



Distributed information-theoretic clustering

Georg Pichler, Pablo Piantanida, Gerald Matz

► To cite this version:

Georg Pichler, Pablo Piantanida, Gerald Matz. Distributed information-theoretic clustering. Information and Inference, 2022, 11 (1), pp.137-166. 10.1093/imaiai/iaab007 . hal-04137003

HAL Id: hal-04137003

<https://hal.science/hal-04137003>

Submitted on 22 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed information-theoretic clustering

GEORG PICHLER*,

Institute of Telecommunications, TU Wien, Vienna, Austria.

*Corresponding author: georg.pichler@ieee.org

PABLO PIANTANIDA

*Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des Signaux et Systèmes,
Gif-sur-Yvette, France and Montreal Institute for Learning Algorithms (Mila), Université de
Montréal, QC, Canada.*

pablo.piantanida@centralesupelec.fr

AND

GERALD MATZ

Institute of Telecommunications, TU Wien, Vienna, Austria.

gerald.matz@tuwien.ac.at

[Received on November 29, 2021]

We study a novel multi-terminal source coding setup motivated by the biclustering problem. Two separate encoders observe two i.i.d. sequences X^n and Y^n , respectively. The goal is to find rate-limited encodings $f(x^n)$ and $g(y^n)$ that maximize the mutual information $I(f(X^n); g(Y^n))/n$. We discuss connections of this problem with hypothesis testing against independence, pattern recognition, and the information bottleneck method. Improving previous cardinality bounds for the inner and outer bounds allows us to thoroughly study the special case of a binary symmetric source and to quantify the gap between the inner and the outer bound in this special case. Furthermore, we investigate a multiple description (MD) extension of the Chief Operating Officer (CEO) problem with mutual information constraint. Surprisingly, this MD-CEO problem permits a tight single-letter characterization of the achievable region.

Keywords: source coding; mutual information; information bottleneck; CEO problem

1. Introduction

The recent decades witnessed a rapid proliferation of digital data in a myriad of repositories such as internet fora, blogs, web applications, news, emails and the social media bandwagon. A significant part of this data is unstructured and it is thus hard to extract relevant information. This results in a growing need for a fundamental understanding and efficient methods for analyzing data and discovering valuable and relevant knowledge from it in the form of structured information.

When specifying certain hidden (unobserved) features of interest, the problem then consists of extracting those relevant features from a measurement, while neglecting other, irrelevant features. Formulating these idea in terms of lossy source compression [44], we can quantify the complexity of the representations via its compression rate and the fidelity via the information provided about specific (unobserved) features.

In this paper, we introduce and study the distributed clustering problem from a formal information-theoretic perspective. Given correlated samples X_1, X_2 observed at two different encoders, the aim is

to extract a description from each sample such that the descriptions are maximally informative about each other. In other words, each encoder tries to find a (lossy) description $W_j = f_j(X_j^n)$ of its observation X_j^n subject to a complexity requirement (coding rate), maximizing the mutual information $I(W_1; W_2)$. Our goal is to characterize the optimal tradeoff between the *relevance* (mutual information between the descriptions) and the *complexity* of those descriptions (encoding rate).

1.1 Related work

Biclustering (or *co-clustering*) was first explicitly considered by Hartigan [26] in 1972. A historical overview of biclustering including additional background can be found in [33] Section 3.2.4]. In general, given an $S \times T$ data matrix (a_{st}) , the goal of a biclustering algorithm [32] is to find partitions $\mathcal{B}_k \subseteq \{1, \dots, S\}$ and $\mathcal{C}_l \subseteq \{1, \dots, T\}$, $k = 1 \dots K$, $l = 1 \dots L$ such that all the elements of the ‘biclusters’ $(a_{st})_{s \in \mathcal{B}_k, t \in \mathcal{C}_l}$ are homogeneous in a certain sense. The measure of homogeneity of the biclusters depends on the specific application. The method received renewed attention when Cheng and Church [6] applied it to gene expression data. Many biclustering algorithms have been developed since (e.g., see [48] and the references therein). An introductory overview of clustering algorithms for gene expression data can be found in the lecture notes [45]. The *information bottleneck* (IB) method, which can be viewed as a uni-directional information-theoretic variant of biclustering, was successfully applied to gene expression data as well [46].

In 2003, Dhillon *et al.* [10] adopted an information-theoretic approach to biclustering. They used mutual information to characterize the quality of a biclustering. Specifically, for the special case when the underlying matrix represents the joint probability distribution of two discrete random variables X and Y , i.e., $a_{st} = P\{X = s, Y = t\}$, their goal was to find clustering functions $f: \{1, \dots, S\} \rightarrow \{1, \dots, K\}$ and $g: \{1, \dots, T\} \rightarrow \{1, \dots, L\}$ that maximize $I(f(X); g(Y))$ for specific K and L . This idea was successfully employed in numerous research papers since, e.g., [20, 30, 35, 47], where mutual information is typically estimated from samples.

In the present work, we investigate a theoretical extension of the approach in [10], where we consider blocks of n i.i.d. sources and S_n, K_n, T_n , and L_n scale exponentially in the blocklength n . The resulting *information-theoretic biclustering problem* turns out to be equivalent to hypothesis testing against independence with multi-terminal data compression [23] and to a pattern recognition problem [52]. Both these problems are not yet solved in general (for a survey on the hypothesis testing problem, see [24]). The pattern recognition problem has been extensively studied on doubly symmetric binary and jointly Gaussian sources.

A special case of the information-theoretic biclustering problem is given by the *IB problem*, studied in [18], based on the IB method [49]. This problem is solved in terms of a single-letter characterization and is known to be equivalent to source coding under logarithmic loss. A generalization to multiple terminals, the CEO problem under logarithmic loss [8], is currently only solved under specific Markov constraints.

1.2 Contributions

The aim of this paper is to characterize the achievable region of the information-theoretic biclustering problem, its extensions and special cases, and connect them to known problems in network information theory. This problem is fundamentally different from ‘classical’ distributed source coding problems like distributed lossy compression [13] Chapter 12]. Usually, one aims at reducing redundant information, i.e., information that is transmitted by multiple encoders, as much as possible, while still guaranteeing correct decoding. By contrast, in the biclustering problem we are interested in maximizing this very redundancy.

In this sense, it is complementary to conventional distributed source coding and requires adapted proof techniques.

More specifically, the main contributions are as follows.

- We formally prove the equivalence of the information-theoretic biclustering, the hypothesis testing [23], and the pattern recognition problem [52] (Theorem 3.4) and connect it to the IB problem [18, 49] (Proposition 6.2).
- We extensively study the doubly symmetric binary source (DSBS) as a special case (Section 5). In order to perform this analysis, we require stronger cardinality bounds than the ones usually obtained using the convex cover method [13 Appendix C].
- We are able to improve upon the state-of-the-art cardinality bounding techniques by combining the convex cover method with the perturbation method [13 Appendix C] and leveraging ideas similar to [37], which allow us to restrict our attention to the extreme points of the achievable region. The resulting bounds (Proposition 4.3) allow for the use of binary auxiliary random variables in the case of binary sources.
- We show that [52 Conjecture 1] does not hold (Propositions 5.3 and 5.4).
- Based on a weaker conjecture (Conjecture 5.2), we argue that there is indeed a gap between the outer and the inner bound for a DSBS. (Conjecture 5.1).
- We propose an extension of the CEO problem under an information constraint, studied in [8], which requires *multiple description* (MD) coding [14] (see [21] for applications) to account for the possibility that descriptions are not delivered. Using tools from submodularity theory and convex analysis, we are able to provide a complete single-letter characterization of the resulting achievable region (Theorem 7.3), which has the remarkable feature that it allows to exploit rate that is in general insufficient for successful typicality decoding.

1.3 Notation and conventions

For a total order \sqsubset on a set \mathcal{E} (cf. [41 Definition 1.5]) and $e \in \mathcal{E}$ we will use the notation $\sqsupset e := \{e' \in \mathcal{E} : e' \sqsupset e\}$ and accordingly for \sqsubset, \sqsubset and \sqsubseteq . For example, given the total order \sqsubset on $\{1, 2, 3\}$ with $3 \sqsubset 1 \sqsubset 2$, we have $\sqsupset 3 = \{1, 2\}$, $\sqsupset 1 = \{2\}$ and $\sqsupset 2 = \emptyset$.

We will use the shorthand $[l:k] := \{l, l+1, \dots, k-1, k\}$. The notation $\mathbb{1}_{\mathcal{A}}$, $\overline{\mathcal{A}}$, $\text{conv}(\mathcal{A})$, and $|\mathcal{A}|$ is used for the indicator, topological closure, convex hull, and cardinality of a set \mathcal{A} , respectively. When there is no possibility of confusion we identify singleton set with its element, e.g., we write $\{1, 2, 3\} \setminus 1 = \{2, 3\}$. Let \mathbb{R}_+ and \mathbb{R}_- be the set of non-negative and non-positive reals, respectively.

We denote random quantities and their realizations by capital and lowercase letters, respectively. Furthermore, vectors are indicated by bold-face type and have length n , if not otherwise specified. Random variables are assumed to be supported on finite sets and unless otherwise specified, the same letter is used for the random variable and for its support set, e.g., Y takes values in \mathcal{Y} and X_3 takes values in \mathcal{X}_3 . Given a random variable X , we write p_X for its probability mass function (pmf), where the subscript might be omitted if there is no ambiguity. The notation $X \sim p$ indicates that X is distributed according to p and $X \sim \mathcal{B}(p)$ and $Y \sim \mathcal{U}(\mathcal{Y})$ denote a Bernoulli distributed random variable X with parameter $p \in [0, 1]$ and a uniformly distributed random variable Y on its support set \mathcal{Y} . We use $\mathbb{E}[X]$ and $P\{\mathcal{A}\}$ for the expectation of the random variable X and the probability of an event \mathcal{A} , respectively.




Figure 1: Biclustering of two memoryless sources.

Subscripts indicate parts of vectors, e.g., $\mathbf{x}_{\mathcal{A}} := (x_i)_{i \in \mathcal{A}}$ for a vector $\mathbf{x} = (x_1, x_2, \dots, x_n)$ and $\mathcal{A} \subseteq [1:n]$. We further use the common notation $\mathbf{x}_i^j := \mathbf{x}_{\{i, \dots, j\}}$, $\mathbf{x}^j := \mathbf{x}_1^j$. If a vector is already carrying a subscript, it will be separated by a comma, e.g., $\mathbf{x}_{3,1}^5 = (\mathbf{x}_3)_1^5 = (\mathbf{x}_3)^5$. Let $\mathbf{0}$ denote the all-zeros vector and $\mathbf{e}_i = (e_{i,1}, e_{i,2}, \dots, e_{i,n}) \in \mathbb{R}^n$ the i th canonical base vector, i.e., $e_{i,j} = \mathbb{1}_i(j)$. We use the notation of [9, Chapter 2] for information-theoretic quantities. In particular, given random variables (X, Y, Z) and pmfs p and q , $H(X)$, $H(X|Y)$, $I(X; Y)$, $I(X; Y|Z)$, and $D(p||q)$ denote entropy, conditional entropy, mutual information, conditional mutual information, and Kullback-Leibler divergence, respectively. All logarithms in this paper are to base e and therefore all information theoretic quantities are measured in nats. The notation $h_2(p) := -p \log p - (1-p) \log(1-p)$ is used for the binary entropy function, $a * b := a(1-b) + (1-a)b$ is the binary convolution operation and the symbol \oplus denotes binary addition. The notation $X \ominus Y \ominus Z$ indicates that X , Y , and Z form a Markov chain in this order and $X \perp Y$ denotes that X and Y are independent random variables. Slightly abusing notation we consider \emptyset to be a degenerate random variable that is almost surely a constant, such that, e.g., $X \ominus \emptyset \ominus Y \Leftrightarrow X \perp Y$. To simplify the presentation (cf. [13]) when generating codebooks, we will assume that the codebook size is an integer. We will use superscript to indicate that a relation follows from a specific equation. For example, the inequality $a \stackrel{(42)}{\leq} b$ follows from equation (42).

2. Problem statement

In this section we will introduce the *information-theoretic biclustering problem* (or *biclustering problem* for short) with two sources and provide bounds for its achievable region. A schematic overview of the problem is presented in Fig. 1. Let (X, Y) be two random variables. The random vectors (\mathbf{X}, \mathbf{Y}) consist of n i.i.d. copies of (X, Y) . Given a block length $n \in \mathbb{N}$ and coding rates $R_1, R_2 \in \mathbb{R}_+$, an (n, R_1, R_2) -code consists of two functions $f: \mathcal{X}^n \rightarrow \mathcal{M}_1$ and $g: \mathcal{Y}^n \rightarrow \mathcal{M}_2$ such that the finite sets \mathcal{M}_k satisfy $\log |\mathcal{M}_k| \leq nR_k$, $k \in \{1, 2\}$. Thus, the coding rates R_1 and R_2 limit the complexity of the encoders. In contrast to rate-distortion theory, we do not require a specific distortion measure; rather, we quantify the quality of a code in pure information-theoretic terms, namely via mutual information. The idea is to find functions f and g that extract a compressed version of the common randomness in the observed data \mathbf{X} and \mathbf{Y} . To this end, we use the normalized mutual information $I(f(\mathbf{X}); g(\mathbf{Y}))/n$ to quantify the relevance of the two encodings.

DEFINITION 2.1 A triple (μ, R_1, R_2) is *achievable* if, for some $n \in \mathbb{N}$, there exists an (n, R_1, R_2) -code

(f, g) such that

$$\frac{1}{n} \mathbb{I}(f(\mathbf{X}); g(\mathbf{Y})) \geq \mu. \quad (2.1)$$

The achievable region $\overline{\mathcal{R}}$ is defined as the closure of the set \mathcal{R} of achievable triples.

REMARK 2.1 Note that a standard time-sharing argument shows that $\overline{\mathcal{R}}$ is a convex set (cf. [13, Section 4.4]).

We also point out that stochastic encodings cannot enlarge the achievable region as any stochastic encoding can be represented as the convex combination of deterministic encodings and $\overline{\mathcal{R}}$ is convex.

3. Equivalent problems

The biclustering problem turns out to be equivalent to a hypothesis testing and a pattern recognition problem. In this section we will clarify this equivalence by showing that the multi-letter regions agree. These equivalences will provide us with the achievability of \mathcal{R}_* , the ‘multi-letter’ region \mathcal{R}_* of the biclustering problem.

DEFINITION 3.1 Let \mathcal{R}_* be the set of triples (μ, R_1, R_2) such that there exist $n \in \mathbb{N}$ and random variables U, V satisfying $U \dashv\!\!\!\dashv \mathbf{X} \dashv\!\!\!\dashv \mathbf{Y} \dashv\!\!\!\dashv V$ and

$$nR_1 \geq \mathbb{I}(U; \mathbf{X}), \quad (3.1)$$

$$nR_2 \geq \mathbb{I}(V; \mathbf{Y}), \quad (3.2)$$

$$n\mu \leq \mathbb{I}(U; V). \quad (3.3)$$

Next, we consider the hypothesis testing problem with data compression when testing against independence [23, Section 6] and the pattern recognition problem [52]. For completeness sake we briefly describe the problem setups.

DEFINITION 3.2 (Hypothesis testing against independence) Given the potentially dependent sources (X, Y) , define the independent random variables $(\bar{X}, \bar{Y}) \sim p_X \times p_Y$. An (n, R_1, R_2) hypothesis test consists of an (n, R_1, R_2) -code (f_n, g_n) and a set $\mathcal{A}_n \subseteq \mathcal{M}_1 \times \mathcal{M}_2$, where \mathcal{M}_1 and \mathcal{M}_2 are the ranges of f_n and g_n , respectively. The type I and type II error probabilities of $(f_n, g_n, \mathcal{A}_n)$ are defined as $\alpha_n := \mathbb{P}\{(f_n(\mathbf{X}), g_n(\mathbf{Y})) \in \mathcal{A}_n\}$ and $\beta_n := \mathbb{P}\{(f_n(\bar{\mathbf{X}}), g_n(\bar{\mathbf{Y}})) \notin \mathcal{A}_n\}$, respectively. A triple (μ, R_1, R_2) is *HT-achievable* if, for every $\varepsilon > 0$, there is a sequence of (n, R_1, R_2) hypothesis tests $(f_n, g_n, \mathcal{A}_n)$, $n \in \mathbb{N}$ such that

$$\lim_{n \rightarrow \infty} \alpha_n \leq \varepsilon, \quad (3.4)$$

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log \beta_n \geq \mu. \quad (3.5)$$

Let \mathcal{R}_{HT} denote the set of all HT-achievable triples.

