

CausalXtract: a flexible pipeline to extract causal effects from live-cell time-lapse imaging data

Franck Simon^{1,‡}, Maria Colomba Comes^{2,‡}, Tiziana Tocci^{1,2,‡}, Louise Dupuis¹, Vincent Cabeli¹, Nikita Lagrange¹, Arianna Mencattini², Maria Carla Parrini³, Eugenio Martinelli^{2,*}, Hervé Isambert^{1,*}

¹ CNRS UMR168, Institut Curie, Université PSL, Sorbonne Université, Paris, France

² Department of Electronic Engineering, University of Rome Tor Vergata, Rome, Italy

³ INSERM U830, Institut Curie, Université PSL, Paris, France

‡ these authors contributed equally to this work

* corresponding authors: herve.isambert@curie.fr, martinelli@ing.uniroma2.it

Live-cell microscopy routinely provides massive amount of time-lapse images of complex cellular systems under various physiological or therapeutic conditions. However, this wealth of data remains difficult to interpret in terms of causal effects. Here, we describe CausalXtract, a flexible computational pipeline that discovers causal and possibly time-lagged effects from morphodynamic features and cell-cell interactions in live-cell imaging data. CausalXtract methodology combines network-based and information-based frameworks, which is shown to discover causal effects overlooked by classical Granger and Schreiber causality approaches. We showcase the use of CausalXtract to uncover novel causal effects in a tumor-on-chip cellular ecosystem under therapeutically relevant conditions. In particular, we find that cancer associated fibroblasts directly inhibit cancer cell apoptosis, independently from anti-cancer treatment. CausalXtract uncovers also multiple antagonistic effects at different time delays. Hence, CausalXtract provides a unique computational tool to interpret live-cell imaging data for a range of fundamental and translational research applications.

Live-cell imaging microscopy commonly produces extensive amounts of time-lapse images of cellular systems, which can be segmented to extract morphodynamic features and interactions of individual cells under increasingly complex and physiologically relevant conditions. However, this wealth of information remains largely under-exploited due to a lack of methods and tools able to discover causal effects from spatio-temporal correlations under well-controlled experimental conditions.

CausalXtract addresses this need by integrating an advanced live-cell image feature extraction tool with a reliable and scalable causal discovery method, Fig. 1, in order to learn temporal causal networks from live-cell time-lapse imaging data, Fig. 2.

Results and Discussion

CausalXtract feature extraction and causal discovery modules


CausalXtract’s live-cell image feature extraction module (CellHunter+), Fig. 1b, is based on CellHunter software¹ and consists in three steps: detection, tracking and feature extraction

of live cells within time-lapse video images. First, automatic localization/segmentation of cells (*e.g.* tumor and immune cells) is performed with the Circular Hough Transform (CHT) algorithm² to estimate the cell centers and radii. Second, cell trajectories along the frames are constructed by linking the positions detected at the previous time step through Munkres’ algorithm for Optimal sub-pattern Assignment Problems (OAPs)³. Finally, relevant descriptors related to the shape, motility, and state of the cells, as well as cell-cell interactions are quantified from each cell trajectory (Methods).


CausalXtract’s temporal causal discovery module (tMIIC), Fig. 1c, is adapted from the causal discovery method, MIIC⁴⁻⁶, which learns contemporaneous causal networks (*i.e.* when temporal information is not available) for a broad range of biological or biomedical data, from single-cell transcriptomic and genomic alteration data^{4,7} to medical records of patients^{5,6,8}. Live-cell time-lapse imaging data contain, however, information about cellular dynamics, which can in principle facilitate the discovery of novel cause-effect functional processes, based on the assumption that future events cannot cause past ones. To this end, CausalXtract’s discovery module, tMIIC, reconstructs time-unfolded causal networks, where each variable is represented by several nodes at different relative time points⁹, Fig. 1c. Such a time-unfolded network framework¹⁰⁻¹³ is required to account for the temporal correlation between successive time steps in time series data. We benchmarked tMIIC on synthetic datasets resembling the real-world data of interest analyzed in this study (*i.e.* number of time steps, network size and degree distribution) and found that it matches or outperforms state-of-the-art methods, while running order of magnitudes faster on datasets of biologically relevant size including tens to hundreds of thousands time steps, Supplementary Figs. 1-4.

CausalXtract’s temporal network framework goes beyond the seminal concept of temporal causality originally proposed by Granger¹⁴ for linear time series without reference to graphical models and later extended to non-linear dynamics by Schreiber^{15,16}. In particular, Granger-Schreiber causality is in fact too restrictive and may overlook actual causal effects, that can be uncovered by graph-based causal discovery meth-

a Tumor-on-chip preparation



b CausalXtract's live-cell image feature extraction module (CellHunter+)



c CausalXtract's temporal causal discovery module (tMIIC)

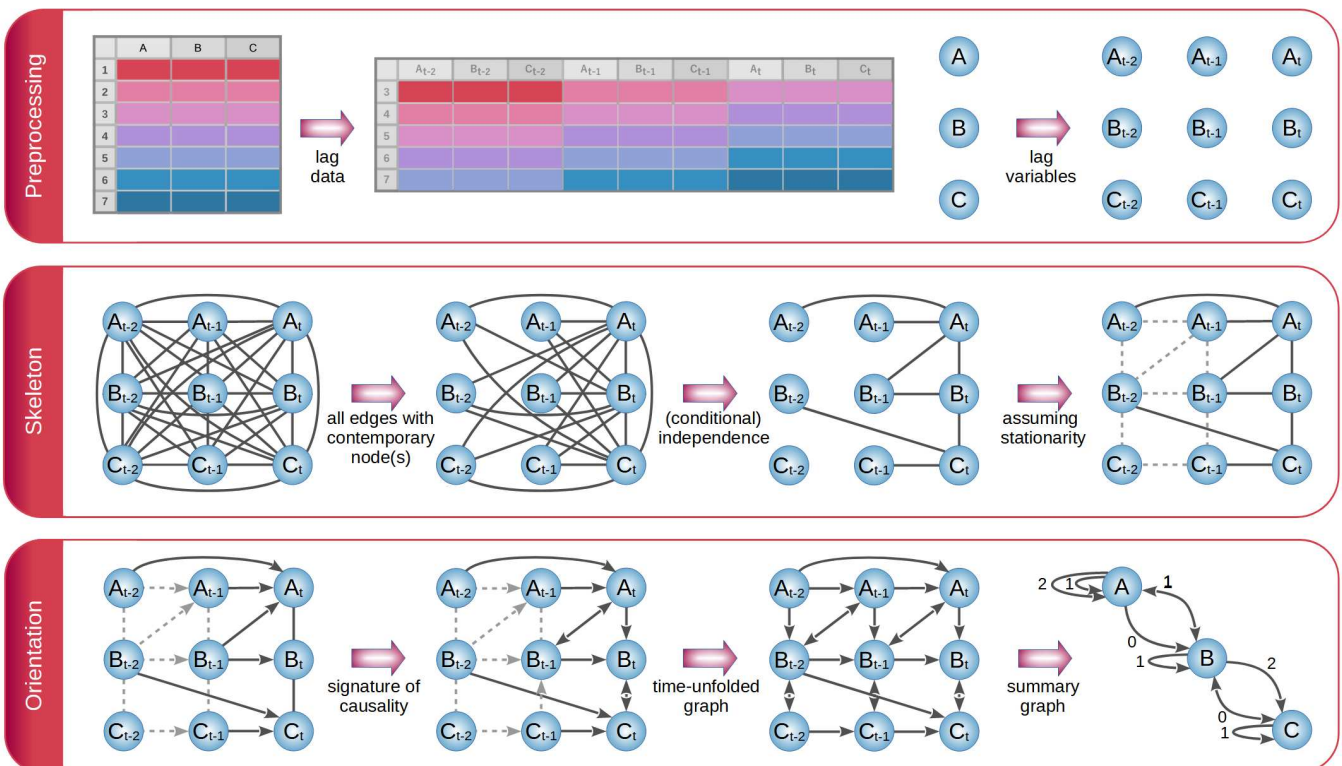


Figure 1: **CausalXtract pipeline.** **a**, Live-cell tumor ecosystem reconstituted *ex vivo*¹ using the tumor-on-chip technology (Methods). **b**, CausalXtract's live-cell image feature extraction module (CellHunter+). The tracking of cancer and immune cells and of their mutual interactions is illustrated in Supplementary Movies 1-3, in absence or presence of cell division and apoptosis event. **c**, CausalXtract's temporal causal discovery module (tMIIC) learns a temporal causal network from the features extracted in (b). See Methods for CausalXtract's implementation details and theoretical foundations. A step-by-step notebook of CausalXtract pipeline is provided with the source code.

ods, Supplementary Fig. 5 (Methods, Theorem 1). In addition, Granger-Schreiber causality has long been known to infer spurious causal associations based on time delays, by excluding the presence of latent common causes *a priori*⁹. CausalXtract circumvents these limitations by combining graph-based and information-based approaches (Methods), while including contemporary and time-delayed effects of unobserved latent variables, that are ubiquitous in cell biology data (*e.g.* the latent effects of cell cycle phases on cellular features and responses).

Application to tumor-on-chip cellular ecosystems

We showcase CausalXtract with the analysis of time-lapse images of a tumor ecosystem reconstituted *ex vivo* using the tumor-on-chip technology, Fig. 1a. These live-cell time-lapse images come from a proof-of-concept study¹ which demonstrated the effects of an anti-cancer drug (the monoclonal antibodies trastuzumab, brand name Herceptin, used to treat HER2+ breast cancers) on a reconstituted tumor microenvironment including cancer cells, immune cells, cancer-associated fibroblasts (CAF), and endothelial cells (Methods). However, a comprehensive extraction and analysis of cellular morphodynamic features and interactions remained unexplored.

To this end, cellular features such as cell geometry, velocity, division, apoptosis, cell-cell transient interactions and persistent contacts were first extracted from the raw images using CausalXtract’s feature extraction module, Fig. 1b and Supplementary Fig. 6. Then, a time-unfolded causal network, Supplementary Fig. 7, and the corresponding summary causal network, Fig. 2a, were reconstructed from extracted cellular features, cell-cell interactions and therapeutic conditions using CausalXtract’s temporal causal discovery module, Fig. 1c.

CausalXtract inferred network, Fig. 2a, uncovers novel biologically relevant findings, in addition to confirming known results from earlier studies. In particular, CausalXtract discovers that CAFs directly inhibit cancer cell apoptosis, independently from anti-cancer treatment, Fig. 2b, while earlier studies reported that CAFs merely reduced the effect of treatment¹. CausalXtract also discovers that treatment increases cancer cell perimeter, Fig. 2c, which has not been reported so far either. In addition, CausalXtract confirms known results from earlier studies. In particular, it recovers that treatment increases cancer cell apoptosis and the number of cancer-immune interactions, as well as decreases the division rate of cancer cells, Fig. 2c. Likewise, CausalXtract recovers that CAFs stimulate cancer cell migration and increase their area, Fig. 2b.

