The Journal of Chemical Physics

RESEARCH ARTICLE | APRIL 08 2025

How local is "local"? Deep learning reveals locality of the induced magnetic field of polycyclic aromatic hydrocarbons

Yair Davidson; Aviad Philipp; Sabyasachi Chakraborty ; Alex M. Bronstein ;
Renana Gershoni-Poranne

Check for updates

J. Chem. Phys. 162, 144101 (2025)
https://doi.org/10.1063/5.0257558

View Online
Export Citation

Articles You May Be Interested In

Predicting bond-currents in polybenzenoid hydrocarbons with an additivity scheme

J. Chem. Phys. (January 2021)

On the additivity of current density in polycyclic aromatic hydrocarbons

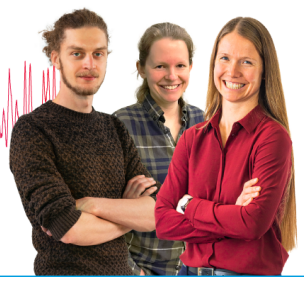
J. Chem. Phys. (July 2009)

The chemical space of B, N-substituted polycyclic aromatic hydrocarbons: Combinatorial enumeration and high-throughput first-principles modeling

J. Chem. Phys. (March 2019)

22 April 2025 09:26:07

Webinar From Noise to Knowledge
May 13th – Register now
Zurich Instruments
Universität Konstanz



AIP Publishing

The Journal of Chemical Physics ARTICLE pubs.aip.org/aip/jcp

How local is "local"? Deep learning reveals locality of the induced magnetic field of polycyclic aromatic hydrocarbons

Cite as: J. Chem. Phys. 162, 144101 (2025); doi: 10.1063/5.0257558
Submitted: 10 January 2025 • Accepted: 19 March 2025 •
Published Online: 8 April 2025

View Online Export Citation CrossMark

Yair Davidson,[1] Aviad Philipp,[1] Sabyasachi Chakraborty,[2] Alex M. Bronstein,[1,3,a)] and Renana Gershoni-Poranne[2,a)]

AFFILIATIONS

[1] Department of Computer Science, Technion—Israel Institute of Technology, Haifa 32000, Israel
[2] Schulich Faculty of Chemistry and the Resnick Sustainability Center for Catalysis, Technion—Israel Institute of Technology, Haifa 32000, Israel
[3] Institute of Science and Technology Austria, Klosterneuburg 3400, Austria

[a)] **Authors to whom correspondence should be addressed:** bron@cs.technion.ac.il and rporanne@technion.ac.il

ABSTRACT

We investigate the locality of magnetic response in polycyclic aromatic molecules using a novel deep-learning approach. Our method employs graph neural networks (GNNs) with a graph-of-rings representation to predict nucleus independent chemical shifts (NICS) in the space around the molecule. We train a series of models, each time reducing the size of the largest molecules used in training. The accuracy of prediction remains high (MAE < 0.5 ppm), even when training the model only on molecules with up to four rings, thus providing strong evidence for the locality of magnetic response. To overcome the known problem of generalization of GNNs, we implement a *k*-hop expansion strategy and succeed in achieving accurate predictions for molecules with up to 15 rings (almost 4 times the size of the largest training example). Our findings have implications for understanding the magnetic response in complex molecules and demonstrate a promising approach to overcoming GNN scalability limitations. Furthermore, the trained models enable rapid characterization, without the need for more expensive DFT calculations.

© 2025 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC) license (https://creativecommons.org/licenses/by-nc/4.0/). https://doi.org/10.1063/5.0257558

22 April 2025 09:26:07

INTRODUCTION

One of the characteristic features of cyclically conjugated (i.e., aromatic) molecules is that they sustain induced magnetic currents when placed into a magnetic field oriented perpendicularly to the molecular plane. This phenomenon is described by the "ring-current model" (RCM) developed by Pauling,[1] Lonsdale,[2] McWeeny,[3] and Pople,[4] which likens an aromatic ring's delocalized π electrons to a wire loop.[5,6] As a result of the induced ring current, an induced magnetic field is generated.[7,8] This induced field causes the typical downfield chemical shift of aromatic protons and is also the basis for the most commonly used method for evaluating aromaticity—NICS (nucleus independent chemical shift).[9–13] NICS is essentially a computational NMR measurement that can be performed at any point desired by the user. By obtaining the chemical shift (δ) at various locations in the vicinity of the molecule, one can indirectly evaluate the direction and magnitude of the induced field and induced current.

The extent to which the magnetic response can be considered "local" has been a subject of ongoing debate. On the one hand, as highlighted by Gomes and Mallion,[14] magnetic properties such as ring currents and induced magnetic fields are inherently non-local. They are influenced by the entire π-electron system of a molecule, not just individual rings. This non-locality is a key reason why magnetic criteria can sometimes give results that seem counterintuitive when compared to other aromaticity measures.[15,16] Aihara concurred, emphasizing that larger circuits in polycyclic systems can contribute significantly to magnetic properties. This entails that the magnetic response in one part of a molecule can be influenced by distant parts of the π system.[17]

© Author(s) 2025

On the other hand, a very common approach to analyzing polycyclic aromatic systems (PASs) is to break them down into smaller components. Perhaps the most well-known example of doing so is the Clar formulation, which identifies disjoint sets of 6π electrons (i.e., Clar sextets) within a polybenzenoid hydrocarbon (PBH),[18,19] an idea subsequently generalized by Glidewell and Lloyd to PASs, in a broader sense.[20] Randić later elaborated on this approach, introducing the concept of conjugated circuits.[21–23] Although it was originally developed to calculate resonance energies, later work by Mandado showed that the conjugated circuit concept could also be used to determine ring current intensities.[24] Based on the same idea—breaking PASs into various smaller components—additional schemes have been developed to predict different manifestations of magnetic response, each one using a different decomposition framework. Our group has developed additivity schemes for predicting both NICS-XY-scans[25,26] and bond currents[27] of PASs using their mono-, bi-, and tricyclic subunits. Tiznado and co-workers recently used similar subunits but a different approach for decomposing the current density into chemically meaningful contributions.[28,29]

The aforementioned schemes and their ability to predict the magnetic behavior of molecules larger than any of the defined subunits strongly imply that the magnetic response of PASs is dominated by *local* effects, rather than global or large-circuit effects. Otherwise, simply adding one ring (i.e., going from n to $n + 1$ rings) in a PAS could result in entirely different and unpredictable behavior. In the majority of molecules, this does not appear to be the case. Nevertheless, the extent of localization remains unclear, as does its dependence on the geometry of the local environment. In addition to the deeper understanding that this may provide, knowing the effect of different substructures may be instrumental for designing PASs with the desired magnetic behavior.

In this work, we aim to address these questions using a deep-learning approach. To this end, we design and implement graph neural net (GNN)-based models that use our previously developed graph-of-rings chemical representation,[30] to predict the NICS values in the vicinity of a PAS. The use of deep-learning models provides us with a new approach for artificially breaking down the molecules and analyzing the importance of the smaller components. As opposed to conventional computational chemistry, in which the smaller subunits need to be valid molecules themselves, in this approach, the model is agnostic to the chemical validity of the substructure. Thus, by using different substructures of the molecule as input to the models, we elucidate the contributions of various structural motifs to the chemical shielding values. We investigate whether these values are indeed sensitive to structural changes at distant locations, and we determine the extent of the locality. To the best of our knowledge, this approach has not been used before in this context and so it provides a complementary perspective to the existing methods. We emphasize that our questions focus on the physical observable of the magnetic response of a PAS under an external magnetic field. We do not attempt to characterize the aromaticity of the molecules, which is a hotly debated topic in and of itself, and one which is much more open to interpretation.