DEFINITION 3.3 (Pattern recognition) Let $(\mathbf{X}(i), \mathbf{Y}(i))$ be n i.i.d. copies of (X, Y) , independently generated for each $i \in \mathbb{N}$. A triple (μ, R_1, R_2) is said to be *PR-achievable* if, for any $\varepsilon > 0$, there is some $n \in \mathbb{N}$, such that there exists an (n, R_1, R_2) -code (f, g) and a function $\phi: (\mathcal{M}_1)^{e^{n\mu}} \times \mathcal{M}_2 \rightarrow [1: e^{n\mu}]$ with

$$\mathbb{P}\{W = \phi(C, g(\mathbf{Y}(W)))\} \geq 1 - \varepsilon, \quad (3.6)$$

where $C := f(\mathbf{X}(i))_{i \in [1: e^{n\mu}]}$ is the compressed codebook and $(\mathbf{X}(i), \mathbf{Y}(i))_{i \in \mathbb{N}} \perp W \sim \mathcal{U}([1: e^{n\mu}])$. Let \mathcal{R}_{PR} denote the set of all PR-achievable triples.

REMARK 3.1 The variant of the inner bound for the pattern recognition problem stated in [52] Theorem 1] is flawed. To see this, note that (using the notation of [52]) the point $(R_x = 0, R_y = b, R_c = b)$ is contained in \mathcal{R}_{in} (choose $U = V = \emptyset$) for any $b > 0$ even if the random variables X and Y are independent. But this point is clearly not achievable in general. However, the region \mathcal{R}_{in} defined in the right column of [52] p. 303] coincides with our findings and the proof given in [52] Appendix A] holds for this region.

The biclustering, hypothesis testing and pattern recognition problems are equivalent in the sense that their ‘multi-letter’ regions agree. The proof of this result is given in Appendix A.1

THEOREM 3.4 $\overline{\mathcal{R}} = \overline{\mathcal{R}_*} = \overline{\mathcal{R}_{\text{HT}}} = \overline{\mathcal{R}_{\text{PR}}}$.

4. Bounds on the achievable region

The following inner and outer bound on the achievable region follow from the corresponding results on the hypothesis testing and pattern recognition problems.

THEOREM 4.1 We have $\mathcal{R}_i \subseteq \overline{\mathcal{R}}$ where

$$\mathcal{R}_i := \bigcup_{U, V} \{(\mu, R_1, R_2) : R_1 \geq I(U; X), R_2 \geq I(V; Y), \mu \leq I(U; V)\}, \quad (4.1)$$

with auxiliary random variables U, V satisfying $U \dashv\dashv X \dashv\dashv Y \dashv\dashv V$.

THEOREM 4.2 We have $\mathcal{R} \subseteq \mathcal{R}_o$, where

$$\mathcal{R}_o := \bigcup_{U, V} \{(\mu, R_1, R_2) : R_1 \geq I(U; X), R_2 \geq I(V; Y), \mu \leq I(U; X) + I(V; Y) - I(UV; XY)\}, \quad (4.2)$$

with U and V any pair of random variables satisfying $U \dashv\dashv X \dashv\dashv Y$ and $X \dashv\dashv Y \dashv\dashv V$.

Using Theorem 3.4 Theorem 4.1 follows either from [23] Corollary 6], or from [52] Appendix A]. Theorem 4.2 follows from [52] Appendix B] using Theorem 3.4

The main differences between the outer and the inner bound lie in the Markov conditions, a phenomenon that also occurs with Berger–Tung type bounds (cf. [50] Chapter 7] or [13] Section 12.2]). Note that \mathcal{R}_o and \mathcal{R}_i would coincide if the Markov condition $U \dashv\dashv X \dashv\dashv Y \dashv\dashv V$ were imposed in the definition of \mathcal{R}_o . The region \mathcal{R}_o is convex since a time-sharing variable can be incorporated into U and V . The inner bound \mathcal{R}_i , however, can be improved by convexification.

Numerical evaluation of \mathcal{R}_o and \mathcal{R}_i requires the cardinalities of the auxiliary random variables to be bounded. We therefore complement Theorems 4.1 and 4.2 with the following result, whose proof is provided in Appendix A.2

PROPOSITION 4.3 Let \mathcal{S}_o and \mathcal{S}_i be defined like \mathcal{R}_o and \mathcal{R}_i , respectively, but with the additional cardinality bounds $|\mathcal{U}| \leq |\mathcal{X}|$ and $|\mathcal{V}| \leq |\mathcal{Y}|$. We then have $\text{conv}(\mathcal{S}_o) = \mathcal{R}_o$ and $\text{conv}(\mathcal{S}_i) = \text{conv}(\mathcal{R}_i) \subseteq \overline{\mathcal{R}}$.

REMARK 4.1 Note that $\text{conv}(\mathcal{S}_i)$ can be explicitly expressed as

$$\text{conv}(\mathcal{S}_i) = \bigcup_{U, V, Q} \{(\mu, R_1, R_2) : R_1 \geq I(U; X|Q), R_2 \geq I(V; Y|Q), \mu \leq I(U; V|Q)\}, \quad (4.3)$$

where U , V , and Q are random variables such that $p_{X,Y,U,V,Q} = p_Q p_{X,Y} p_{U|X,Q} p_{V|Y,Q}$, $|\mathcal{U}| \leq |\mathcal{X}|$, $|\mathcal{V}| \leq |\mathcal{Y}|$, and $|\mathcal{Q}| \leq 3$.

The cardinality bound $|\mathcal{Q}| \leq 3$ follows directly from the strengthened Carathéodory theorem [12 Theorem 18(ii)] because $\text{conv}(\mathcal{R}_i)$ is the convex hull of a connected set in \mathbb{R}^3 .

Note that the cardinality bounds in this result are tighter than the usual bounds obtained with the convex cover method [13 Appendix C], where the cardinality has to be increased by one. We will exploit this fact with binary sources in Section 5 to show that binary auxiliary random variables suffice. The smaller cardinality bounds come at the cost of convexification for the outer bound since in contrast to \mathcal{R}_o , the region \mathcal{S}_o is not necessarily convex.

A tight bound on the achievable region can be obtained if μ is not greater than the Gács-Körner common information (cf. [17][51][54]) of X and Y , as stated in the following corollary.

COROLLARY 4.1 If $Y = \zeta_1(X) = \zeta_2(Y)$ is common to X and Y in the sense of [51] and $0 \leq \mu \leq H(Y)$ then $(\mu, R_1, R_2) \in \overline{\mathcal{R}}$ if and only if $\mu \leq \min\{R_1, R_2\}$.

Proof. Theorem 4.2 entails $\mu \leq \min\{R_1, R_2\}$ for any $(\mu, R_1, R_2) \in \overline{\mathcal{R}}$. With $U = V = Y$, Theorem 4.1 implies $(H(Y), H(Y), H(Y)) \in \overline{\mathcal{R}}$. Using time-sharing with $\mathbf{0} \in \overline{\mathcal{R}}$ we obtain $(\mu, \mu, \mu) \in \overline{\mathcal{R}}$ for $0 \leq \mu \leq H(Y)$ and hence $(\mu, R_1, R_2) \in \overline{\mathcal{R}}$ if $\mu \leq \min\{R_1, R_2\}$. \square

5. Doubly symmetric binary source

In this section, we analyze the achievable region for a DSBS. The same region (cf. Theorem 3.4) was previously analyzed in [52] in the context of a pattern recognition problem. We obtain additional results, disproving [52 Conjecture 1]. In particular, we conjecture that there is a gap between the inner bound $\text{conv}(\mathcal{S}_i)$ and the outer bound \mathcal{R}_o for the DSBS. To support this conjecture, we analyze a region \mathcal{S}_b , previously introduced by the authors of [52], with the property that $\mathcal{S}_b \subseteq \mathcal{S}_i$. However, we prove $\text{conv}(\mathcal{S}_b) \neq \mathcal{R}_o$ and subsequently conjecture that $\text{conv}(\mathcal{S}_b) = \text{conv}(\mathcal{S}_i)$, based on numerical evidence.

For this section, let $(X, Y) \sim \text{DSBS}(p)$ be a DSBS [13 Example 10.1] with parameter $p \in [0, 1]$, i.e., $X \sim \mathcal{B}(\frac{1}{2})$, $X \perp N \sim \mathcal{B}(p)$, and $Y := X \oplus N$. The cardinality bounds in Proposition 4.3 will enable us to use binary auxiliary random variables.

Subsequently we will provide evidence, supporting the following conjecture.

CONJECTURE 5.1. *There exists $p \in [0, 1]$, such that $\text{conv}(\mathcal{S}_i) \neq \mathcal{R}_o$.*


Let \mathcal{S}_b be defined as

$$\mathcal{S}_b := \bigcup_{0 \leq \alpha, \beta \leq \frac{1}{2}} \left\{ (\mu, R_1, R_2) : \begin{aligned} R_1 &\geq \log 2 - h_2(\alpha), \\ R_2 &\geq \log 2 - h_2(\beta), \\ \mu &\leq \log 2 - h_2(\alpha * p * \beta) \end{aligned} \right\}. \quad (5.1)$$

By choosing $U = X \oplus N_1$ and $V = Y \oplus N_2$, where $N_1 \sim \mathcal{B}(\alpha)$ and $N_2 \sim \mathcal{B}(\beta)$ are independent of (X, Y) and of each other, it follows that $\mathcal{S}_b \subseteq \mathcal{S}_i$. To illustrate the tradeoff between complexity (R_1, R_2) and relevance (μ) , the boundary of \mathcal{S}_b is depicted in Fig. 2 for $p = 0.1$.

Based on numerical experiments, we conjecture the following.

CONJECTURE 5.2. *For the source $(X, Y) \sim \text{DSBS}(p)$ with $p \in [0, 1]$ we have $\text{conv}(\mathcal{S}_i) = \text{conv}(\mathcal{S}_b)$, or equivalently $\mathcal{S}_i \subseteq \text{conv}(\mathcal{S}_b)$.*

Figure 2: Boundary of \mathcal{S}_b for $p = 0.1$.

The natural, stronger conjecture that $\mathcal{S}_b = \mathcal{S}_i$ already appeared in [52, Conjecture 1, Eq. (14)]. However, there is the following counterexample [5].

PROPOSITION 5.3 For the source $(X, Y) \sim \text{DSBS}(0)$ we have $\mathcal{S}_b \neq \mathcal{S}_i$.

Proof. For $a \in [0, 1]$ we define (U, V) by the binary channels depicted in Fig. 3, satisfying





Figure 3: Random variables for the proof of Proposition 5.3

$U \circlearrowleft X \circlearrowleft Y \circlearrowleft V$. We obtain $(\mu, R, R) \in \mathcal{S}_i$ with $R = I(U; X) = I(V; Y) = h_2(\frac{a}{2}) - \frac{1}{2}h_2(a)$ and $\mu = I(U; V) = 2R - a \log(2)$. For $a = 0.8$ we have $\mu \approx 0.291103$ and $R \approx 0.42281$. On the other hand, we obtain $\mu_b := \max\{\hat{\mu} : (\hat{\mu}, R, R) \in \mathcal{S}_b\} < 0.285594$ using (5.1) with $\alpha = \beta \approx 0.07658$. As $\mu_b < \mu$ we have $(\mu, R, R) \notin \mathcal{S}_b$.

This argument can be verified numerically using interval arithmetic [34]. Code written in the Octave Programming Language [11] using its interval package [27] can be found at [38]. \square

Note that Proposition 5.3 does not impact Conjecture 5.2 as it concerns the case $p = 0$. For $p = 0$ we have $X = Y$ and Corollary 4.1 implies $\bar{\mathcal{R}} = \{(\mu, R_1, R_2) : R_1, R_2 \geq 0 \text{ and } \mu \leq \min\{R_1, R_2, \log 2\}\}$. It is easily verified that $\bar{\mathcal{R}} = \text{conv}(\mathcal{S}_b)$ and thus Conjecture 5.2 holds for $p = 0$ by Proposition 4.3.

In fact, it can be shown that the entire statement [52, Conjecture 1] does not hold. The second part

Figure 4: Numeric evaluation of $\hat{\mu}_o$ and $\hat{\mu}_b$ for $p = 0.1$.

[52] Conjecture 1, Eq. (15)] claims that $\text{conv}(\mathcal{S}_b) = \mathcal{R}_o$. However, we shall construct a counterexample to this claim, showing that Conjecture 5.1 follows directly from Conjecture 5.2

PROPOSITION 5.4 For the source $(X, Y) \sim \text{DSBS}(0.1)$, we have $\text{conv}(\mathcal{S}_b) \neq \mathcal{R}_o$.

To prove Proposition 5.4 we will construct a point $(\mu, R, R) \in \mathcal{R}_o$ that satisfies $(\mu, R, R) \notin \text{conv}(\mathcal{S}_b)$. To this end, define the concave functions $\hat{\mu}_b(R) := \max\{\mu : (\mu, R, R) \in \text{conv}(\mathcal{S}_b)\}$ and $\hat{\mu}_o(R) := \max\{\mu : (\mu, R, R) \in \mathcal{R}_o\}$ for $R \in [0, \log 2]$. In order to show $\text{conv}(\mathcal{S}_b) \neq \mathcal{R}_o$, it suffices to find $\hat{R} \in [0, \log 2]$ with $\hat{\mu}_b(\hat{R}) < \hat{\mu}_o(\hat{R})$.

We can numerically compute an upper bound for the function $\hat{\mu}_b$. For $\alpha, \beta \in [0, \frac{1}{2}]$, we calculate

$$\tilde{R}_1 := \log 2 - h_2(\alpha), \quad (5.2)$$

$$\tilde{R}_2 := \log 2 - h_2(\beta), \text{ and} \quad (5.3)$$

$$\tilde{\mu} := \log 2 - h_2(\alpha * p * \beta) \quad (5.4)$$

on a suitably fine grid and upper bound the upper concave envelope of the implicitly defined function $\tilde{\mu}(\tilde{R}_1, \tilde{R}_2)$. Evaluating it at $R = \tilde{R}_1 = \tilde{R}_2$ yields an upper bound for $\hat{\mu}_b(R)$.

On the other hand, we can obtain a lower bound for $\hat{\mu}_o$ by computing (4.2) for specific pmfs that satisfy the Markov constraints in Theorem 4.2. Note that based on the cardinality bound in Proposition 4.3 we can restrict the auxiliary random variables U and V to be binary. We randomly sample the binary pmfs that satisfy the Markov constraints in Theorem 4.2 (but not necessarily the long Markov chain $U \dashv X \dashv Y \dashv V$) and in doing so encounter points strictly above the graph of $\hat{\mu}_b$. Figure 4 shows the resulting bounds for $p = 0.1$ in the vicinity of $R = \log 2$. Albeit small, there is clearly a gap between $\hat{\mu}_b$ and $\hat{\mu}_o$ outside the margin of numerical error.

Proof of Proposition 5.4 We observed the largest gap between the two bounds at a rate of $\hat{R} \approx 0.675676$. The particular distribution of (U, V) at this rate, resulting from optimizing over the distributions that satisfy the Markov constraints in Theorem 4.2 is given in Table 1 for reference. Note that this is an exact

u	v	x	y	$P\{U = u, V = v X = x, Y = y\}$
0	0	0	0	0.995358146217353406525
0	0	0	1	0.00249767559844423319075
0	0	1	0	0.002498344003957310643325
0	0	1	1	0.00034313919194834475
0	1	0	0	0.002142603857654094275
0	1	0	1	0.99500307447656326760925
0	1	1	0	0.000000905921035188556675
0	1	1	1	0.00215611073304415445
1	0	0	0	0.002142603857654094275
1	0	0	1	0.00000157432654826600925
1	0	1	0	0.995002406071050190156675
1	0	1	1	0.00215611073304415445
1	1	0	0	0.000356646067338404925
1	1	0	1	0.00249767559844423319075
1	1	1	0	0.002498344003957310643325
1	1	1	1	0.99534463934196334635

Table 1: Distribution resulting from random search.

conditional pmf (i.e., not a numerical approximation) that satisfies the Markov chains $U \dashv X \dashv Y$ and $X \dashv Y \dashv V$. It achieves $I(V; Y) + I(U; X) - I(UV; XY) \approx 0.35972$ which is $\Delta \approx 1.985673 \cdot 10^{-4}$ above the inner bound, thus proving Proposition 5.4. Using interval arithmetic [34] this claim can be verified numerically. Code written in the Octave Programming Language [11] using its interval package [27] can be found at [38]. It uses the distribution given in Table 1. \square

We firmly believe that a tight characterization of the achievable region requires an improved outer bound. However, using current information theoretic tools, it appears very challenging to find a manageable outer bound based on the full Markov chain $U \dashv X \dashv Y \dashv V$.

