

# Electron-Passing Neural Networks for Atomic Charge Prediction in Systems with Arbitrary Molecular Charge

Derek P. Metcalf,<sup>†</sup> Andy Jiang,<sup>†</sup> Steven A. Spronk,<sup>‡</sup> Daniel L. Cheney,<sup>‡</sup> and C. David Sherrill<sup>\*,†</sup>

*†Center for Computational Molecular Science and Technology, School of Chemistry and Biochemistry and School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332-0400, USA*

*‡Molecular Structure and Design, Bristol Myers Squibb Company, P. O. Box 5400, Princeton, NJ 08543, USA*

*\* E-mail: sherrill@gatech.edu.*

## Abstract

Atomic charges are critical quantities in molecular mechanics and molecular dynamics, but obtaining these quantities requires heuristic choices based on atom-typing or relatively expensive quantum mechanical computations to generate a density to be partitioned. Most machine learning efforts in this domain ignore total molecular charges, relying on overfitting and arbitrary rescaling in order to match the total system charge. Here we introduce the electron-passing neural network (EPNN), a fast, accurate neural network atomic charge partitioning model that conserves total molecular charge by construction. EPNNs predict atomic charges very similar to those obtained by partitioning quantum mechanical densities, but at such a small fraction of the cost that they

can be easily computed for large biomolecules. Charges from this method may be used directly for molecular mechanics, as features for cheminformatics, or as input to any neural network potential.

## Introduction

Atomic charge partitioning is the process by which portions of the electron density are assigned to atomic nuclei. This procedure is critical for evaluating electrostatic interactions between atoms and molecules with molecular mechanics (MM) and molecular dynamics (MD). A number of approaches exist to partition an electron density computed from a quantum mechanical (QM) method.<sup>1–3</sup> However, the QM computation is much more time-consuming than the MM computation, and hence obtaining the charges from QM is not practical in normal MM/MD applications, except for those involving replicas of only a few distinct small molecules whose charges can be determined once prior to the MM/MD procedure. Similarly, QM is generally unsuitable for determining charges that might be needed in high-throughput computational screening applications. Atomic charges can instead be assigned heuristically, such as with formal charges, but a more sophisticated approach is to tabulate atom-types and choose charges for each atom-type to best reproduce an experimental or high accuracy computable quantity. However, these approaches typically suffer from the inability to encode geometric dependence such as how the geometry of a system varies during an MD simulation.<sup>4</sup> To circumvent these issues, neural network atomic charge (and higher-order multipole) approximations have been formulated for at least the last twelve years,<sup>5,6</sup> providing fast atomic charge estimates and a means to encode the geometric dependence of charges without relying on a costly QM-based procedure. More recent approaches leverage the latest machine learning frameworks and atomic featurizations, enabling training and inference on much more diverse chemical datasets.<sup>4,7,8</sup> At present, however, these methods have no way of encoding total system charges, so are relegated to predictions on systems of the same charge as the training examples, or they require *ad hoc* scaling of atomic charges. This

forces the user to construct individual models for each charge state, which is inefficient and intractable when considering systems with large total charges. Simply scaling charges *ad hoc* has no physical basis and relies entirely on learning the correlation between geometries and atomic charges in the training set. This may be practical for some applications but is generally an instance of overfitting. One recent work describes AIMNet-ME, an architecture that encodes total system charge during iterative message-passing updates by allocating individual atomic feature vectors for each possible total charge state.<sup>9</sup> In this way, the model can predict properties while being aware of total system charges, unlike previous methods. While this showed good performance for the small systems studied, it remains unclear if this approach is transferable to large systems or to systems with charges besides -1, 0, or 1 due to the requirement of different atomic embeddings for each charge state. The fourth-generation high-dimensional neural network potential (4G-HDNNP) approach from Behler *et al.*<sup>10</sup> uses a similar iterative approach, but chooses symmetry function descriptors as input and utilizes an electronegativity based charge equilibration scheme instead of the purely neural approach of this work and AIMNet-ME. Another recent work dubbed BpopNN incorporates even more physical information, choosing electronic populations that minimize the predicted energy of a system subject to a constrained functional form, and is validated on lithium hydrogen clusters.<sup>11</sup>

Here, we introduce an alternative message-passing neural network model for determining atomic charges of neutral or charged molecules with arbitrary charges, that is applicable to large molecular systems. In this initial study, we target the modeling of atomic charges in protein-ligand systems relevant for drug design efforts. Since atomic charges are not a physical observable, there is no unique way to partition electrons to atomic nuclei. Instead, work in this area seeks to predict charges that reproduce observables, such as energies. We choose to model charges from the minimal-basis iterative stockholder (MBIS) approach, one popular way to partition quantum mechanical densities, validated by its ability to reproduce the electrostatic potential around a molecule and the strength of intermolecular electrostatic

interactions.<sup>3</sup> Our model is trained on a collection of 3503 neutral and charged molecular fragments relevant for proteins or drug-like molecules, and we demonstrate high accuracy when predicting charges on different neutral and charged systems. We demonstrate the computational efficiency of the approach by applying it to the Galectin-3C protein with 2220 atoms.

## Methods

### Message-Passing Neural Networks (MPNNs)

The MPNN is a variety of graph neural network described by Gilmer *et al.* as a general framework, unifying many existing schemes to learn from graph structured data.<sup>12</sup> The MPNN has found success in application to chemical problems, where a graph is defined by the tuple  $\mathcal{G} = (\mathbf{H}, \Omega)$ , with nodes  $\mathbf{H} \in \mathbb{R}^{N \times d_h}$  as the set of  $N$  atom centers with  $d_h$  corresponding features per atom. Edges  $\Omega \in \mathbb{R}^{N \times N \times d_e}$  contain  $d_e$  features per atom pair.

