

On the information-theoretic structure of distributed measurements

David Balduzzi*

Max Planck Institute for Intelligent Systems, Tübingen, Germany

david.balduzzi@tuebingen.mpg.de

The internal structure of a measuring device, which depends on what its components are and how they are organized, determines how it categorizes its inputs. This paper presents a geometric approach to studying the internal structure of measurements performed by distributed systems such as probabilistic cellular automata. It constructs the quale, a family of sections of a suitably defined presheaf, whose elements correspond to the measurements performed by all subsystems of a distributed system. Using the quale we quantify (i) the information generated by a measurement; (ii) the extent to which a measurement is context-dependent; and (iii) whether a measurement is decomposable into independent submeasurements, which turns out to be equivalent to context-dependence. Finally, we show that only indecomposable measurements are more informative than the sum of their submeasurements.

1 Introduction

Any classical physical system (by which we simply mean any deterministic function) can be taken as a measuring apparatus or input/output device. For example, a thermometer takes inputs from the atmosphere and outputs numbers on a digital display. The thermometer categorizes inputs by temperature and is blind to, say, differences in air pressure.

Classical measurements are formalized as follows:

Definition 1. Given a classical physical system with state space X , a *measuring device* is a function $f : X \rightarrow \mathbb{R}$. The output $r \in \mathbb{R}$ is the *reading* and the pre-image $f^{-1}(r) \subset X$ is the *measurement*.

From this point of view a thermometer and a barometer are two functions, $T : X \rightarrow \mathbb{R}$ and $B : X \rightarrow \mathbb{R}$, mapping the state space X of configurations (positions and momenta) of atmospheric particles to real numbers. When the thermometer outputs 2° , it specifies that the atmospheric configuration was in the pre-image $T^{-1}(2^\circ)$ which, assuming the thermometer perfectly measures temperature, is *exactly* characterized as atmospheric configurations with temperature 2° . Similarly, the pre-images generated by the barometer group atmospheric configurations by pressure.

The classical definition of measurement takes a thermometer as a monolithic object described by a single function from atmospheric configurations to real numbers. The internal structure of the thermometer – that is composed of countless atoms and molecules arranged in an extremely specific manner – is swept under the carpet (or, rather, into the function).

This paper investigates the structure of measurements performed by *distributed* systems. We do so by adapting Definition 1 to a large class of systems that contains networks of Boolean functions [18], Conway’s game of life [9, 12] and Hopfield networks [3, 15] as special cases.

Our motivation comes from prior work investigating information processing in discrete neural networks [6, 7]. The brain X can be thought of as an enormously complicated measuring device $S \times X \xrightarrow{f} X$

*I thank Dominik Janzing, Sanjeevi Krishnan and the reviewers for useful comments.

mapping sensory states $s \in S$ and prior brain states $x \in X$ to subsequent brain states. Analyzing the functional dependencies implicit in cortical computations reduces to analyzing how the measurements performed by the brain are composed out of submeasurements by subdevices such as individual neurons and neuronal assemblies. The cortex is of particular interest since it seemingly effortlessly integrates diverse contextual data into a unified gestalt that determines behavior. The measurements performed by different neurons appear to interact in such a way that they generate more information jointly than separately. To improve our understanding of how the cortex integrates information we need to a formal language for analyzing how context affects measurements in distributed systems.

As a first step in this direction, we develop methods for analyzing the geometry of measurements performed by functions with overlapping domains. We propose, roughly speaking, to study context-dependence in terms of the geometry of intersecting pre-images. However, since we wish to work with both probabilistic and deterministic systems, things are a bit more complicated.

We sketch the contents of the paper. Section §2 lays the groundwork by introducing the category of stochastic maps Stoch . Our goal is to study finite set valued functions and conditional probability distributions on finite sets. However, rather than work with sets, functions and conditional distributions, we prefer to study stochastic maps (Markov matrices) between function spaces on sets. We therefore introduce the faithful functor \mathcal{V} taking functions on sets to Markov matrices:

$$[f : X \rightarrow Y] \mapsto [\mathcal{V}f : \mathcal{V}X \rightarrow \mathcal{V}Y],$$

where $\mathcal{V}X$ is functions from X to \mathbb{R} . Conditional probability distributions $p(y|x)$ can also be represented using stochastic maps.

Working with linear operators instead of set-valued functions is convenient for two reasons. First, it unifies the deterministic and probabilistic cases in a single language. Second, the dual T^\natural of a stochastic map T provides a symmetric treatment of functions and their corresponding inverse image functions. Recall the inverse of function $f : X \rightarrow Y$ is $f^{-1} : Y \rightarrow \underline{2}^X$, which takes values in the *powerset* of X , rather than X itself. Dualizing a stochastic map flips the domain and range of the original map, without introducing any new objects:

$$[f^{-1} : Y \rightarrow \underline{2}^X] \text{ corresponds to } [(\mathcal{V}f)^\natural : \mathcal{V}Y \rightarrow \mathcal{V}X], \quad (1)$$

see Proposition 2.

Section §3 introduces distributed dynamical systems. These extend probabilistic cellular automata by replacing cells (space coordinates) with occasions (spacetime coordinates: cell k at time t). Inspired by [1, 14], we treat distributed systems as collections of stochastic maps between function spaces so that processes (stochastic maps) take center stage, rather than their outputs. Although the setting is abstract, it has the advantage that it is *scalable*: using a coarse-graining procedure introduced in [5] we can analyze distributed systems at any spatiotemporal granularity.

Distributed dynamical systems provide a rich class of toy universes. However, since these toy universes do not contain conscious observers we confront Bell's problem [8]: "What exactly qualifies some physical [system] to play the role of 'measurer'?" In our setting, where we do not have to worry about collapsing wave-functions or the distinction between macroscopic and microscopic processes, the solution is simple: *every* physical system plays the role of measurer. More precisely, we track measurers via the category $\text{Sys}_{\mathbf{D}}$ of subsystems of \mathbf{D} . Each subsystem \mathbf{C} is equipped with a mechanism $m_{\mathbf{C}}$ which is constructed by gluing together the mechanisms of the occasions in \mathbf{C} and averaging over extrinsic noise.

Measuring devices are typically analyzed by varying their inputs and observing the effect on their outputs. By contrast this paper fixes the output and *varies the device over all its subdevices* to obtain

a family of submeasurements parametrized by all subsystems in $\text{Sys}_{\mathbf{D}}$. The internal structure of the measurement performed by \mathbf{D} is then studied by comparing submeasurements.

We keep track of submeasurements by observing that they are sections of a suitably defined presheaf. Sheaf theory provides a powerful machinery for analyzing relationships between objects and subobjects [19], which we adapt to our setting by introducing the structure presheaf \mathcal{F} , a contravariant functor from $\text{Sys}_{\mathbf{D}}$ to the category of measuring devices $\text{Meas}_{\mathbf{D}}$ on \mathbf{D} . Importantly, \mathcal{F} is *not* a sheaf: although the gluing axiom holds, uniqueness fails, see Theorem 4. This is because the restriction operator in Meas is (essentially) marginalization, and of course there are infinitely many joint distributions $p(x, y)$ that yield marginals $p(x)$ and $p(y)$.

Section §4 adapts Definition 1 to distributed systems and introduces the simplest quantity associated with a measurement: effective information, which quantifies its precision, see Proposition 5. Crucially, effective information is *context-dependent* – it is computed relative to a baseline which may be completely uninformative (the so-called null system) or provided by a subsystem.