REMARK 5.1 Recently, Kumar and Courtade introduced a conjecture [7, 31] concerning Boolean functions that maximize mutual information. Their work was inspired by a similar problem in computational biology [29]. A weaker form of their conjecture [7, Section IV, 2)], which was solved in [39], corresponds to a zero-rate/one-bit variant of the binary example studied here.

6. The information bottleneck

The information-theoretic problem posed by the IB method [49] can be obtained as a special case from the biclustering problem. We will introduce the problem setup and subsequently show how it can be derived as a special case of Definition 2.1. Note that the definition slightly differs from [10, Definition 1]. However, the achievable region is identical.

DEFINITION 6.1 A pair (μ, R_1) is *IB-achievable* if, for some $n \in \mathbb{N}$, there exists $f: \mathcal{X}^n \rightarrow \mathcal{M}_1$ with $\log |\mathcal{M}_1| \leq nR_1$ and

$$\mu \leq \frac{1}{n} I(f(\mathbf{X}); \mathbf{Y}). \quad (6.1)$$

Let \mathcal{R}_{IB} be the set of all IB-achievable pairs.

PROPOSITION 6.2 For a pair (μ, R_1) , the following are equivalent:

1. $(\mu, R_1) \in \overline{\mathcal{R}_{\text{IB}}}$.
2. $(\mu, R_1, \log|\mathcal{Y}|) \in \overline{\mathcal{R}}$.
3. There exists a random variable U such that $U \dashv\!\!\!\dashv X \dashv\!\!\!\dashv Y$, $I(X;U) \leq R$ and $I(Y;U) \geq \mu$.

Proof. The equivalence $1 \Leftrightarrow 2$ holds as Definition 2.1 collapses to Definition 6.1 for $R_2 = \log|\mathcal{Y}|$. To show $2 \Leftrightarrow 3$ apply Theorems 4.1 and 4.2 with $V = Y$. \square

The tradeoff between ‘relevance’ and ‘complexity’ can equivalently be characterized by the IB function (cf. [8, 18]) $\mu_{\text{IB}}(R) := \sup\{\mu : (\mu, R) \in \overline{\mathcal{R}_{\text{IB}}}\}$. Proposition 6.2 provides

$$\mu_{\text{IB}}(R) = \max_{\substack{U : I(U;X) \leq R \\ U \dashv\!\!\!\dashv X \dashv\!\!\!\dashv Y}} I(U;Y). \quad (6.2)$$

Interestingly, the function (6.2) is the solution to a variety of different problems in information theory. As mentioned in [18], (6.2) is the solution to the problem of loss-less source coding with one helper [2, 58]. Witsenhausen and Wyner [55] investigated a lower bound for a conditional entropy when simultaneously requiring another conditional entropy to fall below a threshold. Their work was a generalization of [57] and furthermore related to [2, 3, 53, 56]. The conditional entropy bound in [55] turns out to be an equivalent characterization of (6.2). Furthermore, μ_{IB} characterizes the optimal error exponent, when testing against independence with one-sided data compression [1 Theorem 2]. Also in the context of gambling in the horse race market, (6.2) occurs as the maximum incremental growth in wealth when rate-limited side-information is available to the gambler [15 Theorem 3].

7. Multiple description CEO problem

In [8 Appendix B] Courtade and Weissman considered a multi-terminal extension of the IB problem, as introduced in Section 6 the CEO problem under an information constraint. Analogous to how the IB problem is a special case of the biclustering problem (cf. Proposition 6.2), this CEO problem presents a special case of a multi-terminal generalization of the biclustering problem [40]. Under a conditional independence assumption, Courtade and Weissman were able to provide a single letter characterization of the achievable region. In what follows we will extend this result, by incorporating MD coding for the CEO problem. Loosely speaking, we require the CEO to obtain valuable information from the message of just one agent alone. Surprisingly, this extension also permits a single-letter characterization under the same conditional independence assumption.

In what follows, let $(X_{\mathcal{J}}, Y)$ be $J+1$ random variables, where $\mathcal{J} := [1:J]$, satisfying the Markov chain $X_j \dashv\!\!\!\dashv Y \dashv\!\!\!\dashv X_{\mathcal{J} \setminus j}$ for every $j \in \mathcal{J}$. An $(n, R_{\mathcal{J}})$ -code $f_{\mathcal{J}}$ consists of J functions $f_j : \mathcal{X}_j^n \rightarrow \mathcal{M}_j$ with $\log|\mathcal{M}_j| \leq nR_j$ for every $j \in \mathcal{J}$.

DEFINITION 7.1 A point $(v_0, v_{\mathcal{J}}, R_{\mathcal{J}}) \in \mathbb{R}^{2J+1}$ is *MI-achievable* if for some $n \in \mathbb{N}$ there exists an $(n, R_{\mathcal{J}})$ -code $f_{\mathcal{J}}$ such that

$$\frac{1}{n} I(f_1(\mathbf{X}_1), f_2(\mathbf{X}_2), \dots, f_J(\mathbf{X}_J); \mathbf{Y}) \geq v_0 \text{ and} \quad (7.1)$$

$$\frac{1}{n} I(f_j(\mathbf{X}_j); \mathbf{Y}) \geq v_j, \quad \text{for all } j \in \mathcal{J}. \quad (7.2)$$

Denote the set of all MI-achievable points by \mathcal{R}_{MI} .

To shorten notation we will introduce the set of random variables

$$\mathcal{P}_* := \left\{ U_{\mathcal{J}}, Q : p_{QU_{\mathcal{J}}X_{\mathcal{J}}Y} = p_Q \cdot p_{X_{\mathcal{J}}Y} \cdot \prod_{j \in \mathcal{J}} p_{U_j|X_jQ} \right\}. \quad (7.3)$$

DEFINITION 7.2 For a total order¹ \sqsubset on \mathcal{J} and a set $\mathcal{I} \subseteq \mathcal{J}$, let the region $\mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})} \subseteq \mathbb{R}^{2J+1}$ be the set of tuples $(v_0, v_{\mathcal{J}}, R_{\mathcal{J}})$ such that there exist random variables $(U_{\mathcal{J}}, \emptyset) \in \mathcal{P}_*$ with

$$R_j \geq I(U_j; X_j | U_{\sqsubset j}), \quad j \in \mathcal{J}, \quad (7.4)$$

$$R_j \geq I(U_j; X_j), \quad j \in \mathcal{I}, \quad (7.5)$$

$$v_j \leq I(U_j; Y), \quad j \in \mathcal{J}, \quad (7.6)$$

$$v_j \leq I(U_j; Y | U_{\sqsubset j}), \quad j \notin \mathcal{I}, \quad (7.7)$$

$$v_0 \leq I(U_{\mathcal{J}}; Y). \quad (7.8)$$

REMARK 7.1 The purpose of the order \sqsubset is to determine the order of the messages for successive decoding. Equivalently, Definition 7.2 could be rephrased in terms of a permutation of \mathcal{J} in place of a total order.

We are now able to state the single-letter characterization of $\overline{\mathcal{R}_{\text{MI}}}$, the proof of which is provided in Appendix A.3

THEOREM 7.3 We have $\overline{\mathcal{R}_{\text{MI}}} = \text{conv} \left(\bigcup_{\sqsubset, \mathcal{I}} \mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})} \right)$, where the union is over all total orders on \mathcal{J} and all sets $\mathcal{I} \subseteq \mathcal{J}$.

For $J = 2$, we have $\overline{\mathcal{R}_{\text{MI}}} = \text{conv} \left(\mathcal{R}_{\text{MI}}^{(1)} \cup \mathcal{R}_{\text{MI}}^{(2)} \cup \mathcal{R}_{\text{MI}}^{(3)} \right)$, where $(v_0, v_{\mathcal{J}}, R_{\mathcal{J}}) \in \mathcal{R}_{\text{MI}}^{(i)}$ iff, for some $(U_{\mathcal{J}}, \emptyset) \in \mathcal{P}_*$, the following inequalities are satisfied

$$\begin{array}{lll} \mathcal{R}_{\text{MI}}^{(1)} : & \mathcal{R}_{\text{MI}}^{(2)} : & \mathcal{R}_{\text{MI}}^{(3)} : \\ v_1 \leq I(Y; U_1) & v_1 \leq I(Y; U_1 | U_2) & v_1 \leq I(Y; U_1) \end{array} \quad (7.9)$$

$$v_2 \leq I(Y; U_2 | U_1) \quad v_2 \leq I(Y; U_2) \quad v_2 \leq I(Y; U_2) \quad (7.10)$$

$$v_0 \leq I(Y; U_1 U_2) \quad v_0 \leq I(Y; U_1 U_2) \quad v_0 \leq I(Y; U_1 U_2) \quad (7.11)$$

$$R_1 \geq I(U_1; X_1) \quad R_1 \geq I(U_1; X_1 | U_2) \quad R_1 \geq I(U_1; X_1) \quad (7.12)$$

$$R_2 \geq I(U_2; X_2 | U_1) \quad R_2 \geq I(U_2; X_2) \quad R_2 \geq I(U_2; X_2). \quad (7.13)$$

REMARK 7.2 Note that the total available rate of encoder 2 is $R_2 = I(X_2; U_2 | U_1)$ to achieve a point in $\mathcal{R}_{\text{MI}}^{(1)}$. Interestingly, this rate is in general less than the rate required to ensure successful typicality decoding of U_2 . However, $v_2 = I(Y; U_2 | U_1)$ can still be achieved.

REMARK 7.3 On the other hand, fixing the random variables U_1, U_2 in the definition of $\mathcal{R}_{\text{MI}}^{(i)}$ shows another interesting feature of this region. The achievable values for v_1 and v_2 vary across $i \in \{1, 2, 3\}$ and hence do not only depend on the chosen random variables U_1 and U_2 , but also on the specific rates R_1 and R_2 .

¹ For the notation regarding total orders refer to Section 1.3

It is worth mentioning that by setting $v_j = 0$ for $j = 1, 2, \dots, J$, the region $\overline{\mathcal{R}_{\text{MI}}}$ reduces to the rate region in [8, Appendix B].

The following proposition shows that $\mathcal{R}_{\text{MI}}^{(\square, \mathcal{I})}$ is computable, at least in principle. The given cardinality bound is not optimal, but it implies $\mathcal{R}_{\text{MI}}^{(\square, \mathcal{I})} = \overline{\mathcal{R}_{\text{MI}}^{(\square, \mathcal{I})}}$.

PROPOSITION 7.4 The region $\mathcal{R}_{\text{MI}}^{(\square, \mathcal{I})}$ remains unchanged if the cardinality bound $|\mathcal{U}_j| \leq |X_j| + 4^J$ is imposed for every $j \in \mathcal{J}$.

The proof of Proposition 7.4 is provided in Appendix A.4

8. Summary and discussion

We introduced a multi-terminal generalizations of the IB problem, termed information-theoretic biclustering. Interestingly, this problem is related to several other problems at the frontier of statistics and information theory and offers a formidable mathematical complexity. Indeed, it is fundamentally different from ‘classical’ distributed source coding problems where the encoders usually aim at reducing, as much as possible, redundant information among the sources while still satisfying a fidelity criterion. In the considered problem, however, the encoders are interested in maximizing precisely such redundant information.

While an exact characterization of the achievable region is mathematically very challenging and still remains elusive, we provided outer and inner bounds to the set of achievable rates. We thoroughly studied the special case of two symmetric binary sources for which novel cardinality bounding techniques were developed. Based on numerical evidence we formulated a conjecture that entails an explicit expression for the inner bound. This conjecture provides strong evidence that our inner and outer bounds do not meet in general. We firmly believe that an improved outer bound, satisfying the adequate Markov chains, is required for a tight characterization of the achievable region.

Furthermore we considered an MD CEO problem which surprisingly permits a single-letter characterization of the achievable region. The resulting region has the remarkable feature that it allows to exploit rate that is in general insufficient to guarantee successful typicality encoding.

The interesting challenge of the biclustering problem lies in the fact that one needs to bound the mutual information between two arbitrary encodings solely based on their rates. Standard information-theoretic manipulations seem incapable of handling this requirement well.

Funding

Wiener Wissenschafts-, Forschungs- und Technologiefonds (ICT12-054, ICT15-119); This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 797805.

Acknowledgment

The authors would like to thank Shlomo Shamai (Shitz) and Emre Telatar for insightful discussions regarding the binary example. We would also like to thank Christian Chapman from the School of Electrical, Computer and Energy Engineering at the Arizona State University for providing the counterexample used in the proof of Proposition 5.3

Data availability statement

No new data were generated or analysed in support of this review.

A. Proofs

A.1 Proof of Theorem 3.4

To prove $\mathcal{R} \subseteq \mathcal{R}_*$, assume $(\mu, R_1, R_2) \in \mathcal{R}$ and choose n , f and g according to Definition 2.1. Defining $U := f(\mathbf{X})$ and $V := g(\mathbf{Y})$ yields inequalities (3.1)–(3.3) and satisfies the required Markov chain.

The inclusions $\overline{\mathcal{R}}_* \subseteq \overline{\mathcal{R}}_{\text{HT}}$ and $\overline{\mathcal{R}}_* \subseteq \overline{\mathcal{R}}_{\text{PR}}$ follow by applying the achievability results [23 Corollary 6] and [52 Theorem 1], respectively, to the vector source (\mathbf{X}, \mathbf{Y}) .

Assuming $(\mu, R_1, R_2) \in \mathcal{R}_{\text{HT}}$, choose an arbitrary $\varepsilon > 0$ and pick an (n, R_1, R_2) hypothesis test $(f_n, g_n, \mathcal{A}_n)$ such that $\alpha_n \leq \varepsilon$ and $-\log \beta_n \geq n(\mu - \varepsilon)$. We apply the log-sum inequality [9 Theorem 2.7.1] and obtain for any $\varepsilon' > 0$, provided that ε is small enough and n is large enough,

$$I(f(\mathbf{X}); g(\mathbf{Y})) \geq (1 - \alpha_n) \log \frac{1 - \alpha_n}{\beta_n} + \alpha_n \log \frac{\alpha_n}{1 - \beta_n} \geq n(\mu - \varepsilon'), \quad (\text{A.1})$$

which implies $(\mu, R_1, R_2) \in \overline{\mathcal{R}}$.

Similarly, assume $(\mu, R_1, R_2) \in \mathcal{R}_{\text{PR}}$ and for an arbitrary $\varepsilon > 0$ and sufficiently large $n \in \mathbb{N}$ pick an (n, R_1, R_2) -code (f, g) and ϕ satisfying (3.6). Then (using the notation of Definition 3.3),

$$I(f(\mathbf{X}); g(\mathbf{Y})) = I(C; g(\mathbf{Y}(W)) | W) \quad (\text{A.2})$$

$$\geq I(C; W | g(\mathbf{Y}(W))) \quad (\text{A.3})$$

$$\geq n\mu - H(W | \phi(C, g(\mathbf{Y}(W)))) \quad (\text{A.4})$$

$$\stackrel{(3.6)}{\geq} n\mu - h_2(\varepsilon) - \varepsilon n\mu. \quad (\text{A.5})$$

The equality in (A.2) holds as $\mathbf{X}(i) \perp \mathbf{Y}(j)$ for $i \neq j$, (A.3) follows from $W \perp C$, and (A.4) follows from $W \perp \mathbf{Y}(W)$, the fact that $H(W) = n\mu$, and the data processing inequality [9 Theorem 2.8.1]. Fano's inequality [9 Theorem 2.10.1] was used in (A.5). This shows $(\mu, R_1, R_2) \in \overline{\mathcal{R}}$ as ε was arbitrary.

A.2 Proof of Proposition 4.3

We start with the proof of $\text{conv}(\mathcal{S}_0) = \mathcal{R}_0$.

