

RESEARCH ARTICLE | SEPTEMBER 04 2026

Investigating causality between principal components in protein dynamics

Special Collection: [Festschrift in honor of Gerhard Hummer: Molecular Machines, Membranes, and Mechanisms: The Physics of Life at Work](#)

Debarshi Banerjee ; Ali Hassanali ; Alessandro Laio  



J. Chem. Phys. 165, 094122 (2026)

<https://doi.org/10.1063/5.0349449>



Articles You May Be Interested In

Towards a robust approach to infer causality from molecular dynamics simulations


J. Chem. Phys. (June 2025)

An iterative and automatic collective variable optimization scheme via unsupervised feature selection with CUR matrix decomposition

J. Chem. Phys. (May 2025)

Log-periodic oscillations as real-time signatures of hierarchical dynamics in proteins

J. Chem. Phys. (February 2024)



Boost Your Lab's Performance in Quantum Research



[Find out more](#)

We support your application with

- Performant instrumentation
- Intuitive software
- In-depth application support

Investigating causality between principal components in protein dynamics

Cite as: J. Chem. Phys. 165, 094122 (2026); doi: 10.1063/5.0349449

Submitted: 23 June 2026 • Accepted: 14 August 2026 •

Published Online: 4 September 2026



Debarshi Banerjee,^{1,2}  Ali Hassanali,^{2,a)}  and Alessandro Laio^{1,2,b)} 

AFFILIATIONS

¹ Scuola Internazionale Superiore di Studi Avanzati (SISSA), Via Bonomea 265, 34136 Trieste, Italy

² International Centre for Theoretical Physics (ICTP), Strada Costiera 11, 34151 Trieste, Italy

Note: This paper is part of the Special Topic on the Festschrift in Honor of Gerhard Hummer: Molecular Machines, Membranes, and Mechanisms: The Physics of Life at Work.

^{a)} Electronic mail: ahassana@ictp.it

^{b)} Author to whom correspondence should be addressed: laio@sissa.it

ABSTRACT

Principal component analysis (PCA) is widely used to characterize collective protein motions from molecular dynamics (MD) simulations. While PCA identifies the dominant modes of structural fluctuation, it does not reveal whether different principal components (PCs) causally influence each other. Here, we investigate this question using a recently introduced causal-discovery framework [V. D. Tatro *et al.*, Proc. Natl. Acad. Sci. U. S. A. **121**, e2317256121 (2024)], which allows us to infer putative causal asymmetries between high-dimensional time series. We apply this approach to long-timescale MD trajectories of two proteins. By analyzing relationships among PCs, we construct directed networks describing how PCs influence one another across time scales. These directional relationships, whose existence is a necessary condition for the presence of a causal link, are not captured by conventional covariance-based analyses and provide information that is complementary to PCA and time-lagged independent component analysis. Our results suggest that our causal inference approach can uncover previously hidden aspects of the dynamical organization of protein motions and offer a new perspective on this very popular class of collective variables.

© 2026 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0349449>

I. INTRODUCTION

Understanding how collective motions in proteins are organized across multiple time scales remains a central challenge in molecular biophysics.^{1–9} Molecular dynamics (MD) simulations now routinely generate trajectories of sufficient length and resolution to characterize fluctuations ranging from fast local motions to rare conformational transitions,^{10–18} enabling the investigation of both thermodynamic and kinetic properties. In particular, equilibrium trajectories of small proteins extending to biologically relevant time scales are now available,^{12,19,20} providing unprecedented opportunities to study the relationship between structure, dynamics, and function.

One of the most widely used approaches for analyzing protein dynamics in the folded state is Principal Component Analysis (PCA).^{21–32} In its standard formulation, PCA is performed by computing the covariance matrix of the positions of a selected subset

of protein atoms (for example, backbone α carbon atoms) after removing overall translational and rotational motions. The principal components (PCs) are the eigenvectors of this covariance matrix, while the corresponding eigenvalues quantify the variance along each component.^{26,28,30,32–35} When only a small number of eigenvalues account for a large fraction of the total variance, protein motions can be represented in a reduced-dimensional basis spanned by the leading PCs. In this basis, the covariance matrix is diagonal by construction, meaning that fluctuations along different PCs are uncorrelated at zero time lag. Owing to its ability to provide a compact description of collective structural fluctuations and to identify large-amplitude motions, PCA has become a standard tool for characterizing conformational landscapes and functionally relevant motions in proteins.^{29,30,32,34,36–46}

When the focus shifts from structural characterization to dynamical behavior, Time-lagged Independent Component Analysis (TICA) provides a natural extension of PCA.^{47–50} While principal

components are uncorrelated at zero lag time, correlations generally arise between them when the system is observed at a finite time delay, τ . $C(\tau)$ denotes the time-lagged covariance matrix between the centered and aligned atomic coordinates at times 0 and τ . In the PCA basis, the corresponding lagged covariance matrix is given by $\Gamma(\tau) = C^{-1}(0)C(\tau)$.^{47,49,51,52} This matrix would remain diagonal for all values of τ only in the idealized case of purely harmonic dynamics. For realistic protein motions, $\Gamma(\tau)$ is generally non-diagonal, reflecting dynamical couplings between different PCs. The eigenvectors of $\Gamma(\tau)$, commonly referred to as time-lagged independent components (ICs), identify collective coordinates associated with the slowest dynamical processes. As a result, they are widely employed to characterize long-timescale conformational changes, construct kinetic models, and define collective variables for enhanced-sampling simulations.^{49–51,53–56}

These developments raise a natural question: can collective structural fluctuations identified by PCA exert asymmetric dynamical influences on one another? More specifically, can principal components extracted from protein MD trajectories *cause* each other?^{57–59} In the language of causal inference, stating that, for example, PC_2 is a putative cause of PC_4 means that knowledge of the present state of PC_2 improves the prediction of the future evolution of PC_4 beyond what can be inferred from the history of PC_4 alone.⁶⁰ A natural first attempt to address this question would be to examine the time-lagged covariance matrix between PCs, $\Gamma(\tau)$. If, for instance, the off-diagonal element $\Gamma_{24}(\tau)$ is large, one may conclude that PC_2 and PC_4 are dynamically coupled. However, dynamical coupling does not necessarily imply causal influence. In particular, causal relationships are generally directional: it may happen that PC_2 influences the future evolution of PC_4 , while the converse is not true. Such asymmetries, if present, cannot be inferred from $\Gamma(\tau)$ alone. For reversible equilibrium dynamics, the transfer operator is self-adjoint with respect to the equilibrium measure. Consequently, TICA identifies slow decorrelating coordinates but cannot assign a directional driver–response interpretation to couplings between pairs of PCs.