METHODOLOGY

In the following subsections, we describe the dataset construction, grid generation, and chemical shielding computations, followed by the various design aspects of the machine-learning model.

Molecular datasets

Just as benzene is considered the prototypical aromatic molecule, so are PBHs often considered to be the model systems for the much broader class of PASs. Therefore, we elected to use a curated dataset of *cata*-condensed PBHs (cc-PBHs) as our test cases in this study. These systems have been extensively investigated, in particular in the context of their local vs global effects.[16,31–33] To investigate the magnetic field around a cc-PBH, we curated a dataset of 301 cc-PBHs with varying numbers of benzene rings (n_{rings}), ranging in size from $n_{\text{rings}} = 1$ to $n_{\text{rings}} = 15$ and encompassing a variety of structural features—linear, angular, branched, planar, non-planar (representative examples are shown in Fig. 1; the entire dataset is provided in the supplementary material–Figs. S1–S3). The dataset includes all isomers for $n_{\text{rings}} \leq 7$ and representative isomers for $8 \leq n_{\text{rings}} \leq 15$ (see Table I for further details).

For the PBHs with $1 \leq n_{\text{rings}} \leq 10$, geometries were obtained from the COMPAS-1D dataset developed by Wahab *et al.*[34] The geometries of the larger PBHs (with $11 \leq n_{\text{rings}} \leq 15$) were optimized at the same level of theory as the COMPAS-1D molecules, using the B3LYP[35] functional with the def2-SVP basis set[36] and with Grimme's D3 dispersion correction[37,38] with Becke–Johnson damping,[39] without any symmetry constraints. All the calculations were performed with ORCA software, version 5.0.3.[40,41] The optimized geometries were confirmed to be minima on the respective potential energy surfaces with no imaginary frequencies.

Grid generation and NICS calculation

The nuanced nature of the magnetic field around a PBH warrants a high-resolution description of the surrounding space, which

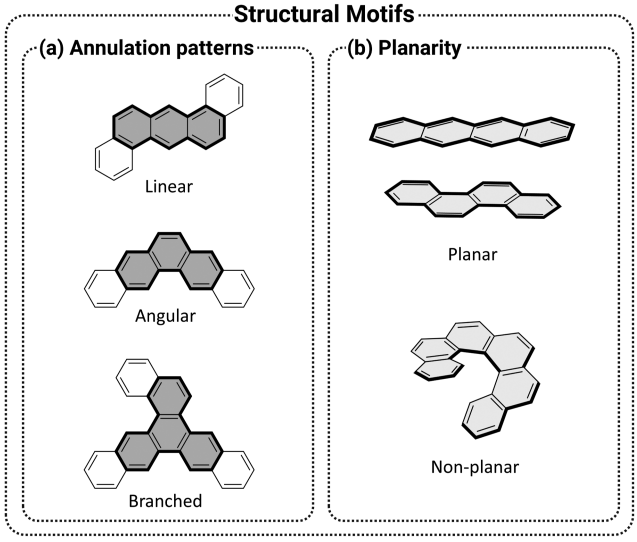


FIG. 1. (a) Representative cc-PBHs depicting: (a) linear, angular, and branched annulation patterns (the respective components are highlighted in bold and colored in gray); (b) planar and non-planar geometries.

22 April 2025 09:26:07

© Author(s) 2025

TABLE I. Number of cc-PBHs in the dataset, divided according to the number of rings. The cc-PBHs used for training the respective models are indicated with a ✓.

	No. cc-PBH	Model			
No. Rings	isomers	7	6	5	4
1	1	✓	✓	✓	✓
2	1	✓	✓	✓	✓
3	2	✓	✓	✓	✓
4	5	✓	✓	✓	✓
5	12	✓	✓	✓	···
6	37	✓	✓	···	···
7	123	✓	···	···	···
8	20	···	···	···	···
9	20	···	···	···	···
10	20	···	···	···	···
11	12	···	···	···	···
12	12	···	···	···	···
13	12	···	···	···	···
14	12	···	···	···	···
15	12	···	···	···	···

is sufficiently broad to account for the induced field at short, intermediate, and long distances from the molecular structure. To obtain this, it is necessary to construct a grid of probe points and calculate the induced magnetic field at each of these points. Such evaluations are essentially identical to performing a NICS (nucleus independent chemical shift)[9,11,13,42] calculation at each point; therefore, throughout this text, we refer to the obtained results as NICS values. For evaluating aromaticity, NICS values are often calculated above/below the molecular plane. However, in this work, we are interested in analyzing the behavior of the induced field in the entire space around the molecule. Therefore, we construct a grid of probe points (also known in the NICS terminology as a "ghost atom" or "bq") around the molecule. The approach of encasing a molecule in a grid of probe points to obtain chemical insights into their magnetic properties has been previously used. Some recent examples have been reported by Kilymis *et al.*,[43] Orozco-Ic *et al.*,[44] and Sanchez-Sanz *et al.*[45]

To calculate the NICS values in the space around each molecule, each molecule is first reoriented to align the molecular plane as closely as possible with the *xy* plane. To ensure full encapsulation of the molecule, the size of the bounding box is determined based on the molecule's size, i.e., an adaptive grid. Along the *x* and *y* axes, the spans are determined by calculating the distance between the farthest two atoms and extending this range by ~1.0 Å in each direction. Along the *z* axis, the span is determined based on the number of rings in the molecule (n_{rings}) and has a fixed size of either ±4.0 Å (for $1 \leq n_{\text{rings}} \leq 7$) or ±5.5 Å (for $n_{\text{rings}} > 7$) from the *xy* plane. In all Cartesian directions, grid points are placed at intervals of 0.5 Å, and any point within 0.5 Å of a C or H atom is removed. As a result of these definitions, depending on the size and (non-)planarity of the molecule, the grid may contain anywhere between 5k probes (for benzene) to 50k probes (for a 15-ring PBH). Figure 2 depicts a representative example of a PBH encapsulated in such a grid.

A probe is placed at each grid point and the isotropic chemical shielding is calculated with the gauge-including atomic orbital (GIAO) method,[46–49] using the CAM-B3LYP functional[50] and def2-SVP[36] basis set (for C and H atoms) and aug-cc-pVDZ[51] basis set (for probes) as implemented in the Gaussian 09, Revision D.01 suite

