

A Class of Nonbinary Symmetric Information Bottleneck Problems

Michael Dikshtein and Shlomo Shamai (Shitz)

Technion–Israel Institute of Technology

Department of Electrical and Computer Engineering, Haifa 3200003, Israel

Email: {michaeldic@campus, sshlomo@ee}.technion.ac.il

Abstract—We study two dual settings of information processing. Let $Y \rightarrow X \rightarrow W$ be a Markov chain with fixed joint probability mass function P_{XY} and a mutual information constraint on the pair (W, X) . For the first problem, known as Information Bottleneck, we aim to maximize the mutual information between the random variables Y and W , while for the second problem, termed as Privacy Funnel, our goal is to minimize it. In particular, we analyze the scenario for which X is the input, and Y is the output of modulo-additive noise channel. We provide analytical characterization of the optimal information rates and the achieving distributions.

I. INTRODUCTION

Let (X, Y) be a pair of random variables specified by a fixed bivariate distribution P_{XY} , of cardinality $|\mathcal{X}| = n$, and respectively $|\mathcal{Y}| = m$. Consider all random variables W satisfying the Markov chain $Y \rightarrow X \rightarrow W$ subject to a constraint on the mutual information of the pair (X, W) . We consider here two extremes of the information processing problem, the Information Bottleneck (IB) function and the Privacy Funnel (PF).

The IB optimization problem, introduced by Tishby et al. [1], is defined as

$$\begin{aligned} R_{P_{XY}}^{\text{IB}}(C) \triangleq & \underset{P_{W|X}}{\text{maximize}} \quad I(Y; W) \\ & \text{subject to} \quad I(X; W) \leq C. \end{aligned} \quad (1)$$

This problem is illustrated in Figure 1. In our study we aim to determine the maximum value and characterize the achieving conditional distribution $P_{W|X}$ (test channels) of (1) for a class of symmetric channels $P_{Y|X}$, and constraints C .

The motivation to study such a model is as follows. Consider a latent random variable Y , which constitutes the Markov chain $Y \rightarrow X \rightarrow W$ and represents a source of information. The user observes a noisy version of Y , i.e., X , and then tries to compress the observed noisy data such that its reconstructed version, W , will be comparable under the maximum mutual information metric to the original data Y . Thus, (1) is essentially a remote source coding problem [2], choosing the distortion measure as the logarithmic-loss. Here W represents the noisy version (X) of the source (Y) with a constrained number of bits ($I(X; W) \leq C$), and the goal is to maximize the relevant information in W regarding Y (measured by the mutual information between Y and W). In the standard IB terminology, $I(X; W)$ is referred to as the complexity of W , and $I(Y; W)$ is referred to as the relevance of W .




Fig. 1: Block diagram of the Information Bottleneck function.

For the particular case where (Y, X, W) are discrete random variables, an optimal $P_{W|X}$ can be found by iteratively solving a set of self-consistent equations [1]. A generalized Blahuto-Arimoto algorithm [3] was proposed to solve those equations. The optimal test-channel $P_{W|X}$ was characterized using a variation principle in [1]. A particular case of deterministic mappings from X to W was considered in [4], and algorithms that find those mappings were described. Unfortunately, since the underlying optimization problem in (1) is not convex, there are no theoretical guarantees for convergence of the proposed iterative algorithms.

There are two cases for which the solution of (1) is thoroughly characterized. The first one, considered in [5], is where the pair (X, Y) is a Doubly Symmetric Binary Source (DSBS) with transition probability p . It was shown that the optimal test channel $P_{W|X}$ is a BSC with transition probability $h_2^{-1}(1 - C)$ where $h_2(\cdot)$ is binary entropy function and $h_2^{-1}(\cdot)$ its inverse. The converse can be established by applying Mrs. Gerber's Lemma [6]. This setting was also solved as an example in [7, Section IV.A]. The optimality of BSC test-channel extends also to a Binary Memoryless Symmetric (BMS) channel [8, Ch. 4] from X to Y , as [9, Theorem 2] implies.