Finally entanglement, introduced in §5, quantifies the obstruction (in bits) to decomposing a measurement into independent submeasurements. It turns out, see discussion after Theorem 10, that entanglement quantifies the extent to which a measurement is context-dependent – the extent to which contextual information provided by one submeasurement is useful in understanding another. Theorem 9 shows that a measurement is more precise than the sum of its submeasurements *only if* entanglement is non-zero. Precision is thus inextricably bound to context-dependence and indecomposability. The failure of unique descent is thus a feature, not a bug, since it provides “elbow room” to build measuring devices that are *not* products of subdevices.

Space constraints prevent us from providing concrete examples; the interested reader can find these in [5–7]. Our running examples are the deterministic set-valued functions

$$f : X \rightarrow Y \quad \text{and} \quad g : X \times Y \rightarrow Z$$

which we use to illustrate the concepts as they are developed.

2 Stochastic maps

Any conditional distribution $p(y|x)$ on finite sets X and Y can be represented as a matrix as follows. Let $\mathcal{V}X = \{\varphi : X \rightarrow \mathbb{R}\}$ denote the vector space of real valued functions on X and similarly for Y . $\mathcal{V}X$ is equipped with Dirac basis $\{\delta_x : X \rightarrow \mathbb{R} | x \in X\}$, where

$$\delta_x(x') = \begin{cases} 1 & \text{if } x = x' \\ 0 & \text{else.} \end{cases}$$

Given a conditional distribution $p(y|x)$ construct matrix m_p with entry $p(y|x)$ in column δ_x and row δ_y . Matrix m_p is *stochastic*: it has nonnegative entries and its columns sum to 1. Alternatively, given a stochastic matrix $m : \mathcal{V}X \rightarrow \mathcal{V}Y$, we can recover the conditional distribution. The Dirac basis induces Euclidean metric

$$\langle \bullet | \bullet \rangle : \mathcal{V}X \otimes \mathcal{V}X \rightarrow \mathbb{R} : \langle \sum \alpha_x \delta_x | \sum \beta_x \delta_x \rangle = \sum \alpha_x \beta_x \quad (2)$$

which identifies vector spaces with their duals $\mathcal{V}X \approx (\mathcal{V}X)^*$. Let $p_m(y|x) := \langle \delta_y | m(\delta_x) \rangle$.

Definition 2. The *category of stochastic maps* Stoch has function spaces $\mathcal{V}X$ for objects and stochastic matrices $m : \mathcal{V}X \rightarrow \mathcal{V}Y$ with respect to Dirac bases for arrows. We identify $(\mathcal{V}X)^*$ with $\mathcal{V}X$ using the Dirac basis without further comment below.

Definition 3. The *dual* of surjective stochastic map $m : \mathcal{V}X \rightarrow \mathcal{V}Y$ is the composition $m^\natural := \mathcal{V}Y \xrightarrow{m^* \circ ren} \mathcal{V}X$, where *ren* is the unique map making diagram

$$\begin{array}{ccc} (\mathcal{V}Y)^* & \xrightarrow{m^*} & (\mathcal{V}X)^* \\ ren \uparrow & & \omega_X \downarrow \\ (\mathcal{V}Y)^* & \xrightarrow{\omega_Y} & \mathbb{R} \end{array}$$

commute. Precomposing m^* with *ren* renormalizes¹ its columns to sum to 1. The stochastic dual of a stochastic transform is stochastic; further, if m is stochastic then $(m^\natural)^\natural = m$.

Category **Stoch** is described in terms of braid-like generators and relations in [11]. A more general, but also more complicated, category of conditional distributions was introduced by Giry [13], see [20].

Example 1 (deterministic functions). Let **FSet** be the category of finite sets. Define faithful functor $\mathcal{V} : \mathbf{FSet} \rightarrow \mathbf{Stoch}$ taking set X to $\mathcal{V}X$ and function $f : X \rightarrow Y$ to stochastic map $\mathcal{V}f : \mathcal{V}X \rightarrow \mathcal{V}Y : \delta_x \mapsto \delta_{f(x)}$. It is easy to see that $\mathcal{V}(X \times Y) = \mathcal{V}X \otimes \mathcal{V}Y$ and $\mathcal{V}(X \cup Y) = \mathcal{V}X \times \mathcal{V}Y$.

We introduce special notation for commonly used functions:

- *Set inclusion.* For any inclusion $i : X \hookrightarrow Y$ of sets, let $\iota := \mathcal{V}i : \mathcal{V}X \rightarrow \mathcal{V}Y$ denote the corresponding stochastic map. Two important examples are
 - *Point inclusion.* Given $x \in X$ define $\iota_x : \mathbb{R} \rightarrow \mathcal{V}X : 1 \mapsto \delta_x$.
 - *Diagonal map.* Inclusion $\Delta : X \hookrightarrow X \times X : x \mapsto (x, x)$ induces $\iota_\Delta : \mathcal{V}X \rightarrow \mathcal{V}X \otimes \mathcal{V}X : \delta_x \mapsto \delta_x \otimes \delta_x$.
- *Terminal map.* Let $\omega_X : \mathcal{V}X \rightarrow \mathbb{R} : \delta_x \mapsto 1$ denote the terminal map induced by $X \rightarrow \{\bullet\}$.
- *Projection.* Let $\pi_{X,Y,X} : \mathcal{V}X \otimes \mathcal{V}Y \rightarrow \mathcal{V}X : \delta_x \otimes \delta_y \mapsto \delta_x$ denote the projection induced by $pr_{X \times Y, X} : X \times Y \rightarrow X : (x, y) \mapsto x$.

Proposition 1 (dual is Bayes over uniform distribution). *The dual of a stochastic map applies Bayes rule to compute the posterior distribution $\langle m^\natural(\delta_y) | \delta_x \rangle = p_m(x|y)$ using the uniform probability distribution.*

Proof: The uniform distribution is the dual $\omega_X^\natural : \mathbb{R} \rightarrow \mathcal{V}X : 1 \mapsto \frac{1}{|X|} \sum_x \delta_x$ of the terminal map $\omega_X : \mathcal{V}X \rightarrow \mathbb{R}$. It assigns equal probability $p_{\omega^\natural}(x) = \frac{1}{|X|}$ to all of X 's elements, and can be characterized as the maximally uninformative distribution [16]. Let $m : \mathcal{V}X \rightarrow \mathcal{V}Y$. The normalized transpose is

$$m^\natural(\delta_y) = \sum_x \frac{p_m(y|x)}{\sum_{x'} p_m(y|x')} \delta_x = \sum_x \frac{p_m(y|x) \cdot p_{\omega^\natural}(x)}{\sum_{x'} p_m(y|x') p_{\omega^\natural}(x')} \delta_x = \sum_x p_m(x|y) \cdot \delta_x. \blacksquare$$

Remark 1. Note that $p_m(x|y) := \langle m^\natural(\delta_y) | \delta_x \rangle \neq \langle \delta_y | m(\delta_x) \rangle =: p_m(y|x)$. Dirac's bra-ket notation must be used with care since stochastic matrices are not necessarily symmetric [10].

Corollary 2 (preimages). *The dual $(\mathcal{V}f)^\natural : \mathcal{V}Y \rightarrow \mathcal{V}X$ of stochastic map $\mathcal{V}f : \mathcal{V}X \rightarrow \mathcal{V}Y$ is conditional distribution*

$$p_{\mathcal{V}f}(x|y) = \begin{cases} \frac{1}{|f^{-1}(y)|} & \text{if } f(x) = y \\ 0 & \text{else.} \end{cases} \quad (3)$$

¹If m is not surjective, i.e. if one of the rows has all zero entries, then the renormalization is not well-defined.