For fixed random variables (X, Y) define the set of pmfs (with finite, but arbitrarily large support)

$$\mathcal{Q} := \{p_{U, X, Y, V} : U \text{ --- } X \text{ --- } Y, X \text{ --- } Y \text{ --- } V\}, \quad (\text{A.6})$$

and the compact set of pmfs with fixed alphabet size

$$\mathcal{Q}(a, b) := \{p_{U, X, Y, V} \in \mathcal{Q} : |\mathcal{U}| = a, |\mathcal{V}| = b\}. \quad (\text{A.7})$$

Define the continuous vector valued function $\mathbf{F} := (F_1, F_2, F_3)$ as

$$F_1(p_{U, X, Y, V}) := I(X; U) + I(Y; V) - I(UV; XY), \quad (\text{A.8})$$

$$F_2(p_{U, X, Y, V}) := I(U; X), \quad (\text{A.9})$$

$$F_3(p_{U, X, Y, V}) := I(V; Y). \quad (\text{A.10})$$

We can now write $\mathcal{R}_0 = \mathbf{F}(\mathcal{Q}) + \mathcal{O}$ and $\mathcal{S}_0 = \mathbf{F}(\mathcal{Q}(|\mathcal{X}|, |\mathcal{Y}|)) + \mathcal{O}$ where $\mathcal{O} := (\mathbb{R}_- \times \mathbb{R}_+ \times \mathbb{R}_+)$. Since \mathcal{R}_0 is convex, we may define the extended real function $\psi(\boldsymbol{\lambda}) := \inf_{\mathbf{x} \in \mathcal{R}_0} \boldsymbol{\lambda} \cdot \mathbf{x}$ and obtain [22 Theorem 2.2, 3.]

$$\overline{\text{conv}(\mathcal{R}_0)} = \overline{\mathcal{R}_0} = \bigcap_{\boldsymbol{\lambda} \in \mathbb{R}^3} \left\{ \mathbf{x} \in \mathbb{R}^3 : \boldsymbol{\lambda} \cdot \mathbf{x} \geq \psi(\boldsymbol{\lambda}) \right\}. \quad (\text{A.11})$$

From the definition of \mathcal{R}_0 , we have $\psi(\boldsymbol{\lambda}) = -\infty$ if $\boldsymbol{\lambda} \notin \mathcal{O}$, and $\psi(\boldsymbol{\lambda}) = \inf_{\mathbf{p} \in \mathcal{Q}} \boldsymbol{\lambda} \cdot \mathbf{F}(\mathbf{p})$ otherwise. This shows, that

$$\overline{\mathcal{R}_0} = \bigcap_{\boldsymbol{\lambda} \in \mathcal{O}} \left\{ \mathbf{x} \in \mathbb{R}^3 : \boldsymbol{\lambda} \cdot \mathbf{x} \geq \psi(\boldsymbol{\lambda}) \right\} \quad (\text{A.12})$$

and using the same argument, one can also show that

$$\overline{\text{conv}(\mathcal{S}_0)} = \bigcap_{\boldsymbol{\lambda} \in \mathcal{O}} \left\{ \mathbf{x} \in \mathbb{R}^3 : \boldsymbol{\lambda} \cdot \mathbf{x} \geq \tilde{\psi}(\boldsymbol{\lambda}) \right\}, \quad \tilde{\psi}(\boldsymbol{\lambda}) = \min_{\mathbf{p} \in \mathcal{Q}(|\mathcal{X}|, |\mathcal{Y}|)} \boldsymbol{\lambda} \cdot \mathbf{F}(\mathbf{p}). \quad (\text{A.13})$$

We shall now prove that $\psi(\boldsymbol{\lambda}) = \tilde{\psi}(\boldsymbol{\lambda})$ for $\boldsymbol{\lambda} \in \mathcal{O}$. For arbitrary $\boldsymbol{\lambda} \in \mathcal{O}$ and $\delta > 0$, we can find random variables $(\tilde{U}, X, Y, \tilde{V}) \sim \tilde{\mathbf{p}} \in \mathcal{Q}$ with $\boldsymbol{\lambda} \cdot \mathbf{F}(\tilde{\mathbf{p}}) \leq \psi(\boldsymbol{\lambda}) + \delta$. By compactness of $\mathcal{Q}(a, b)$ and continuity of \mathbf{F} , there is $\mathbf{p} \in \mathcal{Q}(|\mathcal{U}|, |\mathcal{V}|)$ with

$$\boldsymbol{\lambda} \cdot \mathbf{F}(\mathbf{p}) = \min_{\hat{\mathbf{p}} \in \mathcal{Q}(|\mathcal{U}|, |\mathcal{V}|)} \boldsymbol{\lambda} \cdot \mathbf{F}(\hat{\mathbf{p}}) \leq \boldsymbol{\lambda} \cdot \mathbf{F}(\tilde{\mathbf{p}}) \leq \psi(\boldsymbol{\lambda}) + \delta. \quad (\text{A.14})$$

We now show that there exists $\hat{\mathbf{p}} \in \mathcal{Q}(|\mathcal{X}|, |\mathcal{Y}|)$ with

$$\boldsymbol{\lambda} \cdot \mathbf{F}(\hat{\mathbf{p}}) = \boldsymbol{\lambda} \cdot \mathbf{F}(\mathbf{p}). \quad (\text{A.15})$$

As a consequence of the inequalities $F_1 \leq F_2$ and $F_1 \leq F_3$ we have $\boldsymbol{\lambda} \cdot \mathbf{F}(\mathbf{p}) = 0$ if $\lambda_1 + \max\{\lambda_2, \lambda_3\} \geq 0$. Thus, we only need to show (A.15) for $\boldsymbol{\lambda} \in \mathcal{O}$ with $\lambda_1 + \lambda_2 < 0$ and $\lambda_1 + \lambda_3 < 0$. To this end we use the perturbation method [19][28] and perturb \mathbf{p} , obtaining the candidate

$$(\hat{U}, X, Y, \hat{V}) \sim \hat{\mathbf{p}}(u, x, z, v) = \mathbf{p}(u, x, z, v)(1 + \varepsilon\phi(u)). \quad (\text{A.16})$$

We require

$$1 + \varepsilon\phi(u) \geq 0, \quad \text{for every } u \in \mathcal{U}, \quad (\text{A.17})$$

$$\mathbb{E}[\phi(U)] = 0, \quad (\text{A.18})$$

$$\mathbb{E}[\phi(U)|X=x, Y=z] = 0, \quad \text{if } \mathbf{p}(x, z) > 0. \quad (\text{A.19})$$

The conditions (A.17) and (A.18) ensure that $\hat{\mathbf{p}}$ is a valid pmf and (A.19) implies $\hat{\mathbf{p}} \in \mathcal{Q}$. Observe that there is an $\varepsilon_0 > 0$ for any ϕ , such that (A.17) is satisfied for $\varepsilon \in [-\varepsilon_0, \varepsilon_0]$. Furthermore, (A.19) is equivalent to

$$\mathbb{E}[\phi(U)|X=x] = 0, \quad \text{for every } x \in \mathcal{X} \quad (\text{A.20})$$

due to the Markov chain $U \rightarrow X \rightarrow Y$. Note also that (A.20) already implies (A.18). If $|\mathcal{U}| \geq |\mathcal{X}| + 1$ there is a non-trivial solution to (A.20) which means there exists $\phi \neq 0$ such that (A.17)-(A.19) are

satisfied. We have

$$\begin{aligned}\boldsymbol{\lambda} \cdot \mathbf{F}(\hat{\mathbf{p}}) &= \lambda_1 \left[\mathbf{I}(X; U) - \mathbf{I}(UV; XY) + \mathbf{H}(Y) + \varepsilon \mathbf{H}_\phi(U) - \varepsilon \mathbf{H}_\phi(UX) \right. \\ &\quad \left. - \varepsilon \mathbf{H}_\phi(UV) + \varepsilon \mathbf{H}_\phi(UXYV) + \mathbf{H}(\hat{V}) - \mathbf{H}(Y\hat{V}) \right] \\ &\quad + \lambda_2 [\mathbf{I}(X; U) + \varepsilon \mathbf{H}_\phi(U) - \varepsilon \mathbf{H}_\phi(UX)] \\ &\quad + \lambda_3 [\mathbf{H}(Y) + \mathbf{H}(\hat{V}) - \mathbf{H}(Y\hat{V})].\end{aligned}\tag{A.21}$$

Here, we used the shorthand $\mathbf{H}_\phi(UX) := -\sum_{u,x} p(u,x) \phi(u) \log p(u,x)$ and analogous for other combinations of random variables. By (A.14), we have $\frac{\partial^2}{\partial \varepsilon^2} \boldsymbol{\lambda} \cdot \mathbf{F}(\hat{\mathbf{p}})|_{\varepsilon=0} \geq 0$. Observe that

$$\frac{\partial}{\partial \varepsilon} (\mathbf{H}(\hat{V}) - \mathbf{H}(Y\hat{V})) = \frac{\partial}{\partial \varepsilon} \sum_{z,v} \hat{p}(z,v) \log \frac{\hat{p}(z,v)}{\hat{p}(v)}\tag{A.22}$$

$$= \sum_{z,v} \frac{\partial \hat{p}(z,v)}{\partial \varepsilon} \log \frac{\hat{p}(z,v)}{\hat{p}(v)} + \frac{\partial \hat{p}(z,v)}{\partial \varepsilon} - \frac{\hat{p}(z,v) \frac{\partial \hat{p}(v)}{\partial \varepsilon}}{\hat{p}(v)}\tag{A.23}$$

and consequently,

$$\frac{\partial^2}{\partial \varepsilon^2} \boldsymbol{\lambda} \cdot \mathbf{F}(\hat{\mathbf{p}}) = (\lambda_1 + \lambda_3) \frac{\partial^2}{\partial \varepsilon^2} (\mathbf{H}(\hat{V}) - \mathbf{H}(Y\hat{V}))\tag{A.24}$$

$$\begin{aligned}&= (\lambda_1 + \lambda_3) \sum_{z,v} \left(\frac{\partial \hat{p}(z,v)}{\partial \varepsilon} \right)^2 \frac{1}{\hat{p}(z,v)} \\ &\quad - 2 \frac{\partial \hat{p}(z,v)}{\partial \varepsilon} \frac{\partial \hat{p}(v)}{\partial \varepsilon} \frac{1}{\hat{p}(v)} + \left(\frac{\partial \hat{p}(v)}{\partial \varepsilon} \right)^2 \frac{\hat{p}(z,v)}{\hat{p}(v)^2}.\end{aligned}\tag{A.25}$$

Here we already used that $\frac{\partial^2 \hat{p}(v)}{\partial \varepsilon^2} \equiv \frac{\partial^2 \hat{p}(z,v)}{\partial \varepsilon^2} \equiv 0$. It is straightforward to calculate

$$\frac{\partial \hat{p}(v)}{\partial \varepsilon} = p(v) \mathbb{E}[\phi(U) | V = v],\tag{A.26}$$

$$\frac{\partial \hat{p}(z,v)}{\partial \varepsilon} = p(z,v) \mathbb{E}[\phi(U) | V = v, Y = z],\tag{A.27}$$

$$\hat{p}(z,v)|_{\varepsilon=0} = p(z,v),\tag{A.28}$$

$$\hat{p}(v)|_{\varepsilon=0} = p(v),\tag{A.29}$$

and thus, taking into account that $\lambda_1 + \lambda_3 < 0$,

$$0 \geq \sum_{z,v} p(z,v) \left(\mathbb{E}[\phi(U) | V = v, Y = z] - \mathbb{E}[\phi(U) | V = v] \right)^2,\tag{A.30}$$

which implies for any $(z,v) \in \mathcal{Y} \times \mathcal{V}$ with $p(z,v) > 0$,

$$\sum_u p(u|z,v) \phi(u) = \sum_u p(u|v) \phi(u).\tag{A.31}$$

From (A.31) we can conclude

$$H(\hat{V}) - H(Y\hat{V}) = \sum_{z,v} \hat{p}(z,v) \log \frac{\hat{p}(z,v)}{\hat{p}(v)} \quad (\text{A.32})$$

$$= \sum_{z,v,u} p(u,z,v)(1 + \varepsilon\phi(u)) \log \frac{\sum_{\hat{u}} p(\hat{u},z,v)(1 + \varepsilon\phi(\hat{u}))}{\sum_{\hat{u}} p(\hat{u},v)(1 + \varepsilon\phi(\hat{u}))} \quad (\text{A.33})$$

$$= \sum_{z,v,u} p(u,z,v)(1 + \varepsilon\phi(u)) \log \frac{p(z,v)(1 + \varepsilon\sum_{\hat{u}} p(\hat{u}|z,v)\phi(\hat{u}))}{p(v)(1 + \varepsilon\sum_{\hat{u}} p(\hat{u}|v)\phi(\hat{u}))} \quad (\text{A.34})$$

$$\stackrel{(\text{A.31})}{=} \sum_{z,v,u} p(u,z,v)(1 + \varepsilon\phi(u)) \log \frac{p(z,v)}{p(v)} \quad (\text{A.35})$$

$$= \sum_{z,v} p(z,v) \log \frac{p(z,v)}{p(v)} + \varepsilon \sum_{z,v,u} \phi(u) p(u,z,v) \log \frac{p(z,v)}{p(v)} \quad (\text{A.36})$$

$$= H(V) - H(YV) + \varepsilon H_\phi(V) - \varepsilon H_\phi(YV), \quad (\text{A.37})$$

where we used

$$H_\phi(V) := - \sum_{u,v} p(u,v) \phi(u) \log p(v), \quad (\text{A.38})$$

$$H_\phi(YV) := - \sum_{u,z,v} p(u,z,v) \phi(u) \log p(z,v). \quad (\text{A.39})$$

Substituting in (A.21) shows that $\boldsymbol{\lambda} \cdot \mathbf{F}(\hat{p})$ is linear in ε . And by the optimality of p it must be constant. We may now choose ε maximal, i.e., such that there is at least one $u \in \mathcal{U}$ with $p(u)(1 + \varepsilon\phi(u)) = 0$. This effectively reduces the cardinality of $\hat{\mathcal{U}}$ by at least one and may be repeated until $|\hat{\mathcal{U}}| = |\mathcal{X}|$ (as then $\phi \equiv 0$). The same process can be carried out for V and yields $\hat{p} \in \mathcal{Q}(|\mathcal{X}|, |\mathcal{Y}|)$ such that (A.15) holds.

Using (A.14) and (A.15) we obtain

$$\psi(\boldsymbol{\lambda}) \leq \tilde{\psi}(\boldsymbol{\lambda}) \leq \boldsymbol{\lambda} \cdot \mathbf{F}(\hat{p}) \leq \psi(\boldsymbol{\lambda}) + \delta. \quad (\text{A.40})$$

As $\delta > 0$ was arbitrary, we proved $\psi(\boldsymbol{\lambda}) = \tilde{\psi}(\boldsymbol{\lambda})$, which implies $\overline{\mathcal{R}_0} = \overline{\text{conv}(\mathcal{S}_0)}$ using (A.12) and (A.13). We find $\mathcal{R}_0 = \text{conv}(\mathcal{S}_0)$ by writing

$$\overline{\mathcal{R}_0} = \overline{\text{conv}(\mathcal{S}_0)} = \overline{\text{conv}(\mathbf{F}(\mathcal{Q}(|\mathcal{X}|, |\mathcal{Y}|)) + \mathcal{O})} \quad (\text{A.41})$$

$$= \text{conv}(\mathcal{S}_0) = \text{conv}(\mathbf{F}(\mathcal{Q}(|\mathcal{X}|, |\mathcal{Y}|)) + \mathcal{O}) \quad (\text{A.42})$$

$$\subseteq \mathbf{F}(\mathcal{Q}) + \mathcal{O} \quad (\text{A.43})$$

$$= \mathcal{R}_0, \quad (\text{A.44})$$

where (A.42) follows from Lemma A.3. The relation (A.43) is a consequence of $\mathcal{Q}(|\mathcal{X}|, |\mathcal{Y}|) \subseteq \mathcal{Q}$ and the convexity of $\mathbf{F}(\mathcal{Q})$.