Interestingly, CausalXtract identifies also multiple and possibly antagonistic effects with different time delays. For instance, CausalXtract recovers several antagonistic relations between morphodynamic features such as cell division and eccentricity or cell apoptosis and area, Fig. 2d. Indeed, the late phases of cell division are associated to a marked increase in eccentricity (red edge) but preceded by a net decrease in eccentricity, two to three hours before cytokinesis (blue edges), once the decision to divide has been made (*i.e.* the probable latent cause) and the cell is actually duplicating its biological materials (prophase), Fig. 2d. Likewise, the area change upon apoptosis is predicted to first decrease soon after apoptosis (blue edge) before eventually increasing upon cell lysis (red edge), Fig. 2d. These results are robust to variations in sampling rate, Supplementary Fig. 8.

All in all, CausalXtract is a flexible pipeline which uncovers novel and possibly time-lagged causal relations between cellular features under controlled conditions (*e.g.* drug). CausalXtract uniquely combines live-cell feature extraction with information theory and causal discovery approaches. It consists of two interdependent computational modules, conceived to warrant interoperability with alternative live-cell segmentation and tracking methods or alternative temporal causal discovery methods.

CausalXtract opens up new avenues to analyze live-cell imaging data for a range a fundamental and translational research applications, such as the use of tumor-on-chips to screen immunotherapy responses on patient-derived tumor samples. With the advent of virtually unlimited live-cell image data, flexible hypothesis-free interpretation methods are much needed¹⁷ and we believe that CausalXtract can bring unique insights based on causal discovery to interpret such information-rich live-cell imaging data.

Materials and Methods

Tumor-on-chip preparation and live-cell microscopy

Videos analyzed in the present study refer to biological experiments emulating a 3D breast tumor ecosystem¹. All tumor-on-chip experiments have a central endothelium compartment containing endothelial cells (primary human umbilical vein endothelial cells, HUVECs) and two lateral chambers filled with biomimetic hydrogel (collagen type I at 2.3 mg/mL) seeded with cancer cells (HER2+ breast cancer BT474 cell line) and immune cells (peripheral blood mononuclear cells, PBMCs) from healthy donors, Fig. 1a. Four experimental conditions were considered depending on the presence or absence of breast cancer-associated fibroblasts (CAF cell line Hs578T) and drug treatment (trastuzumab, Herceptin). Videos were acquired by inverted motorized Leica microscopes with a frame rate of 2 minutes for up to 48h (1440 frames). Fig. 1b shows a crop frame with cancer cells, PBMCs and CAFs. Each video was cropped into multiple small 300×300 pixel videos (referred to as crops in the following), each of which represented a field of view at subsequent time frames containing a “main” cancer cell (MCC) initially placed at the center of the image, some PBMC immune cells, other cancer cells and possibly CAFs within the surrounding of the MCC depending on the experimental conditions. 36 video crops of up to 1440 frames were analyzed (46,935 frames in total) corresponding to 9 video crops per experimental conditions.

CausalXtract’s live-cell image feature extraction module

The live-cell image feature extraction module (CellHunter+), Fig. 1b, extends the CellHunter software¹ and consists in three steps: detection, tracking and feature extraction of live cells within time-lapse video images. First, cell detection is based on the segmentation of circular-shaped objects using CHT² with radii set around the theoretical radii of the two cell populations ($r_{im} = 4$ px for immune cells and $r_{ca} = 14$ px for MCCs with a pixel resolution 1 px = $0.645 \mu\text{m}^1$). Then, cell tracking is performed by linking cells detected at the i^{th} frame to cells located at the $(i + 1)^{th}$ frame within a maximum distance from the detected cell candidate. While the motions of both MCCs and immune cells resemble random walks with time-varying drift and volatility, these two cell types exhibit different motility characteristics¹. Hence, different maximum distances are considered for the two cell populations: it was set to 40 px for MCCs and to 20 px for immune cells. For each cell population, an OAP using the Munkres algorithm³ is solved: the globally best possible pairing among located objects is based on an assignment cost equal to the inverse of the distance between pairs of cell candidates at the i^{th} and $(i + 1)^{th}$ frames. Cell appearing/disappearing and cell overlaps due to projection errors of the 3D scene in the 2D domain are also handled. Finally, cellular morphodynamic features and cell-cell interaction features are extracted at successive positions along each

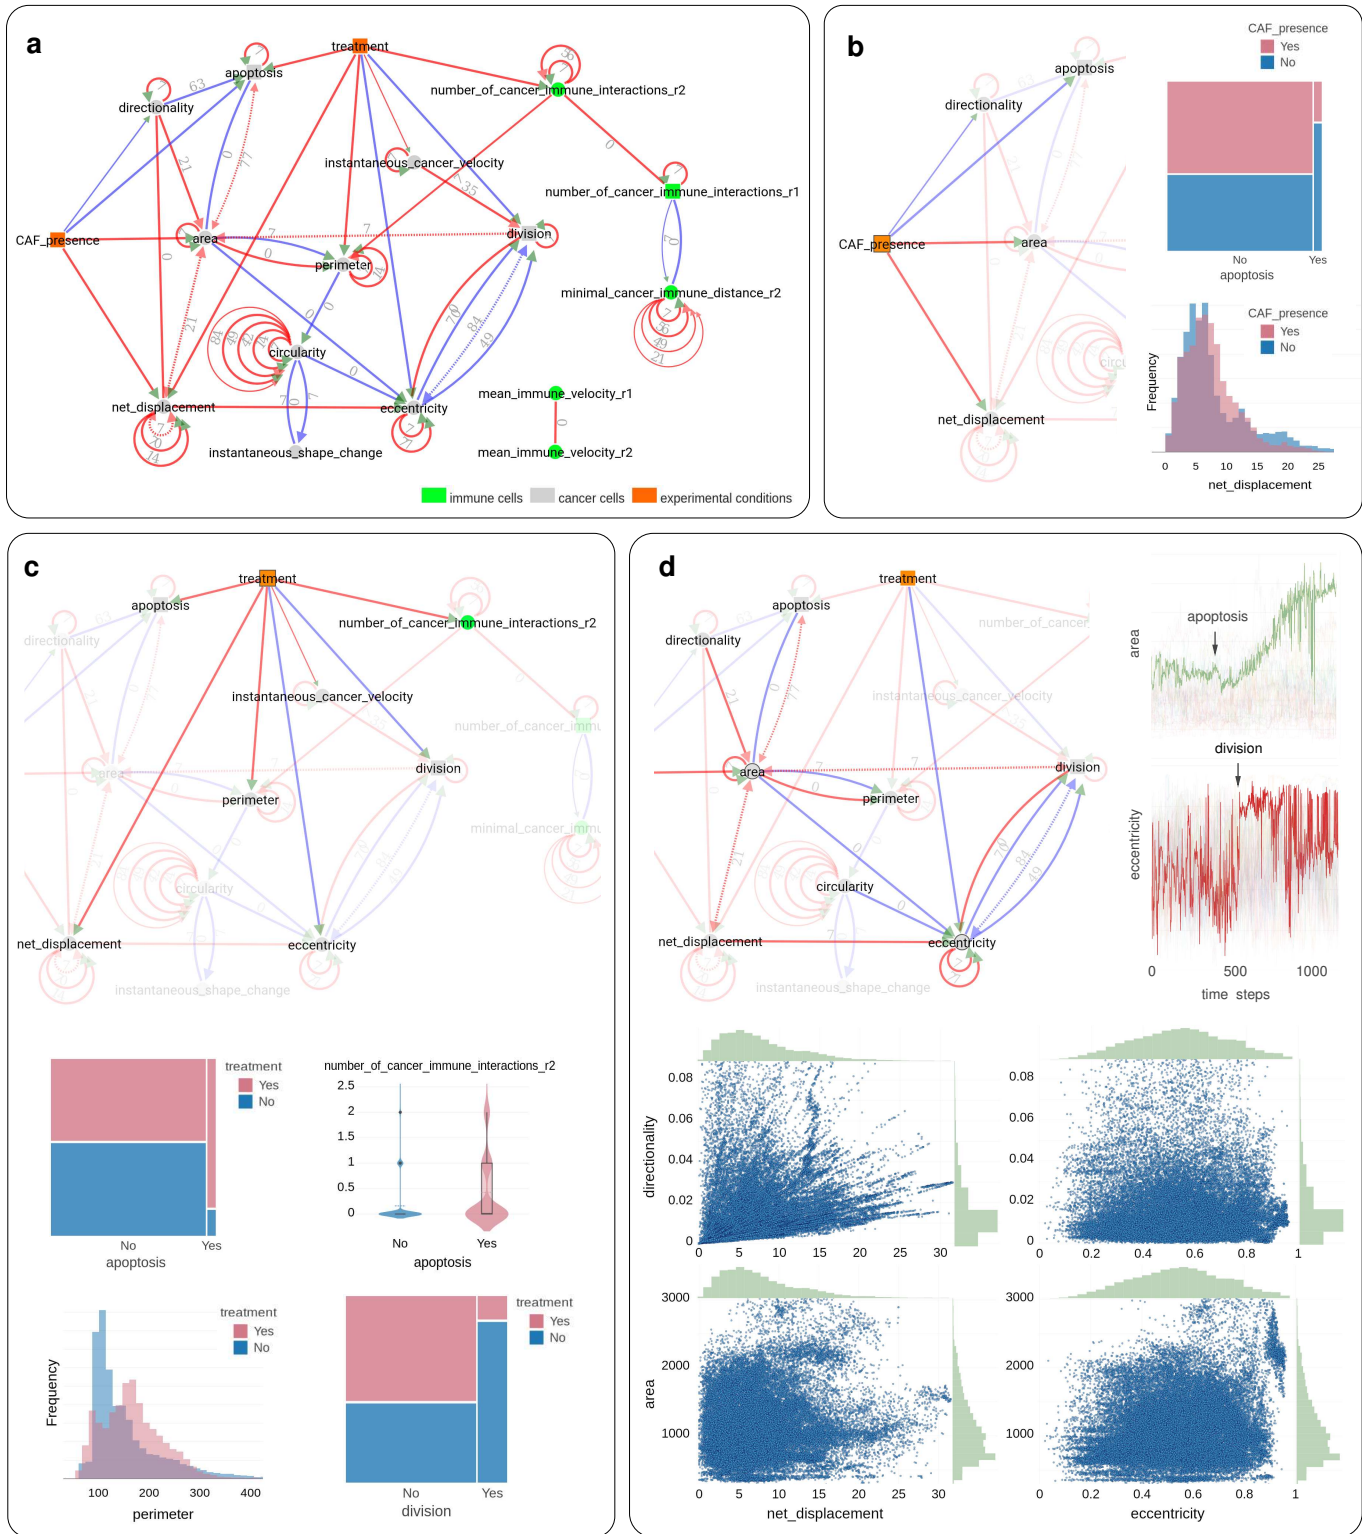


Figure 2: Application of CausalXtract to time-lapse images of tumor ecosystems reconstituted ex vivo¹. **a**, Summary causal network inferred by CausalXtract. The underlying time-unfolded causal network is shown on Supplementary Fig. 7. Red (resp. blue) edges correspond to positive (resp. negative) associations. Bidirected dashed edges represent the effect of unobserved (latent) common causes. Annotations on edges correspond to time delays in time-steps (1 ts = 2 min). The inferred network is largely robust to variations in sampling rate ($\delta\tau$) and maximum lag (τ), Supplementary Fig. 8. Here $\delta\tau = 7$ ts and $\tau = 84$ ts are chosen automatically by CausalXtract, Supplementary Fig. 8b. **b**, The CAF presence subnetwork highlighting the direct causal effects of CAFs on cancer cells. In particular, CausalXtract uncovers that CAFs directly inhibit cancer cell apoptosis independently from treatment, which has not been reported so far. **c**, The treatment subnetwork highlighting the direct causal effects of treatment on cancer cells. In particular, CausalXtract uncovers that treatment increases cancer cell perimeter, which has not been reported either. **d**, The eccentricity-area subnetwork highlighting multiple direct and possibly antagonistic time-lagged effects, notably, between cell division and eccentricity and between cell apoptosis and area, as discussed in main text.

trajectory. For each MCC, 15 descriptors were extracted, Supplementary Fig. 6, and classified into four main categories: cell shape, motility, state, and interaction descriptors.

