Check for updates

RESEARCH ARTICLE

Local invertibility and sensitivity of atomic structure-feature mappings

[version 1; peer review: 2 approved]

Sergey N. Pozdnyakov[1], Liwei Zhang[2], Christoph Ortner[2], Gábor Csányi[3], Michele Ceriotti[1]

[1]Laboratory of Computational Science and Modelling, Institute of Materials, Federal Institute of Technology (EPFL) CH-1015 Lausanne, Lausanne, 1015, Switzerland
[2]Department of Mathematics, University of British Columbia, Vancouver, British Columbia, V6T 1Z2, Canada
[3]Engineering Laboratory, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, UK

v1 **First published:** 20 Oct 2021, **1**:126
https://doi.org/10.12688/openreseurope.14156.1
Latest published: 20 Oct 2021, **1**:126
https://doi.org/10.12688/openreseurope.14156.1

Abstract

Background: The increasingly common applications of machine-learning schemes to atomic-scale simulations have triggered efforts to better understand the mathematical properties of the mapping between the Cartesian coordinates of the atoms and the variety of representations that can be used to convert them into a finite set of symmetric *descriptors* or *features*.
Methods: Here, we analyze the sensitivity of the mapping to atomic displacements, using a singular value decomposition of the Jacobian of the transformation to quantify the sensitivity for different configurations, choice of representations and implementation details.
Results: We show that the combination of symmetry and smoothness leads to mappings that have singular points at which the Jacobian has one or more null singular values (besides those corresponding to infinitesimal translations and rotations). This is in fact desirable, because it enforces physical symmetry constraints on the values predicted by regression models constructed using such representations. However, besides these symmetry-induced singularities, there are also spurious singular points, that we find to be linked to the *incompleteness* of the mapping, i.e. the fact that, for certain classes of representations, structurally distinct configurations are not guaranteed to be mapped onto different feature vectors. Additional singularities can be introduced by a too aggressive truncation of the infinite basis set that is used to discretize the representations.
Conclusions: These results exemplify the subtle issues that arise when constructing symmetric representations of atomic structures, and provide conceptual and numerical tools to identify and investigate them in both benchmark and realistic applications.

Open Peer Review

Approval Status ✓ ✓

	1	2
version 1 20 Oct 2021	✓ view	✓ view

1. **Reinhard Maurer**, Technical University of Munich, Munich, Germany
University of Warwick, Coventry, UK

2. **Bastiaan Braams**, Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands

Any reports and responses or comments on the article can be found at the end of the article.

Keywords
atomistic simulations, machine-learning, structural descriptors



This article is included in the European Research Council (ERC) gateway.



This article is included in the Horizon 2020 gateway.



This article is included in the Computational Chemistry collection.

Corresponding author: Michele Ceriotti (michele.ceriotti@epfl.ch)

Author roles: Pozdnyakov SN: Data Curation, Formal Analysis, Software, Visualization, Writing – Review & Editing; **Zhang L**: Formal Analysis; **Ortner C**: Conceptualization, Formal Analysis, Writing – Original Draft Preparation, Writing – Review & Editing; **Csányi G**: Conceptualization, Writing – Original Draft Preparation, Writing – Review & Editing; **Ceriotti M**: Conceptualization, Data Curation, Formal Analysis, Visualization, Writing – Original Draft Preparation, Writing – Review & Editing

Competing interests: No competing interests were disclosed.

Grant information: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement Nos 677013 and 101001890). MC and SNP acknowledge support from the Swiss National Science Foundation [Project No. 200021-182057] and from the NCCR MARVEL, funded by the Swiss National Science Foundation (SNSF). CO is supported by Leverhulme Research Project Grant RPG-2017-191 and by the Natural Sciences and Engineering Research Council of Canada (NSERC) [funding reference number IDGR019381].
The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2021 Pozdnyakov SN *et al*. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Pozdnyakov SN, Zhang L, Ortner C *et al.* **Local invertibility and sensitivity of atomic structure-feature mappings [version 1; peer review: 2 approved]** Open Research Europe 2021, **1**:126 https://doi.org/10.12688/openreseurope.14156.1

First published: 20 Oct 2021, **1**:126 https://doi.org/10.12688/openreseurope.14156.1

Plain language summary

The transformation of the atomic coordinates of a molecule or a material to a more symmetric mathematical form is the first step in the application of machine-learning techniques to atomic-scale simulations. The properties of such mappings – e.g. how much the symmetric descriptors change upon deformations of the original structure – have a knock-on effect on any model built on top of it. We study some of the most widely used families of descriptors, revealing how both their fundamental nature and the details of how they are implemented affect the sensitivity of the descriptors and thus the accuracy of the subsequent machine-learning models.

Introduction

There has been a tidal wave of interest in the last decade in applying machine learning tools to atomistic modelling problems. See for example the recent thematic issue of Chemical Reviews for a collection of review articles[1]. In this first "heroic phase" of the development of this field, authors used a wide variety of encodings of atomic structure and regression methods to make models. While it was widely recognized that it is advantageous to encode the physical symmetries of translation, rotation and permutation invariance into *descriptors* of atomic structure, there was little enthusiasm (or opportunity) to rigorously evaluate desirable general properties of different descriptors, as well as compare them with one another, independently of the regression methods and specific applications.

A number of recent papers have taken on such challenges. It is now understood that many of the descriptors based on the local neighborhood density are equivalent in the limit of high resolution[2] and can be derived from body-ordered expansions of a suitably defined atomic density[3,4].

In the present paper we will continue our theoretical investigation of representations of *local atomic environments*, A_i, given in terms of a feature vector $\boldsymbol{\xi} = \boldsymbol{\xi}(A_i) = \{\xi_q(A_i)\}_{q=1\ldots n_{\rm feat}}$ that is invariant under rotations, reflections and permutations of like atoms. Such a representation immediately leads to the question of whether the atomic environment A_i can be reconstructed from the features $\boldsymbol{\xi}$, up to symmetries. In the context of representing atomic environments (and global structures) this was first explored in some detail in Ref. 5 where it was immediately observed that the invariance of the features under said symmetries makes this a formidable theoretical challenge. For example, it was shown that any descriptor based on three-body features cannot uniquely identity *every* configuration containing four or more neighbors, while there are configurations with seven or more neighbors that cannot be distinguished by any descriptor based on four-body features. That is, pairs of atomic environments can be constructed that are indistinguishable under these descriptors. These observations apply to the vast majority of descriptors used in the field, including in particular[6,7]; see Ref. 2 for an extensive discussion.

This challenge points to the fundamental question of under which conditions the feature vector $\boldsymbol{\xi}$ is a *coordinate system*, or in other words, whether its image is a smooth manifold. Since regression, classification and reconstruction tasks are primarily undertaken in feature space, this would be a highly desirable property for the performance of algorithms employed in such tasks. Aside from the injectivity of $\boldsymbol{\xi}$ alluded to in the previous paragraph, and which encapsulates the global structure, the immediate next question is to understand its *local* structure. That is, we will investigate whether $\boldsymbol{\xi}$ provides a local coordinate system, i.e. is locally smooth and invertible. This can be adressed by studying the *sensitivity* of $\boldsymbol{\xi}$ through properties of its Jacobian matrix. In the context of atomic structures, similar studies were first undertaken in Ref. 8,9, where the sensitivity of ξ was related to the accuracy of machine-learning models based on the features, and in Ref. 10, where it was used to identify regions with near-constant value of the features ξ_q.

The purpose of the present work is to explore in more detail the issues of sensitivity, local invertibility, and stability of descriptors constructed from symmetrized ν-correlations of the local atomic density. In high symmetry configurations a *natural* loss of sensitivity occurs for all smooth and invariant features (we explain this in detail in the main text), which may even be beneficial for regression tasks. A more disturbing observation, analogous to the degenerate pairs discovered in Ref. 5, is that certain popular descriptors also exhibit a loss of sensitivity in *non-symmetric* configurations. We will demonstrate that spurious singularities (loss of sensitivity in non-symmetric structures) arise by two mechanisms: (1) a lack of numerical resolution in the discretization of the atomic density, which is of course easily remedied; and (2) as intersections of degenerate pair manifolds, which are more fundamental and non-trivial to remove or rule out.

Theory

Symmetry-adapted ν-correlations

Our main focus is on a class of representations that correspond to ν-point correlations of the atom density, $\rho_i(\mathbf{r})$, expressed relative to an environment A_i, a finite neighborhood centered on the i-th atom. The density can be written, using the notation introduced in Refs. 3,11, and formalized in Ref. 2, as

$$\langle \mathbf{x} | \rho_i \rangle = \sum_{j \in A_i} \langle \mathbf{x} | \mathbf{r}_{ji}; g \rangle, \tag{1}$$

where $\langle \mathbf{x}|\mathbf{r}_{ji};\ g\rangle \equiv g(\mathbf{x} - \mathbf{r}_{ji})$ is a Gaussian of width σ_a centered on the vector $\mathbf{r}_{ji} = \mathbf{r}_j - \mathbf{r}_i$ that separates the central atom i and its j-th neighbor, with j running over the indices of atoms within A_i. The symmetrized ν-point correlations are obtained by integrating over rotations (and inversion) tensor products of this density

$$\langle \mathbf{x}_1; \ldots \mathbf{x}_\nu | \overline{\rho_i^{\otimes \nu}} \rangle = \int \mathrm{d}\hat{R}\, \langle \mathbf{x}_1 | \hat{R} | \rho_i \rangle \cdots \langle \mathbf{x}_\nu | \hat{R} | \rho_i \rangle . \tag{2}$$

These definitions are quite abstract, but encompass a majority of the representations that have been used in the application of machine learning to atomistic problems – including atom-centered symmetry functions[6,12,13], smooth overlap of atomic positions (SOAP) powerspectrum[14], bispectrum[7], Faber-Christensen-Huang-von Lilienfeld (FCHL) descriptors[15], all of which are limited to low correlation orders, as well as representations for which ν can be increased systematically, such as the moment tensor potential (MTP)[16], the atomic cluster expansion (ACE)[4] and the N-body iterative contraction of equivariants (NICE)[17].

In practical use, the density-correlation is discretized into a feature vector using a basis, and the choice of basis can have an impact on how effectively structural information can be stored in the corresponding feature vector (see e.g. Ref. 18). However, we shall first focus on the nature of density-correlation representations in the complete basis set limit. For instance, it was recently shown that three- and four-body correlations (corresponding to ν = 2, 3) are *incomplete*, i.e. it is possible to find pairs of *degenerate* environments, A_i^+ and A_i^- that are not related by symmetry, but have the same $|\overline{\rho_i^{\otimes 2}}\rangle$ or $|\overline{\rho_i^{\otimes 3}}\rangle$ representation[5]. It is important to stress that this is true of any descriptor that is a discretized version of these low-order representations, and that non-linear models built on top of such descriptors, no matter how complicated or sophisticated, cannot eliminate this fundamental shortcoming.

