Quantifying Feature Interaction

Zahra Pisheh and Ricardo Vilalta

University of Houston, Department of Computer Science

PATTERN ANALYSIS LAB
finding meaningful patterns in scientific data

Definition

A group of attributes X is involved in a k-way interaction when we cannot reconstruct their relationship merely with ℓ-way interactions, ℓ < k. A good example could be XOR.

Application

- Feature selection Algorithms
- Gene–gene interactions
- Defining Consciousness: The measure of integrated information, is an attempt to a quantify the magnitude of conscious experience.

Measurements

Mutual Information quantifies the amount of information obtained about one feature, through the other feature.

$$MI(X;Y) = \sum_{x}\sum_{y} p(x,y) log2(\frac{p(x,y)}{p(x)p(y)})$$

2-way interaction

interaction between two attributes quantifies with mutual information:

$$I(X:Y) = MI(X;Y)$$

It is Always zero or positive. It is zero if and only if the two attributes are Independent.

3-way Interaction

We could wonder how much uncertainty about A remains after having obtained the knowledge of B and or how C affects the interaction between A and B.

$$I(A;B|C) = \sum_{a,b,c} p(a,b,c) log2(\frac{p(a,b|c)}{p(a|c)p(b|c)})$$

k-way Interaction

Interaction information among attributes can be understood as the amount of information that is common to all the attributes, but not present in any subset.

$$I(S) \triangleq -\sum_{\Gamma \subseteq S} (-1)^{|S|-|\Gamma|} H(\Gamma) = I(S \backslash X|X) - I(S \backslash X), \qquad X \in S$$

I(S)> 0 -> Synergy

I(S)<0 -> Redundancy

Examples

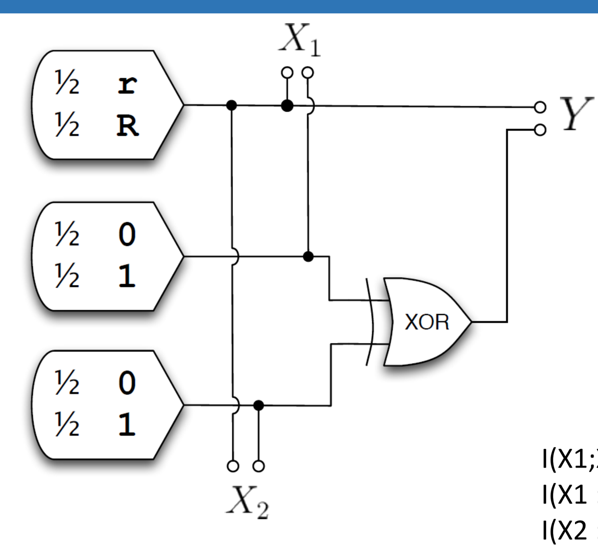

I(X1;X2 :Y) = 2
I(X1 :Y) = 1
I(X2 :Y) = 1

X1	X2	Y
r0	r0	**r0**
r0	r1	**r1**
r1	r0	**r1**
r1	r1	**r0**
R0	R0	**R0**
R0	R1	**R1**
R1	R0	**R1**
R1	R1	**R0**

Figure1.Example RDNXOR. This is the canonical example of redundancy and synergy coexisting.

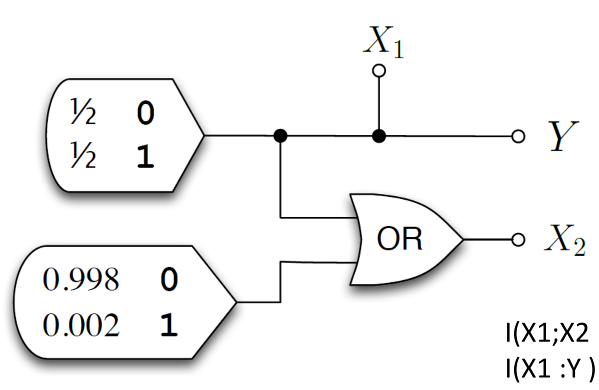

I(X1;X2 :Y) = 1
I(X1 :Y) = 1
I(X2 :Y) = 0:99

X1	X2	Y
0	0	**0**
0	1	**0**
1	1	**1**

Figure2. An ideal measure would detect that all of the information X2 specifies about Y is also specified by X1.

