses separate parameters for each
 532 atom type, which could become infeasible for more diverse datasets. A unified basis set for all atom
 533 types could be employed in the future to mitigate this problem.

534 An efficient representation of the density is crucial for basis-set-based models like BOA. Another
 535 avenue for future work is therefore to improve the basis set used. BOA currently uses a fixed basis
 536 set of uncontracted Gaussian-type basis functions as a base and learns a radial correction factor.
 537 To achieve more flexibility, also the exponents of the Gaussian basis functions could be learned,
 538 as done in (Fu et al., 2024). In BOA, the overlap and Coulomb matrices used internally could be
 539 adapted during training, which could be achieved using differentiable quantum chemistry packages
 like PySCFAD (Zhang & Chan, 2022).

540 REFERENCES
541

542 Dylan M. Anstine, Roman Zubatyuk, and Olexandr Isayev. AIMNet2: a neural network potential
543 to meet your neutral, charged, organic, and elemental-organic needs. *Chemical Science*, 16(23):
544 10228–10244, June 2025. ISSN 2041-6539. doi: 10.1039/D4SC08572H. URL <https://pubs.rsc.org/en/content/articlelanding/2025/sc/d4sc08572h>. Publisher:
545 The Royal Society of Chemistry.
546

547 Ilyes Batatia, David P. Kovacs, Gregor Simm, Christoph Ortner, and Gabor Csanyi. MACE:
548 Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force
549 Fields. *Advances in Neural Information Processing Systems*, 35:11423–11436, December
550 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/4a36c3c51af11ed9f34615b81edb5bbc-Abstract-Conference.html.
551

552 Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M. Elena, Dávid P. Kovács, Janosh Riebesell,
553 Xavier R. Advincula, Mark Asta, Matthew Avaylon, William J. Baldwin, Fabian Berger, Noam
554 Bernstein, Arghya Bhowmik, Filippo Bigi, Samuel M. Blau, Vlad Cărăre, Michele Ceriotti, Sang-
555 gyu Chong, James P. Darby, Sandip De, Flaviano Della Pia, Volker L. Deringer, Rokas Elijošius,
556 Zakariya El-Machachi, Fabio Falconni, Edvin Fako, Andrea C. Ferrari, John L. A. Gardner,
557 Mikolaj J. Gawkowski, Annalena Genreith-Schriever, Janine George, Rhys E. A. Goodall, Jonas
558 Grandel, Clare P. Grey, Petr Grigorev, Shuang Han, Will Handley, Hendrik H. Heenen, Kersti
559 Hermansson, Christian Holm, Cheuk Hin Ho, Stephan Hofmann, Jad Jaafar, Konstantin S. Jakob,
560 Hyunwook Jung, Venkat Kapil, Aaron D. Kaplan, Nima Karimitari, James R. Kermode, Panagiotis
561 Kourtis, Namu Kroupa, Jolla Kullgren, Matthew C. Kuner, Domantas Kuryla, Guoda Liepuoniute,
562 Chen Lin, Johannes T. Margraf, Ioan-Bogdan Magdău, Angelos Michaelides, J. Harry
563 Moore, Aakash A. Naik, Samuel P. Niblett, Sam Walton Norwood, Niamh O'Neill, Christoph
564 Ortner, Kristin A. Persson, Karsten Reuter, Andrew S. Rosen, Louise A. M. Rosset, Lars L.
565 Schaaf, Christoph Schran, Benjamin X. Shi, Eric Sivonxay, Tamás K. Stenczel, Viktor Svahn,
566 Christopher Sutton, Thomas D. Swinburne, Jules Tilly, Cas van der Oord, Santiago Vargas, Eszter
567 Varga-Umbrich, Tejs Vegge, Martin Vondrák, Yangshuai Wang, William C. Witt, Thomas Wolf,
568 Fabian Zills, and Gábor Csányi. A foundation model for atomistic materials chemistry, September
569 2025. URL <http://arxiv.org/abs/2401.00096>. arXiv:2401.00096 [physics].
570

571 Mihail Bogojeski, Leslie Vogt-Maranto, Mark E. Tuckerman, Klaus-Robert Müller, and Kieron
572 Burke. Quantum chemical accuracy from density functional approximations via ma-
573 chine learning. *Nature Communications*, 11(1):5223, October 2020. ISSN 2041-1723.
574 doi: 10.1038/s41467-020-19093-1. URL <https://www.nature.com/articles/s41467-020-19093-1>. Publisher: Nature Publishing Group.
575

576 Felix Brockherde, Leslie Vogt, Li Li, Mark E. Tuckerman, Kieron Burke, and Klaus-Robert Müller.
577 Bypassing the Kohn-Sham equations with machine learning. *Nature Communications*, 8(1):872,
578 October 2017. ISSN 2041-1723. doi: 10.1038/s41467-017-00839-3. URL <https://www.nature.com/articles/s41467-017-00839-3>. Publisher: Nature Publishing Group.
579

580 Chaoran Cheng and Jian Peng. Equivariant Neural Operator Learning with Graphon Convolu-
581 tion. November 2023. URL <https://openreview.net/forum?id=EjiA3uWpnc¬eId=fqtWXxtsdg>.
582

583 Bruno Cuevas-Zuviría and Luis F. Pacios. Machine Learning of Analytical Electron Density in
584 Large Molecules Through Message-Passing. *Journal of Chemical Information and Modeling*, 61
585 (6):2658–2666, June 2021. ISSN 1549-9596. doi: 10.1021/acs.jcim.1c00227. URL <https://doi.org/10.1021/acs.jcim.1c00227>. Publisher: American Chemical Society.
586

587 Beatriz G. del Rio, Brandon Phan, and Rampi Ramprasad. A deep learning framework to emulate
588 density functional theory. *npj Computational Materials*, 9(1):158, August 2023. ISSN 2057-
589 3960. doi: 10.1038/s41524-023-01115-3. URL <https://www.nature.com/articles/s41524-023-01115-3>. Publisher: Nature Publishing Group.
590

591 Jonas Elsborg, Luca Thiede, Alán Aspuru-Guzik, Tejs Vegge, and Arghya Bhowmik. ELECTRA:
592 A Cartesian Network for 3D Charge Density Prediction with Floating Orbitals, May 2025. URL
593 <http://arxiv.org/abs/2503.08305>. arXiv:2503.08305 [cs].

594 Pol Febrer, Peter Bjørn Jørgensen, Miguel Pruneda, Alberto García, Pablo Ordejón, and Arghya
 595 Bhowmik. Graph2Mat: universal graph to matrix conversion for electron density prediction.
 596 *Machine Learning: Science and Technology*, 6(2):025013, April 2025. ISSN 2632-2153. doi: 10.
 597 1088/2632-2153/adc871. URL <https://dx.doi.org/10.1088/2632-2153/adc871>.
 598 Publisher: IOP Publishing.