Overlap matrix representations

We also consider a second class of representations of structures that are derived from the eigenvalues of generalized distance matrices[19], and which are far less well understood both theoretically and in applications. Because the distance matrix is non-linearly transformed, and combined with angular terms into a form that resembles orbital overlap matrices from quantum chemistry, this representation is usually referred to as "overlap matrix fingerprints" (OMFP). It was observed numerically that OMFPs are able to distinguish the degenerate atomic environment pairs identified in Ref. 5, which are indistinguishable by low-order correlation features. To the best of our knowledge, however, no theoretical framework exists to explain this, nor any results suggesting that this is a general property and that no degenerate structures exist for OMFP descriptors.

To construct a set of OMFP, one begins by specifying an overlap matrix, $\mathbf{T}$, analogous to a non-orthogonal tight-binding model with N_{orb} *artificial* orbitals per atom. For an atomic environment A_i that has N atoms, $\mathbf{T}$ has $N \times N_{\mathrm{orb}}$ rows and columns, and has blocks

$$T_{jj'} = f_{\mathrm{cut}}(r_{ij}) f_{\mathrm{cut}}(r_{ij'}) t(\mathbf{r}_{jj'}) \in \mathbb{C}^{N_{\mathrm{orb}} \times N_{\mathrm{orb}}}, \tag{3}$$

where t is some non-linear transformation of the interatomic distances. The OMFP is then the ordered spectrum $\{\tau_k\}_{k=1}^{N \times N_{\mathrm{orb}}}$ of $\mathbf{T}$. If $\mathbf{T}$ is covariant to rotations then the spectrum will be invariant. In our computational experiments in the present work we will use this form, as proposed in 19, in particular using N_{orb} = 4 (s and p orbitals). In order to ensure consistency, we compute OMFP with the same implementation used in previous studies[8,19].

Due to eigenvalue crossings the spectrum is non-smooth: the derivatives of the ordered set of eigenvalues with respect to atomic positions are discontinuous. These discontinuities can be removed by an averaging procedure[8] or – more simply – by projecting the spectrum onto a basis, e.g. of polynomials,

$$\langle n | A_i; \mathbf{T} \rangle \coloneqq \sum_k (\tau_k)^n = \mathrm{tr}(\mathbf{T}^n). \tag{4}$$

More importantly for us, expressing OMFPs in a polynomial basis reveals that the feature $\langle n|A_i; \mathbf{T} \rangle$ can be explicitly written as an n-correlation, because it is given by the sum of products of elements of $\mathbf{T}$ with n factors (see Section 5.3 and Equation (91) in Ref. 2 for the details). That is, the OMFP contain *some* of the density correlation features up to correlation order N. The open question is to understand whether these have genuinely high correlation order or are actually just high-order polynomials of low correlation-order features such as in the example discussed in Sec. VI.C of Ref. 2. To the best of our knowledge, all applications of OMFP thus far have used the ordered spectrum directly, hence there is no established choice of a basis to obtain smooth features (using the monomials in (4) would lead to numerically instabilities). In what follows we will use, as a proof of concept only, a naive basis of sine and cosine functions

$$\begin{aligned} \langle (2q) | A_i; \mathbf{T} \rangle &\coloneqq \sum_k \cos(q\tau_k/K) \\ \langle (2q+1) | A_i; \mathbf{T} \rangle &\coloneqq \sum_k \sin(q\tau_k/K), \end{aligned} \tag{5}$$

where K is a parameter of the order of the range spanned by the spectrum. This enables us to benchmark the impact of the lack of smoothness of the ordered spectrum in our tests. A more careful assessment of different bases is left for future investigation.

Sensitivity and the Jacobian

We want to assess whether the mapping from structure to features results, at least locally, in a *coordinate system*, i.e., whether it is possible to relate changes in the features to changes in the structure in a one-to-one manner. This question is directly connected to the question of the sensitivity of the features to a deformation of the structure, which has been the subject of recent investigation[8,9]. The central quantity that contains the answer to these questions is the Jacobian $\mathbf{J}$, which for the environment-centered density $|\rho_i\rangle$ reads

$$J_{j\alpha,\mathbf{x}}(\rho_i) = \frac{\partial}{\partial r_j^\alpha} \langle \mathbf{x} | \rho_i \rangle = \frac{\partial}{\partial r_j^\alpha} \langle \mathbf{x} | \mathbf{r}_{ji}; g \rangle = \frac{x^\alpha - r_{ji}^\alpha}{\sigma_a^2} \langle \mathbf{x} | \mathbf{r}_{ji}; g \rangle \equiv \langle \mathbf{x} | \mathbf{r}_{ji}; \partial_\alpha g \rangle, \tag{6}$$

where j enumerates the neighbors of the central atom, α indicates one of the x, y or z Cartesian coordinates, and the notation $\partial_\alpha g$ indicates the derivative of a 3D Gaussian with respect to the α direction. $J_{j\alpha,\mathbf{x}}$ is an infinite-dimensional operator, a generalization of the Jacobian matrix, that for a finite feature vector $\boldsymbol{\xi}$ of length n_{feat} contains $3N$ rows, corresponding to the coordinates of the N neighboring atoms, and n_{feat} columns, corresponding to the number of indices that enumerate the features for a given discretization of the representation. Note that this is the transpose of the most common definition of the Jacobian, which we chose due to the analogy with a design matrix where each column corresponds to a feature.

The singular values, $s_k \geq 0$, $k = 1, \ldots, 3N$, of the Jacobian matrix (or operator) $\mathbf{J}$ can be obtained by performing a singular value decomposition, or by computing the square root of the eigenvalues of $\mathbf{J}\mathbf{J}^*$ (the "sensitivity matrix" of Ref. 8), and identify the principal modes of variation of the representation. These singular values indicate how much the features change when the atoms are distorted by an infinitesimal amount according to the displacement patterns associated with the corresponding singular vector (i.e., eigenvector of $\mathbf{J}\mathbf{J}^*$).

Sensitivity, local invertibility and global invertibility

Consider a smooth descriptor $\boldsymbol{\xi} : \mathbb{R}^{3N} \to \mathbb{R}^d$, $d \geq 3N$, with Jacobian $\mathbf{J} : \mathbb{R}^{3N} \to \mathbb{R}^{3N \times d}$. We use $\mathbf{u} \equiv \{\mathbf{r}_{ji}\}_{j=1}^N$ to indicate the Cartesian coordinates of the neighbors within A_i. If $\mathbf{J}(\mathbf{u}_0)$ has full rank $3N$, then the image of $\boldsymbol{\xi}$ is a smooth manifold locally around $\boldsymbol{\xi}(\mathbf{u}_0)$ and the mapping $\boldsymbol{\xi}$ is locally invertible around $\mathbf{u}_0$. In other words, $\boldsymbol{\xi}$ is a coordinate system, locally around $\mathbf{u}_0$. We can also say in this case that $\boldsymbol{\xi}$ is *sensitive to small perturbations*. The degree of sensitivity may be measured in terms of the singular values of $\mathbf{J}$.

By contrast, in Ref. 5 we explored the seemingly much more stringent condition of *global* invertibility of a descriptor, i.e., injectivity of the mapping $\mathbf{u} \mapsto \boldsymbol{\xi}(\mathbf{u})$ on the set of all admissible configurations (atomic environments).

Sensitivity of ρ_i and its discretization

Although this work is primarily concerned with sensitivity of symmetry-adapted features, it is nevertheless instructive to first consider the sensitivity of ρ_i itself. We will briefly summarize the effect that the smearing σ_a and the discretization have on the sensitivity of $\langle \mathbf{x}|\rho_i\rangle$, which highlight important considerations for practical implementation. For the sake of clarity of presentation we will consider the specific case when g is a Gaussian, but our analysis applies whenever g is analytic and rapidly decaying, i.e. a "smeared Dirac delta".

To that end, we compute the scalar product of two rows in the Jacobian of the neighbor density (6):

$$\int \mathrm{d}\mathbf{x}\, \langle \partial_{\alpha'} g; \mathbf{r}_{j'i}|\mathbf{x}\rangle \langle \mathbf{x}|\mathbf{r}_{ji}; \partial_\alpha g\rangle = \frac{e^{-\frac{(\mathbf{r}_{ji}-\mathbf{r}_{j'i})^2}{4\sigma_a^2}}}{(4\pi\sigma_a^2)^{3/2}} \frac{2\sigma_a^2\delta_{\alpha\alpha'} - (r_{j'i}^{\alpha'} - r_{ji}^{\alpha'})(r_{j'i}^{\alpha} - r_{ji}^{\alpha})}{4\sigma_a^4}. \tag{7}$$

In the limit in which the contributions to the density from individual atoms do not overlap ($\mathbf{r}_{jj'}/\sigma_a \to \infty$), the Gaussian term tends to $\delta_{jj'}$, and the scalar product reduces to $\delta_{jj'}\delta_{\alpha\alpha'}/(16\pi^{3/2}\sigma_a^5)$. Thus, the rows of $\mathbf{J}$ are orthogonal in this limit, and all the singular values of $\mathbf{J}$ are equal to $1/\sqrt{16\pi^{3/2}\sigma_a^5}$. In particular we obtain that the condition number of the Jacobian $\mathrm{cond}(\mathbf{J}(\rho_i)) \to 1$, which is indeed the strongest notion of sensitivity we can hope for.

In other words, in this regime where all atoms do not overlap relative to the smearing parameter σ_a, $|\rho_i\rangle$ is equally sensitive to the displacement of each neighbor, independent of the structure considered.

