

Comment on “Manifolds of quasi-constant SOAP and ACSF fingerprints and the resulting failure to machine learn four body interactions”

Sergey N. Pozdnyakov,¹ Michael J. Willatt,¹ Albert P. Bartók,²
Christoph Ortner,³ Gábor Csányi,⁴ and Michele Ceriotti^{1, *}

¹*Laboratory of Computational Science and Modelling, Institute of Materials,
Ecole Polytechnique Fédérale de Lausanne, Lausanne 1015, Switzerland*

²*Department of Physics and Warwick Centre for Predictive Modelling,
School of Engineering, University of Warwick, Coventry CV4 7AL, United Kingdom*

³*Mathematics Institute, University of Warwick, Coventry CV4 7AL, United Kingdom*

⁴*Engineering Laboratory, University of Cambridge,
Trumpington Street, Cambridge CB2 1PZ, United Kingdom*

(Dated: February 13, 2023)

The “quasi-constant” SOAP and ACSF fingerprint manifolds recently discovered by Parsaeifard and Goedecker[1] are closely related to the degenerate pairs of configurations, that are a known shortcoming of all low-body-order atom-density correlation representations of molecular structures. Configurations that are rigorously singular – that we demonstrate can only occur in finite, discrete sets, and not as a continuous manifolds – determine the complete failure of machine-learning models built on this class of descriptors. The “quasi-constant” manifolds, on the other hand, exhibit low but non-zero sensitivity to atomic displacements. As a consequence, for any such manifold, it is possible to optimize the model parameters and the training set to mitigate their impact on learning, even though this is often impractical and it is preferable to use descriptors that avoid both the exact singularities and the associated numerical instability.

The paper by Parsaeifard and Goedecker[1] presents an interesting observation on the numerical behaviour of smooth overlap of atomic positions (SOAP) and atom-centered symmetry functions (ACSF) representations. Manifolds along which the representation exhibits low sensitivity to atomic displacements (identified by small singular values of the Jacobian \mathbf{J}) are described. It is shown that the performance of machine-learning (ML) models based on these features is poor when the target has a four-body character: for the model used in Ref. 1, it appears that it is numerically impossible to learn along the manifold, even by explicitly adding it to the training set. Representations which are inherently high-order, such as overlap matrix fingerprints, do not exhibit these numerical issues. The significance of these empirical observations can be better appreciated by putting them in the context of the fundamental properties of atom-correlation based representations. Here we show that: (1) Low-sensitivity regions are a direct manifestation of the degenerate manifolds discussed in Ref. 2, which lead to configurations where the Jacobian is *exactly* singular[3]; (2) The extent to which these singular configurations induce a region of low sensitivity depends on the numerical details of the implementation; (3) Low sensitivity does not altogether prevent model fitting in the interpolative regime, although it affects significantly the accuracy and extrapolative power of the model.

Given that ML interatomic potentials for molecules and especially materials are usually built as a sum of atomic contributions [6], the geometry of atom-centered

environments are described by symmetrized representations. Many of these can be understood in terms of correlations of the neighbor density [7], and classified depending on the order of correlations ν that they incorporate: representations that have the same ν may differ in computational efficiency and numerical stability, but encode the same type of information. Recently, a fundamental limitation of representations that rely on 3-body correlations (such as SOAP[8] and ACSF[9, 10]) was revealed.[2] Pairs of environments exist which share *identical* 3-body representations even though the configurations differ. The existence of indistinguishable configurations affects the expressive power of models based on those features, which will be incapable of predicting distinct values for the corresponding atom-centered properties, *even if both structures are used during training*.

One can understand the origin of these degeneracies by noting that, in the complete basis set limit, $\nu = 2$ features contain information on the distances between the central atom and two of its neighbors, as well as the angle between them. Thus, they contain enough information to compute the set of entries of the Gram matrix for the environment. Much as it is the case for distance-based matrix representations of molecules[11], the Gram matrix allows to reconstruct a point cloud. Crucially, however, $\nu = 2$ features are invariant to atom index permutations, and do not allow to recover the position of the entries in the matrix, similar to what is seen when symmetrizing distance-matrix representations by sorting their entries[12]. Most of the time, shuffling the entries of the Gram matrix leads to point clouds that can only be embedded in a dimension higher than 3. Degenerate structures occur when two or more permutations result in distinct 3D point clouds. Incidentally, this discussion