599 Xiang Fu, Andrew Scott Rosen, Kyle Bystrom, Rui Wang, Albert Musaelian, Boris Kozinsky, Tess
 600 Smidt, and Tommi Jaakkola. A Recipe for Charge Density Prediction. November 2024. URL
 601 <https://openreview.net/forum?id=b7REKaNUTv>.

602 Mario Geiger and Tess Smidt. e3nn: Euclidean Neural Networks, July 2022. URL <http://arxiv.org/abs/2207.09453> [cs].

603 Sheng Gong, Tian Xie, Taishan Zhu, Shuo Wang, Eric R. Fadel, Yawei Li, and Jeffrey C.
 604 Grossman. Predicting charge density distribution of materials using a local-environment-
 605 based graph convolutional network. *Physical Review B*, 100(18):184103, November 2019.
 606 doi: 10.1103/PhysRevB.100.184103. URL <https://link.aps.org/doi/10.1103/PhysRevB.100.184103>. Publisher: American Physical Society.

607 P. Hohenberg and W. Kohn. Inhomogeneous Electron Gas. *Physical Review*, 136(3B):B864–B871,
 608 November 1964. doi: 10.1103/PhysRev.136.B864. URL <https://link.aps.org/doi/10.1103/PhysRev.136.B864>. Publisher: American Physical Society.

609 Clemens Isert, Kenneth Atz, José Jiménez-Luna, and Gisbert Schneider. QMugs, quantum me-
 610 chanical properties of drug-like molecules. *Scientific Data*, 9(1):273, June 2022. ISSN 2052-
 611 4463. doi: 10.1038/s41597-022-01390-7. URL <https://www.nature.com/articles/s41597-022-01390-7>. Publisher: Nature Publishing Group.

612 Peter Bjørn Jørgensen and Arghya Bhowmik. Equivariant graph neural networks for fast electron
 613 density estimation of molecules, liquids, and solids. *npj Computational Materials*, 8(1):183,
 614 August 2022. ISSN 2057-3960. doi: 10.1038/s41524-022-00863-y. URL <https://www.nature.com/articles/s41524-022-00863-y>. Publisher: Nature Publishing Group.

615 Seongsu Kim and Sungsoo Ahn. Gaussian Plane-Wave Neural Operator for Electron Density Es-
 616 timation. June 2024. URL <https://openreview.net/forum?id=TvoG41N1Y3>.

617 Tsz Wai Ko, Jonas A. Finkler, Stefan Goedecker, and Jörg Behler. Accurate Fourth-Generation
 618 Machine Learning Potentials by Electrostatic Embedding. *Journal of Chemical Theory and Com-
 619 putation*, 19(12):3567–3579, June 2023. ISSN 1549-9618. doi: 10.1021/acs.jctc.2c01146. URL
 620 <https://doi.org/10.1021/acs.jctc.2c01146>. Publisher: American Chemical So-
 621 ciety.

622 Teddy Koker, Keegan Quigley, Eric Taw, Kevin Tibbetts, and Lin Li. Higher-order equivariant
 623 neural networks for charge density prediction in materials. *npj Computational Materials*, 10(1):
 624 161, July 2024. ISSN 2057-3960. doi: 10.1038/s41524-024-01343-1. URL <https://www.nature.com/articles/s41524-024-01343-1>. Publisher: Nature Publishing Group.

625 Boris Kozinsky, Albert Musaelian, Anders Johansson, and Simon Batzner. Scaling the Lead-
 626 ing Accuracy of Deep Equivariant Models to Biomolecular Simulations of Realistic Size. In
 627 *Proceedings of the International Conference for High Performance Computing, Networking,
 628 Storage and Analysis*, SC '23, pp. 1–12, New York, NY, USA, November 2023. Association
 629 for Computing Machinery. ISBN 979-8-4007-0109-2. doi: 10.1145/3581784.3627041. URL
 630 <https://dl.acm.org/doi/10.1145/3581784.3627041>.

631 Chenghan Li, Or Sharir, Shunyue Yuan, and Garnet Kin-Lic Chan. Image super-resolution in-
 632 spired electron density prediction. *Nature Communications*, 16(1):4811, May 2025. ISSN 2041-
 633 1723. doi: 10.1038/s41467-025-60095-8. URL <https://www.nature.com/articles/s41467-025-60095-8>. Publisher: Nature Publishing Group.

634 Yi-Lun Liao, Brandon M. Wood, Abhishek Das, and Tess Smidt. EquiformerV2: Improved
 635 Equivariant Transformer for Scaling to Higher-Degree Representations. October 2023. URL
 636 <https://openreview.net/forum?id=mCOBKZmrzD>.

648 Peter Lippmann, Gerrit Gerhartz, Roman Remme, and Fred A Hamprecht. Beyond
 649 Canonicalization: How Tensorial Messages Improve Equivariant Message Pass-
 650 ing. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Con-
 651 ference on Representation Learning*, volume 2025, pp. 88067–88087, 2025. URL
 652 [https://proceedings.iclr.cc/paper_files/paper/2025/file/
 653 db7534a06ace69f4ec95bc89e91d5dbb-Paper-Conference.pdf](https://proceedings.iclr.cc/paper_files/paper/2025/file/db7534a06ace69f4ec95bc89e91d5dbb-Paper-Conference.pdf).

654 Erpai Luo, Xinran Wei, Lin Huang, Yunyang Li, Han Yang, Zaishuo Xia, Zun Wang, Chang Liu, Bin
 655 Shao, and Jia Zhang. Efficient and Scalable Density Functional Theory Hamiltonian Prediction
 656 through Adaptive Sparsity. June 2025. URL <https://openreview.net/forum?id=K3lykWhXON¬eId=k94UW6cbpQ>.

657 Ilan Mitnikov and Joseph Jacobson. E3STO: Orbital Inspired SE(3)-Equivariant Molecular Repre-
 658 sentation for Electron Density Prediction, October 2024. URL <http://arxiv.org/abs/2410.06119> [physics].

659 Saro Passaro and C. Lawrence Zitnick. Reducing SO(3) Convolutions to SO(2) for Efficient
 660 Equivariant GNNs. In *Proceedings of the 40th International Conference on Machine Learning*,
 661 pp. 27420–27438. PMLR, July 2023. URL [https://proceedings.mlr.press/v202/
 662 passaro23a.html](https://proceedings.mlr.press/v202/passaro23a.html). ISSN: 2640-3498.

