

Unsupervised Learning via Total Correlation Explanation

Greg Ver Steeg

University of Southern California
 Information Sciences Institute
 gregv@isi.edu

Abstract

Learning by children and animals occurs effortlessly and largely without obvious supervision. Successes in automating supervised learning have not translated to the more ambiguous realm of unsupervised learning where goals and labels are not provided. Barlow (1961) suggested that the signal that brains leverage for unsupervised learning is dependence, or redundancy, in the sensory environment. Dependence can be characterized using the information-theoretic multivariate mutual information measure called total correlation. The principle of Total Correlation Explanation (CorEx) is to learn representations of data that “explain” as much dependence in the data as possible. We review some manifestations of this principle along with successes in unsupervised learning problems across diverse domains including human behavior, biology, and language.

1 Introduction

The brain is an information-processing chunk of meat with some amazing properties. Linsker, inventor of the InfoMax principle, made a statement thirty years ago that still rings true today, with some caveats.

A young animal or child perceives and identifies features in its environment in an apparently effortless way. No presently known algorithms even approach this flexible, general-purpose perceptual capability. Discovering the principles that may underlie perceptual processing is important both for neuroscience and for the development of synthetic perceptual systems. [Linsker, 1988]

Perception problems like visual object recognition that were once the exclusive domain of humans are now routinely carried out by computers. Surprisingly, this success did not come from profound understanding of the information processing principles in the brain, but more from brute force scaling of deep supervised optimization algorithms trained on large, labeled image datasets [Bengio *et al.*, 2013]. On the other hand, supervised deep learners are not as flexible or general purpose as Linsker envisioned because they rely heavily on the availability and quality of the training labels. Moreover,

the resulting representations can be brittle and hard to interpret [Szegedy *et al.*, 2014]. Reducing the reliance on labels would greatly broaden the scope of these methods.

Unfortunately, successes in supervised learning have not translated to unsupervised learning [Bengio *et al.*, 2012]. The inherent ambiguity and open-ended nature of the unsupervised problem makes it difficult to brute force a solution [Minsky, 1961]. Learning principles that are general and flexible enough to apply across diverse domains, like human cognition, are needed. Barlow (1961) posited that the signal that the brain leverages for effective learning in the absence of direct supervision are redundancies or dependencies observed in the sensory environment. Investigating how the brain uses this redundancy has motivated several influential ideas [Barlow *et al.*, 1989; Simoncelli and Olshausen, 2001].

This paper describes an information-theoretic approach to formalizing Barlow’s idea called Total Correlation Explanation (CorEx). CorEx constructs a hierarchy of latent factors that progressively explain more dependencies in the observations as measured by multivariate information, also called total correlation. “Explanation” here is meant in the statistical sense that conditioned on the latent factors, Y , the observed random variables X_1, \dots, X_n will be statistically independent. While some learning approaches will learn to memorize even random, independent noise, from the CorEx perspective a lack of relationships in the data implies that there is nothing to learn. We will review the CorEx principle in the context of related information-theoretic methods then demonstrate its power and versatility on a wide range of unsupervised learning problems from human behavior to biology.

2 Information Principles for Learning

Claude Shannon launched the field of information theory with a seminal paper in 1948. Consider a random variable, X , that can take values like x with probability $p(X = x)$ (a coin flip, for example). Shannon defined the entropy of this random variable as $H(X) \equiv \langle -\log p(x) \rangle$ where brackets indicate expectation values over random variables. Shannon showed that this is the unique measure (up to scaling) that satisfies reasonable axioms and bounds the best rate of lossless compression. For a deeper intuition, see [DeDeo, 2015].

Given two random variables, X_1 and X_2 , the mutual information is just the difference between the sum of individual entropies and the entropy of the variables considered jointly

as a single system,

$$I(X_1; X_2) \equiv H(X_1) + H(X_2) - H(X_1, X_2).$$

Shannon demonstrated that this quantity bounds the number of messages that can be reliably sent over a noisy channel.