Shape descriptors. The active contour algorithm implemented in Matlab¹⁸ was used to segment the MCC boundaries on each video crop frame. Taking as input a frame representing the i^{th} snapshot of the t^{th} MCC, it returns a binary image, where the MCC is represented by a white region. From the binary image, the shape properties of the region occupied by each MCC were extracted using the Matlab *regionprops* algorithm. The resulting descriptors of the extracted shape are listed below:

- *area* indicates the number of pixels composing the region. The equivalent diameter of the t^{th} MCC in the i^{th} frame is defined as $d_i^t = \sqrt{4 \cdot \text{area}}/\pi$.
- *perimeter* represents the distance along the MCC boundary.
- *circularity* is defined as $4 \cdot \text{area} \cdot \pi / \text{perimeter}^2$, which is equal to 1 when the region is perfectly circular.
- *eccentricity* denotes the eccentricity of the ellipse with the same second moments as the region. The value is equal to 1 when the region is a line and to 0 when the region is a circle.
- *instantaneous shape change* is defined as, $|d_i^t - d_{i-1}^t|$, corresponding to the difference in absolute value of the equivalent diameters between the i^{th} and $(i-1)^{\text{th}}$ frames of the t^{th} MCC.

Motility descriptors. The positions $p_i^t = (x_i^t, y_i^t)$ and p_{i-1}^t of the t^{th} MCC in the i^{th} and $(i-1)^{\text{th}}$ frames were compared using the Euclidean distance $d(\cdot)$ to define the following motility parameters:

- *instantaneous cancer velocity*¹⁹ is defined as $d(p_i^t, p_{i-1}^t)/\Delta t$, where Δt is the time interval between two consecutive frames.
- *net displacement*¹⁹ indicates the resultant distance between the initial and current positions of the t^{th} MCC, $d(p_i^t, p_1^t)$.
- *directionality*¹⁹ is defined as the ratio of net displacement, $d(p_i^t, p_1^t)$, and curvilinear distance, $\sum_{k=2}^i d(p_k^t, p_{k-1}^t)$. It measures the persistence of motion and ranges from 0 for confined cells to 1 for cells moving perfectly straight in one direction.

State descriptors. They record apoptosis or division events:

- *apoptosis* indicates if the MCC has died during the experiment. It is set to ‘No’ as long as the cell has not died and becomes ‘Yes’ for the remaining frames after the cell undergoes apoptosis.
- *division* indicates if the MCC has divided during the experiment. It is set to ‘No’ as long as the cell has not divided and becomes ‘Yes’ for the remaining frames after the cell divides.

Interaction descriptors. Interactions between MCCs and immune cells were defined with respect to two radii around each MCC, $r_1 = r_{\text{im}} + r_{\text{ca}} + 2 = 20 \text{ px}$ and $r_2 = 2 \times (r_{\text{im}} + r_{\text{ca}}) = 36 \text{ px}^1$. Hence, r_1 refers to MCC and immune cells in actual physical contact, while r_2 refers to MCC and immune cells in close vicinity. Then, for each sample the following interaction features were defined:

- *number of cancer-immune interactions* (r_2) corresponds to the number of immune cells within the interaction radius r_2 around the MCC on that frame.
- *number of cancer-immune interactions* (r_1) corresponds to the number of immune cells in close contact with the MCC on that frame.
- *minimal cancer-immune distance* (r_2) is the minimum distance between the MCC and the immune cells within a radius r_2 .
- *mean immune velocity* (r_2) is the mean instantaneous velocity norm of the immune cells within the interaction radius r_2 around the MCC.
- *mean immune velocity* (r_1) is the mean instantaneous velocity norm of the immune cells in close contact with the MCC.

Overview of causal discovery methods for non-temporal data

Traditional causal discovery methods^{20,21} aim to learn causal networks from datasets of independent samples by proceeding through successive steps. They first learn structural constraints in the form of unconditional or conditional independence between variables and remove the corresponding edges from an initial fully connected network. The second step then consists in orienting some of the retained edges based on the signature of causality in observational data. This corresponds to orienting three-variable “v-structure” motifs as, $X \rightarrow Z \leftarrow Y$, whenever the edge $X - Y$ has been removed without conditioning on the variable Z , which implies that Z cannot be a cause of X nor Y . This does not guarantee, however, that X (or Y) is an actual cause of Z , which also requires to rule out the possibility that the edge between X and Z (or Y and Z) might originate from a latent common cause, L , unobserved in the dataset, *i.e.* $X \leftarrow L \rightarrow Z$. In addition, classical causal discovery methods are prone to spurious conditional independences, which lead to many false negative edges and limit the accuracy of inferred orientations. The recent causal discovery method, MIIC⁴⁻⁶, which combines constraint-based and information-based principles, learns more robust causal graphical models by first collecting iteratively significant information contributors before assessing conditional independences. In practice, MIIC’s strategy limits spurious conditional independences which improves its edge sensitivity and orientation reliability compared to traditional constraint-based methods⁴⁻⁶. In addition, MIIC can handle missing data⁵ and also heterogeneous multimodal data, by analyzing continuous and categorical variables on the same footing, based on a mutual information supremum principle for finite dataset^{5,6}. Last, MIIC distinguishes genuine causal relations from putative and latent causal effects⁶, that are ubiquitous in real-world applications.

CausalXtract’s causal discovery module for time series data

In order to analyze time series datasets, CausalXtract’s causal discovery module (tMIIC) aims to learn a time-unfolded graph, \mathcal{G}_t , where each variable is represented by a series of nodes associated to its value at different relative time points, Fig. 1c. Such a time-unfolded network framework¹⁰⁻¹³ is required to account for the temporal correlation between successive samples in time series data. Assuming that the dynamics can be considered stationary (see Benchmarking of CausalXtract’s causal discovery module section, below), the time-unfolded graph, \mathcal{G}_t , should be translationally invariant over time and can be assigned a periodic structure *a priori*. In addition, \mathcal{G}_t can be restricted to a few time steps from the running time, t , back to a maximum time lag, $t - \tau$, since nodes at future time points ($t' > t$) cannot *a priori* influence the observed data at current or previous time points ($t' \leq t$), Fig. 1c. The maximum time lag τ should be chosen so as to have little effect on the final graphical model, which can be achieved for instance by setting τ to twice the average relaxation time of the variables of the dataset. In practice, we may also limit the number of time points ν in \mathcal{G}_t by introducing a time increment $\delta\tau$ between consecutive time points, which leads to $\nu = \tau/\delta\tau$ time-lagged layers in \mathcal{G}_t .

Such a compact periodic graphical representation over a sliding temporal window is learned with tMIIC, which extends MIIC causal discovery method to analyze time series data. First, tMIIC identifies all necessary edges involving at least one contemporaneous node at time t , Fig. 1c. Once these time-lagged and contemporaneous necessary edges have been identified, they are simply duplicated at earlier time points to enforce the translational invariance of \mathcal{G}_t skeleton. Time-lagged edges are then pre-oriented with a first arrowhead pointing towards the future, considering that current time points cannot cause earlier events. Then, contemporaneous and time-lagged edges can be further oriented using MIIC orientation probability scores applied to \mathcal{G}_t , which may also uncover a second arrowhead (backward in time) for time-lagged edges. This corresponds to time-lagged latent causal effects from unobserved common causes, Fig. 1c.

Learning such structural models including latent variables from

time series data was first proposed for time-lagged effects¹⁰ and subsequently extended to contemporaneous effects¹¹ by adapting the constraint-based FCI method allowing for latent variables²¹. While traditional constraint-based methods suffer from poor recall, the recent PCMCI¹² / PCMCI+²² method improves recall by introducing ad hoc conditioning rules for auto-correlated time series. By contrast, tMIIC does not require any ad hoc conditioning rules, as it relies on the same robust information-theoretic strategy as MIIC to limit spurious independence and improve edge recall. tMIIC also captures time-lagged and contemporaneous effects due to latent variables.

Relation to Granger-Schreiber temporal causality

The concept of temporal causality was originally formulated by Granger¹⁴ without reference to any graphical model by comparing linear autoregression with or without past values of possible causal variables. This was later extended to non-linear relations by Schreiber^{15,16} using the notion of Transfer Entropy, $T_{X \rightarrow Y}$, which can be expressed in terms of multivariate conditional information,

$$T_{X \rightarrow Y} = I(Y_t; \mathbf{X}_{t' < t} | \mathbf{Y}_{t' < t}) \quad (1)$$

where $\mathbf{X}_{t' < t}$ and $\mathbf{Y}_{t' < t}$ denote the sets of variables, $X_{t'}$ and $Y_{t'}$, taken at earlier time points t' than t .

While Eq. 1 is asymmetric upon X/Y permutation, a simple comparison of Transfer Entropy asymmetry (*e.g.* $T_{X \rightarrow Y} > T_{Y \rightarrow X} \geq 0$) does not necessarily translate into causal direction as this asymmetry is also expected for non-causal relations. Interestingly, this is in fact the absence of Transfer Entropy in one direction (*e.g.* $T_{Z \rightarrow X} \approx 0$) which suggests the possibility of a causal relation in the opposite direction, $X \rightarrow Z$, as in the case of v-structures in graph-based causal discovery methods, provided that a latent common cause can be excluded between the two variables (as discussed above).

We clarify in Theorem 1 below this relation between temporal causality without reference to any structural model (Eq. 1) and structural causality entailed by time-unfolded causal graphical models (\mathcal{G}_t). This highlights the common foundations of temporal and structural causalities beyond their seemingly unrelated definitions.

Theorem 1. [$T_{Y \rightarrow X} = 0$ implies temporal ($2 \text{ var} + t$) v-structures] If X_t is adjacent to Y_t in \mathcal{G}_t and $T_{Y \rightarrow X} = I(X_t; \mathbf{Y}_{t' < t} | \mathbf{X}_{t' < t}) = 0$, then for all $Y_{t'}$ adjacent to Y_t in \mathcal{G}_t , with $t' < t$, there is a temporal ($2 \text{ var} + t$) v-structure, $Y_{t'} \rightarrow Y_t \leftarrow X_t$, in \mathcal{G}_t , Supplementary Fig. 5a.