663 Chen Qian, Valdas Vitartas, James Kermode, and Reinhard J. Maurer. Equivariant Electronic Hamil-
 664 tonian Prediction with Many-Body Message Passing, August 2025. URL <http://arxiv.org/abs/2508.15108> [cond-mat].

665 Zhuoran Qiao, Matthew Welborn, Animashree Anandkumar, Frederick R. Manby, and Thomas F.
 666 Miller, III. OrbNet: Deep learning for quantum chemistry using symmetry-adapted atomic-orbital
 667 features. *The Journal of Chemical Physics*, 153(12):124111, September 2020. ISSN 0021-9606.
 668 doi: 10.1063/5.0021955. URL <https://doi.org/10.1063/5.0021955>.

669 Zhuoran Qiao, Anders S. Christensen, Matthew Welborn, Frederick R. Manby, Anima Anandkumar,
 670 and Thomas F. Miller. Informing geometric deep learning with electronic interactions to acceler-
 671 ate quantum chemistry. *Proceedings of the National Academy of Sciences*, 119(31):e2205221119,
 672 2022. doi: 10.1073/pnas.2205221119. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2205221119>. eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.2205221119>.

673 Joshua A Rackers, Lucas Tecot, Mario Geiger, and Tess E Smidt. A recipe for cracking the quantum
 674 scaling limit with machine learned electron densities. *Machine Learning: Science and Technol-
 675 ogy*, 4(1):015027, February 2023. ISSN 2632-2153. doi: 10.1088/2632-2153/acb314. URL
 676 <https://dx.doi.org/10.1088/2632-2153/acb314>. Publisher: IOP Publishing.

677 Raghunathan Ramakrishnan, Pavlo O. Dral, Matthias Rupp, and O. Anatole von Lilienfeld. Quan-
 678 tum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1(1):140022,
 679 August 2014. ISSN 2052-4463. doi: 10.1038/sdata.2014.22. URL <https://www.nature.com/articles/sdata201422>. Publisher: Nature Publishing Group.

680 Roman Remme, Tobias Kaczun, Tim Ebert, Christof A. Gehrig, Dominik Geng, Gerrit Gerhartz,
 681 Marc K. Ickler, Manuel V. Klockow, Peter Lippmann, Johannes S. Schmidt, Simon Wagner,
 682 Andreas Dreuw, and Fred A. Hamprecht. Stable and Accurate Orbital-Free Density Func-
 683 tional Theory Powered by Machine Learning. *Journal of the American Chemical Society*,
 684 147(32):28851–28859, August 2025. ISSN 0002-7863. doi: 10.1021/jacs.5c06219. URL
 685 <https://doi.org/10.1021/jacs.5c06219>. Publisher: American Chemical Society.

686 Lars Ruddigkeit, Ruud van Deursen, Lorenz C. Blum, and Jean-Louis Reymond. Enumeration of
 687 166 Billion Organic Small Molecules in the Chemical Universe Database GDB-17. *Journal of
 688 Chemical Information and Modeling*, 52(11):2864–2875, November 2012. ISSN 1549-9596. doi:
 689 10.1021/ci300415d. URL <https://doi.org/10.1021/ci300415d>. Publisher: Ameri-
 690 can Chemical Society.

691 Guillem Simeon and Gianni De Fabritiis. TensorNet: Cartesian Tensor Representations for Ef-
 692 ficient Learning of Molecular Potentials. November 2023. URL <https://openreview.net/forum?id=BEH1PdBZ2e>.

702 Feitong Song and Ji Feng. NeuralSCF: Neural network self-consistent fields for density func-
 703 tional theory, June 2024. URL <http://arxiv.org/abs/2406.15873>. arXiv:2406.15873
 704 [physics].

705 Qiming Sun, Xing Zhang, Samragni Banerjee, Peng Bao, Marc Barbry, Nick S. Blunt, Nikolay A.
 706 Bogdanov, George H. Booth, Jia Chen, Zhi-Hao Cui, Janus J. Eriksen, Yang Gao, Sheng Guo, Jan
 707 Hermann, Matthew R. Hermes, Kevin Koh, Peter Koval, Susi Lehtola, Zhendong Li, Junzi Liu,
 708 Narbe Mardirossian, James D. McClain, Mario Motta, Bastien Mussard, Hung Q. Pham, Artem
 709 Pulkin, Wirawan Purwanto, Paul J. Robinson, Enrico Ronca, Elvira R. Sayfutyarova, Maximil-
 710 ian Scheurer, Henry F. Schurkus, James E. T. Smith, Chong Sun, Shi-Ning Sun, Shiv Upadhyay,
 711 Lucas K. Wagner, Xiao Wang, Alec White, James Daniel Whitfield, Mark J. Williamson, Sebas-
 712 tian Wouters, Jun Yang, Jason M. Yu, Tianyu Zhu, Timothy C. Berkelbach, Sandeep Sharma,
 713 Alexander Yu. Sokolov, and Garnet Kin-Lic Chan. Recent developments in the PySCF program
 714 package. *The Journal of Chemical Physics*, 153(2):024109, July 2020. ISSN 0021-9606. doi:
 715 10.1063/5.0006074. URL <https://doi.org/10.1063/5.0006074>.

716 Ethan M. Sunshine, Muhammed Shuaibi, Zachary W. Ulissi, and John R. Kitchin. Chemical Proper-
 717 ties from Graph Neural Network-Predicted Electron Densities. *The Journal of Physical Chemistry*
 718 C, 127(48):23459–23466, December 2023. ISSN 1932-7447. doi: 10.1021/acs.jpcc.3c06157.
 719 URL <https://doi.org/10.1021/acs.jpcc.3c06157>. Publisher: American Chemi-
 720 cal Society.

721 Richard Van Noorden. These are the most-cited research papers of all time. *Nature*, 640(8059):591,
 722 2025.

723 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N
 724 Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates,
 725 Inc., 2017. URL https://papers.nips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fdbd053c1c4a845aa-Abstract.html.

726 Zun Wang, Chang Liu, Nianlong Zou, He Zhang, Xinran Wei, Lin Huang, Lijun Wu, and Bin
 727 Shao. Infusing Self-Consistency into Density Functional Theory Hamiltonian Prediction via
 728 Deep Equilibrium Models, October 2024. URL <http://arxiv.org/abs/2406.03794>.
 arXiv:2406.03794 [cs].

729 Florian Weigend and Reinhart Ahlrichs. Balanced basis sets of split valence, triple zeta valence
 730 and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy. *Physical*
 731 *Chemistry Chemical Physics*, 7(18):3297–3305, August 2005. ISSN 1463-9084. doi: 10.1039/
 732 B508541A. URL <https://pubs.rsc.org/en/content/articlelanding/2005/cp/b508541a>. Publisher: The Royal Society of Chemistry.