The definition of mutual information between two parties can be generalized easily to n parties. We refer to this generalization as TC for total correlation following its historical introduction [Watanabe, 1960].

$$\begin{aligned} TC(X_1, \dots, X_n) &\equiv \sum_{i=1}^n H(X_i) - H(X_1, \dots, X_n) \\ &= D_{KL} \left(p(x_1, \dots, x_n) \parallel \prod_{i=1}^n p(x_i) \right). \end{aligned}$$

In the following, we will often shorten $X \equiv X_1, \dots, X_n$. We also wrote TC in terms of KL-divergence, D_{KL} . $TC(X) = 0$ if and only if all the X_i are independent. We can define conditional TC as $TC(X|Y) = D_{KL}(p(x|y) \parallel \prod_i p(x_i|y))$ and this quantity will be zero if and only if X_i are independent conditioned on Y .

Although Shannon gave a general yet precise definition of information, he warned that using everyday words for terminology can be misleading [Shannon, 1956]. If we consider building a representation of observations, X , as $Y = f(X)$, it seems quite intuitive to maximize the “information” that Y has about X by maximizing mutual information, $I(X; Y)$. Indeed, this is the substance of the InfoMax principle [Linsker, 1988; Bell and Sejnowski, 1995]. But the information is maximized if we simply memorize the data. Typically, InfoMax is invoked to maximize mutual information under some simplicity constraints. Still, it seems counter-intuitive that adding more resources should degrade learning performance. This phenomenon has been observed using InfoMax for clustering: more data leads to worse clustering [Ver Steeg *et al.*, 2014].

An opposite point of view is taken in the information bottleneck [Tishby *et al.*, 2000]. In that case, $I(X; Y)$ is actually minimized under the intuition that we should compress X as much as possible. A trade-off is formulated that we should compress Y while maintaining information about some labels, Z . While this approach is natural, it requires labeled data.

Unsupervised approaches have also been motivated from the point of view that the brain has to compress information. In particular, Barlow points out that in a compressed representation neurons should fire independently [Barlow *et al.*, 1989] (since correlations signal redundancy). In this spirit, independent component analysis (ICA) seeks to minimize $TC(Y)$, where Y is a representation of the inputs, X [Hyvärinen and Oja, 2000]. Other work suggests that efficient coding of information in the brain should additionally require that the firing of individual neurons be sparse [Simoncelli and Olshausen, 2001]. While ICA has led to many useful results, the CorEx principle suggests a modification with several benefits.

These learning principles are summarized in Table 1. While this list may look old-fashioned to some readers, these ideas are regularly invoked in modern papers on deep learning [Vincent *et al.*, 2008; Dinh *et al.*, 2014; Tishby and Zaslavsky, 2015; Kolchinsky *et al.*, 2017; Achille and Soatto, 2016;

Table 1: Comparison of information principles for learning.

Method	Objective
InfoMax	$\max_Y I(Y; X)$
Info bottleneck	$\min_Y I(Y; X) - \beta I(Y; Z)$
ICA	$\min_Y TC(Y)$
Common information	$\min_Y TC(X Y)$
CorEx (1 layer)	$\min_Y TC(X Y) + TC(Y)$
(1 layer, alternate form)	$\min_Y \sum_j I(Y_j; X) - \sum_i I(X_i; Y)$

Chen *et al.*, 2016b]. For brevity, other notable ideas have been omitted in this discussion including learning based on “coarse-graining” observations [Goerg and Shalizi, 2012; Wolpert *et al.*, 2014], Jaynes’ maximum entropy principle [Jaynes, 2003], integrated information theory [Oizumi *et al.*, 2014], and common information [Ver Steeg *et al.*, 2017].