The primary goal of the MPNN is to propagate information about neighboring atoms via a *message-passing* step and accumulate that information with an *update*, which can modify the node features and/or edge features. After several such message-passes and updates, a *readout* step is performed to evaluate the graph for some property. In chemistry, the molecular energy is a key quantity, so the readout typically consists of a node-wise evaluation function to acquire “atomic energies,” which are then summed to recover an estimate of the molecular energy. In principle, this general approach can be used to regress toward any system-wide, atomic, or atomic-pairwise property, provided appropriate labels. Below are details of a typical MPNN implementation; for connections to other works the authors refer the reader to Reference 12.

Initial states of the graph  $\mathcal{G}^0$  can be chosen intuitively and are problem-dependent. For molecular graphs, atomic initial states are often chosen as  $\mathbf{h}_v^0 = (OHE(Z_v), [0, \dots, 0])$ , where  $OHE(Z_v)$  is a one-hot encoding of the element of atom  $v$ , and  $(.,.)$  denotes concatenation.

The additional zeroes pad the initial state to length  $d_h$  and are updated during message-passing. Initial edge states are often defined as the interatomic distance projected on a set of radial basis functions, such as a set of Gaussian functions  $\mathbf{e}_{vw}^0 = \{e^{-\eta_j(r_{vw}-\mu_j)^2}\}_j$ ,  $j \in \{1, \dots, d_e\}$  for atom pair  $v, w$ .  $\eta$  and  $\mu$  are parameters of the Gaussian distribution which may differ for each  $j$ , and  $r_{vw}$  is the Euclidian distance between atoms  $v$  and  $w$ . This initialization step is illustrated in Fig. 1 A.

Message passes and updates are performed  $T$  times, as defined by the user. A message from atom  $w$  to  $v$  at step  $t$  is generated by  $M_t(\mathbf{h}_v^t, \mathbf{h}_w^t, \mathbf{e}_{vw}^t) = NN_t(\mathbf{h}_v^t, \mathbf{h}_w^t, \mathbf{e}_{vw}^t)$  with step-specific dense feed-forward neural network  $NN_t$ . All messages to atom  $v$  are accumulated by a symmetric function, in this case a sum:

$$m_v^{t+1} = \sum_{w \neq v} M_t(\mathbf{h}_v^t, \mathbf{h}_w^t, \mathbf{e}_{vw}^t) \quad . \quad (1)$$

Node feature vectors  $\mathbf{h}_v^t$  update according to

$$\mathbf{h}_v^{t+1} = U_t(\mathbf{h}_v^t, m_v^{t+1}) = NN^U(\mathbf{h}_v^t, m_v^{t+1}) \quad , \quad (2)$$

for example, where  $NN^U$  is typically a small, dense feed-forward neural network. Edges may update analogously but typically do not. These message-passing and update steps are represented by Fig. 1 B and C, respectively, and loop  $T$  times. In a typical MPNN, a readout function is used to evaluate some target property. In the case of molecular energy, atomic energies are usually from neural network evaluations of the final node features and accumulated with a sum:

$$R = \sum_v^N NN^R(\mathbf{h}_v^T) \quad , R \in \mathbb{R}, \quad (3)$$

where  $NN^R$  is another small dense feed-forward neural network. In this work we forgo the

explicit readout function  $R$  and instead use the learned atomic features  $\mathbf{h}_v^T$  as input to an “electron-passing phase” described below.

## Electron-Passing Neural Networks (EPNNs)

The EPNN introduced here is a modified MPNN able to predict atomic charges (node-level property) while conserving the total charge (graph-level property) of a molecule or system. This is accomplished by initializing the system with a set of atomic charges  $\{q_v^0\} \forall v \in \mathcal{G}$  such that the total charge  $Q^G$  is correct, then only updating charges with operations that conserve the total charge. This charge initialization step is illustrated in Fig. 1 D, where the example molecular charge -1 is uniformly distributed across atoms in a toy system.

In lieu of a simple readout function to recover a graph-level property, an electron-passing phase is initiated to recover node-level properties subject to the constant charge constraint. The electron-passing functions  $\sigma$  are different per electron passing step  $s$  as

$$\sigma_s(q_v^s, q_w^s, \mathbf{h}_v^T, \mathbf{h}_w^T, \mathbf{e}_{vw}) = NN_s(q_v^s, q_w^s, \mathbf{h}_v^T, \mathbf{h}_w^T, \mathbf{e}_{vw}) - NN_s(q_w^s, q_v^s, \mathbf{h}_w^T, \mathbf{h}_v^T, \mathbf{e}_{vw}), \quad (4)$$

describing an electron pass from atom  $w$  to atom  $v$ . Electron passes from all atoms are accumulated and atomic charges are updated with

$$q_v^{s+1} = \sum_{w \neq v}^N \sigma_s(q_v^s, q_w^s, \mathbf{h}_v^T, \mathbf{h}_w^T, \mathbf{e}_{vw}) f_c(\mathbf{e}_{vw}). \quad (5)$$