Proof: By the proof of Proposition 1

$$(\mathcal{V}f)^\natural(\delta_y) = \frac{1}{|f^{-1}(y)|} \sum_{\{x|f(x)=y\}} \delta_x. \blacksquare$$

The support of $p_{\mathcal{V}f}(X|y)$ is $f^{-1}(y)$. Elements in the support are assigned equal probability, thereby treating them as an undifferentiated list. Dual $(\mathcal{V}f)^\natural$ thus generalizes the inverse image $f^{-1} : Y \rightarrow \underline{2}^X$. Conveniently however, the dual $(\mathcal{V}X)^\natural$ simply flips the domain and range of $\mathcal{V}f$, whereas the inverse image maps to powerset $\underline{2}^X$, an entirely new object.

Corollary 3 (marginalization with respect to uniform distribution). *Precomposing $\mathcal{V}X \otimes \mathcal{V}Y \xrightarrow{m} \mathcal{V}Z$ with the dual π_X^\natural to $\mathcal{V}X \otimes \mathcal{V}Y \xrightarrow{\pi_X} \mathcal{V}X$ marginalizes $p_m(z|x, y)$ over the uniform distribution on Y .*

Proof: By Corollary 2 we have $\pi_X^\natural : \mathcal{V}X \rightarrow \mathcal{V}X \otimes \mathcal{V}Y : \delta_y \mapsto \frac{1}{|Y|} \sum_{y \in Y} \delta_x \otimes \delta_y$. It follows immediately that

$$p_{m \circ \pi_X^\natural}(z|x) = \frac{1}{|Y|} \sum_{y \in Y} p_m(z|x, y). \blacksquare$$

Precomposing with π_X^\natural treats inputs from Y as extrinsic noise. Although duals can be defined so that they implement Bayes' rule with respect to other probability distributions, this paper restricts attention to the simplest possible renormalization of columns, Definition 2. The uniform distribution is convenient since it uses minimal prior knowledge (it depends only on the number of elements in the set) to generalize pre-images to the stochastic case, Proposition 2.

3 Distributed dynamical systems

Probabilistic cellular automata provide useful toy models of a wide range of physical and biological systems. A cellular automaton consists of a collection of cells, each equipped with a mechanism whose output depends on the prior outputs of its neighbors. Two important examples are

Example 2 (Conway's game of life). The cellular automaton is a grid of deterministic cells with outputs $\{0, 1\}$. A cell outputs 1 at time t iff: (i) three of its neighbors outputted 1s at time $t - 1$ or (ii) it and two neighbors outputted 1s at $t - 1$. Remarkably, a sufficiently large game of life grid can implement any deterministic computation [9].

Example 3 (Hopfield networks). These are probabilistic cellular automata [3, 15], again with outputs $\{0, 1\}$. Cell n_k fires with probability proportional to

$$p(n_{k,t} = 1 | n_{\bullet, t-1}) \propto \exp \left[\frac{1}{T} \sum_{j \rightarrow k} \alpha_{jk} \cdot n_{j, t-1} \right].$$

Temperature T controls network stochasticity. Attractors $\{\xi^1, \dots, \xi^N\}$ are embedded into a network by setting the connectivity matrix as $\alpha_{jk} = \sum_{\mu=1}^N (2\xi_j^\mu - 1)(2\xi_k^\mu - 1)$.

It is useful to take a finer perspective on cellular automata by decomposing them into spacetime coordinates or occasions [5]. An occasion $v_l = n_{i,t}$ is a cell n_i at a time point t . Two occasions are linked $v_k \rightarrow v_l$ if there is a connection from v_k 's cell to v_l 's (because they are neighbors or the same cell) and their time coordinates are $t - 1$ and t respectively for some t , so occasions form a directed graph. More generally:

Definition 4. A distributed dynamical system \mathbf{D} consists of the following data:

- D1. Directed graph.** A graph $G_{\mathbf{D}} = (V_{\mathbf{D}}, E_{\mathbf{D}})$ with a finite set of vertices or occasions $V_{\mathbf{D}} = \{v_1 \dots v_n\}$ and edges $E_{\mathbf{D}} \subset V_{\mathbf{D}} \times V_{\mathbf{D}}$.
- D2. Alphabets.** Each vertex $v_l \in V_{\mathbf{D}}$ has finite alphabet A_l of outputs and finite alphabet $S_l := \prod_{k \in \text{src}(l)} A_k$ of inputs, where $\text{src}(l) = \{v_k | v_k \rightarrow v_l\}$.
- D3. Mechanisms.** Each vertex v_l is equipped with stochastic map $m_l : \mathcal{V}S_l \rightarrow \mathcal{V}A_l$.

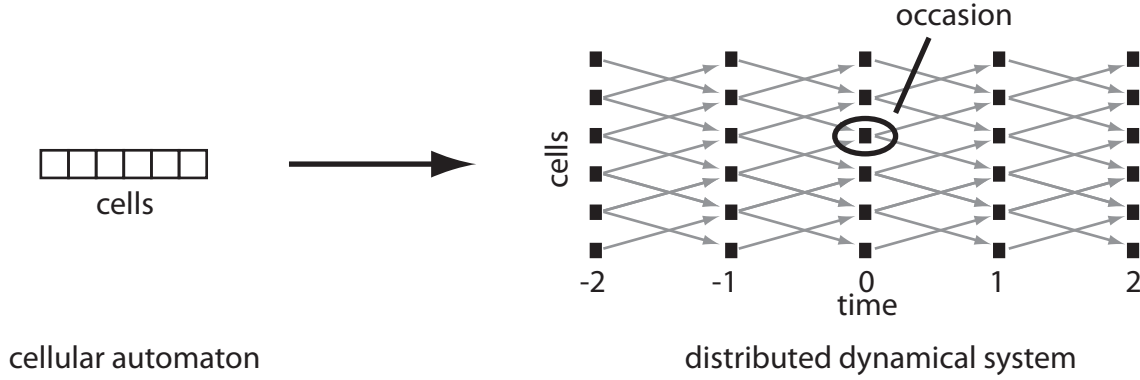


Figure 1: Mapping a cellular automaton to a distributed dynamical system.

Taking any cellular automaton over a finite time interval $[t_\alpha, t_\omega]$ initializing the mechanisms at time t_α with fixed values (initial conditions) or probability distributions (noise sources) yields a distributed dynamical system, see Fig. 1. Each cell of the original automaton corresponds to a series of occasions in the distributed dynamical system, one per time step.

Cells with memory – i.e. whose outputs depend on their neighbors outputs over multiple time steps – receive inputs from occasions more than one time step in the past. If a cell’s mechanism changes (learns) over time then different mechanisms are assigned to the cell’s occasions at different time points.

The sections below investigate the compositional structure of measurements: how they are built out of submeasurements. Technology for tracking subsystems and submeasurements is therefore necessary. We introduce two closely related categories:

Definition 5. The *category of subsystems* $\text{Sys}_{\mathbf{D}}$ of \mathbf{D} is a Boolean lattice with objects given by sets of ordered pairs of vertices $\mathbf{C} \in \underline{2}^{V_{\mathbf{D}} \times V_{\mathbf{D}}}$ and arrows given by inclusions $i_{12} : \mathbf{C}_1 \hookrightarrow \mathbf{C}_2$. The initial and terminal objects are $\perp_{\mathbf{D}} = \emptyset$ and $\top_{\mathbf{D}} = V_{\mathbf{D}} \times V_{\mathbf{D}}$.