733 Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco S Cohen. 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric
 734 Data. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates,
 735 Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/488e4104520c6aab692863cc1dba45af-Abstract.html>.

736 Brandon M. Wood, Misko Dzamba, Xiang Fu, Meng Gao, Muhammed Shuaibi, Luis Barroso-
 737 Luque, Kareem Abdelmaqsoud, Vahe Gharakhanyan, John R. Kitchin, Daniel S. Levine, Kyle
 738 Michel, Anuroop Sriram, Taco Cohen, Abhishek Das, Ammar Rizvi, Sushree Jagriti Sahoo,
 739 Zachary W. Ulissi, and C. Lawrence Zitnick. UMA: A Family of Universal Models for Atoms,
 740 June 2025. URL <http://arxiv.org/abs/2506.23971>. arXiv:2506.23971 [cs].

741 Zilong Yuan, Zechen Tang, Honggeng Tao, Xiaoxun Gong, Zezhou Chen, Yuxiang Wang, He Li,
 742 Yang Li, Zhiming Xu, Minghui Sun, Boheng Zhao, Chong Wang, Wenhui Duan, and Yong Xu.
 743 Deep learning density functional theory Hamiltonian in real space, July 2024. URL <http://arxiv.org/abs/2407.14379>. arXiv:2407.14379 [physics].

744 He Zhang, Chang Liu, Zun Wang, Xinran Wei, Siyuan Liu, Nanning Zheng, Bin Shao, and Tie-
 745 Yan Liu. Self-Consistency Training for Density-Functional-Theory Hamiltonian Prediction. June
 746 2024a. URL <https://openreview.net/forum?id=Vw4Yar2fmW>.

756 He Zhang, Siyuan Liu, Jiacheng You, Chang Liu, Shuxin Zheng, Ziheng Lu, Tong Wang, Nanning
757 Zheng, and Bin Shao. Overcoming the barrier of orbital-free density functional theory for molec-
758 ular systems using deep learning. *Nature Computational Science*, 4(3):210–223, March 2024b.
759 ISSN 2662-8457. doi: 10.1038/s43588-024-00605-8. URL <https://www.nature.com/articles/s43588-024-00605-8>. Publisher: Nature Publishing Group.

760
761 Xing Zhang and Garnet Kin-Lic Chan. Differentiable quantum chemistry with PySCF for molecules
762 and materials at the mean-field level and beyond. *The Journal of Chemical Physics*, 157(20):
763 204801, November 2022. ISSN 0021-9606. doi: 10.1063/5.0118200. URL <https://doi.org/10.1063/5.0118200>.
764
765

766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

810 A APPENDIX
811812 A.1 ABLATION STUDIES
813814 A.1.1 RADIAL CORRECTION AND SMOOTH ABSOLUTE VALUE
815

816 An ablation study of the radial correction factor and taking the smooth absolute value of one of the
817 basis functions in each pair in the expansion of the density is performed. The results are shown
818 in Table 3. For each configuration, three training runs are performed and the mean and standard
819 error are reported. Both the radial correction and taking the absolute value significantly improve the
820 performance, with the radial correction having the larger effect. The best performance is achieved
821 when both techniques are used together.

822
823 **Table 3: Ablation study of taking the absolute value of one of the basis functions in the pair**
824 **and applying the radial correction.** The small BOA version is trained on the QM9 VASP dataset.
825 Errors are reported as NMAE [%].

	Absolute Value	Radial Correction	NMAE [%]
	✗	✗	0.204 ± 0.002
	✓	✗	0.167 ± 0.004
	✗	✓	0.1423 ± 0.0006
	✓	✓	0.1381 ± 0.0003

833
834 A.1.2 CHOICE OF BASIS SET
835

836 The impact of the choice of basis sets is studied on two additional smaller basis sets, `def2-SVP`
837 and `def2-TZVP` (Weigend & Ahlrichs, 2005). All hyperparameters other than the used basis are
838 chosen as in BOA small and the models are trained on the QM9 VASP data. The results are shown
839 in Table 4. Smaller basis sets still perform reasonably well, with the medium size basis `def2-`
840 `TZVP` still surpassing the previous state of the art. Still, using larger basis sets improves the NMAE
841 significantly.

842
843
844 **Table 4: Basis set ablation.** All settings other than the basis set are kept constant to the small BOA
845 version and the models are trained on the QM9 VASP dataset. Three training runs are performed for
846 each configuration and the mean and standard error are reported.

	NMAE [%]	Number of basis functions
<code>def2-SVP</code>	0.194 ± 0.001	103
<code>def2-TZVP</code>	0.1504 ± 0.0003	192
<code>def2-QZVPPD</code>	0.1381 ± 0.0003	374

853
854 A.1.3 QUADRATIC EXPANSION
855

856 To study the impact of the quadratic expansion we train additional models using the same hy-
857 perparameters as BOA small. When evaluating the density on the grid we drop $\hat{g}_a^{(r)}(\mathbf{r})$ and
858 $g_{abo}^{(r)}(\mathbf{r})$ and expand the density linearly as $\rho(\mathbf{r}) = \sum_{a \in \mathcal{N}} \hat{g}_a^{(l)}(\mathbf{r}) + \sum_{(a,b) \in \mathcal{E}_c} \sum_o^{N^o} g_{abo}^{(l)}(\mathbf{r}) =$
859 $\sum_{a \in \mathcal{N}} \sum_{\mu} \hat{g}_{a\mu}^{(l)} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) + \sum_{(a,b) \in \mathcal{E}_c} \sum_o^{N^o} \sum_{\mu} g_{ab\mu}^{(l)} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a)$ where the sum over b can
860 be carried out before evaluating the basis functions $\omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a)$ on the grid. Three models were
861 trained on the QM9 VASP data using this linear expansion and evaluated on the test set. The re-
862 sults are shown in Table 5. The quadratic expansion significantly outperforms the linear expansion,
863 showing the importance of the quadratic expansion for accurate density predictions in BOA.

864
 865 **Table 5: Quadratic expansion ablation.** All settings other than the expansion type are kept constant
 866 to the small BOA version and the models are trained on the QM9 VASP dataset. Three training runs
 867 are performed for each configuration and the mean and standard error are reported.

NMAE [%]	
Linear Expansion	0.2716 ± 0.0007
Quadratic Expansion	0.1381 ± 0.0003

874 A.1.4 RADIAL CUTOFF 875

876 We conducted an additional experiment with larger cutoffs, specifically a message passing cutoff of
 877 $r_{mp} = 8\text{\AA}$ (instead of $r_{mp} = 6\text{\AA}$) and an edge feature cutoff of $r_e = 4\text{\AA}$ (instead of $r_e = 3\text{\AA}$) were
 878 chosen. Other than the cutoffs the BOA Small settings are used and three models are trained on
 879 the QM9 VASP data. Mean and standard error on the test set are reported in Table 6. Using larger
 880 cutoffs results in a small but consistent improvement in performance.

