THE DUAL INFORMATION BOTTLENECK

Anonymous authors
Paper under double-blind review

ABSTRACT

The Information-Bottleneck (IB) framework suggests a general characterization of optimal representations in learning, and deep learning in particular. It is based on the optimal trade off between the representation complexity and accuracy, both of which are quantified by mutual information. The problem is solved by alternating projections between the *encoder* and *decoder* of the representation, which can be performed locally at each representation level. The framework, however, has practical drawbacks, in that mutual information is notoriously difficult to handle at high dimension, and only has closed form solutions in special cases. Further, because it aims to extract representations which are minimal sufficient statistics of the data with respect to the desired label, it does not necessarily optimize the actual prediction of unseen labels. Here we present a formal dual problem to the IB which has several interesting properties. By switching the order in the KL-divergence between the representation decoder and data, the optimal decoder becomes the geometric rather than the arithmetic mean of the input points. While providing a good approximation to the original IB, it also preserves the form of exponential families, and optimizes the mutual information on the predicted label rather than the desired one. We also analyze the critical points of the dual IB and discuss their importance for the quality of this approach.

1 INTRODUCTION

1.1 THE INFORMATION BOTTLENECK METHOD

The Information Bottleneck (IB) method (Tishby et al., 1999), is an information-theoretic framework for describing efficient representations of an "input" random variable X (input patterns), for predicting an "output" variable Y (desired label). In this setting the joint distribution of X and Y, $p(\mathrm{x}, \mathrm{y})$ defines the problem, or rule, and the training data are a finite sample from this distribution. In general, we assume that $p(\mathrm{y} \mid \mathrm{x})$ is strictly stochastic, and hence bounded away from $\{0, 1\}$[1]. The representation variable $\hat{\mathrm{X}}$ is in general a stochastic function of X which forms a Markov chain $\mathrm{Y} \to \mathrm{X} \to \hat{\mathrm{X}}$, and only depends on Y through the input X. We call the map $p(\hat{\mathrm{x}} \mid \mathrm{x})$ the *encoder* of the representation and denote by $p(\mathrm{y} \mid \hat{\mathrm{x}})$ the *Bayes optimal decoder* for this representation; i.e., the best possible prediction of the *desired label* Y from the representation $\hat{\mathrm{X}}$.

The IB trade off between the encoder and decoder mutual information values is defined by the minimization of the Lagrangian:

$$\mathcal{F}[p_\beta(\hat{\mathrm{x}} \mid \mathrm{x}); p_\beta(\hat{\mathrm{x}}); p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})] = I(\mathrm{X}; \hat{\mathrm{X}}) - \beta I(\mathrm{Y}; \hat{\mathrm{X}}) \,, \tag{1}$$

independently over the convex sets of the normalized distributions, $\{p_\beta(\hat{\mathrm{x}} \mid \mathrm{x})\}$, $\{p_\beta(\hat{\mathrm{x}})\}$ and $\{p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})\}$, given a positive Lagrange multiplier β. As shown in (Tishby et al., 1999; Shamir et al., 2010), this is a natural generalization of the classical concept of *Minimal Sufficient Statistics* (Cover & Thomas, 2006), where the estimated parameter is replaced by the output variable Y and *exact* statistical sufficiency is characterized by the mutual information equality: $I(\hat{\mathrm{X}}; \mathrm{Y}) = I(\mathrm{X}; \mathrm{Y})$. The minimality of the statistics is captured by the minimization of $I(\mathrm{X}; \hat{\mathrm{X}})$, due to the Data Processing Inequality (DPI). However, non-trivial minimal sufficient statistics only exist for very special parametric distributions known as exponential families (Brown, 1986). Thus in general, the IB relaxes

[1]This may seem like a limitation of the method, but it can be shown that in the deterministic limit the IB is well defined. This stochastic assumption simplifies the analysis but does not pose any real limitation.