In practical simulations, the density ρ_i will be discretized by projecting it onto an orthogonal basis $\{\phi_q\}$,

$$\langle q|\rho_i\rangle = \int \phi_q(\mathbf{x}) \langle \mathbf{x}|\rho_i\rangle \,\mathrm{d}\mathbf{x} = \sum_j \int \phi_q(\mathbf{x}) g(\mathbf{x}-\mathbf{r}_{ji})\,\mathrm{d}\mathbf{x},$$

and denote the finite feature vector with $\boldsymbol{\xi} = \{\xi_q\}_{q=1}^{n_{\mathrm{feat}}} = \{\langle q|\rho_i\rangle\}_{q=1}^{n_{\mathrm{feat}}}$ where q ranges over a finite index-set. We must now consider whether the perfect sensitivity of ρ_i (in the regime $\mathbf{r}_{jj'}/\sigma \to \infty$) survives under this discretization. A straightforward calculation yields

$$J_{j\alpha,q}(\boldsymbol{\xi}) = \frac{\partial \langle q|\rho_i\rangle}{\partial r_j^\alpha} = \int \frac{\partial \rho_i}{\partial r_j^\alpha}\phi_q(\mathbf{x})\,\mathrm{d}\mathbf{x} = \int \frac{\partial g(\mathbf{x}-\mathbf{r}_j)}{\partial r_j^\alpha}\phi_q(\mathbf{x})\,\mathrm{d}\mathbf{x} \equiv \langle q|\mathbf{r}_{ji}; \partial_\alpha g\rangle. \tag{8}$$

Observe that $J_{j\alpha,q}$ are the L^2-projection coefficients of $\frac{\partial \rho_i}{\partial r_j^\alpha}$, hence it is natural to consider the projection, Π, over the *finite* basis set

$$\Pi\left[\frac{\partial \rho_i}{\partial r_j^\alpha}\right](\mathbf{x}) := \sum_{q=1}^{n_{\mathrm{feat}}} \langle q|\mathbf{r}_{ji}; \partial_\alpha g\rangle \phi_q(\mathbf{x}) = \sum_{q=1}^{n_{\mathrm{feat}}} J_{j\alpha,q}(\boldsymbol{\xi})\phi_q(\mathbf{x}). \tag{9}$$

With this definition we obtain that we can rewrite the inner product of two rows of $J(\boldsymbol{\xi})$, which is now a sum over the features, as

$$(JJ^*)_{j\alpha,j'\alpha'} = \sum_{q=1}^{n_{\mathrm{feat}}} J_{j\alpha,q} J^*_{j'\alpha',q} = \int \Pi\left[\frac{\partial \rho_i}{\partial r_j^\alpha}\right]\Pi\left[\frac{\partial \rho_i}{\partial r_{j'}^{\alpha'}}\right], \tag{10}$$

which is now identical to (7) except for the projection error. Standard approximation theory results (e.g., 20) imply that this error will decay with a rate that depends on the smoothness of $\frac{\partial \rho_i}{\partial r_j^\alpha}$, i.e., on the smoothness of g, and on the choice of basis ϕ_q.

For simplicity let us assume that g is analytic (though a Gaussian would be entire and yield even stronger results), and that ϕ_q is a basis of polynomials which is the most common choice in this field $(\phi_q(x\hat{\mathbf{x}}) = R_n(x)Y_l^m(\hat{\mathbf{x}}))$, then the error will be exponentially small. Taking into account also the rescaling of space via the smearing width σ_a, and the domain encoded in the cut-off radius, one can obtain that

$$\left\| \Pi\left[\frac{\partial \rho_i}{\partial r_j^\alpha}\right] - \frac{\partial \rho_i}{\partial r_j^\alpha} \right\| \lesssim \exp(-c\sigma_a \min(n_{\max}, l_{\max})), \tag{11}$$

for a constant c that depends on g and has unit of inverse distance. From this, we can conclude again that $\mathrm{cond}(\mathbf{J}(\boldsymbol{\xi})) \to 1$ at an exponential rate as the discretization parameter $\min(n_{\max}, l_{\max})$ increases. Recalling the condition number estimate in terms of smearing width (7) and a brief argument detailed in the *Extended data*[21] yields the combined estimate

$$\mathrm{cond}(\mathbf{J}) - 1 \lesssim e^{-\frac{r_0^2}{8\sigma_a^2}} + e^{-c\sigma_a \min(n_{\max}, l_{\max})}, \tag{12}$$

where $r_0 = \min_{j,j'} r_{jj'}$. These estimates are illustrated, confirmed, and explored more quantitatively in Figure 4.

While (7) is exact, experience from approximation theory is that our estimates are close to sharp. This strongly suggests that to obtain a discretization of ρ_i one must first choose σ_a such that $r_0/\sigma_a \gg 1$, and then choose the discretization parameters such that $\min(n_{\max}, l_{\max}) \gg 1/(c\sigma_a)$. In this regime, we expect that $\mathrm{cond}(\mathbf{J}(\boldsymbol{\xi}))$ will be close to one. However if these requirements are not satisfied then we would expect poor sensitivity encoded in the fact that singular values of $\mathbf{J}$ will be close to zero. More generally, (12) strongly hints that there is an optimal balance between the smearing and discretization parameters: given a smearing width σ_a there is a minimal resolution that is required but increasing it may not produce a descriptor with more uniform sensitivity. Vice versa, given a minimal distance r_0 and a resolution $\min(n_{\max}, l_{\max})$ of the basis, there is an optimal choice of smearing width that minimizes the condition number. See in particular the right-hand panel of Figure 4 for a quantitative illustration of this effect.

For the remainder of the theoretical discussion we shall assume that σ_a is chosen sufficiently small and $\min(n_{\max}, l_{\max})$ sufficiently large so that the resolution of ρ_i will not affect the results.

Loss of sensitivity after symmetrization

The case of the symmetrized density correlations is more complicated. Given that $|\overline{\rho_i^{\otimes \nu}}\rangle$ features are invariant with respect to rotations, the Jacobian has three singular values associated with a rigid rotation of the environment. When investigating the singular behavior of $\mathbf{J}$ it is then useful to work in a basis of atomic displacements that removes rigid rotations. The translational symmetry of $|\rho_i\rangle$, which is carried over to $|\overline{\rho_i^{\otimes \nu}}\rangle$, is taken care of by discarding from the Jacobian the row associated with the central atom i. If one considers $\nu = 1$ features that only depend on the distances of the neighbors to the central atom, it is clear that perturbations of the structure in which each neighbor is moved without changing r_{ji} will not result in a change of the feature values, and that $\mathbf{J}$ has at most N non-zero singular values, reflecting the fact that a two-body description of the environment is highly incomplete.

In the remainder of the paper we perform a numerical and geometric analysis of the behavior of $\mathbf{J}$, focusing in particular on $\nu = 2$ features, corresponding to distances and angles at the central atom. We will see that in certain symmetric configurations all symmetric features will have "degenerate" directions corresponding to singular values $s_k = 0$ – and that this may in fact be beneficial in terms of using these features to learn structure-property relations. However, more importantly, we will construct examples of configurations without such natural symmetries where two-correlations still have degenerate directions and thus any descriptor based on two-correlations inherits this degeneracy. In particular this means that the descriptor does not define a local coordinate system, which can have severe consequences for reconstructing atomic environments and for regression tasks.

We say that these structures or environments are *linearly degenerate*, or *linearly degenerate singularities*, to distinguish them from the discrete pairs of degenerate structures discussed in Ref. 5. By contrast, when a zero singular value is caused purely by a symmetry in the structure we will call such a structure a *symmetric singularity*.

Pedagogical example

As a pedagogical example consider the case of a single particle on the real axis, described by its coordinate x. Suppose we are interested in properties of this particle that are invariant under reflection, i.e. $O(1)$ symmetry. In this case we may decide to choose $\xi = x^2$ as a feature (or, 1-dimensional feature vector) describing the position. Clearly, knowledge of ξ allows us to reconstruct x up to a sign (the reflection), even near $x = 0$ where the feature ξ is singular, i.e. $\partial_x \xi = 0$. This is a *symmetric singularity* since it is induced by the symmetry group and generic to *all* symmetric functions. In that point (and only in that point) we cannot even invert the descriptor

locally, i.e., ξ does not supply a local coordinate system. However, if y is a reflection-symmetric property, then its Taylor expansion,

$$y(x) \sim c_0 + \frac{c_2}{2} x^2 + \frac{c_4}{4!} x^4 + \dots$$

has only even terms, and hence can be expressed as a smooth function of ξ, $y(x) = \tilde{y}(\xi)$. This suggests that ξ is well-suited for representing structure-property relationships.

Now suppose we make the less fortunate choice $\xi = \cos(x)$. By analogy with the power spectrum descriptor, ξ is not injective, i.e. we can find pairs of structures (in fact infinite tuples) that map to the same descriptor value, e.g., if $x_\pm := \pi \pm \epsilon$ then $\xi(x_+) = \xi(x_-)$. As this degenerate pair meets at $x_0 := \pi$ we obtain a *linearly degenerate singularity* expressed by the fact that $\partial_x \xi(x_0) = 0$. Of course we could have immediately seen this root of $\partial_x \xi$. However, we emphasize the intersection of degenerate pairs as it appears to be the *generic mechanism* underlying such linear degeneracies even in the much more complex case of symmetry adapted ν-correlations.

A general symmetric property y of course need not be symmetric about $x_0 = \pi$ and hence has a general Taylor expansion,

$$y(x) \sim a_0 + a_1 (x - x_0) + \frac{a_2}{2} (x - x_0)^2 + \frac{a_3}{6} (x - x_0)^3 + \dots$$

If $(a_1, a_3, \dots) \neq 0$, it will be impossible to represent $y(x) = \tilde{y}(\xi)$ in a (potentially small) neighborhood of the degenerate point x_0.

Symmetries and singularities of density-correlation features

We now proceed to demonstrate the general concepts we introduced in the previous section for an actual atomistic system, and for a (highly converged) discretization of different classes of representations. As an exemplar system we use C-centered environments of CH_4 configurations, which were used in Ref. 5 to construct concrete realizations of the degeneracies that are observed for low-body-order density correlation features. We consider different types of configurations to illustrate the various cases in which null singular values of the Jacobian can appear. We compare $\nu = 2$ features, $\nu = 3$ features, as well as OMFP and their projection on a smooth basis. We also plot the change in energy of the system as a function of molecular distortions, as an indication of the behavior one should expect for a typical molecular modeling target.

Directionally-resolved sensitivity analysis

For a feature vector containing a finite number n_{feat} of components, describing an environment with N neighbors, the Jacobian is a $3N \times n_{\text{feat}}$ matrix. Its singular value decomposition, $\mathbf{J} = \mathbf{U}\text{Diag}(\mathbf{s})\mathbf{V}^T$ identifies the principal modes of variation of the representation. $\mathbf{s}$ is a vector containing n_{J} principal values (usually $3N$, assuming the typical case in which $n_{\text{feat}} \gg 3N$) that indicate how much the various features change when the atoms are distorted by an infinitesimal amount according to the displacement patterns associated with each left singular vector contained in the columns of $\mathbf{U}$. The columns in the matrix $\mathbf{V}$ describe what feature distortion pattern is associated with each principal component; note that this construction implies that if $n_{\text{feat}} \gg 3N$, the right singular vectors $\mathbf{V}$ span a subspace of dimension smaller than n_{feat}, and so there are some changes in the feature vectors that cannot be realized by distortions of the structure. A crucial observation is that any orthogonal transformation of the features changes the right singular vectors $\mathbf{V}$, but not $\mathbf{s}$ or $\mathbf{U}$. Thus, in the complete basis set limit, one can characterize the sensitivity of representations in a way that is independent of the choice of basis, and so the analysis we carry out in this section is (largely) independent on the details of the discretization.

