

Training collective variables for enhanced sampling via neural networks based discriminant analysis

LUIGI BONATI⁽¹⁾(²)(³)

⁽¹⁾ *Department of Physics, ETH - Zurich, Switzerland*

⁽²⁾ *Institute of Computational Science, Università della Svizzera italiana
Lugano, Switzerland*

⁽³⁾ *Atomistic Simulations, Italian Institute of Technology - Genova, Italy*

received 18 January 2021

Summary. — A popular way to accelerate the sampling of rare events in molecular dynamics simulations is to introduce a potential that increases the fluctuations of selected collective variables. For this strategy to be successful, it is critical to choose appropriate variables. Here we review some recent developments in the data-driven design of collective variables, which combine Fisher’s discriminant analysis and neural networks. This approach allows to compress the fluctuations of metastable states into a low-dimensional representation. We illustrate through different applications the effectiveness of this method in accelerating the sampling, while also identifying the physical descriptors that undergo the most significant changes in the process.

Atomistic simulations, and in particular Molecular Dynamics (MD), play an important role in many fields of science. However, their application is often hindered by the presence of high free energy barriers between metastable states, resulting in transitions that occur on timescales too long to simulate [1]. Several important phenomena fall into the category of rare events, including phase transitions and chemical reactions, as well as biological processes such as protein folding and ligand binding.

Accelerating the sampling with collective variables. – Several enhanced sampling methods have been devised to extend the scope of MD simulations. A large family of these techniques relies on the identification of a small set of order parameters \mathbf{s} , called collective variables (CVs), which are functions of the atomic positions \mathbf{R} [1]. The potential energy surface $U(\mathbf{R})$ is modified by the addition of an external bias potential $V(\mathbf{s}(\mathbf{R}))$ which enhances the fluctuations of the selected variables by lowering the energetic barriers. This increases the sampling of the transition region, in turn accelerating the occurrence of rare events. A popular method is metadynamics [2], in which the bias potential $V(\mathbf{s})$ is expressed as a sum of repulsive Gaussians that are deposited at the points visited in the CVs space. In this way, the system is discouraged from re-exploring configurations that have already been visited and can escape local minima. A recent evolution of metadynamics, called on-the-fly probability enhanced sampling (OPES), builds the bias potential from an on-the-fly estimate of the probability distribution $P(\mathbf{s})$ [3].

For these methods to be successful, the collective variables must be related to the slow relaxation modes of the system or, in other terms, to the transition pathways connecting local minima. Since this information is usually not known beforehand, several approaches

have been proposed to design CVs using information limited to the metastable states. Machine Learning (ML) methods have risen to prominence in recent years, and there is no shortage of applications to atomistic simulations [4, 5] and enhanced sampling [6-8]. Both supervised [9] and unsupervised methods [10, 11] have been applied to the search of data-driven CVs. We discuss here in detail the approaches that exploit Fisher’s linear discriminant analysis [12, 13], and in particular recent advances in which this statistical method has been coupled with neural networks [14].

Linear Discriminant Analysis. – Linear Discriminant Analysis (LDA) [15] is a supervised method, first developed by Fisher, which is routinely used to perform classification and dimensionality reduction tasks. LDA seeks for a linear projection of the input features into a lower-dimensional space such that the classes are maximally separated. We illustrate here for simplicity the case of two classes, *i.e.* two metastable states A and B . First, we choose a set of descriptors \mathbf{d} that can be, for instance, distances, angles, dihedrals or coordination numbers. We run short MD simulations in basin A and B and compute the values of such descriptors. We use the average values $\boldsymbol{\mu}_A$, $\boldsymbol{\mu}_B$ and the variance matrices $\boldsymbol{\Sigma}_A$, $\boldsymbol{\Sigma}_B$ of the descriptors to characterize the basins and define the within-scatter matrix as $\mathbf{S}_w = \mathbf{S}_A + \mathbf{S}_B$ and the between-scatter matrix as $\mathbf{S}_b = (\boldsymbol{\mu}_A - \boldsymbol{\mu}_B)(\boldsymbol{\mu}_A - \boldsymbol{\mu}_B)^T$. Fisher’s criterion states that the linear combination $\bar{\mathbf{w}}$ that optimally separates classes is the one for which the ratio of the two scattering matrices in the projected space is maximal, *i.e.* $\bar{\mathbf{w}} = \arg \max_{\mathbf{w}} \frac{\mathbf{w}^T \mathbf{S}_b \mathbf{w}}{\mathbf{w}^T \mathbf{S}_w \mathbf{w}}$.

From a practical perspective, we can recast the above task as a generalized eigenvalue problem: $\mathbf{S}_b \bar{\mathbf{w}} = u \mathbf{S}_w \bar{\mathbf{w}}$, where the eigenvalue u measures the degree of separation of the two states along the corresponding eigenvector $\bar{\mathbf{w}}$. Once $\bar{\mathbf{w}}$ is found, the one-dimensional projection $s = \bar{\mathbf{w}}^T \mathbf{d}$ can be used as CV [12, 13].