In order to prove $\text{conv}(\mathcal{S}_i) = \text{conv}(\mathcal{R}_i)$, we will only show the cardinality bound $|\mathcal{U}| \leq |\mathcal{X}|$. The corresponding bound for $|\mathcal{V}|$ follows analogously. We note that the weaker bounds $|\mathcal{U}| \leq |\mathcal{X}| + 1$ and $|\mathcal{V}| \leq |\mathcal{Y}| + 1$ can be obtained directly using the convex cover method [13, Appendix C], [2, 59]. Define the continuous vector-valued function

$$\mathbf{F}(p_{\tilde{U}, \tilde{X}, \tilde{Y}, \tilde{V}}) := (I(\tilde{U}; \tilde{V}), I(\tilde{X}; \tilde{U}), I(\tilde{Y}; \tilde{V})), \quad (\text{A.45})$$

and the compact, connected sets of pmfs

$$\mathcal{Q} := \left\{ p_{\tilde{U}, \tilde{X}, \tilde{Y}, \tilde{V}} : p_{\tilde{U}, \tilde{X}, \tilde{Y}, \tilde{V}} = p_{\tilde{U}|X} p_{X,Y} p_{\tilde{Y}|Y}, \tilde{U} = [0:|\mathcal{X}|], \tilde{V} = [0:|\mathcal{Y}|] \right\}, \quad (\text{A.46})$$

$$\tilde{\mathcal{Q}} := \left\{ p_{\tilde{U}, \tilde{X}, \tilde{Y}, \tilde{V}} \in \mathcal{Q} : \tilde{U} = [1:|\mathcal{X}|] \right\}. \quad (\text{A.47})$$

To complete the proof of the proposition, it suffices to show

$$\text{conv}(\mathbf{F}(\mathcal{Q})) \subseteq \text{conv}(\mathbf{F}(\tilde{\mathcal{Q}})), \quad (\text{A.48})$$

since we then have with $\mathcal{O} := (\mathbb{R}_- \times \mathbb{R}_+ \times \mathbb{R}_+)$,

$$\text{conv}(\mathcal{R}_i) = \text{conv}(\mathbf{F}(\mathcal{Q}) + \mathcal{O}) \quad (\text{A.49})$$

$$= \text{conv}(\mathbf{F}(\mathcal{Q})) + \mathcal{O} \quad (\text{A.50})$$

$$\stackrel{(\text{A.48})}{\subseteq} \text{conv}(\mathbf{F}(\tilde{\mathcal{Q}})) + \mathcal{O} \quad (\text{A.51})$$

$$= \text{conv}(\mathbf{F}(\tilde{\mathcal{Q}}) + \mathcal{O}) \quad (\text{A.52})$$

$$= \text{conv}(\mathcal{S}_i), \quad (\text{A.53})$$

where (A.50) and (A.52) follow from Lemma A.3. The region $\mathbf{F}(\mathcal{Q}) \subseteq \mathbb{R}^3$ is compact [41 Theorem 4.22]. Therefore, its convex hull $\text{conv}(\mathbf{F}(\mathcal{Q}))$ is compact [4 Corollary 5.33] and can be represented as an intersection of halfspaces in the following manner [22 Proposition 2.2, 3.]: Defining $V(\boldsymbol{\lambda}) := \max_{\mathbf{x} \in \mathbf{F}(\mathcal{Q})} \boldsymbol{\lambda} \cdot \mathbf{x}$ for $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \lambda_3) \in \mathbb{R}^3$, we have

$$\text{conv}(\mathbf{F}(\mathcal{Q})) = \bigcap_{\boldsymbol{\lambda} \in \mathbb{R}^3} \left\{ \mathbf{x} \in \mathbb{R}^3 : \boldsymbol{\lambda} \cdot \mathbf{x} \leq V(\boldsymbol{\lambda}) \right\}. \quad (\text{A.54})$$

With the same reasoning we obtain

$$\text{conv}(\mathbf{F}(\tilde{\mathcal{Q}})) = \bigcap_{\boldsymbol{\lambda} \in \mathbb{R}^3} \left\{ \mathbf{x} \in \mathbb{R}^3 : \boldsymbol{\lambda} \cdot \mathbf{x} \leq \tilde{V}(\boldsymbol{\lambda}) \right\}, \quad (\text{A.55})$$

where $\tilde{V}(\boldsymbol{\lambda}) := \max_{\mathbf{x} \in \mathbf{F}(\tilde{\mathcal{Q}})} \boldsymbol{\lambda} \cdot \mathbf{x}$. We next show $\tilde{V}(\boldsymbol{\lambda}) \geq V(\boldsymbol{\lambda})$ which already implies (A.48) due to (A.54) and (A.55).

Let $\mathbf{t} = (t_x)_{x \in \mathcal{X} \setminus x_0}$ be $|\mathcal{X}| - 1$ test functions $t_x(p_{\tilde{X}}) := p_{\tilde{X}}(x)$ for all but one $x \in \mathcal{X}$. Choose any $\boldsymbol{\lambda} \in \mathbb{R}^3$ and fix $(U, X, Y, V) \sim p \in \mathcal{Q}$ that achieve $\boldsymbol{\lambda} \cdot \mathbf{F}(p) = V(\boldsymbol{\lambda})$. Define the continuous function

$$f(p_{\tilde{X}}) := \lambda_1(H(V) - H(\tilde{V})) + \lambda_2(H(X) - H(\tilde{X})) + \lambda_3 I(Y; V) \quad (\text{A.56})$$

where $(\tilde{V}, \tilde{Y}, \tilde{X}) \sim p_{V|Y} p_{Y|X} p_{\tilde{X}}$. The point $((p_X(x))_{x \in \mathcal{X} \setminus x_0}, V(\boldsymbol{\lambda}))$ lies in the convex hull of the compact [36 Theorem 26.5] and connected [41 Theorem 4.22] set $(\mathbf{t}, f)(\mathcal{Q})$. Therefore, by the strengthened Carathéodory theorem [12 Theorem 18(ii)], $|\mathcal{X}|$ points suffice, i.e., there exists a random variable U' with $|U'| = |\mathcal{X}|$ and thus $p_{U', X, Y, V} \in \tilde{\mathcal{Q}}$, such that $\mathbb{E}_{U'}[f(p_{X|U'}(\cdot | U'))] = \boldsymbol{\lambda} \cdot \mathbf{F}(p_{U', X, Y, V}) = V(\boldsymbol{\lambda})$. This shows $\tilde{V}(\boldsymbol{\lambda}) \geq V(\boldsymbol{\lambda})$.

By applying the same reasoning to V , one can show that $|\mathcal{V}| = |\mathcal{Y}|$ also suffices.

A.3 Proof of Theorem 7.3

We will prove Theorem 7.3 by showing an inner and an outer bound (Lemmas A.1 and A.2 respectively) and subsequently prove tightness.

LEMMA A.1 We have $\mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})} \subseteq \overline{\mathcal{R}_{\text{MI}}}$ for any $\mathcal{I} \subseteq \mathcal{J}$ and any total order \sqsubset on \mathcal{J} .

Proof. In part, the proof of this lemma closely follows the proof of [25] Theorem 1]. We will use $\mathcal{T}_\varepsilon(X)$ to denote the ε typical sequences [25] Section III].

Pick a total order \sqsubset on \mathcal{J} , a set $\mathcal{I} \subseteq \mathcal{J}$, $(U_{\mathcal{J}}, \emptyset) \in \mathcal{P}_*$, and $(v_0, v_{\mathcal{J}}, R_{\mathcal{J}})$ satisfying (7.4)-(7.8). We will use typicality coding and deterministic binning to obtain a code. Letting $\hat{R}_j = I(U_j; X_j | U_{\sqsubset j})$, we verify for $\mathcal{A} = \sqsubset j$, $j \in \mathcal{J}$, and any $\mathcal{B} \subseteq \mathcal{A}$, that

$$\sum_{j \in \mathcal{B}} \hat{R}_j = \sum_{j \in \mathcal{B}} I(U_j; X_j | U_{\sqsubset j}) \quad (\text{A.57})$$

$$\geq \sum_{j \in \mathcal{B}} I(U_j; X_B | U_{\sqsubset j}, U_{\mathcal{A} \setminus \mathcal{B}}) \quad (\text{A.58})$$

$$= I(U_B; X_B | U_{\mathcal{A} \setminus \mathcal{B}}). \quad (\text{A.59})$$

Following the proof of [25] Theorem 1] and applying the conditional typicality lemma [25] Lemma 3.1.(iv)], we can thus for any $\varepsilon > 0$ and n large enough obtain an $(n, \hat{R}_{\mathcal{J}} + \varepsilon)$ -code $\hat{f}_{\mathcal{J}}$ and for any $\mathcal{A} = \sqsubset j$, $j \in \mathcal{J}$, a decoding function $g_{\mathcal{A}}$, such that, $P\{\mathcal{S}_{\mathcal{A}}\} \geq 1 - \varepsilon$, where $\mathcal{S}_{\mathcal{A}}$ is the ‘success’ event

$$\mathcal{S}_{\mathcal{A}} := \{(\mathbf{Y}, \mathbf{X}_{\mathcal{J}}, g_{\mathcal{A}} \circ \hat{f}_{\mathcal{A}}(\mathbf{X}_{\mathcal{A}})) \in \mathcal{T}_\varepsilon(Y, X_{\mathcal{J}}, U_{\mathcal{A}})\}. \quad (\text{A.60})$$

For $j \notin \mathcal{I}$, we set $f_j := \hat{f}_j$, but for $j \in \mathcal{I}$, we let f_j be typicality encoding *without* binning, in total yielding an $(n, R_{\mathcal{J}} + \varepsilon)$ -code. Moreover, for n large enough and $j \in \mathcal{I}$, we find decoding functions g_j , such that $P\{\mathcal{S}_j\} \geq 1 - \varepsilon$ also for the ‘success’ events

$$\mathcal{S}_j := \{(\mathbf{Y}, \mathbf{X}_{\mathcal{J}}, g_j \circ f_j(\mathbf{X}_j)) \notin \mathcal{T}_\varepsilon(Y, X_{\mathcal{J}}, U_j)\}. \quad (\text{A.61})$$

To shorten notation, let $W_j = f_j(\mathbf{X}_j)$ and $\hat{W}_j := \hat{f}_j(\mathbf{X}_j)$ for $j \in \mathcal{J}$. Pick an arbitrary² $\varepsilon' > 0$. Provided that n is large enough and ε small enough, we have for any $\mathcal{A} = \sqsubset j$ ($j \in \mathcal{J}$),

$$\frac{1}{n} I(\mathbf{Y}; W_{\mathcal{A}}) \geq \frac{1}{n} I(\mathbf{Y}; \hat{W}_{\mathcal{A}}) \quad (\text{A.62})$$

$$\geq \frac{1}{n} I(\mathbf{Y}; g_{\mathcal{A}}(\hat{W}_{\mathcal{A}})) \quad (\text{A.63})$$

$$= H(Y) - \frac{1}{n} H(\mathbf{Y} | g_{\mathcal{A}}(\hat{W}_{\mathcal{A}})) \quad (\text{A.64})$$

$$\geq H(Y) - \frac{1}{n} H(\mathbf{Y}, \mathbf{1}_{\mathcal{S}_{\mathcal{A}}} | g_{\mathcal{A}}(\hat{W}_{\mathcal{A}})) \quad (\text{A.65})$$

$$= H(Y) - \frac{1}{n} H(\mathbf{1}_{\mathcal{S}_{\mathcal{A}}} | g_{\mathcal{A}}(\hat{W}_{\mathcal{A}})) - \frac{1}{n} H(\mathbf{Y} | g_{\mathcal{A}}(\hat{W}_{\mathcal{A}}), \mathbf{1}_{\mathcal{S}_{\mathcal{A}}}) \quad (\text{A.66})$$

$$\geq H(Y) - \varepsilon' - \frac{1}{n} (1 - \varepsilon) H(\mathbf{Y} | g_{\mathcal{A}}(\hat{W}_{\mathcal{A}}), \mathcal{S}_{\mathcal{A}}) - \varepsilon H(Y) \quad (\text{A.67})$$

²In what follows, we will routinely merge expressions that can be made arbitrarily small (for n large and ε sufficiently small) and bound them by ε' .

$$\geq H(Y) - \varepsilon' - \frac{1}{n} H(\mathbf{Y} | g_{\mathcal{A}}(\hat{W}_{\mathcal{A}}), \mathcal{S}_{\mathcal{A}}) \quad (\text{A.68})$$

$$\geq H(Y) - \varepsilon' - \frac{1}{n} \sum_{\mathbf{u}_{\mathcal{A}}} P\{g_{\mathcal{A}}(\hat{W}_{\mathcal{A}}) = \mathbf{u}_{\mathcal{A}} | \mathcal{S}_{\mathcal{A}}\} \log |\mathcal{T}_{\varepsilon}(Y | U_{\mathcal{A}} = \mathbf{u}_{\mathcal{A}})| \quad (\text{A.69})$$

$$\geq H(Y) - H(Y | U_{\mathcal{A}}) - \varepsilon' \quad (\text{A.70})$$

$$= I(U_{\mathcal{A}}; Y) - \varepsilon'. \quad (\text{A.71})$$

Here, (A.62) and (A.63) follow from the data processing inequality [9 Theorem 2.8.1], we applied the entropy bound [9 Theorem 2.6.4] in (A.69) and the cardinality bound for the set of conditionally typical sequences [25 Lemma 3.1.(v)] in (A.70). In particular, for $\mathcal{A} = \mathcal{J}$ we obtain

$$\frac{1}{n} I(\mathbf{Y}; W_{\mathcal{J}}) \geq I(U_{\mathcal{J}}; Y) - \varepsilon' \stackrel{(7.8)}{\geq} v_0 - \varepsilon'. \quad (\text{A.72})$$

For $j \in \mathcal{J}$ and $\mathcal{A} = \square j$ we obtain the following chain of inequalities, where (A.77) and (A.78) will be justified subsequently.

$$\frac{1}{n} I(\mathbf{Y}; W_j) \geq \frac{1}{n} I(\mathbf{Y}; \hat{W}_j) \geq \frac{1}{n} I(\mathbf{Y}; \hat{W}_j | \hat{W}_{\mathcal{A}}) \quad (\text{A.73})$$

$$= \frac{1}{n} I(\mathbf{Y}; \hat{W}_j \hat{W}_{\mathcal{A}}) - \frac{1}{n} I(\mathbf{Y}; \hat{W}_{\mathcal{A}}) \quad (\text{A.74})$$

$$\stackrel{(A.71)}{\geq} I(U_{\mathcal{A}} U_j; Y) - \varepsilon' - \frac{1}{n} I(\mathbf{Y}; \hat{W}_{\mathcal{A}}) \quad (\text{A.75})$$

$$= I(U_{\mathcal{A}} U_j; Y) - \varepsilon' - \frac{1}{n} I(\mathbf{X}_{\mathcal{A}}; \hat{W}_{\mathcal{A}}) + \frac{1}{n} I(\mathbf{X}_{\mathcal{A}}; \hat{W}_{\mathcal{A}} | \mathbf{Y}) \quad (\text{A.76})$$

$$\geq I(U_{\mathcal{A}} U_j; Y) - \varepsilon' - I(X_{\mathcal{A}}; U_{\mathcal{A}}) + H(X_{\mathcal{A}} | Y) - \frac{1}{n} H(\mathbf{X}_{\mathcal{A}} | \hat{W}_{\mathcal{A}}, \mathbf{Y}) \quad (\text{A.77})$$

$$\geq I(U_{\mathcal{A}} U_j; Y) - \varepsilon' - I(X_{\mathcal{A}}; U_{\mathcal{A}}) + H(X_{\mathcal{A}} | Y) - H(X_{\mathcal{A}} | U_{\mathcal{A}}, Y) \quad (\text{A.78})$$

$$= I(U_j; Y | U_{\mathcal{A}}) - \varepsilon' \quad (\text{A.79})$$

$$\stackrel{(7.7)}{\geq} v_j - \varepsilon'. \quad (\text{A.80})$$

Equality in (A.76) follows from the Markov chain $\hat{W}_{\mathcal{A}} \rightarrow \mathbf{X}_{\mathcal{A}} \rightarrow \mathbf{Y}$. In (A.77) we used that for ε small enough and n large enough, we have

$$\frac{1}{n} I(\mathbf{X}_{\mathcal{A}}; \hat{W}_{\mathcal{A}}) = \frac{1}{n} H(\hat{W}_{\mathcal{A}}) \quad (\text{A.81})$$

$$\leq \frac{1}{n} \sum_{j \in \mathcal{A}} H(\hat{W}_j) \quad (\text{A.82})$$

$$\leq \sum_{j \in \mathcal{A}} \left(I(U_j; X_j | U_{\square j}) + \varepsilon \right) \quad (\text{A.83})$$

$$\leq I(U_{\mathcal{A}}; X_{\mathcal{A}}) + \varepsilon', \quad (\text{A.84})$$

where (A.82) follows from the chain rule for entropy [9 Theorem 2.2.1] and the data processing inequality [9 Theorem 2.8.1] and (A.83) follows from the entropy bound [9 Theorem 2.6.4] and the

fact that $f_{\mathcal{J}}$ is an $(n, \hat{R}_{\mathcal{J}} + \varepsilon)$ -code. The inequality (A.78) can be derived similar to (A.71) as for n large enough and ε small enough,

$$\frac{1}{n}H(\mathbf{X}_{\mathcal{A}}|\hat{W}_{\mathcal{A}}, \mathbf{Y}) \leq \frac{1}{n}H(\mathbf{X}_{\mathcal{A}}|g_{\mathcal{A}}(\hat{W}_{\mathcal{A}}), \mathbf{Y}) \quad (\text{A.85})$$

$$\leq \frac{1}{n}H(\mathbf{X}_{\mathcal{A}}, \mathbb{1}_{S_{\mathcal{A}}} | g_{\mathcal{A}}(\hat{W}_{\mathcal{A}}), \mathbf{Y}) \quad (\text{A.86})$$

$$\leq \varepsilon' + \frac{1}{n}H(\mathbf{X}_{\mathcal{A}} | g_{\mathcal{A}}(\hat{W}_{\mathcal{A}}), \mathbf{Y}, S_{\mathcal{A}}) \quad (\text{A.87})$$

$$\leq \varepsilon' + \frac{1}{n} \sum_{\mathbf{u}_{\mathcal{A}}, \mathbf{y}} \mathbb{P}\{g_{\mathcal{A}}(\hat{W}_{\mathcal{A}}) = \mathbf{u}_{\mathcal{A}}, \mathbf{Y} = \mathbf{y} | S_{\mathcal{A}}\} \log |\mathcal{T}_{\varepsilon}(X_{\mathcal{A}} | U_{\mathcal{A}}, Y = \mathbf{u}_{\mathcal{A}}, \mathbf{y})| \quad (\text{A.88})$$

$$\leq \varepsilon' + H(X_{\mathcal{A}} | U_{\mathcal{A}}, Y). \quad (\text{A.89})$$