If a structure A_i is distorted according to a small Cartesian displacement $d\mathbf{u}$, the features change according to

$$d\langle q|A_i\rangle = \sum_k (d\mathbf{u} \cdot \mathbf{U}_k) s_k V_{qk}. \tag{13}$$

The magnitude of change of a feature indexed by q is given by the projection of the Cartesian displacement on the left singular vector $\mathbf{U}_k$, scaled by the associated singular value s_k, and then multiplied by the right singular vectors $\mathbf{V}_k$. Only the latter term depend on the discretization of $|A_i\rangle$, while in the complete basis set limit the left singular vectors $\mathbf{U}_k$ and the singular values s_k can be converged with respect to n_{max} and l_{max}. Here we use an optimized radial basis[22] with $n_{\text{max}} = 20$ components, built using the implementation in librascal[23] as the principal components

computed based on 1000 structures from the random CH_4 dataset, and starting from 100 Discrete Variable Representation (DVR) features, a large angular cutoff $l_{max} = 20$ and a very sharp density smearing $\sigma_a = 0.05$ Å to approach the complete basis limit of the density correlation features. In order to reduce the dependence of the sensitivity analysis on the discretization of the features we do not plot an arbitrary component of $\langle q|A_i\rangle$, but write the Cartesian displacement around the reference structure in terms of distortions Δ_k projected along the left singular vectors

$$\Delta\mathbf{u} = \sum_k \Delta_k \mathbf{U}_k(A_i), \tag{14}$$

and report on the finite changes of the features projected on the right singular vectors

$$\Delta\langle k|A_i(\Delta\mathbf{u})\rangle \equiv \sum_q V_{qk}(A_i)[\langle q|A_i(\Delta\mathbf{u})\rangle - \langle q|A_i\rangle]. \tag{15}$$

Equation (15) can also be used when the features that are being probed differ by those that define the left singular directions and $\Delta\mathbf{u}$ – for instance, in what follows we deform the structure along left principal directions for $|A;\overline{\rho_i^{\otimes 2}}\rangle$, but then inspect the change in $|A;\overline{\rho_i^{\otimes 3}}\rangle$ and OMFP. It suffices to define the directional derivative of the features along the Cartesian displacement associated with one of the left singular vectors, and normalize it

$$\tilde{V}_{qk} = \frac{\partial\langle q|A_i(\Delta\mathbf{u})\rangle}{\partial\Delta_k}; \quad \tilde{\mathbf{V}}_k \leftarrow \tilde{\mathbf{V}}_k / \left|\tilde{\mathbf{V}}_k\right|.$$

A subtle but important aspect is that symmetry-invariant features have some "trivial" zeros among the singular values, associated with rigid translations and rotations of the environment. These singularities should be resolved as a preliminary step, because otherwise non-trivial $s_k = 0$ directions would mix with the trivial ones, obfuscating the analysis. Singularities associated with a rigid translation of an environment are easily eliminated by only considering the Jacobian for the N atoms that are not the environment center. Rotations require more attention. A basis of $3N - 3$ displacements that are orthogonal to each other and to the displacements corresponding to infinitesimal rigid rotations of the environment around its center can be built based on purely geometric arguments, and the Jacobian should then be projected in this basis. The left singular vectors can be converted to build full Cartesian displacements by multiplying them by the transpose of the rotation-less basis matrix.

General structure

The left column of Figure 2 shows the changes in features associated with the largest and smallest singular values of the $\nu = 2$ Jacobian for a CH_4 structure corresponding to the optimal tetrahedral geometry, with H atoms distorted at random by 0.3Å. In the small-displacement regime, the changes in the representation $|A_i\rangle$ are linear, and the slope corresponds to the associated singular value. The energy has a non-zero gradient along the two directions, and the relative slope along the two directions reflects, at least qualitatively, the trend seen for all the choices of features except for the smooth OMFP.

Symmetric structure

Consider an environment $A_i = \{\mathbf{r}_{ji}\}_j$ that is left unchanged by application of a symmetry operation $\hat{S}$, so that $\hat{S}A_i = A_i$, with atoms indices mapped as $j \to \hat{S}j$, i.e. $\hat{S}\mathbf{r}_{ji} = \mathbf{r}_{\hat{S}ji}$. A general infinitesimal distortion of the atoms $d\mathbf{u}$ will generate an environment $A_i(d\mathbf{u}) = \mathbf{r}_{ji} + d\mathbf{u}_j$ that is not left invariant by the symmetry. If one can find a displacement pattern such that

$$\hat{S}d\mathbf{u}_j = -d\mathbf{u}_{\hat{S}j} \tag{17}$$

then $\hat{S}A_i(d\mathbf{u}) = A_i(-d\mathbf{u})$. Any symmetric set of features should be equal for $A_i(d\mathbf{u})$ and $A_i(-d\mathbf{u})$, meaning that the feature gradient along $d\mathbf{u}$ has to be zero: the Jacobian must have one additional singular value equal to zero – or more if several orthogonal displacements satisfying (17) can be found. A couple of simple examples of these symmetric distortions are shown in Figure 1b.

In the middle column of Figure 2 we show the sensitivity of different kinds of representations for two types of distortion of a minimum-energy CH_4 structure. The red curve corresponds to a symmetric breathing mode, in which all CH bonds are stretched in phase. This corresponds to the only non-zero singular value of the Jacobian for this configuration, and all descriptors change linearly along this direction. The blue curve, instead, corresponds to a stretch mode in which two CH bonds contract, and two dilate by the same amount. This deformation is symmetric in the sense of Equation (17), which leads to a symmetric behavior of the descriptors response. The case of OMFP merits further discussion: due to the eigenvalue sorting that guarantees atom index permutation invariance, the response has a cusp for $\Delta_2 = 0$. A similar cusp will be present for every symmetric deformation, which in

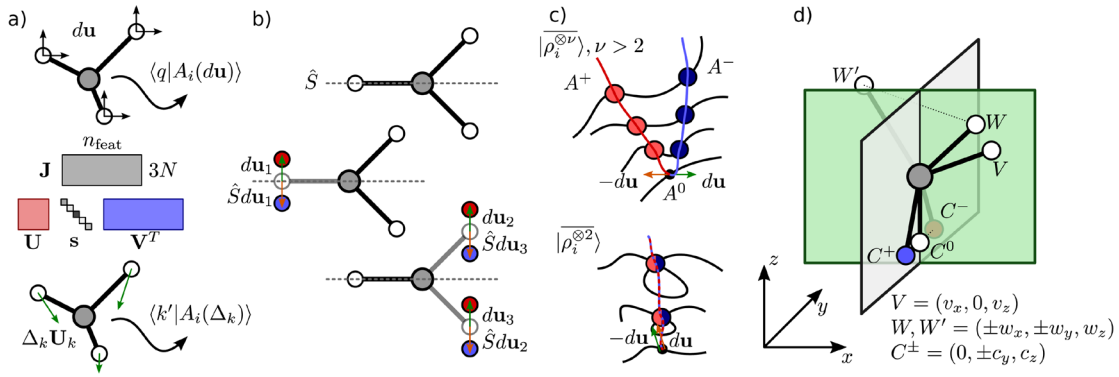


Figure 1. (**a**) Construction of a basis-set independent sensitivity analysis of the representation: linear perturbations of the feature vector are described by a Jacobian **J** that couples changes in the 3*N* Cartesian coordinates to changes in the features; right singular vectors project the basis-dependent features into a 3*N*-dimensional space that describes the intrinsic sensitivity of the representation. (**b**) Examples of a symmetric environment, and two deformations that lead to symmetric pairs of structures, and should therefore be associated with a zero singular value. (**c**) Schematic of the origin of spurious zero singular values associated with pairs of degenerate environments $A^\pm$: if the twin manifolds meet, their intersection produce a "degenerate singular point" A^0. A small deformation that moves along $A^\pm$ leads to finite separation (and finite feature derivatives) in a non-degenerate feature space, but not when using an incomplete representation. (**d**) An example of the construction of a degenerate singularity A^0 based on the same family of degenerate structures discussed in Ref. 5.

the case of the minimum-energy CH_4 structure implies all but one degree of freedom. Fortunately, it is simple to solve this problem: when using *smooth* OMFP, computed projecting the spectrum on a smooth basis following Equation (5), the cusp is removed, and the descriptors have zero gradient as required by the combination of symmetry and smoothness.

A quadratic behavior along the symmetric direction is consistent with the changes in energy, which follow a parabolic trend. For this minimum-energy configuration, the trend is parabolic along *both* directions. For the red curve, however, this quadratic behavior is a consequence of the CH bond length being the equilibrium one: if one changed it, or re-computed this same configuration at a different level of theory, there would be a non-zero slope for the energy around $\Delta_1 = 0$. The blue curve, instead, is a consequence of symmetry, and the trend would be quadratic for this stretch deformation even if the CH bonds were not the minimum-energy ones – provided that they were all equal.

Degenerate structure

For $\nu = 2$ features (three-body atom-centered symmetry functions (ACSF), SOAP powerspectrum . . .) one can also find configurations with zero singular values that are not associated with a symmetry, but rather with degenerate structures of the kinds discussed in Ref. 5. In short, one can find pairs of environments, A_i^+ and A_i^-, that have exactly the same $|\overline{\rho_i^{\otimes 2}}\rangle$ features, although they are not equivalent and in general have different physical properties. If the environments $A_i^\pm$ belong to manifolds $\mathbb{A}_i^\pm$ that cross then this gives rise to a "doubly-degenerate" environments A_i^0, for which there is at least one direction such that $A_i^0(d\mathbf{u}) \in \mathbb{A}_i^+$ and $A_i^0(-d\mathbf{u}) \in \mathbb{A}_i^-$. Thus, the value of the $|A^0(d\mathbf{u}); \overline{\rho_i^{\otimes 2}}\rangle$ features is identical to those of $|A^0(-d\mathbf{u}); \overline{\rho_i^{\otimes 2}}\rangle$. As a consequence, the Jacobian $\mathbf{J}(A_i^0)$ has a *spurious* zero among its singular values, leading to quadratic (or higher) variation of $\langle k|A_i(\Delta\mathbf{u})\rangle$ when displacing the atoms along the coordinate associated with it. This behavior is very clear in the right column of Figure 2. For a structure sitting at the intersection of the $\mathbb{A}_i^+$ and $\mathbb{A}_i^-$ manifolds, $\nu = 2$ density features have a spurious zero singular value, and the corresponding component of $|\overline{\rho_i^{\otimes 2}}\rangle$ change quadratically when displacing the atoms along the left singular vector. The pathological nature of this behavior is apparent in the fact that both OMFP and the $\nu = 3$ density representation vary with non-zero slope along the same geometric deformation, and so does a physical property such as the molecular energy. This shortcoming of low-body-order features is a direct consequence of the lack of injectivity discussed in Ref. 5.