Charge passes  $\sigma_s$  are scaled by a cutoff function  $f_c(\mathbf{e}_{vw}) = \max(0, \mathbf{e}_{vw} - \epsilon)$ , where  $\epsilon$  is a small tolerance (1e-5) to smooth any discontinuities resulting from atoms entering and leaving the defined range. To ensure conservation of electrons, the passing function must be antisymmetric with respect to permuting the two atoms

$$\sigma_s(q_v^s, q_w^s, \mathbf{h}_v^T, \mathbf{h}_w^T, \mathbf{e}_{vw}) = -\sigma_s(q_w^s, q_v^s, \mathbf{h}_w^T, \mathbf{h}_v^T, \mathbf{e}_{vw}), \quad (6)$$

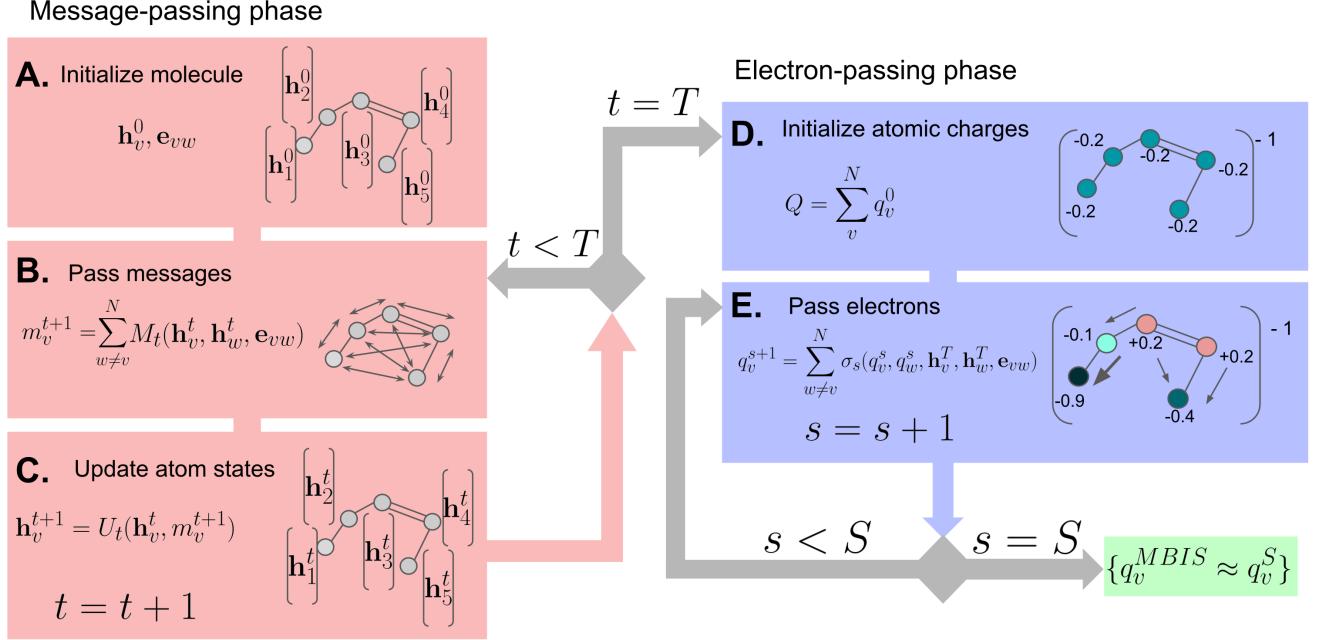


Figure 1: A schematic for electron-passing neural networks. **A.** Initial states are constructed for each atom  $\mathbf{h}_v^0$  and the edges between atoms  $\mathbf{e}_{vw}$ . Atom initial states are simply one-hot encodings of the atomic identity and edge initial states are interatomic distances projected on radial basis functions. **B.** Messages are generated by iteration-dependent neural networks  $M_t$  and are accumulated according to Eq. 1. **C.** Atomic states  $\mathbf{h}_v^t$  are updated according to Eq. 2. If the user-specified  $T$  iterations have elapsed, the algorithm proceeds to the electron-passing phase. Otherwise, message passing continues. **D.** Initial atomic charges  $q_v^0$  come from distributing the known total charge  $Q$  uniformly across atoms, for example. **E.** Electrons are passed between atoms for user-specified  $S$  iterations according to Eqs. 4 and 5. After  $S$  iterations, the EPNN charges  $q_v^S$  should approximate the reference charges  $q_v^{MBIS}$ .

a condition enforced by the functional form of equation 4. The electron-passing step is illustrated in Fig. 1 E, where charges are iteratively updated at each step  $s < S$  according to equations 4 and 5. After  $S$  electron-passing steps have elapsed, the predicted atomic charges are read directly from  $q_v^S$ . Unlike during message-passing, where node features  $\mathbf{h}_v$  are updated, electron-passing only updates node charges  $q_v$ . During training, these charges can be compared with any target charge partitioning such as MBIS, or used in a composite way to reproduce a computed quantity, such as electrostatic interaction energies. Discrepancies from target quantities can be used as an error and backpropagated to optimize the collection of neural networks that define the EPNN. Total message-passing steps  $T$  and electron-passing steps  $S$  are treated as hyperparameters and are optimized for validation set accuracy. We found  $T, S = 3$  were optimal in our studies.