22 April 2025 09:26:07

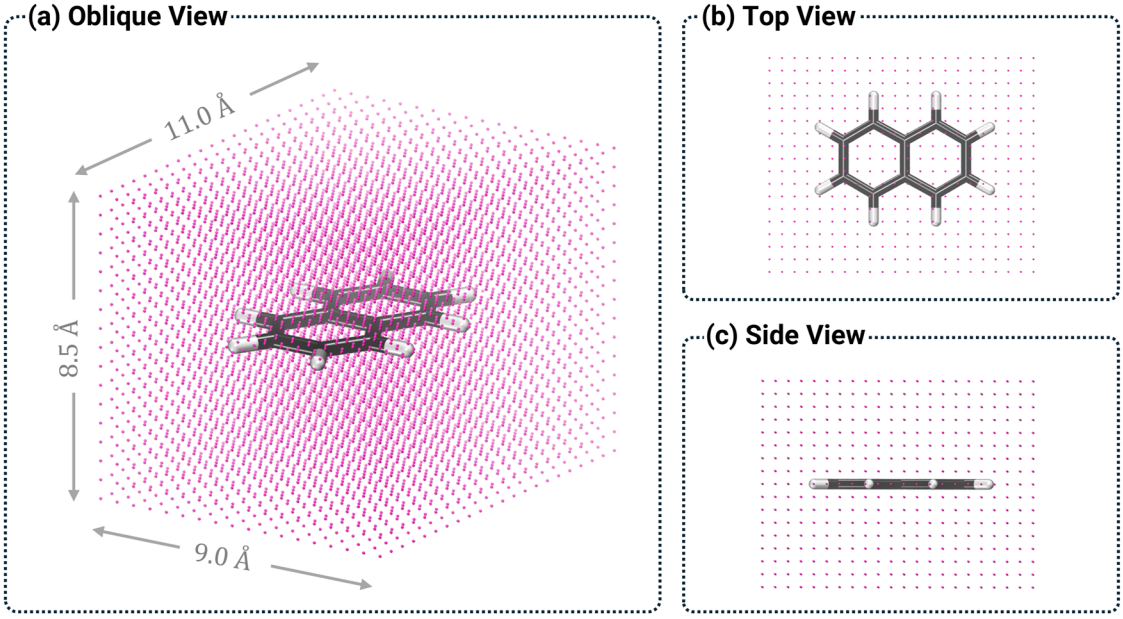


FIG. 2. Representative PBH (naphthalene) encapsulated in a three-dimensional grid of 6678 probe points: (a) oblique view, (b) view from the top, and (c) view from the side.

© Author(s) 2025

of programs.[52] Following convention, we report the negative of the isotropic chemical shielding as the NICS value.

Molecular graph representation

As established in our previous work,[30] we represent each PBH molecule as a graph of rings (GOR) $\mathcal{G}(\mathbf{X})$, as illustrated in Fig. 3(a). Each node in the GOR represents a benzene ring, and the columns of the matrix $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_n)$ represent the location of the individual ring centroids. Two nodes are connected by an edge if and only if the rings are fused; this connectivity can often be inferred from $\mathbf{X}$ alone. This choice is motivated by the two main advantages of the GOR: (a) it substantially reduces the complexity of the chemical representation, allowing models to train faster and with less data; (b) we have previously shown that the chemical properties and behavior of PBHs are primarily determined at the ring-level.[30,34,53]

Because we are interested in predicting the NICS scalar field $\delta(\mathbf{y})$ surrounding the molecule, we receive the location $\mathbf{y}$ of a probe point as the input and augment the node features of each of the rings in $\mathcal{G}(\mathbf{X})$ to hold the distance between the probe and the ring center, as well as the distance between the probe and the ring plane (note that ring planes may differ in non-planar molecules). This results in a feature vector of the form $(\mathbf{x}_i, \|\mathbf{x}_i - \mathbf{y}\|, |\hat{\mathbf{n}}_i \cdot (\mathbf{x}_i - \mathbf{y})|)$, where $\hat{\mathbf{n}}_i$ denotes the ith ring unit normal vector, as visualized in Fig. 3(b). We denote the augmented GOR as $\mathcal{G}(\mathbf{X}, \mathbf{y})$, emphasizing its dependence on the probe location $\mathbf{y}$. The inclusion of the normal information in the form of the point-to-plane distance is to enable us to represent anisotropic fields. Importantly, this representation is key to subsequent interpretation, as it allows us to connect each probe to the individual rings.

Prediction model

For the prediction of the scalar field property $\delta(\mathbf{y})$ at a given probe location $\mathbf{y} \in \mathbb{R}^3$, we adopt the well-established SE(3)-transformer architecture, which was designed for 3D point clouds and graphs[54] and produces node-level outputs that are equivariant to the special Euclidean group (rotation and translation). Because we are interested in a single scalar output for every molecule-probe pair, $(\mathcal{G}(\mathbf{X}), \mathbf{y})$, represented as the augmented GOR $\mathcal{G}(\mathbf{X}, \mathbf{y})$, we add a final permutation-invariant aggregation layer to achieve SE(3) invariance. The application of the model to the augmented GOR is denoted as $\delta(\mathbf{y}) \approx g(\mathcal{G}(\mathbf{X}, \mathbf{y}))$. Furthermore, given m probes arranged into a matrix $\mathbf{Y} = (\mathbf{y}_1, \ldots, \mathbf{y}_m)$, we denote by $g(\mathcal{G}(\mathbf{X}, \mathbf{Y})) = (g(\mathcal{G}(\mathbf{X}, \mathbf{y}_1)), \ldots, g(\mathcal{G}(\mathbf{X}, \mathbf{y}_m)))$. Our model's architecture guarantees special Euclidean invariance of the modeled field, $g(\mathcal{G}(\mathbf{\Pi X T}, \mathbf{Y T})) = g(\mathcal{G}(\mathbf{X}, \mathbf{Y}))$ for every rotation and translation $\mathbf{T} \in \mathrm{SE}(3)$, as well as invariance to the permutation of the rings, $\mathbf{\Pi}$.

Graph neural networks such as the one proposed here are known for their poor generalization of graphs of different sizes compared to those observed at training.[55] To mitigate this phenomenon and, specifically, to be able to generalize to larger molecules than those seen at training, we propose to localize the model's input to a subgraph of the full respective GOR.

A k-hop expansion subgraph (henceforth, "k-graph") of a node i in $\mathcal{G}$ is the subgraph that contains the nodes distanced at most k nodes from i, where the distance is understood as the number of edges in the shortest path connecting between the nodes. The k-graph may be considered analogous to extended circular fingerprints, with the center of the circle being a specific node on the GOR and the radius extending to neighboring rings rather than atoms.

Given a probe $\mathbf{y}$, we first determine its nearest ring using the point-to-point distance,

$$i^*(\mathbf{y}) = \arg\min_i \|\mathbf{x}_i - \mathbf{y}\|$$

and then construct the k-graph around i^*. We denote the k-graph constructed in this manner by $\mathcal{G}_k(\mathbf{X}, \mathbf{y})$ and refer to Fig. 4 for a visualization.

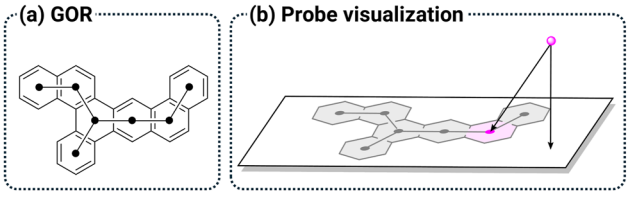


FIG. 3. (a) Graph of rings (GORs) corresponding to a sample molecular system. The circles represent nodes, and the lines between them are the edges. (b) The distances from a random probe to the ring and to the molecular plane in $\mathbb{R}^3$ used as node features.