2.1 Total Correlation Explanation

One drawback of InfoMax, ICA, and others is that they are intrinsically shallow. For example, a transformation into independent components using ICA does not provide intermediate representations. Even if intermediate steps are used to get the independent components, there is no reason for these terms to be meaningful since they are not represented in the objective. It would be nice to have a hierarchy of abstraction where lower layers capture local relationships and higher layers reflect more global relationships.

One Layer We will begin with a shallow formulation of CorEx, to see how it compares to other learning principles, then we will show that it has a natural hierarchical extension. Assume that each factor, Y_j , (also called neurons or hidden units) is a function of the data, X , $Y_j = f_j(X)$. More generally, it could be drawn from a probabilistic function, $Y_j \sim p(Y_j|X)$. The Y ’s are generated according to a graphical model like the left side of Fig. 1 (with $r = 1$ for now). Under what conditions can we interpret these Y ’s as latent factors that generate the data? In other words, when can we flip the arrows to get the graphical model on the right where the Y ’s generate the dependence in X ? The graphical model on the right side of Fig. 1 is equivalent to a set of conditional independence relationships that can be summarized by saying $TC(X|Y) + TC(Y) = 0$ [Ver Steeg and Galstyan, 2017]. In other words, each layer explains the correlations in the layer below or is independent. This is one way to write the CorEx objective (for one layer).

$$\min_{p(y_j|x)} TC(X|Y) + TC(Y)$$

The objective is non-negative and the global minimum occurs at zero, in which case we can *flip the arrows* to interpret Y ’s as generating the dependence in X . Looking at Table 1, we see that this optimization has an ICA term plus another term that demands that Y ’s make the X ’s conditionally independent. With some manipulation, it can be seen that this optimization is also equivalent to the following optimization [Ver Steeg and Galstyan, 2015].

$$\min_{p(y_j|x)} \sum_j I(Y_j; X) - \sum_i I(X_i; Y)$$




Figure 1: A hierarchical representation where each layer is a probabilistic function of the layer below it. We deduce a condition under which we can flip the arrows and interpret the constructed factors, Y , as generating the dependence in the data.

Again looking at Table 1, this alternate form shows the similarity to the information bottleneck. Instead of requiring labels, Z , we simply compress X into each latent factor Y_j while maintaining relevance about each of our observed variables, X_i . From this point of view, we can view CorEx as a special unsupervised version of the information bottleneck.

Multiple Layers Now we extend to the hierarchical case. We start by building a hierarchy with r layers as on the left side of Fig. 1. For simplicity we define the variables at layer k as Y^k and we define $Y^0 \equiv X$ and $Y^{r+1} = 0$ (a constant). Now let each layer explain dependence in the layer below.

$$\min_{p(y_j^k|x)} \sum_{k=0}^r TC(Y^k|Y^{k+1}) \quad (1)$$

Again, this quantity is non-negative and has a global minimum at zero. At the global minimum, we can flip the arrows and interpret the latent factors as a generative model for the dependence in X . We can again re-write the objective as a sum of bottleneck-like optimizations at each layer of the hierarchy.

$$\max_{p(y_j^k|x)} \sum_{k=0}^r \left(\sum_i I(Y_i^k; Y^{k+1}) - \sum_j I(Y_j^{k+1}; Y^k) \right) \quad (2)$$

An advantage of writing the objective in this second form is that it can be directly plugged into a useful inequality [Ver Steeg and Galstyan, 2015].

$$TC(X) \geq \sum_{k=0}^r \left(\sum_i I(Y_i^k; Y^{k+1}) - \sum_j I(Y_j^{k+1}; Y^k) \right) \quad (3)$$

$TC(X)$ is the amount of dependence observed in the data. For high-dimensional systems, this is hard to estimate. However, the bound in Eq. 3 allows us to solve a hierarchy of optimization problems giving progressively tighter lower bounds on the dependence in X . Because each layer directly contributes to the lower bound on $TC(X)$, we can quantify the value of depth and stop adding layers to the hierarchy when the lower bound stops increasing. Local optima of this non-convex optimization can be found using algorithms with low computational and sample complexity.