Importantly, the electron-passing phase is functionally independent from the message-passing phase. This means any atomic featurization can be used before electron-passing, not just traditional message-passing as shown in this study. Simpler descriptors such as symmetry functions could act equivalently as inputs, and the iterative nature of the electron-passing phase partially overcomes the locality of the symmetry functions.<sup>13</sup>

## Data Collection

For our model, we choose to infer charges from the popular minimal-basis iterative stockholder (MBIS) method<sup>3</sup> in a supervised manner. Our training and validation data come from three sources: a subset of 1338 small neutral molecules from the QM9 dataset,<sup>14</sup> a set of 62 hand-curated anionic and cationic molecules relevant to drug discovery shown in Figures 2 and 3, respectively, and a 2979 sidechain molecular dimer subset from the sidechain-sidechain interaction (SSI) dataset.<sup>15</sup> SSI contains paired neutral and charged sidechain monomers, including cationic arginine and lysine and anionic aspartate and glutamate. The variety in training and validation data attempts to ensure sound predictions on pharmacologically relevant small molecules and proteins without sacrificing accuracy on neutral systems. The

MBIS atomic charges were computed for each system using the HORTON software package<sup>16</sup> to partition densities from density functional theory (DFT) computations with the PBE0 functional<sup>17</sup> and aug-cc-pVDZ basis<sup>18</sup> performed in the Psi4 electronic structure package.<sup>19,20</sup> Additional computational details and all structures and charges are included in the supporting information.

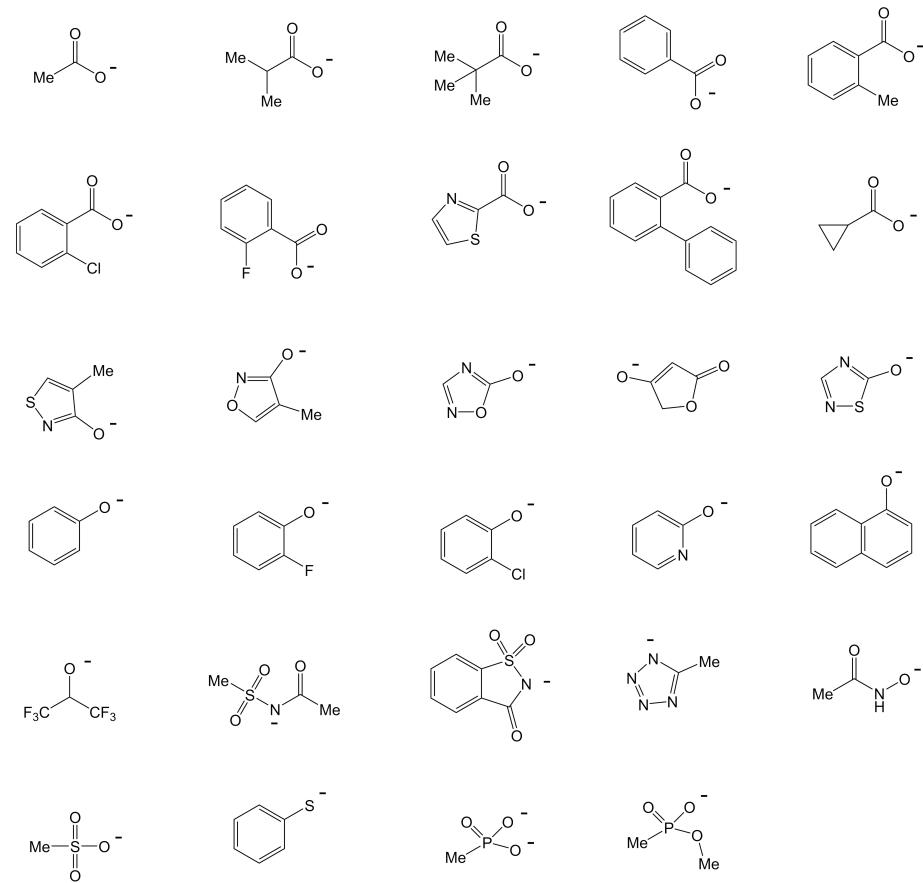
To capture transferability to large systems, we constructed two test cases from the Galectin 3C protein.<sup>21</sup> First, the entire protein was extracted from the PDB entry 6QLP. This geometry was used for model evaluation but we did not compute its charges with the MBIS procedure. Next a three-residue, 80-atom fragment was extracted from the protein, its cleaved bonds capped with hydrogens, and its charges computed with MBIS as above. This particular fragment was chosen for illustrative purposes since it has a negative charge, unlike the positively charged protein.

## Results and Discussion

### Small Molecule Performance

Of the 4379 structures in the combined QM9 subset, SSI dataset, and the manually curated set representing pharmaceutically relevant molecules, a random 80% were used for training and the remainder for validation. Validation performance relative to the target MBIS values is shown in Figure 4, where each dataset is shown separately and charge states are indicated by color. Performance on this validation set is competitive with other charge models while utilizing relatively little data.<sup>4,8,9</sup> Noteworthy is the relative challenge in localizing ions, especially anions. The larger errors in these systems are likely a result of larger charge magnitudes (since the atomic charges in molecular ions can more easily have magnitudes near 1), but also suffer from data sparsity. The data sparsity concern is especially apparent in the anions summarized in Figure 2, where many examples are wholly unique molecular environments compared to the training set. Of note are phosphorus-containing species and