the minimal sufficiency problem to a continuous family of representations $\hat{\mathrm{X}}$ which are characterized by the trade off between compression, $I(\mathrm{X};\hat{\mathrm{X}}) \equiv I_{\mathrm{X}}$, and accuracy, $I(\mathrm{Y};\hat{\mathrm{X}}) \equiv I_{\mathrm{Y}}$, along a convex line in the *Information-Plane* (I_{Y} vs. I_{X}). When the rule $p(\mathrm{x},\mathrm{y})$ is strictly stochastic, the convex optimal line is smooth and each point along the line is uniquely characterized by the value of β. We can then consider the optimal representations $\hat{\mathrm{x}} = \hat{\mathrm{x}}(\beta)$ as encoder-decoder pairs: $(p_\beta(\mathrm{x}|\hat{\mathrm{x}}), p_\beta(\mathrm{y}|\hat{\mathrm{x}}))^2$ - a point in the continuous manifold defined by the Cartesian product of these distribution simplexes. We also consider a small variation of these representations, $\delta\hat{\mathrm{x}}$, as an infinitesimal change in this (encoder-decoder) continuous manifold (not necessarily on the optimal line(s)).

1.2 IB AND RATE-DISTORTION THEORY

The IB optimization trade off can be considered as a generalized rate-distortion problem (Cover & Thomas, 2006) with the distortion function between a data point, x and a representation point $\hat{\mathrm{x}}$ taken as the KL-divergence between their predictions of the desired label y:

$$d_{\mathrm{IB}}(\mathrm{x},\hat{\mathrm{x}}) = D[p(\mathrm{y}\mid\mathrm{x})||p_\beta(\mathrm{y}|\hat{\mathrm{x}})] = \sum_{\mathrm{y}} p(\mathrm{y}\mid\mathrm{x})\log\frac{p(\mathrm{y}\mid\mathrm{x})}{p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})}\ . \tag{2}$$

The expected distortion $\langle d_{\mathrm{IB}}(\mathrm{x},\hat{\mathrm{x}})\rangle_{p_\beta(\mathrm{x},\hat{\mathrm{x}})}$ for the optimal decoder is simply the label-information loss: $I(\mathrm{X};\mathrm{Y}) - I(\hat{\mathrm{X}};\mathrm{Y})$, using the Markov chain condition. Thus minimizing the expected IB distortion is equivalent to maximizing $I(\hat{\mathrm{X}};\mathrm{Y})$, or minimizing equation 1. Minimizing this distortion is equivalent to minimizing the $\log$-loss or the cross-entropy loss, as done in most deep learning applications, and it upper-bounds other loss functions such as the $\mathcal{L}_1$-loss (due to the Pinsker inequality, or (Painsky & Wornell, 2018)). The Pinsker inequality shows that the relative entropy gives a symmetric upper bound to the $\mathcal{L}_1$-loss, or variation distance, $D[p||q] \geq \frac{1}{2\log 2}\|p-q\|_1^2$.

1.3 THE IB EQUATIONS

For discrete X and Y, a necessary condition for the IB (local) minimization is given by the three self-consistent equations for the optimal encoder-decoder pairs, known as *the* IB *equations*:

$$\begin{cases} (i) & p_\beta(\hat{\mathrm{x}}\mid\mathrm{x}) = \frac{p(\hat{\mathrm{x}})}{Z(\mathrm{x};\beta)} e^{-\beta D[p(\mathrm{y}|\mathrm{x})\|p_\beta(\mathrm{y}|\hat{\mathrm{x}})]} \\ (ii) & p_\beta(\hat{\mathrm{x}}) = \sum_{\mathrm{x}} p_\beta(\hat{\mathrm{x}}\mid\mathrm{x})p(\mathrm{x}) \\ (iii) & p_\beta(\mathrm{y}\mid\hat{\mathrm{x}}) = \sum_{\mathrm{x}} p(\mathrm{y}\mid\mathrm{x})p_\beta(\mathrm{x}\mid\hat{\mathrm{x}}) \end{cases} . \tag{3}$$

Iterating these equations is a generalized, Blahut-Arimoto, alternating projection algorithm (Tusnady & Csiszar, 1984; Cover & Thomas, 2006) and it converges to a stationary point of the Lagrangian, equation 1 (Tishby et al., 1999). Notice that the minimizing decoder, (equation 3-(iii)), is precisely the *Bayes optimal decoder* for the representation $\hat{\mathrm{x}}(\beta)$, given the Markov chain conditions.