* michele.ceriotti@epfl.ch




FIG. 1. (a) Low-sensitivity manifolds computed starting at different distances from a first-order degenerate CH_4 environment, A^0 . The discrete degenerate pair manifolds A^\pm are also shown. The manifold with $c = 1$ is graphically indistinguishable from the manifold used by Parsaeifard and Goedecker, although it depends slightly on the SOAP implementation and parameters. The others start from a point with a distance to A^0 scaled by the given factor c . (b) Condition number (CN) of \mathbf{J} for structures along the low-sensitivity manifolds. Full (dashed) lines correspond to a smearing $\sigma = 0.1$ (0.5) Å. Even though the CN along the manifold varies depending on σ , the geometries are very similar. SOAP features are computed using librascal[4], with $n_{\max} = l_{\max} = 8$, using optimal radial basis functions[5].

proves that a continuous manifold along which $\nu = 2$ features are *exactly* constant cannot exist: for a finite number of neighbors there can be only a finite number of permutations of the entries in the Gram matrix of an environment and hence a finite set of degenerate structures sharing the same $\nu = 2$ features. The largest known set of discrete degeneracies contains three structures.[2]

Getting back to the discrete degenerate pairs of Ref. 2, one can find that matching pairs are arranged along two continuous manifolds, A^\pm : at their intersection one finds structures A^0 for which \mathbf{J} has a “spurious” singular value – the features change as the square of the atomic displacement away from A^0 , even if A^0 is not symmetric.[3] Both “zero-order” degenerate pairs A^\pm and “first-order” singular configurations A^0 are intrinsic problems of any 3-body representation. The “quasi-constant” manifolds observed by Parsaeifard and Goedecker are a direct consequence of the structures being geometrically close to a first-order singularity, the spurious zero-sensitivity structure at the intersection of the manifolds of degenerate pairs. As shown in Fig. 1a, one can find a pair of discrete

degeneracies that meet close to the centre of the low-sensitivity manifold. Given that there is an entire manifold of first-order singularities,[13] one can expect to find low-sensitivity regions for several other molecular geometries. For each A^0 , many “low-sensitivity orbits” can be found by following the procedure outlined in Ref. 1 starting from points in the vicinity of the singularity. The level of sensitivity depends on the exact starting point, with the condition number of \mathbf{J} usually increasing as one approaches A^0 , where \mathbf{J} is exactly singular. This observation also explains the strong dependency of the sensitivity on the parameters of the representation, which is noted in Ref. 1. The use of large Gaussian smearing σ , or a small basis set that is not capable of discretizing accurately the neighbor-density correlations, means that the presence of a first-order singularity affects a larger region of configurations, with high condition numbers observed also further from A^0 (Fig. 1b).

Since the “quasi-constant” manifold is close, but does not intersect, the discrete degeneracies, it should be possible to distinguish structures along it, and train a model for the four-body potential, provided that one uses a high-resolution representation, and includes structures close to the manifold in the training set. To verify this, we perform similar numerical experiments to those discussed in Ref. 1, using an orbit starting at the same distance from A_0 as for the manifold discussed by Parsaeifard and Goedecker, following the lowest-sensitivity direction for SOAP features with $n_{\max} = 8$ optimal ra-

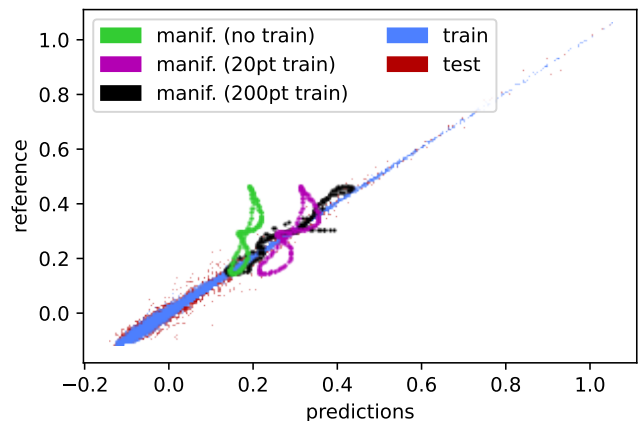


FIG. 2. Parity plot for neural-network models based on SOAP features, for a manifold built following the protocol from Parsaeifard and Goedecker, with SOAP features having $\sigma = 0.3$ Å, $l_{\max} = 8$ angular channels and $n_{\max} = 8$ optimal radial basis functions. Blue and red points correspond to training and validation sets. The green points are the predictions of for the quasi-degenerate manifold of a model trained without including any manifold configuration. The purple points are the predictions for a model that includes 10% (20 points) of the manifold configurations in the training set. The black points are the predictions for a model that includes the full manifold (200 points) during training. Changes for the validation error between models are small.