Addressing this question requires tools that go beyond covariance or cross-correlation based analyses. Several approaches have been proposed to infer causality from molecular simulations. Early studies employed multivariate autoregressive models and analyses inspired by Granger causality to identify directed relationships between coordinates or residue fluctuations.^{61–64} An information-theoretic approach, dubbed transfer entropy,^{65,66} has also been used to probe directional information transfer and allosteric communication in proteins.^{67–69} More broadly, causal inference for time series has developed into rich theoretical frameworks aimed at predicting information asymmetries, conditional dependencies, and structural causal graphs.^{57,59,60,70} In practice, however, applying these approaches to atomistic MD remains challenging because the statistical errors of causal estimators are typically large when the causal links are weak.

To address these difficulties, we recently introduced the Imbalance Gain (IG), a statistical test designed to identify putative causal couplings in time-dependent data.^{71–73} The IG can be viewed as a proxy for the transfer entropy^{65,66} and can also be efficiently estimated for noisy and high-dimensional data.⁷² It is also related to Granger causality and related methods, but, unlike those approaches, it does not require fitting an explicit dynamical model.

Recent studies have shown that putative causal couplings identified through the IG can robustly reveal genuine causal relationships even in equilibrium molecular dynamics, despite the time-reversible nature of the microscopic equations of motion and the stationarity of the underlying distribution.⁷³ The same framework has also uncovered non-trivial causal relationships involving translational and rotational degrees of freedom in liquid water.⁷⁴ It is important to note that the IG and other similar approaches infer the presence of *putative* causal links from observational data. The presence of an asymmetric dynamical coupling between two variables is a *necessary* condition for the presence of a causal link between the two variables. However, this condition is not sufficient, since the two variables could be simultaneously caused by a third unobserved variable. Revealing the presence of confounders and building a genuine causal graph^{57,59,70} in molecular modeling remains an open challenge.

In the present work, we use the IG framework to investigate whether the principal components describing collective protein motions exhibit putative causal links. We analyze two long-timescale trajectories generated by Shaw Research: ubiquitin (PDB ID: 1UBQ)²⁰ and NTL9 (PDB ID: 2HBA).¹² NTL9 is a fast-folding protein whose trajectory samples large conformational changes and multiple folding and unfolding events,¹² whereas ubiquitin represents a prototypical native-state system whose equilibrium fluctuations span timescales from picoseconds to milliseconds and involve subtle long-range collective motions.²⁰ By comparing these two systems, we assess whether putative causal links among PCs emerge, and whether such links are controlled by variance, relaxation timescale, or spatial collectivity. We find that both proteins exhibit a causal structure but with distinct physical character: in ubiquitin, the directionality of the putative causal link originates from a highly collective, protein-wide mode and affects only the next few leading components, whereas in NTL9, a localized slow mode centered on the β -turn directionally couples to several faster modes, including more delocalized principal components.

II. METHODS

Within the Imbalance Gain (IG) framework,⁷² a putative causal link between two (possibly multidimensional) collective variables X and Y is estimated in terms of predictive asymmetry between the two time series $X(t)$ and $Y(t)$. The IG is based on the Information Imbalance⁷¹ (II), a statistical framework that provides a nonparametric way to compare the information content of different representations of the same data by exploiting the relative ordering of distances in the two spaces, X and Y . In particular, the II quantifies how well neighborhoods defined in X are preserved in Y . Denoting by r_{ij}^X the distance rank of configuration j among the neighbors of i in space X , the II of X with respect to Y is defined as

$$\Delta(X \rightarrow Y) = \frac{2}{N} \langle r_{ij}^Y | r_{ij}^X \leq k \rangle = \frac{2}{N^2 k} \sum_{i,j: r_{ij}^X \leq k} r_{ij}^Y, \quad (1)$$

where N is the total number of configurations and k is a small number (here set to 50, which is about 2% of N) used to restrict the average to the closest neighbors in space X . $\Delta(X \rightarrow Y)$ yields a value between 0 and 1. Small values of $\Delta(X \rightarrow Y)$ indicate that neighborhoods in X are well preserved in Y , implying that X contains most of

the information needed to reconstruct Y . Values closer to 1 instead indicate that X is minimally informative with respect to Y . The II framework has been used to infer non-trivial relationships between collective variables in a wide variety of applications in molecular systems.^{75–78}

The presence of a putative causal link is then inferred by estimating the II between collective variables at different times. Qualitatively, one estimates whether the inclusion of a candidate driver variable improves the prediction of a target variable at a future time.^{72,73} Since the information imbalance is computed over a set of points, we suppose to have access to multiple replicas of the trajectory. If the available data consist of a single trajectory, an ensemble of replicas is constructed by dividing the trajectory into non-overlapping chunks. For a target variable $Y(0)$ and a candidate driver $X(0)$, one compares the II between the present state of the target and its future, $\Delta(Y(0) \rightarrow Y(\tau))$, with the II obtained when also considering the putative driver in the present, $\Delta([\alpha X(0), Y(0)] \rightarrow Y(\tau))$, where $\alpha > 0$ is a variational parameter. In these equations, $Y(0)$, $X(0)$, and $[\alpha X(0), Y(0)]$ are shorthand notations for the Euclidean distance between pairs of replicas. For example, $[\alpha X(0), Y(0)] = (\alpha^2 \|X_i(0) - X_j(0)\|^2 + \|Y_i(0) - Y_j(0)\|^2)^{\frac{1}{2}}$, where the subscripts i and j label different replicas.