CorEx hierarchically decomposes multivariate information in X in terms of contributions from latent factors at each layer of a hierarchy. This decomposition can be viewed as a generalization of the hierarchical decomposition introduced by Watanabe [Watanabe, 1960].

Incremental CorEx Just as PCA has an incremental version where one component is extracted at a time, the incremental version of the CorEx principle is the information sieve [Ver Steeg and Galstyan, 2016]. For the incremental version, we consider the special case where Y is one-dimensional so that $TC(Y) = 0$ and the objective in Eq. 1 reduces to $\min_Y TC(X|Y)$. After learning one Y that optimizes the objective, we transform the data (a kind of information-theoretic orthogonalization) so that we can learn another factor that extracts more dependence. The sieve optimization is a dual formulation of the optimization defining the Wyner common information [Wyner, 1975; Op’t Veld and Gastpar, 2016] and can be viewed as a decomposition of common information [Ver Steeg *et al.*, 2017].

2.2 Implementations of the CorEx Principle

We briefly review implementations of the CorEx principle, including their applicability and limitations. Code is available at <http://github.com/gregversteeg>.

- **CorEx** [Ver Steeg and Galstyan, 2014]: The original implementation is restricted to discrete variables and tree structured latent factors. The functionality is subsumed by other versions.
- **bio_corex** [Ver Steeg and Galstyan, 2015; Pepke and Ver Steeg, 2017]: The most flexible version includes augmentations that were designed for challenges in the biology domain, though there is nothing specific to biology in the implementation. This version handles discrete and continuous variables, overlapping latent factor structure, and missing data. Although CorEx has intrinsically low sample complexity, some biology data is severely under-sampled. This version implements a Bayesian smoothing of the marginal parameter estimates that reduces the appearance of spurious correlations [Pepke and Ver Steeg, 2017]. This version runs quickly for problems with up to thousands of variables.
- **discrete_sieve** [Ver Steeg and Galstyan, 2016]: This is the first implementation of the information sieve for discrete variables. In the discrete case, the information orthogonalization required by the sieve is extremely challenging. The approach is impractical for most real world problems.
- **LinearSieve** [Ver Steeg *et al.*, 2017]: The linear version of the information sieve (for continuous variables) is fast and practical for finding the top components explaining the most correlation in data. Like other incremental methods, repeated application eventually introduces numerical instability.
- **LinearCorEx** [Ver Steeg and Galstyan, 2017]: The fastest, non-sparse version of CorEx assumes latent factors are linear functions of the inputs. Linear CorEx is very effective for covariance estimation and subspace clustering in high-dimensional, under-sampled data.
- **corex_topic** [Reing *et al.*, 2016; Gallagher *et al.*, 2016]: This implementation exploits sparsity for major speed-ups, easily handling hundreds of thousands of variables depending on the sparsity of the data. While this version was developed for topic modeling using binary bag of words data, it can be applied to any binary data.

This implementation includes a semi-supervised option.

3 Applications

The most successful applications of the CorEx principle involve high-dimensional data in complex domains that are difficult to model a priori. In practice, the learned hierarchical representation is used in several ways. First, the structure induces a hierarchical clustering of input variables. Second, the latent factors at each layer constitute a reduced dimensionality representation of the input data. Third, individual factors often disentangle true latent factors of variation in the data. Finally, like PCA, we can rank factors according to their value. While PCA finds components that explain the most *variance*, in CorEx we find factors that explain the most *dependence*.

Social Science Latent factor models like the “Big 5” personality factors are popular in social science because of their simplicity and interpretability. Given data from an online survey with the Big 5 questions, CorEx was able to perfectly reconstruct the Big 5 factors while standard methods failed [Ver Steeg and Galstyan, 2014]. CorEx outperforms traditional factor methods when the number of variables is larger than the number of samples [Ver Steeg *et al.*, 2017], as is increasingly common in social science experiments.