The second case, first considered in [10], is where (X, Y) are jointly Gaussian. It was shown that the optimal distribution of (Y, X, W) is also jointly Gaussian. The optimality of the Gaussian test-channel can be proved using conditional Entropy Power Inequality [11, Ch. 2]. It can also be established using I-MMSE and Single Crossing Property [12]. Moreover, under the I-MMSE framework, the proof can be easily extended to Jointly Gaussian Random Vectors (\mathbf{X}, \mathbf{Y}) [13].

The IB method can also be seen as a variation on some closely related problems in the Information Theory literature. A bound on the conditional entropy for a pair of discrete random variables subject to entropy constraint has been consid-

ered in [7] as a method to characterize common information [14]. A method based on convex analysis was proposed to find the achieving distributions and several important examples were given. We will show that the problem addressed in [7] is equivalent to (1).

The problem of Common Reconstruction (CR) [15] is a different type of source coding with side-information, a.k.a. Wyner-Ziv coding [6]. In [15] the distortion was measured with a log-loss merit, and the encoder is required to perfectly reconstruct decoder's sequence. It can be shown that for the CR, the resulting single-letter rate-distortion region is equivalent to IB.

The problem of Information Combining [16] was analyzed in the context of check nodes in LDPC decoding. Two extremes were considered in form of maximization and minimization of mutual information for the binary X setting [9]. It can be shown that the first extreme is equivalent to PF, while the second recovers the IB setting. A recent comprehensive tutorial on the IB method and related problems is given in [5].

Applications of IB methods in Machine Learning are detailed in [17]. Furthermore, the IB methodology connects to many timely aspects, such as Capital Investment [18], Distributed Learning [19], Deep Learning [20], and Convolutional Neural Networks [21].

The PF, which was first introduced in [22], is a dual problem to the IB method. In contrast to IB problem, the goal in PF is to minimize $I(Y;W)$ over all test-channels $P_{W|X}$ subject to $I(X;W) \geq C$. To be more formal, the PF function, $R^{\text{PF}} : [0, H(X)] \rightarrow \mathbb{R}_+$ is defined as

$$R_{P_{xy}}^{\text{PF}}(C) \triangleq \underset{P_{W|X}}{\text{minimize}} \quad I(W;Y) \quad (2)$$

subject to $I(X;W) \geq C$.

Note that since the objective function is a convex function of $P_{W|X}$, taking the constraint here with reverse inequality, i.e. $I(X;W) \leq C$, will induce a trivial solution, i.e. taking X and W independent.

PF is directly connected to Information Combining [9], [16]. For example, if the channel from X to Y is a BMS, then by [9], $P_{W|X}$ is a Binary Erasure Channel (BEC). A rather intriguing example is the setting where the pair (X,Y) are jointly Gaussian, where the result of the minimization is zero, since one can use the channel from X to W to describe the less significant bits of X [23]. Furthermore, the additive noise Helper problem studied in [24], is directly linked to the PF. By reformulating the former as an information combining problem, the solution follows directly as was shown in [23].

In this work we address the input symmetric nonbinary setting for the IB and PF functions. We will find conditions on the bivariate source (X,Y) for which the stochastic encoder from X to W can be completely characterized, thus extending the binary examples from [7], [9] and [5]. Omitted proofs are at the arXiv version of this paper [25].

II. NOTATIONS AND BASIC PROPERTIES

We denote by Δ_n the n dimensional probability simplex, $\mathbf{q} \in \Delta_n$ the marginal probability vector of X, and T the

transition matrix from X to Y, i.e.,

$$T_{ij} \triangleq P(Y = i|X = j), \quad 1 \leq i \leq m, 1 \leq j \leq n. \quad (3)$$

We further rewrite (1) with explicit dependence on \mathbf{q} and T as $R_T(\mathbf{q}, C) = R(C) = R_{P_{xy}}(C)$. The entropy of an n -ary probability vector $\mathbf{p} \in \Delta_n$ is denoted by $h_n(\mathbf{p})$.

The following tight cardinality bound was established in [26]. It was actually already proved for the corresponding dual problem, namely the IB Lagrangian, in [27]. But since $R_T(\mathbf{q}, C)$ is generally not a strictly convex function of C , the result in [27] cannot be directly applied for our problem (1).

Lemma 1 ([26, Th. 9]): The optimization over W in (1) can be restricted to $|W| \leq n$.

