


Pointwise Partial Information Decomposition Using Specificity and Ambiguity Lattices

4th Int. Elec. Conf. on Entropy and Its Applications (ECEA-4)

Conor Finn
Mikhail Prokopenko
Joseph T. Lizier (presenting)

November 10, 2017



Information decomposition

Goal: identify **redundant**, **unique** and **complementary (synergistic)** components of MI

Consider three random variables S_1 , S_2 and T

I Aim: predict T using S_1 and S_2


I Mutual information captures

$$I(T; S_1) = R(T : S_1; S_2) + U(T : S_1 \cap S_2)$$

$$I(T; S_2) = R(T : S_1; S_2) + U(T : S_2 \cap S_1)$$

I Joint mutual information captures

$$I(T; S_1; S_2) = R(T : S_1; S_2) + U(T : S_1 \cap S_2) + U(T : S_2 \cap S_1) + C(T : S_1; S_2)$$



Information decomposition: examples

1. **Unique information** $U(T : S_1 \cap S_2)$

UNQ				RDN			
p	s ₁	s ₂	t	p	s ₁	s ₂	t
1/2	0	0	0	1/2	0	0	0
1/2	1	0	1	1/2	1	1	1

2. **Redundant information** $R(T : S_1; S_2)$

XOR				AND			
p	s ₁	s ₂	t	p	s ₁	s ₂	t
1/4	0	0	0	1/4	0	0	0
1/4	0	1	1	1/4	0	1	0
1/4	1	0	1	1/4	1	0	0
1/4	1	1	0	1/4	1	1	1

3. **Synergistic information** $C(T : S_1; S_2)$

4. In general, all types of information are present

Information decomposition

Goal: identify **redundant**, **unique** and **complementary (synergistic)** components of MI

We have 3 equations in 4 unknowns!

Consider three random variables S_1 , S_2 and T

I Aim: predict T using S_1 and S_2


I Mutual information captures

$$I(T; S_1) = R(T : S_1; S_2) + U(T : S_1 | S_2)$$

$$I(T; S_2) = R(T : S_1; S_2) + U(T : S_2 | S_1)$$

I Joint mutual information captures

$$I(T; S_1; S_2) = R(T : S_1; S_2) + U(T : S_1 | S_2) + U(T : S_2 | S_1) + C(T : S_1; S_2)$$



Partial Information Decomposition (Williams and Beer, 2010)

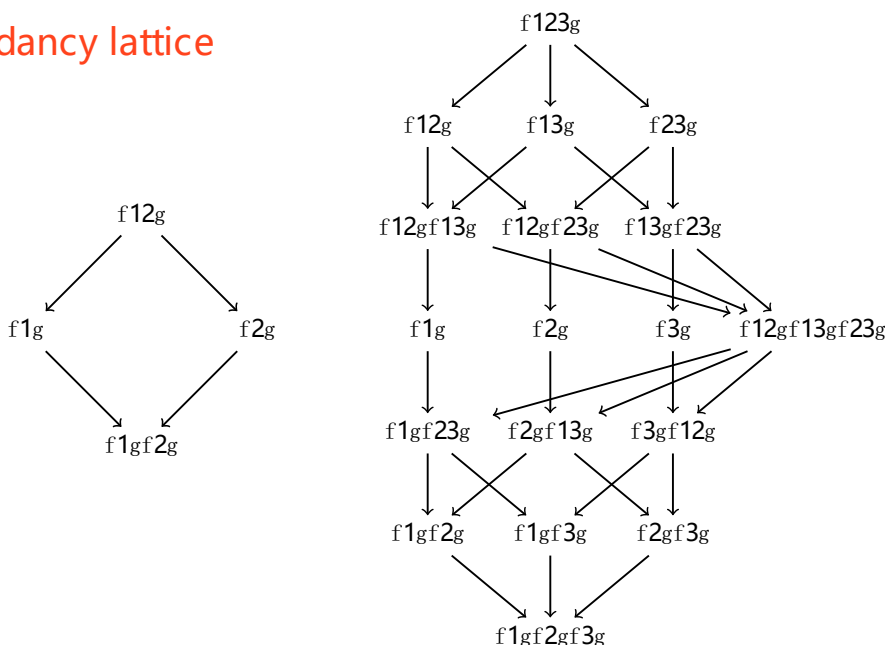
I Axiomatic framework extending this decomposition to arbitrary number of sources