Neural network based discriminant analysis. – While LDA provides a robust method for finding low-dimensional representations, it suffers from several limitations, the most important being that it can only provide a linear combination of the given descriptors. This precludes the application of the LDA CV to complex systems, where several descriptors need to be combined in a nonlinear way to capture the dynamics of the process. For this reason, we proposed to extend it with neural networks (NNs), similarly to what has been done in the field of computer science for classification [16]. This is achieved by using a NN to perform a nonlinear featurization of the input descriptors, and to compress them into a lower-dimensional space. LDA is then performed in this latent space, and the eigenvalue of the generalized eigenvalue equation is used as loss function to maximize the discriminative power of the NN. Further details are reported in [14]. In fig. 1 we exemplify the behavior of LDA and Deep-LDA in the simple case of two-dimensional Gaussians. We observe that, due to the nonlinear transformation performed by the NN, the Deep-LDA CV discriminates the two states better than its linear version.

Before moving on to applications, we summarize the main steps of this method from a practical viewpoint. Once a set of descriptors has been chosen, the first step is to simulate short MD trajectories in the metastable states and compute the descriptors’ values. The NN can be trained on these datasets using the Pytorch library, following the tutorial on Github [17]. Finally, the CV can be used along with enhanced sampling methods through the open-source library PLUMED [18], taking advantage of an interface we developed to import Pytorch models. Codes and input files are available also on the PLUMED-NEST [19] repository with `plumID:20.004`.

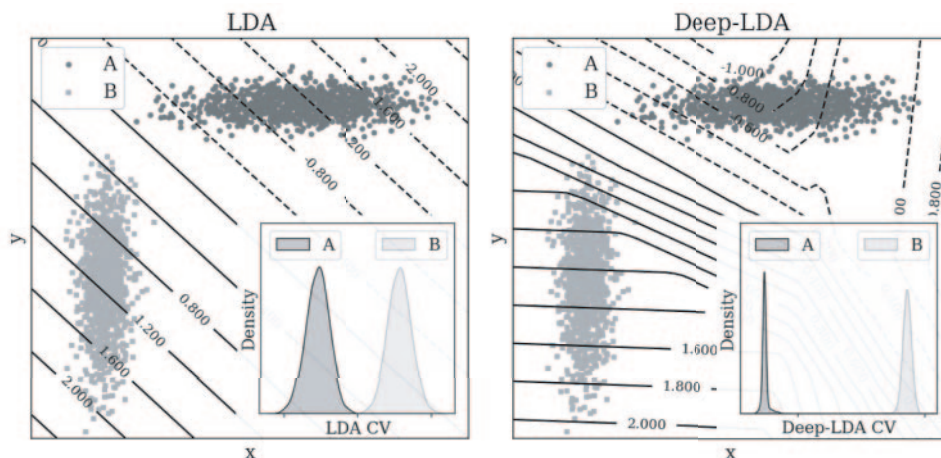


Fig. 1. – Comparison of LDA (left) and Deep-LDA (right) for a 2D Gaussian toy model. The points represent the two states, while the lines correspond to the CV isolines. In the case of LDA, the CV projection, orthogonal to the isolines, is a straight line, while for Deep-LDA it connects the fluctuations of the two states in a nonlinear fashion. In the insets we report the density profile along the CVs, to highlight the different discriminative power.

Enhancing the sampling with Deep-LDA. – As an application of the Deep-LDA CV, we examine the chemical reaction between vinyl alcohol and formaldehyde which has been studied in ref. [14]. Since the free energy barrier between reactants and products is on the order of 170 kJ/mol ($68 k_B T$ at room temperature), it is impractical to observe any transition in a MD simulation without resorting to enhanced sampling methods.

The set of descriptors consists of all 40 interatomic distances, whose values are measured in short simulations (5 ps) of the reactants and products. The Deep-LDA CV is trained on these datasets, using a NN with 3 hidden layers and $\{24, 12, 4\}$ nodes per layer. Further computational details can be found in ref. [14]. The CV is used to accelerate sampling in combination with the OPES method, which leads to the observation of several transitions between the two states. This allows the free energy profile of the reaction to be computed with great accuracy (fig. 2(a)). This sampling efficiency implies that the CV, even if based only on the configurations of the metastable states, is able to capture the concerted reaction mechanism.