The Imbalance Gain (IG) is then defined as

$$IG_{X \rightarrow Y}(\tau) = \frac{\Delta(Y(0) \rightarrow Y(\tau)) - \min_{\alpha} \Delta([\alpha X(0), Y(0)] \rightarrow Y(\tau))}{\Delta(Y(0) \rightarrow Y(\tau))}. \quad (2)$$

The optimal value of α that maximizes the IG for a given time lag τ is determined by searching over a grid of 500 evenly spaced values for $\alpha \in [0, 10]$. If $IG = 0$, adding the information about $X(0)$ does not improve the prediction of $Y(\tau)$, since the best possible predictor is $Y(0)$ alone. A positive value of the $IG_{X \rightarrow Y}(\tau)$ indicates instead that the inclusion of $X(0)$ improves the prediction of the future state of Y , providing evidence for a putative causal link at lag τ . As already discussed, the observation of $IG > 0$ between two variables does not

imply that a causal link between the two exists, since both variables could be caused by a third unobserved variable. However, if one finds that $IG = 0$ (or small, in practical applications), one can conclude that the causal link does not exist.⁷² Since the IG is estimated using distances between configurations, it naturally accommodates high-dimensional descriptors and avoids the computational problems that are encountered when one attempts to infer the causal links from transfer entropy, which requires estimating conditional probability distributions in high-dimensional spaces.^{65,66}

We also quantify a putative causal link between two variables from the integral of the IG over time (IG Total),

$$IGT_{X \rightarrow Y} = \int_0^{\infty} IG_{X \rightarrow Y}(\tau') d\tau'. \quad (3)$$

In the present work, we apply this framework to PCs obtained from protein dynamics. Treating individual PCs as time-dependent variables, we evaluate the lag-dependent IG between pairs of components to assess whether specific modes systematically improve the prediction of others. This allows us to infer the presence of putative causal links in principal component space. To simplify the visualization and the discussion, we restrict our analysis to the top five PCs.

III. RESULTS

A. Ubiquitin

In the first application, we analyze the native-state dynamics of ubiquitin using the 1 ms molecular dynamics trajectory reported in Ref. 20. As discussed in that work, the trajectory exhibits a mild conformational transition at ~ 0.65 ms. Since our goal is to characterize causal relationships between principal components within a single native-state basin, we restricted the analysis to the first 0.65 ms of the simulation. This choice avoids mixing fluctuations associated with distinct conformational substates and ensures that the resulting PCs describe a well-defined region of the native-state ensemble. The eigenvalue spectrum and the explained variance of the PCs are shown in the [supplementary material](#), Fig. S1.




FIG. 1. Ubiquitin: (a) the autocorrelation functions of the top 5 PCs, $PC_{i=1,5}$, which are colored red, blue, green, orange, and purple, respectively, with shaded areas denoting error bars. (b) The pairwise IG curves over time between $PC_1 \leftrightarrow PC_{2,3,4,5}$. The IG values are on the Y-axis, and time lag (τ) in nanoseconds is on the X-axis. Pink lines denote the direction $PC_1 \rightarrow PC_{2,3,4,5}$, whereas green lines denote the reverse direction. Shaded areas denote error bars computed over 20 independent realizations.

We first characterize the intrinsic relaxation timescales of the leading modes. In Fig. 1(a), we report the autocorrelation functions of the first five principal components. All five modes decorrelate within ~ 500 ns, which is substantially shorter than the analyzed trajectory segment. We determine the relaxation times from the integral of the autocorrelation function. This indicates that the selected portion of the simulation is sufficiently long to sample the dominant fluctuations associated with these PCs. Among the modes considered, PC₁ is the slowest, while PC₂ and PC₃ decorrelate on comparable, shorter timescales.

Previous studies on equilibrium MD simulations have used apparent asymmetries in time-lagged two-body cross-correlation functions, $\mathbb{E}[X_0 Y_\tau]$, to infer causal relationships.^{67,79} However, for equilibrium dynamics satisfying stationarity and time reversibility (as is the case for MD simulations), these correlation functions are invariant under exchange of the two variables. In particular, $\mathbb{E}[X_0 Y_\tau] = \mathbb{E}[X_{-\tau} Y_0] = \mathbb{E}[X_\tau Y_0]$, where the first equality follows from stationarity and the second from time-reversibility. Therefore, any observed asymmetry in such correlation functions cannot be interpreted as evidence of directional causality. Instead, it must arise from statistical errors, numerical violations of time reversibility introduced by the integration scheme, or perturbations associated with thermostats and barostats. The IG (or transfer entropy) approach is instead asymmetric and directional and can be viewed conceptually as a many-body cross-correlation function where the information of the caused variable at the present time breaks the symmetry otherwise seen in two-body cross-correlation functions.^{72,73}

Therefore, we next examine putative causal links between the leading modes using the imbalance gain, instead of using two-body

cross-correlation functions between the PCs. In Fig. 1(b), we show the pairwise IG curves between PC₁ and PC_{2,3,4,5}. The strongest asymmetry is observed for the pairs PC₁ $\leftrightarrow$ PC₂ and PC₁ $\leftrightarrow$ PC₃, where the dominant direction of information transfer is PC₁ $\rightarrow$ PC_{2,3}. This indicates that knowledge of PC₁ improves the prediction of the future evolution of PC₂ and PC₃ more than the reverse. The asymmetry becomes weaker for PC₄, where the driving effect from PC₁ remains detectable but is considerably more modest. For PC₅, the IG signal is substantially reduced and becomes only weakly discernible. Consistently, when lower-variance PCs are considered, the directional signal progressively decreases and eventually approaches zero, as shown in the additional pairwise IG curves reported in the [supplementary material](#) (see Fig. S7). Overall, these results suggest that the dynamical influence of PC₁ is primarily exerted on the other dominant modes and progressively fades for less important PCs. This behavior is consistent with our empirical observations from both Langevin-based models as well as simple molecular systems, in which the slowest collective variable typically appears as a putative causal driver.^{72,73}

