
Which bits went where? Past and future transfer entropy decomposition with the information bottleneck

Kieran A. Murphy

Dept. of Bioengineering, University of Pennsylvania
kieranm@seas.upenn.edu

Zhuowen Yin

Dept. of Bioengineering, University of Pennsylvania
zhuowen@seas.upenn.edu

Dani S. Bassett

Dept. of Bioengineering, University of Pennsylvania
Dept. of Electrical & Systems Engineering, University of Pennsylvania
Dept. of Neurology, Perelman School of Medicine, University of Pennsylvania
Dept. of Psychiatry, Perelman School of Medicine, University of Pennsylvania
Dept. of Physics & Astronomy, University of Pennsylvania
The Santa Fe Institute
Montreal Neurological Institute, McGill University
dsb@seas.upenn.edu

Abstract

Whether the system under study is a shoal of fish, a collection of neurons, or a set of interacting atmospheric and oceanic processes, transfer entropy measures the flow of information between time series and can detect possible causal relationships. Much like mutual information, transfer entropy is generally reported as a single value summarizing an amount of shared variation, yet a more fine-grained accounting might illuminate much about the processes under study. Here we propose to decompose transfer entropy and localize the bits of variation on both sides of information flow: that of the originating process's past and that of the receiving process's future. We employ the information bottleneck (IB) to compress the time series and identify the transferred entropy. We apply our method to decompose the transfer entropy in several synthetic recurrent processes and an experimental mouse dataset of concurrent behavioral and neural activity. Our approach highlights the nuanced dynamics within information flow, laying a foundation for future explorations into the intricate interplay of temporal processes in complex systems.

1 Introduction

Causality forms the backbone of explanations in science. Detecting cause and effect from observational data alone is generally impossible, as there can always be unobserved shared causes. However, certain signatures have been identified that can indicate *possible* causality, such as Granger causality [1], transfer entropy [2], and directed information [3]. Measuring such signatures can be a powerful step toward understanding a complex system, with applications as broad as groups of fish [4], neural signals [5, 6], and Earth-scale climate processes [7].

Transfer entropy quantifies the additional information a source process shares with a target process after accounting for the information contained in the target process's history [2, 8]. As a conditional

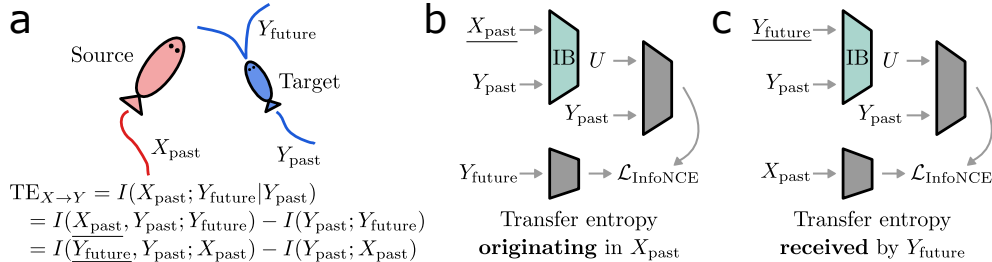


Figure 1: **Localizing transfer entropy in the past and future.** (a) If the past of a large red fish (the source) helps predict the future of a small blue fish (the target) after accounting for its past, then there is positive transfer entropy. Equivalently, if the future of the blue fish helps infer the past of the red fish after accounting for the blue fish’s past, then there is positive transfer entropy. (b) A machine learning scheme to extract the transfer entropy from the source’s past, using the information bottleneck (IB). (c) An analogous scheme to extract the transfer entropy from the target’s future.

mutual information, transfer entropy is nonparametric and generalizes to arbitrary probability distributions and relationships between variables, in contrast to Granger causality; transfer entropy and Granger causality are equivalent for Gaussian variables [9]. On the other hand, conditional mutual information can be challenging to estimate from limited data samples [10, 6].