Information Decomposition

Information can be redundant, unique, or synergistic
PI-region represents an irreducible nonnegative slice of the mutual information
I(*X*1...*n* :*Y*) that is either:

1. Redundant. Information carried by a singleton predictor as well as available
somewhere else. For *n* = 2: {1,2}. For *n* = 3: {1,2}, {1,3}, {2,3}, {1,2,3}, {1,23}, {2,13}, {3,12}.

2. Unique. Information carried by exactly one singleton predictor and is available nowhere else. For *n* = 2: {1}, {2}. For *n* = 3: {1}, {2}, {3}.

3. Synergistic. Any and all information in I(*X*1...*n* :*Y*) that is not carried by a singleton predictor. *n* = 2: {12}. For *n* = 3: {13}, {23}, {123}, {12,13},{12,23}, {13,23}, {12,13,23}.

Diagrams

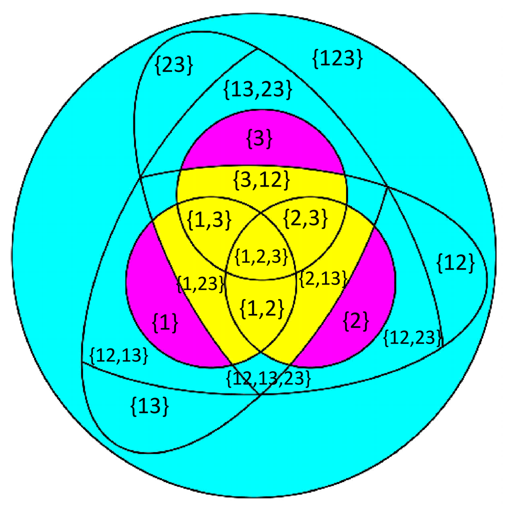

Figure3. Each PI-region represents nonnegative information about Y . A PI-region's color represents whether its information is redundant (yellow), unique (magenta), or synergistic (cyan).

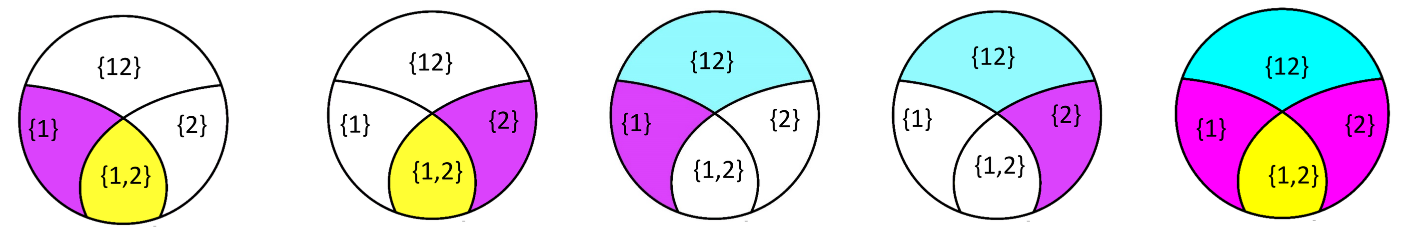

Figure4. PI-diagrams for n = 2 representing standard informational quantities.

Desired Properties

- *(S0) Weak Symmetry*
- *(M0) Weak Monotonicity*
- *(SR) Self-Redundancy*
- *(SR) Self-Redundancy*
- *(M1) Strong Monotonicity*
- *(LP) Local Positivity*
- *(TM) Target Monotonicity*

Path Forward

- Study the decomposition of Mutual information for more than three features
- Better understanding of the feature relationships in the presence of higher-order dependencies and synergy
- Define a robust measure for synergy and redundancy
- Develop a feature selection algorithm considering feature interaction

References

(1) Jakulin, Aleks, and Ivan Bratko. *Analyzing attribute dependencies*. Springer Berlin Heidelberg, 2003.
(2) (Griffith, Virgil. *Quantifying synergistic information*. Diss. California Institute of Technology, 2014.
(3) Griffith, Virgil, and Tracey Ho. "Quantifying Redundant Information in Predicting a Target Random Variable." *Entropy* 17.7 (2015): 4644-4653.