881
 882 **Table 6: Radial cutoff ablation.** All settings other than the cutoffs are kept constant to the small
 883 BOA version and the models are trained on the QM9 VASP dataset. Three training runs are per-
 884 formed for each configuration and the mean and standard error are reported.

NMAE [%]	
Standard cutoffs	0.1381 ± 0.0003
Larger cutoffs	0.1343 ± 0.0007

893 A.2 DFT INITIALIZATION AND COULOMB ENERGY ERROR 894

895 To evaluate the quality of the predicted electron densities on metrics closer to practical applications,
 896 we calculate the Coulomb energy error and investigate the savings in KS-DFT calculations when us-
 897 ing the predicted densities as initialization. Both of these evaluations are performed on the full QM9
 898 PySCF test set of 10000 molecules using a BOA large model. BOA achieves a mean absolute error
 899 of 66 meV between the Coulomb energy calculated on the predicted densities and the converged
 900 label densities, comparing favorably to the 167 meV reported for the ResNet model (Li et al., 2025).

901 KS-DFT calculations using our densities as initial guesses and using a standard superposition of
 902 atomic densities as initial guesses are performed. The same settings as in the data generation de-
 903 scribed by Li et al. (2025) are used. With these settings the BOA initialization reduces the mean
 904 number of SCF iterations needed from 15.7 to 10.2, a reduction by 35%. This is a result competitive
 905 with the 35% reduction reported by Li et al. (2025) on the same data.

906 907 A.3 EFFICIENCY 908

909 We evaluate the efficiency of ELECTRA [2], SCDP [1], standard BOA, and BOA with `def2-TZVP`
 910 basis models. We note that there is no difference in inference efficiency between the small and large
 911 settings, since only the number of probe points and the batch size during training are changed, and
 912 we therefore do not need to differentiate between them here. All models were evaluated on a 40GB
 913 A100 GPU using a block size maximizing the VRAM usage. The time per molecule is evaluated
 914 over the whole QM9 Vasp test set and the mean and standard deviation over these molecules are
 915 reported in Table 7. While BOA is roughly 2 times slower than SCDP, using the smaller `def2-`
 916 `TZVP` mitigates this effect to some extend and achieves a smaller error than the previous state of
 917 the art with efficiency comparable to SCDP. The efficiency/accuracy trade-off is easily steerable in
 BOA by employing basis sets of different size.

918
919
920
921
922

Table 7: Time Efficiency. The time per molecule needed during inference is evaluated on the QM9 VASP test set for ELECTRA (Elsborg et al., 2025), SCDP (Fu et al., 2024), standard BOA and BOA with `def2-TZVP` basis. The mean and standard deviation over all molecules in the test set are reported.

923
924
925
926
927
928
929
930

Model	Time per molecule [s]
ELECTRA	0.14 ± 0.03
SCDP	0.58 ± 0.16
BOA	1.27 ± 0.27
BOA <code>def2-TZVP</code>	0.64 ± 0.15

A.4 ERROR DISTRIBUTION

The distribution of the electron density prediction errors on QM9 VASP is analyzed in more detail. Figures 5A and 5B show 2D histograms of the absolute errors of the predicted density for BOA large and SCDP (Fu et al., 2024), binned by the distance to the nearest and second-nearest atom. In each bin the absolute errors are summed. The errors are evaluated on 100 randomly sampled molecules and normalized by the total error of the SCDP model over all evaluated grid points. Additionally, Figure 5C shows the label electron density distribution binned in the same way and normalized by the total density over all evaluated grid points. Both models show the largest errors in the regions with high electron density. In these critical regions with the nearest atom at a distance of less than approximately 1.0 Å and the second-nearest atom at a distance of 0.5 Å to 1.5 Å, the BOA model, however, shows significantly smaller errors than SCDP.

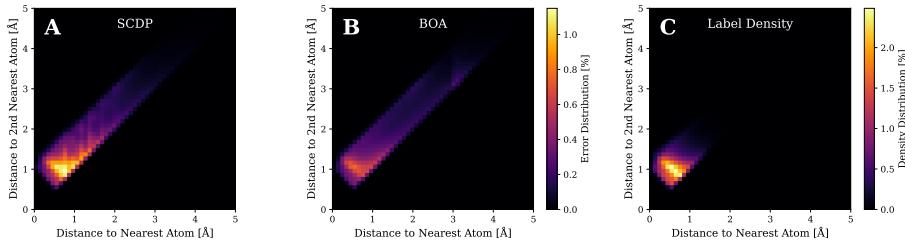


Figure 5: Error distribution analysis. The distribution of the absolute errors of the predicted electron density on the QM9 VASP dataset is shown as a function of the distance to the nearest and second-nearest atom. The errors are evaluated on 100 randomly sampled molecules and normalized by the total error of the SCDP model (Fu et al., 2024). (A) shows the error distribution of BOA large, (B) shows the error distribution of SCDP, and (C) shows the label electron density distribution, normalized by the total density over all evaluated grid points. Both models show the largest errors in the regions with high electron density, but BOA shows significantly smaller errors in these critical regions.

A.5 FUNCTION MESSAGE PASSING DERIVATION

We show that the messages $m_{abm\mu}$ are the coefficients in the basis of node a that represent the feature functions $h_{bn}(\mathbf{r})$ of node b best in the least squares sense. The message passing step therefore solves the optimization problem

$$\min_{m_{abm\mu}} \left\| h_{bm}(\mathbf{r}) - \sum_{\mu} m_{abm\mu} \omega_{\mu}^{Z_a} (\mathbf{r} - \mathbf{r}_a) \right\|^2 \quad (22)$$

$$\Rightarrow \min_{m_{abm\mu}} \left\| \sum_{\mu} h_{bm\mu} \omega_{\mu}^{Z_b} (\mathbf{r} - \mathbf{r}_b) - \sum_{\mu} m_{abm\mu} \omega_{\mu}^{Z_a} (\mathbf{r} - \mathbf{r}_a) \right\|^2. \quad (23)$$

972 Taking the derivative with respect to $m_{abm\mu}$ and setting it to zero results in
 973

$$974 \quad 0 = 2 \left(\sum_{\mu} h_{bm\mu} \omega_{\mu}^{Z_b}(\mathbf{r} - \mathbf{r}_b) - \sum_{\mu} m_{abm\mu} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) \right) \quad (24)$$

$$977 \quad \Rightarrow \sum_{\mu} h_{bm\mu} \omega_{\mu}^{Z_b}(\mathbf{r} - \mathbf{r}_b) = \sum_{\mu} m_{abm\mu} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a). \quad (25)$$

