

Data-Driven and Model-Based Causal Discovery in Sparse Multivariate Time Series: A Comparative Framework for Sustainable Development Systems

Joao Pita Costa[†]
IRCAI
Institut Jozef Stefan
Ljubljana
joao.pitacosta@ijs.si

Andrej Srakar
IRCAI
Institut Jozef Stefan
Ljubljana
andrej.srakar@ijs.si

Oleksandra Topal
Department of AI
Institut Jozef Stefan
Ljubljana
oleksandra.topal@ijs.si

Jan Šturm
Department of AI
Institut Jozef Stefan
Ljubljana
jan.sturm@ijs.si

Abstract / Povzetek

This paper investigates Sustainable Development Goal (SDG) indicators as sparse multivariate time series, focusing on the identification and interpretation of relationships between development dimensions over time. Three complementary approaches are explored: SHAP-based explainable machine learning to identify influential indicators, Transfer Entropy to assess directional information flows, and StreamStory to characterize temporal dynamics through multivariate states and transitions. The study investigates whether causal structures inferred from empirical time-series data can be validated through complementary analytical approaches and compared with model-derived causal structures. Particular attention is given to the degree of agreement between methods, as converging evidence may provide stronger support for inferred relationships and reveal previously unrecognized system dynamics. It resolves into a problem of SDG causal discovery under small- n , multivariate, nonlinear, and incomplete time-series conditions. The results contribute to a comparative framework for exploring, validating, and interpreting causal patterns in sparse multivariate development data, while highlighting the methodological challenges associated with causal discovery from short, incomplete, and highly interdependent time series.

Keywords / Ključne besede

SDG; Causal Discovery; Multivariate Time Series; Transfer Entropy; Explainable Artificial Intelligence.

1 Introduction

The UN SDGs provide a comprehensive framework for understanding global development through 17 interconnected goals and a large system of statistical indicators. These indicators cover diverse dimensions including poverty, health, education, infrastructure, economic development, inequality, and

environmental sustainability. Data are collected and maintained by international organizations such as the United Nations and the World Bank, as well as aggregated repositories such as Our World in Data. Most SDG indicators are available as longitudinal observations across countries, providing an important empirical basis for studying development trajectories over time. However, the same characteristics that make these datasets valuable for monitoring also make them particularly challenging for causal analysis.

SDG indicators are typically observed at annual frequency, frequently contain missing observations, have heterogeneous temporal coverage, and are strongly interdependent across sectors and geographical contexts. These characteristics create a fundamental methodological problem for the direct application of causal discovery techniques [1]. In contrast to typical causal discovery, an SDG time series may contain only a few observations for a country, with substantial gaps between them and changes in measurement practices over time.

Traditional causal inference relies heavily on predefined structural causal models or expert-provided Directed Acyclic Graphs (DAGs) (Pearl, 2009). However, in complex, high-dimensional applications typically encountered in knowledge discovery and data mining, complete domain expertise is rarely available, making manual graph construction error-prone or impossible. Causal discovery bridges this fundamental gap by operating as a data-driven inverse problem; it automates the learning of the underlying causal skeleton and directions directly from observational or temporal datasets (Vowels et al., 2022; Zhang et al., 2025). Understanding its advantages is of great importance for causality research at present.

The problem is further complicated by the multivariate nature of sustainable development. SDG indicators do not evolve independently: economic, social, environmental, health, and infrastructure variables can influence one another, while common external factors can simultaneously affect multiple indicators. Development systems may also contain delayed effects and feedback mechanisms. The resulting time series can therefore exhibit nonlinear dependencies, temporal autocorrelation, common trends, and bidirectional relationships. The original analysis similarly identifies SDG systems as short, noisy, highly interdependent, nonlinear, and potentially characterized by feedback loops. These properties create opportunities for causal discovery, but also increase the risk of

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Information Society 2026, 5–9 October 2026, Ljubljana, Slovenia

© 2025 Copyright held by the owner/author(s).

<http://doi.org/10.70314/is.2026.sikdd.122>

interpreting statistical dependence or temporal precedence as evidence of causation.

The present study builds on this perspective by treating the SDG indicator system as a collection of sparse multivariate time series and examining three complementary analytical approaches: SHAP for interpretable predictive influence, Transfer Entropy for directional temporal dependencies, and StreamStory for multivariate state transitions [4]. This provides a richer econometric and statistical overview about the relationships studied.