Axioms (PID)

- (1) Symmetry: $R(T : S_1; \dots; S_n)$ is invariant under permutations of the S_i 's
- (2) Monotonicity: $R(T : S_1; \dots; S_n) \geq R(T : S_1; \dots; S_{n-1})$
- (3) Self-redundancy: $R(T : S_i) = I(T; S_i)$

I Yields a **redundancy lattice**

Redundancy lattice



Partial Information Decomposition (Williams and Beer, 2010)

- I Axiomatic framework extending this decomposition to arbitrary number of source

Axioms (PID)

- (1) Symmetry: $R(T : S_1; \dots; S_n)$ is invariant under permutations of the S_i 's
- (2) Monotonicity: $R(T : S_1; \dots; S_n) \geq R(T : S_1; \dots; S_{n-1})$
- (3) Self-redundancy: $R(T : S_i) = I(T; S_i)$

- I Yields a **redundancy lattice**
- I Still no accepted, compatible definition of unique, redundant, or synergistic information

Pointwise information theory

- I From four postulates, Fano (1961) derived the **pointwise** mutual information

$$i(x; y) = \log \frac{p(y|x)}{p(y)} = \log \frac{p(x|y)}{p(x)} \geq 0$$

- I Corollaries: (average) mutual information, pointwise entropy and (Shannon) entropy

Pointwise information decomposition

- I Pointwise decomposition for each realisation

$$i(t; s_1) = r(t : s_1; s_2) + u(t : s_1 n s_2)$$

$$i(t; s_2) = r(t : s_1; s_2) + u(t : s_2 n s_1)$$

$$i(t; s_1 s_2) = r(t : s_1; s_2) + u(t : s_1 n s_2) + u(t : s_2 n s_1) + c(t : s_1; s_2)$$

- I Should be able to take the expectation over all realisations

$$R(T : S_1; S_2) = \mathbb{E} [r(t : s_1; s_2)] \quad U(T : S_1 n S_2) = \mathbb{E} [u(t : s_1 n s_2)]$$

$$C(T : S_1; S_2) = \mathbb{E} [c(t : s_1; s_2)] \quad U(T : S_2 n S_1) = \mathbb{E} [u(t : s_2 n s_1)]$$

- I This should recover the (average) information decomposition

$$I(T; S_1) = R(T : S_1; S_2) + U(T : S_1 n S_2)$$

$$I(T; S_2) = R(T : S_1; S_2) + U(T : S_2 n S_1)$$

$$I(T; S_1 S_2) = R(T : S_1; S_2) + U(T : S_1 n S_2) + U(T : S_2 n S_1) + C(T : S_1; S_2)$$

Motivation: PWUNQ

I Consider PWUNQ:

p	s ₁	s ₂	t	i(t; s ₁)	i(t; s ₂)	i(t; s ₁ s ₂)	r	u ₁	u ₂	c
1/4	0	1	1	0	1	1	0	0	1	0
1/4	1	0	1	1	0	1	0	1	0	0
1/4	0	2	2	0	1	1	0	0	1	0
1/4	2	0	2	1	0	1	0	1	0	0
Expected values				1/2	1/2	1	0	1/2	1/2	0

I According to I_{min} Williams and Beer (2010), \hat{U} of Bertschinger et al. (2014), S_{VK} of Griffith and Koch (2014) and I_{red} of Harder et al. (2013)

$$R = \langle r \rangle = 1/2 \text{ bit} \notin 0 \text{ bit}$$

Pointwise Partial Information Decomposition

Axioms (PPID)

- (1) Symmetry: $r(t : s_1; \dots; s_n)$ is invariant under permutations of the s_i 's
- (2) Monotonicity: $r(t : s_1; \dots; s_n) \geq r(t : s_1; \dots; s_{n-1})$
- (3) Self-redundancy: $r(t : s_i) = i(t; s_i)$

I Problems:


1. Pointwise mutual information is not non-negative
2. Still no clear definition of redundant information

What is pointwise information $i(t; s_1)$?