The putative causal links between pairs of PCs are summarized in Fig. 2 using the normalized IGT values defined in Eq. (3). The graph highlights the dominant links between PC₁ and PC_{2,3,4,5}. In agreement with the time-resolved IG curves, the strongest directed links correspond to PC₁ $\rightarrow$ PC₂ and PC₁ $\rightarrow$ PC₃, confirming that the leading mode acts as the main driver of the next two dominant modes. A weaker directional asymmetry is also observed for the PC₁ $\leftrightarrow$ PC_{4,5} links, which remain biased toward the PC₁ $\rightarrow$ PC_{4,5} direction. However, the magnitude of these links is markedly smaller than for PC₂ and PC₃, indicating that the causal influence of PC₁ rapidly weakens as one moves toward lower-variance PCs. As we




FIG. 2. Ubiquitin: putative causal links between PC₁ and the remaining leading principal components. Nodes represent the first five PCs, and directed edges indicate the normalized integrated imbalance gain, IGT, defined in Eq. (3). For each pair, the edge length and width, as well as the numerical label, report the relative strength of the directional information transfer, with larger values corresponding to stronger causal links. The strongest links are observed from PC₁ toward PC₂ and PC₃, whereas the links involving PC₄ and PC₅ are substantially weaker. For each PC, the residues with the largest contributions to the corresponding eigenvector are highlighted in blue on the protein structure, as quantified using the Brunschweiler collectivity index. The arrows in blue show the corresponding directions of the eigenvector for the highlighted residues, while the rest of the protein is shown in gray.

will see in Sec. III B, however, this is not necessarily a universal rule. In addition, Fig. S9 of the [supplementary material](#) shows that PC_2 exhibits appreciable asymmetric IG with both PC_3 and PC_4 . In these cases, the dominant direction of information transfer is $PC_2 \rightarrow PC_{3,4}$, while the corresponding reverse links are noticeably weaker. By contrast, the IG between PC_2 and PC_5 is much smaller and only weakly directional. The same trend persists for lower-variance PCs, for which the causal signal rapidly decreases and eventually becomes negligible.

In Fig. 2, we also highlight the residues that contribute most strongly to each principal component. These contributions were quantified using the Brüschweiler collectivity index, which quantifies the degree to which the collective motions are delocalized in the protein.^{80,81} Details are provided in Fig. S5 of the [supplementary material](#). This analysis shows that PC_1 is the most delocalized mode among those considered, with a collectivity index of ~ 0.8 , indicating that residues distributed across the entire protein contribute substantially to this collective fluctuation. By contrast, PC_2 – PC_5 are more localized, with collectivity indices in the range 0.4–0.6. The dominant role of PC_1 would suggest a scenario in which more collective protein motions have a causal influence on more spatially localized fluctuations. In Sec. III B, we will show that this scenario does not hold in general.

B. NTL9

In the second application, we analyze folded-state segments of the NTL9 trajectory. In Ref. 12, NTL9 was simulated for ~ 3 ms, during which the protein undergoes multiple reversible folding–unfolding transitions. Within these simulations, the protein remains folded for several extended intervals; in particular, we identified four independent folded segments longer than 200 μ s. The results discussed in the following are obtained by analyzing one representative folded segment. The same analysis performed on the remaining folded segments yields qualitatively consistent results, as reported in the [supplementary material](#) (see the section “NTL9

results for another folded trajectory fragment” and Figs. S13–S17 of the [supplementary material](#)). This allows us to focus specifically on causal relationships between PCs within the folded-state dynamics, rather than on the much larger structural changes associated with folding and unfolding. The eigenvalue spectrum and the explained variance of the PCs are shown in Fig. S3 of the [supplementary material](#).

Figure 3(a) shows the autocorrelation functions of the first five principal components. In contrast to ubiquitin, where the leading modes decorrelate on more comparable sub-microsecond timescales, NTL9 displays a clearer separation between the slowest mode and the remaining PCs. In particular, PC_1 decorrelates on a timescale of ~ 1 μ s, whereas $PC_{2,3,4,5}$ relax substantially faster, with characteristic timescales of roughly 200 ns. This separation indicates that PC_1 captures a distinct slow fluctuation within the folded-state ensemble of NTL9, while the remaining leading PCs describe faster modes embedded within the same folded basin.

We then examine whether this separation of timescales is reflected in the directionality of information transfer between the leading PCs. In Fig. 3(b), we report the pairwise IG curves between PC_1 and $PC_{2,3,4,5}$. In all cases, the curves show a pronounced asymmetry, with the $PC_1 \rightarrow PC_{2,3,4,5}$ direction consistently dominating over the reverse direction. Thus, knowledge of the slow coordinate PC_1 substantially improves the prediction of the future evolution of the faster leading PCs, whereas the converse effect is much weaker. Compared with ubiquitin, where the directional influence of PC_1 is strongest for PC_2 and PC_3 but rapidly decreases for lower-variance modes, the NTL9 folded segment exhibits a more systematic and persistent directional coupling from PC_1 toward all the other PCs considered (see Fig. S8 of the [supplementary material](#)).

This trend is summarized in Fig. 4, where we report the putative causal links constructed from the normalized IGT values defined in Eq. (3). The graph confirms the behavior observed in the time-resolved IG curves: all links involving PC_1 are strongly asymmetric and predominantly oriented from PC_1 toward $PC_{2,3,4,5}$. By contrast, when PC_2 is taken as the reference mode, the links with $PC_{3,4,5}$ are




FIG. 3. NTL9: (a) the autocorrelation functions of the top 5 PCs, $PC_{i=1,5}$, which are colored red, blue, green, orange, and purple, respectively, with shaded areas denoting error bars. (b) The pairwise IG curves over time between $PC_1 \leftrightarrow PC_{2,3,4,5}$. The IG values are on the Y-axis, and time lag (τ) in nanoseconds is on the X-axis. Pink lines denote the direction $PC_1 \rightarrow PC_{2,3,4,5}$, whereas green lines denote the reverse direction. Shaded areas denote error bars computed over 20 independent realizations.