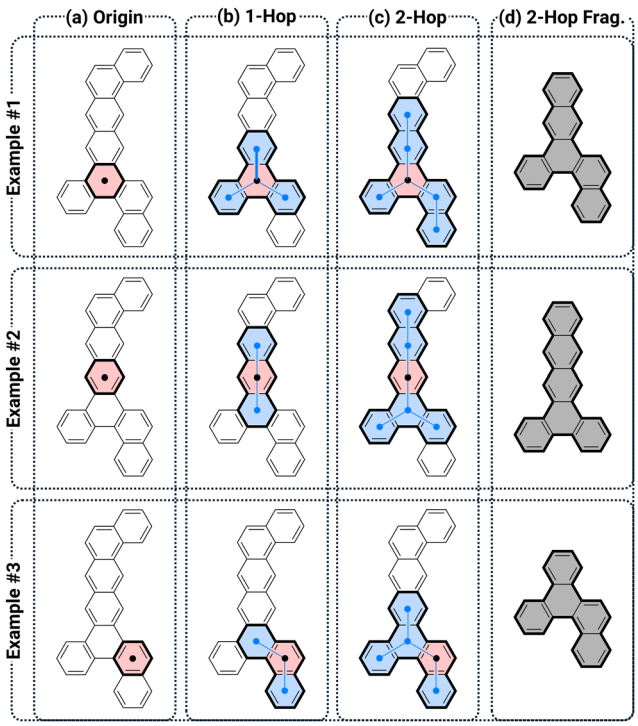


FIG. 4. Three examples of k-graph generation within a single molecule. (a) The node of interest, marked in black with the respective ring in pink. (b) The sub-graphs generated by the 1-hop (blue). (c) The sub-graphs generated by the 2-hop (blue). (d) The molecular fragment corresponding to the 2-hop subgraph (gray).

22 April 2025 09:26:07

© Author(s) 2025

RESULTS AND DISCUSSION

The main objective of this work was to investigate whether the overall magnetic behavior of the molecule is emergent or rather is a superposition of individual local contributions. If the former is true, then partial knowledge of the structure is insufficient to predict the behavior of a given PBH, and trends seen in small molecules are not necessarily transferable to larger ones. Conversely, if the latter is true, then the properties of smaller molecules can be used to predict larger ones. Furthermore, the NICS values at different points in the vicinity of a given polycyclic molecule should be predictable, even if the entire molecular structure is unknown. Regardless of the result, answering this question would provide insight into the behavior of these molecules. However, from a practical perspective, the latter case is preferred, as it enables much faster generation of magnetic data for larger molecules. This can be used to screen molecules but can also inform the inverse design of molecules with desired magnetic response properties.

In the following sections, we first describe the training and validation of the constructed models, their generalization to larger molecules, and the chemical insight that can be extracted.

Initial model training

Our first model was trained only on molecules containing $n_{\text{rings}} \leq 7$ and was, therefore, termed *model*-7. This subset of molecules contains 180 structures, which we surmised would provide sufficient data for training. Training was performed using a [0.8, 0.1, 0.1] split of the 180-molecule subset, where the respective fractions were used for training, testing, and cross-validation for hyperparameter tuning.

To assess the predictive performance of the trained *model*-7, we performed two types of evaluation: in-distribution and out-of-distribution.

In-distribution results

The *model*-7 achieves a remarkable MAE of 0.096 ± 0.051 ppm (computed as the average of MAEs over all molecules). To provide context, we note that Flaig *et al.* benchmarked several wave-function and DFT methods and reported errors of 0.15–0.36 ppm and 2.5–8.6 ppm for the ^{1}H and ^{13}C chemical shifts, respectively.[56] In addition, it is worth noting that previously reported machine learning models for predicting ^{13}C chemical shifts gave an error of ~2–4 ppm.[57–60] For NICS values, which cannot be benchmarked against the experiment, it is generally accepted that the margin of error of DFT calculations is ~1–2 ppm.[13] Therefore, any effort to further improve the accuracy of *model*-7 was deemed unnecessary.

Figure 5 shows a representative example of the results of the model in comparison with the computed DFT results for the same molecule, demonstrating the successful replication of the NICS values by the model.

Out-of-distribution results

Having established that *model*-7 predicts the NICS values well for the systems with $n_{\text{rings}} \leq 7$, we next turned to evaluate its performance on larger molecules, i.e., on molecules out of the distribution of the training data. To this end, we tested the model on increasingly larger molecules, from $n_{\text{rings}} = 8$ to $n_{\text{rings}} = 15$. As the respective MAEs for these evaluations show (Table II), *model*-7 performs well for the smaller test molecules ($n_{\text{rings}} = 8, 9$) but its performance steadily decreases with the size of the molecule, reaching MAEs of

22 April 2025 09:26:07

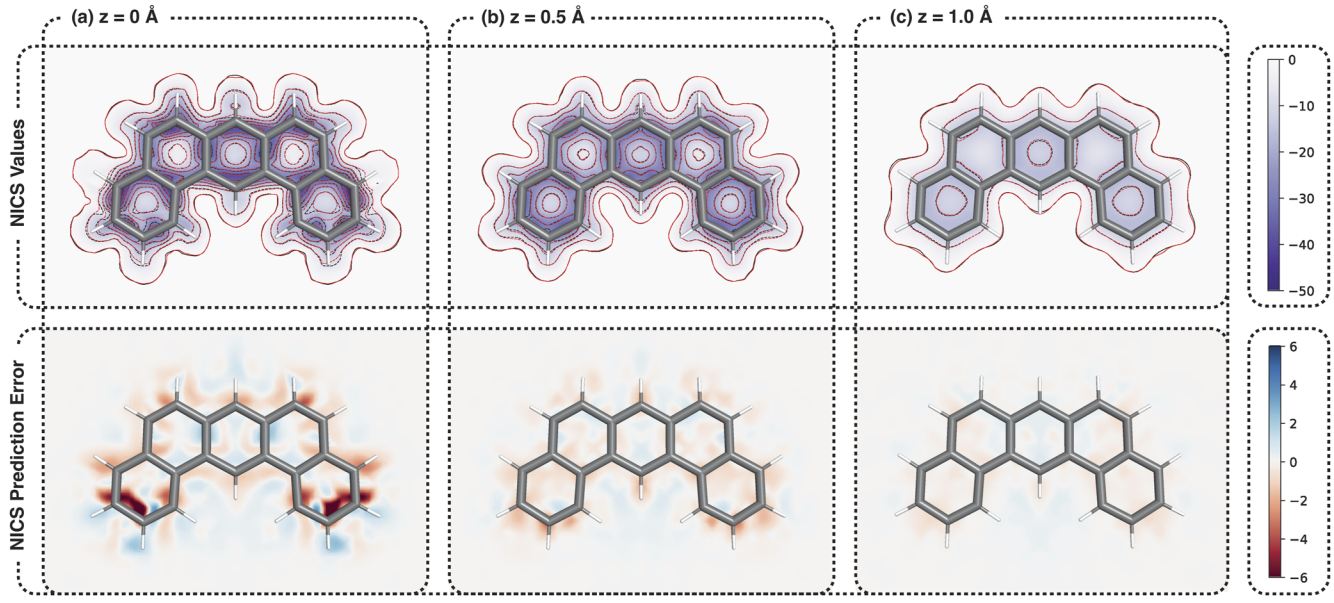


FIG. 5. Cross sections of the 3D NICS values of a sample in-distribution molecule at heights of (a) 0, (b) 0.5, and (c) 1 Å from the molecular plane. Top panel: DFT-computed values (black contour and color image) and values predicted by *model*-7 (red contour); bottom panel: prediction error for each of the plots above, respectively. All the values are in ppm.

© Author(s) 2025

TABLE II. Results of *model*-7 on test molecules, divided by the size of the molecules. The mean is first computed per molecule, and then the mean and standard deviation are computed over all molecules.