Remark 2. Subsystems are defined as ordered pairs of vertices, rather than subgraphs of the directed graph of \mathbf{D} . Pairs of occasions that are not connected by edges are *ineffective*; they do not contribute to the information-processing performed by the system. We include them in the formalism precisely to make their lack of contribution explicit, see Remark 3.

Let $\text{src}(\mathbf{C}) = \{v_k | (v_k, v_l) \in \mathbf{C}\}$ and similarly for $\text{trg}(\mathbf{C})$. Set the input alphabet of \mathbf{C} as the product of the output alphabets of its source occasions $S^{\mathbf{C}} = \prod_{\text{src}(\mathbf{C})} A_k$ and similarly the output alphabet of \mathbf{C} as the product of the output alphabets of its target occasions $A^{\mathbf{C}} = \prod_{\text{trg}(\mathbf{C})} A_l$.

Definition 6. The *category of measuring devices* $\text{Meas}_{\mathbf{D}}$ on \mathbf{D} has objects $\text{Hom}_{\text{Stoch}}(\mathcal{V}A^{\mathbf{C}}, \mathcal{V}S^{\mathbf{C}})$ for

$\mathbf{C} \in \underline{2}^{V_{\mathbf{D}} \times V_{\mathbf{D}}}$. For $\mathbf{C}_1 \hookrightarrow \mathbf{C}_2$ define arrow

$$r_{21} : \text{Hom}(\mathcal{V}A^{\mathbf{C}_2}, \mathcal{V}S^{\mathbf{C}_2}) \rightarrow \text{Hom}(\mathcal{V}A^{\mathbf{C}_1}, \mathcal{V}S^{\mathbf{C}_1})$$

$$\left[\mathcal{V}A^{\mathbf{C}_2} \xrightarrow{T} \mathcal{V}S^{\mathbf{C}_2} \right] \mapsto \left[\mathcal{V}A^{\mathbf{C}_1} \xrightarrow{\pi_A^{\natural}} \mathcal{V}A^{\mathbf{C}_2} \xrightarrow{T} \mathcal{V}S^{\mathbf{C}_2} \xrightarrow{\pi_S^{\natural}} \mathcal{V}S^{\mathbf{C}_1} \right],$$

where π_A and π_S are shorthands for projections as in Definition 1.

The reason for naming $\text{Meas}_{\mathbf{D}}$ the category of measuring devices will become clear in §4 below. The two categories are connected by contravariant functor \mathcal{F} :

Theorem 4 (structure presheaf). *The structure presheaf \mathcal{F} taking*

$$\mathcal{F}_{\mathbf{D}} : \text{Sys}_{\mathbf{D}}^{\text{op}} \rightarrow \text{Meas}_{\mathbf{D}} : \mathbf{C} \mapsto \text{Hom}(\mathcal{V}A^{\mathbf{C}}, \mathcal{V}S^{\mathbf{C}}) \text{ and } i_{12} \mapsto r_{21}$$

satisfies the gluing axiom but has non-unique descent.

Proof: Functor \mathcal{F} is trivially a presheaf since it is contravariant. It is an *interesting* presheaf because the gluing axiom holds.

For gluing we need to show that for any collection $\{\mathbf{C}_j\}_{j=1}^l$ of subsystems and sections $\mathbf{m}_j^{\natural} \in \mathcal{F}_{\mathbf{D}}(\mathbf{C}_j)$ such that $r_{j,ji}(\mathbf{m}_j^{\natural}) = r_{i,ji}(\mathbf{m}_i^{\natural})$ for all i, j there exists section $\mathbf{m}^{\natural} \in \mathcal{F}_{\mathbf{D}}(\bigcup_{j=1}^l \mathbf{C}_j)$ such that $r_i(\mathbf{m}^{\natural}) = \mathbf{m}_i^{\natural}$ for all i . This reduces to finding a conditional distribution that causes diagram

$$\begin{array}{ccc} ? & \longrightarrow & \mathbf{m}_i \\ \downarrow & & \downarrow \\ \mathbf{m}_j & \longrightarrow & \mathbf{m}_{ij} \end{array}$$

in $\text{Meas}_{\mathbf{D}}$ to commute. The vertices are conditional distributions and the arrows are marginalizations, so rewrite as

$$\begin{array}{ccc} ? & \longrightarrow & p(x, y|u, w) \\ \downarrow & & \downarrow \\ p(x, z|v, w) & \longrightarrow & p(x|w), \end{array}$$

where $p(x|w) = \sum_{v,z} p(x, z|v, w) p^{\max H}(v)$ and similarly for the vertical arrow. It is easy to see that

$$p(x, y, z|u, v, w) := \frac{p(x, y|u, w) p(x, z|v, w)}{p(x|w)}$$

satisfies the requirement.

For \mathcal{F} to be a sheaf it would also have to satisfy unique descent: the section satisfying the gluing axiom must not only *exist* for any collection $\{\mathbf{C}_j\}_{j=1}^l$ with compatible restrictions but must also be *unique*. Descent in \mathcal{F} is not unique because there are many distributions satisfying the requirement above: strictly speaking r is a marginalization operator rather than restriction. For example, there are many distributions $p(y, z)$ that marginalize to give $p(y)$ and $p(z)$ besides the product distribution $p(y)p(z)$. ■

The structure presheaf \mathcal{F} depends on the graph structure and alphabets; mechanisms play no role. We now construct a family of sections of \mathcal{F} using the mechanisms of \mathbf{D} 's occasions. Specifically, given a

subsystem $\mathbf{C} \in \text{Sys}_{\mathbf{D}}$, we show how to glue its occasions' mechanisms together to form joint mechanism $\mathbf{m}_{\mathbf{C}}$. The mechanism $\mathbf{m}_{\mathbf{D}} = \mathbf{m}_{\top}$ of the entire system \mathbf{D} is recovered as a special case.

In general, subsystem \mathbf{C} is not isolated: it receives inputs along edges contained in \mathbf{D} but *not* in \mathbf{C} . Inputs along these edges cannot be assigned a fixed value since in general there is no preferred element of A_l . They also cannot be ignored since \mathbf{m}_l is defined as receiving inputs from all its sources. Nevertheless, the mechanism of \mathbf{C} should depend on \mathbf{C} alone. We therefore treat edges not in \mathbf{C} as sources of extrinsic noise by marginalizing with respect to the uniform distribution as in Corollary 3.

For each vertex $v_l \in \text{trg}(\mathbf{C})$ let $S_l^{\mathbf{C}} = \prod_{k \in \text{src}(l) \cap \text{src}(\mathbf{C})} A_k$. We then have projection $\pi_l : \mathcal{V}S_l \rightarrow \mathcal{V}S_l^{\mathbf{C}}$. Define

$$\mathbf{m}_l^{\mathbf{C}} := \left[\mathcal{V}S_l^{\mathbf{C}} \xrightarrow{\pi_l^{\dagger}} \mathcal{V}S_l \xrightarrow{\mathbf{m}_l} \mathcal{V}A_l \right]. \quad (4)$$

It follows immediately that \mathbf{C} is itself a distributed dynamical system defined by its graph, whose alphabets are inherited from \mathbf{D} and whose mechanisms are constructed by marginalizing.

Next, we tensor the mechanisms of individual occasions and glue them together using the diagonal map $\Delta : S^{\mathbf{C}} \rightarrow \prod_{v_l \in \text{trg}(\mathbf{C})} S_l^{\mathbf{C}}$. The diagonal map used here² generalizes $X \xrightarrow{\Delta} X \times X$ and removes redundancies in $\prod_l S_l^{\mathbf{C}}$, which may, for example, include the same source alphabets many times in different factors.