980 Multiplying both sides with $\omega_{\nu}^{Z_a}(\mathbf{r} - \mathbf{r}_a)$ and integrating over \mathbf{r} results in
 981

$$982 \quad \int d\mathbf{r} \sum_{\mu} h_{bm\mu} \omega_{\mu}^{Z_b}(\mathbf{r} - \mathbf{r}_b) \omega_{\nu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) = \int d\mathbf{r} \sum_{\mu} m_{abm\mu} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) \omega_{\nu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) \quad (26)$$

$$985 \quad \Rightarrow \sum_{\mu} W_{\nu\mu}^{ab} h_{bm\mu} = \sum_{\mu} W_{\nu\mu}^{aa} m_{abm\mu} \quad (27)$$

$$987 \quad \Rightarrow m_{abm\mu} = \sum_{\nu} (W^{aa})_{\mu\nu}^{-1} \sum_{\kappa} W_{\nu\kappa}^{ab} h_{bm\kappa}, \quad (28)$$

989 which is the solution to the optimization problem in Equation 23 and exactly what is implemented
 990 in the message passing step.
 991

993 A.6 EQUIVARIANCE OF BOA

995 To understand the equivariance of the different BOA layers, it is necessary to understand how the
 996 basis functions $\omega_{\mu}^Z(\mathbf{r})$ transform under rotations. The angular part of the basis functions is given by
 997 spherical harmonics, which are known to transform under irreducible representations of the rotation
 998 group $\text{SO}(3)$. These irreducible representations are indexed by l with dimension $2l + 1$. The basis
 999 functions can now be grouped according to their l value, i.e. there are groups of $2l + 1$ basis functions
 1000 that transform under the irreducible representation indexed by l . A basis function $\omega_m^l(\mathbf{r})$ in the group
 1001 indexed by l transforms under a rotation $\mathbf{R} \in \text{SO}(3)$ as

$$1002 \quad \omega_m^l(\mathbf{R}\mathbf{r}) = \sum_{m'} D_{mm'}^l(\mathbf{R}) \omega_{m'}^l(\mathbf{r}), \quad (29)$$

1005 where $D_{mm'}^l(\mathbf{R})$ is the Wigner D-matrix. A set of basis functions $\omega_{\mu}^Z(\mathbf{r})$ can therefore be separated
 1006 into irreducible representations

$$1008 \quad \omega^Z(\mathbf{r}) = \left(\underbrace{\omega_{-l_1}^{l_1}(\mathbf{r}), \dots, \omega_{l_1}^{l_1}(\mathbf{r}), \omega_{-l_2}^{l_2}(\mathbf{r}), \dots, \omega_{l_2}^{l_2}(\mathbf{r}), \dots}_{2l_1 + 1} \underbrace{\omega_{-l_2}^{l_2}(\mathbf{r}), \dots, \omega_{l_2}^{l_2}(\mathbf{r}), \dots}_{2l_2 + 1} \right)^T, \quad (30)$$

1012 resulting in a block diagonal transformation behavior under rotations
 1013

$$1014 \quad \omega^Z(\mathbf{R}\mathbf{r}) = \mathbf{D}^Z(\mathbf{R}) \omega^Z(\mathbf{r}), \quad (31)$$

1016 where $\mathbf{D}^Z(\mathbf{R})$ is a block diagonal matrix with blocks given by the Wigner D-matrices corresponding
 1017 to the l values of the basis functions. The Wigner D-matrices are orthogonal matrices, so that $\mathbf{D}^Z(\mathbf{R})$
 1018 is also an orthogonal matrix, i.e. $(\mathbf{D}^Z(\mathbf{R}))^T = (\mathbf{D}^Z(\mathbf{R}))^{-1}$.

1019 BOA employs the overlap and Coulomb matrices between the basis functions of different nodes,
 1020 which are defined as
 1021

$$1022 \quad W_{\mu\nu}^{ab} = \int d\mathbf{r} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) \omega_{\nu}^{Z_b}(\mathbf{r} - \mathbf{r}_b), \quad (32)$$

$$1024 \quad C_{\mu\nu}^{ab} = \int d\mathbf{r} \int d\mathbf{r}' \frac{\omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{r}_a) \omega_{\nu}^{Z_b}(\mathbf{r}' - \mathbf{r}_b)}{\|\mathbf{r} - \mathbf{r}'\|}. \quad (33)$$

1026 Considering a rotation \mathbf{R} of the molecule geometry, without changing the orientation of the basis
 1027 functions, the overlap matrix transforms as
 1028

$$1029 \bar{W}_{\mu\nu}^{ab} = \int d\mathbf{r} \omega_{\mu}^{Z_a}(\mathbf{r} - \mathbf{R}\mathbf{r}_a) \omega_{\nu}^{Z_b}(\mathbf{r} - \mathbf{R}\mathbf{r}_b) = \int d\mathbf{r}' \omega_{\mu}^{Z_a}(\mathbf{R}\mathbf{r}' - \mathbf{R}\mathbf{r}_a) \omega_{\nu}^{Z_b}(\mathbf{R}\mathbf{r}' - \mathbf{R}\mathbf{r}_b) \quad (34)$$

$$1030 = \int d\mathbf{r}' \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) \omega_{\mu'}^{Z_a}(\mathbf{r}' - \mathbf{r}_a) \sum_{\nu'} D_{\nu\nu'}^{Z_b}(\mathbf{R}) \omega_{\nu'}^{Z_b}(\mathbf{r}' - \mathbf{r}_b) \quad (35)$$

$$1031 = \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) \sum_{\nu'} D_{\nu\nu'}^{Z_b}(\mathbf{R}) \int d\mathbf{r}' \omega_{\mu'}^{Z_a}(\mathbf{r}' - \mathbf{r}_a) \omega_{\nu'}^{Z_b}(\mathbf{r}' - \mathbf{r}_b) \quad (36)$$

$$1032 = \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) \sum_{\nu'} D_{\nu\nu'}^{Z_b}(\mathbf{R}) W_{\mu'\nu'}^{ab}. \quad (37)$$

$$1033 \Rightarrow \bar{W}^{ab} = \mathbf{D}^{Z_a}(\mathbf{R}) \mathbf{W}^{ab} (\mathbf{D}^{Z_b}(\mathbf{R}))^T, \quad (38)$$

1034 where \bar{W} is the transformed matrix. This notation will be used throughout this section. Completely
 1035 analogously, the Coulomb matrix transforms
 1036

$$1037 \bar{\mathbf{C}}^{ab} = \mathbf{D}^{Z_a}(\mathbf{R}) \mathbf{C}^{ab} (\mathbf{D}^{Z_b}(\mathbf{R}))^T. \quad (39)$$