No. rings	MAE (ppm)	k-graph MAE (ppm)
1–7	0.112 ± 0.057	N/A
8	0.152 ± 0.040	0.124 ± 0.053
9	0.596 ± 0.226	0.155 ± 0.062
10	1.182 ± 0.321	0.183 ± 0.073
11	1.571 ± 0.327	0.202 ± 0.066
12	2.141 ± 0.443	0.208 ± 0.066
13	2.526 ± 0.608	0.267 ± 0.063
14	3.143 ± 0.563	0.248 ± 0.043
15	3.822 ± 0.969	0.219 ± 0.047

3.143 ± 0.563 and 3.822 ± 0.969, respectively, when presented with the largest molecules in our dataset (n_{rings} = 14, 15). One could interpret these results simply as evidence for the non-locality of the magnetic response. However, we hypothesized that the underlying problem is related to the model itself, rather than the chemical behavior of the molecules. As has been repeatedly reported in the literature, GNNs have difficulty generalizing to larger graphs than those used in training.[61–63]

Solving the generalization problem

Although at this time, there is no universal solution to the generalization problem, we envisioned that our particular dataset and chemical question may provide a unique opportunity for a simple solution, due to two characteristics.

First, in our representation, it is possible to identify a *central* node that has the strongest effect on the property at probe point $\mathbf{y}$, for any given input $\mathcal{G}(\mathbf{X},\mathbf{y})$. This is simply the node closest to $\mathbf{y}$. Second, as described above, the induced magnetic field—and hence also the NICS value—is a result of the induced current and its strength at a given point is directly related to its distance from the current. Indeed, several years ago, our group showed that the NICS values can be replicated from bond currents using the Biot–Savart law,[27] which provides a mathematical basis for this distance-dependent relationship. Thus, we assume that, for any given point in space, the strongest effects should be made by the rings that are closest to that point.

To validate this assumption, we conducted a test using n_{rings} = 15 molecules. We modified the input given to *model*-7, such that for each probe point, we found the closest node $\mathbf{x}_{i^*}$ and identified the *k*-graph around $\mathbf{x}_{i^*}$. This *k*-graph was then associated with the probe as the input for *model*-7, rather than providing the entire molecular structure as input. This entailed that, for the same molecule, different probe points could have different structural inputs (see Fig. 4 for an example of such a case). In addition, depending on the proximity of the closest node $\mathbf{x}_{i^*}$ to the edge of the molecule, different *k*-graphs may have varying numbers of rings despite having the same *k* (compare example No. 3 to example Nos. 1 and 2 in Fig. 4). We performed this *k*-graph generation for a series of *k*-hops, using $k = 1, \ldots, 6$, which allowed us to investigate the relationship between the model's performance and the size of the sub-graph.

The results for the different *k*'s for the n_{rings} = 15 systems can be seen in Fig. 6. We observe that for MAEs calculated over all probes in the test molecules (red trace), the errors increase continuously as the *k* value is increased. However, distinguishing between two different types of probes reveals a more nuanced result and illustrates the dependence on the size of the molecules in the training set. For clarity, we will denote the maximal size of the training molecules as n_{max} (recall that the largest molecules in the training set have n_{rings} = 7). The plots clearly show the difference in prediction performance for probe points with *k*-graph < n_{max} (blue trace) vs those with *k*-graph > n_{max} (green trace). While the former show very low MAEs that steadily decrease with the size of *k*, the latter show markedly larger MAEs that grow rapidly with the size of *k*. That is, for the same exact probe points, *model*-7 has a poorer performance when it associates those points with larger molecular structures. In other words, as long as the probe point is associated with a *k*-graph that is within the distribution of the training data, the model predicts well, even if the molecule itself is outside of the distribution. This suggests that the error is indeed attributable to the well-documented generalization difficulty in GNNs.

Following this initial test, we generalized this approach to all molecules by implementing a simple strategy: given a molecule with a number of rings that exceeds n_{max}, we first created the *k*-graph $\mathcal{G}_k(\mathbf{X},\mathbf{y})$ as previously explained, where *k* is chosen to be the largest natural number such that the number of nodes in the generated *k*-graph $\mathcal{G}_k(\mathbf{X},\mathbf{y})$ does not exceed n_{max}. Then, we applied the trained *model*-7 to $\mathcal{G}_k(\mathbf{X},\mathbf{y})$. Applying this simple algorithm led to a substantial improvement in performance (right-hand column in Table II) and to two main observations. First, we observe that the MAEs of the larger molecules are substantially reduced; for example, for the n_{rings} = 15 molecules, it drops from 3.822 ± 0.969 ppm to a mere 0.219 ± 0.047 ppm (lower than the expected error of the ground-truth DFT calculations). Second, we note that the MAEs are now more or less consistent over a wide range of molecular sizes;

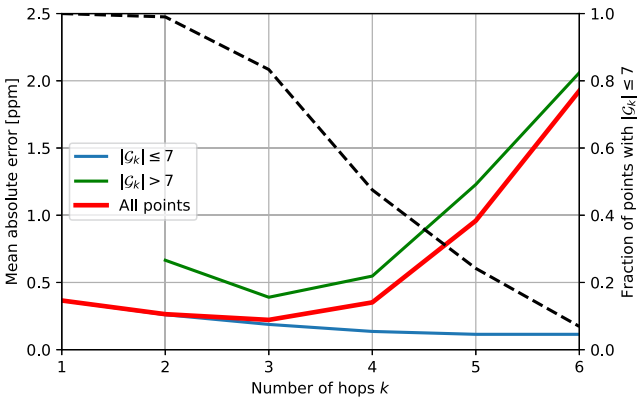


FIG. 6. MAE of *model*-7 on out-of-distribution n_{rings} = 15 molecules when using the *k*-graph as input, plotted as a function of *k*. Red: MAE over all probe points across all molecules. Green: MAE when taking only *k*-graphs > n_{max}. Blue: MAE when taking only *k*-graphs < n_{max}. Black (dashed): the fraction of probe points for which the *k*-graphs < n_{max}.

22 April 2025 09:26:07

© Author(s) 2025

this is yet another evidence that the issue was model-related and not based on chemical behavior. Overall, the model demonstrates the ability to predict the NICS values around molecules containing up to n_{rings} = 15, even though it was trained only on molecules containing up to n_{rings} = 7. As such, this provides evidence for locality.

The extent of locality

As mentioned above, the results of *model*-7 in the previous sections support the hypothesis that the magnetic response is determined by local effects. The natural follow-up is then: how local is "local"? Within the framework of this study, it is possible to address this question by reducing the size of the molecules used in training and, therefore, also reducing the *k*-hop that is used as input. In our initial evaluation, we used an arbitrary cutoff point of n_{rings} = 7, but perhaps a similar predictive ability can be achieved with even smaller subunits? If so, this would not only support the idea of locality but also provide a better quantitative measure of its extent.

To investigate this and further study the scaling laws of the *k*-hop expansion method, we proceeded to train a series of additional models, named "*model*-n_{max}," for n_{max} = 3–6, where n_{max} refers to the size of the largest molecules included in the training set (Table I details the subset of data used for training each of the respective models). Then, for each n_{max}, we used the respective trained *model*-n_{max} to predict the NICS values around all out-of-distribution molecules, i.e., those having $n_{rings} > n_{max}$. All the predictions were made using the approach described above, that is, using *k*-graphs of up to n_{max} rings only.