Let mechanism $\mathbf{m}_{\mathbf{C}}$ be

$$\mathbf{m}_{\mathbf{C}} := \left[\mathcal{V}S^{\mathbf{C}} \xrightarrow{\Delta} \bigotimes_{v_l \in \text{trg}(\mathbf{C})} \mathcal{V}S_l^{\mathbf{C}} \xrightarrow{\otimes_{v_l \in \text{trg}(\mathbf{C})} \mathbf{m}_l^{\mathbf{C}}} \mathcal{V}A^{\mathbf{C}} \right]. \quad (5)$$

The dual of $\mathbf{m}_{\mathbf{C}}$ is

$$\mathbf{m}_{\mathbf{C}}^{\dagger} := \left[\mathcal{V}A^{\mathbf{C}} \rightarrow \mathcal{V}S^{\mathbf{C}} \right]. \quad (6)$$

Finally, we find that we have constructed a family of sections of \mathcal{F} :

Definition 7. The quale $\mathbf{q}_{\mathbf{D}}$ is the family of sections of \mathcal{F} constructed in Eqs. (4), (5) and (6)

$$\mathbf{q}_{\mathbf{D}} := \left\{ \mathbf{m}_{\mathbf{C}}^{\dagger} \in \mathcal{F}(\mathbf{C}) = \text{Hom} \left(\mathcal{V}A^{\mathbf{C}}, \mathcal{V}S^{\mathbf{C}} \right) \mid \mathbf{C} \in \text{Sys}_{\mathbf{D}} \right\}.$$

The construction used to glue together the mechanism of the entire system can also be used to construct the mechanism of any subsystem, which provides a window – the quale – into the compositional structure of distributed processes.

4 Measurement

This section adapts Definition 1 to distributed stochastic systems. The first step is to replace elements of state space X with stochastic maps $d_{in} : \mathbb{R} \rightarrow \mathcal{V}S^{\mathbf{D}}$, or equivalently probability distributions on $S^{\mathbf{D}}$, which are the system's inputs. Individual elements of $S^{\mathbf{D}}$ correspond to Dirac distributions.

Second, replace function $f : X \rightarrow \mathbb{R}$ with mechanism $\mathbf{m}_{\mathbf{D}} : \mathcal{V}S^{\mathbf{D}} \rightarrow \mathcal{V}A^{\mathbf{D}}$. Since we are interested in the compositional structure of measurements we also consider submechanisms $\mathbf{m}_{\mathbf{C}}$. However, comparing mechanisms requires that they have the same domain and range, so we extend $\mathbf{m}_{\mathbf{C}}$ to the entire system as follows

$$\mathbf{m}_{\mathbf{C}} = \mathcal{V}S^{\mathbf{D}} \xrightarrow{\pi} \mathcal{V}S^{\mathbf{C}} \xrightarrow{\mathbf{m}_{\mathbf{C}}} \mathcal{V}A^{\mathbf{C}} \xrightarrow{\pi^{\dagger}} \mathcal{V}A^{\mathbf{D}}. \quad (7)$$

²which is surjective in the sense that all rows contain non-zero entries

We refer to the extension as $\mathfrak{m}_{\mathbf{C}}$ by abuse of notation. We extend mechanisms implicitly whenever necessary without further comment. Extending mechanisms in this way maps the quale into a cloud of points in $\text{Hom}(\mathcal{V}A^{\mathbf{D}}, \mathcal{V}S^{\mathbf{D}})$ labeled by objects in $\text{Sys}_{\mathbf{D}}$.

In the special case of the initial object $\perp_{\mathbf{D}}$, define

$$\mathfrak{m}_{\perp} = \mathcal{V}S^{\mathbf{D}} \xrightarrow{\omega} \mathbb{R} \xrightarrow{\omega^{\natural}} \mathcal{V}A^{\mathbf{D}}.$$

Remark 3. Subsystems differing by non-existent edges (Remark 2) are mapped to the same mechanism by this construction, thus making the fact that the edges do not exist explicit within the formalism.

Composing an input with a submechanism yields an output $d_{out} := \mathfrak{m}_{\mathbf{C}} \circ d_{in} : \mathbb{R} \rightarrow \mathcal{V}A^{\mathbf{D}}$, which is a probability distribution on $A^{\mathbf{D}}$. We are now in a position to define

Definition 8. A *measuring device* is the dual $\mathfrak{m}_{\mathbf{C}}^{\natural}$ to the mechanism of a subsystem. An *output* is a stochastic map $d_{out} : \mathbb{R} \rightarrow \mathcal{V}A^{\mathbf{D}}$. A *measurement* is a composition $\mathfrak{m}_{\mathbf{C}}^{\natural} \circ d_{out} : \mathbb{R} \rightarrow \mathcal{V}S^{\mathbf{D}}$.

Recall that stochastic maps of the form $\mathbb{R} \rightarrow \mathcal{V}X$ correspond to probability distributions on X . Outputs as defined above are thus probability distributions on $A^{\mathbf{D}}$, the output alphabet of \mathbf{D} . Individual elements of $A^{\mathbf{D}}$ are recovered as Dirac vectors: $\mathbb{R} \xrightarrow{\delta_a} \mathcal{V}A^{\mathbf{D}}$.

Definition 9. The *effective information* generated by \mathbf{C}_1 in the *context* of subsystem $\mathbf{C}_2 \subset \mathbf{C}_1$ is

$$ei(\mathfrak{m}_{\mathbf{C}_2} \rightarrow \mathfrak{m}_{\mathbf{C}_1}, d_{out}) := H \left[\mathfrak{m}_{\mathbf{C}_1}^{\natural} \circ d_{out} \left\| \mathfrak{m}_{\mathbf{C}_2}^{\natural} \circ d_{out} \right. \right]. \quad (8)$$

The *null context*, corresponding to the empty subsystem $\perp = \emptyset \subset V_{\mathbf{D}} \times V_{\mathbf{D}}$, is a special case where $\mathfrak{m}_{\mathbf{C}}^{\natural} \circ d_{out}$ is replaced by the uniform distribution $\omega_{\mathbf{D}}^{\natural}$ on $S^{\mathbf{D}}$. To simplify notation define

$$ei(\mathfrak{m}_{\mathbf{C}}, d_{out}) := ei(\mathfrak{m}_{\perp} \rightarrow \mathfrak{m}_{\mathbf{C}}, d_{out}).$$

Here, $H[p||q] = \sum_i p_i \log_2 \frac{p_i}{q_i}$ is the Kullback-Leibler divergence or relative entropy [17]. Eq. (8) expands as

$$ei(\mathfrak{m}_{\mathbf{C}_2} \rightarrow \mathfrak{m}_{\mathbf{C}_1}, d_{out}) = \sum_{s \in S^{\mathbf{D}}} \left\langle \mathfrak{m}_{\mathbf{C}_1}^{\natural} \circ d_{out} \left| \delta_s \right. \right\rangle \cdot \log_2 \frac{\left\langle \mathfrak{m}_{\mathbf{C}_1}^{\natural} \circ d_{out} \left| \delta_s \right. \right\rangle}{\left\langle \mathfrak{m}_{\mathbf{C}_2}^{\natural} \circ d_{out} \left| \delta_s \right. \right\rangle}. \quad (9)$$

When $d_{out} = \delta_a$ for some $a \in A^{\mathbf{D}}$ we have

$$ei(\mathfrak{m}_{\mathbf{C}_2} \rightarrow \mathfrak{m}_{\mathbf{C}_1}, \delta_a) = \sum_{s \in S^{\mathbf{D}}} p_{\mathfrak{m}_{\mathbf{C}_1}}(s|a) \cdot \log_2 \frac{p_{\mathfrak{m}_{\mathbf{C}_1}}(s|a)}{p_{\mathfrak{m}_{\mathbf{C}_2}}(s|a)}. \quad (10)$$

Definition 8 requires some unpacking. To relate it to the classical notion of measurement, Definition 1, we consider system $\mathbf{D} = \left\{ v_X \xrightarrow{f} v_Y \right\}$ where the alphabets of v_X and v_Y are the sets $A_{v_X} = X$ and $A_{v_Y} = Y$ respectively, and the mechanism of v_Y is $\mathfrak{m}_Y = \mathcal{V}f$. In other words, system \mathbf{D} corresponds to a single deterministic function $f : X \rightarrow Y$.