1.4 CRITICAL POINTS AND CRITICAL SLOWING-DOWN

One of the most interesting aspects of the IB equations is the existence of critical points along the optimal line of solutions in the *Information-Plane* (i.e. the information curve). At these points the representations change topology and cardinality (number of clusters) (Zaslavsky & Tishby, 2019; Parker et al., 2003) and they form the skeleton of the information curve and representation space. Under the strict stochastic assumption, the information-curve is a smooth function of the Lagrange multiplier β, but its derivative may not be smooth. Critical points are bifurcations of the solutions, which are values of β for which two different solutions (representations) co-exist. To identify such points we perform a perturbation analysis of the IB equations, as in (Zaslavsky & Tishby, 2019). Taking a small perturbation of the representation, denoted for brevity by $\delta\hat{\mathrm{x}}$, the changes in the $\log$ encoder and $\log$ decoder that satisfy equation 3 for a given β can be determined through the nonlinear eigenvalues problems:

$$\left[I - \beta C^{\mathrm{IB}}_{\mathrm{xx}'}(\hat{\mathrm{x}},\beta)\right]\frac{\partial\log p_\beta(\mathrm{x}'\mid\hat{\mathrm{x}})}{\partial\hat{\mathrm{x}}} = 0\ , \quad \left[I - \beta C^{\mathrm{IB}}_{\mathrm{yy}'}(\hat{\mathrm{x}},\beta)\right]\frac{\partial\log p_\beta(\mathrm{y}'\mid\hat{\mathrm{x}})}{\partial\hat{\mathrm{x}}} = 0, \tag{4}$$

[2]Here we use the *inverse encoder*, which is in the fixed dimension simplex of distributions over X.

with the two square matrices C defined by:

$$C^{\mathrm{IB}}_{\mathrm{xx}'}(\hat{\mathrm{x}},\beta)=\sum_{\mathrm{y}}p(\mathrm{y}\mid\mathrm{x})\frac{p_\beta(\mathrm{x}'\mid\hat{\mathrm{x}})}{p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})}p(\mathrm{y}\mid\mathrm{x}')\ ,\quad C^{\mathrm{IB}}_{\mathrm{yy}'}(\hat{\mathrm{x}},\beta)=\sum_{\mathrm{x}}p(\mathrm{y}\mid\mathrm{x})\frac{p_\beta(\mathrm{x}\mid\hat{\mathrm{x}})}{p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})}p(\mathrm{y}'\mid\mathrm{x}). \quad (5)$$

As shown in (Zaslavsky & Tishby, 2019), these two matrices have the same eigenvalues and have non-trivial eigenvectors (i.e., different co-existing optimal representations) precisely at the critical values of β, the bifurcation points of the IB solution. At these points the cardinality of the representation $\hat{\mathrm{X}}$ (the number of "IB-clusters") changes due to splits of clusters, resulting in topological phase transitions in the encoder. These critical points form the "skeleton" of the topology of the optimal representations. Between critical points the optimal representations change continuously (with β).

The important computational consequence of critical points is known as *critical slowing down* (Tredicce et al., 2004). For binary Y, near a critical point the convergence time, τ_β, of the iterations of equation 3 scales like: $\tau_\beta \sim 1/(1-\beta\lambda_2)$, where λ_2 is the second eigenvalue of either $C^{\mathrm{IB}}_{\mathrm{yy}'}$ or $C^{\mathrm{IB}}_{\mathrm{xx}'}$. At criticality, $\lambda_2(\hat{\mathrm{x}})=\beta^{-1}$ and the number of iterations diverges. This phenomenon dominates any local minimization of equation 3 which is based on alternate encoder-decoder optimization.

The discussed phenomenons of the IB are graphically demonstrated in Figure 1. The figure shows the bifurcation diagram, the non-trivial eigenvalues of $C^{\mathrm{IB}}_{\mathrm{yy}'}$ and convergence time of the IB iterations, for a simple problem with 4 critical points.