Consider modeling the movement of the blue fish in Fig. 1a so as to forecast its future given its past. If the red fish causes the blue fish to alter its behavior, then the accuracy of the forecast will improve when information about the past of the red fish is incorporated into the model. The converse is not necessarily true—that an improvement in accuracy implies a causal relationship—but quantifying information flow in this way is often the closest to causality one can get with observational data alone (i.e., without having the ability to intervene in the data collection process) [11, 9].

While the above framing in terms of forecasting with and without the source’s past is the common way to present transfer entropy, there is an equivalent alternative that we will leverage in this work. Again employing the fish of Fig. 1a, consider the task of inferring the red fish’s movement given the blue fish’s movement over the same time frame. If the accuracy of the model improves when incorporating the future of the blue fish, there is information flow from the past of the red fish to the future of the blue fish. With these two equivalent expressions for transfer entropy, we can focus on the effect of incorporating either the past of the red fish or the future of the blue fish into a predictive model. The net information change in either case must be equal to the transfer entropy, but the originating and terminating variation can look different on either side of the information flow. Our goal in this work is to identify the variation on either side by using the information bottleneck.

The information bottleneck (IB) is a method to isolate targeted variation in a learned compression scheme [12, 13], allowing for control over learned representations [14] and fine-grained inspection of *what* the information is [15]. Here we propose an IB scheme to encapsulate the transfer entropy from either side of an information flow into a learned compression space. The encapsulated information can then be decomposed in terms of multiple sources via the distributed IB [15, 16], inspected as a local quantity varying in time [17], and used as a targeted representation for downstream tasks [18].

2 Approach

Transfer entropy measures information flow from a *source* time series to a *target* time series and is the information gained in either of two equivalent scenarios (Fig. 1a): (i) the information that the source’s past adds to the target’s past about the target’s future, and (ii) the information that the target’s future adds to the target’s past about the source’s past. We will set up both scenarios as a constrained communication problem with the information bottleneck [12] in order to identify and then decompose the bits of transferred entropy in the source’s past and in the target’s future.

Let X_t and Y_t represent two random processes, the source and target, and let t be a discrete index for time. We assume the processes are stationary and will use the same time horizon, τ , into the past and future. For shorthand, $Y_{-\tau:0}$ and $Y_{0:\tau}$ (using python indexing for τ steps into the past and future, respectively [8]) will be written as Y_{past} and Y_{future} , respectively.

The forecasting information inherent to the target process is the mutual information [19] between the recent past and the near future, $I(Y_{\text{past}}; Y_{\text{future}})$. The transfer entropy from source to target is the increase in forecasting information when incorporating the recent history of the source,

$$\text{TE}_{X \rightarrow Y} = I(X_{\text{past}}, Y_{\text{past}}; Y_{\text{future}}) - I(Y_{\text{past}}; Y_{\text{future}}). \quad (1)$$

Equivalently, the transfer entropy can be found in the target’s future as the amount of information added about the source’s past,

$$\text{TE}_{X \rightarrow Y} = I(Y_{\text{future}}, Y_{\text{past}}; X_{\text{past}}) - I(Y_{\text{past}}; X_{\text{past}}). \quad (2)$$

Eqns. 1 and 2 can be seen as the difference in predictive power between two extremes of a continuous spectrum parameterized by the amount of information incorporated about X_{past} and Y_{future} , respectively. We will traverse the spectrum by compressing the added piece of information with the information bottleneck (Fig. 1b). Importantly, to compress the added information requires Y_{past} as context in both cases, which can simply be passed along as an additional input to the encoder. Letting $U = f(X_{\text{past}}, Y_{\text{past}})$ be the compression of the additional information, and β be an information cost, minimizing the following Lagrangian with respect to U traverses the spectrum relevant to Eqn. 1,

$$\mathcal{L} = \beta I(X_{\text{past}}, Y_{\text{past}}; U) - I(U, Y_{\text{past}}; Y_{\text{future}}). \quad (3)$$