Proof: if $T_{Y \rightarrow X} = I(X_t; \mathbf{Y}_{t' < t} | \mathbf{X}_{t' < t}) = 0$, then all pairs ($X_t, Y_{t'}$) should be unconnected (assuming ‘faithfulness’, *i.e.* no coincidental cancellation of effects) and all unshielded triples $Y_{t'} - Y_t - X_t$ should be temporal v-structures, $Y_{t'} \rightarrow Y_t \leftarrow X_t$, as $Y_t \notin \mathbf{X}_{t' < t}$ in $T_{Y \rightarrow X}$ \square

Note, however, that the converse of Theorem 1 is not true: a temporal v-structure does not imply a vanishing Transfer Entropy, as shown with the counterexample in Supplementary Fig. 5b. As a result, the presence of a temporal v-structure, $Y_{t'} \rightarrow Y_t \leftarrow X_t$ in \mathcal{G}_t , does not necessarily imply a vanishing transfer entropy, $T_{Y \rightarrow X} = 0$, as long as there remains an edge between any $Y_{t''}$ and X_t , as in the example in Supplementary Fig. 5b. Hence, Granger-Schreiber causality is in fact too restrictive and may miss actual causal effects, which can be uncovered by structural causal discovery methods like tMIIC. In addition, Granger-Schreiber causality is also known to infer spurious causal associations by excluding the presence of latent common causes *a priori*. By contrast, CausalXtract’s causal discovery module includes time-delayed as well as synchronous effects originating from unobserved latent variables, as discussed above.

Benchmarking of CausalXtract’s causal discovery module

The performance of CausalXtract’s causal discovery module (tMIIC) has been assessed using Tigramite package²², which provides different methods to learn temporal causal networks from time series data. We compared tMIIC to two methods capable of orienting contemporaneous edges (PC and PCMCI+) and tested three different kernels for estimating mutual information (Parcorr, GPDC and KNN).

Benchmark networks and datasets have been chosen to resemble the real-world data analyzed in this study (*i.e.* similar number of time steps, network size and degree distribution) and include a large range of linear and non-linear relations between variables.

A first series of datasets was generated for a 15 node benchmark network (Supplementary Fig. 1a) with linear combinations of contributions inspired by the Tigramite package, Supplementary Table 1. Running times and scores (Precision, Recall, F-score) have been averaged over 10 datasets (Supplementary Fig. 1b) and show that tMIIC scores are at par with PC and PCMCI+ using GPDC or KNN kernels but that tMIIC runs orders of magnitude faster, which enables to use tMIIC on much larger datasets of biological interest including a few tens or hundreds of thousands samples. Only PC or PCMCI+ using ParCorr kernel match tMIIC running speed but with significantly lower scores, as Fscores level off around 0.6-0.7 at large sample size, while tMIIC Fscore exceeds 0.9 (Supplementary Fig. 1b).

Importantly, increasing the number of time-lagged layers from $\tau = 2$ (as in the actual model, Supplementary Fig. 1a) to 5 or 10 layers in the inferred time-unfolded network (Supplementary Fig. 2) leads to very similar network reconstructions for simulated stationary data. This demonstrates tMIIC insensitivity to an overestimated maximum lag for the reconstituted network. Interestingly, however, when the generated data is no longer stationary, increasing the number of layers leads to multiple self-loops at non-stationary variables, whilst the rest of the network remains relatively unaffected (Supplementary Fig. 3). It demonstrates that CausalXtract’s causal discovery module is robust to the presence of non-stationary variables but requires long-time range interactions, and therefore multiple time-lagged layers, to account for these non-stationary dynamics at specific variables. This striking observation on benchmark networks is also consistent with the multiple self-loops observed for a number of non-stationary variables in the real-world application on cellular ecosystems, Fig. 2a and Supplementary Fig. 6.

A second series of more complex datasets was also generated for another 15 node benchmark network (Supplementary Fig. 4a) with non-linear combinations of contributors, Supplementary Table 2. Here, tMIIC tends to outperform both PC and PCMCI+, in terms of Recall and Fscores, while remaining orders of magnitude faster compared to GPDC and KNN kernels. Only PC or PCMCI+ using ParCorr kernel match tMIIC running speed but with significantly lower scores (*i.e.* Fscores level off around 0.4-0.5 at large sample size, while tMIIC Fscore exceeds 0.8). This demonstrates that CausalXtract’s causal discovery module (tMIIC) is both a reliable and scalable method to discover complex temporal causal relations in very large time series datasets including a few hundred thousand samples.

Data availability

The original live-cell time-lapse image data and extracted crops are available at: <https://doi.org/10.5281/zenodo.7755699>.

Code availability

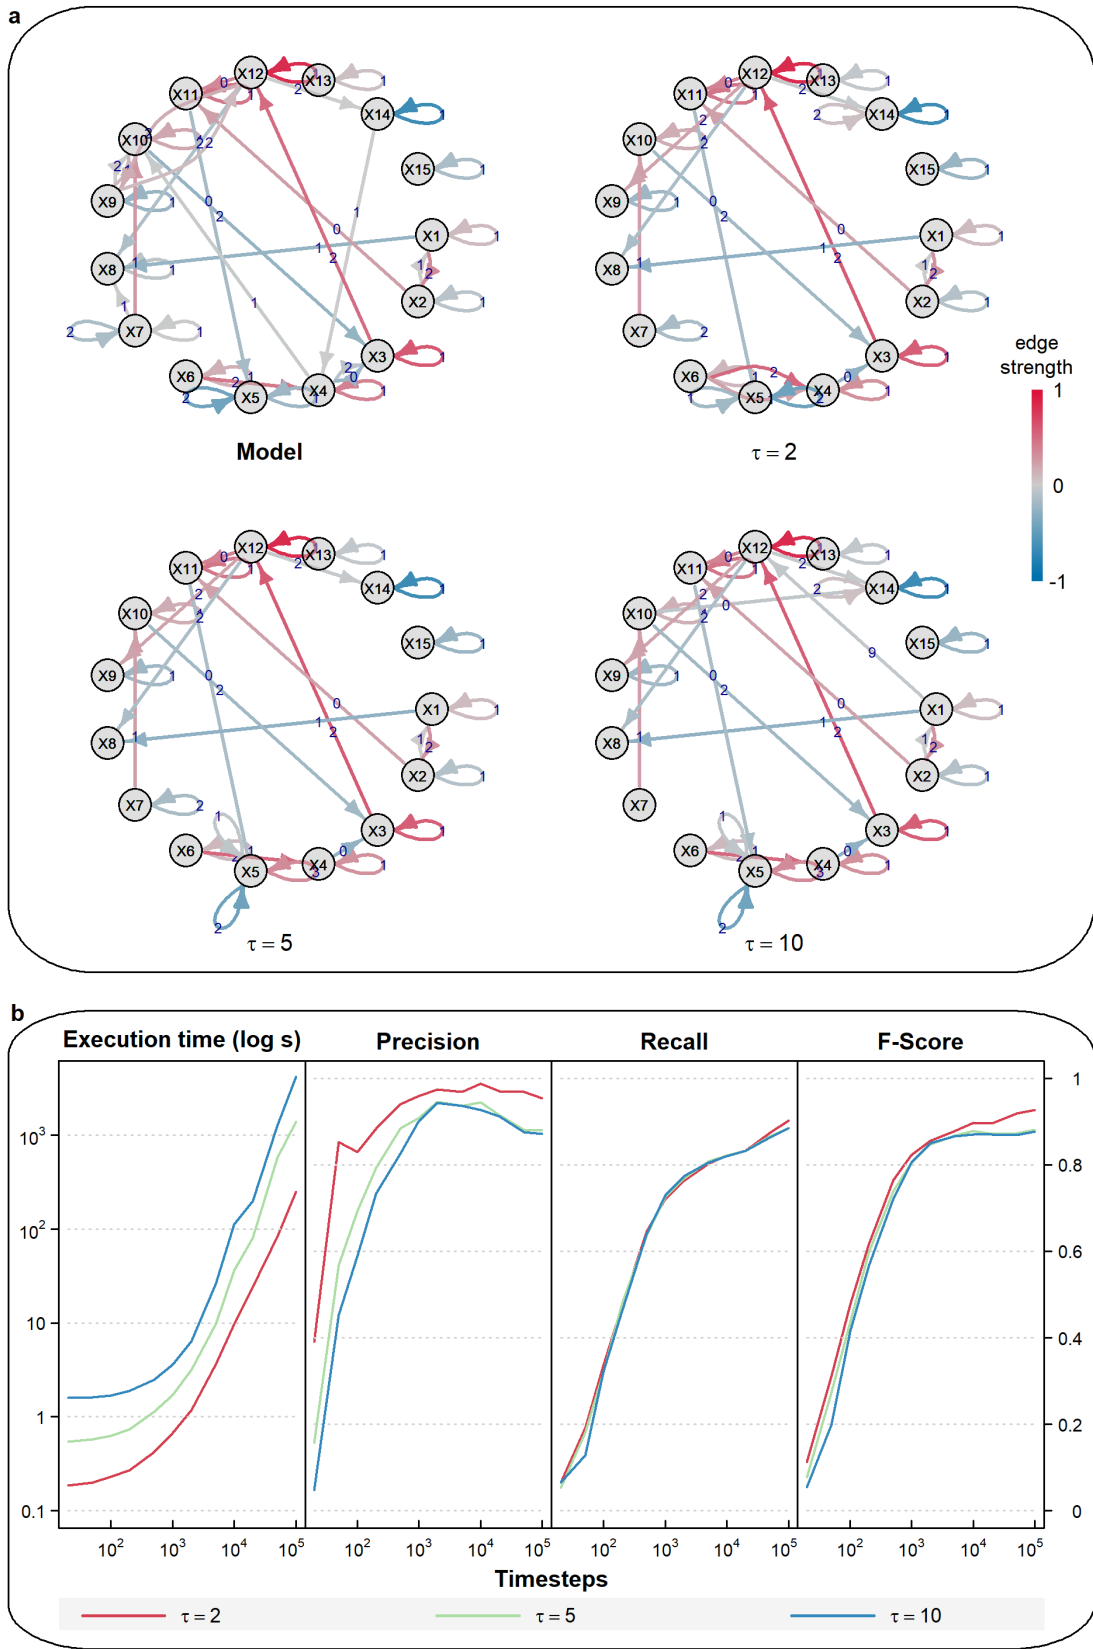
The source code of the CausalXtract pipeline is available at: <https://github.com/miicTeam/CausalXtract>. It includes a demo R markdown notebook of CausalXtract pipeline, which reproduces step-by-step the results reported in the manuscript, Fig. 2, starting from the original live-cell time-lapse images of the tumor-on-chip ecosystem, Fig. 1a. Tigramite package used for benchmark comparison is available at: <https://github.com/jakobrunge/tigramite>