FIG. 4. NTL9: putative causal links between PC_1 and the remaining leading principal components. Nodes represent the first five PCs, and directed edges indicate the normalized integrated Imbalance Gain, IGT, defined in Eq. (3). For each pair, the edge length and width, as well as the numerical label, report the relative strength of the directional information transfer, with larger values corresponding to stronger causal links. PC_1 displays strong directional influence toward all other leading PCs considered, identifying it as a dominant slow driver within the folded-state dynamics of NTL9. For each PC, the residues with the largest contributions to the corresponding eigenvector are highlighted in blue on the protein structure, as quantified using the Brüschweiler collectivity index. The arrows in blue show the corresponding directions of the eigenvector for the highlighted residues, while the rest of the protein is shown in gray.

largely reciprocal and of comparable magnitude, as shown in Fig. S10 of the [supplementary material](#). Taken together, these results identify PC_1 as a dominant slow mode that directionally conditions the dynamics of several faster principal components within the folded state of NTL9.

We again use the Brüschweiler collectivity index^{80,81} to identify the residues that contribute most strongly to each principal component; these residues are highlighted in blue in Fig. 4. Additional details for this particular protein are provided in Fig. S6 of the [supplementary material](#). This analysis reveals a trend that differs markedly from ubiquitin. In ubiquitin, PC_1 is the most delocalized mode, with contributions distributed across the protein. By contrast, in NTL9, PC_1 is comparatively localized, with a collectivity index of ~ 0.4 . The collectivity indices of $PC_{2,3,4,5}$ are 0.52, 0.38, 0.29, and 0.65, respectively. Structurally, PC_1 is primarily localized around the β -turn region of the protein, and PC_2 , PC_3 , and PC_4 also involve motions broadly concentrated in this region, with additional secondary contributions. This suggests that the strong information transfer from PC_1 to these modes reflects a local hierarchy of motions, in which a slow fluctuation centered around the β -turn organizes faster fluctuations in the same structural region.

A particularly interesting case is PC_5 , which is substantially more delocalized than the other higher-order modes considered here. Despite this difference in spatial character, the $PC_1 \rightarrow PC_5$ link remains strongly asymmetric. This observation indicates that directional information transfer between PCs is not necessarily determined by localization or delocalization. In the case of NTL9, a localized slow fluctuation can carry predictive information about a more delocalized mode involving other parts of the protein. This result highlights the nontrivial character of information transfer in the PC representation and suggests that the present framework may

provide a useful route for identifying dynamical couplings between motions occurring in distinct regions of a protein, with potential relevance for future studies of allosteric communication.

IV. DISCUSSION

In this work, we have investigated the presence of putative causal links between principal components extracted from long equilibrium protein trajectories. Principal component analysis is most often used as a dimensionality-reduction tool, where the leading modes are interpreted in terms of variance, collectivity, and structural deformation. Here, we have taken a complementary perspective and asked whether the present value of one PC improves the prediction of the future evolution of another PC. By using the Imbalance Gain (IG) approach,⁷² we obtain a time-lagged and directional measure of information transfer between collective modes. Importantly, the directional information transfer identified by our approach cannot be identified by the TICA analysis, since the time-delayed covariance matrix between PCs is symmetric if the dynamics satisfies detailed balance, as is the case for molecular dynamics.

The two systems studied here display qualitatively distinct patterns of directional coupling. In ubiquitin, the leading principal component, which is also the slowest one, is a highly delocalized, protein-wide mode and exhibits the strongest information transfer toward PC_2 and PC_3 . The influence of PC_1 becomes progressively weaker and more symmetric for higher-order PCs, indicating that in this case, the dominant collective fluctuation primarily influences a few leading modes. This behavior is consistent with a picture in which a broad native-state fluctuation provides a slowly varying dynamical context for more localized motions, while lower-variance modes become increasingly weakly coupled to the dominant PC.

In NTL9, PC₁ is clearly the slowest mode and displays a pronounced directional influence on all the other leading PCs considered. Unlike ubiquitin, however, this dominant mode is not the most delocalized one. Instead, PC₁ is localized around the β -turn region, while some of the modes receiving information from it, in particular, PC₅, are more spatially extended. This result shows that directional information transfer between PCs is not simply determined by the collectivity or variance of the corresponding modes. A localized slow fluctuation can carry predictive information about faster and more delocalized motions, suggesting that the causal organization of PC space can reveal dynamical relationships that are not apparent from the PCA spectrum or collectivity indices alone.

The relationship between the free-energy landscape and causal coupling remains an important open question. While the FES of the PCs of ubiquitin and NTL9 do not reveal distinct metastable substates within the folded basin (see Figs. S11 and S12 of the [supplementary material](#)), they exhibit different degrees of asymmetry and confinement. In particular, in ubiquitin, PC₁ displays a slightly more rugged profile, whereas the remaining PCs are characterized by relatively smooth, valley-like free-energy landscapes. Instead, in NTL9, the free energy rises more steeply along PC₁, with no signs of ruggedness. As shown previously,⁷³ asymmetries in the FES can bias dynamical couplings between collective variables. However, these free energy landscapes show that for the systems analyzed in this work, the landscape by itself is not a sufficient predictor of the asymmetries. Future studies that condition causality analyses on different metastable states, particularly along folding–unfolding landscapes, may provide further insight into how state-dependent relaxation processes govern causal relationships in biomolecular dynamics.