As we have already mentioned, the IB function defined in (1) is closely related to the Conditional Entropy Bound (CEB) problem studied in [7], which is given by

$$F_T(\mathbf{q}, x) \triangleq \underset{W \rightarrow X \rightarrow Y}{\text{minimize}} \quad H(Y|W) \quad (4)$$

subject to $H(X|W) \geq x$.

Remark 1: Note that originally in [7] the conditional entropy constraint was given with equality, and equivalence to the inequality setting was established in [7, Theorem 2.5]. It turns out that the aforementioned problem is closely connected to the IB function.

Proposition 2.1: The IB function defined in (1) is equivalent to the CEB function defined in (4).

The latter result implies that we can utilize the properties of $F_T(\mathbf{q}, x)$ developed in [7] for our problem in a straightforward manner, an aspect that we will heavily rely on in Section III.

In a very similar manner to Proposition 2.1, we can redefine the Privacy Funnel problem defined in (2) as follows.

$$F_T^{\text{PF}}(\mathbf{q}, x) \triangleq \underset{P_{W|X}}{\text{maximize}} \quad H(Y|W) \quad (5)$$

subject to $H(X|W) \leq x$.

We have the following characterization of $F_T^{\text{PF}}(\mathbf{q}, x)$.

Theorem 1: The function $F_T^{\text{PF}}(\mathbf{q}, \cdot)$ is concave on the compact convex domain $\{x : 0 \leq x \leq h_n(\mathbf{q})\}$ and for each (\mathbf{q}, x) , the maximum is attained with W taking at most $n + 1$ values. The proof of this theorem is similar to [7, Theorem 2.3] and is omitted here due to space limitations.

III. THE SYMMETRIC INFORMATION BOTTLENECK

In this section we will give a characterization of the achieving conditional distributions and the value of the problem defined in (1) for specific class of input symmetric channels. We begin with the definitions of symmetric group of permutation, symmetry group of stochastic matrix and input symmetric channel [7].

Definition 1: Let \mathcal{S}_n denote the representation of the symmetric group of permutation of n objects by the $n \times n$ permutation matrices. Let $\mathcal{S}_n \times \mathcal{S}_m$ be the representation of the direct product group by the pairs (G, Π) , $G \in \mathcal{S}_n$; $\Pi \in \mathcal{S}_m$ with the composition $(G_1, \Pi_1)(G_2, \Pi_2) = (G_1 G_2, \Pi_1 \Pi_2)$.

For an $m \times n$ stochastic matrix T , (an n input, m output channel), let \mathcal{G} be the set $\{(G, \Pi) \in \mathcal{S}_n \times \mathcal{S}_m | TG = \Pi T\}$,

and let \mathcal{G}_i (\mathcal{G}_o) be the projections of \mathcal{G} on the first (second) factor. If $TG_1 = \Pi_1 T$, $TG_2 = \Pi_2 T$, then $TG_1 G_2 = \Pi_1 \Pi_2 T$ which shows that \mathcal{G} , \mathcal{G}_i , \mathcal{G}_o are subgroups of the finite groups $\mathcal{S}_n \times \mathcal{S}_m$, \mathcal{S}_n , \mathcal{S}_m respectively. \mathcal{G} is the symmetry group of T , \mathcal{G}_i (\mathcal{G}_o) is the input (output) symmetry group.

The channel defined by T will be called input (output) symmetric if \mathcal{G}_i (\mathcal{G}_o) is transitive (a subgroup of \mathcal{S}_n is transitive if each element of $\{1, \dots, n\}$ can be mapped to every other element of $\{1, \dots, n\}$ by some member of the subgroup). T is said to be symmetric if both \mathcal{G}_i and \mathcal{G}_o are transitive.

We also define the set of (\mathbf{q}, C) for which we will have a complete characterization of the achieving distributions.

Definition 2: Assume that $\{G_\alpha\}_{\alpha=1}^n \in \mathcal{S}_n$ is a set of n distinct elements. Let $\phi(\mathbf{p}, \lambda) \triangleq h_m(T\mathbf{p}) - \lambda h_n(\mathbf{p})$ and $\mathbf{p}^* = \arg\min_{\mathbf{p} \in \Delta_n} \phi(\mathbf{p}, \lambda)$. We define the following set for any $\lambda \in [0, 1]$:

$$\mathcal{Q} \triangleq \left\{ (\mathbf{q}, C) : \mathbf{q} = \sum_{\alpha=1}^n w_\alpha G_\alpha \mathbf{p}^*, \mathbf{w} \in \Delta_n, C = 1 - h_n(\mathbf{p}^*) \right\}. \quad (6)$$

Equipped with this definition we are ready to state our main theorem here.