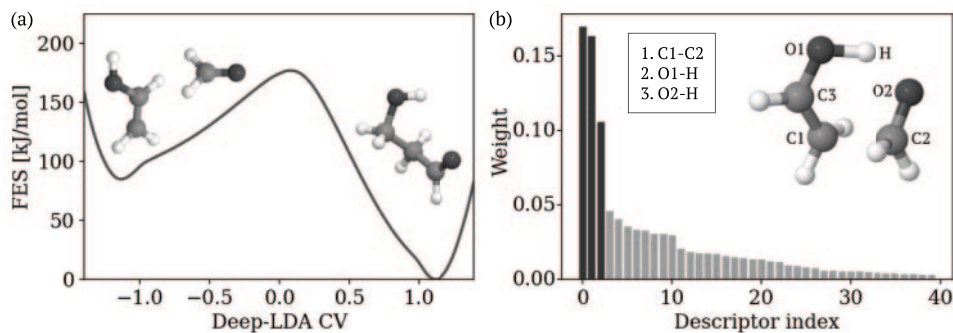


Fig. 2. – (a) Free energy as a function of Deep-LDA CV computed from the OPES simulation. The error, calculated with a block average, is about 1 kJ/mol . (b) Contributions of the descriptors to the CV, computed as in ref. [14], where the three dominant descriptors are highlighted.

Extracting physical insights from the CVs. – A notable feature of these data-driven CVs is their ability to highlight the most important descriptors, thereby deepening the physical understanding of the process. We illustrate this property with two examples.

The first is the chemical reaction described above: in fig. 2(b) the descriptors are ranked according to their contribution to the CV, from which we can recognize three dominant distances. This fact is in accordance with the chemical changes that the system undergoes during the reaction, in which a bond is formed between the two carbon atoms and, simultaneously, a proton transfer occurs between the oxygen atoms.

The second example is the study of the binding between a calixarene host and a set of ligands, reported in [20]. The Deep-LDA CV is used to include the effect of water in the study of the binding process, which is otherwise typically modeled using information limited to the ligand and the host. The analysis of the descriptors' weights allows to understand the role of water in binding and unbinding events, and the inclusion of this contribution in the CV set results in a very efficient sampling of the process [20].

Conclusions. – The combination of Fisher's linear discriminant and neural networks allows the design of data-driven collective variables using information limited to local minima. This approach has proven effective in accelerating the sampling of different types of rare events and identifying the most relevant atomic descriptors. It represents a promising tool for studying complex systems, and can be used as an initial guess for methods designed to extract CVs from enhanced sampling simulations.

* * *

The author thanks Dr. Michele Invernizzi for carefully reading the manuscript. The research was supported by the NCCR MARVEL, funded by the Swiss National Science Foundation, and European Union Grant No. ERC-2014-AdG-670227/VARMET.

REFERENCES

- [1] VALSSON O., TIWARY P. and PARRINELLO M., *Annu. Rev. Phys. Chem.*, **67** (2016) 159.
- [2] LAIO A. and PARRINELLO M., *Proc. Natl. Acad. Sci. U.S.A.*, **99** (2002) 20.
- [3] INVERNIZZI M. and PARRINELLO M., *J. Phys. Chem. Lett.*, **11** (2020) 7.
- [4] NOÉ F. *et al.*, *Annu. Rev. Phys. Chem.*, **71** (2020) 361.
- [5] BONATI L. and PARRINELLO M., *Phys. Rev. Lett.*, **121** (2018) 26.
- [6] SIDKY H., CHEN W. and FERGUSON A. L., *Mol. Phys.*, **118** (2020) 5.
- [7] WANG Y., RIBEIRO J. M. L. and TIWARY P., *Curr. Opin. Struct. Biol.*, **61** (2020) 139.
- [8] BONATI L., ZHANG Y. and PARRINELLO M., *Proc. Natl. Acad. Sci. U.S.A.*, **116** (2019) 36.
- [9] SULTAN M. M. and PANDE V. S., *J. Chem. Phys.*, **149** (2018) 9.
- [10] CHEN W. and FERGUSON A. L., *J. Comput. Chem.*, **39** (2018) 25.
- [11] SCHÖBERL M., ZABARAS N. and KOUTSOURELAKIS P.-S., *J. Chem. Phys.*, **150** (2019) 2.
- [12] MENDELS D., PICCINI G. and PARRINELLO M., *J. Phys. Chem. Lett.*, **9** (2018) 11.
- [13] PICCINI G., MENDELS D. and PARRINELLO M., *J. Chem. Theory Comput.*, **14** (2018) 10.
- [14] BONATI L., RIZZI V. and PARRINELLO M., *J. Phys. Chem. Lett.*, **11** (2020) 8.
- [15] WELLING M., Dep. Comput. Sci. Univ. Toronto. (2005).
- [16] DORFER M., KELZ R. and WIDMER, G., in *International Conference on Learning Representations, ICLR 2016*, arXiv:1511.04707[CS.LG].
- [17] <https://github.com/luigibonati/data-driven-CVs>.
- [18] TRIBELLO G. A. *et al.*, *Comput. Phys. Commun.*, **185** (2014) 2.
- [19] THE PLUMED CONSORTIUM, *Nat. Methods*, **16** (2019) 8.
- [20] RIZZI V., BONATI L., ANSARI N. and PARRINELLO M., *Nat. Commun.*, **12** (2021) 93.