It is important to remark that PCA and the IG between PCs probe two different aspects. PCA identifies directions of large-amplitude structural variance, whereas IG asks whether fluctuations along one such direction improve the prediction of another at later times. The first criterion is geometric and instantaneous, while the second is dynamical and time-lagged. Therefore, the most informative causal driver need not be the PC with the largest variance or the most spatially delocalized eigenvector. In ubiquitin, these properties approximately coincide: the most collective mode is also the dominant source of directional information transfer. In NTL9, they separate: the slowest and most predictive mode is relatively localized, yet it influences both localized and more distributed fluctuations. This distinction highlights the value of studying causal structure in PC space rather than relying exclusively on variance, relaxation times, or spatial collectivity. In the systems analyzed so far, we observe several cases in which the IGT values are broadly symmetric between pairs of PCs (e.g., between PC₂ $\leftrightarrow$ PC_{3,4,5} in NTL9, see Fig. S10 of the [supplementary material](#)). However, we have not yet identified a case in which the dominant causal direction is inverted relative to the ordering of explained variance; that is, we do not observe lower-variance PCs acting as stronger causal drivers of higher-variance PCs.

The present results suggest several natural extensions. One direction is to move from pairwise causal graphs toward multivariate formulations in which the influence of a PC is assessed while accounting for the simultaneous effects of other modes. The approach introduced in Ref. 82 might allow us to perform this analysis, with the caveat that the statistical uncertainty of the inferred

collective links should be carefully controlled. Performing this analysis would help distinguish whether the observed links correspond to direct dynamical couplings or are mediated by intermediate PCs. A second direction is to connect PC-level causality to residue-level or structural mechanisms, for example, by combining IG with mode localization, contact maps, or residue-wise dynamical descriptors. This would be particularly useful for systems such as NTL9, where a localized slow fluctuation appears to influence more extended protein motions.

Our results suggest an interpretation of the directional coupling in terms of the residues that contribute to each PC. In simple terms, there are certain residue fluctuations that can have an asymmetric influence on other residue fluctuations. Our analysis should be considered a first step toward an approach for identifying those residues. The causal asymmetry observed between PCs also suggests a route toward experimentally testable interventions in protein dynamics. Because each PC can be decomposed into residue-level contributions, residues with large contributions to a putative driver PC constitute candidate sites for localized perturbations whose effects may subsequently emerge in the motion described by a driven PC. In this sense, the inferred directionality predicts how perturbing one collective mode may alter the future evolution of another and identifies the timescale over which that response becomes detectable. This perspective could be applied to proteins with known allosteric regulation, ligand binding, or conformational switching, where directional couplings between collective modes may provide a compact description of how local perturbations propagate through the protein.^{4,67,83–92} Applied without prior assumptions about known allosteric sites or pathways, the approach could help identify protein regions that may initiate an allosteric response and characterize how the resulting dynamical influence propagates through the structure. Similarly, an optical probe placed on specific sites on the protein surface could provide an experimentally accessible measure of a particular asymmetric influence between different protein regions. Monitoring its response after the photoexcitation of a chromophore would offer a way to test the predicted directional coupling, and our approach could guide the selection of perturbation and observation sites in future experiments.

SUPPLEMENTARY MATERIAL

The [supplementary material](#) includes a discussion of the different MD trajectories that are analyzed, a description of the Brüschweiler collectivity index, results for NTL9 from another trajectory fragment that show the consistency of the results, [supplementary material](#) figures that show the PCA eigenvalue spectra, the cosine similarity of PCs between two halves of the analyzed trajectories to show consistency of the PCs across both halves, the Brüschweiler collectivity index plots, the IG curves from PC₁ $\leftrightarrow$ PC_{6,7,8,9}, and the IG curves and corresponding IGT graphs from PC₂ $\leftrightarrow$ PC_{3,4,5} for both systems.

ACKNOWLEDGMENTS

D.B. and A.H. thank the European Commission for funding from ERC Grant No. HyBOP 101043272. This work was partially funded by NextGenerationEU through the Italian National Center for HPC, Big Data, and Quantum Computing (Grant No.

CN00000013 received by A.L.). The authors thank Shaw Research for sharing their ubiquitin trajectory from Ref. 20 and their NTL9 trajectory from Ref. 12. The authors also acknowledge the use of AI tools to improve the writing and readability of the manuscript.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Debarshi Banerjee: Conceptualization (equal); Formal analysis (lead); Investigation (equal); Methodology (equal); Software (lead); Writing – original draft (equal); Writing – review & editing (equal). **Ali Hassanali:** Conceptualization (equal); Investigation (equal); Methodology (equal); Supervision (equal); Writing – original draft (equal); Writing – review & editing (equal). **Alessandro Laio:** Conceptualization (equal); Investigation (equal); Methodology (equal); Supervision (equal); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are openly available at https://github.com/debarshibanerjee/pca_causality_proteins. The trajectories were provided by Shaw Research.