**Drug-like anionic fragments**



**Anionic fragments of proteins, phosphate groups, and sulfate counter ion**

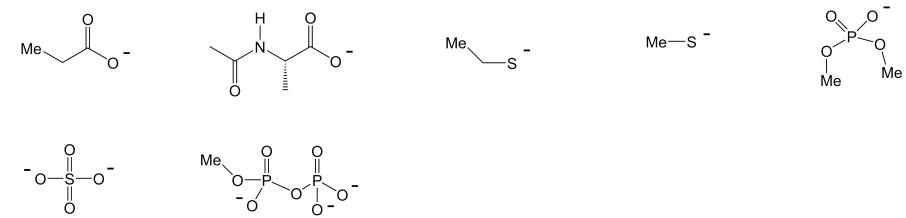
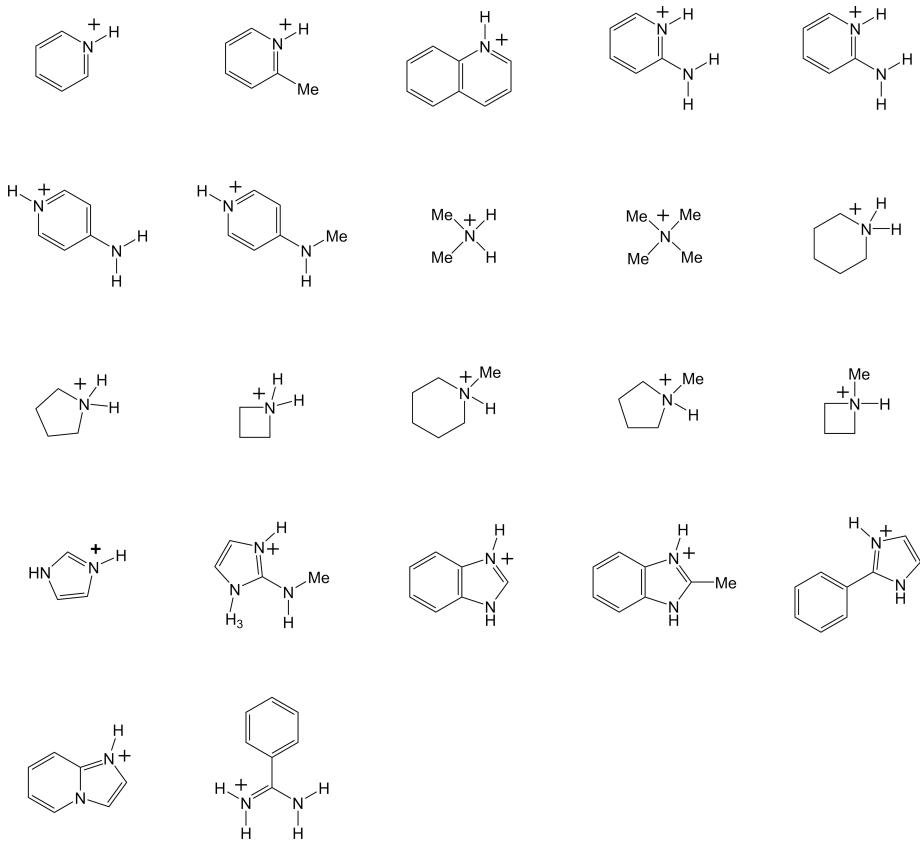


Figure 2: Small anionic molecules relevant to drug discovery. Data preparation and validation information and molecular geometries are included in the supporting information.

**Drug-like cationic fragments**



**Protein cationic fragments**

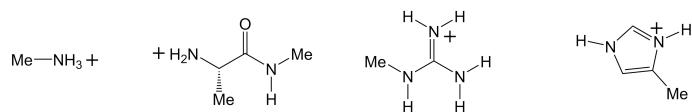


Figure 3: Small cationic molecules relevant to drug discovery. Data preparation and validation information and molecular geometries are included in the supporting information.

multiply-charged species, which are very sparse and have large average charge magnitudes, respectively. As larger datasets for charged systems become available, we believe predictions on anions will approach those of cations and neutral systems. Predictions on the SSI dataset are especially accurate, much of which has to do with shared monomers between the training and validation sets. Nonetheless, prediction errors are smaller than the variance due to charge transfer in a typical charge-neutral or charge-charge interaction in the SSI dataset, so the models present good transferability.

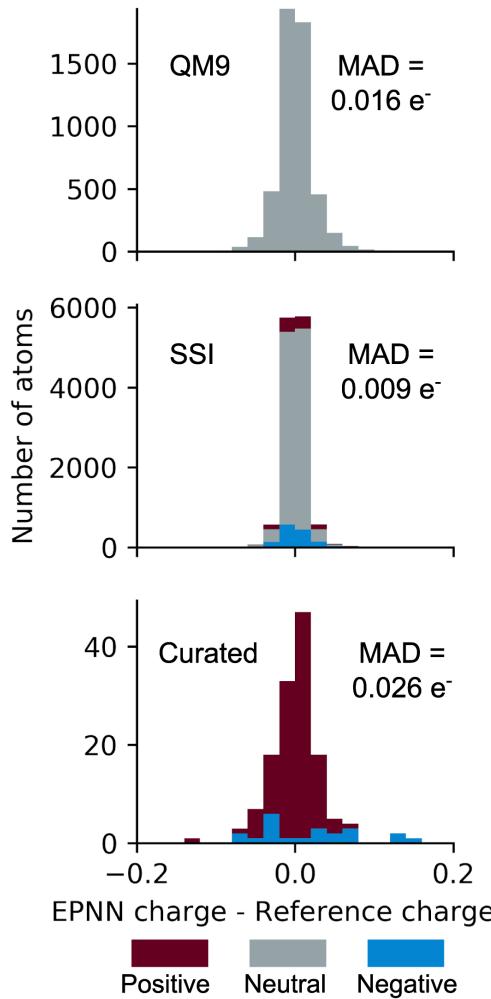


Figure 4: Summary errors from the validation set, differentiated by source data set (QM9, SSI, Curated) and by the charge state of the molecular system.