In the limit $\beta \rightarrow \infty$, no information is used about the source and we recover the self-forecasting information $I(Y_{\text{past}}; Y_{\text{future}})$. At the other extreme, $\beta \rightarrow 0$, all information from the source is used. The difference in the loss between these extremes is the transfer entropy. In practice, we decrease β logarithmically over the course of a single training run to obtain the full spectrum of transferred information; the transfer entropy is obtained by the difference between the endpoints of the trajectory. To optimize with machine learning, we employ the same information cost as a β -VAE [20, 13], and a lower bound on the forecasting information via InfoNCE [21],

$$\mathcal{L} = \beta \mathbb{E}[D_{\text{KL}}(p(u|x_{\text{past}}, y_{\text{past}})||r(u))] - \mathbb{E}[\mathcal{L}_{\text{NCE}}(f(u, y_{\text{past}}), g(y_{\text{future}}))]. \quad (4)$$

The first expectation is over $p(x_{\text{past}}, y_{\text{past}})$ and the second is over $p(u|x_{\text{past}}, y_{\text{past}})p(x_{\text{past}}, y_{\text{past}}, y_{\text{future}})$. $r(u)$ is a prior distribution, taken to be the standard normal $\mathcal{N}(\mathbf{0}, \mathbf{1})$, and $f(\cdot)$ and $g(\cdot)$ are functions parameterized by neural networks. While Eqn. 4 targets the transfer entropy in X_{past} , the transfer entropy in Y_{future} is found analogously with the compression variable $U = f(Y_{\text{future}}, Y_{\text{past}})$ (Fig. 1c).

Without loss of generality, assume the source and target are compositions of processes—e.g., multiple physiological signals from each fish in Fig. 1a—so that $X_t = (X_t^1, \dots, X_t^N)$. For added interpretability around the extracted information, we compress each component process and/or each point in time separately by distributing an information bottleneck to each as a separate random variable [15, 16]. The Lagrangian corresponding to Eqn. 3 is

$$\mathcal{L} = \beta \sum_i^N [I(X_{\text{past}}^i, Y_{\text{past}}; U^i)] - I(U, Y_{\text{past}}; Y_{\text{future}}). \quad (5)$$

We note important distinctions from the recently proposed Transfer Entropy Bottleneck [18]: we propose a simpler joint encoding strategy, a decomposition of the future transfer entropy, and a distributed IB scheme for more granularity when analyzing information flow between processes.

3 Results and Discussion

We will now work through synthetic examples to build intuition for the transfer entropy decomposition, and then analyze an experimental dataset from a group of mice where neural activity and behavioral data were collected simultaneously. All quantities are in bits unless specified otherwise. Implementation specifics can be found in the Appendix.

Synthetic Boolean networks. For the systems of binary-valued processes in Fig. 2, we generated trajectories of length 10^4 and trained the end-to-end setups shown in Fig. 1b&c with a time horizon $\tau = 3$. A classic example regarding transfer entropy [8] (slightly modified) is shown in Fig. 2a&b. Blue and orange are random processes that determine the next timestep of the green process through




Figure 2: **Transferred entropy in binary-valued recurrent networks.** (a) *Left:* The update rule for four processes, where nodes without inputs (blue, orange) are randomly sampled at each timestep. The source and target processes are indicated by the shaded boxes and marked X and Y , respectively. *Middle:* Distributed information planes that visualize the decomposition of transfer entropy in the source’s past and the target’s future. *Right:* The share of transfer entropy residing in different timesteps of the source’s past (top) and target’s future (bottom) (taken from the rightmost point of the trajectories in the middle). (b) Same as panel a, but with different target processes. (c) Same as panel a, but with randomly generated connection weights and an integrate-and-fire scheme.

an XOR interaction, and the red process stores the most recent state of green. There is positive transfer entropy when considering the blue and orange processes together as a single composite source, and either or both of the green and red processes as targets. With the distributed IB, we bottlenecked each process and each timestep to identify the origin and the terminus of the transferred entropy.