1038 Similarly, one can consider a function $f(\mathbf{r}) = \sum_{\mu} f_{\mu} \omega_{\mu}^Z(\mathbf{r})$. The rotated function $\bar{f}(\mathbf{r})$ represented
 1039 in the same basis with rotated molecule geometry is given by
 1040

$$1041 \bar{f}(\mathbf{r}) = f(\mathbf{R}^{-1}\mathbf{r}) = \sum_{\mu} f_{\mu} \omega_{\mu}^Z(\mathbf{R}^{-1}\mathbf{r}) = \sum_{\mu} f_{\mu} \sum_{\mu'} D_{\mu\mu'}^Z(\mathbf{R}^{-1}) \omega_{\mu'}^Z(\mathbf{r}) = \sum_{\mu'} \bar{f}_{\mu'} \omega_{\mu'}^Z(\mathbf{r}), \quad (40)$$

$$1042 \Rightarrow \bar{f}_{\mu'} = \sum_{\mu} D_{\mu\mu'}^Z(\mathbf{R}^{-1}) f_{\mu}. \quad (41)$$

$$1043 \Rightarrow \bar{\mathbf{f}} = (\mathbf{D}^Z(\mathbf{R}^{-1}))^T \mathbf{f} = \mathbf{D}^Z(\mathbf{R}) \mathbf{f}. \quad (42)$$

1044 A.6.1 MESSAGE PASSING

1045 The overlap integrals $o_{abm\mu}$ defined in Equation 9 transform as
 1046

$$1047 \bar{o}_{abm\mu} = \sum_{\nu} \bar{W}_{\mu\nu}^{ab} \bar{h}_{bm\nu} = \sum_{\nu} \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) \sum_{\nu'} D_{\nu\nu'}^{Z_b}(\mathbf{R}) W_{\mu'\nu'}^{ab} \sum_{\nu''} D_{\nu\nu''}^{Z_b}(\mathbf{R}) h_{bm\nu''} \quad (43)$$

$$1048 = \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) o_{abm\mu'}, \quad (44)$$

1049 leading to the transformation behavior
 1050

$$1051 \bar{m}_{abm\mu} = \sum_{\nu} (\bar{W}^{aa})_{\mu\nu}^{-1} \bar{o}_{abm\nu} = \sum_{\nu} \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) (W^{aa})_{\mu'\nu}^{-1} \sum_{\nu'} D_{\nu\nu'}^{Z_b}(\mathbf{R}) o_{abm\nu'} \quad (45)$$

$$1052 = \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) m_{abm\mu'} \quad (46)$$

1053 of the messages $m_{abm\mu}$. The overlap integrals α_{abm} are scalars because
 1054

$$1055 \bar{\alpha}_{abmn} = \sum_{\mu\nu} \bar{h}_{am\mu} \bar{W}_{\mu\nu}^{ab} \bar{h}_{bn\nu} \quad (47)$$

$$1056 = \sum_{\mu\nu} \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) h_{am\mu'} \sum_{\mu''} D_{\mu\mu''}^{Z_a}(\mathbf{R}) \sum_{\nu'} D_{\nu\nu'}^{Z_b}(\mathbf{R}) W_{\mu'\nu'}^{ab} \sum_{\nu''} D_{\nu\nu''}^{Z_b}(\mathbf{R}) h_{bn\nu''} \quad (48)$$

$$1057 = \sum_{\mu'} h_{am\mu'} \sum_{\nu'} W_{\mu'\nu'}^{ab} h_{bn\nu'} = \alpha_{abmn}. \quad (49)$$

1080 For scalar features any MLP is equivariant, since scalars are invariant under rotations. The attention
 1081 weights $\tilde{\alpha}_{abmn}$ are therefore also scalars. The updated node features $\tilde{h}_{am\mu}$ transform as
 1082

$$1083 \tilde{h}_{am\mu} = \sum_{b \in \mathcal{N}_{mp}(a)} \sum_n \tilde{\alpha}_{abmn} \bar{m}_{abn\mu} \quad (50)$$

$$1086 = \sum_{b \in \mathcal{N}_{mp}(a)} \sum_n \tilde{\alpha}_{abmn} \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) m_{abn\mu'} \quad (51)$$

$$1088 = \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) \tilde{h}_{am\mu'}, \quad (52)$$

1090 showing that the message passing layer is equivariant.

1093 A.6.2 NONLINEARITY

1094 The nonlinearity is based on calculating scalar quantities, which are then transformed by an MLP,
 1095 and weighing the original features with these scalars. The scalar features l_{amn} invariant under
 1096 rotations because
 1097

$$1098 \bar{l}_{amn} = \sum_{\mu\nu} \bar{h}_{am\mu} \bar{C}_{\mu\nu}^{aa} \bar{h}_{an\nu} \quad (53)$$

$$1100 = \sum_{\mu\nu} \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) h_{am\mu'} \sum_{\mu''} D_{\mu\mu''}^{Z_a}(\mathbf{R}) \sum_{\nu'} D_{\nu\nu'}^{Z_a}(\mathbf{R}) C_{\mu''\nu'}^{aa} \sum_{\nu''} D_{\nu\nu''}^{Z_a}(\mathbf{R}) h_{an\nu''} \quad (54)$$

$$1103 = \sum_{\mu'} h_{am\mu'} \sum_{\nu'} C_{\mu'\nu'}^{aa} h_{an\nu'} = l_{amn}. \quad (55)$$

1105 Therefore, the weights w_{amn} are also scalars and the updated node features $\tilde{h}_{am\mu}$ transform as
 1106

$$1107 \tilde{h}_{am\mu} = \sum_n \bar{w}_{amn} \bar{h}_{an\mu} = \sum_n w_{amn} \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) h_{an\mu'} \quad (56)$$

$$1110 = \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) \tilde{h}_{am\mu'}, \quad (57)$$

1112 showing that the nonlinearity is equivariant.