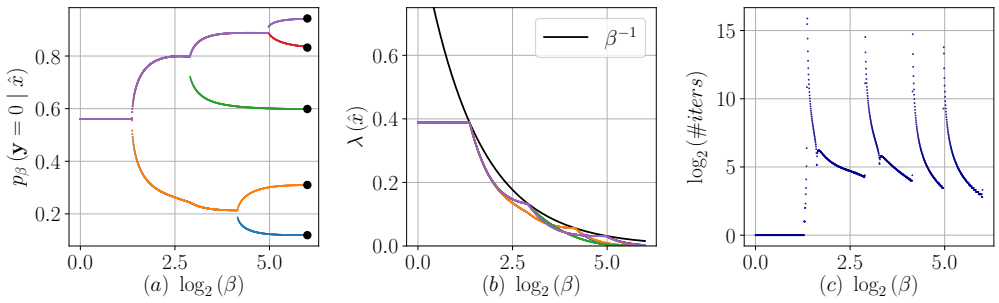

Figure 1: An example of IB solutions presenting the critical points and the *critical slowing down* of the alternating projection algorithm. (a) The algorithm's solution to $p_\beta(\mathbf{y}=0\mid\hat{\mathrm{x}})$ as a function of β. The black dots depict the input distribution $p(\mathbf{y}=0\mid\mathrm{x})$. (b) The second eigenvalue of $C_{\mathrm{yy}'}(\hat{\mathrm{x}};\beta)$, $\lambda_2(\hat{\mathrm{x}})$, along with β^{-1} as a function of β. (c) Convergence time of the algorithm as a function of β.

1.5 Contribution of this work

Supervised learning is generally separated into two phases: the training phase, where the internal representations are formed from the training data, and the prediction phase, where these representations are used to predict labels of new input patterns. In our Markov chain description we need to add another variable, $\hat{\mathrm{Y}}$, the *predicted label* which obtains the same values as Y, but distributed differently:

$$\overbrace{\mathrm{Y}\to\underbrace{\mathrm{X}\to\hat{\mathrm{X}}_\beta}_{\text{prediction}}}^{\text{training}}\to\hat{\mathrm{Y}}\,. \quad (6)$$

The left-hand part of this chain describes the representation training, and the right-hand part is the Maximum Likelihood (ML) prediction using these representations (Slonim et al., 2006). So far the prediction variable $\hat{\mathrm{Y}}$ has not been part of the IB optimization problem. It has been implicitly assumed that the *Bayes optimal decoder*, $p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})$, which minimizes the IB distortion at a given β, is also the best choice for making predictions of $\hat{\mathrm{Y}}$ from the representation $\hat{\mathrm{X}}_\beta$ through the right-hand Markov chain. That is, denoting $p_\beta(\hat{\mathrm{y}}\mid\hat{\mathrm{x}})\equiv p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})$ the ML decoder is the mixture over the

internal representations:

$$p_\beta(\hat{\mathrm{y}} \mid \mathrm{x}) \equiv \sum_{\hat{\mathrm{x}}} p_\beta(\hat{\mathrm{y}} \mid \hat{\mathrm{x}}) p_\beta(\hat{\mathrm{x}} \mid \mathrm{x}). \tag{7}$$

This, however, is not necessarily optimal. For example due to finite sample, in which it can be very different from the one obtained from the full distribution (Shamir et al., 2010).

Here we show that by merely switching the order of the arguments in the KL-divergence of the IB distortion, namely:

$$d_{\mathrm{dualIB}}(\mathrm{x}, \hat{\mathrm{x}}) = D[p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) || p(\mathrm{y} \mid \mathrm{x})] = \sum_{\mathrm{y}} p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) \log \frac{p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})}{p(\mathrm{y} \mid \mathrm{x})} \,, \tag{8}$$

which in geometric terms is known as the *dual* distortion problem (Felice & Ay, 2019), we obtain several interesting results: (i) The decoder changes from the arithmetic to the geometric mean of the cluster data distributions; (ii) it preserves exponential form of the original data distribution, if one exists, for all values of β; (iii) it preserves the low dimensional sufficient statistics of the data, making the scaling to large problems much easier; (iv) it (variationally) optimizes the predicted label information $I(\mathrm{X}; \hat{\mathrm{Y}})$, and with it the desired and predicted label information $I(\mathrm{Y}; \hat{\mathrm{Y}})$.