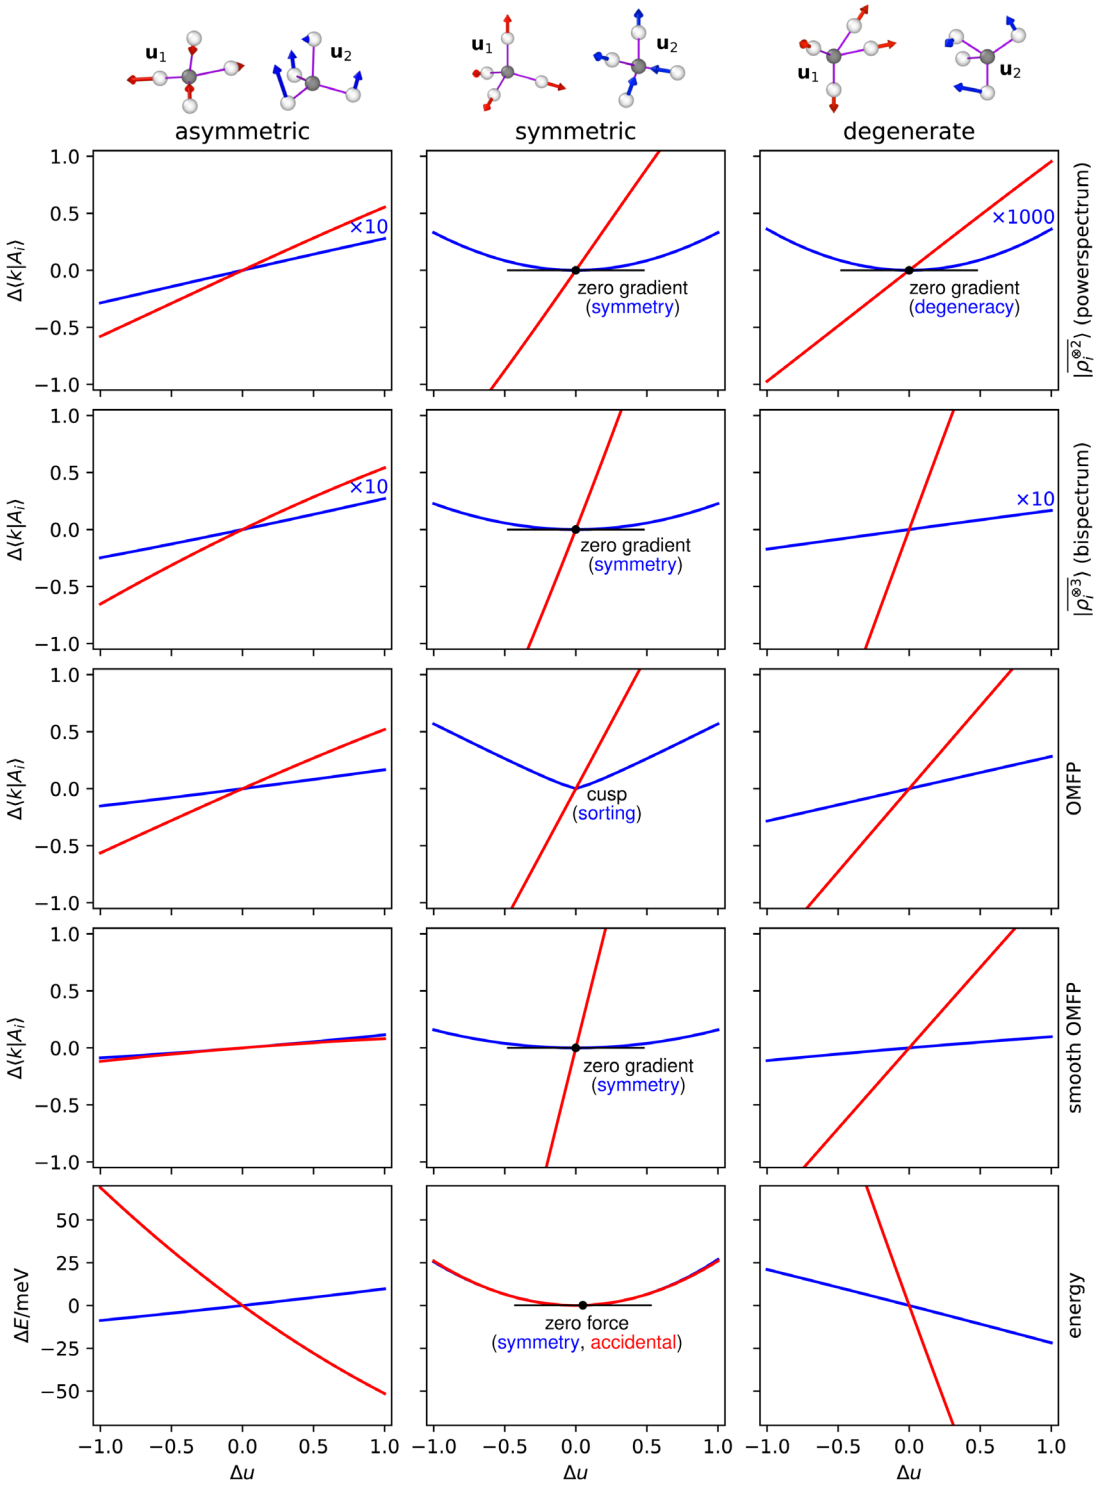


Figure 2. Changes in the components of features and properties when displacing the H atoms of a CH_4 molecule along selected principal directions. Results are shown for distortions along two principal directions (determined for $\nu = 2$ SOAP features), corresponding to the displacement patterns indicated at the top of the figure, where the color of arrows matches the color of curves. The columns correspond, left to right, the top and bottom singular values of a distorted, asymmetric CH_4 structure; the symmetric and one asymmetric breathing mode of the ideal methane structure; a "doubly-degenerate" structure at the intersection of the two branches of the manifolds discussed in Ref. 5. Rows show, top to bottom, projected variations of $\nu = 2$, $\nu = 3$, OMFP and smooth OMFP features, and the total energy of the molecule.

Similar to the order-zero degeneracies discussed in Ref. 5, information on the features centered on other atoms may be sufficient to break the tie, and lift the singularity, resulting in non-zero gradient of models of molecular properties. Any set of discrete degenerate structures can generate a spurious singularity whenever the degenerate manifolds cross. This is also true for the construction, discussed in Ref. 5, of a pair of structures that have the same $\nu = 3$ features and so also the bispectrum, and similar four-body features, can have linearly degenerate structures. We could not determine – but cannot rule out – the existence of additional manifolds, or isolated structures, which have a spurious zero in the eigenspectrum of $\mathbf{J}$ and are not the result of the intersection of those associated with discrete degeneracies.

Accuracy of regression models

To investigate the relationship between the mathematical properties of the structure-feature mapping and the behaviour of a regression model that uses them, we fit the energy of the example CH_4 configurations shown in Figure 2. We chose kernel ridge regression (KRR) based on a squared exponential kernel as a simple model with universal approximation properties. The elements of the kernel matrix between training configurations are defined as:

$$K_{ii'} = \mathrm{k}(A_i, A_i') \equiv e^{-\gamma \sum_q (\xi_q(A_i) - \xi_q(A_i'))^2}, \tag{18}$$

where $\xi_q(A_i)$ indicates one of the entries in the feature vector associated with an environment A_i. The representation may be either a discretization of $\left|\rho^{\otimes 2}\right\rangle$, $\left|\rho^{\otimes 3}\right\rangle$ or a set of OMFP.

The predictions for a new environment A_* can be written as

$$E(A_*) = \sum_i \mathrm{k}(A_*, A_i) b_i, \quad \mathbf{b} = (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{E} \tag{19}$$

where $\mathbf{K}$ is the kernel between train configurations, and $\mathbf{E}$ is the vector of the energies associated with the training structures. The model is entirely determined by the training set, the features, and two hyperparameters – the scale of the squared exponential kernel γ and the regularization σ^2. We optimize γ and σ^2 by grid search, minimizing the error of the predictions on the structure that are not used for training (further details are given in the *Extended data*[21]). Even though this procedure implies some information leaking from the test structures, doing this consistently for all models ensures a fair, and deterministic, comparison.

The overall errors of the different models are shown in Table 1, and the errors along the two one-dimensional cuts associated with the displacement patterns in Figure 2 are plotted in Figure 3. For the generic, asymmetric structure the three representations perform similarly. Errors are larger in the direction of $\mathbf{u}_1$ – which also

Table 1. Accuracy of regression models for the energy across a grid of 21 × 21 points along the two selected directions, as a function of the regular *n* × *n* sub-grid used for training. The same structures, distortions and representations are used as those shown in Figure 2. Errors are given in μeV (such small values are possible because we are effectively fitting a 2D function with a range of a few 10s of meV) and correspond to root mean square errors on the test points. OMFP = overlap matrix fingerprints.

Structure and grid	$\overline{\lvert\rho_i^{\otimes 2}\rangle}$	$\overline{\lvert\rho_i^{\otimes 3}\rangle}$	OMFP	s-OMFP
asymmetric, 3 × 3	89	48	230	16
asymmetric, 7 × 7	0.3	5.6	0.06	0.1
symmetric, 3 × 3	49	140	570	48
symmetric, 7 × 7	5.8	5.4	37	5.3
degenerate, 3 × 3	9600	150	6.5	1.4
degenerate, 7 × 7	5700	0.42	0.11	0.06