dial functions,[5] $l_{\max} = 8$, $\sigma = 0.3 \text{ \AA}$, which leads to a condition number comparable to the manifold in Fig. 1 of Ref. 1. We use a simple dense neural network with two inner layers with 80 neurons in each, similar to that discussed in Ref. 1. Hyperbolic tangents are used as activation, and group norms[14] inserted to accelerate convergence (Fig. 2). SOAP features *do* perform rather poorly in the extrapolative regime, which is unsurprising given that the “quasi-constant” manifold is in the close vicinity of the intersection of the manifolds of degenerate pairs, far from the training set, and that the target is specifically built as a pure 4-body term, which is intrinsically challenging for a 3-body representation. The SOAP-based NN we use here can approximate the energy along the “quasi-constant” manifolds, if a few reference configurations from it are included as part of the training set (Fig. 2, purple points). Learning points close to a singularity, however, is *harder*, and even including 200 points from the manifold is not sufficient to achieve perfect interpolation (black points). Manifolds that are even closer to the singularity, or that are built with higher Gaussian smearing, require even denser sampling and/or more flexible models. Increasing the resolution (e.g. taking $\sigma = 0.1 \text{ \AA}$ as in Fig. 5 of Ref. 1) reduces the ill-conditioning, and would facilitate learning. Even though the purely interpolative regime makes this an artificial exercise, which does not address the underlying problems of low-order atom-centered representations, it does prove that the quasi-constant nature of the manifold does not prevent learning *in principle*, contrary to the case of exact degeneracies.

Numerical stability and extrapolative power are certainly important in practice, although different from the outright failure of constructing a regression model. A more realistic analysis of regression performance would require working with configurations and regression targets that are relevant for a practical application, for which a careful tuning of the sensitivity of the representation often improves the accuracy of models[15–17]. Furthermore, even though 3-body features *do* exhibit very severe pathological behaviour in the form of degenerate pairs of configurations, this does not necessarily lead to catastrophic failures of the corresponding ML potentials. The saving grace, that appears to underpin the enormous success of ML models based on $\nu = 1, 2$ features, is the fact that most models are built using *multiple atoms as centers*, which resolve the degenerate pairs discussed in Ref. 2 – and thus largely address the underlying cause of the low sensitivity observed by Parsaeifard and Goedecker. [1] The benchmarks in Ref. 2 provide an estimate of the combined impact of *all types of degeneracies* for a challenging dataset of 7 million CH_4 configurations, whose atomization energies span a range of several tens of

eV. When using a multi-center model, the limiting accuracy that can be reached by a NN model based on SOAP features is of the order of 1 kcal/mol.

The problem of constructing a complete and concise set of descriptors for atomic configurations is not entirely solved yet. Low-body order features are known to present pathological behavior for some structures, and can be systematically improved by incorporating higher-order terms. Some schemes, such as MTP[18], ACE[19], NICE[20] give in principle access to a *linearly complete* representation. Others, such as OM fingerprints[21] do not generate a complete linear basis,[22] but can incorporate enough higher-order terms to resolve all known degeneracies (see Section 5 of Ref. 22 for a more thorough discussion).

It is important to distinguish fundamental shortcomings connected with the geometric nature of the density correlations from the numerical instabilities they cause for neighboring configurations. For the former, SOAP, ACSF, or any $\nu = 2$ construction such as DeepMD[23] or the 3-body FCHL[24] would fail in exactly the same way regardless of the hyperparameters and the implementation details, and lead to the impossibility of learning properties even for structures included in the train set. For the latter, tuning the hyperparameters of the representation, choosing a flexible regression model and enlarging the reference data set can mitigate the impact of the instabilities, even though impractical amounts of training data and unusual training protocols might be needed to achieve accurate predictions. Incorporating higher-body order terms, either with systematically-improvable density-correlation schemes[18–20] or with OM fingerprints[21] should be attempted any time one suspects that these effects are impacting the performance of a ML model in a practical application.