Acknowledgements

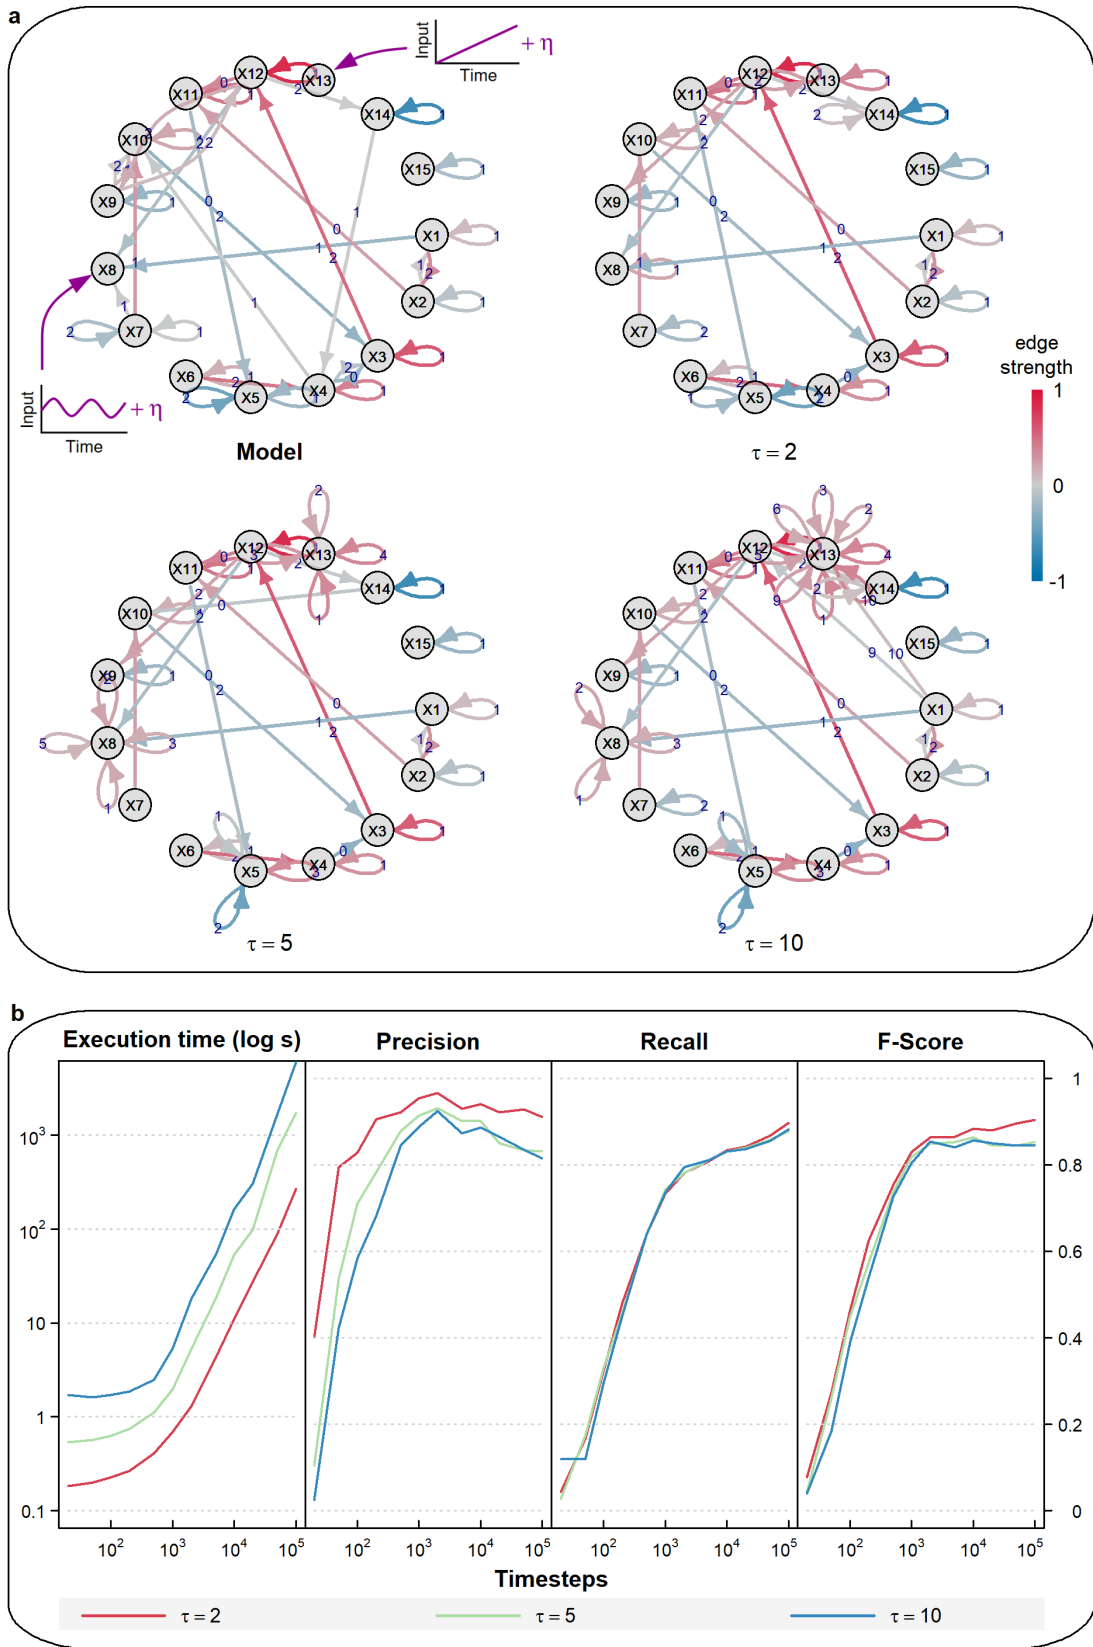
This work was supported by ITMO Cancer (grant No 20CM106) and the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 847718. LD acknowledges support from AMX PhD fellowship, VC from ARC foundation and NL from CNRS-Imperial College joint PhD programme.

References

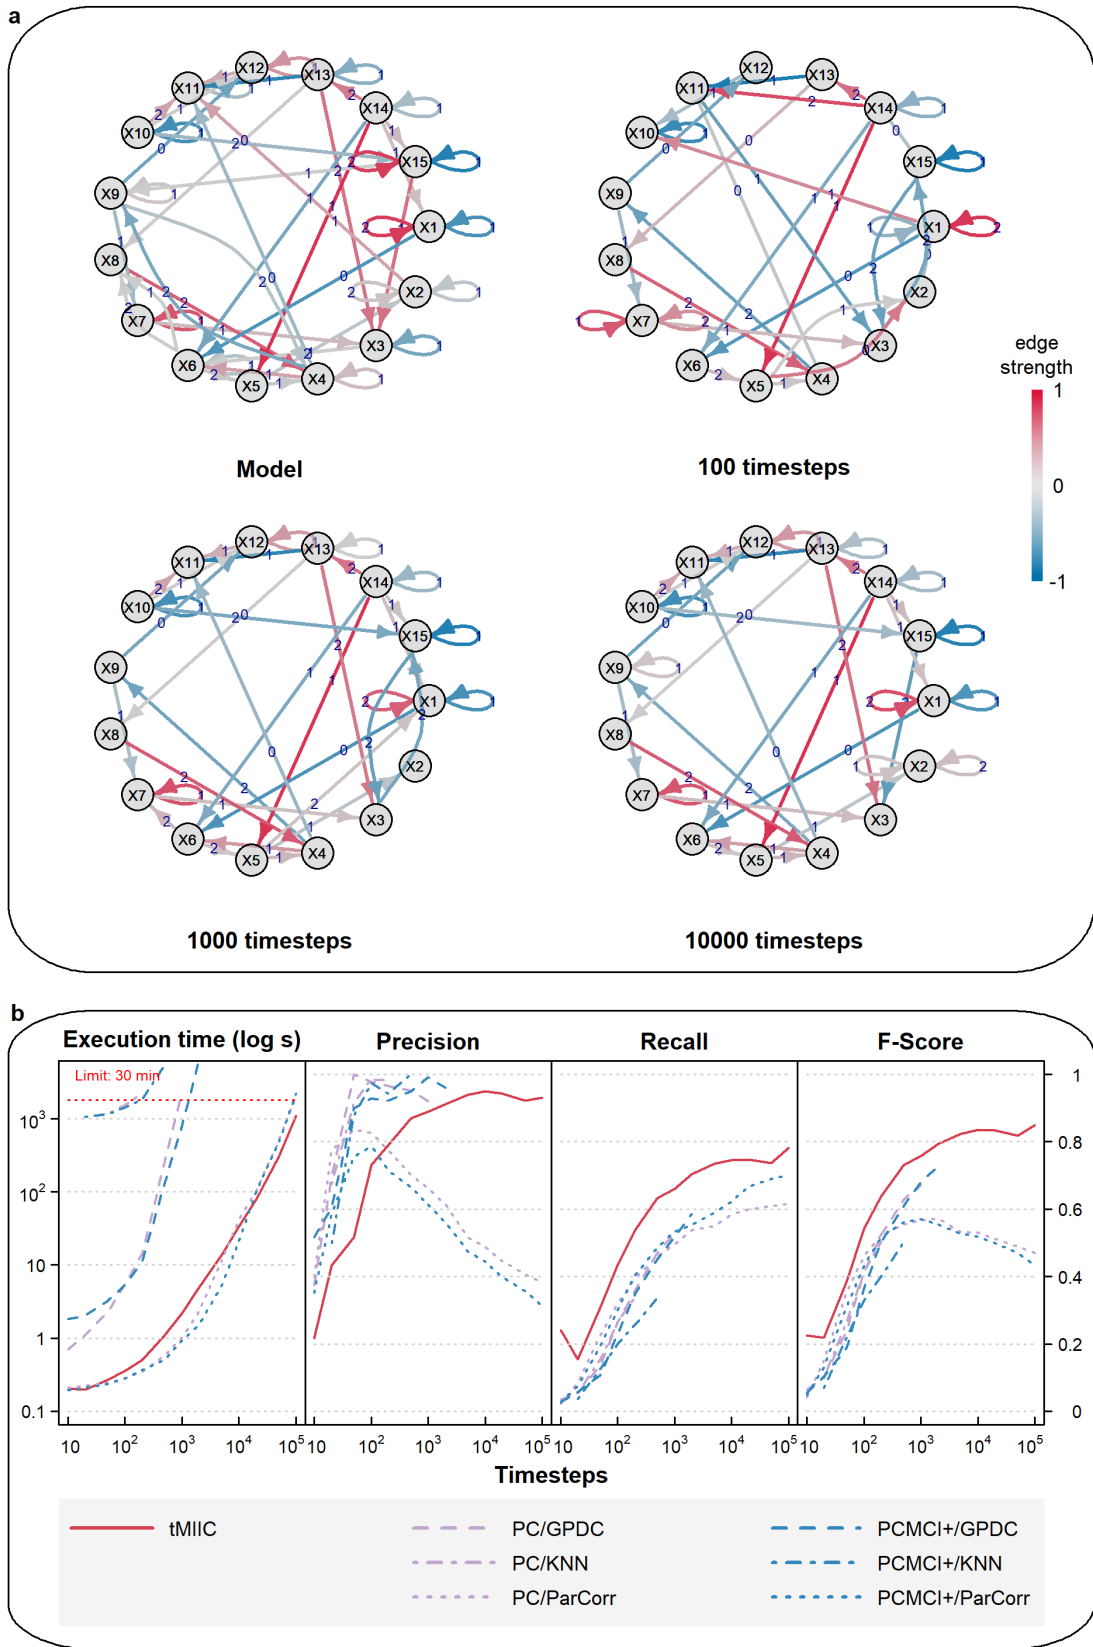
1. Nguyen, M. *et al.* Dissecting Effects of Anti-cancer Drugs and Cancer-Associated Fibroblasts by On-Chip Reconstitution of Immunocompetent Tumor Microenvironments. *Cell Reports* **25**, 3884–3893.e3 (2018).
2. Davies, E. *Machine vision* 3rd ed. (Morgan Kaufmann, 2004).
3. Munkres, J. Algorithms for the assignment and transportation problems. *J. Soc. Ind. Appl. Math.* **5**, 32–38 (1957).
4. Verny, L., Sella, N., Affeldt, S., Singh, P. P. & Isambert, H. Learning causal networks with latent variables from multivariate information in genomic data. *PLoS Comput. Biol.* **13**, e1005662 (2017).
5. Cabeli, V. *et al.* Learning clinical networks from medical records based on information estimates in mixed-type data. *PLoS Comput. Biol.* **16**, e1007866 (2020).
6. Da Câmara Ribeiro-Dantas, M. *et al.* Learning interpretable causal networks from very large datasets, application to 400,000 medical records of breast cancer patients. Preprint at <https://doi.org/10.48550/arXiv.2303.06423> (2023).
7. Desterke, C. *et al.* Inferring Gene Networks in Bone Marrow Hematopoietic Stem Cell-Supporting Stromal Niche Populations. *iScience* **23**, 101222 (2020).
8. Sella, N. *et al.* Interactive exploration of a global clinical network from a large breast cancer cohort. *npj Digital Med* **5**, 113 (2022).
9. Assaad, C., Devijver, E. & Gaussier, E. Survey and Evaluation of Causal Discovery Methods for Time Series. *Journal of Artificial Intelligence Research* **73**, 767–819 (2022).
10. Entner, D. & Hoyer, P. On Causal Discovery from Time Series Data using FCI. *Proceedings of the 5th European Workshop on Probabilistic Graphical Models, PGM 2010* (2010).
11. Malinsky, D. & Spirtes, P. *Causal Structure Learning from Multivariate Time Series in Settings with Unmeasured Confounding in Proceedings of 2018 ACM SIGKDD Workshop on Causal Discovery, CD@KDD 2018* (eds Le, T. D., Zhang, K., Kiciman, E., Hyvärinen, A. & Liu, L.) **92** (2018), 23–47.
12. Runge, J., Nowack, P., Kretschmer, M., Flaxman, S. & Sejdinovic, D. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances* **5** (2019).
13. Runge, J. *et al.* Inferring causation from time series in Earth system sciences. *Nat. Commun.* **10**, 2553 (2019).
14. Granger, C. W. J. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* **37**, 424 (1969).
15. Schreiber, T. Measuring Information Transfer. *Physical Review Letters* **85**, 461–464 (2000).
16. Barnett, L., Barrett, A. B. & Seth, A. K. Granger Causality and Transfer Entropy Are Equivalent for Gaussian Variables. *Physical Review Letters* **103** (2009).
17. Driscoll, M. K. & Zaritsky, A. Data science in cell imaging. *Journal of Cell Science* **134**, jcs254292 (2021).
18. Chan, T. F. & Vese, L. A. Active contours without edges. *IEEE Trans. Image Process.* **10**, 266–277 (2001).
19. Masuzzo, P., Van Troys, M., Ampe, C. & Martens, L. Taking aim at moving targets in computational cell migration. *Trends Cell Biol.* **26**, 88–110 (2016).
20. Pearl, J. *Causality* (Cambridge university press, 2009).
21. Spirtes, P., Glymour, C. N., Scheines, R. & Heckerman, D. *Causation, prediction, and search* (MIT press, 2000).
22. Runge, J. *Discovering contemporaneous and lagged causal relations in autocorrelated nonlinear time series datasets in Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence* (eds Peters, J. & Sontag, D.) **124** (2020), 1388–1397.