Proposition 5 (classical measurement). *The measurement $(\mathcal{V}f)^{\natural} \circ \delta_y$ performed when deterministic function $f : X \rightarrow Y$ outputs y is equivalent to the preimage $f^{-1}(y)$. Effective information is $ei(\mathcal{V}f, \delta_y) = \log_2 \frac{|X|}{|f^{-1}(y)|}$.*


Proof: By Corollary 2 measurement $(\mathcal{V}f)^\natural \circ \delta_y$ is conditional distribution

$$p_{\mathcal{V}f}(x|y) = \begin{cases} \frac{1}{|f^{-1}(y)|} & \text{if } f(x) = y \\ 0 & \text{else.} \end{cases}$$

which generalizes the preimage. Effective information follows immediately. ■

Effective information can be interpreted as quantifying a measurement's precision. It is high if few inputs cause f to output y out of many – i.e. $f^{-1}(y)$ has few elements relative to $|X|$ – and conversely is low if many inputs cause f to output y – i.e. if the output is relatively insensitive to changes in the input. Precise measurements say a lot about what the input could have been and conversely for vague measurements with low ei .

The point of this paper is to develop techniques for studying measurements constructed out of two or more functions. We therefore present computations for the simplest case, distributed system $X \times Y \xrightarrow{g} Z$, in considerable detail. Let \mathbf{D} be the graph



with obvious assignments of alphabets and the mechanism of v_Z as $\mathbf{m}_Z = \mathcal{V}g$. To make the formulas more readable let $\mathbf{m}_{XY} = \mathcal{V}g$, $\mathbf{m}_{X\bullet} = \mathcal{V}g \circ \pi_{XY,X}^\natural$ and $\mathbf{m}_{\bullet Y} = \mathcal{V}g \circ \pi_{XY,Y}^\natural$. We then obtain lattice

$$\begin{array}{ccc} & \top = \mathbf{m}_{XY} & \\ ei(\mathbf{m}_{X\bullet} \rightarrow \mathbf{m}_{XY}) \nearrow & & \nwarrow ei(\mathbf{m}_{\bullet Y} \rightarrow \mathbf{m}_{XY}) \\ \mathbf{m}_{X\bullet} & ei(\mathbf{m}_{XY}) & \mathbf{m}_{\bullet Y} \\ ei(\mathbf{m}_{X\bullet}) \nwarrow & & \nearrow ei(\mathbf{m}_{\bullet Y}) \\ & \mathbf{m}_\perp & \end{array} \quad (11)$$

The remainder of this section and most of the next analyzes measurements in the lattice.

Proposition 6 (partial measurement). *The measurement performed on X when $g : X \times Y \rightarrow Z$ outputs z , treating Y as extrinsic noise, is conditional distribution*

$$p(x|z) = \begin{cases} \frac{|g_{x \times Y}^{-1}(z)|}{|g^{-1}(z)|} & \text{if } g(x, y) = z \text{ for some } y \in Y \\ 0 & \text{else,} \end{cases} \quad (12)$$

where $g_{x \times Y}^{-1}(z) := pr_Y(g^{-1}(z) \cap \{x\} \times Y)$. The effective information generated by the partial measurement is

$$ei(\mathbf{m}_{X\bullet}^\natural, \delta_z) = \log_2 |X| + \sum_{x \in X} p(x|z) \cdot \log_2 p(x|z). \quad (13)$$

Proof: Treating Y as a source of extrinsic noise yields $\mathcal{V}X \xrightarrow{\pi^\natural} \mathcal{V}X \otimes \mathcal{V}Y \xrightarrow{\mathcal{V}g} \mathcal{V}Z$ which takes $\delta_x \mapsto \frac{1}{|Y|} \sum_{y \in Y} \delta_{g(x,y)}$. The dual is

$$\mathbf{m}_{X\bullet}^\natural = \pi_{XY,X} \circ (\mathcal{V}g)^\natural : \delta_z \mapsto \sum_{x \in X} \frac{|g_{x \times Y}^{-1}(z)|}{|g^{-1}(z)|} \cdot \delta_x.$$

The computation of effective information follows immediately. ■

A partial measurement is precise if the preimage $g^{-1}(z)$ has small or empty intersection with $\{x\} \times Y$ for most x , and large intersection for few x .

Propositions 5 and 6 compute effective information of a measurement relative to the null context provided by complete ignorance (the uniform distribution). We can also compute the effective information generated by a measurement in the context of a submeasurement:

Proposition 7 (relative measurement). *The information generated by measurement $X \times Y \xrightarrow{g} Z$ in the context of the partial measurement where Y is unobserved noise, is*

$$ei(\mathfrak{m}_{X^\bullet} \rightarrow \mathfrak{m}_{XY}, \delta_z) = \sum_{x \in X} \frac{g_{x \times Y}^{-1}(z)}{g^{-1}(z)} \log_2 \frac{|Y|}{g_{x \times Y}^{-1}(z)}. \quad (14)$$

Proof: Applying Propositions 5 and 6 obtains

$$ei(\mathfrak{m}_{X^\bullet} \rightarrow \mathfrak{m}_{XY}, \delta_z) = \sum_{(x,y) \in g^{-1}(z)} \frac{1}{|g^{-1}(z)|} \log_2 \left[\frac{1}{|g^{-1}(z)|} \cdot \frac{|g^{-1}(z)| \cdot |Y|}{|g_{x \times Y}^{-1}(z)|} \right]$$

which simplifies to the desired expression. ■

To interpret the result decompose $X \times Y \xrightarrow{g} Z$ into a family of functions $\mathcal{R} = \{Y \xrightarrow{g_{x \times Y}} Z | x \in X\}$ labeled by elements of X , where $g_{x \times Y}(y) := g(x, y)$. The precision of the measurement performed by $g_{x \times Y}$ is $\log_2 \frac{|Y|}{g_{x \times Y}^{-1}(z)}$. It follows that the precision of the relative measurement, Eq. (14), is the expected precision of the measurements performed by family \mathcal{R} taken with respect to the probability distribution $p(x|z) = \frac{g_{x \times Y}^{-1}(z)}{g^{-1}(z)}$ generated by the noisy measurement.