Theorem 2: Assume that T is input symmetric stochastic matrix with input symmetry group \mathcal{G}_i of order n . Then for every $(\mathbf{q}, C) \in \mathcal{Q}$ defined in (6), the optimal test-channel from W to X is a modulo-additive channel.

Note if \mathbf{q} is uniform over n , then it always in \mathcal{Q} , as taking \mathbf{w} to be uniform over n , we obtain

$$\mathbf{q} = \sum_{\alpha=1}^n w_\alpha G_\alpha \mathbf{p}^* = \frac{1}{n} \sum_{\alpha=1}^n G_\alpha \mathbf{p}^* = \mathbf{u}_n, \quad (7)$$

where \mathbf{u}_n is an n -ary uniform probability vector. This fact induces the following corollary.

Corollary 3.1: Assume that T is input symmetric stochastic matrix with input symmetry group \mathcal{G}_i of order n and X is uniformly distributed over n . Then for every $C \in [0, \log n]$, the test-channel from W to X is a modulo-additive noise channel and W is uniform over n .

A particular case for which T is input symmetric, is when the channel from X to Y is a modulo-additive noise channel, i.e., there exist a random variable Z , with probability vector \mathbf{z} such that $Y = X \oplus Z$, where \oplus is modulo n addition. An equivalent representation of the modulo-additive noise channel is using circulant matrix. A circulant matrix $A \in M_n(\mathbb{F})$ [28, p. 33] has the form

$$A = \begin{pmatrix} a_1 & a_2 & \cdots & a_n \\ a_n & a_1 & a_2 & \cdots & a_{n-1} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ a_2 & a_3 & \cdots & a_n & a_1 \end{pmatrix}, \quad (8)$$

i.e., the entries in each row are cyclic permutations of those in the first. In this case we have the following corollary.

Corollary 3.2: If $T = A$ as defined in (8), then some modulo additive test channel from W to X achieves $R_A(\mathbf{q}, C)$.

In particular, there exists an n -ary random variable V , with $H(V) = \log n - C$, such that $X = W \oplus V$ achieves $R_A(\mathbf{q}, C)$.

Although this result greatly simplifies the optimization space, it does not give a precise analytical solution to the problem. In the following subsection, we provide an example, for which the achieving distribution and the objective function value can be fully characterized.

A. Hamming Channels

Let $T = T_\alpha = \alpha I_n + (1 - \alpha)n^{-1}E_n$, where I_n is the $n \times n$ identity matrix, E_n the all ones matrix, and $0 \leq \alpha \leq 1$. The channel with transition matrix T_α is called a Hamming channel with parameter α . Note that T_α is in particular a circulant matrix, therefore by Corollary 3.2 the optimal channel from W to X is a modulo-additive channel. Thus, (4) can be reformulated as follows.

$$F_T(\mathbf{q}, x) \triangleq \begin{aligned} & \underset{\mathbf{v} \in \Delta_n}{\text{minimize}} && h_n(T_\alpha \mathbf{v}) \\ & \text{subject to} && h_n(\mathbf{v}) \geq x. \end{aligned} \quad (9)$$

The optimization problem defined in (9) is identical to the problem considered in [29]. Furthermore, it was solved for the Hamming channel and the achieving distribution was found.

Lemma 2 ([29, Lemma 7]): For $n \times n$ Hamming channel T_α the solution to (9) is attained for

$$\mathbf{v} = \beta \mathbf{e} + (1 - \beta)\mathbf{u}_n. \quad (10)$$

where \mathbf{e} is any standard basis vector of Δ_n .

Since \mathbf{v} is determined by a single parameter β and satisfies $h_n(\mathbf{v}) = \log n - C$, we can find β explicitly as follows:

$$\begin{aligned} C &= \log n - h_n(\mathbf{v}) \\ &= \frac{n-1}{n}(1-\beta)\log(1-\beta) + \frac{\beta n + 1 - \beta}{n}\log(\beta n + 1 - \beta) \\ &\triangleq g_n(\beta). \end{aligned}$$

Thus, β can be recovered from C as $\beta = g_n^{-1}(C)$. In summary, we have the following theorem.