1114 A.6.3 L2 NORMALIZATION

1116 The L2 normalization of the features is done per node and feature channel, i.e. for each a and m .
 1117 The normalization factor is a scalar because
 1118

$$1119 \bar{n}_{am} = \sqrt{\sum_{\mu\nu} \bar{h}_{am\mu} \bar{W}_{\mu\nu}^{aa} \bar{h}_{am\nu}} \quad (58)$$

$$1122 = \sqrt{\sum_{\mu\nu} \sum_{\mu'} D_{\mu\mu'}^{Z_a}(\mathbf{R}) h_{am\mu'} \sum_{\mu''} D_{\mu\mu''}^{Z_a}(\mathbf{R}) \sum_{\nu'} D_{\nu\nu'}^{Z_a}(\mathbf{R}) W_{\mu''\nu'}^{aa} \sum_{\nu''} D_{\nu\nu''}^{Z_a}(\mathbf{R}) h_{am\nu''}} \quad (59)$$

$$1125 = \sqrt{\sum_{\mu'} h_{am\mu'} \sum_{\nu'} W_{\mu'\nu'}^{aa} h_{am\nu'}} = n_{am}. \quad (60)$$

1127 Therefore, the normalized features $\tilde{h}_{am\mu}$ are straightforwardly equivariant.

1130 A.6.4 EDGE UPDATE

1132 Exactly the same arguments as for the L2 normalization and the nonlinearity can be used to show
 1133 that $o_{abmn}^{(n)}$ and $o_{abmn}^{(e)}$ are scalars. Therefore, also the generated weights $s_{ab}^{(\cdot,\cdot)}$ and $\tilde{w}_{abmn}^{(\cdot,\cdot)}$ are scalars,
 making the update step in Equation 21 equivariant.

1134
1135

A.7 HYPERPARAMETERS AND TRAINING DETAILS

1136
1137
1138
1139
1140
1141
1142
1143

The model is trained in two stages, the pre-training of the initial guess followed by the full training. During training an exponential moving average (EMA) of the previous weights is used. For testing the epoch that performed best in validation is chosen in all experiments. A mean absolute error loss on the electron density values on the grid is used. When evaluating the density on the grid, edge function values are set to 0 if both of the edge functions have a distance of more than 3 Å from the grid point. The electron density prediction on the grid is multiplied with the standard deviation of the labels calculated over the whole training set. The predicted density is additionally multiplied with 0.1, as this stabilized training in initial experiments.

1144
1145

A.7.1 MODEL HYPERPARAMETERS

1146
1147
1148
1149

The detailed list of hyperparameters of the model is given in Table 8. Settings changed for the pre-training of the initial guess are listed in Table 9. One special layer is the MLP for the edge update,

1150
1151

Table 8: The hyperparameters on the QM9 dataset using the small settings. Differences on other datasets are in the text.

1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178

Hyperparameter	Value
Batch Size	12
Initial Learning Rate	0.001
Learning Rate Schedule	CosineAnnealingLR
T_{\max}	500000 steps
Optimizer	Adam
$\beta_1, \beta_2, \epsilon$	(0.9, 0.999, 10^{-8})
Weight Decay	0.0
Max Training Steps	500000
Gradient Clipping Value	0.5
EMA Decay	0.995
# BOA blocks	4
# hidden channels N^c	32
# edge functions after partial channel mean	8
MLP activation	SiLU
Initial Guess Channels	1
Absolute Value Scale λ	1000
scale factor	0.1
Probed Points	5000
Message passing cutoff r_{mp}	6 Å
Edge function cutoff r_e	3 Å
Basis Set	def2-QZVPPD

Table 9: The pretraining hyperparameters that differ from the default settings in Table 8.

1179
1180
1181
1182
1183
1184
1185
1186
1187

Hyperparameter	Value
Initial Learning Rate	0.001
Learning Rate Scheduler	CosineAnnealingWarmRestarts
T_0	250
T_{mult}	1
# Steps	1000

which consists of a two linear layers, with a SiLU activation and a LayerNorm between them. The first layer has $2 \cdot (N^c)^2$ input and output neurons, and the second has $4 \cdot N^c \cdot (N^c + 1)$ output neurons.

1188 The MLP used in the nonlinearity consists of two linear layers with $(N^c)^2$ input and output neurons,
 1189 separated by first a SiLU activation and then a LayerNorm.

1190
 1191 The Gaussian radial embedding employed in the radial correction uses 50 Gaussians evenly spaced
 1192 between $0a_0$ and $3a_0$ (with a_0 being the Bohr radius). Therefore, 50 values are fed into the following
 1193 MLP, which consists of first a linear layer with dimension 50×50 then SiLU activation and a linear
 1194 layer with dimension $50 \times N^{\text{exp}}$. Here N^{exp} is the number of different exponents used in the radial
 1195 parts of the basis functions, i.e. the number of unique radial functions. For basis functions belonging
 1196 to the same irreducible representations of $\text{SO}(3)$, and therefore having the same radial part, the same
 1197 correction factor is used.

1198
 1199 **Differences between BOA sizes** Going from the small to the large version of BOA, the number of
 1200 randomly probed points per molecule is increased from 5000 to 6000. Additionally, the batch size
 1201 is increased from 12 to 24 which also leads to an increase in "epochs", since the maximum number
 1202 of steps remains at 500000.

1203
 1204 **Differences between datasets** On the MD datasets, the number of maximum training steps was
 1205 reduced to 200000.

A.7.2 BASIS SET

1206
 1207 The number of basis functions per atom type and shell of `def2-QZVPPD` basis set is given in Table
 1208 10.

1209
 1210 Table 10: Uncontracted (primitive) shells per ℓ in `def2-QZVPPD` and total spherical functions.

Element	0	1	2	3	4	Total (sph.)
H	7	4	2	1	0	36
C	16	8	4	2	1	83
N	16	8	4	2	1	83
O	16	9	4	2	1	86
F	16	9	4	2	1	86

A.8 SMOOTH ABSOLUTE VALUE OF ONE OF THE EDGE FUNCTIONS

1221
 1222 Taking the smooth absolute value

$$|x|_s = \begin{cases} \frac{\lambda}{2}x^2, & \text{if } |x| < \frac{1}{\lambda} \\ |x| - \frac{1}{2\lambda}, & \text{otherwise} \end{cases} \quad (61)$$

1223
 1224 of $g_{abo}^{(r)}(\mathbf{r})$ resulted in slightly improved performance compared to keeping the function un-
 1225 modified, see Appendix Section A.1.1. The factor λ steers the scale of the smooth abso-
 1226 lute value function and is here chosen as $\lambda = 1000$. The density is then given by $\hat{\rho}(\mathbf{r}) =$
 $\sum_{(a,b) \in \mathcal{E}_e} \sum_o^{N^o} g_{abo}^{(l)}(\mathbf{r}) \left| g_{abo}^{(r)}(\mathbf{r}) \right|_s + \sum_{a \in \mathcal{N}} \hat{g}_a^{(l)}(\mathbf{r}) \left| \hat{g}_a^{(r)}(\mathbf{r}) \right|_s$.

A.9 LLM USAGE

1235
 1236 Large Language Models (LLMs) were used to assist with writing (formulation and wording), coding
 1237 (AI autocomplete), and gathering resources in the preparation of this article.