2 Data and Methods

Our framework is intended to provide a more cautious and technically grounded approach to causal discovery in sustainable development systems, where limited observations, missing data, nonlinear interactions, and feedback mechanisms constrain the application of conventional methods.

The dataset considered in this study spans 1987–2023, giving a maximum of approximately 289 indicators with annual observations¹ for a complete country-level series², mostly sourced at [5] and [6], while substantial portions of the underlying indicator matrix are incomplete. Consequently, methods designed for longer or denser time series cannot necessarily be transferred directly to this setting without examining their statistical stability and sensitivity to preprocessing choices. Missing data represent an additional and particularly important challenge. Gaps in SDG datasets are not necessarily random: reporting capacity, national statistical systems, indicator definitions, and geographical or socioeconomic characteristics can influence whether an observation is available. The resulting missingness can therefore introduce systematic biases rather than simply reducing sample size. Moreover, imputation is not neutral from a causal perspective. Reconstructing missing observations may smooth temporal trajectories, alter autocorrelation structures, or introduce relationships between variables that were not directly observed. The study therefore treats completeness, missingness patterns, and imputation as explicit components of the analytical pipeline rather than as a purely technical preprocessing step.

These limitations are particularly relevant when applying established causal-discovery methods. Granger causality provides a well-established framework for testing whether the past of one time series improves the prediction of another, but its conventional formulations are sensitive to assumptions concerning stationarity and model specification. The original study found that direct application of Granger causality was problematic for the SDG data because many series were non-stationary. Other causal discovery approaches, including Structural Equation Models (SEM), the PC and FCI algorithms, and LiNGAM, impose assumptions concerning linearity, acyclicity, independence, or stationarity that may not correspond well to complex development systems with feedback and nonlinear interactions. Information-theoretic approaches provide an alternative perspective. Transfer Entropy (TE), introduced by Schreiber, measures the directed information transferred from

one time series to another while accounting for information contained in the target's own history³. Formally, it is introduced as:

$$TE_{X \rightarrow Y} = I(Y_{t+1}; X_t^{(k)} | Y_t^{(l)})$$

where the analysis estimates conditional mutual information I from short empirical sequences X and Y at a given fixed time t .

Unlike conventional correlation measures, TE can capture nonlinear and time-lagged dependencies without requiring a specific parametric model. The previous study therefore used TE to examine directional relationships between influential SDG indicators and target variables, including bidirectional comparisons to distinguish potential upstream and downstream relationships. However, the flexibility of TE does not eliminate the fundamental limitations imposed by short time series. Estimation of information-theoretic quantities can itself be sensitive to sample size, history length, estimator choice, noise, and preprocessing. In the previous implementation, for example, TE was calculated independently for each country using a fixed history length of four observations. This raises an important methodological question: to what extent can the resulting directional relationships be considered robust when the available empirical evidence is intrinsically limited?

A complementary perspective, building on the work published in [2] and [3] focusing SDG 1, is provided by explainable machine learning. SHAP (SHapley Additive exPlanations) can quantify the contribution of individual features to the predictions of a machine-learning model³ (see Figure 2) and thereby provide an interpretable representation of predictive relevance. In the previous study, Random Forest models were used to reduce the original high-dimensional indicator space, after which SHAP was used to identify the most influential indicators for each target variable. This approach is useful because it distinguishes the question of which variables are informative for prediction from the more demanding question of which variables participate in directional temporal relationships. Importantly, the two notions need not coincide (see Figure 1).




Figure 1: SHAP (SHapley Additive exPlanations) values attribute to each feature the change in the expected model prediction when conditioning on that feature (source: Lundberg and Lee, 2017).

The previous results already indicated that some indicators with high SHAP importance had relatively weak Transfer Entropy, whereas some indicators with moderate predictive importance exhibited stronger temporal information transfer. This divergence provides an important motivation for comparing complementary analytical approaches rather than relying on a single causal-discovery method.

¹ The detailed list of SDG indicator timeseries collected <https://github.com/IRCAI-SDGobservatory/data/blob/main/IRCAI-SDGobservatory-Indicators.csv>

² Data sources analysis available at IRCAI's SDG Observatory: <https://sdg-observatory.ircai.org/index.php/sdg-16/>.