I The surprise of the posterior compared to the surprise of the prior for event $t; s_1$


$$\text{Prior } p(t) \square! p(t|s_1) \text{ Posterior}$$

I Finn et al. (2017)—this change is ultimately derived from **exclusions**



$$\text{where } t^c = fT \text{ ntg and } s_1^c = fS_1 \text{ ns}_1\text{g}$$

Motivation




I The exclusions differ, but yet
 $i(t; s_1) = i(t; s_2) = 4 \rightarrow 3$ bit


I Pointwise MI is not injective

I Same info \$ same exclusions


Two types of exclusions



Purely informative exclusion



General case



Purely misinformative exclusions

I Idea: split the pointwise MI into two components


$$i(s! t) = i^+(s! t) \square i^\square(s! t)$$

I In Finn et al. (2017) we proved that

(Specificity) $i^+(s_1! t) = h(s_1)$

(Ambiguity) $i^\square(s_1! t) = h(s_1|t)$

Specificity and ambiguity decomposition



$$i(s_1! t) = i(s_2! t) = \log \frac{4}{3} \text{ bit}$$

$$i_+(s_1! t) = \log \frac{4}{3} \text{ bit}$$

$$i_+(s_2! t) = \log \frac{8}{3} \text{ bit}$$

$$i_\square(s_1! t) = 0 \text{ bit}$$

$$i_\square(s_2! t) = 1 \text{ bit}$$

PPID using Specificity and Ambiguity

Axioms (PPID using Specificity and Ambiguity)

- (1) Symmetry: $r^\square(t : s_1; \dots; s_n)$ is invariant under permutations of the s_i 's
- (2) Monotonicity: $r^\square(t : s_1; \dots; s_n) \square r^\square(t : s_1; \dots; s_{n-1})$
- (3) Self-redundancy: $r^\square(t : s_i) = i^\square(t; s_i)$

- I Yields two redundancy lattices: the specificity and ambiguity lattices
- I No longer have the non-negativity problem
- I Still need a measure of redundant information on each lattice

PPID using Specificity and Ambiguity

Additional Axiom (PPID using Specificity and Ambiguity)

- (4) Pointwise event space: $r^\square(t : s_1; \dots; s_n)$ depend only on the size of informative and misinformative exclusions

- I Akin to Assumption ($\square\square$) of Bertschinger et al. (2014)! cannot infer complementary info (synergy) without full joint distribution.
- I Leads us to define the redundant specificity and redundant ambiguity

$$r_{\min}^+(s_1; \dots; s_k | t) = \min_{s_j} h(s_j) \quad r_{\min}^\square(s_1; \dots; s_k | t) = \min_{s_j} h(s_j | t)$$

Example: PWUNQ

p	s ₁	s ₂	t	i ₁ ⁺	i ₁ [□]	i ₂ ⁺	i ₂ [□]	i ₁₂ ⁺	i ₁₂ [□]	r ⁺	u ₁ ⁺	u ₂ ⁺	c ⁺	r [□]	u ₁ [□]	u ₂ [□]	c [□]
1/4	0	1	1	1	1	2	1	2	1	1	0	1	0	1	0	0	0
1/4	1	0	1	2	1	1	1	2	1	1	1	0	0	1	0	0	0
1/4	0	2	2	1	1	2	1	2	1	1	0	1	0	1	0	0	0
1/4	2	0	2	2	1	1	1	2	1	1	1	0	0	1	0	0	0
Expected values				3/2	1	3/2	1	2	1	1	1/2	1/2	0	1	0	0	0