The results of these evaluations are summarized in Table III, which details the MAEs on the out-of-distribution molecules for each of the *model*-n_{max} variants, respectively. In addition, Fig. 7(a) displays the MAEs for these out-of-distribution molecules, divided according to molecule size. Remarkably, even though for each probe point, the *model*-n_{max} "sees" only a fraction of the entire molecular structure, the predicted NICS values are all within (or better than) expected DFT-level accuracy.

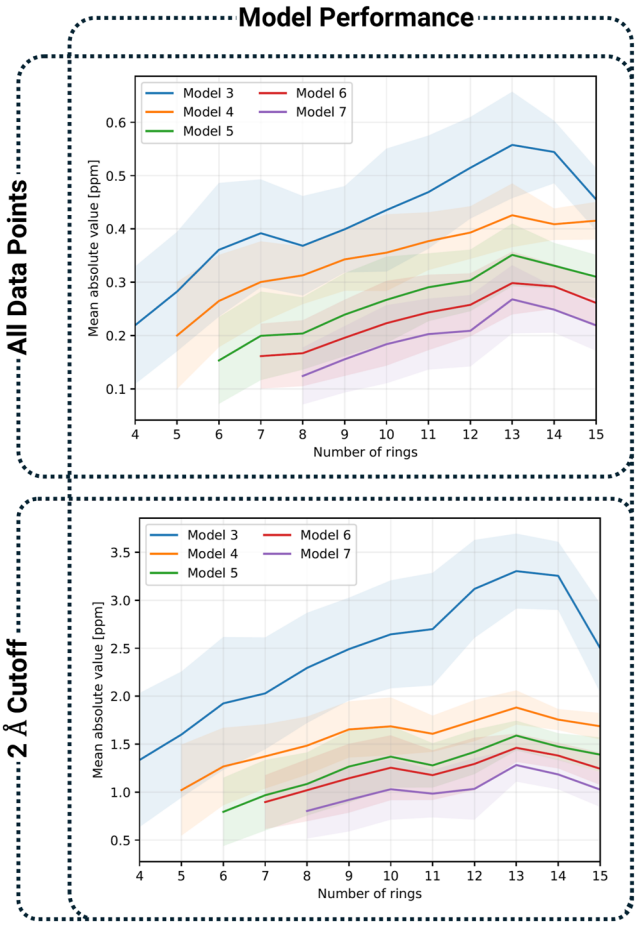


FIG. 7. MAEs of all *model*-n_{max} using k-hop expansion on out-of-distribution molecules, plotted as a function of molecule size. Top panel: MAEs calculated on all data points. Bottom panel: MAEs calculated on all data points up to 2 Å from ring centers.

TABLE III. Results of *model*-3–*model*-7 on test molecules, divided by the size of the molecules. The mean is first computed per molecule, and then the mean and standard deviation are computed over all molecules.

	MAE of model (ppm)				
No. rings	3-model	4-model	5-model	6-model	7-model
4	1.335 ± 0.698				
5	1.600 ± 0.656	1.021 ± 0.474			
6	1.924 ± 0.692	1.267 ± 0.403	0.794 ± 0.355		
7	2.029 ± 0.582	1.372 ± 0.335	0.967 ± 0.367	0.896 ± 0.280	
8	2.294 ± 0.574	1.484 ± 0.302	1.085 ± 0.328	1.019 ± 0.322	0.803 ± 0.286
9	2.490 ± 0.535	1.653 ± 0.292	1.266 ± 0.375	1.144 ± 0.358	0.919 ± 0.329
10	2.644 ± 0.562	1.685 ± 0.301	1.369 ± 0.322	1.254 ± 0.337	1.030 ± 0.316
11	2.699 ± 0.586	1.608 ± 0.188	1.278 ± 0.231	1.177 ± 0.259	0.983 ± 0.246
12	3.118 ± 0.510	1.744 ± 0.214	1.419 ± 0.230	1.293 ± 0.274	1.033 ± 0.320
13	3.303 ± 0.390	1.882 ± 0.178	1.589 ± 0.154	1.462 ± 0.147	1.282 ± 0.172
14	3.254 ± 0.355	1.755 ± 0.109	1.474 ± 0.139	1.381 ± 0.133	1.184 ± 0.154
15	2.502 ± 0.461	1.686 ± 0.137	1.391 ± 0.180	1.244 ± 0.168	1.026 ± 0.179

22 April 2025 09:26:07

© Author(s) 2025

The remarkably good performance led us to hypothesize that perhaps the MAEs obtained in this manner are biased toward lower values by probe points located at large distances from the molecules, which have NICS values close to 0. To provide a more stringent evaluation of the performance, we calculated the MAEs again, this time considering only probe points at distances smaller than 2 Å from the molecule [Fig. 7(b)]. Comparison of the two plots in Fig. 7 clearly shows that MAEs do increase when the farther points are disregarded; nevertheless, for the majority of the models, the MAEs are still satisfactory (smaller than 2.0 ppm). Moreover, it appears that the error begins to plateau, a further indication of non-emergent behavior. The outlier is *model*-3, which shows a substantially larger error and a stronger increase in error with size. This behavior may be explained by the fact that *model*-3 was only trained on four molecules: benzene, naphthalene, anthracene, and phenanthrene. It did not train on triphenylene and, therefore, has difficulty in inferring the field around branching points. Incorporation of the branching motif in *model*-4 shows a significant improvement in prediction ability. For further discussion on this point, we refer the reader to the supplementary material.

The systematic improvement seen in the subsequent models can be attributed to two possible sources: (a) the smaller the model, the fewer training data it trained on (for example, the training set of *model*-4 consisted of only nine molecules, a substantially smaller number than the 144-molecule training set of *model*-7); (b) the smaller the model, the more the test data were out of its training data distribution.

CONCLUSIONS

In this study, we employed a novel deep-learning approach to investigate the locality of magnetic response in polybenzenoid hydrocarbons (PBHs). By implementing a *k*-hop expansion strategy, we successfully overcame the common challenge of poor generalization in GNNs and set the foundation for the current investigation in doing so. Our models accurately predicted NICS values for molecules significantly larger than those in their respective training sets, with consistently low mean absolute errors across a wide range of molecular sizes. Achieving such successful performance using only local structural information supports the hypothesis that the overall magnetic behavior of PBHs emerges from a superposition of individual local contributions rather than from global effects. In particular, our results suggest that a minimal library of building blocks, comprising PBHs of up to only $n_{\text{rings}} = 4$, is sufficient to obtain satisfactory prediction of the induced magnetic field around much larger PBHs.

From a conceptual perspective, our findings contribute to the ongoing debate about local vs global effects in aromatic systems, providing a quantitative framework for understanding the extent of magnetic response locality in PBHs. Reflecting back on the literature debate that compelled us to embark on this study, the conclusions of this work dovetail with those of Bultinck and co-workers from 2006 and 2007. Following careful and insightful analyses,[15,16] the authors reconciled the discrepancies between aromaticity indices by demonstrating that the magnetic criterion (to which NICS belongs) has a different underlying definition of "local aromaticity" than the other metrics. Whereas most indices attempt to evaluate the benzenoid character of a single ring, the magnetic criterion inherently considers larger circuits. Our current results fully concur with this conclusion: in showing that tricyclic (or tetracyclic, in the case of branched molecules) systems are required to predict the behavior of larger systems, we have substantiated the notion that the magnetic criterion does not evaluate local aromaticity at the individual ring level. Furthermore, we have provided a quantitative evaluation of the extent of the locality.