Theorem 3: Assume that T is a Hamming channel with parameter α , then $R_T(\mathbf{u}_n, C)$ is attained with a Hamming channel with parameter $\beta = g_n^{-1}(C)$ and is given by

$$R_T(\mathbf{u}_n, C) = \frac{1+(n-1)\alpha\beta}{n}\log(1+(n-1)\alpha\beta) + \frac{1-\alpha\beta}{n}\log(1-\alpha\beta). \quad (11)$$

B. Examples

Now let us consider two special cases.

1) *BMS:* Assume that the channel from X to Y is a BMS channel. Let \mathbf{z} be an m -ary probability vector and G_m be the $m \times m$ anti-diagonal matrix with unit entries. The respective transition matrix in this case is $T = [\mathbf{z}, G_m \mathbf{z}]$. Note that

$$G_m T = [G_m \mathbf{z}, G_m G_m \mathbf{z}] = [\mathbf{z} G_2 \mathbf{z}] = T G_2. \quad (12)$$

Therefore, T is input symmetric stochastic matrix with input symmetry group \mathcal{G}_i of order 2. Thus, since the only binary-input binary-output symmetric channel is a BSC, combining with Theorem 2, we recover the following result from [9].

Corollary 3.3 ([9, Theorem 2]): Given that the channel from X to Y is a BMS, then BSC channel from X to W maximizes $I(W; Y)$.

The latter result can also be deduced from [30].

2) *Ternary-Input Ternary-Output (TITO) Circulant Matrix*: The general TITO Circulant Matrix is defined as follows:

$$T = \begin{pmatrix} 1 - \alpha - \beta & \alpha & \beta \\ \beta & 1 - \alpha - \beta & \alpha \\ \alpha & \beta & 1 - \alpha - \beta \end{pmatrix}. \quad (13)$$

We can further ask if there are values of C such that $R(C)$ can be achieved with W taking at most two points. The following corollary states the opposite.

Corollary 3.4: The minimum cardinality of W that achieves $R(C)$ is exactly 3 for $C \neq 0$.

C. Numerical Simulation

We proceed to verify Theorem 3 via numerical optimization for $n = 3$. Since V is independent of the choice of α , we fix





Fig. 2: (a) Optimal \mathbf{v} for $\alpha = 0.5$ vs C . (b) $R_{T_\alpha}(C)$ vs C for $\alpha = 0.5$.

$\alpha = 0.5$ and compare it with respect to the value of C . Figure 2 shows the probability vector V and $R_{T_\alpha}(C)$ for various values of α . We observe that the numerical optimization agrees with theoretical arguments of Theorem 3.

IV. THE SYMMETRIC PRIVACY FUNNEL

In this section we consider a special symmetric setting for the PF problem (5) for which the transition matrix from X to Y is an input symmetric stochastic matrix as defined in Definition 1.

Theorem 4: Let T be an input symmetric stochastic matrix with input symmetry group \mathcal{G}_i of order n , and X be a uniformly distributed random variable. Let $(G_1 = I, G_2, \dots, G_n) \in \mathcal{G}_i$. Furthermore, denote by $(\mathbf{p}^*, \lambda^*)$ a pair for which

$$\phi(\mathbf{u}, \lambda^*) = \phi(\mathbf{p}^*, \lambda^*) \geq \phi(\mathbf{p}, \lambda^*) \quad \forall \mathbf{p} \in \Delta_n. \quad (14)$$

Then, for every $C \leq C^* \triangleq \log n - h_n(\mathbf{p}^*)$, the transition matrix from W to X , given by

$$B = (\mathbf{p}^* \quad G_2 \mathbf{p}^* \quad \dots \quad G_n \mathbf{p}^* \quad \mathbf{u}), \quad (15)$$

achieves (2). Moreover,

$$R_{\text{px}}^{\text{PF}}(C) = C \cdot \frac{\log n - h_n(T\mathbf{p}^*)}{\log n - h_n(\mathbf{p}^*)}. \quad (16)$$

Also, (15) implies that the transition matrix from X to W is a class of noisy n -ary symmetric erasure channel.