- I Recombining the average specificities and average ambiguities yields the PID

$$R(T : S_1; S_2) = 1 \square 1 = 0 \text{ bit} \quad U(T : S_1 \text{ n } S_2) = 1 \Rightarrow 0 = 1 \Rightarrow \text{bit}$$

$$C(T : S_1; S_2) = 0 \square 0 = 0 \text{ bit} \quad U(T : S_2 \text{ n } S_1) = 1 \Rightarrow 0 = 1 \Rightarrow \text{bit}$$
- I Matches the PPID suggested earlier

Example: XOR

p	s ₁	s ₂	t	i ₁ ⁺	i ₁ ⁻	i ₂ ⁺	i ₂ ⁻	i ₁₂ ⁺	i ₁₂ ⁻	r ⁺	u ₁ ⁺	u ₂ ⁺	c ⁺	r ⁻	u ₁ ⁻	u ₂ ⁻	c ⁻
1/4	0	0	0	1	1	1	1	2	1	1	0	0	1	1	0	0	0
1/4	0	1	1	1	1	1	1	2	1	1	0	0	1	1	0	0	0
1/4	1	0	1	1	1	1	1	2	1	1	0	0	1	1	0	0	0
1/4	1	1	0	1	1	1	1	2	1	1	0	0	1	1	0	0	0
Expected values				1	1	1	1	2	1	1	0	0	1	1	0	0	0

I Recombining the average specificities and average ambiguities yields the PID

$$R(T : S_1; S_2) = 1 \square 1 = 0 \text{ bit} \quad U(T : S_1 \cap S_2) = 0 \square 0 = 0 \text{ bit}$$

$$C(T : S_1; S_2) = 1 \square 0 = 1 \text{ bit} \quad U(T : S_2 \cap S_1) = 0 \square 0 = 0 \text{ bit}$$

I Identifies redundancy due to shared knowledge from Bertschinger et al. (2014)

Further properties

I Can be interpreted in terms of Kelly gambling (e.g. difference in unique information is difference in winnings of two players)

I Is generalisable beyond two sources

I Has a target chain rule (in net):

$$r_{\min} \square S_1; S_2! \ t_1; t_2 \square = r_{\min} \square S_1; S_2! \ t_1 \square + r_{\min} \square S_1; S_2! \ t_2 \square t_1 \square$$

I Provides consistent decomposition of two-bit copy I (f S₁; S₂g S₁; S₂) via target chain rule

I Many echoes of intuition behind previous measures, e.g. (Bertschinger et al., 2014; Ince, 2017; Williams and Beer, 2010), but is fully pointwise and component-wise.

References

I Entropy special issue on Information Decomposition: <http://bit.ly/infodecomp-2017>

I Conor Finn, Mikhail Prokopenko, and Joseph T. Lizier. Pointwise Partial Information Decomposition Using Specificity and Ambiguity Lattices 2017. To appear.

Nils Bertschinger, Johannes Rauh, Eckehard Olbrich, Jürgen Jost, and Nihat Ay. Quantifying unique information. *Entropy*, 16(4):2161–2183, 2014.

Robert Fano. *Transmission of Information*. The MIT Press, 1961.

Conor Finn, Mikhail Prokopenko, and Joseph T. Lizier. Decomposing pointwise information into directed positive and negative components. 2017. To appear.

Virgil Griffith and Christof Koch. Quantifying synergistic mutual information. In Mikhail Prokopenko, editor, *Guided Self-Organization: Inception*, volume 9 of *Emergence, Complexity and Computation*, pages 159–190. Springer Berlin Heidelberg, 2014. ISBN 978-3-642-53733-2.

Malte Harder, Christoph Salge, and Daniel Polani. Bivariate measure of redundant information. *Physical Review E*, 87(1):012130, 2013.

Robin Ince. Measuring multivariate redundant information with pointwise common change in surprisal. *Entropy*, 19(7):318, 2017. doi: 10.3390/e19070318.

Claude E Shannon. A mathematical theory of communication. *Bell Syst. Tech. J.*, 27:623–656, 1948.

Paul L Williams and Randall D Beer. Nonnegative decomposition of multivariate information. *arXiv preprint arXiv:1004.2515*, 2010.