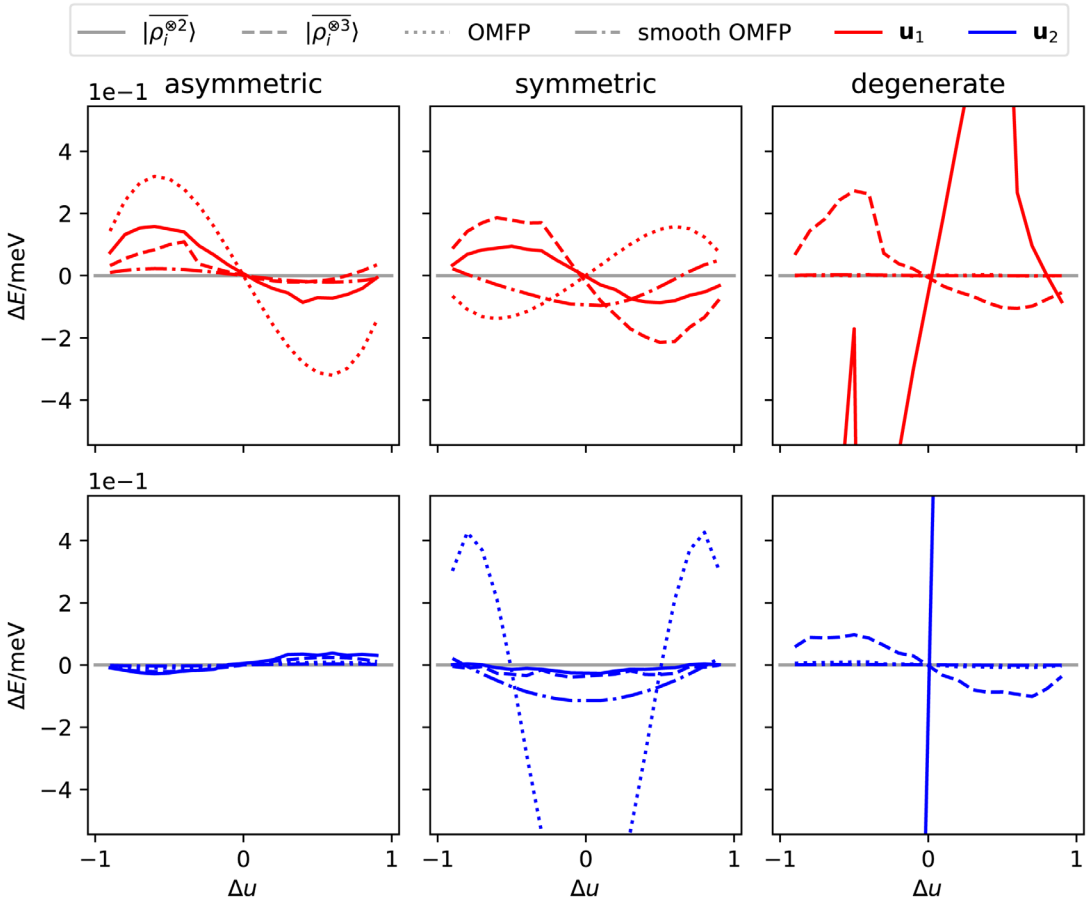


Figure 3. Error in fitting the energy of a CH_4 molecule when the atoms are displaced along the directions show in Figure 2. All models are trained using a 3 × 3 grid of points along $\mathbf{u}_1$ and $\mathbf{u}_2$, and the line dashing indicates the features used for the model. The two rows correspond to cuts along the directions corresponding to the two separate components.

corresponds to the highest singular value for $|\overline{\rho_i^{\otimes 2}}\rangle$ – than in the direction of $\mathbf{u}_2$ – which we selected as the lowest singular value. For the symmetric CH_4 structure, the global minimum of the methane molecule, the energy model displays a very interesting behaviour. The two deformations correspond to similar displacement-energy curves (Figure 2, bottom panel), but the model errors are very different. This is because the symmetric breathing mode $\mathbf{u}_1$ is an "accidental" energy minimum, and so the model predictions yield a small, but non-zero, force for $\Delta_1 = 0$. On the other hand, $\mathbf{u}_2 = 0$ is required to be an extremal point due to symmetry, and so all symmetric features predict, by construction, a symmetric curve with no force for $\Delta_2 = 0$ and achieve a much lower prediction error. The error of the OMFP-based model is considerably higher than the others. This is because of the non-smooth behavior in $\Delta_2 = 0$, which means that even if the curve is symmetric, the force in $\Delta_2 \to 0^{\pm}$ is not zero. Indeed, the smooth version of the OMFP has performance comparable to those of the power spectrum. Finally, displacements around the structure associated with a degenerate singularity result in a potential energy surface that cannot be fitted using $|\overline{\rho_i^{\otimes 2}}\rangle$ features. Even though being incapable of reproducing the non-zero force for $\Delta_2 = 0$ is a serious shortcoming of $|\overline{\rho_i^{\otimes 2}}\rangle$, the effect is dwarfed by the impact of the discrete degeneracies: the model displays an enormous error also for $\Delta_2 = \pm 1$, *even though those displacements corresponds to two of the training points*. The OMFP model performs better than $|\overline{\rho_i^{\otimes 3}}\rangle$ features, which are, however, perfectly capable of resolving the degeneracy. When increasing the training grid to 7 × 7 points, all features except for $|\overline{\rho_i^{\otimes 2}}\rangle$ converge, with only a small residual error.

Finite basis set

In the discussion this far we have been careful to use a highly-converged discretization of the density-correlation features, so that our numerical results are representative of the intrinsic nature of the representation, rather than of specific parameters, or implementation choices. Considerations on the "physical" and "degenerate" singularities, and the smoothness of the structure-feature mapping, applies equally well to SOAP, ACSF, FCHL, etc. However, implementation details – in particular the type and size of the basis used to discretize density correlations – do matter, as they affect the condition number of the Jacobian, possibly introducing numerical instabilities and near singularities. This is easy to see by considering an overly parsimonious discretization, in which the number of features is smaller than $3N - 3$: then, $\mathbf{J}$ has a rank that is insufficient to fully characterize structural distortions.

We can then consider a more realistic example, taking a database of CH_4 environments with the H atoms randomly distributed in a sphere of 3Å radius around the central carbon, and assessing the condition number of the Jacobian for the expansion coefficients of the C-centered density, as a function of the basis type and size. As shown in Figure 4, a large basis set size is needed to approach the theoretical condition number, which should be one for the non-symmetrized density expansion coefficients as discussed in *Sensitivity of ρ_i and its discretization*. Furthermore, the nature of the radial basis is very important. A basis optimized to maximize the information content for a given n_{max} approaches the limiting value for $n_{max} \approx 8$, even though a very large l_{max} has to be used, in addition. A Gaussian type orbitals (GTO) basis converges more slowly, and the DVR basis leads, for $n_{max} = 4$, to a mean condition number above 10^6. These numerical results closely match but quantify more precisely the theoretical predictions made in *Sensitivity of ρ_i and its discretization*, and further strengthen the importance of implementation details: even in a case in which one would expect a perfectly-conditioned Jacobian, near-singularities can arise when using a small, or sub-optimal, discretization of the representation. The right panel of Figure 4 demonstrates the role of the Gaussian width as well as its interplay with the discretization in determining the condition number of $\mathbf{J}$. As predicted theoretically by Equation 12, the resolution of the individual particles and thus the condition number initially improves with decreasing smearing width, but eventually the numerical discretization is no longer sufficient to resolve the narrow Gaussian, at which point the condition number increases again. The optimal value of σ_a is shifted to the left as the discretization is refined. This analysis supports – at least in terms of achieving a uniform sensitivity of the density expansion coefficients to atomic displacements – the choice of a Gaussian smearing of the density that is about half of the minimal interatomic separation, a practice that is often adopted by practitioners when constructing Gaussian approximation potentials[24]. It moreover highlights the importance of the discretization parameters when a small smearing width is employed.

Results for density-correlation features are qualitatively similar (which is unsurprising, given that they are computed from combinations of density expansion coefficients) but deserve further comments. As shown in Figure 5, the mean condition number of $\nu = 2$ features is higher than of $|\rho_i\rangle$, and we could not reduce it below 10-20. At the same time, the dependence of the condition number on n_{max} is sharper than for the density expansion coefficients: the presence of products of radial functions, which is implicit in the definition of $\langle n_1; n_2; l|\overline{\rho_i^{\otimes 2}}\rangle$, implies that fewer starting basis functions are needed to reach a satisfactory description of radial degrees of freedom. Similar observations can be made for the $\nu = 3$ correlations: the condition number decreases quickly with both n_{max} and l_{max} – the

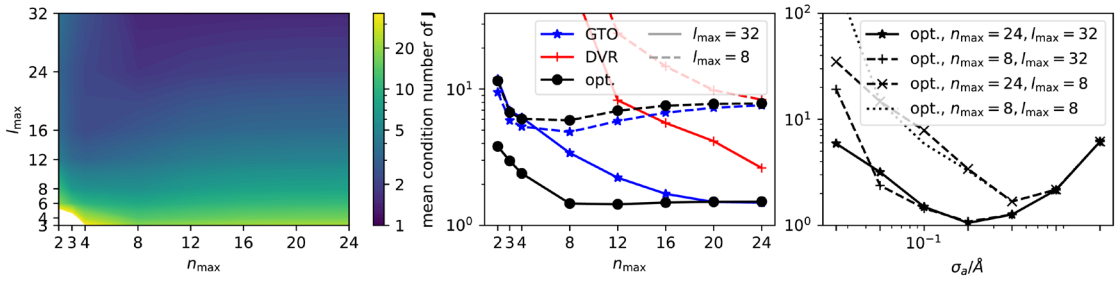


Figure 4. Mean condition number (CN) for the Jacobian of $|\rho_i\rangle$ features, computed from random CH_4 configurations[25], for C-centered environments. Left: mean CN as a function of n_{max} and l_{max}, for an optimal radial basis[22] and a density smearing σ_a = 0.1Å. Center: mean CN as a function of n_{max}, for two selected values of l_{max} and comparing three different choices of basis, for σ_a = 0.1Å. Right: mean CN as a function of σ_a, using an optimal radial basis for different values of (n_{max}, l_{max}).

correlations involve products of multiple radial and angular terms – but it is very difficult to reduce the mean condition number below ≈ 10. The question is whether the higher condition number of $\nu = 2, 3$ invariant features is due to the amplification of the singular values associated with the fact that density correlations are products of the expansion coefficients, or is associated with the proximity to the geometric singularities discussed in *Symmetries and singularities of density-correlation features*. To elucidate this point, we show in Figure 6 an analysis of the condition numbers computed for a random subset of 100 representative CH_4 structures extracted from the dataset in 25. For a given discretization, the condition number of the symmetrized features is consistently higher than that of the underlying $|\rho_i\rangle$ expansion coefficients. However, $\nu = 3$ features have consistently a *lower* condition number than the powerspectrum features, which indicates that the relationship between body order and the condition number is not as simple as if it reflected just the products that enter the definition of high-ν representations. In addition, $\nu = 2, 3$ features have a more pronounced high-condition number tail. The corresponding structures have a low condition number for $\mathbf{J}(|\rho_i\rangle)$ (Figure 6, center), suggesting that the large condition number is not a consequence of the discretization, but is linked to proximity to symmetric (or degenerate) singular points. To determine the nature of the singularities we use a similar strategy to that adopted in Ref. 5 to study discrete power spectrum degeneracies. We compare the condition number of $\mathbf{J}(|\overline{\rho_i^{\otimes 2}}\rangle)$ against that of $\mathbf{J}(|\overline{\rho_i^{\otimes 3}}\rangle)$. Structures that approach a degenerate singularity should have a large $\mathrm{cond}(\mathbf{J}(|\overline{\rho_i^{\otimes 2}}\rangle))$, and a much smaller $\mathrm{cond}(\mathbf{J}(|\overline{\rho_i^{\otimes 3}}\rangle))$. As shown in Figure 6 (right), all the structures we considered with a high condition number for

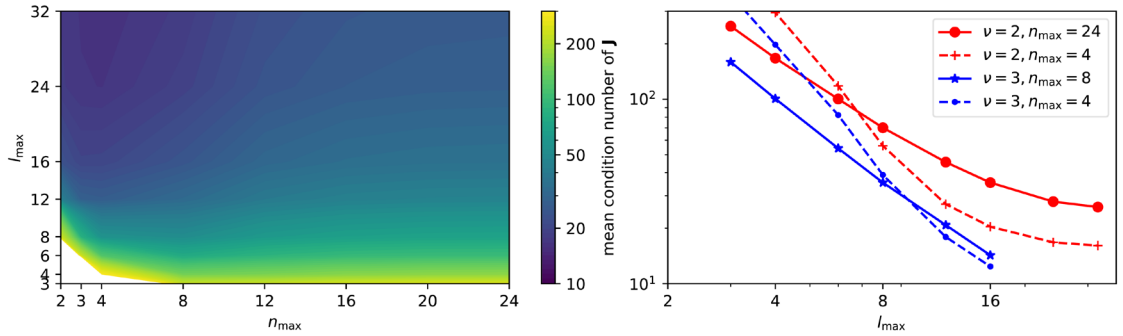


Figure 5. Mean condition number (CN) for the Jacobian of $|\overline{\rho_i^{\otimes \nu}}\rangle$ features, computed from random CH_4 configurations[25], for C-centered environments. Left: mean CN as a function of n_{max} and l_{max}, for smooth overlap of atomic positions (SOAP, $\nu = 2$) features, an optimal radial basis[22] and and a density smearing $\sigma_a = 0.1$Å. Right: mean CN as a function of l_{max}, for different values of n_{max}, comparing $\nu = 2$ and $\nu = 3$ features.