³ SDG application by IRCAI/E3 with implementation available at <https://python-visualisation.midas.ijs.si/world-map/sdg-16/>




Figure 2: Screenshot of the implementation of SHAP to identify relations between SDG indicators (left), and of the implementation of Transfer Entropy to explore relations of dependency of SDG indicators at country level (right).

A third perspective is offered by StreamStory⁴, which abstracts multivariate time series into discrete states and represents their evolution as transitions within a Markov model (see the SDG-focused Streamstory in Figure 3 and the data exploration tool in Figure 4). Rather than examining individual indicator relationships, this approach captures recurring configurations of multiple indicators and the transitions between them. Such a representation can provide a system-level view that complements feature-level predictive importance and pairwise directional information. In particular, relationships identified through SHAP or Transfer Entropy can be examined in the context of broader transitions between developmental states. The previous study proposed this comparison as a way of examining whether strong information-transfer relationships correspond to observable changes in multivariate system configurations.

While SHAP tells us about model-dependent predictive contribution, whereas TE investigates conditional information transfer through time. StreamStory then operates at a higher level, modelling transitions between multivariate states. This provides us with a preliminary technical comparison of three different representations of dependency. These approaches suggest that causal analysis of SDG indicators should not be reduced to the application of a single causal-discovery algorithm. SHAP addresses predictive relevance, Transfer Entropy addresses directional information flow, and StreamStory provides a representation of system-level temporal dynamics. Their outputs therefore describe different aspects of the same underlying process. The central methodological challenge is consequently not only whether an individual method can produce a causal relationship, but whether independently derived representations provide consistent evidence for that relationship.

This issue becomes particularly important when empirical causal structures are compared with causal structures derived from an explicit model of the system. A model encodes assumptions about how variables interact and provides a representation of expected mechanisms, while empirical time-series analysis provides evidence about patterns actually present in observations. Agreement between the two can provide converging evidence for a relationship. Conversely,

disagreement can be scientifically informative. A relationship detected consistently in empirical data but absent from the model may indicate an incomplete or incorrectly specified mechanism, while a model-predicted relationship not supported by the data may indicate that the mechanism is weak, masked by noise or confounding, operating at another temporal scale, or not represented adequately in the observations. Thus, comparison between data-derived and model-derived causal structures can be treated not only as a validation exercise but also as a mechanism for identifying potentially new system dynamics.




Figure 3: Screenshot of the implementation of Markov chains through Streamstory for SDG analysis




Figure 4: Screenshot of multilevel analysis of time-series data through the SDG Observatory

3 Discussion of Main Results

The results demonstrate that no single analytical approach is sufficient to characterize the dynamics of sparse SDG time-series data. The three methods provide complementary perspectives on the same system. SHAP identifies indicators that are most influential for predicting the target variables, while Transfer Entropy provides information about the direction and strength of temporal dependencies. StreamStory complements these feature-level analyses by representing the evolution of the multivariate system through states and transitions. An important result is the partial divergence between predictive importance and directional information transfer. Indicators with high SHAP importance do not necessarily exhibit strong Transfer Entropy, while some indicators with moderate predictive importance show stronger temporal information transfer. This suggests that predictive

⁴ SDG focused instance with implementation available at <https://sirius.ijs.si/sdg/>

relevance should not be interpreted directly as causal relevance and supports the use of multiple complementary methods when analysing SDG dynamics.

Transfer Entropy (TE) quantifies directional temporal dependence and is applied to multiple SDG indicators and evaluated in both directions, $X \rightarrow Y$ and $Y \rightarrow X$, allowing potential drivers, consequences, and feedback relationships to be identified.

For the SDG analysis, TE is generalized from a simple two-variable relationship into a framework for exploring relationships within a sparse multivariate system. Rather than proposing a new TE formulation, the study extends its application to a sparse multivariate SDG setting, combining TE with SHAP-based predictive relevance and StreamStory-based state transitions to provide complementary evidence for the interpretation and validation of potential causal relationships. The Transfer Entropy results further indicate that relationships between SDG indicators can be directional and potentially bidirectional, with information flowing across different SDG domains. However, given the short and incomplete nature of the available time series, these relationships should be interpreted as evidence of potential causal influence rather than definitive causal effects.