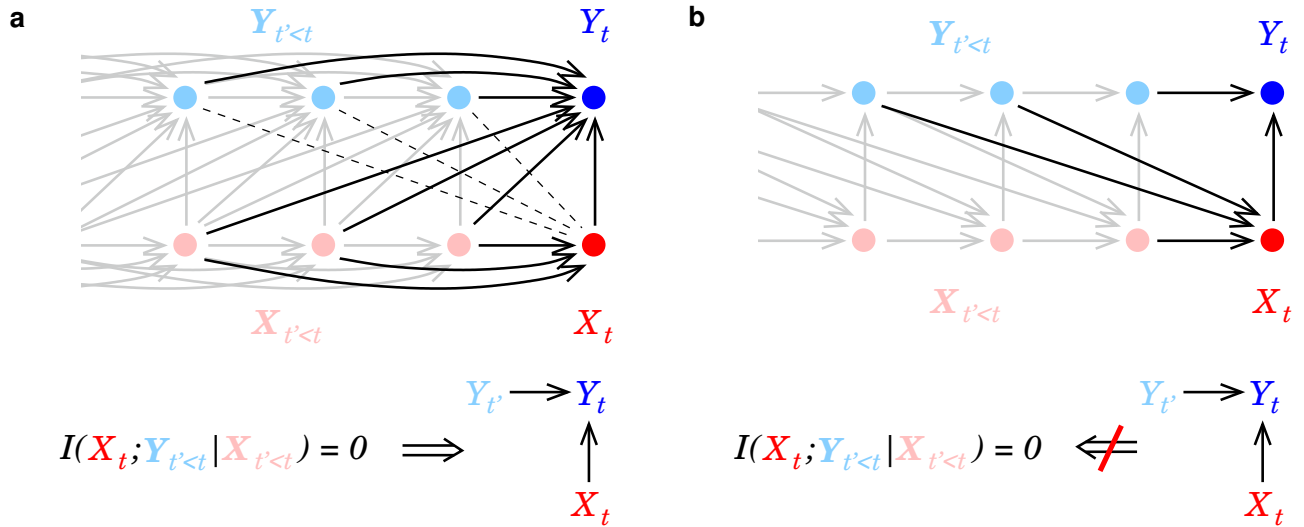
Supplementary Fig. 2: **CausalXtract insensitivity to an overestimated maximum lag τ .** **a**, Example of a temporal causal network model with a maximum lag $\tau = 2$. Corresponding temporal causal networks inferred by CausalXtract's causal discovery module (tMIIC), from 1,000 time step stationary time series (Supplementary Table 1), while assuming different maximum lags $\tau = 2, 5$ or 10. **b**, Running times and scores (Precision, Recall, Fscore) of tMIIC temporal causal network reconstructions for $\tau = 2, 5$ or 10, averaged over ten stationary time series of 10 to 10⁵ time steps. Overestimating the maximum lag τ has little impact on the reconstructed networks, as long as the time series are stationary, as demonstrated in Supplementary Fig. 3.