SUPPORTING MATERIALS

We provide a docker file that uses SOAP features computed from librascal[4, 25] to reconstruct the quasi-constant orbits, and train a simple NN based on PyTorch[26].

ACKNOWLEDGMENTS

MJW, SP and MC acknowledge funding by the Swiss National Science Foundation (Project No. 200021-182057) and by the Swiss Platform for Advanced Scientific Computing (PASC). CO acknowledges support by the Natural Sciences and Engineering Research Council of Canada (NSERC) [funding reference number GR019381].

[1] B. Parsaeifard and S. Goedecker, J. Chem. Phys. (2022), 10.1063/5.0070488.

[2] S. N. Pozdnyakov, M. J. Willatt, A. P. Bartók, C. Ortner, G. Csányi, and M. Ceriotti, Phys. Rev. Lett. **125**, 166001

- (2020).
- [3] S. N. Pozdnyakov, L. Zhang, C. Ortner, G. Csányi, and M. Ceriotti, *Open Res Europe* **1**, 126 (2021).
- [4] F. Musil, M. Veit, A. Goscinski, G. Fraux, M. J. Willatt, M. Stricker, and M. Ceriotti, *J. Chem. Phys.* **154**, 114109 (2021).
- [5] A. Goscinski, F. Musil, S. Pozdnyakov, J. Nigam, and M. Ceriotti, *J. Chem. Phys.* **155**, 104106 (2021).
- [6] J. Behler, *Chem. Rev.* **121**, 10037 (2021).
- [7] M. J. Willatt, F. Musil, and M. Ceriotti, *J. Chem. Phys.* **150**, 154110 (2019).
- [8] A. P. Bartók, R. Kondor, and G. Csányi, *Phys. Rev. B* **87**, 184115 (2013).
- [9] J. Behler and M. Parrinello, *Phys. Rev. Lett.* **98**, 146401 (2007).
- [10] J. Behler, *The Journal of Chemical Physics* **134**, 074106 (2011).
- [11] M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, *Phys. Rev. Lett.* **108**, 058301 (2012).
- [12] J. E. Moussa, *Phys. Rev. Lett.* **109**, 059801 (2012).
- [13] Note that the descriptors are *not* constant along this manifold, and at each point the singular direction is orthogonal to the manifold.
- [14] Y. Wu and K. He, arxiv:1803.08494 (2018).
- [15] B. Huang and O. A. von Lilienfeld, *The Journal of Chemical Physics* **145**, 161102 (2016).
- [16] M. J. Willatt, F. Musil, and M. Ceriotti, *Phys. Chem. Chem. Phys.* **20**, 29661 (2018).
- [17] M. A. Caro, *Phys. Rev. B* **100**, 024112 (2019).
- [18] A. V. Shapeev, *Multiscale Model. Simul.* **14**, 1153 (2016).
- [19] R. Drautz, *Phys. Rev. B* **99**, 014104 (2019).
- [20] J. Nigam, S. Pozdnyakov, and M. Ceriotti, *J. Chem. Phys.* **153**, 121101 (2020).
- [21] A. Sadeghi, S. A. Ghasemi, B. Schaefer, S. Mohr, M. A. Lill, and S. Goedecker, *J. Chem. Phys.* **139**, 184118 (2013).
- [22] F. Musil, A. Grisafi, A. P. Bartók, C. Ortner, G. Csányi, and M. Ceriotti, *Chem. Rev.* **121**, 9759 (2021).
- [23] L. Zhang, J. Han, H. Wang, R. Car, and W. E, *Phys. Rev. Lett.* **120**, 143001 (2018).
- [24] F. A. Faber, A. S. Christensen, B. Huang, and O. A. Von Lilienfeld, *J. Chem. Phys.* **148**, 241717 (2018).
- [25] F. Musil, M. Veit, T. Junge, M. Stricker, A. Goscinski, G. Fraux, and M. Ceriotti, “LIBRASCAL,” <https://github.com/cosmo-epfl/librascal> (2020).
- [26] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshain, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019) pp. 8024–8035.