Gene Expression If DNA is the software for our biology, then gene expression tells us which code is running. New sequencing methods allow us to read expression levels for thousands of genes but, unfortunately, only a tiny fraction of these processes are understood. Applying CorEx to gene expression reveals a rich array of strong biological signals. Gene clusters discovered by CorEx exhibit four times more enrichment with respect to known biology in gene ontology databases than standard approaches [Pepke and Ver Steeg, 2017]. The CorEx hierarchy also seems to accurately reflect biological organization. For instance, two low-level clusters, immune cell activation and inflammatory response, are combined in a higher-level group related to immune signaling.

Neuroscience Next generation imaging technology gives detailed views of individual brains but even the largest studies can only afford to image a relatively small number of brains. Small sample sizes with high-dimensional data has led to low statistical power in most studies and a crisis of replication in the field [Button *et al.*, 2013]. Exploratory data analysis using CorEx can help identify only the strongest, most robust dependencies in the data, even with small sample sizes. Early results have identified well-known relationships along with novel, biologically plausible candidate effects that increase predictive power [Madsen *et al.*, 2015; Madsen *et al.*, 2016; Zavaliangos-Petropulu *et al.*, 2017]. Moreover, CorEx outperforms ICA for disentangling spatial modes in imaging data [Ver Steeg *et al.*, 2017].

Language Applied to bag of words vectors of text data, CorEx representations can be interpreted as hierarchical topic models [Ver Steeg and Galstyan, 2014; Hodas *et al.*, 2015]. Latent tree based methods like CorEx outperform LDA [Chen *et al.*, 2016a] on several measures of topic quality. A semi-supervised version further increases the interpretability of topics by allowing us to “anchor” some latent factors to designated words of interest [Reing *et al.*, 2016;

Gallagher *et al.*, 2016].

Finance The stock market exhibits a high degree of dependence and measuring this dependence is important for quantifying risk. Applied to monthly returns on the S&P 500, CorEx discovers a hierarchy of latent factors related to industry sectors [Ver Steeg and Galstyan, 2015]. Using a linear version of CorEx allows us to estimate covariance matrices and this approach outperforms state-of-the-art methods like GLASSO for under-sampled, high-dimensional stock market data [Ver Steeg and Galstyan, 2017].

4 Conclusion

Learning useful representations of observations without supervision across diverse domains is a challenging, unsolved problem. While successes in supervised learning have not immediately translated to comparable results for unsupervised learning, the engineering achievements that have driven those successes now allow us to define flexible, powerful architectures and quickly and easily optimize them under a wide variety objectives [TensorFlow, 2015]. Taking information-theoretic learning principles and applying them within these powerful frameworks has already generated many compelling results [Chen *et al.*, 2016b; Kolchinsky *et al.*, 2017; Achille and Soatto, 2016]. Applying the CorEx principle in the same context is a logical step for future work.

Unlike similar learning principles, CorEx naturally decomposes information in a hierarchical way. This hierarchical structure has numerical advantages since the objective is less reliant on back-propagation for training. Credit for predicting a correct label does not have to be assigned to intermediate representations as in supervised learning [Minsky, 1961], instead each latent factor in the representation has a quantifiable information value within the hierarchical decomposition. Besides the technical advantages, many applications show that individual latent factors learned via CorEx reflect diverse and meaningful structure in real-world datasets. We hope that ongoing CorEx developments continue to accelerate the discovery of knowledge through unsupervised learning.