In the special case of $g : X \times Y \rightarrow Z$ relative precision is simply the difference of the precision of the larger and smaller subsystems:

Corollary 8 (comparing measurements).

$$ei(\mathfrak{m}_{X^\bullet} \rightarrow \mathfrak{m}_{XY}, \delta_z) = ei(\mathfrak{m}_{XY}, \delta_z) - ei(\mathfrak{m}_{X^\bullet}, \delta_z)$$

Proof: Applying Propositions 5, 6, 7 and simplifying obtains

$$\begin{aligned} ei(\mathfrak{m}_{XY}, \delta_z) - ei(\mathfrak{m}_{X^\bullet}, \delta_z) &= \log_2 \frac{|X| \cdot |Y|}{|g^{-1}(z)|} - \sum_x \frac{|g_{x \times Y}^{-1}(z)|}{|g^{-1}(z)|} \log_2 \frac{|X| \cdot |g_{x \times Y}^{-1}(z)|}{|g^{-1}(z)|} \\ &= \log_2 \frac{|Y|}{|g^{-1}(z)|} + \sum_{(x,y) \in g^{-1}(z)} \frac{1}{|g^{-1}(z)|} \log_2 \frac{|g^{-1}(z)|}{|g_{x \times Y}^{-1}(z)|} \\ &= ei(\mathfrak{m}_{X^\bullet} \rightarrow \mathfrak{m}_{XY}, \delta_z). \quad \blacksquare \end{aligned}$$

5 Entanglement

The proof of Theorem 4 showed the structure presheaf has non-unique descent, reflecting the fact that measuring devices do not necessarily reduce to products of subdevices. Similarly, as we will see, measurements do not in general decompose into independent submeasurements. Entanglement, γ , quantifies how far a measurement diverges in bits from the product of its submeasurements. It turns out that $\gamma > 0$ is necessary for a system to generate more information than the sum of its components: non-unique descent

thus provides “room at the top” to build systems that perform more precise measurements collectively than the sum of their components.

Entanglement has no direct relation to quantum entanglement. The name was chosen because of a formal resemblance between the two quantities, see Supplementary Information of [7].

Definition 10. Entanglement over partition $\mathcal{P} = \{M_1 \dots M_m\}$ of $\text{src}(\mathbf{m}_{\mathbf{D}})$ is

$$\gamma(\mathbf{m}_{\mathbf{D}}, \mathcal{P}, d_{out}) = H \left[\mathbf{m}_{\mathbf{D}}^{\dagger} \circ d_{out} \parallel \bigotimes_{i=1}^m \pi_j \circ \mathbf{m}_j^{\dagger} \circ d_{out} \right]$$

where $\pi_j : \mathcal{V}S^{\mathbf{D}} \rightarrow \mathcal{V}S^{M_j}$ and $\mathbf{m}_j = \{(k, l) \in \mathbf{m}_{\mathbf{D}} \mid k \in M_j\}$.

Projecting via π_j marginalizes onto the subspace $\mathcal{V}S^{M_j}$. Entanglement thus compares the measurement performed by the entire system with submeasurements over the decomposition of the source occasions into partition \mathcal{P} .

Theorem 9 (effective information decomposes additively when entanglement is zero).

$$\gamma(\mathbf{m}_{\mathbf{D}}, \mathcal{P}, d_{out}) = 0 \implies ei(\mathbf{m}_{\mathbf{D}}, d_{out}) = \sum_{i=1}^m ei(\mathbf{m}_j, d_{out}).$$

Proof: Follows from the observations that (i) $H[p \parallel p_1 \otimes p_2] = 0$ if and only if $p = p_1 \otimes p_2$; (ii) $H[p_1 \otimes p_2 \parallel q_1 \otimes q_2] = H[p_1 \parallel q_1] + H[p_2 \parallel q_2]$; and (iii) the uniform distribution on \mathbf{D} is a tensor of uniform distributions on subsystems of \mathbf{D} . ■

The theorem shows the relationship between effective information and entanglement. If a system generates more information “than it should” (meaning, more than the sum of its subsystems), then the measurements it generates are entangled. Alternatively, only indecomposable measurements can be more precise than the sum of their submeasurements.

We conclude with some detailed computations for $X \times Y \xrightarrow{g} Z$, Diagram (11). Let $\mathcal{P} = \{X|Y\}$.

Theorem 10 (entanglement and effective information for $g : X \times Y \rightarrow Z$).

$$\begin{aligned} \gamma(\mathbf{m}_{XY}, \mathcal{P}, \delta_z) &= \sum_{(x,y) \in g^{-1}(z)} \frac{1}{|g^{-1}(z)|} \log_2 \frac{|g^{-1}(z)|}{|g_{X \times Y}^{-1}(z)| \cdot |g_{X \times Y}^{-1}(z)|} \\ &= ei(\mathbf{m}_{XY}, \delta_z) - ei(\mathbf{m}_{X\bullet}, \delta_z) - ei(\mathbf{m}_{\bullet Y}, \delta_z). \end{aligned}$$

Proof: The first equality follows from Propositions 5 and 6

$$\gamma(\mathbf{m}_{XY}, \mathcal{P}, \delta_z) = \sum_{(x,y) \in g^{-1}(z)} = \sum_{(x,y) \in g^{-1}(z)} \frac{1}{|g^{-1}(z)|} \log_2 \left[\frac{1}{|g^{-1}(z)|} \cdot \frac{|g^{-1}(z)|}{|g_{X \times Y}^{-1}(z)|} \cdot \frac{|g^{-1}(z)|}{|g_{X \times Y}^{-1}(z)|} \right].$$

From the same propositions it follows that $ei(\mathbf{m}_{XY}, \delta_z) - ei(\mathbf{m}_{X\bullet}, \delta_z) - ei(\mathbf{m}_{\bullet Y}, \delta_z)$ equals

$$\begin{aligned} &\log_2 \frac{|X| \cdot |Y|}{|g^{-1}(x)|} - \sum_x \frac{|g_{X \times Y}^{-1}(z)|}{|g^{-1}(z)|} \log_2 \frac{|X| \cdot |g_{X \times Y}^{-1}(z)|}{|g^{-1}(z)|} - \sum_y \frac{|g_{X \times Y}^{-1}(z)|}{|g^{-1}(z)|} \log_2 \frac{|Y| \cdot |g_{X \times Y}^{-1}(z)|}{|g^{-1}(z)|} \\ &= \log_2 \frac{1}{g^{-1}(z)} - \sum_{(x,y) \in g^{-1}(z)} \frac{1}{|g^{-1}(z)|} \cdot \log_2 \frac{|g_{X \times Y}^{-1}(z)|}{|g^{-1}(z)|} \cdot \frac{|g_{X \times Y}^{-1}(z)|}{|g^{-1}(z)|}. \end{aligned}$$

Entanglement quantifies how far the size of the pre-image of $g^{-1}(z)$ deviates from the sizes of its $X \times y$ and $x \times Y$ slices as x and y are varied. ■

By Corollary 8 entanglement also equals $ei(m_{X\bullet} \rightarrow m_{XY}, \delta_z) - ei(m_{\bullet Y}, \delta_z)$. In Diagram (11) entanglement is the vertical arrow minus both arrows at the bottom, or the difference between opposing diagonal arrows. Note that the diagonal arrows from left to right are constructed by adding edge $v_Y \rightarrow v_Z$ to the null system and the subsystem $m_{X\bullet} = \{v_X \rightarrow v_Z\}$ respectively. Entanglement is the difference between the information generated by the diagonal arrows. It quantifies the difference between the information $\{v_Y \rightarrow v_Z\}$ generates in two different contexts.

Corollary 11 (characterization of disentangled set-valued functions). *Function $X \times Y \xrightarrow{g} Z$ performs a disentangled measurement when outputting z iff*

$$g^{-1}(z) = g_{x \times Y}^{-1}(z) \times g_{X \times y}^{-1}(z)$$

for any x, y such that $g(x, y) = z$.