The difference in total information $I(U, Y_{\text{past}}, Y_{\text{future}})$ between the start and end of optimization is the transfer entropy, and therefore the same when focusing on the transfer entropy’s origin in X_{past} and its destination in Y_{future} . However, the trajectories are qualitatively different for the past and future decompositions due to the nature of interactions between the constituent processes [16]. In Fig. 2a, the single bit of transfer entropy starts as two bits in the blue and orange sources’ past at timestep $t - 1$, and then ends as one bit in the future of the red process at timestep $t + 1$. Note that the green and red process contain the same information (with a delay of one), meaning the transferred entropy equivalently resides in the green process at timestep t . When repeating the distributed IB optimization for Fig. 2a, the transferred entropy was localized in the green process at timestep t in half of the optimizations and in the red process at timestep $t + 1$ in the other half. The degeneracy resolves if we ignore the green process (Fig. 2b), and the transfer entropy increases to two bits. Four bits originate in the blue and orange processes and then terminate in the red process over the future two timesteps.

The benefit of machine learning becomes more apparent when decomposing information flow in more complicated processes. In Fig. 2c, a set of connected processes includes two sources (blue and orange, random), two hidden states (grey), and two targets (green and red). Nodes with connections from the previous timestep follow an integrate-and-fire update rule, meaning they will fire if the sum of their inputs, weighted by +1 or -1 as indicated by the solid and dashed lines, is greater than zero. The share of information from different timesteps showcases the nontrivial recurrence established by the randomly selected connection scheme and provides clues about the underlying process interactions.

Experimental mouse data. We next analyzed calcium imaging and spontaneous behavioral data from six mice publicly uploaded with Benisty et al. [22], with a time horizon of 1 second (10 samples). A rough schematic of the experimental setup is shown in Fig. 3a.

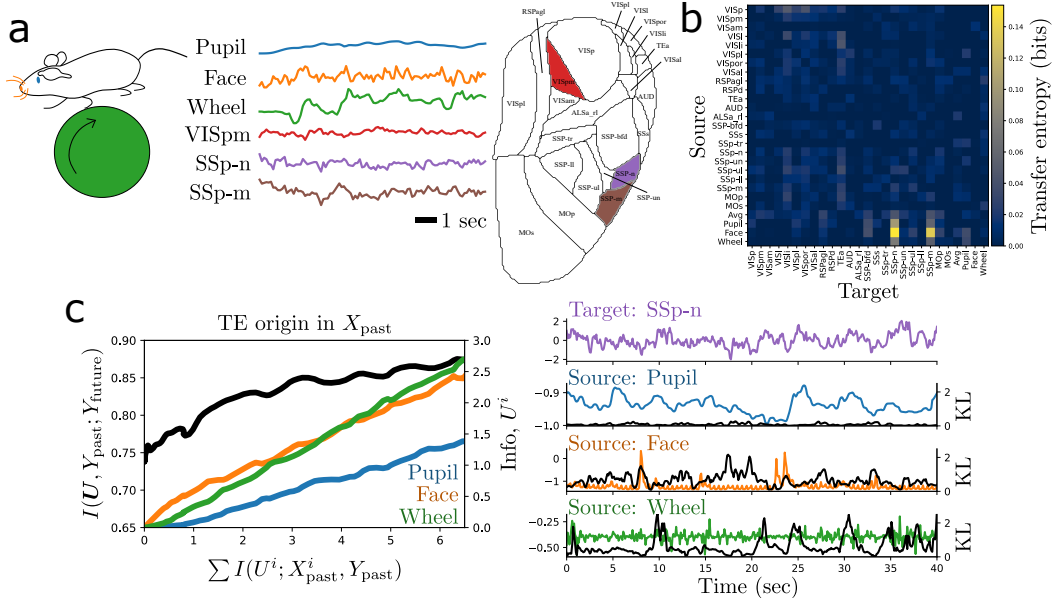


Figure 3: **Transfer entropy between brain and behavior.** (a) Concurrent neural and behavioral recordings were taken of six mice; example time series shown on the right with the brain regions shown with matching colors in the atlas. (b) Pairwise transfer entropy between the 23 brain regions, the reference average signal (Avg), and three behavioral streams. (c) Transfer entropy decomposition from behaviors to the purple region from a, the primary somatosensory area for the nose (SSp-n). The instantaneous Kullback-Leibler (KL) cost in natural units (nats) per channel (black) is shown concurrently with the raw time series (colored).