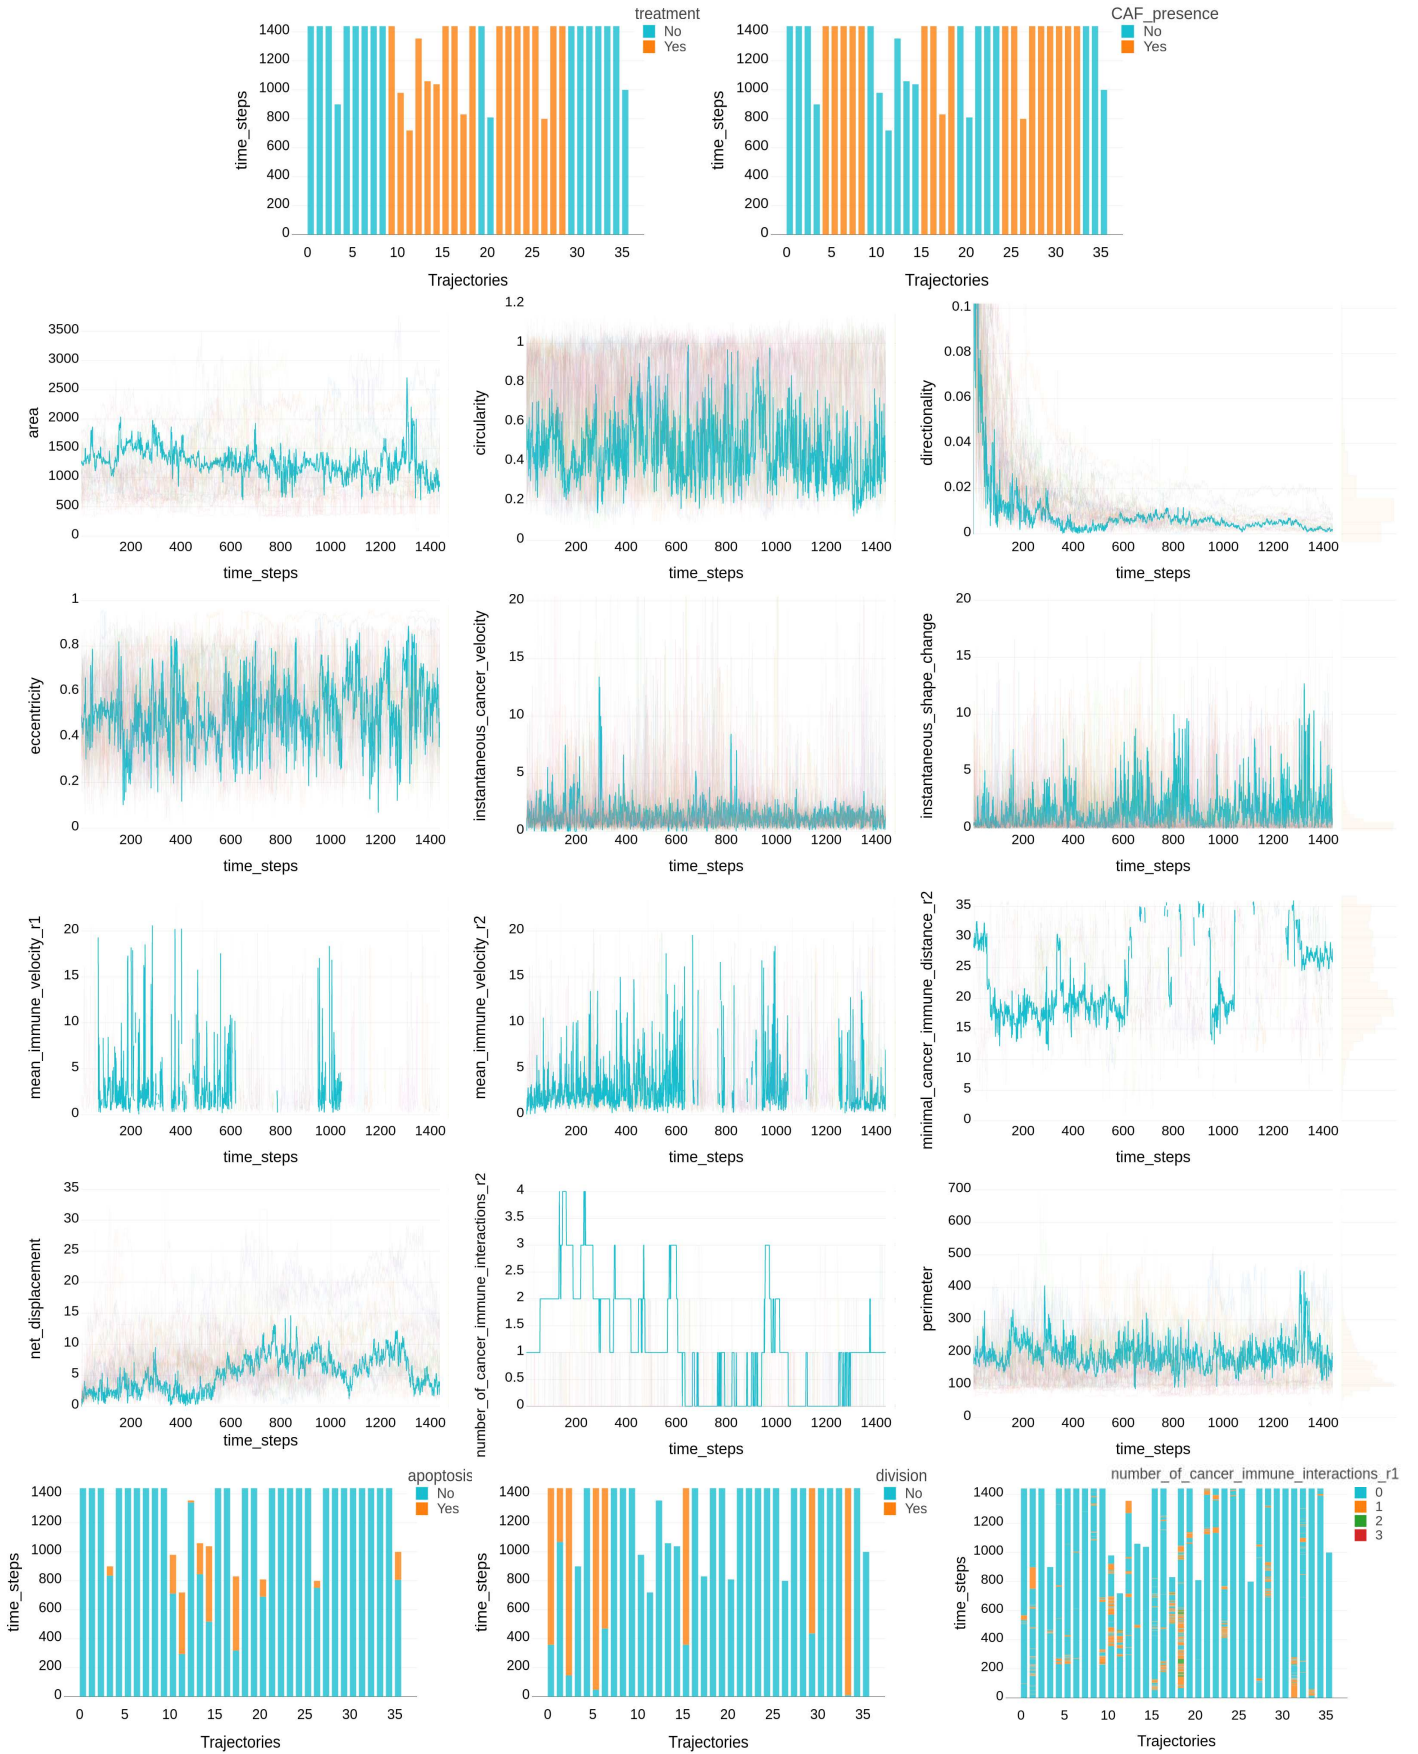
Supplementary Fig. 3: **CausalXtract sensitivity to non-stationary variables.** **a**, Example of a temporal causal network model ($\tau = 2$) with a low frequency periodic input ($T = 100$) applied to X8 and a time-linear trend applied to X13. Corresponding temporal causal networks inferred by tMIIC from 1,000 time step time series (Supplementary Table 1) including non-stationary inputs to X8 and X13. Increasing the maximum lag from $\tau = 2$ to $\tau = 5$ or 10 leads to the appearance of multiple self-loops, which result from the non-stationary dynamics of X8 and X13, whilst the rest of the network remains largely unaffected. **b**, Running times and scores (Precision, Recall, Fscore ignoring X8 and X13 self-loops) of tMIIC causal network reconstructions for $\tau = 2, 5$ or 10, averaged over ten time series of 10 to 10^5 time steps.





Supplementary Fig. 4: **Benchmark assessment of CausalXtract's causal discovery module (tMIIC) using more complex time series datasets.** **a**, Example of a 15 node causal network to generate more complex benchmark time series datasets based on non-linear combinations of contributions, Supplementary Table 2. Examples of temporal causal networks reconstructed by tMIIC based on 100, 1,000 or 10,000 simulated time steps. **b**, Running times and scores (Precision, Recall, Fscore) averaged over 10 datasets and compared to PC and PCMCI+ methods using different kernels (GPDC, KNN, ParCorr); tMIIC outperforms both PC and PCMCI+, in terms of Recall and Fscores, while running orders of magnitude faster, except for the ParCorr kernel, which leads, however, to significantly lower scores at large sample size.



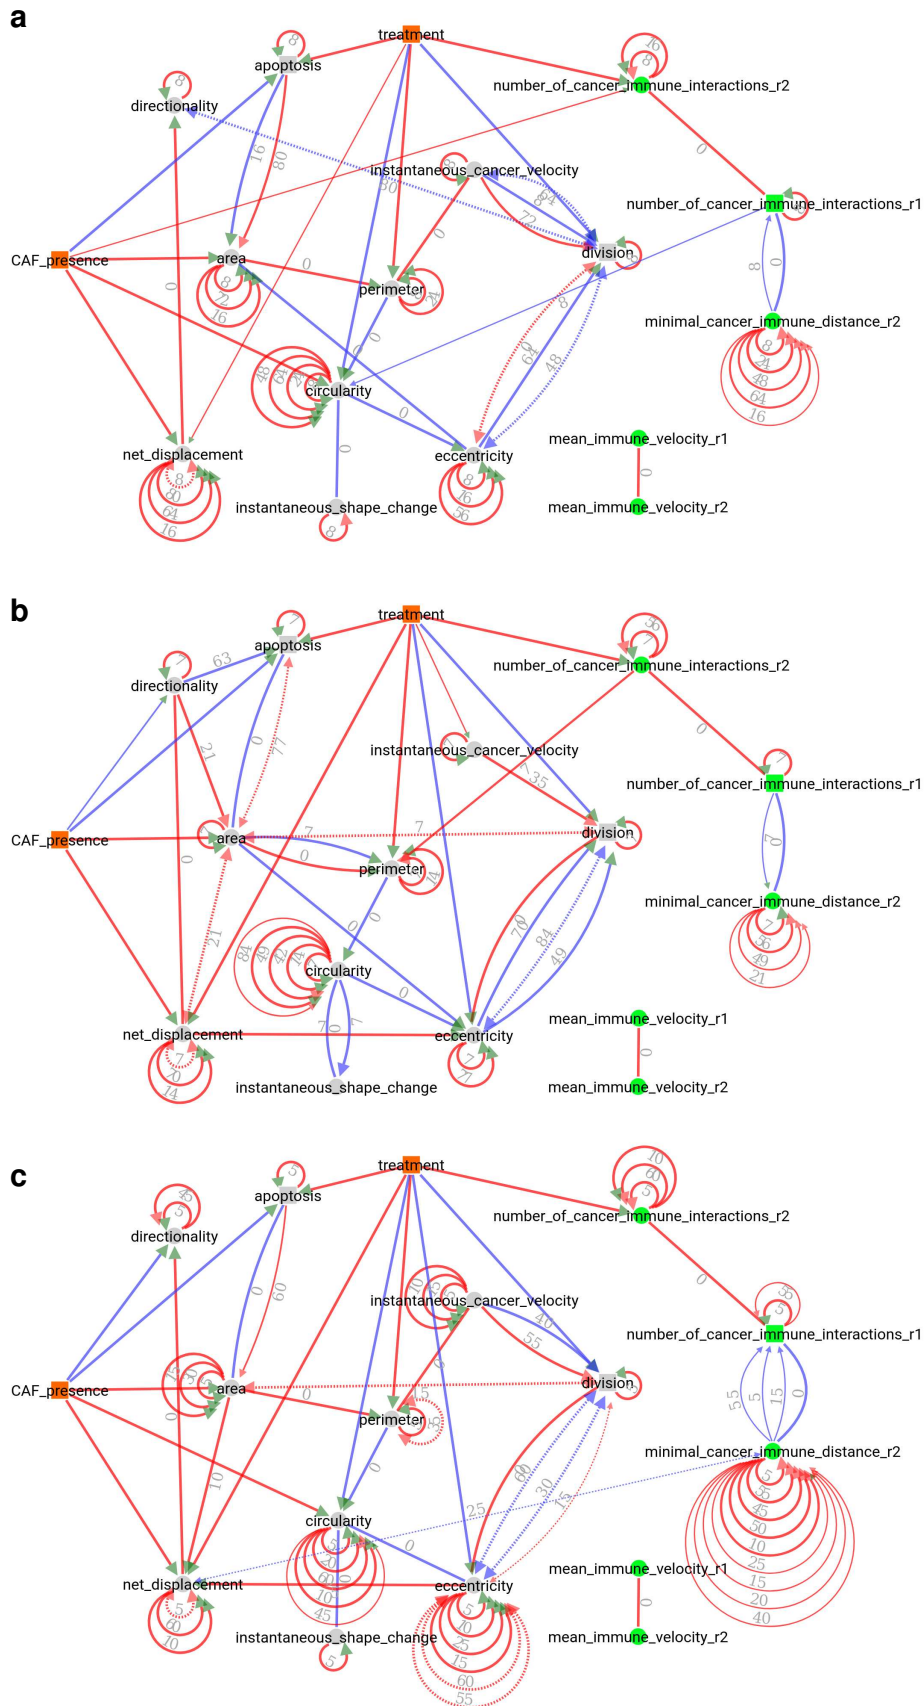
Supplementary Fig. 5: **Time-unfolded causal network framework and relation to Granger-Schreiber temporal causality.** **a**, A vanishing Transfer Entropy, *i.e.* $T_{Y \rightarrow X} = I(X_t; Y_{t' < t} | X_{t' < t}) = 0$, implies *i*) the absence of (dashed) edge between X_t and any $Y_{t'}$, with $t' < t$, and *ii*) if X_t is adjacent to Y_t , the presence of temporal (2-variable + time) v-structures, $Y_{t'} \rightarrow Y_t \leftarrow X_t$, for all $Y_{t'}$ adjacent to Y_t , with $t' < t$ (Methods, Theorem 1). These results can be readily extended to include the presence of other observed variables, $V_{t' \leq t}$, by redefining Transfer Entropy as, $T_{Y \rightarrow X} = I(X_t; Y_{t' < t} | X_{t' < t}, V_{t' \leq t})$, which discards contributions from indirect paths through other observed variables, $V_{t' \leq t}$. **b**, By contrast, the presence of a temporal (2-variable + time) v-structure, $Y_{t'} \rightarrow Y_t \leftarrow X_t$ does not imply a vanishing Transfer Entropy, as long as there remains an edge between any $Y_{t'' < t}$ and X_t . It implies that Granger-Schreiber temporal causality is in fact too restrictive and may overlook actual causal effects, which can be uncovered by graph-based causal discovery methods like CausalXtract's causal discovery module (tMIIC). Hence, CausalXtract's time-unfolded network framework, combining graph-based and information-based approaches, sheds light on the common foundations of the seemingly unrelated graph-based causality and Granger-Schreiber temporal causality, while clarifying their actual differences and limitations.