## Charge Transfer and Polarization

Understanding the change in atomic charges in response to nearby molecules is critical to explain many condensed phase phenomena. The SSI dataset, which enumerates interactions between protein sidechains, constitutes a significant fraction of the training and validation data for this charge model. By examining interactions between charged and uncharged sidechains, we can immediately quantify the combined charge transfer and polarization effects. An example of glutamic acid and glutamine from the validation set is illustrated in Figure 5, where the model is able to replicate the charge distribution of each monomer in vacuum, as well as in the presence of another side chain, with reasonable accuracy. The difference in a monomer’s charges as one moves from the gas phase to the dimer complex, which we may denote as  $\Delta q_i$  for each atom  $i$ , is due to charge transfer to/from the other monomer and polarization of the monomer’s electrons, induced by the presence of the opposing monomer. Most of the error in the ML prediction of  $\Delta q_i$  is due to inaccuracy when predicting gas-phase monomeric charges, resulting in a slightly exaggerated view of polarization and charge transfer. Nonetheless, the accuracy of the model is able to qualitatively capture even very subtle effects, like in this charge-neutral interaction.

## Protein Validation

To show application to large systems and the extensivity of our charge model with respect to total system charge, we conducted an illustrative study using the Galectin 3C protein.<sup>21</sup> To reduce the protein to a manageable size in order to perform the reference MBIS charge partitioning, we isolated three connected residues and capped the ends to retain sensible bonding. This subsystem has net charge -1, which is fed directly as input along with the Cartesian coordinates of the system. The subsystem is comprised of 80 atoms, larger than any molecule in the training set. Just as with the smaller molecules, the charge is able to self-organize and accurately match the target MBIS partitioning. This comparison is illustrated in Figure 6A. Next, the rest of the protein was added, increasing the net charge to +2, a

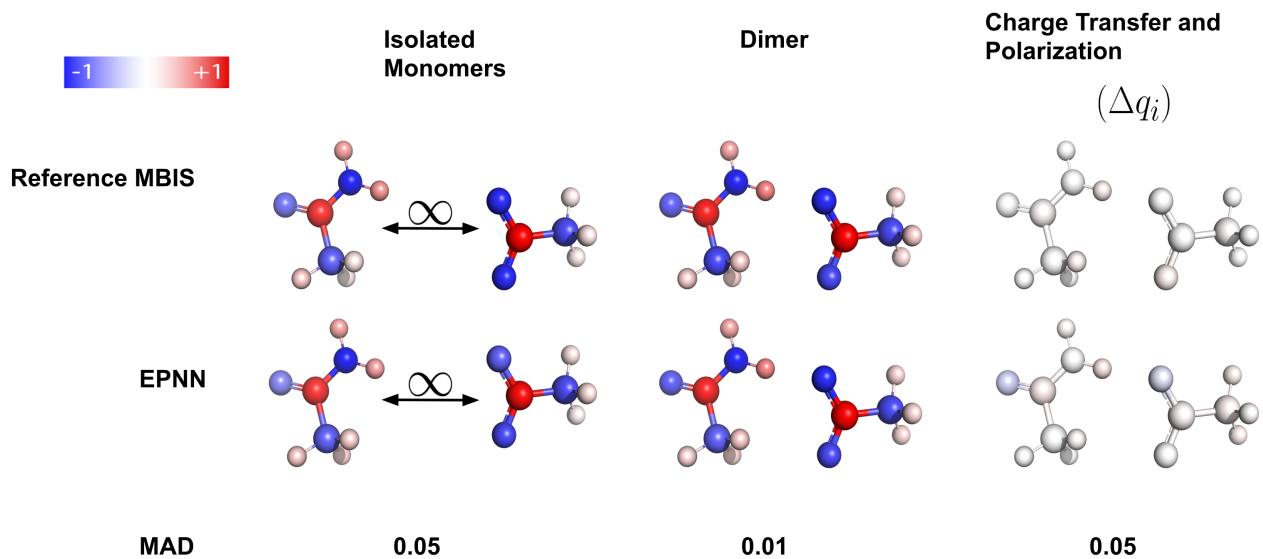


Figure 5: An example polarization and charge transfer interaction between glutamine (left) and glutamic acid (right) sidechains from the SSI dataset.

3-electron change. Of note is the fact 32 atoms in this structure are formally charged, and just tend to cancel, but the net charge is the only input required. The system was again evaluated with the EPNN. The model produces very similar predictions for the subsystem, highlighting its extensivity in system size and total charge magnitude. Indeed, the only large variations in charge from the subsystem computation are in atoms whose bonding was compromised by the capping procedure. A visual comparison is shown in Figure 6B, where charges for the entire protein are approximated but the subset from the fragment study are highlighted.

On a single 8-core Intel Core i7-9800X CPU, the 80-atom subsystem DFT computation took 19 minutes. By comparison, the EPNN evaluation of this system on the same hardware took 0.16 seconds. The 2220-atom supersystem EPNN evaluation took just 62 seconds.