Note that the optimization procedure in (14) is performed once for every $C \in [0, \log n - h_n(\mathbf{p}^*)]$. Moreover, for $C \in [0, \log n - h_n(\mathbf{p}^*)]$, the optimal test-channel from X to W is no longer symmetric as we show using a numerical example. We now provide some examples that illustrate Theorem 4.

A. Examples

We begin with the simplest scenario where X is a binary random variable. Plugging this choice in Theorem 4 and noting that $\mathbf{p}^* = \mathbf{e}$ in this case, results in the following corollary.

Corollary 4.1: Assume that the channel from X to Y is a BMS, then, BEC test-channel $P_{W|X}$ with parameter $\epsilon = 1 - C$ minimizes $I(Y; W)$ subject to $I(X; W) = C$.

Note that this result recovers [9, Theorem 1], but here with only one-sided symmetry restriction.

We further illustrate Theorem 4 using numerical optimization for a particular choice of the channel from X to Y being a symmetric TITO with parameters $(\alpha, \beta) = (0.1, 0.05)$, as defined in (13). For this choice of channel parameters, $C^* = 0.59$. In Figure 3 we compare the results of global optimization solution of (2) versus choosing $P_{W|X}$ be the respective optimal input-symmetric channel as described in Theorem 4 for various values of C . We observe that our results from Theorem 4 agree with the brute-force numerical optimization for all values of $C \in [0, C^*]$. For values greater than C^* , we observe that the curve which is restricted to input symmetric transition matrices from X to W , is sub-optimal. In this region of link capacity, the numerical optimization achieves lower rates. By carefully observing the numerical solution, one can notice that the optimal test-channel in this region is no longer input symmetric.




Fig. 3: Optimal R for $\alpha = 0.1$, $\beta = 0.05$ vs C

V. OUTLOOK

As said, the Information Bottleneck and Privacy Funnel are two dual optimization problems which have been applied in a variety of emerging applications such as Deep Neural Networks, Privacy Algorithms, and design of Polar Codes [17]. It also interesting to consider rather more classical use-cases, i.e, multi-user channel capacity and Noisy Source Coding problems. A comprehensive summary of the different relations between the IB and Privacy Funnel problems has been presented in [26].

ACKNOWLEDGMENT

This work has been supported by the European Union's Horizon 2020 Research And Innovation Programme, grant agreement no. 694630 and by the WIN consortium via the Israel minister of economy and science.

REFERENCES

- [1] N. Tishby, F. C. N. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. 37th Annu. Allerton Conf. Commun. Control Comput.*, Sep. 1999, p. 368–377.
- [2] J. Wolf and J. Ziv, "Transmission of noisy information to a noisy receiver with minimum distortion," *IEEE Trans. Inf. Theory*, vol. 16, pp. 406–411, Jul. 1970.
- [3] S. Arimoto, "An algorithm for computing the capacity of arbitrary discrete memoryless channels," *IEEE Trans. Inf. Theory*, vol. 18, pp. 14–20, Jan. 1972.
- [4] N. Slonim, "The information bottleneck: Theory and applications," Ph.D. dissertation, Hebrew University of Jerusalem, Jerusalem, Israel, 2002.
- [5] A. Zaidi, I. E. Aguerri, and S. S. (Shitz), "On the Information Bottleneck Problems: Models, Connections, Applications and Information Theoretic Views," *Entropy*, vol. 22, no. 2, 2020.
- [6] A. Wyner and J. Ziv, "A theorem on the entropy of certain binary sequences and applications I," *IEEE Trans. Inf. Theory*, vol. 19, pp. 769–772, 1973.
- [7] H. S. Witsenhausen and A. D. Wyner, "A Conditional Entropy Bound for a Pair of Discrete Random Variables," *IEEE Trans. Inf. Theory*, vol. 21, no. 5, pp. 493–501, Sep. 1975.
- [8] T. Richardson and R. Urbanke, *Modern Coding Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2008.
- [9] I. Sutskever, S. Shamaï, and J. Ziv, "Extremes of information combining," *IEEE Trans. Inf. Theory*, vol. 51, no. 4, pp. 1313–1325, Apr. 2005.