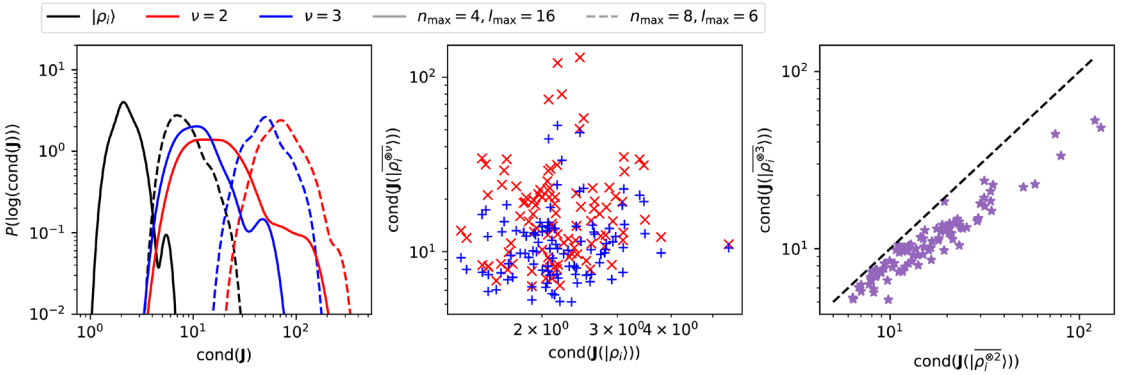


Figure 6. A comparison of the condition numbers (CNs) of the Jacobian of $|\rho_i\rangle$, $|\overline{\rho_i^{\otimes 2}}\rangle$, $|\overline{\rho_i^{\otimes 3}}\rangle$, computed from random CH_4 configurations[25], for C-centered environments. Left: histogram of the CN of different structures for typical discretization parameters ($n_{max} = 8$, $l_{max} = 6$; dashed lines), and parameters that lead to the best CN for $\nu = 3$ features ($n_{max} = 4$, $l_{max}16$; full lines). Center: parity plot of the condition number for 100 random structures, comparing the values for density expansion coefficients (x-axis), and those for $\nu = 2, 3$ density correlations (y-axis), using $n_{max} = 4$, $l_{max} = 16$ consistently. Right: parity plot of the condition number for the same 100 random structures, comparing the values for $\nu = 2$ (x-axis), and those for $\nu = 3$ (y-axis) density correlations.

powerspectrum features also have a comparatively high condition number for the bispectrum features, which indicates that they are (close to) symmetric singularities, and that none of the 100 random structures we considered are close to a degenerate singularity. Overall, this numerical analysis supports the analytical estimates in *Symmetries and singularities of density-correlation features*: the basis set and the density smearing must be converged (and balanced) to obtain neighbor density coefficients with a uniform sensitivity to deformations. For some configurations, symmetric representations exhibit (near) singular behavior, which is, at least for this dataset, consistent with physical constraints and not a manifestation of pathological behavior.

Conclusions

We have investigated, both analytically and numerically, the sensitivity to structural distortions of a family of symmetric representations of atomic environments that can be interpreted as a hierarchy of ν-point correlations of the neighbor density, and compared it with that of OMFP, an alternative set of features that incorporates *some* of the components at each order of correlation. Some of the results we show are intrinsic properties of the mathematical structure that underlies density-correlation representations, and independent of the discrete basis used to compute them as a finite-dimensional vector. In the limit of a complete basis, and for a sharp density, the mapping from coordinates to *translationally* invariant environment features has uniform sensitivity with respect to the displacement of any of the neighbors, as measured in terms of the Jacobian. Enforcing $O(3)$ group symmetries, however, introduces some symmetric singular points – configurations for which symmetry implies that the Jacobian has fewer than $3N - 3$ non-zero singular values. These symmetric singularities are physically meaningful, and beneficial to the regression of properties of the structures that are bound by symmetry to have critical points for those configurations. For low orders of correlation, we also find that the symmetric density correlation features present *degenerate singularities*, which are directly linked to the discrete degenerate structures that have been recently shown to be a manifestation of the incompleteness of the structure-feature mapping[5]. These singularities are unphysical, and detrimental to the construction of energy models – although potentially less so than the discrete degenerate manifolds they derive from.

It is important to stress that, even though both kinds of singularities can occur arranged along continuous manifolds, these are distinct from the manifolds of "quasi-constant features" observed in Ref. 10. In all cases we have observed so far the features are *not* constant along the manifold. Instead, structure along the manifold only share the lack of invertibility of the linearized structure-feature mapping, along one or more directions *orthogonal to the manifold*. Higher-order correlations, either from a systematic expansion or from OMFP, cure these problems. In the latter case the lack of smoothness in the feature mapping has a negative impact on the accuracy of models close to symmetric structures, which is easily cured by projecting the spectrum on a smooth basis.

We also consider how these general results change when using a *finite* discretization of the features. We find that the use of a small basis means that even the Jacobian of the mapping between atomic coordinates and neighbor density coefficients – that ought to have uniform sensitivity – can have a large condition number, which depends on the type and number of basis functions as well as on the smearing of the atom density, with optimal results obtained for an intermediate smearing of a fraction of the minimum interatomic separation. Symmetrized features computed with a limited basis, too sharp or too broad Gaussian smearing, inherit the anisotropic response to atomic deformations from the expansion coefficients they are built from. Even when using converged density coefficients, however, configurations with a high condition number can be found, corresponding to structures that are close to one of the symmetric singularities. We find instead that – at least in the benchmark dataset we consider – degenerate singularities are much rarer.

This *local* analysis of the nature of the structure-features mapping complements the geometric arguments of our previous work, and clarifies that the presence of structures or manifolds in which the Jacobian is low rank (or with a very large condition number) can be physical, but also the result of artifacts introduced by the incompleteness of the low-order descriptors or the excessive truncation of the basis used to discretize the density-correlation features. While we focus on applications to chemical and materials sciences, our results have obvious implications for any framework based on a description of point clouds, and the increasingly popular equivariant neural network frameworks. The challenges we discuss in the present work also appear in other contexts, for example in signal processing in relation to the "phase recovery" problem[26,27] (see also 28 for an extensive review) or in distance geometry where a closely related challenge is the "unassigned distance geometry problem"[29,30]. From a mathematical perspective, there are several open questions, which are closely related to those that arise for discrete degeneracies. The existence of classes of degenerate singularities different from those we

present, the frequency of problematic structures in realistic scenarios, whether and how many degeneracies remain as the body order of the descriptors is increased, and more generally the role of symmetries in changing the topology of the structure-feature mapping are all aspects for which further investigation is needed, and for which this work provides useful conceptual and numerical tools.

Data availability

Underlying data

The randomly-distorted CH_4 structures used in *Finite basis set* are a random representative subset of those distributed in the following data record.

Materials Cloud: Randomly-displaced methane configurations. https://doi.org/10.24435/materialscloud:qy-dp[25].

This project contains the following underlying data:

- methane.extxyz.gz - 7732488 random methane molecules along with their dft energies and forces in the extended xyz format

The data used to generate Figure 2, including molecular geometries, reference energies and projected feature displacements, are stored in the following data record.

Materials Cloud: Sensitivity benchmarks of structural representations for atomic-scale machine learning. https://doi.org/10.24435/materialscloud:7z-g6[31].

This project contains the following underlying data:

- *.chemiscope.json.gz. (2D manifolds corresponding to the plots in Figure 2)

Data are available under the terms of the Creative Commons Attribution-NonCommercial 4.0 International license (CC BY-NC 4.0).

Extended data

Figshare: Local invertibility and sensitivity of atomic structure-feature mappings, Supporting information. https://doi.org/10.6084/m9.figshare.16734562.v1[21].

This project contains the following extended data:

- SI.pdf (text file containing derivations and model details)

Data are available under the terms of the Creative Commons Attribution 4.0 International license (CC-BY 4.0).

Software availability

Powerspectrum and bispectrum features used in this work were generated using librascal, a library to compute representations for atomic-scale machine learning.

Source code available from: https://github.com/cosmo-epfl/librascal/tree/open-research-europe-2021

Archived source code at time of publication: https://doi.org/10.5281/zenodo.5547856[32]

License: GNU Lesser General Public License v3.0 (LGPL 3.0)

Acknowledgements

We would like to thank Stefan Goedecker and Behnam Parsaeifard for sharing the code we used to compute OMFP.