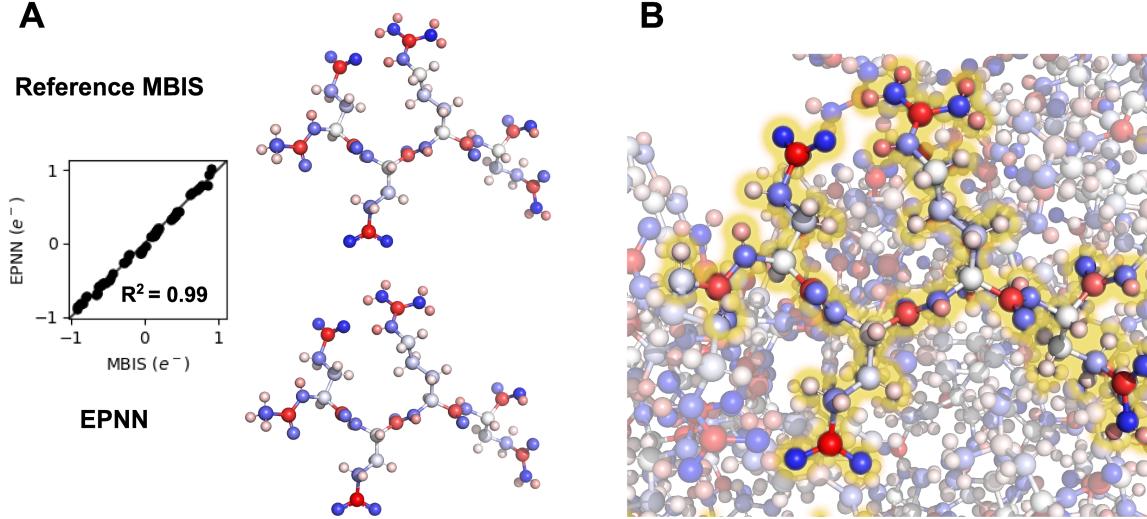


Figure 6: **A.** A comparison between reference MBIS charges and predicted EPNN charges on an anionic fragment of Galectin 3C. Blue indicates negatively charged atoms and red indicates positively charged. **B.** EPNN charge predictions on the entire 2220-atom Galectin 3C protein, net charge +2. The fragment from **A** is backlit to show similarity in charge prediction despite system size and charge changes.

## Conclusions

A neural network scheme is introduced which maps directly from Cartesian coordinates and total system charge to an atomic charge partitioning. Unlike other works, this architecture conserves total charge by construction, requiring a single model and no arbitrary scaling. Additionally, the model does not grow or need additional encoding or training data in each specific charge state to express arbitrary total charges. This architecture, dubbed an electron-passing neural network (EPNN), attains high accuracy on charged and neutral drug- and protein-like molecules with a small training set. Subtle charge transfer and polarization effects due to intermolecular interactions are reproduced with good fidelity. The quality of charges estimated with the EPNN is independent of system size: accurate charge predictions on a protein fragment persist when the entire protein of over 2000 atoms and different charge state is evaluated. With appropriate distance cutoff criteria, the computational complexity of the EPNN is linear in number of atoms. Future work includes the generation of more data for charged systems, including new chemical environments and non-equilibrium structures.

The electron-passing layer is a simple, modular unit that can be used with any atomic featurization as input, making it ideal to use in conjunction with existing machine learning models for energy or other properties. The physically meaningful atomic charge output can be used for molecular mechanics, machine-learned potentials, or as features for cheminformatics. As large datasets including charged systems become publicly available, we hope electron-passing layers can be trained for systems including different elements and a wider range of chemical systems.

## Acknowledgement

The authors gratefully acknowledge financial support from Bristol Myers Squibb, and from the U.S. National Science Foundation through grant CHE-1955940. D.P.M. gratefully acknowledges support from the U.S. National Science Foundation through a Graduate Research Fellowship (grant DGE-1650044). This research was supported in part by research cyberinfrastructure resources and services provided by the Partnership for an Advanced Computing Environment (PACE) at the Georgia Institute of Technology, and by the Georgia Tech Hive Cluster, funded by the U.S. National Science Foundation through grant MRI-1828187.

## Supporting Information Available

The Supporting Information includes computational details in data collection and architectural details of the model. It is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jcim.XXXX>. All code and data used for this study are available at <https://github.com/derekmetcalf/epnn>. This material is available free of charge via the Internet at <http://pubs.acs.org/>.

## References

- (1) Hirshfeld, F. L. Bonded-Atom Fragments for Describing Molecular Charge Densities. *Theor. Chem. Acc.* **1977**, *44*, 129–138.
- (2) Maslen, E.; Spackman, M. Atomic Charges and Electron Density Partitioning. *Aust. J. of Phys.* **1985**, *38*, 273–287.
- (3) Verstraelen, T.; Vandenbrande, S.; Heidar-Zadeh, F.; Vanduyfhuys, L.; Speybroeck, V. V.; Waroquier, M.; Ayers, P. W. Minimal Basis Iterative Stockholder: Atoms in Molecules for Force-Field Development. *J. Chem. Theory Comput.* **2016**, *12*, 3894–912.
- (4) Bleiziffer, P.; Schaller, K.; Riniker, S. Machine Learning of Partial Charges Derived from High-Quality Quantum-Mechanical Calculations. *J. Chem. Inf. Model.* **2018**, *58*, 579–590.
- (5) Darley, M. G.; Handley, C. M.; Popelier, P. L. A. Beyond Point Charges: Dynamic Polarization from Neural Net Predicted Multipole Moments. *J. Chem. Theory Comput.* **2008**, *4*, 1435–1448.
- (6) Artrith, N.; Morawietz, T.; Behler, J. High-Dimensional Neural-Network Potentials for Multicomponent Systems: Applications to Zinc Oxide. *Phys. Rev. B* **2011**, *83*, 153101.
- (7) Zubatyuk, R.; Smith, J. S.; Leszczynski, J.; Isayev, O. Accurate and Transferable Multitask Prediction of Chemical Properties with an Atoms-In-Molecules Neural Network. *Sci. Adv.* **2019**, *5*, eaav6490.
- (8) Sifain, A. E.; Lubbers, N.; Nebgen, B. T.; Smith, J. S.; Lokhov, A. Y.; Isayev, O.; Roitberg, A. E.; Barros, K.; Tretiak, S. Discovering a Transferable Charge Assignment Model Using Machine Learning. *J. Phys. Chem. Lett.* **2018**, *9*, 4495–4501.