For $j \in \mathcal{I}$, we have similarly to (A.71) that

$$\frac{1}{n}I(\mathbf{Y}; W_j) \geq \frac{1}{n}I(\mathbf{Y}; g_j(W_j)) \quad (\text{A.90})$$

$$= H(Y) - \frac{1}{n}H(\mathbf{Y} | g_j(W_j)) \quad (\text{A.91})$$

$$\geq H(Y) - \frac{1}{n}H(\mathbf{Y}, \mathbb{1}_{S_j} | g_j(W_j)) \quad (\text{A.92})$$

$$\geq H(Y) - \frac{1}{n}H(\mathbb{1}_{S_j}) - \frac{1}{n}H(\mathbf{Y} | g_j(W_j), \mathbb{1}_{S_j}) \quad (\text{A.93})$$

$$\geq H(Y) - \varepsilon' - \frac{1}{n}H(\mathbf{Y} | g_j(W_j), S_j) \quad (\text{A.94})$$

$$\geq H(Y) - \varepsilon' - \frac{1}{n} \sum_{\mathbf{u}_j} \mathbb{P}\{g_{\mathcal{A}}(W_j) = \mathbf{u}_j | S_j\} \log |\mathcal{T}(Y | U_j = \mathbf{u}_j)| \quad (\text{A.95})$$

$$\geq H(Y) - \varepsilon' - H(Y | U_j) \quad (\text{A.96})$$

$$= I(U_j; Y) - \varepsilon' \stackrel{(7.6)}{\geq} v_j - \varepsilon'. \quad (\text{A.97})$$

Combining (A.72), (A.80) and (A.97) we see that $(v_0 + \varepsilon', v_{\mathcal{J}} + \varepsilon', R_{\mathcal{J}} + \varepsilon) \in \mathcal{R}_{\text{MI}}$, completing the proof as $\varepsilon, \varepsilon'$ were arbitrary. \square

LEMMA A.2 If $(v_0, v_{\mathcal{J}}, R_{\mathcal{J}}) \in \mathcal{R}_{\text{MI}}$, then for all $i \in \mathcal{J}$ and $\mathcal{A} \subseteq \mathcal{J}$,

$$\sum_{j \in \mathcal{A}} R_j - v_i \geq I(X_{\mathcal{A}}; U_{\mathcal{A}} | YQ) - I(Y; U_{i \setminus \mathcal{A}} | Q), \quad (\text{A.98})$$

$$\sum_{j \in \mathcal{A}} R_j \geq I(X_{\mathcal{A}}; U_{\mathcal{A}} | YQ), \text{ and} \quad (\text{A.99})$$

$$\sum_{j \in \mathcal{A}} R_j - v_0 \geq I(X_{\mathcal{A}}; U_{\mathcal{A}} | YQ) - I(Y; U_{\mathcal{J} \setminus \mathcal{A}} | Q), \quad (\text{A.100})$$

for some random variables $(U_{\mathcal{J}}, Q) \in \mathcal{P}_*$.

Proof. For $(v_0, v_{\mathcal{J}}, R_{\mathcal{J}}) \in \mathcal{R}_{\text{MI}}$ we apply Definition 7.1 choosing an $(n, R_{\mathcal{J}})$ -code $f_{\mathcal{J}}$ for $X_{\mathcal{J}}$ and define $U_j := f_j(\mathbf{X}_j)$ for $j \in \mathcal{J}$. In the following, let either $\mathcal{A} = \mathcal{J}$, or $\mathcal{A} = \{j\}$ for $j \in \mathcal{J}$. Slightly abusing

notation, we define $v_{\mathcal{A}} := v_j$ for $\mathcal{A} = \{j\}$ and $v_{\mathcal{A}} := v_0$ for $\mathcal{A} = \mathcal{J}$. We thus have

$$\frac{1}{n} \mathbf{I}(U_{\mathcal{A}}; \mathbf{Y}) \geq v_{\mathcal{A}}. \quad (\text{A.101})$$

With $U_{j,i} := (U_j, \mathbf{X}_{j,1}^{i-1})$ and $Q_i := (\mathbf{Y}^{i-1}, \mathbf{Y}_{i+1}^n)$ we have for every $\mathcal{B} \subseteq \mathcal{J}$,

$$n \sum_{j \in \mathcal{B}} R_j \geq \mathbf{H}(U_{\mathcal{B}}) \quad (\text{A.102})$$

$$= \mathbf{I}(U_{\mathcal{B}}; \mathbf{X}_{\mathcal{B}}) \quad (\text{A.103})$$

$$= \mathbf{I}(U_{\mathcal{B}}; \mathbf{X}_{\mathcal{B}}, \mathbf{Y}) \quad (\text{A.104})$$

$$= \mathbf{I}(U_{\mathcal{B}}; \mathbf{Y}) + \mathbf{I}(U_{\mathcal{B}}; \mathbf{X}_{\mathcal{B}} | \mathbf{Y}) \quad (\text{A.105})$$

$$= \mathbf{I}(U_{\mathcal{A}} U_{\mathcal{B}}; \mathbf{Y}) - \mathbf{I}(U_{\mathcal{A} \setminus \mathcal{B}}; \mathbf{Y} | U_{\mathcal{B}}) + \mathbf{I}(U_{\mathcal{B}}; \mathbf{X}_{\mathcal{B}} | \mathbf{Y}) \quad (\text{A.106})$$

$$= \mathbf{I}(U_{\mathcal{A}}; \mathbf{Y}) + \mathbf{I}(U_{\mathcal{B} \setminus \mathcal{A}}; \mathbf{Y} | U_{\mathcal{A}}) - \mathbf{I}(U_{\mathcal{A} \setminus \mathcal{B}}; \mathbf{Y} | U_{\mathcal{B}}) + \mathbf{I}(U_{\mathcal{B}}; \mathbf{X}_{\mathcal{B}} | \mathbf{Y}) \quad (\text{A.107})$$

$$\stackrel{(\text{A.101})}{\geq} n v_{\mathcal{A}} + \mathbf{I}(U_{\mathcal{B} \setminus \mathcal{A}}; \mathbf{Y} | U_{\mathcal{A}}) - \mathbf{I}(U_{\mathcal{A} \setminus \mathcal{B}}; \mathbf{Y} | U_{\mathcal{B}}) + \mathbf{I}(U_{\mathcal{B}}; \mathbf{X}_{\mathcal{B}} | \mathbf{Y}) \quad (\text{A.108})$$

$$\geq n v_{\mathcal{A}} - \mathbf{I}(U_{\mathcal{A} \setminus \mathcal{B}}; \mathbf{Y}) + \mathbf{I}(U_{\mathcal{B}}; \mathbf{X}_{\mathcal{B}} | \mathbf{Y}) \quad (\text{A.109})$$

$$= \sum_{i=1}^n [v_{\mathcal{A}} - \mathbf{I}(U_{\mathcal{A} \setminus \mathcal{B}}; Y_i | \mathbf{Y}^{i-1}) + \mathbf{I}(U_{\mathcal{B}}; X_{\mathcal{B},i} | \mathbf{Y} \mathbf{X}_{\mathcal{B}}^{i-1})] \quad (\text{A.110})$$

$$\geq \sum_{i=1}^n [v_{\mathcal{A}} - \mathbf{I}(U_{\mathcal{A} \setminus \mathcal{B},i}; Y_i | Q_i) + \mathbf{I}(U_{\mathcal{B}}; X_{\mathcal{B},i} | \mathbf{Y} \mathbf{X}_{\mathcal{B}}^{i-1})] \quad (\text{A.111})$$

$$= \sum_{i=1}^n [v_{\mathcal{A}} - \mathbf{I}(U_{\mathcal{A} \setminus \mathcal{B},i}; Y_i | Q_i) + \mathbf{I}(U_{\mathcal{B},i}; X_{\mathcal{B},i} | Y_i Q_i)]. \quad (\text{A.112})$$

The result now follows by a standard time-sharing argument. Note that the required Markov chains and the independence constraints are satisfied. \square

The following result is a simple corollary of Lemma A.2 and will suffice for us.

COROLLARY A.1 For any $(v_0, v_{\mathcal{J}}, R_{\mathcal{J}}) \in \mathcal{R}_{\text{MI}}$ there are random variables $(U_{\mathcal{J}}, Q) \in \mathcal{P}_*$ with

$$R_j \geq 0, \quad \text{for all } j \in \mathcal{J}, \quad (\text{A.113})$$

$$\sum_{j \in \mathcal{A}} R_j - v_0 \geq \mathbf{I}(X_{\mathcal{A}}; U_{\mathcal{A}} | Y Q) - \mathbf{I}(Y; U_{\mathcal{J} \setminus \mathcal{A}} | Q), \quad \text{for all } \mathcal{A} \subseteq \mathcal{J}, \quad (\text{A.114})$$

$$R_j - v_j \geq \mathbf{I}(X_j; U_j | Y Q), \quad \text{for all } j \in \mathcal{J}, \quad (\text{A.115})$$

$$v_j \leq \mathbf{I}(Y; U_j | Q), \quad \text{for all } j \in \mathcal{J}. \quad (\text{A.116})$$

In the following proof, we will make use of some rather technical results on convex polyhedra, derived in Appendix A.5

Proof of Theorem 7.3 Assume $(v_0, v_{\mathcal{J}}, R_{\mathcal{J}}) \in \mathcal{R}_{\text{MI}}$. We can then find $(U_{\mathcal{J}}, Q) \in \mathcal{P}_*$ such that (A.113)–(A.116) hold. We define $(\tilde{v}_0, \tilde{v}_{\mathcal{J}}) := -(v_0, v_{\mathcal{J}})$ to simplify notation. It is straightforward to check that the inequalities (A.113)–(A.116) define a sequence of closed convex polyhedra $\mathcal{H}^{(j)}$ in the variables $(R_{\mathcal{J}}, \tilde{v}_0, \tilde{v}_{\mathcal{J}})$ that satisfy Assumptions 1 and 2 of Lemma A.8. $\mathcal{H}^{(0)}$ is defined by (A.113) and (A.114) alone, and for $j \in [0:J]$ the polyhedron $\mathcal{H}^{(j)}$ is given in the $K + j$ variables $(R_{\mathcal{J}}, v_0, v_1, v_2, \dots, v_j)$ by

adding constraints (A.115) and (A.116) for each $j \in [1:j]$. The set $\mathcal{H}^{(0)}$ is a supermodular polyhedron [16, Section 2.3] in the K variables $(R_{\mathcal{J}}, \tilde{v}_0)$ on $(\mathcal{K}, 2^{\mathcal{K}})$ with rank function

$$\vartheta(\mathcal{A}) = \begin{cases} 0, & K \notin \mathcal{A}, \\ I(X_{\mathcal{A} \setminus K}; U_{\mathcal{A} \setminus K} | YQ) - I(Y; U_{\mathcal{J} \setminus \mathcal{A}} | Q), & K \in \mathcal{A}, \end{cases} \quad (\text{A.117})$$

where supermodularity follows via standard information-theoretic arguments. By the extreme point theorem [16, Theorem 3.22], every extreme point of $\mathcal{H}^{(0)}$ is associated with a total order \sqsubset on \mathcal{K} . Such an extreme point is given by

$$\begin{aligned} R_j^{(\sqsubset)} &= 0 \text{ for } j \sqsubset K, \\ R_j^{(\sqsubset)} &= I(U_j; X_j | U_{\sqsubset j} Q) \text{ for } j \sqsupset K, \\ v_0^{(\sqsubset)} &= I(Y; U_{\sqsubset K} | Q) - I(Y; U_{\sqsubset K} | YQ). \end{aligned} \quad (\text{A.118})$$

Assumption 3 of Lemma A.8 is now verified by

$$R_j^{(\sqsubset)} \leq I(X_j; U_j | YQ) + I(Y; U_j | Q) = I(X_j; U_j | Q). \quad (\text{A.119})$$

By applying Lemma A.8 we find that every extreme point of $\mathcal{H}^{(j)}$ is given by a subset $\mathcal{I} \subseteq \mathcal{J}$ and an order \sqsubset of \mathcal{K} as

$$R_j^{(\sqsubset, \mathcal{I})} = I(X_j; U_j | Q), \quad j \in \mathcal{I}, \quad (\text{A.120})$$

$$R_j^{(\sqsubset, \mathcal{I})} = 0, \quad j \notin \mathcal{I} \text{ and } j \sqsubset K, \quad (\text{A.121})$$

$$R_j^{(\sqsubset, \mathcal{I})} = I(U_j; X_j | U_{\sqsubset j} Q), \quad j \notin \mathcal{I} \text{ and } j \sqsupset K, \quad (\text{A.122})$$

$$v_j^{(\sqsubset, \mathcal{I})} = I(U_j; Y | Q), \quad j \in \mathcal{I}, \quad (\text{A.123})$$

$$v_j^{(\sqsubset, \mathcal{I})} = -I(U_j; X_j | YQ), \quad j \notin \mathcal{I} \text{ and } j \sqsubset K, \quad (\text{A.124})$$

$$v_j^{(\sqsubset, \mathcal{I})} = I(U_j; Y | U_{\sqsubset j} Q), \quad j \notin \mathcal{I} \text{ and } j \sqsupset K, \quad (\text{A.125})$$

$$v_0^{(\sqsubset, \mathcal{I})} = I(Y; U_{\sqsubset K} | Q) - I(Y; U_{\sqsubset K} | YQ). \quad (\text{A.126})$$

For each $q \in \mathcal{Q}$ with $P\{Q=q\} > 0$ let the point $(v_0^{(\sqsubset, \mathcal{I}, q)}, v_{\mathcal{J}}^{(\sqsubset, \mathcal{I}, q)}, R_{\mathcal{J}}^{(\sqsubset, \mathcal{I}, q)})$ be defined by (A.120)–(A.126) but given $\{Q=q\}$. By substituting $U_j \rightarrow \emptyset$ if $j \notin \mathcal{I}$ and $j \sqsubset K$, we see that $(v_0^{(\sqsubset, \mathcal{I}, q)}, v_{\mathcal{J}}^{(\sqsubset, \mathcal{I}, q)}, R_{\mathcal{J}}^{(\sqsubset, \mathcal{I}, q)}) \in \mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})}$ and consequently $(v_0^{(\sqsubset, \mathcal{I})}, v_{\mathcal{J}}^{(\sqsubset, \mathcal{I})}, R_{\mathcal{J}}^{(\sqsubset, \mathcal{I})}) \in \text{conv}(\mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})})$. Defining the orthant $\mathcal{O} := \{(v_0, v_{\mathcal{J}}, R_{\mathcal{J}}) : v_0 \leq 0, v_{\mathcal{J}} \leq \mathbf{0}, R_{\mathcal{J}} \geq \mathbf{0}\}$, this implies

$$(v_{\mathcal{K}}, R_{\mathcal{J}}) \in \text{conv}\left(\bigcup_{\sqsubset, \mathcal{I}} \text{conv}(\mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})})\right) + \mathcal{O} = \text{conv}\left(\bigcup_{\sqsubset, \mathcal{I}} \mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})}\right) + \text{conv}(\mathcal{O}) \quad (\text{A.127})$$

$$= \text{conv}\left(\bigcup_{\sqsubset, \mathcal{I}} \mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})} + \mathcal{O}\right) \quad (\text{A.128})$$

$$= \text{conv}\left(\bigcup_{\sqsubset, \mathcal{I}} \mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})}\right), \quad (\text{A.129})$$

where (A.128) follows from [43, Theorem 1.1.2] and in (A.129) we used that $\mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})} + \mathcal{O} = \mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})}$ by definition. Together with Lemma A.1 and the convexity of \mathcal{R}_{MI} we obtain

$$\mathcal{R}_{\text{MI}} \subseteq \text{conv} \left(\bigcup_{\sqsubset, \mathcal{I}} \mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})} \right) \subseteq \overline{\mathcal{R}_{\text{MI}}}. \quad (\text{A.130})$$

Note that $\overline{\mathcal{R}_{\text{MI}}}$ is convex by a time-sharing argument.