References

1. Ceriotti M, Clementi C, von Lilienfeld OA: **Introduction: Machine Learning at the Atomic Scale.** *Chem Rev.* 2021; **121**(16): 9719–9721.
PubMed Abstract | Publisher Full Text

2. Musil F, Grisafi A, Bartók AP, *et al.*: **Physics-Inspired Structural Representations for Molecules and Materials.** *Chem Rev.* 2021; **121**(16): 9759–9815.
PubMed Abstract | Publisher Full Text

3. Willatt MJ, Musil F, Ceriotti M: **Atom-density representations for machine learning.** *J Chem Phys.* 2019; **150**(15): 154110.
PubMed Abstract | Publisher Full Text

4. Drautz R: **Atomic cluster expansion for accurate and transferable interatomic potentials.** *Phys Rev B.* 2019; **99**(1): 014104.
Publisher Full Text

5. Pozdnyakov SN, Willatt MJ, Bartók AP, *et al.*: **Incompleteness of Atomic Structure Representations.** *Phys Rev Lett.* 2020; **125**(16): 166001.
PubMed Abstract | Publisher Full Text

6. Behler J, Parrinello M: **Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces.** *Phys Rev Lett.* 2007; **98**(14): 146401.
PubMed Abstract | Publisher Full Text

7. Bartók AP, Payne MC, Kondor R, *et al.*: **Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons.** *Phys Rev Lett.* 2010; **104**(13): 136403.
PubMed Abstract | Publisher Full Text

8. Parsaeifard B, De DS, Christensen AS, *et al.*: **An assessment of the structural resolution of various fingerprints commonly used in machine learning.** *Mach Learn Sci Technol.* 2020.
Publisher Full Text

9. Onat B, Ortner C, Kermode JR: **Sensitivity and dimensionality of atomic environment representations used for machine learning interatomic potentials.** *J Chem Phys.* 2020; **153**(14): 144106.
PubMed Abstract | Publisher Full Text

10. Parsaeifard B, Goedecker S: **Manifolds of quasi constant soap and acsf fingerprints**. *arxiv: 2102.06915.* 2021.
Reference Source

11. Willatt MJ, Musil F, Ceriotti M: **Feature optimization for atomistic machine learning yields a data-driven construction of the periodic table of the elements.** *Phys Chem Chem Phys.* 2018; **20**(47): 29661–29668.
PubMed Abstract | Publisher Full Text

12. Smith JS, Isayev O, Roitberg AE: **ANI-1: An extensible neural network potential with DFT accuracy at force field computational cost.** *Chem Sci.* 2017; **8**(4): 3192–3203.
PubMed Abstract | Publisher Full Text | Free Full Text

13. Zhang L, Han J, Wang H, *et al.*: **Deep Potential Molecular Dynamics: A Scalable Model with the Accuracy of Quantum Mechanics.** *Phys Rev Lett.* 2018; **120**(14): 143001.
PubMed Abstract | Publisher Full Text

14. Bartók AP, Kondor R, Csányi G: **On representing chemical environments.** *Phys Rev B.* 2013; **87**(18): 184115.
Publisher Full Text

15. Faber FA, Christensen AS, Huang B, *et al.*: **Alchemical and structural distribution based representation for universal quantum machine learning.** *J Chem Phys.* 2018; **148**(24): 241717.
PubMed Abstract | Publisher Full Text

16. Shapeev AV: **Moment Tensor Potentials: A Class of Systematically Improvable Interatomic Potentials.** *Multiscale Model Simul.* 2016; **14**(3): 1153–1173.
Publisher Full Text

17. Nigam J, Pozdnyakov S, Ceriotti M: **Recursive evaluation and iterative contraction of *N*-body equivariant features.** *J Chem Phys.* 2020; **153**(12): 121101.
PubMed Abstract | Publisher Full Text

18. Goscinski A, Fraux G, Imbalzano G, *et al.*: **The role of feature space in atomistic learning.** *Mach Learn Sci Technol.* 2021; **2**: 025028.
Publisher Full Text

19. Zhu L, Amsler M, Fuhrer T, *et al.*: **A fingerprint based metric for measuring similarities of crystalline structures.** *J Chem Phys.* 2016; **144**(3): 034203.
PubMed Abstract | Publisher Full Text

20. Powell MJD: **Approximation Theory and Methods**. Cambridge University Press, 1981.
Reference Source

21. Pozdnyakov S, Zhang L, Ortner C, *et al.*: **Local invertibility and sensitivity of atomic structure-feature mappings, Supporting information**. *Figshare*. 2021.
http://www.doi.org/10.6084/m9.figshare.16734562.v1

22. Goscinski A, Musil F, Pozdnyakov S, *et al.*: **Optimal radial basis for density-based atomic representations.** *J Chem Phys.* 2021; **155**(10): 104106.
PubMed Abstract | Publisher Full Text

23. Musil F, Veit M, Goscinski A, *et al.*: **Efficient implementation of atom-density representations.** *J Chem Phys.* 2021; **154**(11): 114109.
PubMed Abstract | Publisher Full Text

24. Deringer VL, Bartók AP, Bernstein N, *et al.*: **Gaussian Process Regression for Materials and Molecules.** *Chem Rev.* 2021; **121**(16): 10073–10141.
PubMed Abstract | Publisher Full Text | Free Full Text

25. Pozdnyakov S, Willatt M, Ceriotti M: **Dataset: Randomly-displaced methane configurations**. *Materials Cloud*. 2020.
http://www.doi.org/10.24435/materialscloud:qy-dp

26. Yellott JI, Iverson GJ: **Uniqueness properties of higher-order autocorrelation functions.** *J Opt Soc Am A. JOSAA.* 1992; **9**(3): 388–404.
Publisher Full Text

27. Kakarala R: **The Bispectrum as a Source of Phase-Sensitive Invariants for Fourier Descriptors: A Group-Theoretic Approach.** *J Math Imaging Vis.* 2012; **44**(3): 341–353.
Publisher Full Text

28. Uhrin M: **Through the eyes of a descriptor: Constructing complete, invertible, descriptions of atomic environments**. 2021.
Reference Source

29. Boutin M, Kemper G: **On reconstructing *n*-point configurations from the distribution of distances or areas.** *Adv Appl Math.* 2004; **32**(4): 709–735.
Publisher Full Text

30. Duxbury PM, Granlund L, Gujarathi SR, *et al.*: **The unassigned distance geometry problem.** *Discrete Appl Math.* 2016; **204**: 117–132.
Publisher Full Text

31. Pozdnyakov S, Ceriotti M: **Dataset: Sensitivity benchmarks of structural representations for atomic-scale machine learning**. *Materials Cloud*. 2021.
http://www.doi.org/10.24435/materialscloud:7z-g6

32. Fraux G: **cosmo-epfl/librascal: Version used in the "Local invertibility and sensitivity of atomic structure-feature mappings" paper**. *Zenodo*. 2021.
http://www.doi.org/10.5281/zenodo.5547856

Open Peer Review

Current Peer Review Status: ✔ ✔

Version 1

Reviewer Report 17 November 2021

https://doi.org/10.21956/openreseurope.15263.r27875

© **2021 Braams B.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

✔ **Bastiaan Braams**
The Dutch National Center for Mathematics and Computer Science, Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands

The research article under review concerns issues of uniqueness and conditioning of atomic structure feature mappings that are used in machine-learning of molecular energies and force fields. The paper elaborates on an earlier more accessible presentation by an overlapping group of authors [1] (ref. [5] in the present article) and I recommend to readers to start with [5] as an introduction to the present article. The feature maps are classified with respect to the degree of the involved correlations: one-point (\nu=1) features involving distances to the central atom only, two-point features involving distances and plane angles, three-point features also solid angles, and so on; and note that the points in this language are arbitrary points in the neighborhood of the central atom, they are not neighboring atoms. For any fixed degree \nu the correlation depends on \nu spatial coordinates and the level of discretization is another parameter; in the paper the complete basis set limit is emphasized at first and then numerical experiments are provided for specific discretizations using the potential energy surface of methane (CH4) as the primary example. Altogether the article offers an important contribution to the understanding of accuracy of machine-learned atomic force fields based on atomic neighborhood feature mappings.

In motivating the study the authors emphasize the diversity of approaches relying on local feature mappings, and it may be useful to point out that not all machine learning methods for large molecules and materials fall in that category. In particular I highlight the SchNet approach [2] that is a forerunner of a large class of methods for machine learned force fields under the umbrella name of graph convolutional neural networks (GCNN); see the review [3] and many references therein. In the GCNN approach the identity of atoms is retained and one works with functions of the local internuclear distances. As noted also in the present paper, the matrix of internuclear distances provides a good coordinate system as the Jacobian of this feature vector is singular only when it is mandated by symmetry. Replacing each internuclear distance rij by some f(rij) where f is monotone, smooth and compactly supported, going to zero at some modest distance, one obtains a sparse edge-weighted graph with the number of edges bounded linearly in the number of atoms. This representation seems to avoid the possible ill-conditioning associated with the atomic

feature mappings. (For present purposes we ignore issues of long-range electrostatics; they are a special concern in any case.) However, a comparison of methods relying on atomic feature mappings with methods from the GCNN class is outside the scope of the present work.

References

1. Pozdnyakov SN, Willatt MJ, Bartók AP, Ortner C, et al.: Incompleteness of Atomic Structure Representations.*Phys Rev Lett*. 2020; **125** (16): 166001 PubMed Abstract | Publisher Full Text
2. Schütt KT, Sauceda HE, Kindermans PJ, Tkatchenko A, et al.: SchNet - A deep learning architecture for molecules and materials.*J Chem Phys*. 2018; **148** (24): 241722 PubMed Abstract | Publisher Full Text
3. Fung V, Zhang J, Juarez E, Sumpter B: Benchmarking graph neural networks for materials chemistry. *npj Computational Materials*. 2021; **7** (1). Publisher Full Text

Is the work clearly and accurately presented and does it cite the current literature?
Yes

Is the study design appropriate and does the work have academic merit?
Yes

Are sufficient details of methods and analysis provided to allow replication by others?
Yes

If applicable, is the statistical analysis and its interpretation appropriate?
Yes

Are all the source data underlying the results available to ensure full reproducibility?
Yes

Are the conclusions drawn adequately supported by the results?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Plasma Physics, Chemical Physics, Atomic and Molecular Data, Computational Science

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Reviewer Report 16 November 2021

https://doi.org/10.21956/openreseurope.15263.r27874

© 2021 Maurer R. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Reinhard Maurer
[1] Chair for Theoretical Chemistry, Technical University of Munich, Munich, Germany
[2] Department of Chemistry, University of Warwick, Coventry, UK

The manuscript studies important mathematical properties of atomic symmetry functions that are commonly used in atomistic machine learning applications. A sensitivity analysis of common atomic structure-feature mappings is introduced to assess if such mappings retain their sensitivity to capture changes in atomic configuration in scenarios with high symmetry and when a finite basis representation is applied. The theory and analysis are described in great detail, which I have found to be highly educational.

The authors show that symmetry-induced singularities and lack of local invertibility can limit the ability to represent energy landscapes and lead to errors, which reveals interesting variations in the quality of different features in Figure 3. Particularly the failure of the power spectrum to describe discplacements of a degenerate singularity structure is fascinating.

The manuscript offers a comprehensive and compelling analysis of fundamental mathematical aspects, while also discussing how artefacts can materialise in realistic calculations where finite truncated basis sets are employed. It is also commendable that all underlying data and the employed software has been made available.

I recommend acceptance of the study. My only suggestion would be to revisit some of the figure captions to better describe individual subpanels (as for example done in Figure 1).

Is the work clearly and accurately presented and does it cite the current literature?
Yes

Is the study design appropriate and does the work have academic merit?
Yes

Are sufficient details of methods and analysis provided to allow replication by others?
Yes

If applicable, is the statistical analysis and its interpretation appropriate?
Yes

Are all the source data underlying the results available to ensure full reproducibility?
Yes

Are the conclusions drawn adequately supported by the results?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: electronic structure theory, materials modelling

I confirm that I have read this submission and believe that I have an appropriate level of

expertise to confirm that it is of an acceptable scientific standard.