REFERENCES

- 1 M. Karplus and J. A. McCammon, *Nat. Struct. Biol.* **9**, 646 (2002).
- 2 K. Henzler-Wildman and D. Kern, *Nature* **450**, 964 (2007).
- 3 D. D. Boehr, R. Nussinov, and P. E. Wright, *Nat. Chem. Biol.* **5**, 789 (2009).
- 4 J. Guo and H.-X. Zhou, *Chem. Rev.* **116**, 6503 (2016).
- 5 K. Nam and M. Wolf-Watz, *Struct. Dyn.* **10**, 014301 (2023).
- 6 R. Grandori, *Front. Biophys.* **1**, 1219843 (2023).
- 7 J. Ono, Y. Matsumura, T. Mori, and S. Saito, *J. Phys. Chem. B* **128**, 20 (2023).
- 8 G. Haran, *Biophys. J.* **123**, 321a (2024).
- 9 C. E. Hatton and P. Mehrabi, *Biophys. Rev.* **16**, 563 (2024).
- 10 D. E. Shaw, R. O. Dror, J. K. Salmon, J. P. Grossman, K. M. Mackenzie, J. A. Bank, C. Young, M. M. Deneroff, B. Batson, K. J. Bowers *et al.*, in *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis* (ACM, 2009), pp. 1–11.
- 11 J. L. Klepeis, K. Lindorff-Larsen, R. O. Dror, and D. E. Shaw, *Curr. Opin. Struct. Biol.* **19**, 120 (2009).
- 12 K. Lindorff-Larsen, S. Piana, R. O. Dror, and D. E. Shaw, *Science* **334**, 517–520 (2011).
- 13 R. O. Dror, R. M. Dirks, J. P. Grossman, H. Xu, and D. E. Shaw, *Annu. Rev. Biophys.* **41**, 429 (2012).
- 14 T. J. Lane, D. Shukla, K. A. Beauchamp, and V. S. Pande, *Curr. Opin. Struct. Biol.* **23**, 58 (2013).
- 15 S. A. Hollingsworth and R. O. Dror, *Neuron* **99**, 1129 (2018).
- 16 D. J. Huggins, P. C. Biggin, M. A. Dämgén, J. W. Essex, S. A. Harris, R. H. Henchman, S. Khalid, A. Kuzmanic, C. A. Laughton, J. Michel *et al.*, *WIREs Comput. Mol. Sci.* **9**, e1393 (2019).
- 17 D. E. Shaw, P. J. Adams, A. Azaria, J. A. Bank, B. Batson, A. Bell, M. Bergdorf, J. Bhatt, J. A. Butts, T. Correia *et al.*, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (ACM, 2021), pp. 1–11.
- 18 A. P. Bhati, A. Hoti, A. Potterton, M. K. Bieniek, and P. V. Coveney, *J. Chem. Theory Comput.* **19**, 3359 (2023).
- 19 D. E. Shaw, P. Maragakis, K. Lindorff-Larsen, S. Piana, R. O. Dror, M. P. Eastwood, J. A. Bank, J. M. Jumper, J. K. Salmon, Y. Shan, and W. Wrighers, *Science* **330**, 341 (2010).
- 20 K. Lindorff-Larsen, P. Maragakis, S. Piana, and D. E. Shaw, *J. Phys. Chem. B* **120**, 8313–8320 (2016).
- 21 A. Kitao, F. Hirata, and N. Gö, *Chem. Phys.* **158**, 447 (1991).
- 22 T. Ichiye and M. Karplus, *Proteins: Struct., Funct., Bioinf.* **11**, 205 (1991).
- 23 A. E. García, *Phys. Rev. Lett.* **68**, 2696 (1992).
- 24 A. Amadei, A. B. M. Linssen, and H. J. C. Berendsen, *Proteins: Struct., Funct., Bioinf.* **17**, 412 (1993).
- 25 S. Hayward, A. Kitao, and N. Gö, *Proteins: Struct., Funct., Bioinf.* **23**, 177 (1995).
- 26 S. Hayward and N. Gö, *Annu. Rev. Phys. Chem.* **46**, 223 (1995).
- 27 M. A. Balsera, W. Wrighers, Y. Oono, and K. Schulten, *J. Phys. Chem.* **100**, 2567 (1996).
- 28 A. Kitao and N. Gö, *Curr. Opin. Struct. Biol.* **9**, 164 (1999).
- 29 G. G. Maisuradze, A. Liwo, and H. A. Scheraga, *J. Mol. Biol.* **385**, 312 (2009).
- 30 C. C. David and D. J. Jacobs, *Protein Dynamics: Methods and Protocols* (Springer, 2013), pp. 193–226.
- 31 R. Cossio-Pérez, J. Palma, and G. Pierdominici-Sottile, *J. Chem. Inf. Model.* **57**, 826 (2017).
- 32 A. Kitao, *J* **5**, 298 (2022).
- 33 M. Buchner, B. M. Ladanyi, and R. M. Stratt, *J. Chem. Phys.* **97**, 8522 (1992).
- 34 L. M. Horstink, R. Abseher, M. Nilges, and C. W. Hilbers, *J. Mol. Biol.* **287**, 569 (1999).
- 35 H. J. Berendsen and S. Hayward, *Curr. Opin. Struct. Biol.* **10**, 165 (2000).
- 36 P.-L. Chau, D. M. F. van Aalten, R. P. Bywater, and J. B. C. Findlay, *J. Comput.-Aided Mol. Des.* **13**, 11 (1999).
- 37 Y. Mu, P. H. Nguyen, and G. Stock, *Proteins: Struct., Funct., Bioinf.* **58**, 45 (2005).
- 38 F. Capozzi, C. Luchinat, C. Micheletti, and F. Pontiggia, *J. Proteome Res.* **6**, 4245 (2007).
- 39 A. Altis, P. H. Nguyen, R. Hegger, and G. Stock, *J. Chem. Phys.* **126**, 244111 (2007).
- 40 F. Pontiggia, A. Zen, and C. Micheletti, *Biophys. J.* **95**, 5901 (2008).
- 41 A. Altis, M. Otten, P. H. Nguyen, R. Hegger, and G. Stock, *J. Chem. Phys.* **128**, 245102 (2008).
- 42 G. Morra, R. Potestio, C. Micheletti, and G. Colombo, *PLoS Comput. Biol.* **8**, e1002433 (2012).
- 43 G. G. Maisuradze, A. Liwo, P. Senet, and H. A. Scheraga, *J. Chem. Theory Comput.* **9**, 2907 (2013).
- 44 F. Sittel, A. Jain, and G. Stock, *J. Chem. Phys.* **141**, 014111 (2014).
- 45 M. Ernst, F. Sittel, and G. Stock, *J. Chem. Phys.* **143**, 244114 (2015).
- 46 M. Post, S. Wolf, and G. Stock, *J. Chem. Phys.* **150**, 204110 (2019).
- 47 L. Molgedey and H. G. Schuster, *Phys. Rev. Lett.* **72**, 3634 (1994).
- 48 Y. Naritomi and S. Fuchigami, *J. Chem. Phys.* **134**, 065101 (2011).
- 49 G. Pérez-Hernández, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, *J. Chem. Phys.* **139**, 015102 (2013).
- 50 C. R. Schwantes and V. S. Pande, *J. Chem. Theory Comput.* **9**, 2000 (2013).
- 51 F. Noé and C. Clementi, *J. Chem. Theory Comput.* **11**, 5002 (2015).
- 52 G. Pérez-Hernández and F. Noé, *J. Chem. Theory Comput.* **12**, 6118 (2016).
- 53 F. Nüske, B. G. Keller, G. Pérez-Hernández, A. S. J. S. Mey, and F. Noé, *J. Chem. Theory Comput.* **10**, 1739 (2014).
- 54 M. M. Sultan and V. S. Pande, *J. Chem. Theory Comput.* **13**, 2440 (2017).
- 55 S. Schultze and H. Grubmüller, *J. Chem. Theory Comput.* **17**, 5766 (2021).
- 56 L. Bonati, G. Piccini, and M. Parrinello, *Proc. Natl. Acad. Sci. U. S. A.* **118**, e2113533118 (2021).
- 57 J. Pearl, *Causality: Models, Reasoning and Inference*, 2nd ed. (Cambridge University Press, 2009).
- 58 P. Spirtes, *Causation, Prediction and Search* (MIT Press, 2000).
- 59 J. Runge, *Chaos* **28**, 075310 (2018).
- 60 C. W. J. Granger, *Econometrica* **37**, 424–438 (1969).
- 61 A. Gorecki, J. Trylska, and B. Lesyng, *EPL* **75**, 503 (2006).
- 62 H. Kamberaj and A. van der Vaart, *Biophys. J.* **97**, 1747 (2009).