It remains to show that $\text{conv} \left(\bigcup_{\sqsubset, \mathcal{I}} \mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})} \right)$ is closed. Using Proposition 7.4 we can write $\mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})} = \mathbf{F}^{(\sqsubset, \mathcal{I})}(\mathcal{P}'_*) + \mathcal{O}$, where $\mathcal{P}'_* := \{\mathbf{p}_{Y, X_{\mathcal{J}}, U_{\mathcal{J}}} : (U_{\mathcal{J}}, \emptyset) \in \mathcal{P}_*, |\mathcal{U}_j| = |\mathcal{X}_j| + 4^j, j \in \mathcal{J}\}$ is a compact subset of the probability simplex and $\mathbf{F}^{(\sqsubset, \mathcal{I})}$ is a continuous function, given by the definition of $\mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})}$, (7.4)–(7.8). We can thus write

$$\text{conv} \left(\bigcup_{\sqsubset, \mathcal{I}} \mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})} \right) = \text{conv} \left(\bigcup_{\sqsubset, \mathcal{I}} \mathbf{F}^{(\sqsubset, \mathcal{I})}(\mathcal{P}'_*) + \mathcal{O} \right), \quad (\text{A.131})$$

which is closed by Lemma A.3 □

A.4 Proof of Proposition 7.4

Pick arbitrary $j, j \in \mathcal{J}$. For nonempty $\mathcal{B} \subseteq \mathcal{J}$ with $j \in \mathcal{B}$ we can write $H(X_j|U_{\mathcal{B}}) = \mathbb{E}_{U_j}[f_{j, \mathcal{B}}(\mathbf{p}_{X_j|U_j}(\cdot|U_j))]$ as well as $H(Y|U_{\mathcal{B}}) = \mathbb{E}_{U_j}[g_{\mathcal{B}}(\mathbf{p}_{X_j|U_j}(\cdot|U_j))]$, where

$$f_{j, \mathcal{B}}(\mathbf{p}_{X_j|U_j}(\cdot|u_j)) := H(X_j|U_{\mathcal{B} \setminus j}, U_j = u_j), \quad (\text{A.132})$$

$$g_{\mathcal{B}}(\mathbf{p}_{X_j|U_j}(\cdot|u_j)) := H(Y|U_{\mathcal{B} \setminus j}, U_j = u_j). \quad (\text{A.133})$$

Observe that $f_{j, \mathcal{B}}$ and $g_{\mathcal{B}}$ are continuous functions of $\mathbf{p}_{X_j|U_j}(\cdot|u_j)$. Apply the support lemma [13, Appendix C] with the functions $f_{j, \mathcal{B}}$ and $g_{\mathcal{B}}$ for all $j \in \mathcal{J}$, $j \in \mathcal{B} \subseteq \mathcal{J}$, and $|\mathcal{X}_j| - 1$ test functions, which guarantee that the marginal distribution \mathbf{p}_{X_j} does not change. We obtain a new random variable \hat{U}_j with $H(X_j|U_{\mathcal{B} \setminus j} \hat{U}_j) = H(X_j|U_{\mathcal{B}})$ and $H(Y|U_{\mathcal{B} \setminus j} \hat{U}_j) = H(Y|U_{\mathcal{B}})$. By rewriting (7.4)–(7.8) in terms of conditional entropies, it is evident that the defining inequalities for $\mathcal{R}_{\text{MI}}^{(\sqsubset, \mathcal{I})}$ remain the same when replacing U_j by \hat{U}_j . The support of \hat{U}_j satisfies the required cardinality bound³

$$|\hat{\mathcal{U}}_j| \leq |\mathcal{X}_j| - 1 + J2^{J-1} + 2^{J-1} \quad (\text{A.134})$$

$$\leq |\mathcal{X}_j| + 4^J. \quad (\text{A.135})$$

The same process is repeated for every $j \in \mathcal{J}$.

A.5 Results on convex polyhedra

We start this appendix with a simple lemma, which will be used in several proofs.

LEMMA A.3 For a compact set $\mathcal{C} \subseteq \mathbb{R}^n$ and a closed, convex set $\mathcal{B} \subseteq \mathbb{R}^n$,

$$\mathcal{A} := \text{conv}(\mathcal{C} + \mathcal{B}) = \text{conv}(\mathcal{C}) + \mathcal{B} = \overline{\mathcal{A}}. \quad (\text{A.136})$$

³There are J ways to choose j and 2^{J-1} ways to choose \mathcal{B} .

Proof. We have $\mathcal{A} = \text{conv}(\mathcal{C}) + \text{conv}(\mathcal{B}) = \text{conv}(\mathcal{C}) + \mathcal{B}$ by [43 Theorem 1.1.2] and the convexity of \mathcal{B} . Note that $\text{conv}(\mathcal{C})$ is compact by [22 Theorem 2.3.4]. \mathcal{A} is the sum of a compact set and a closed set and, hence, closed [42 Exercise 1.3(e)]. \square

Let \mathcal{H} be the convex polyhedron $\mathcal{H} := \{\mathbf{x} \in \mathbb{R}^n : \mathbf{A}\mathbf{x} \geq \mathbf{b}\}$ for an $m \times n$ matrix $\mathbf{A} = (\mathbf{a}_{(1)}, \mathbf{a}_{(2)}, \dots, \mathbf{a}_{(m)})^T$ and $\mathbf{b} \in \mathbb{R}^m$, where $\mathbf{a}_{(j)}^T$ is the j th row of \mathbf{A} . In this section we will use the notation of [22]. In particular, we shall call a closed convex set *line-free* if it does not contain a (straight) line. The *characteristic cone* of a closed convex set \mathcal{C} is defined as $\text{cc}(\mathcal{C}) := \{\mathbf{y} : \mathbf{x} + \lambda\mathbf{y} \in \mathcal{C} \text{ for all } \lambda \geq 0\}$ ($\mathbf{x} \in \mathcal{C}$ arbitrary) and $\text{ext}(\mathcal{C})$ is the set of all *extreme points* of \mathcal{C} , i.e., points $\mathbf{x} \in \mathcal{C}$ that cannot be written as $\mathbf{x} = \lambda\mathbf{y} + (1 - \lambda)\mathbf{z}$ with $\mathbf{y}, \mathbf{z} \in \mathcal{C}$, $\mathbf{y} \neq \mathbf{z}$ and $\lambda \in (0, 1)$.

LEMMA A.4 A point \mathbf{y} is in $\text{cc}(\mathcal{H})$ if and only if $\mathbf{A}\mathbf{y} \geq \mathbf{0}$.

Proof. If $\mathbf{A}\mathbf{y} \geq \mathbf{0}$, $\mathbf{x} \in \mathcal{H}$ and $\lambda \geq 0$, $\mathbf{A}(\mathbf{x} + \lambda\mathbf{y}) \geq \mathbf{A}\mathbf{x} \geq \mathbf{b}$. On the other hand, for $\mathbf{a}_{(j)}^T \mathbf{y} < 0$, we have $\mathbf{a}_{(j)}^T (\mathbf{x} + \lambda\mathbf{y}) < b_j$ for $\lambda > \frac{b_j - \mathbf{a}_{(j)}^T \mathbf{x}}{\mathbf{a}_{(j)}^T \mathbf{y}} > 0$. \square

LEMMA A.5 If, for every $i \in [1:n]$, there exists $j \in [1:m]$ such that $\mathbf{e}_i = \mathbf{a}_{(j)}$ and for every $j \in [1:m]$, $\mathbf{a}_{(j)} \geq \mathbf{0}$, then \mathcal{H} is line-free and $\text{cc}(\mathcal{H}) = \mathbb{R}_+^n$.

Proof. For any $\mathbf{y} \in \mathbb{R}_+^n$, clearly $\mathbf{A}\mathbf{y} \geq \mathbf{0}$ and hence $\mathbf{y} \in \text{cc}(\mathcal{H})$ by Lemma A.4. If $\mathbf{y} \notin \mathbb{R}_+^n$ we have $y_i < 0$ for some $i \in [1:n]$ and choose $j \in [1:m]$ such that $\mathbf{a}_{(j)} = \mathbf{e}_i$, resulting in $\mathbf{a}_{(j)}^T \mathbf{y} = y_i < 0$. To show that \mathcal{H} is line-free assume that $\mathbf{x} + \lambda\mathbf{y} \in \mathcal{H}$ for all $\lambda \in \mathbb{R}$. This implies $\pm\mathbf{y} \in \text{cc}(\mathcal{H})$, i.e., $\mathbf{y} = \mathbf{0}$. \square

DEFINITION A.1 A point \mathbf{x} is on an *extreme ray* of the cone $\text{cc}(\mathcal{H})$ if the decomposition $\mathbf{x} = \mathbf{y} + \mathbf{z}$ with $\mathbf{y}, \mathbf{z} \in \text{cc}(\mathcal{H})$ implies that $\mathbf{y} = \lambda\mathbf{z}$ for some $\lambda \in \mathbb{R}$.

It is easy to see that the points on extreme rays of \mathcal{O} are given by $\mathbf{x} = \lambda\mathbf{e}_i$ for $\lambda \geq 0$ and $i \in [1:n]$.

Define $\mathcal{A}(\mathbf{x}) := \{j \in [1:m] : \mathbf{a}_{(j)}^T \mathbf{x} = b_j\}$. We say that exactly n_0 linearly independent inequalities are satisfied with equality at \mathbf{x} , if $\mathbf{A}\mathbf{x} \geq \mathbf{b}$ and $(\mathbf{a}_{(j)})_{j \in \mathcal{A}(\mathbf{x})}$ has rank n_0 .

LEMMA A.6 $\mathbf{x} \in \text{ext}(\mathcal{H})$ if and only if exactly n linearly independent inequalities are satisfied with equality at \mathbf{x} .

Proof. Assuming that less than n linearly independent inequalities are satisfied with equality at \mathbf{x} , we find $\mathbf{0} \neq \mathbf{c} \perp (\mathbf{a}_{(j)})_{j \in \mathcal{A}(\mathbf{x})}$ and thus $\mathbf{x} \pm \varepsilon\mathbf{c} \in \mathcal{H}$ for a small $\varepsilon > 0$, showing that $\mathbf{x} \notin \text{ext}(\mathcal{H})$.

Conversely assume $\mathbf{x} \notin \text{ext}(\mathcal{H})$, i.e., $\mathbf{x} = \lambda\mathbf{x}' + (1 - \lambda)\mathbf{x}''$ for $\lambda \in (0, 1)$ and $\mathbf{x}', \mathbf{x}'' \in \mathcal{H}$, $\mathbf{x}' \neq \mathbf{x}''$. For any $j \in \mathcal{A}(\mathbf{x})$, we then have $\lambda\mathbf{a}_{(j)}^T \mathbf{x}' + (1 - \lambda)\mathbf{a}_{(j)}^T \mathbf{x}'' = b_j$, which implies $\mathbf{a}_{(j)}^T \mathbf{x}' = \mathbf{a}_{(j)}^T \mathbf{x}'' = b_j$ and therefore $\mathbf{0} \neq \mathbf{x}' - \mathbf{x}'' \perp (\mathbf{a}_{(j)})_{j \in \mathcal{A}(\mathbf{x})}$. \square

LEMMA A.7 Assuming that \mathcal{H} is line free and that exactly $n - 1$ linearly independent inequalities are satisfied with equality at \mathbf{x} . Then either $\mathbf{x} = \lambda\mathbf{c} + (1 - \lambda)\mathbf{d}$ where $\lambda \in (0, 1)$ and $\mathbf{c}, \mathbf{d} \in \text{ext}(\mathcal{H})$ or $\mathbf{x} = \mathbf{c} + \mathbf{d}$ where $\mathbf{c} \in \text{ext}(\mathcal{H})$ and $\mathbf{d} \neq \mathbf{0}$ lies on an extreme ray of $\text{cc}(\mathcal{H})$.

Proof. We obtain $\mathbf{0} \neq \mathbf{r} \perp (\mathbf{a}_{(j)})_{j \in \mathcal{A}(\mathbf{x})}$. Define $\lambda_1 := \inf\{\lambda : \mathbf{x} + \lambda\mathbf{r} \in \mathcal{H}\}$ and $\lambda_2 := \sup\{\lambda : \mathbf{x} + \lambda\mathbf{r} \in \mathcal{H}\}$. Clearly $\lambda_1 \leq 0 \leq \lambda_2$. As \mathcal{H} is line free, we may assume without loss of generality $\lambda_1 = -1$ (note that $\mathbf{x} \notin \text{ext}(\mathcal{H})$) and set $\mathbf{c} = \mathbf{x} - \mathbf{r}$. We now have $\mathbf{c} \in \text{ext}(\mathcal{H})$ as otherwise $\mathbf{c} - \varepsilon\mathbf{r} \in \mathcal{H}$ for some small $\varepsilon > 0$.

If $\lambda_2 < \infty$, define $\mathbf{d} = \mathbf{x} + \lambda_2\mathbf{r}$, which yields $\mathbf{d} \in \text{ext}(\mathcal{H})$ and $\mathbf{x} = \lambda\mathbf{c} + (1 - \lambda)\mathbf{d}$ with $\lambda = \frac{\lambda_2}{\lambda_2 + 1}$. Note that $\lambda_2 \neq 0$ as $\mathbf{x} \notin \text{ext}(\mathcal{H})$.

If $\lambda_2 = \infty$ we have $\mathbf{x} - \mathbf{c} = \mathbf{r} \in \text{cc}(\mathcal{H})$. We need to show that \mathbf{r} is also on an extreme ray of $\text{cc}(\mathcal{H})$.

Assuming $\mathbf{r} = \mathbf{r}' + \mathbf{r}''$ with $\mathbf{r}', \mathbf{r}'' \in \text{cc}(\mathcal{H})$ yields $\mathbf{a}_{(j)}^T(\mathbf{r}' + \mathbf{r}'') = 0$, which implies $\mathbf{a}_{(j)}^T \mathbf{r}' = \mathbf{a}_{(j)}^T \mathbf{r}'' = 0$ for every $j \in \mathcal{A}(\mathbf{x})$ by Lemma [A.4] \square

For each $j \in [0:J]$, define the closed convex polyhedron $\mathcal{H}^{(j)} := \{\mathbf{x} \in \mathbb{R}^{K+j} : \mathbf{A}^{(j)} \mathbf{x} \geq \mathbf{b}^{(j)}\}$, where $\mathbf{A}^{(j)}$ is a matrix and $\mathbf{b}^{(j)}$ a vector of appropriate dimension. We make the following three assumptions:

1. $\mathbf{A}^{(j)}$ and $\mathbf{b}^{(j)}$ are defined recursively as

$$\mathbf{A}^{(j)} := \begin{pmatrix} \mathbf{A}^{(j-1)} & \mathbf{0} \\ \mathbf{0}^T & 1 \\ \mathbf{e}_j^T & 1 \end{pmatrix}, \quad \mathbf{b}^{(j)} = \begin{pmatrix} \mathbf{b}^{(j-1)} \\ c_1^{(j)} \\ c_2^{(j)} \end{pmatrix}, \quad (\text{A.137})$$

where $c_1^{(j)}$ and $c_2^{(j)}$ are arbitrary reals.