We estimated the pairwise transfer entropy between all 23 regions, their common average signal (‘Avg’), and the three behavioral time series, one pair at a time using the difference in forecasting InfoNCE values with and without the source (Fig. 3b). The largest pairwise values link behavior to brain region SSp-n, and we decompose the multivariate transfer entropy in Fig. 3c. The pupil time series contained the second largest transfer entropy with SSp-n (Fig. 3b), though its contribution was the smallest when combined with the face and wheel data, suggesting redundant information.

At any point along the information bottleneck spectrum, we can inspect the learned compression scheme. We encoded a stream of the validation data with the three learned encoders at an intermediate value of $I(U, Y_{past}; Y_{future})$ and display the Kullback-Leibler (KL) cost—whose expectation serves as a penalty on the transmitted information in Eqn. 4—at each point in time (Fig. 3c, right). Note each encoding scheme embeds the source past along with the target past as context, meaning the spikes in KL cost could arise from an interaction between the two signals. In this way we obtain a fine-grained picture suggesting where in the source time series the transfer entropy spawns, which could be paired with a local estimate of the transfer entropy [23] to take a microscope to information transfer between processes.

Discussion. The quantification of information flows between parts of a system via transfer entropy is an important step toward a deeper understanding and serves as a screening for possible causation. In this work, we localized transfer entropy on both sides of an information flow: from its origin in the source’s past to its terminus in the target’s future. We note that although both routes to transfer entropy outlined are equivalent, the prediction task involved is qualitatively different. When compressing the source’s past, the prediction task is to forecast the target’s future from its past. By contrast, when compressing the target’s future, the prediction task is to link the target’s past with the source’s past, which will generally share a lower amount of information than the target’s past with its future. The difficulty of optimization for the two formulations may thus be different in practice.

Finally, we note that conditioning on additional processes when computing transfer entropy allows one to exclude the influence of the processes [2, 17]. The proposed framework readily handles such additional processes by appending them to both instances of Y_{past} in the schematics of Fig. 1b, c.

References

- [1] Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438, 1969. 1
- [2] Thomas Schreiber. Measuring information transfer. *Physical review letters*, 85(2):461, 2000. 1, 5
- [3] James Massey et al. Causality, feedback and directed information. In *Proc. Int. Symp. Inf. Theory Applic.(ISITA-90)*, volume 2, 1990. 1
- [4] Sachit Butail, Violet Mwaffo, and Maurizio Porfiri. Model-free information-theoretic approach to infer leadership in pairs of zebrafish. *Physical Review E*, 93(4):042411, 2016. 1
- [5] Leonardo Novelli and Joseph T Lizier. Inferring network properties from time series using transfer entropy and mutual information: Validation of multivariate versus bivariate approaches. *Network Neuroscience*, 5(2):373–404, 2021. 1
- [6] Matthäus Staniek and Klaus Lehnertz. Symbolic transfer entropy. *Physical review letters*, 100(15):158101, 2008. 1, 2
- [7] Jakob Runge, Sebastian Bathiany, Erik Bollt, Gustau Camps-Valls, Dim Coumou, Ethan Deyle, Clark Glymour, Marlene Kretschmer, Miguel D Mahecha, Jordi Muñoz-Marí, et al. Inferring causation from time series in earth system sciences. *Nature communications*, 10(1):2553, 2019. 1
- [8] Ryan G James, Nix Barnett, and James P Crutchfield. Information flows? a critique of transfer entropies. *Physical review letters*, 116(23):238701, 2016. 1, 2, 3
- [9] Lionel Barnett, Adam B Barrett, and Anil K Seth. Granger causality and transfer entropy are equivalent for gaussian variables. *Physical review letters*, 103(23):238701, 2009. 2
- [10] Sudipto Mukherjee, Himanshu Asnani, and Sreeram Kannan. Ccmi : Classifier based conditional mutual information estimation. In Ryan P. Adams and Vibhav Gogate, editors, *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pages 1083–1093. PMLR, 22–25 Jul 2020. URL <https://proceedings.mlr.press/v115/mukherjee20a.html>. 2
- [11] Judea Pearl. *Causality*. Cambridge university press, 2009. 2
- [12] Naftali Tishby, Fernando C. Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000. 2
- [13] Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. Deep variational information bottleneck. *International Conference on Learning Representations (ICLR)*, 2017. 2, 3
- [14] Ziv Goldfeld and Yury Polyanskiy. The information bottleneck problem and its applications in machine learning. *IEEE Journal on Selected Areas in Information Theory*, 1(1):19–38, 2020. 2
- [15] Kieran A Murphy and Dani S. Bassett. Interpretability with full complexity by constraining feature information. In *International Conference on Learning Representations (ICLR)*, 2023. URL https://openreview.net/forum?id=R_OL5mLhsv. 2, 3
- [16] Kieran A Murphy and Dani S Bassett. Information decomposition in complex systems via machine learning. *Proceedings of the National Academy of Sciences*, 121(13):e2312988121, 2024. 2, 3, 4
- [17] Joseph T Lizier, Mikhail Prokopenko, and Albert Y Zomaya. Local information transfer as a spatiotemporal filter for complex systems. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 77(2):026110, 2008. 2, 5
- [18] Damjan Kalajdzievski, Ximeng Mao, Pascal Fortier-Poisson, Guillaume Lajoie, and Blake Aaron Richards. Transfer entropy bottleneck: Learning sequence information transfer. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=kJcWlP7BRs>. 2, 3