Table 1: Comparative interpretation of the three analytical approaches and causal models

Analytical approach	Primary objective	What it can reveal
SHAP	Identify influential indicators for a target	Which indicators contribute most to prediction of the target
Transfer Entropy	Identify directional temporal dependencies	Potential drivers, consequences, feedbacks, nonlinear and lagged relationships
StreamStory	Characterize system-level temporal dynamics	Recurring development states, transitions, anomalies and broader temporal patterns

Table 1 highlights that the three analytical approaches should be understood as complementary rather than interchangeable sources of causal evidence. SHAP provides an interpretable assessment of predictive relevance, Transfer Entropy extends this analysis by examining directional information flow, and StreamStory represents the data as multivariate states and transitions, allowing recurring developmental configurations and changes in system behaviour to be identified.

Overall, the results support the proposed comparative framework: agreement between complementary methods provides stronger evidence for a relationship, while disagreement identifies relationships that require further investigation. In particular, discrepancies between data-derived relationships and model-derived structures can be treated as potentially informative signals of incomplete model assumptions or previously unrecognized system dynamics. The findings therefore suggest that, for sparse SDG time series, causal

discovery is most appropriately viewed as a process of triangulation, validation, and hypothesis generation, rather than as the direct establishment of causal effects.

4 Conclusions & Future Work

This study explored the use of three complementary approaches—SHAP, Transfer Entropy, and StreamStory—for analysing sparse multivariate time series of Sustainable Development Goal indicators. The results illustrate that these methods capture different dimensions of system behaviour: SHAP identifies predictive relevance, Transfer

An important conclusion is that causal discovery in SDG systems should not be considered as the output of a single algorithm. Instead, the combination and comparison of different analytical perspectives can provide a stronger basis for identifying robust relationships and distinguishing potentially meaningful dynamics from relationships that may result from statistical or methodological artefacts. In particular, agreement between approaches can provide additional support for an inferred relationship, while discrepancies can identify relationships requiring further investigation. Future research will therefore focus on two main directions, case-study-based causal discovery, and improving causal methods for specific SDG systems.

Acknowledgments / Zahvala

This research was partially funded by the European Commission's Horizon research and innovation program under grant agreement 10080425 (ELIA).

References / Literatura

- [1] J. Renner, R. Maestre, and I. Ochoa, "R e x: Causal discovery based on machine learning and explainability techniques," *Pattern Recognition*, 2025.
- [2] Luka Urbanč, Marko Grobelnik and Joao Pita Costa (2024) Predicting poverty using regression. Proceedings of the SIKDD.
- [3] Marko Grobelnik and Joao Pita Costa (2024) An AI-driven Observatory Against Poverty. The Future of Life Institute.
- [4] L. Stopar, P. Skraba, M. Grobelnik, and D. Mladenec, "Streamstory: Exploring multivariate time series on multiple scales," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 4, pp. 1788–1802, 2019.
- [5] Our World in Data team, "SDG tracker: Measuring progress to- wards the sustainable development goals," *Our World in Data*, 2023, <https://ourworldindata.org/sdgs>.
- [6] World Bank, "SDG data bank," (2026) available at: [https://databank.worldbank.org/source/sustainable-development-goals-\(sdgs\)](https://databank.worldbank.org/source/sustainable-development-goals-(sdgs)), 2025, accessed: 2026-01-15.
- [7] W. Zhou, S. Yu, and B. Chen, "Causality detection with matrix-based transfer entropy," *Information Sciences*, vol. 613, pp. 357–375, 2022.
- [8] T. Schreiber, "Measuring information transfer," *Physical Review Letters*, vol. 85, no. 2, pp. 461–464, 2000.

Pearl, J. *Causality*. Cambridge University Press, 2009.

Vowels, M. J., Camgoz, N. C., & Bowden, R. (2022). D'ya like DAGs? A survey on structure learning and causal discovery. *ACM Computing Surveys*, 55(4), 1–36.

Zhang, Z., Zhou, Y., & Zhang, K. (2025). Causal Discovery from Temporal Data: An Overview and New Frontiers. *ACM Transactions on Knowledge Discovery from Data*, 19(3), 1–28.