- ⁶³Y. Qi and W. Im, *J. Chem. Theory Comput.* **9**, 3799 (2013).
- ⁶⁴M. Sobieraj and P. Setny, *J. Chem. Theory Comput.* **18**, 1936 (2022).
- ⁶⁵T. Schreiber, *Phys. Rev. Lett.* **85**, 461 (2000).
- ⁶⁶K. Hlavackovaschindler, M. Palus, M. Vejmelka, and J. Bhattacharya, *Phys. Rep.* **441**, 1–46 (2007).
- ⁶⁷A. Hacsuleyman and B. Erman, *PLoS Comput. Biol.* **13**, e1005319 (2017).
- ⁶⁸A. Hacsuleyman and B. Erman, *Proteins: Struct., Funct., Bioinf.* **85**, 1056 (2017).
- ⁶⁹T. Hempel, N. Plattner, and F. Noé, *J. Chem. Theory Comput.* **16**, 2584 (2020).
- ⁷⁰J. Runge, A. Gerhardus, G. Varando, V. Eyring, and G. Camps-Valls, *Nat. Rev. Earth Environ.* **4**, 487 (2023).
- ⁷¹A. Glielmo, C. Zeni, B. Cheng, G. Csányi, and A. Laio, *PNAS Nexus* **1**, pgac039 (2022).
- ⁷²V. D. Totto, G. Fortunato, D. Bueti, and A. Laio, *Proc. Natl. Acad. Sci. U. S. A.* **121**, e2317256121 (2024).
- ⁷³V. Del Totto, D. Banerjee, A. Hassanali, and A. Laio, *J. Chem. Phys.* **162**, 244105 (2025).
- ⁷⁴L. Huet, V. D. Totto, D. Banerjee, A. Laio, and A. A. Hassanali, *arXiv:2604.19491* (2026).
- ⁷⁵E. D. Donkor, A. Laio, and A. Hassanali, *J. Chem. Theory Comput.* **19**, 4596–4605 (2023).
- ⁷⁶E. D. Donkor, A. Offei-Danso, A. Rodriguez, F. Sciortino, and A. Hassanali, *J. Phys. Chem. Lett.* **15**, 3996–4005 (2024).
- ⁷⁷S. Di Pino, D. Banerjee, M. Monti, G. D. Mirón, G. Cassone, and A. Hassanali, *J. Chem. Phys.* **163**, 184505 (2025).
- ⁷⁸D. Banerjee, G. D. Mirón, A. Rodriguez, and A. Hassanali, *arXiv:2605.08381* (2026).
- ⁷⁹S. Dutta, M. Ghosh, and J. Chakrabarti, *Sci. Rep.* **7**, 40439 (2017).
- ⁸⁰R. Brüschweiler, *J. Chem. Phys.* **102**, 3396–3403 (1995).
- ⁸¹J. Prompers, S. Lienin, and R. Brüschweiler, *Biocomputing 2001* (World Scientific, 2000), pp. 79–88.
- ⁸²M. Allione, V. Del Totto, and A. Laio, *Phys. Rev. Lett.* **135**, 047401 (2025).
- ⁸³I. Rivalta, M. M. Sultan, N.-S. Lee, G. A. Manley, J. P. Loria, and V. S. Batista, *Proc. Natl. Acad. Sci. U. S. A.* **109**, E1428 (2012).
- ⁸⁴A. D. Vogt and E. Di Cera, *Biochemistry* **51**, 5894 (2012).
- ⁸⁵A. T. VanWart, J. Eargle, Z. Luthey-Schulten, and R. E. Amaro, *J. Chem. Theory Comput.* **8**, 2949 (2012).
- ⁸⁶S. Hertig, N. R. Latorraca, and R. O. Dror, *PLoS Comput. Biol.* **12**, e1004746 (2016).
- ⁸⁷S. Bowerman and J. Wereszczynski, *Methods Enzymol.* **578**, 429 (2016).
- ⁸⁸I. Rivalta, G. P. Lisi, N.-S. Snoeberger, G. Manley, J. P. Loria, and V. S. Batista, *Biochemistry* **55**, 6484 (2016).
- ⁸⁹J. Wang, A. Jain, L. R. McDonald, C. Gambogi, A. L. Lee, and N. V. Dokholyan, *Nat. Commun.* **11**, 3862 (2020).
- ⁹⁰S. J. Wodak, E. Paci, N. V. Dokholyan, I. N. Berezovsky, A. Horovitz, J. Li, V. J. Hilser, I. Bahar, J. Karanicolas, G. Stock *et al.*, *Structure* **27**, 566 (2019).
- ⁹¹A. Hardie, B. P. Cossins, S. Lovera, and J. Michel, *Commun. Chem.* **6**, 125 (2023).
- ⁹²M. Bernetti, S. Bosio, V. Bresciani, F. Falchi, and M. Masetti, *Curr. Opin. Struct. Biol.* **86**, 102820 (2024).