- [10] G. Chechik, A. Globerson, N. Tishby, and Y. Weiss, "Information Bottleneck for Gaussian Variables," *J. Mach. Learn. Res.*, vol. 6, pp. 165–188, Dec. 2005.
- [11] A. E. Gamal and Y. Kim, *Network Information Theory*. Cambridge University Press, 2011.
- [12] D. Guo, S. Shamaï (Shitz), and S. Verdú, "The Interplay Between Information and Estimation Measures," *Found. Trends Signal Process.*, vol. 6, no. 4, pp. 243–429, 2012.
- [13] R. Bustin, M. Payaro, D. P. Palomar, and S. Shamaï (Shitz), "On MMSE crossing properties and implications in parallel vector Gaussian channels," *IEEE Trans. Inf. Theory*, vol. 59, no. 2, pp. 818–844, Feb. 2013.
- [14] P. Gács and J. Körner, "Common information is far less than mutual information," *Probl. Contr. Inform. Theory*, vol. 2, no. 2, pp. 149–162, 1973.
- [15] Y. Steinberg, "Coding and common reconstruction," *IEEE Trans. Inf. Theory*, vol. 55, no. 11, pp. 4995–5010, Nov. 2009.
- [16] I. Land and J. Huber, "Information Combining," *Found. Trends Commun. Inf. Theory*, vol. 3, no. 3, pp. 227–330, Nov. 2006.
- [17] Z. Goldfeld and Y. Polyanskiy, "The Information Bottleneck Problem and its Applications in Machine Learning," *IEEE J. Sel. Areas Inf. Theory*, vol. 1, no. 1, pp. 19–38, May 2020.
- [18] E. Erkip and T. M. Cover, "The Efficiency of Investment Information," *IEEE Trans. Inf. Theory*, vol. 44, no. 3, pp. 1026–1040, May 1998.
- [19] P. Farajiparvar, A. Beirami, and M. Nokleby, "Information Bottleneck Methods for Distributed Learning," in *Proc. 56th Annu. Allerton Conf. Commun., Control Comput.*, 2018, pp. 24–31.
- [20] R. A. Amjad and B. C. Geiger, "Learning Representations for Neural Network-Based Classification Using the Information Bottleneck Principle," *IEEE Trans. Pattern Anal.*, vol. 42, no. 9, pp. 2225–2239, Sep. 2020.
- [21] S. Yu, K. Wickstrøm, R. Jenssen, and J. C. Príncipe, "Understanding convolutional neural networks with information theory: An initial exploration," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 1, pp. 435–442, Jan. 2021.
- [22] F. du Pin Calmon, A. Makhdoumi, M. Médard, M. Varia, M. Christiansen, and K. R. Duffy, "Principal inertia components and applications," *IEEE Trans. Inf. Theory*, vol. 63, no. 8, pp. 5011–5038, Aug. 2017.
- [23] S. Shamaï, "The information bottleneck: A unified information theoretic view," National Conference on Communications (NCC2021), Jul. 2021, plenary Address.
- [24] S. I. Bross and A. Lapidoth, "The additive noise channel with a helper," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Visby, Sweden, Aug. 2019, pp. 1–5.
- [25] M. Dikshtein and S. Shamaï, "A class of nonbinary symmetric information bottleneck problems," *CoRR*, vol. abs/2110.00985, 2021. [Online]. Available: <https://arxiv.org/abs/2110.00985>
- [26] S. Asoodeh and F. P. Calmon, "Bottleneck Problems: An Information and Estimation-Theoretic View," *Entropy*, vol. 22, no. 11, p. 1325, 2020.
- [27] P. Harremoës and N. Tishby, "The Information Bottleneck Revisited or How to Choose a Good Distortion Measure," in *Proc. 2007 IEEE Int. Symp. Inf. Theory*, Jun. 2007, pp. 566–570.
- [28] R. A. Horn, *Matrix Analysis*, 2nd ed. Cambridge: Cambridge University Press, 2012.
- [29] H. S. Witsenhausen, "Entropy inequalities for discrete channels," *IEEE Trans. Inf. Theory*, vol. 20, no. 5, pp. 610–616, Sep. 1974.
- [30] N. Chayat and S. Shamaï, "Extension of an entropy property for binary input memoryless symmetric channels," *IEEE Trans. Inf. Theory*, vol. 35, no. 5, pp. 1077–1079, 1989.