From a practical viewpoint, the locality principle demonstrated here, and the tools developed in the course of this study, enable rapid and accurate prediction of magnetic properties for large PBHs. Importantly, the methodology introduced in this work can be generalized to larger chemical spaces (e.g., heteroaromatics) and to other types of scalar field properties (e.g., Fukui functions, Hirshfeld surfaces, and electron density), providing new tools for predicting properties and reactivity in these important compounds. Both aspects have potential applications in molecular design and virtual screening.

In conclusion, the current study not only advances our understanding of the magnetic behavior of PBHs but also demonstrates the power of interpretable deep-learning approaches in unraveling complex structure–property relationships in chemistry.

SUPPLEMENTARY MATERIAL

The supplementary material provides computational details, complete list of structures included in the dataset, description of model training, and further analysis of the performance of *model*-3.

ACKNOWLEDGMENT

The authors express their gratitude to Professor Dr. Peter Chen for his continued support. The authors acknowledge the Branco Weiss Fellowship for supporting this research as part of a Society in Science grant and the Israel Science Foundation for financial support (Grant No. 1745/23 to R.G.-P.). R.G.-P. is a Branco Weiss Fellow, a Horev Fellow, and an Alon Scholarship recipient. A.M.B. was supported by the ERC StG EARS and the Israeli Science Foundation.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Yair Davidson: Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Validation (equal); Visualization (equal); Writing – original draft (supporting). **Aviad Philipp**: Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Validation (equal); Visualization (equal); Writing – original draft (supporting). **Sabyasachi Chakraborty**: Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Visualization (equal); Writing – original draft (equal). **Alex M. Bronstein**: Conceptualization (equal); Formal analysis (equal); Funding acquisition (equal); Resources (equal); Supervision (equal);

22 April 2025 09:26:07

© Author(s) 2025

Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **Renana Gershoni-Poranne**: Conceptualization (equal); Formal analysis (equal); Funding acquisition (equal); Project administration (equal); Resources (equal); Supervision (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data generated in the course of this work and underpinning the analyses reported herein are openly available on Gitlab at https://gitlab.com/porannegroup/magnetic_locality. The datasets are provided as *.csv* files. Further description of the data structure is provided on the GitLab repository.

REFERENCES

[1] L. Pauling, "The diamagnetic anisotropy of aromatic molecules," J. Chem. Phys. **4**, 673–677 (1936).
[2] K. Y. Lonsdale, "Magnetic anisotropy and electronic structure of aromatic molecules," Proc. R. Soc. London, Ser. A **159**, 149–161 (1937).
[3] R. McWeeny, "Ring currents and proton magnetic resonance in aromatic molecules," Mol. Phys. **1**, 311–321 (1958).
[4] J. A. Pople, "Molecular orbital theory of aromatic ring currents," Mol. Phys. **1**, 175–180 (1958).
[5] T. Dickens and R. Mallion, "Topological ring-currents in conjugated systems," MATCH Commun. Math. Comput. Chem **76**, 297–356 (2016).
[6] D. Sundholm, M. Dimitrova, and R. J. F. Berger, "Current density and molecular magnetic properties," Chem. Commun. **57**, 12362–12378 (2021).
[7] G. Merino, T. Heine, and G. Seifert, "The induced magnetic field in cyclic molecules," Chem.—Eur. J. **10**, 4367–4371 (2004).
[8] R. Islas, T. Heine, and G. Merino, "The induced magnetic field," Acc. Chem. Res. **45**, 215–228 (2012).
[9] P. v. R. Schleyer, C. Maerker, A. Dransfeld, H. Jiao, and N. J. R. van Eikema Hommes, "Nucleus-independent chemical shifts: A simple and efficient aromaticity probe," J. Am. Chem. Soc. **118**, 6317–6318 (1996).
[10] P. v. R. Schleyer, "Introduction: Aromaticity," Chem. Rev. **101**, 1115–1118 (2001).
[11] Z. Chen, C. S. Wannere, C. Corminboeuf, R. Puchta, and P. v. R. Schleyer, "Nucleus-independent chemical shifts (NICS) as an aromaticity criterion," Chem. Rev. **105**, 3842–3888 (2005).
[12] R. Gershoni-Poranne and A. Stanger, "Magnetic criteria of aromaticity," Chem. Soc. Rev. **44**, 6597–6615 (2015).
[13] R. Gershoni-Poranne and A. Stanger, "NICS–nucleus independent chemical shift," in *Aromaticity: Modern Computational Methods and Applications* (Elsevier, 2021), p. 1.
[14] J. A. N. F. Gomes and R. B. Mallion, "Aromaticity and ring currents," Chem. Rev. **101**, 1349–1384 (2001).
[15] P. Bultinck, S. Fias, and R. Ponec, "Local aromaticity in polycyclic aromatic hydrocarbons: Electron delocalization versus magnetic indices," Chem.—Eur. J. **12**, 8813–8818 (2006).
[16] P. Bultinck, "Critical analysis of the local aromaticity concept in polyaromatic hydrocarbons," Faraday Discuss. **135**, 347–365 (2007).
[17] J. Aihara and H. Kanno, "Local aromaticities in large polyacene molecules," J. Phys. Chem. A **109**, 3717–3721 (2005).
[18] E. Clar, in *Mobile Source Emissions Including Policyclic Organic Species*, edited by D. Rondia, M. Cooke, and R. K. Haroz (Springer Netherlands, Dordrecht, 1983), pp. 49–58.
[19] M. Solà, "Forty years of Clar's aromatic π-sextet rule," Front. Chem. **1**, 66422 (2013).
[20] C. Glidewell and D. Lloyd, "Mndo study of bond orders in some conjugated BI- and tri-cyclic hydrocarbons," Tetrahedron **40**, 4455–4472 (1984).
[21] M. Randić, "Conjugated circuits and resonance energies of benzenoid hydrocarbons," Chem. Phys. Lett. **38**, 68–70 (1976).
[22] I. Gutman and M. Randić, "A correlation between Kekulē valence structures and conjugated circuits," Chem. Phys. **41**, 265–270 (1979).
[23] S. Nikolić, N. Trinajstić, and D. J. Klein, "The conjugated-circuit model," Comput. Chem. **14**, 313–322 (1990).
[24] M. Mandado, "Determination of London susceptibilities and ring current intensities using conjugated circuits," J. Chem. Theory Comput. **5**, 2694–2701 (2009).
[25] R. Gershoni-Poranne, "Piecing it together: An additivity scheme for aromaticity using NICS-XY scans," Chem.—Eur. J. **24**, 4165–4172 (2018).
[26] P. Finkelstein and R. Gershoni-Poranne, "An additivity scheme for aromaticity: The heteroatom case," ChemPhysChem **20**, 1508–1520 (2019).
[27] E. Paenurk, S. Feusi, and R. Gershoni-Poranne, "Predicting bond-currents in polybenzenoid hydrocarbons with an additivity scheme," J. Chem. Phys. **154**, 024110 (2021).
[28] L. Leyva-Parra, I. Casademont-Reig, R. Pino-Rios, L. Ruiz, M. Alonso, and W. Tiznado, "New perspectives on delocalization pathways in aromatic molecular chameleons," ChemPhysChem **25**, e202400271 (2024).
[29] L. Leyva-Parra, R. Pino-Rios, D. Inostroza, M. Solà, M. Alonso, and W. Tiznado, "Aromaticity and magnetic behavior in benzenoids: Unraveling ring current combinations," Chem.—Eur. J. **30**, e202302415 (2024).
[30] T. Weiss, A. Wahab, A. M. Bronstein, and R. Gershoni-Poranne, "Interpretable deep-learning unveils structure–property relationships in polybenzenoid hydrocarbons," J. Org. Chem. **88**, 9645–9656 (2023).
[31] P. W. Fowler and W. Myrvold, "The "anthracene problem": Closed-form conjugated-circuit models of ring currents in linear polyacenes," J. Phys. Chem. A **115**, 13191–13200 (2011).
[32] M. K. Cyrański, B. T. Stępień, and T. M. Krygowski, "Global and local aromaticity of linear and angular polyacenes," Tetrahedron **56**, 9663–9667 (2000).
[33] R. Firouzi and M. Zahedi, "Polyacenes electronic properties and their dependence on molecular size," J. Mol. Struct.: THEOCHEM **862**, 7–15 (2008).
[34] A. Wahab, L. Pfuderer, E. Paenurk, and R. Gershoni-Poranne, "The COMPAS Project: A computational database of polycyclic aromatic systems. Phase 1: Catacondensed polybenzenoid hydrocarbons," J. Chem. Inf. Model. **62**, 3704–3713 (2022).
[35] A. D. Becke, "Density-functional thermochemistry. III. The role of exact exchange," J. Chem. Phys. **98**, 5648–5652 (1993).
[36] F. Weigend and R. Ahlrichs, "Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy," Phys. Chem. Chem. Phys. **7**, 3297–3305 (2005).
[37] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, "A consistent and accurate *ab initio* parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu," J. Chem. Phys. **132**, 154104 (2010).
[38] S. Grimme, S. Ehrlich, and L. Goerigk, "Effect of the damping function in dispersion corrected density functional theory," J. Comput. Chem. **32**, 1456–1465 (2011).
[39] E. R. Johnson and A. D. Becke, "A post-Hartree–Fock model of intermolecular interactions," J. Chem. Phys. **123**, 024101 (2005).
[40] F. Neese, "The ORCA program system," Wiley Interdiscip. Rev.: Comput. Mol. Sci. **2**, 73–78 (2012).
[41] F. Neese, "Software update: The ORCA program system, version 4.0," Wiley Interdiscip. Rev.: Comput. Mol. Sci. **8**, e1327 (2018).
[42] A. Stanger, "NICS—Past and present," Eur. J. Org Chem. **2020**, 3120–3127.
[43] D. Kilymis, A. P. Bartók, C. J. Pickard, A. C. Forse, and C. Merlet, "Efficient prediction of nucleus independent chemical shifts for polycyclic aromatic hydrocarbons," Phys. Chem. Chem. Phys. **22**, 13746–13755 (2020).
[44] M. Orozco-Ic, J. Barroso, N. D. Charistos, A. Muñoz-Castro, and G. Merino, "Consequences of curvature on induced magnetic field: The case of helicenes," Chem.—Eur. J. **26**, 326–330 (2020).
[45] G. Sánchez-Sanz, I. Alkorta, C. Trujillo, and J. Elguero, "A theoretical NMR study of the structure of benzynes and some of their carbocyclic and heterocyclic analogs," Tetrahedron **68**, 6548–6556 (2012).