References

- [Achille and Soatto, 2016] Alessandro Achille and Stefano Soatto. Information dropout: Learning optimal representations through noisy computation. *ArXiv e-prints*, 2016.
- [Barlow *et al.*, 1989] Horace B Barlow, Tej P Kaushal, and Graeme J Mitchison. Finding minimum entropy codes. *Neural Computation*, 1(3):412–423, 1989.
- [Bell and Sejnowski, 1995] Anthony J Bell and Terrence J Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159, 1995.
- [Bengio *et al.*, 2012] Yoshua Bengio, Aaron C Courville, and Pascal Vincent. Unsupervised feature learning and deep learning: A review and new perspectives. *CoRR, abs/1206.5538*, 1, 2012.
- [Bengio *et al.*, 2013] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- [Button *et al.*, 2013] Katherine S Button, John PA Ioannidis, Claire Mokrysz, Brian A Nosek, Jonathan Flint, Emma SJ Robinson, and Marcus R Munafò. Power failure: why small sample size

- undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5):365–376, 2013.
- [Chen *et al.*, 2016a] Peixian Chen, Nevin L Zhang, Leonard KM Poon, and Zhourong Chen. Progressive em for latent tree models and hierarchical topic detection. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [Chen *et al.*, 2016b] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances In Neural Information Processing Systems 29 (NIPS)*, pages 2172–2180, 2016.
- [DeDeo, 2015] Simon DeDeo. Information theory for intelligent people, 2015.
- [Dinh *et al.*, 2014] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- [Gallagher *et al.*, 2016] R.J. Gallagher, K. Reing, D. Kale, and G. Ver Steeg. Anchored correlation explanation: Topic modeling with minimal domain knowledge. *arXiv preprint arXiv:1611.10277*, 2016.
- [Goerg and Shalizi, 2012] G.M. Goerg and C.R. Shalizi. Licors: Light cone reconstruction of states for non-parametric forecasting of spatio-temporal systems. *arXiv:1206.2398*, 2012.
- [Hodas *et al.*, 2015] N. Hodas, G. Ver Steeg, J. Harrison, S. Chikkagoudar, E. Bell, and C. Corley. Disentangling the lexicons of disaster response in twitter. In *The 3rd International Workshop on Social Web for Disaster Management (SWDM'15)*, 2015.
- [Hyvärinen and Oja, 2000] A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural networks*, 13(4):411–430, 2000.
- [Jaynes, 2003] Edwin T Jaynes. *Probability theory: the logic of science*. Cambridge university press, 2003.
- [Kolchinsky *et al.*, 2017] Artemy Kolchinsky, Brendan D Tracey, and David H Wolpert. Nonlinear information bottleneck. *arXiv preprint arXiv:1705.02436*, 2017.
- [Linsker, 1988] Ralph Linsker. Self-organization in a perceptual network. *Computer*, 21(3):105–117, 1988.
- [Madsen *et al.*, 2015] Sarah K. Madsen, Greg Ver Steeg, Adam Mezher, Neda Jahanshad, Talia M. Nir, Xue Hua, Boris A. Gutman, Aram Galstyan, and Paul M. Thompson. Information-theoretic characterization of blood panel predictors for brain atrophy and cognitive decline in the elderly. *IEEE International Symposium on Biomedical Imaging*, 2015.
- [Madsen *et al.*, 2016] Sarah K. Madsen, Greg Ver Steeg, Madelaine Daianu, Adam Mezher, Neda Jahanshad, Talia M. Nir, Xue Hua, Boris A. Gutman, Aram Galstyan, and Paul M. Thompson. Relative value of diverse brain mri and blood-based measures for predicting cognitive decline in the elderly. *SPIE Medical Imaging*, 2016.
- [Minsky, 1961] Marvin Minsky. Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1):8–30, 1961.
- [Oizumi *et al.*, 2014] Masafumi Oizumi, Larissa Albantakis, and Giulio Tononi. From the phenomenology to the mechanisms of consciousness: integrated information theory 3.0. 2014.