(9) Zubatyuk, R.; Smith, J.; Nebgen, B. T.; Tretiak, S.; Isayev, O. Teaching a Neural Network to Attach and Detach Electrons from Molecules. *ChemRxiv* **2020**, *10.26434/chemrxiv.12725276.v1*.

(10) Ko, T. W.; Finkler, J. A.; Goedecker, S.; Behler, J. A Fourth-Generation High-Dimensional Neural Network Potential with Accurate Electrostatics Including Non-local Charge Transfer. *arXiv* **2020**, *abs/2009.06484*.

(11) Xie, X.; Persson, K. A.; Small, D. W. Incorporating Electronic Information into Machine Learning Potential Energy Surfaces via Approaching the Ground-State Electronic Energy as a Function of Atom-Based Electronic Populations. *J. Chem. Theory Comput.* **2020**, *16*, 4256–4270.

(12) Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; Dahl, G. E. Neural Message Passing for Quantum Chemistry. *ArXiv* **2017**, *abs/1704.01212*.

(13) Behler, J. Atom-Centered Symmetry Functions for Constructing High-Dimensional Neural Network Potentials. *J. Chem. Phys.* **2011**, *134*, 074106.

(14) Ramakrishnan, R.; Dral, P. O.; Rupp, M.; von Lilienfeld, O. A. Quantum Chemistry Structures and Properties of 134 Kilo Molecules. *Sci. Data* **2014**, *1*.

(15) Burns, L. A.; Faver, J. C.; Zheng, Z.; Marshall, M. S.; Smith, D. G. A.; Vanommeslaeghe, K.; MacKerell, A. D.; Merz, K. M.; Sherrill, C. D. The BioFragment Database (BFD<sub>b</sub>): An Open-Data Platform for Computational Chemistry Analysis of Noncovalent Interactions. *J. Chem. Phys.* **2017**, *147*, 161727.

(16) T. Verstraelen, P. Tecmer, F. Heidar-Zadeh, K. Boguslawski, M. Chan, Y. Zhao, T.D. Kim, S. Vandenbrande, D. Yang, C.E. González-Espinoza, S. Fias, P.A. Limacher, D. Berrocal, A. Malek, P.W. Ayers HORTON 2.0.1, <http://theochem.github.com/horton/>, 2015.

(17) Perdew, J. P.; Ernzerhof, M.; Burke, K. Rationale for Mixing Exact Exchange with Density Functional Approximations. *J. Chem. Phys.* **1996**, *105*, 9982–9985.

(18) Papajak, E.; Zheng, J.; Xu, X.; Leverentz, H. R.; Truhlar, D. G. Perspectives on Basis Sets Beautiful: Seasonal Plantings of Diffuse Basis Functions. *J. Chem. Theory Comput.* **2011**, *7*, 3027–3034.

(19) Parrish, R. M.; Burns, L. A.; Smith, D. G. A.; Simmonett, A. C.; DePrince, A. E.; Hohenstein, E. G.; Bozkaya, U.; Sokolov, A. Y.; Di Remigio, R.; Richard, R. M.; Gonthier, J. F.; James, A. M.; McAlexander, H. R.; Kumar, A.; Saitow, M.; Wang, X.; Pritchard, B. P.; Verma, P.; Schaefer, H. F.; Patkowski, K.; King, R. A.; Valeev, E. F.; Evangelista, F. A.; Turney, J. M.; Crawford, T. D.; Sherrill, C. D. Psi4 1.1: An Open-Source Electronic Structure Program Emphasizing Automation, Advanced Libraries, and Interoperability. *J. Chem. Theory Comput.* **2017**, *13*, 3185–3197.

(20) Smith, D. G. A.; Burns, L. A.; Simmonett, A. C.; Parrish, R. M.; Schieber, M. C.; Galvelis, R.; Kraus, P.; Kruse, H.; Di Remigio, R.; Alenaizan, A.; James, A. M.; Lehtola, S.; Misiewicz, J. P.; Scheurer, M.; Shaw, R. A.; Schriber, J. B.; Xie, Y.; Glick, Z. L.; Sirianni, D. A.; O'Brien, J. S.; Waldrop, J. M.; Kumar, A.; Hohenstein, E. G.; Pritchard, B. P.; Brooks, B. R.; Schaefer, H. F.; Sokolov, A. Y.; Patkowski, K.; DePrince, A. E.; Bozkaya, U.; King, R. A.; Evangelista, F. A.; Turney, J. M.; Crawford, T. D.; Sherrill, C. D. Psi4 1.4: Open-Source Software for High-Throughput Quantum Chemistry. *J. Chem. Phys.* **2020**, *152*, 184108.

(21) Berman, H.; Westbrook, J.; Feng, Z.; Gilliland, G.; Bhat, T.; Weissig, H.; Shindyalov, I.; Bourne, P. The Protein Data Bank. *Nucleic Acids Res.* **2000**, *28*, 235–242.