22 April 2025 09:26:07

© Author(s) 2025

[46]R. Ditchfield, "Molecular orbital theory of magnetic shielding and magnetic susceptibility," J. Chem. Phys. **56**, 5688–5691 (1972).
[47]K. Wolinski, J. F. Hinton, and P. Pulay, "Efficient implementation of the gauge-independent atomic orbital method for NMR chemical shift calculations," J. Am. Chem. Soc. **112**, 8251–8260 (1990).
[48]J. Gauss, "Effects of electron correlation in the calculation of nuclear magnetic resonance chemical shifts," J. Chem. Phys. **99**, 3629–3643 (1993).
[49]J. R. Cheeseman, G. W. Trucks, T. A. Keith, and M. J. Frisch, "A comparison of models for calculating nuclear magnetic resonance shielding tensors," J. Chem. Phys. **104**, 5497–5509 (1996).
[50]T. Yanai, D. P. Tew, and N. C. Handy, "A new hybrid exchange–correlation functional using the Coulomb-attenuating method (CAM-B3LYP)," Chem. Phys. Lett. **393**, 51–57 (2004).
[51]T. H. Dunning, Jr., "Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen," J. Chem. Phys. **90**, 1007–1023 (1989).
[52]M. J. Frisch *et al.*, *GAUSSIAN 09, Revision D.01*, Gaussian, Inc., Wallingford CT, 2013.
[53]S. Fite, A. Wahab, E. Paenurk, Z. Gross, and R. Gershoni-Poranne, "Text-based representations with interpretable machine learning reveal structure–property relationships of polybenzenoid hydrocarbons," J. Phys. Org. Chem. **36**, e4458 (2023).
[54]F. B. Fuchs, D. E. Worrall, V. Fischer, and M. Welling, "SE(3)-transformers: 3D roto-translation equivariant attention networks," arXiv:2006.10503 (2020).
[55]M. Sypetkowski, F. Wenkel, F. Poursafaei, N. Dickson, K. Suri, P. Fradkin, and D. Beaini, "On the scalability of GNNs for molecular graphs," arXiv:2404.11568 (2024).
[56]D. Flaig, M. Maurer, M. Hanni, K. Braunger, L. Kick, M. Thubauville, and C. Ochsenfeld, "Benchmarking hydrogen and carbon NMR chemical shifts at HF, DFT, and MP2 levels," J. Chem. Theory Comput. **10**, 572–578 (2014).
[57]P. Gao, J. Zhang, Q. Peng, J. Zhang, and V.-A. Glezakou, "General protocol for the accurate prediction of molecular 13C/1H NMR chemical shifts via machine learning augmented DFT," J. Chem. Inf. Model. **60**, 3746–3754 (2020).
[58]W. Gerrard, L. A. Bratholm, M. J. Packer, A. J. Mulholland, D. R. Glowacki, and C. P. Butts, "IMPRESSION – prediction of NMR parameters for 3-dimensional chemical structures using machine learning with near quantum chemical accuracy," Chem. Sci. **11**, 508–515 (2020).
[59]A. Gupta, S. Chakraborty, and R. Ramakrishnan, "Revving up ^{13}C NMR shielding predictions across chemical space: Benchmarks for atoms-in-molecules kernel machine learning with new data for 134 kilo molecules," Mach. Learn.: Sci. Technol. **2**, 035010 (2021).
[60]D. Lemm, G. F. von Rudorff, and O. A. von Lilienfeld, "Impact of noise on inverse design: The case of NMR spectra matching," Digital Discovery **3**, 136–144 (2024).
[61]D. Buffelli, P. Lio, and F. Vandin, "SizeShiftReg: A regularization method for improving size-generalization in graph neural networks," in *Advances in Neural Information Processing Systems* (NeurIPS, 2022), pp. 31871–31885.
[62]B. Bevilacqua, Y. Zhou, and B. Ribeiro, "Size-invariant graph representations for graph classification extrapolations," in *International Conference on Machine Learning* (NeurIPS, 2021), pp. 837–851.
[63]G. Yehudai, E. Fetaya, E. Meirom, G. Chechik, and H. Maron, From Local Structures to Size Generalization in Graph Neural Networks (2021).

22 April 2025 09:26:07

© Author(s) 2025