- [Op't Veld and Gastpar, 2016] Giel Op't Veld and Michael C Gastpar. Caching gaussians: Minimizing total correlation on the gray-wyner network. In *50th Annual Conference on Information Systems and Sciences (CISS)*, 2016.
- [Pepke and Ver Steeg, 2017] Shirley Pepke and Greg Ver Steeg. Comprehensive discovery of subsample gene expression components by information explanation: therapeutic implications in cancer. *BMC medical genomics*, 10(1):12, 2017.
- [Reing *et al.*, 2016] Kyle Reing, David C Kale, Greg Ver Steeg, and Aram Galstyan. Toward interpretable topic discovery via anchored correlation explanation. In *Proceedings of the 2016 ICML Workshop on Human Interpretability in Machine Learning (WHI 2016)*, number arXiv:1606.07043, 2016.
- [Shannon, 1948] C.E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 1948.
- [Shannon, 1956] Claude Shannon. The bandwagon (edtl.). *IRE Transactions on Information Theory*, 1(2):3, 1956.
- [Simoncelli and Olshausen, 2001] Eero Simoncelli and Bruno Olshausen. Natural image statistics and neural representation. *Annu. Rev. Neurosci.*, 24, 2001.
- [Szegedy *et al.*, 2014] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *ICLR*, 2014.
- [TensorFlow, 2015] TensorFlow. Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [Tishby and Zaslavsky, 2015] Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *Information Theory Workshop (ITW), 2015 IEEE*. IEEE, 2015.
- [Tishby *et al.*, 2000] Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv:physics/0004057*, 2000.
- [Ver Steeg and Galstyan, 2014] G. Ver Steeg and A. Galstyan. Discovering structure in high-dimensional data through correlation explanation. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [Ver Steeg and Galstyan, 2015] Greg Ver Steeg and Aram Galstyan. Maximally informative hierarchical representations of high-dimensional data. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015. <http://arxiv.org/abs/1410.7404>.
- [Ver Steeg and Galstyan, 2016] Greg Ver Steeg and Aram Galstyan. The information sieve. In *International Conference on Machine Learning (ICML)*, 2016.
- [Ver Steeg and Galstyan, 2017] Greg Ver Steeg and Aram Galstyan. Low complexity gaussian latent factor models and a blessing of dimensionality. *arXiv:1706.03353 [stat.ML]*, 2017.
- [Ver Steeg *et al.*, 2014] Greg Ver Steeg, Aram Galstyan, Fei Sha, and Simon DeDeo. Demystifying information-theoretic clustering. In *International Conference on Machine Learning*, 2014.
- [Ver Steeg *et al.*, 2017] Greg Ver Steeg, Shuyang Gao, Kyle Reing, and Aram Galstyan. Sifting common information from many variables. *IJCAI*, 2017.
- [Vincent *et al.*, 2008] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103. ACM, 2008.
- [Watanabe, 1960] Satoshi Watanabe. Information theoretical analysis of multivariate correlation. *IBM Journal of research and development*, 4(1):66–82, 1960.
- [Wolpert *et al.*, 2014] David H Wolpert, Joshua A Grochow, Eric Libby, and Simon DeDeo. A framework for optimal high-level descriptions in science and engineering—preliminary report. *arXiv preprint arXiv:1409.7403*, 2014.
- [Wyner, 1975] Aaron D Wyner. The common information of two dependent random variables. *Information Theory, IEEE Transactions on*, 21(2):163–179, 1975.
- [Zavaliangos-Petropulu *et al.*, 2017] Artemis Zavaliangos-Petropulu, Emily L Dennis, Greg Ver Steeg, Talin Babikian, Richard Mink, Christopher Babbitt, Jeffrey Johnson, Christopher C Giza, Robert F Asarnow, and Paul M Thompson. Variable clustering reveals associations between subcortical brain volume and cognitive changes in pediatric traumatic brain injury. In *12th International Symposium on Medical Information Processing and Analysis*. International Society for Optics and Photonics, 2017.