2. Each entry of $\mathbf{A}^{(0)}$ equals 0 or 1 and for all $k \in \mathcal{K}$ at least one row of $\mathbf{A}^{(0)}$ is equal to \mathbf{e}_k^T . Due to Assumption [1] this also implies that each entry of $\mathbf{A}^{(j)}$ is in $\{0, 1\}$ and for all $k \in [1:K+j]$ at least one row of $\mathbf{A}^{(j)}$ is equal to \mathbf{e}_k^T .
3. For any extreme point $\mathbf{x} \in \text{ext}(\mathcal{H}^{(0)})$ and any $j \in \mathcal{J}$, assume $x_j \leq c_2^{(j)} - c_1^{(j)}$.

LEMMA A.8 Under Assumptions [1] to [3] for every $j \in [0:J]$ and every extreme point $\mathbf{y} \in \text{ext}(\mathcal{H}^{(j)})$ there is an extreme point $\mathbf{x} \in \text{ext}(\mathcal{H}^{(0)})$ and a subset $\mathcal{I}_j \subseteq [1:j]$ such that $y_K = x_K$ and for every $j \in \mathcal{J}$,

$$y_j = \begin{cases} x_j, & j \notin \mathcal{I}_j, \\ c_2^{(j)} - c_1^{(j)}, & j \in \mathcal{I}_j, \end{cases} \quad (\text{A.138})$$

and for every $j \in [1:j]$,

$$y_{K+j} = \begin{cases} c_2^{(j)} - x_j, & j \notin \mathcal{I}_j, \\ c_1^{(j)}, & j \in \mathcal{I}_j. \end{cases} \quad (\text{A.139})$$

Proof. For every $j \in \mathcal{J}$, $\mathcal{H}^{(j)}$ is line free by Assumption [2] and Lemma [A.5] and can be written [22, Lemma 6, p. 25] as $\mathcal{H}^{(j)} = \text{cc}(\mathcal{H}^{(j)}) + \text{conv}(\text{ext}(\mathcal{H}^{(j)}))$. Lemma [A.5] also implies $\text{cc}(\mathcal{H})^{(j)} = \mathcal{O}$.

Let us proceed inductively over $j \in [0:J]$. For $j = 0$ the statement is trivial. Given any $\mathbf{y} \in \text{ext}(\mathcal{H}^{(j)})$, we need to obtain $\mathbf{x} \in \text{ext}(\mathcal{H}^{(0)})$ and \mathcal{I}_j such that \mathbf{y} is given according to [A.138] and [A.139]. Let $\mathbf{z} = \mathbf{y}_1^{K+j-1}$ be the truncation of \mathbf{y} . Exactly $K+j$ linear independent inequalities of $\mathbf{A}^{(j)} \mathbf{y} \geq \mathbf{b}^{(j)}$ are satisfied with equality by Lemma [A.6] which is possible in only two different ways:

- *Construction I:* Exactly $K+j-1$ linear independent inequalities of $\mathbf{A}^{(j-1)} \mathbf{z} \geq \mathbf{b}^{(j-1)}$ are satisfied with equality, i.e., $\mathbf{z} \in \text{ext}(\mathcal{H}^{(j-1)})$ by Lemma [A.6] and at least one of

$$y_{K+j} \geq c_1^{(j)}, \quad (\text{A.140})$$

$$y_j + y_{K+j} \geq c_2^{(j)}, \quad (\text{A.141})$$

is satisfied with equality.

As $\mathbf{z} \in \text{ext}(\mathcal{H}^{(j-1)})$, there exists $\mathbf{x} \in \text{ext}(\mathcal{H}^{(0)})$ and \mathcal{I}_{j-1} such that (A.138) holds for $j \in \mathcal{J}$ and (A.139) holds for $j \in [1:j-1]$ by the induction hypothesis. In particular $y_j = x_j$. Assuming that (A.141) holds with equality, we have $y_{K+j} = c_2^{(j)} - x_j$. Thus, the point \mathbf{x} together with $\mathcal{I}_j = \mathcal{I}_{j-1}$ yields \mathbf{y} from (A.138) and (A.139). Equality in (A.140) implies equality in (A.141) by Assumption 3.

- *Construction II:* Exactly $K + j - 2$ linear independent inequalities of $\mathbf{A}^{(j-1)} \mathbf{z} \geq \mathbf{b}^{(j-1)}$ are satisfied with equality and (A.140) and (A.141) are both satisfied with equality as well. Additionally, these $K + j$ inequalities together need to be linearly independent. This can occur in two different ways by Lemma A.7.

Assume $\mathbf{z} = \lambda \mathbf{x} + (1 - \lambda) \mathbf{x}'$ for $\mathbf{x}, \mathbf{x}' \in \text{ext}(\mathcal{H}^{(j-1)})$, $\mathbf{x} \neq \mathbf{x}'$ and $\lambda \in (0, 1)$. This implies $y_{K+j} = c_1^{(j)}$ and $y_j = \lambda x_j + (1 - \lambda) x'_j = c_2^{(j)} - c_1^{(j)}$, which by Assumption 3 already implies $x_j = x'_j = c_2^{(j)} - c_1^{(j)}$. Thus, (A.140) and (A.141) are satisfied (with equality) for every $\lambda \in [0, 1]$ and \mathbf{y} cannot be an extreme point as it can be written as a non-trivial convex combination.

We can thus focus on the second option which is that \mathbf{z} is on an extreme ray of $\mathcal{H}^{(j-1)}$, i.e., $\mathbf{z} = \mathbf{x} + \lambda \mathbf{e}_{j'}$ for some $\mathbf{x} \in \text{ext}(\mathcal{H}^{(j-1)})$, $\lambda > 0$ and $j' \in [1:K+j-1]$. If $j' \neq j$, (A.140) and (A.141) are satisfied for all $\lambda > 0$ and thus \mathbf{y} cannot be an extreme point because it can be written as a non-trivial convex combination. For $j' = j$ the point \mathbf{x} with $\mathcal{I}_j = \mathcal{I}_{j-1} \cup j$ yields the desired extreme point.

□

B. Bibliography

- [1] AHLWEDE, R. & CSISZÁR, I. (1986) Hypothesis Testing with Communication Constraints. *IEEE Trans. Inf. Theory*, **32**(4), 533–542.
- [2] AHLWEDE, R. & KÖRNER, J. (1975) Source Coding with Side Information and a Converse for Degraded Broadcast Channels. *IEEE Trans. Inf. Theory*, **21**(6), 629–637.
- [3] ——— (1977) On the connection between the entropies of input and output distributions of discrete memoryless channels. in *Proc. 5th Conf. Probability Theory, Sep. 1974*, pp. 13–23, Brasov, Romania.
- [4] ALIPRANTIS, C. D. & BORDER, K. C. (2006) *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, 3 edn.
- [5] CHAPMAN, C. (2017) personal communication.
- [6] CHENG, Y. & CHURCH, G. M. (2000) Biclustering of Expression Data. in *Proc. 8th Int. Conf. Intelligent Syst. for Molecular Biology*, vol. 8, pp. 93–103, San Diego, CA.
- [7] COURTADE, T. A. & KUMAR, G. R. (2014) Which Boolean Functions Maximize Mutual Information on Noisy Inputs?. *IEEE Trans. Inf. Theory*, **60**(8), 4515–4525.
- [8] COURTADE, T. A. & WEISSMAN, T. (2014) Multiterminal Source Coding Under Logarithmic Loss. *IEEE Trans. Inf. Theory*, **60**(1), 740–761.

- [9] COVER, T. M. & THOMAS, J. A. (2006) *Elements of Information Theory*. John Wiley & Sons.
- [10] DHILLON, I. S., MALLELA, S. & MODHA, D. S. (2003) Information-theoretic Co-clustering. in *Proc. 9th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, pp. 89–98, Washington, DC.
- [11] EATON, J. W. (1993) GNU Octave. Free Software Foundation.
- [12] EGGLESTON, H. G. (1958) *Convexity*. Cambridge University Press.
- [13] EL GAMAL, A. & KIM, Y.-H. (2011) *Network Information Theory*. Cambridge University Press.
- [14] EL GAMAL, A. A. & COVER, T. M. (1982) Achievable Rates for Multiple Descriptions. *IEEE Trans. Inf. Theory*, **28**(6), 851–857.
- [15] ERKIP, E. & COVER, T. M. (1998) The Efficiency of Investment Information. *IEEE Trans. Inf. Theory*, **44**(3), 1026–1040.
- [16] FUJISHIGE, S. (2005) *Submodular Functions and Optimization*, Annals of Discrete Mathematics. Elsevier Science, 2 edn.
- [17] GÁCS, P. & KÖRNER, J. (1973) Common Information Is Far Less Than Mutual Information. *Problems of Control and Inform. Theory*, **2**, 149–162.
- [18] GILAD-BACHRACH, R., NAVOT, A. & TISHBY, N. (2003) An Information Theoretic Trade-off between Complexity and Accuracy. in *Learning Theory and Kernel Machines*, pp. 595–609. Springer.
- [19] GOHARI, A. A. & ANANTHARAM, V. (2012) Evaluation of Marton’s Inner Bound for the General Broadcast Channel. *IEEE Trans. Inf. Theory*, **58**(2), 608–619.
- [20] GOKCAY, E. & PRINCIPE, J. C. (2002) Information Theoretic Clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **24**(2), 158–171.
- [21] GOYAL, V. K. (2001) Multiple Description Coding: Compression meets the Network. *IEEE Signal Process. Mag.*, **18**(5), 74–93.
- [22] GRÜNBAUM, B. (2003) *Convex Polytopes*. Springer, New York.
- [23] HAN, T. S. (1987) Hypothesis Testing with Multiterminal Data Compression. *IEEE Trans. Inf. Theory*, **33**(6), 759–772.
- [24] HAN, T. S. & AMARI, S. (1998) Statistical Inference under Multiterminal Data Compression. *IEEE Trans. Inf. Theory*, **44**(6), 2300–2324.
- [25] HAN, T. S. & KOBAYASHI, K. (1980) A Unified Achievable Rate Region for a General Class of Multiterminal Source Coding Systems. *IEEE Trans. Inf. Theory*, **26**(3), 277–288.
- [26] HARTIGAN, J. A. (1972) Direct Clustering of a Data Matrix. *Journal of the American Statistical Association*, **67**(337), 123–129.
- [27] HEIMLICH, O. (2015) GNU Octave Interval Package. .

- [28] JOG, V. & NAIR, C. (2010) An information inequality for the BSSC broadcast channel. in *Inform. Theory and Applicat. Workshop (ITA)*, pp. 1–8, San Diego, CA.
- [29] KLOTZ, J. G., KRACHT, D., BOSSERT, M. & SCHÖBER, S. (2014) Canalizing Boolean Functions Maximize Mutual Information. *IEEE Trans. Inf. Theory*, **60**(4), 2139–2147.
- [30] KRASKOV, A. & GRASSBERGER, P. (2009) MIC: Mutual Information Based Hierarchical Clustering. in *Information Theory and Statistical Learning*, pp. 101–123. Springer.
- [31] KUMAR, G. R. & COURTADE, T. A. (2013) Which Boolean Functions are Most Informative?. in *Proc. IEEE Int. Symp. on Inform. Theory*, pp. 226–230, Istanbul, Turkey.
- [32] MADEIRA, S. C. & OLIVEIRA, A. L. (2004) Biclustering Algorithms for Biological Data Analysis: A Survey. *IEEE/ACM Trans. Comput. Biol. Bioinformatics*, **1**(1), 24–45.
- [33] MIRKIN, B. (1996) *Mathematical Classification and Clustering*. Kluwer Academic Publisher.
- [34] MOORE, R. E., KEARFOTT, R. B. & CLOUD, M. J. (2009) *Introduction to Interval Analysis*. SIAM.
- [35] MÜLLER, A. C., NOWOZIN, S. & LAMPERT, C. H. (2012) Information Theoretic Clustering Using Minimum Spanning Trees. in *Joint DAGM (German Association for Pattern Recognition) and OAGM Symposium*, pp. 205–215. Springer.
- [36] MUNKRES, J. R. (2000) *Topology*. Prentice Hall.
- [37] NAIR, C. (2013) Upper concave envelopes and auxiliary random variables. *Int. J. of Advances in Eng. Sciences and Appl. Math.*, **5**(1), 12–20.
- [38] PICHLER, G. (2017) DSBS-MutInf-counterexample. URL: <https://doi.org/10.5281/zenodo.1042588>
- [39] PICHLER, G., MATZ, G. & PIANTANIDA, P. (2016) A Tight Upper Bound on the Mutual Information of Two Boolean Functions. in *Proc. Inform. Theory Workshop*, pp. 16–20, Cambridge, UK.
- [40] PICHLER, G., PIANTANIDA, P. & MATZ, G. (2016) Distributed Information-Theoretic Biclustering. in *Proc. IEEE Int. Symp. on Inform. Theory*, pp. 1083–1087, Barcelona, Spain.
- [41] RUDIN, W. (1976) *Principles of Mathematical Analysis*. McGraw-Hill, 3 edn.
- [42] ——— (1991) *Functional Analysis*. McGraw-Hill, 2 edn.
- [43] SCHNEIDER, R. (2014) *Convex Bodies: The Brunn-Minkowski Theory*. Cambridge University Press, 2 edn.
- [44] SHANNON, C. E. (1993) Coding Theorems for a Discrete Source With a Fidelity Criterion. in *Claude Elwood Shannon: collected papers*, ed. by N. J. A. Sloane, & A. D. Wyner, pp. 325–350. IEEE Press.
- [45] SHARAN, R. (2006) Analysis of Biological Networks: Network Modules – Clustering and Biclustering. lecture notes.

- [46] SLONIM, N., ATWAL, G. S., TKAČIK, G. & BIALEK, W. (2005) Information-based clustering. *Proc. of the Nat. Academy of Sciences of the United States of America*, **102**(51), 18297–18302.
- [47] STEEG, G. V., GALSTYAN, A., SHA, F. & DEDEO, S. (2014) Demystifying Information-Theoretic Clustering. in *International Conference on Machine Learning*, pp. 19–27.
- [48] TANAY, A., SHARAN, R. & SHAMIR, R. (2005) Biclustering Algorithms: A Survey. *Handbook of Computational Molecular Biology*, **9**(1-20), 122–124.
- [49] TISHBY, N., PEREIRA, F. C. & BIALEK, W. (1999) The information bottleneck method. in *Proc. 37th Annual Allerton Conference on Communication, Control, and Computing*, pp. 368–377, Monticello, IL.
- [50] TUNG, S.-Y. (1978) Multiterminal Source Coding. Ph.D. thesis, Cornell University.
- [51] WAGNER, A. B., KELLY, B. G. & ALTUG, Y. (2011) Distributed Rate-distortion with Common Components. *IEEE Trans. Inf. Theory*, **57**(7), 4035–4057.
- [52] WESTOVER, M. B. & O’SULLIVAN, J. A. (2008) Achievable Rates for Pattern Recognition. *IEEE Trans. Inf. Theory*, **54**(1), 299–320.
- [53] WITSENHAUSEN, H. S. (1974) Entropy Inequalities for Discrete Channels. *IEEE Trans. Inf. Theory*, **20**(5), 610–616.
- [54] ——— (1975) On Sequences of Pairs of Dependent Random Variables. *SIAM Journal on Applied Mathematics*, **28**(1), 100–113.
- [55] WITSENHAUSEN, H. S. & WYNER, A. D. (1975) A Conditional Entropy Bound for a Pair of Discrete Random Variables. *IEEE Trans. Inf. Theory*, **21**(5), 493–501.
- [56] WYNER, A. (1973) A Theorem on the Entropy of Certain Binary Sequences and Applications: Part II. *IEEE Trans. Inf. Theory*, **19**(6), 772–777.
- [57] WYNER, A. & ZIV, J. (1973) A Theorem on the Entropy of Certain Binary Sequences and Applications: Part I. *IEEE Trans. Inf. Theory*, **19**(6), 769–772.
- [58] WYNER, A. D. (1975) On Source Coding with Side Information at the Decoder. *IEEE Trans. Inf. Theory*, **21**(3), 294–300.
- [59] WYNER, A. D. & ZIV, J. (1976) The Rate-distortion Function for Source Coding with Side Information at the Decoder. *IEEE Trans. Inf. Theory*, **22**(1), 1–10.