Supplementary Fig. 6: **Time series of cellular features extracted from the tumor ecosystems.** Example of time series of cellular features extracted by CausalXtract’s feature extraction module (CellHunter+) from the tumor ecosystems analyzed in this study, Fig. 1a. It includes two experimental control parameters (*i.e.* treatment and CAF presence) and 15 cellular features extracted every 2 minutes over a period of two days. Continuous features are highlighted for one trajectory (traj.18), while categorical features are shown for all trajectories.

a**b**

Supplementary Fig. 7: **Time-unfolded causal network inferred by CausalXtract.** **a**, Time-unfolded causal network assuming stationary dynamics of cellular ecosystems implying translational time invariance of the inferred causal network. **b**, Only edges involving at least one contemporaneous variables (*i.e.* at time t) need to be tested for conditional independence by tMIIC and the remaining edges are then duplicated at all previous time steps before assigning orientations when time-lagged latent variables are taken into account, Fig. 1c. Variables retaining multiple self-loops with different time-delays correspond to non-stationary variables in Supplementary Fig. 6, in agreement with benchmarks from simulated data including non-stationary variables, Supplementary Fig. 3.



Supplementary Fig. 8: **Robustness of CausalXtract’s temporal causal networks to variations in sampling rate.** Summary causal networks inferred by CausalXtract using different sampling rates ($\delta\tau$). **a**, $\delta\tau = 8$ ts and $\tau = 80$ ts, in time step units (1 ts = 2 min). **b**, $\delta\tau = 7$ ts, and $\tau = 84$ ts, as chosen automatically by CausalXtract based on the average relaxation time across the 15 monitored variables, $\tau_R = 40$ ts, which defines a maximum lag $\tau = 2\tau_R = 80$ ts. Given a total number of (time-lagged and -unlagged) nodes, chosen to be around 200 nodes for computational efficiency, it leads to 13 temporal layers ($\nu + 1 = 200/15 \simeq 13$) and a lag increment $\delta\tau = \tau/\nu \simeq 7$ ts. This summary causal network corresponds to Fig. 2a. **c**, $\delta\tau = 5$ ts and $\tau = 60$ ts, corresponding to $\tau = \nu \cdot \delta\tau$ with $\nu + 1 = 13$ temporal layers, as in (b).

Supplementary Table 1: 15 nodes model.

Nodes

$$\begin{aligned}
 X_t^1 &\leftarrow -0.47 f_2(X_{t-1}^1) + 0.29 f_3(X_{t-1}^2) \times \eta_1 \\
 X_t^2 &\leftarrow 0.49 f_2(X_{t-1}^2) + 0.4 f_1(X_{t-2}^1) + \eta_2 \\
 X_t^3 &\leftarrow 0.56 f_1(X_{t-1}^3) + 0.44 f_4(X_{t-2}^4) - 0.26 f_2(X_{t-2}^{10}) + 0.56 f_2(X_t^4) + \eta_3 \\
 X_t^4 &\leftarrow 0.24 f_3(X_{t-1}^4) - 0.24 f_2(X_{t-2}^6) - 0.12 f_4(X_{t-1}^{14}) \times \eta_4 \\
 X_t^5 &\leftarrow -0.39 f_3(X_{t-1}^5) - 0.42 f_3(X_{t-2}^5) - 0.39 f_3(X_t^{11}) + \eta_5 \\
 X_t^6 &\leftarrow -0.32 f_2(X_{t-1}^6) + \eta_6 \\
 X_t^7 &\leftarrow -0.17 f_4(X_{t-1}^7) - 0.17 f_1(X_{t-2}^7) + \eta_7 \\
 X_t^8 &\leftarrow 0.39 f_4(X_{t-1}^8) - 0.46 f_4(X_{t-1}^7) - 0.39 f_3(X_{t-1}^1) - 0.4 f_3(X_{t-2}^{12}) + \eta_8 \\
 X_t^9 &\leftarrow -0.34 f_1(X_{t-1}^9) + 0.43 f_3(X_{t-2}^{12}) + \eta_9 \\
 X_t^{10} &\leftarrow 0.2 f_1(X_{t-1}^{10}) + 0.18 f_4(X_{t-2}^9) + 0.17 f_1(X_{t-1}^9) + 0.48 f_3(X_{t-1}^7) - 0.26 f_4(X_{t-1}^4) + \eta_{10} \\
 X_t^{11} &\leftarrow 0.41 f_2(X_{t-1}^{11}) + 0.54 f_3(X_t^2) - 0.55 f_2(X_t^{12}) + \eta_{11} \\
 X_t^{12} &\leftarrow -0.45 f_2(X_{t-1}^{12}) - 0.43 f_4(X_{t-2}^3) - 0.17 f_4(X_{t-2}^9) \times \eta_{12} \\
 X_t^{13} &\leftarrow 0.45 f_3(X_{t-1}^{13}) + \eta_{13} \\
 X_t^{14} &\leftarrow 0.28 f_2(X_{t-1}^{14}) + 0.37 f_1(X_{t-2}^{12}) \times \eta_{14} \\
 X_t^{15} &\leftarrow 0.52 f_3(X_{t-1}^{15}) + \eta_{15}
 \end{aligned}$$

Functions

$$\begin{aligned}
 f_1(x) &= x \\
 f_2(x) &= x (1 - 4 e^{-\frac{x^2}{2}}) \\
 f_3(x) &= x (1 - 4 x^3 e^{-\frac{x^2}{2}}) \\
 f_4(x) &= \cos(x)
 \end{aligned}$$

Noises

The η are white noises generated for each node or contribution using a normal distribution:

$$\eta \sim \mathcal{N}(0, 1)$$

Supplementary Table 2: 15 nodes model with combinations.

Nodes

$$\begin{aligned}
X_t^1 &\leftarrow \eta - 0.7 f_6(u(\eta + X_{t-1}^1)) - 0.87 f_5(u(\eta + (X_{t-1}^{14} \times X_{t-2}^1))) \\
X_t^2 &\leftarrow \eta + 0.65 f_1(u(\eta + X_{t-1}^2)) - 0.63 f_3(u(\eta + X_{t-2}^2)) + 0.79 f_3(u(\eta + X_{t-1}^5)) \\
X_t^3 &\leftarrow \eta - 0.76 f_5(u(\eta + X_{t-1}^3)) - 0.59 f_6(u(\eta + X_{t-1}^7)) - 0.85 f_2(u(\eta + X_{t-1}^{15})) \\
&\quad - 0.89 f_5(u(\eta + (X_{t-2}^{13} \times X_{t-1}^7))) \\
X_t^4 &\leftarrow \eta - 0.7 f_6(u(\eta + X_{t-1}^5)) - 0.86 f_2(u(\eta + X_{t-2}^8)) + 0.53 f_1(u(\eta + (X_{t-1}^4 \times X_{t-2}^9))) \\
X_t^5 &\leftarrow \eta + 0.54 f_2(u(\eta + (X_{t-1}^{14} \times X_{t-2}^6))) \\
X_t^6 &\leftarrow \eta - 0.85 f_2(u(\eta + X_{t-1}^6)) - 0.79 f_3(u(\eta + X_{t-2}^3)) + 0.59 f_1(u(\eta + X_{t-1}^4)) \\
&\quad + 0.75 f_3(u(\eta + X_t^1)) + 0.57 f_2(u(\eta + X_{t-1}^{14})) \\
X_t^7 &\leftarrow \eta + 0.74 f_1(u(\eta + X_{t-1}^7)) + 0.54 f_6(u(\eta + X_{t-1}^9)) - 0.53 f_2(u(\eta + (X_{t-1}^9 \times X_{t-1}^7))) \\
X_t^8 &\leftarrow \eta \times (-0.63 f_1(u(\eta + X_{t-1}^6)) + 0.81 f_5(u(\eta + X_t^{13})) + 0.53 f_6(u(\eta + (X_{t-2}^6 \times X_{t-1}^6))) \\
&\quad - 0.69 f_6(u(\eta + (X_t^{13} \times X_{t-1}^6)))) \\
X_t^9 &\leftarrow \eta + 0.79 f_3(u(\eta + X_{t-2}^4)) + 0.69 f_6(u(\eta + (X_{t-1}^9 \times X_{t-1}^{15}))) \\
X_t^{10} &\leftarrow \eta + 0.54 f_6(u(\eta + X_{t-1}^{10})) \\
X_t^{11} &\leftarrow \eta + 0.83 f_6(u(\eta + X_{t-1}^{11})) - 0.76 f_4(u(\eta + X_{t-1}^{13})) - 0.73 f_3(u(\eta + X_{t-1}^2)) \\
&\quad + 0.74 f_2(u(\eta + X_t^4)) - 0.87 f_2(u(\eta + X_{t-2}^{10})) + 0.72 f_4(u(\eta + X_{t-1}^{12})) \\
&\quad - 0.73 f_1(u(\eta + (X_{t-2}^{10} \times X_{t-1}^{13}))) \\
X_t^{12} &\leftarrow \eta + 0.7 f_3(u(\eta + X_{t-1}^{10})) - 0.55 f_5(u(\eta + X_t^9)) - 0.54 f_5(u(\eta + (X_{t-1}^{12} \times X_{t-1}^{10}))) \\
X_t^{13} &\leftarrow \eta - 0.62 f_3(u(\eta + X_{t-2}^{14})) - 0.61 f_1(u(\eta + (X_{t-1}^{13} \times X_{t-2}^{14}))) \\
X_t^{14} &\leftarrow \eta - 0.78 f_6(u(\eta + X_{t-1}^{14})) \\
X_t^{15} &\leftarrow \eta - 0.68 f_4(u(\eta + X_{t-1}^{15})) + 0.85 f_4(u(\eta + X_{t-2}^{15})) - 0.6 f_5(u(\eta + X_{t-2}^{10})) \\
&\quad + 0.68 f_6(u(\eta + X_{t-1}^{14})) + 0.81 f_4(u(\eta + (X_{t-1}^{14} \times X_{t-2}^{10})))
\end{aligned}$$

Functions

$$\begin{aligned}
u(x) &= \max(-1, \min(1, x)) \\
f_1(x) &= x \\
f_2(x) &= x(1 - 4e^{-\frac{x^2}{2}})/1.52387 \\
f_3(x) &= 4x^2 \\
f_4(x) &= 8x^3 \\
f_5(x) &= 16x^4 \\
f_6(x) &= \cos(\pi x)
\end{aligned}$$

Noises

The η are white noises generated for each node or contribution using a normal distribution:

$$\eta \sim \mathcal{N}(0, 0.1)$$