Proof: By Theorem 10 entanglement is zero iff

$$|g^{-1}(z)| = |g_{x \times Y}^{-1}(z)| \cdot |g_{X \times y}^{-1}(z)|$$

for any x, y such that $g(x, y) = z$. This implies the desired result since $g^{-1}(z) \hookrightarrow g_{x \times Y}^{-1}(z) \times g_{X \times y}^{-1}(z)$. ■

Thus, the measurement generated by g is disentangled iff its pre-image $g^{-1}(z)$ satisfies a strong geometric “rectangularity” constraint: that the pre-image decomposes into the product of its $x \times Y$ and $X \times y$ slices for all pairs of slices intersecting $g^{-1}(z)$. The categorizations performed within a disentangled measuring device have nothing to do with each other, so that the device is best considered as two (or more) distinct devices that happen to have been grouped together for the purposes of performing a computation.

Example 4. An XOR-gate $g : X \times Y \rightarrow Z$ outputting 0 generates an entangled measurement. The pre-image is $g^{-1}(0) = \{00, 11\}$ so the XOR-gate generates 1 bit of information about occasions v_X and v_Y . However, the bit is *not localizable*. The measurement generates no information about occasion v_X taken singly, since its output could have been 0 or 1 with equal probability; and similarly for v_Y .

Finally, and unsurprisingly, a function is completely disentangled across all its measurements iff it is a product of two simpler functions:

Corollary 12 (completely disentangled functions are products). *If $X \times Y \xrightarrow{g} Z$ is surjective, then $\gamma(m_{XY}, \mathcal{P}, \delta_z) = 0$ for all $z \in Z$ iff g decomposes into $X \times Y \xrightarrow{g_1 \times g_2} Z_1 \times Z_2 = Z$ for $X \xrightarrow{g_1} Z_1$ and $Y \xrightarrow{g_2} Z_2$.*

Proof: The reverse implication is trivial. In the forward direction, note that $Z = \{g^{-1}(z) | z \in Z\}$ and, by Corollary 11, each pre-image has product structure $g^{-1}(z) = g_{x \times Y}^{-1}(Z) \times g_{X \times y}^{-1}(z)$. Let $Z_1 = \{g_{X \times y}^{-1}(z) | y \in Y \text{ and } z \in Z\}$ and similarly for Z_2 . Define

$$g_1 : X \rightarrow Z_1 : x \mapsto \text{the unique element of form } g_{X \times y}^{-1}(z) \text{ containing it,}$$

and similarly for g_2 . ■

6 Discussion

This paper developed techniques for analyzing the internal structure of distributed measurements. We introduced entanglement, which quantifies the extent to which a measurement is indecomposable. Entanglement can be shown to quantify context-dependence. Moreover, positive entanglement is necessary

for a system to generate more information than the sum of its subsystems. Along the way, we constructed the quale, which geometrically represents the compositional structure of a distributed measurement. The information-theoretic approach developed here is dual, in a precise sense, to the algorithmic perspective on computation. Studying duals m^{\dagger} instead of mechanisms m shifts the focus from *what* the algorithm does to *how* it does it: instead of analyzing rules we analyze functional dependencies.

The intuition driving the paper is that the structure presheaf \mathcal{F} is an information-theoretic analogue of a tangent space. A particle moving in a manifold X defines a vector field – a section of the tangent space to X , which is a sheaf. The tangent vector at a point depends on the particle’s location at “nearby time-points”: it is computed by taking the limit of difference in positions at t and $t+h$ as $h \rightarrow 0$. Similarly, a system performing a measurement generates a quale, a section of the structure presheaf consisting of “nearby counterfactuals”. The quale is computed by applying Bayes’ rule to determine which inputs could have led to the output.³ How far this analogy can be developed remains to be seen.

Entanglement can be loosely considered as an information-theoretic analogue of curvature: the extent to which interactions within a system “warp” sections of \mathcal{F} away from a product structure. A related approach to geometrically analyzing the complexity of interactions was proposed in [4]. In fact, this project began as an attempt to reformulate [7] in terms of sheaf cohomology using ideas from [4]. We failed at the first step since the structure presheaf is not a sheaf. However, the failure was instructive since it is precisely the *obstruction* to forming a sheaf that is of interest since it is the obstruction (entanglement) that quantifies indecomposability and context-dependence, and only systems whose measurements are entangled are able to generate more information than the sum of their subsystems.

References

- [1] Samson Abramsky & Bob Coecke (2009): *Categorical Quantum Mechanics*. In K Engesser, D M Gabbay & D Lehmann, editors: *Handbook of Quantum Logic and Quantum Structures: Quantum Logic*, Elsevier.
- [2] S Amari & H Nagaoka (2000): *Methods of Information Geometry*. *AMS Translations of Mathematical Monographs* 191, Oxford University Press.
- [3] DJ Amit (1989): *Modelling brain function: the world of attractor neural networks*. Cambridge University Press.
- [4] N Ay, E Olbrich, N Bertschinger & J Jost (2006): *A unifying framework for complexity measures of finite systems*. In: *Proceedings of ECCS06, European Complex Systems Society*, Oxford, UK, pp. ECCS06–174.
- [5] David Balduzzi (2011): *Detecting emergent processes in cellular automata with excess information*. preprint .
- [6] David Balduzzi & Giulio Tononi (2008): *Integrated Information in Discrete Dynamical Systems: Motivation and Theoretical Framework*. *PLoS Comput Biol* 4(6), p. e1000091, doi:10.1371/journal.pcbi.1000091.
- [7] David Balduzzi & Giulio Tononi (2009): *Qualia: the geometry of integrated information*. *PLoS Comput Biol* 5(8), p. e1000462, doi:10.1371/journal.pcbi.1000462.
- [8] J S Bell (1990): *Against ‘Measurement’*. *Physics World* August, pp. 33–40.
- [9] ER Berlekamp, JC Conway & RK Guy (1982): *Winning Ways for your Mathematical Plays*. 2, Academic Press.
- [10] P A M Dirac (1958): *The Principles of Quantum Mechanics*. Oxford University Press.
- [11] Tobias Fritz (2009): *A presentation of the category of stochastic matrices*. *arXiv:0902.2554v1* .
- [12] Martin Gardner (1970): *Mathematical Games - The Fantastic Combinations of John Conway’s New Solitaire Game, Life*. *Scientific American* 223, pp. 120–123.

³A counterfactual input is “nearby” to an output if it causes (leads to) that output.

- [13] M Giry (1981): *A categorical approach to probability theory*. In B Banaschewski, editor: *Categorical Aspects of Topology and Analysis*, Springer.
- [14] G 't Hooft (1999): *Quantum gravity as a dissipative deterministic system*. *Classical and Quantum Gravity* 16(10).
- [15] JJ Hopfield (1982): *Neural networks and physical systems with emergent computational properties*. *Proc. Nat. Acad. Sci.* 79, pp. 2554–2558.
- [16] E T Jaynes (1957): *Information theory and statistical mechanics*. *Phys. Rev.* 106(4), pp. 620–630.
- [17] E T Jaynes (1985): *Entropy and Search Theory*. In CR Smith & WT Grandy, editors: *Maximum-entropy and Bayesian Methods in Inverse Problems*, Springer.
- [18] Stuart Kauffman, Carsten Peterson, Björn Samuelsson & Carl Troein (2003): *Random Boolean network models and the yeast transcriptional network*. *Proc Natl Acad Sci U S A* 100(25), pp. 14796–9, doi:10.1073/pnas.2036429100.
- [19] S MacLane & Ieke Moerdijk (1992): *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Springer.
- [20] P Panangaden (1998): *Probabilistic relations*. In C Baier, M Huth, M Kwiatkowska & M Ryan, editors: *PROBMIV'98*, pp. 59–74.