- [19] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 1999. [3](#)
- [20] Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations (ICLR)*, 2017. [3](#)
- [21] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [3](#)
- [22] Hadas Benisty, Daniel Barson, Andrew H Moberly, Sweyta Lohani, Lan Tang, Ronald R Coifman, Michael C Crair, Gal Mishne, Jessica A Cardin, and Michael J Higley. Rapid fluctuations in functional connectivity of cortical networks encode spontaneous behavior. *Nature Neuroscience*, 27(1):148–158, 2024. [4](#)
- [23] Ramón Martínez-Cancino, Arnaud Delorme, Johanna Wagner, Kenneth Kreutz-Delgado, Roberto C Sotero, and Scott Makeig. What can local transfer entropy tell us about phase-amplitude coupling in electrophysiological signals? *Entropy*, 22(11):1262, 2020. [5](#)

Parameter	Value
Bottleneck MLP architecture	[64 dense leaky_ReLU]
Bottleneck embedding space dimension	8
Y_{past} encoder architecture	[128 LSTM tanh, 64 dense leaky_ReLU]
Y_{past} embedding space dimension	32
Encoder MLP architecture (to shared embedding space)	[256 dense leaky_ReLU, 256 dense leaky_ReLU]
Predicted quantity encoder architecture	[128 LSTM tanh, 64 dense leaky_ReLU]
InfoNCE similarity metric $s(u, v)$	Euclidean squared
Batch size	128
Optimizer	Adam
Learning rate	3×10^{-4}
β_{initial}	5×10^{-5}
β_{final}	3
Annealing steps	2×10^4
Horizon length	3

Table 1: Training parameters for synthetic experiments.

Parameter	Value
Bottleneck MLP architecture	[256 dense leaky_ReLU]
Bottleneck embedding space dimension	8
Y_{past} encoder architecture	[128 LSTM tanh, 128 LSTM tanh, 256 dense leaky_ReLU]
Y_{past} embedding space dimension	32
Encoder MLP architecture (to shared embedding space)	[256 dense leaky_ReLU]
Predicted quantity encoder architecture	[128 LSTM tanh, 128 LSTM tanh, 256 dense leaky_ReLU]
InfoNCE similarity metric $s(u, v)$	Euclidean squared
Batch size	128
Optimizer	Adam
Learning rate	3×10^{-4}
β_{initial}	10^{-3}
β_{final}	1
Annealing steps	2×10^4
Horizon length	10

Table 2: Training parameters for mouse data analysis.

A Implementation specifics

All experiments were implemented in Tensorflow and run on a single computer with a 12 GB GeForce RTX 3060 GPU. An iPython notebook to run the synthetic examples (Fig. 2) is included with the submission. Training hyperparameters and architecture details are shown in Tables 1 and 2.