The remainder of the paper is organized as follows. We solve the dualIB problem in §2. We discuss the dualIB's critical points §2.1 and provide a comparison to the original IB in §2.2. Next, in §3 we focus on the special case of data from exponential families. We conclude in §4 with further extensions and possible applications to deep learning.

2 THE DUALIB: MAXIMIZING THE PREDICTION INFORMATION

In the dualIB framework we consider the "full" learning Markov chain, equation 6. That is, given the "input" random variable X (input patterns), the "output" random variable Y (desired label) and the "representation" $\hat{\mathrm{X}}$, the ML optimal "predicted label" $\hat{\mathrm{Y}}$ is given by the mixture distribution, equation 7.[3]

The dualIB optimization can be written as the following rate-distortion problem:

$$\mathcal{F}^*[p_\beta(\hat{\mathrm{x}} \mid \mathrm{x}); p_\beta(\hat{\mathrm{x}}); p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})] = I(\mathrm{X}; \hat{\mathrm{X}}) + \beta \langle d_{\mathrm{dualIB}}(\mathrm{x}, \hat{\mathrm{x}}) \rangle_{p(\mathrm{x}, \hat{\mathrm{x}})} \,, \tag{9}$$

with the average distortion given in terms of mutual information on $\hat{\mathrm{Y}}$, $I(\hat{\mathrm{X}}; \hat{\mathrm{Y}})$ and $I(\mathrm{X}; \hat{\mathrm{Y}})$:

$$\begin{aligned} \langle d_{\mathrm{dualIB}}(\mathrm{x}, \hat{\mathrm{x}}) \rangle_{p(\mathrm{x}, \hat{\mathrm{x}})} &= \underbrace{I(\hat{\mathrm{X}}; \hat{\mathrm{Y}}) - I(\mathrm{X}; \hat{\mathrm{Y}})}_{(a)} + \underbrace{\langle D[p_\beta(\hat{\mathrm{y}} \mid \mathrm{x}) \| p(\hat{\mathrm{y}} \mid \mathrm{x})] \rangle_{p(\mathrm{x})}}_{(b)} \\ &\geq I(\hat{\mathrm{X}}; \hat{\mathrm{Y}}) - I(\mathrm{X}; \hat{\mathrm{Y}}) \,. \end{aligned} \tag{10}$$

This is similar to the known IB relation: $\langle d_{\mathrm{IB}}(\mathrm{x}, \hat{\mathrm{x}}) \rangle_{p(\mathrm{x}, \hat{\mathrm{x}})} = I(\mathrm{Y}; \mathrm{X}) - I(\mathrm{Y}; \hat{\mathrm{X}})$ with an extra positive term (b).

Both terms, (a) and (b), vanish precisely when $\hat{\mathrm{X}}$ is a sufficient statistic for X with respect to $\hat{\mathrm{Y}}$, since we can then reverse the order of X and $\hat{\mathrm{X}}$ in the Markov chain (equation 6). This replaces the roles of Y and $\hat{\mathrm{Y}}$ as the variable for which $\hat{\mathrm{X}}_\beta$ are approximately minimally sufficient statistics. In that sense the dualIB shifts the emphasis from the training phase to the prediction phase. This implies that minimizing the dualIB functional maximizes the mutual information between Y and $\hat{\mathrm{Y}}$, $I(\mathrm{Y}; \hat{\mathrm{Y}})$, as well as the mutual information $I(\mathrm{X}; \hat{\mathrm{Y}})$. This is illustrated in figure $4(a), (b)$.

The next theorem states the form of the solutions of *the Dual Information Bottleneck*:

Theorem 1. *The minima of equation 9, can be obtained by generalized Blahut-Arimoto iterations between the encoder and the decoder as in the original* IB*, with the following modifications: (i) Replace the distortion by its dual in the encoder update; (ii) Update the decoder by the encoder's "geometric" mean of the data distributions* $p(\mathrm{y} \mid \mathrm{x})$.

[3] We abuse the notation $p(\hat{\mathrm{y}} \mid \mathrm{x}) \equiv p(\mathrm{y} \mid \mathrm{x})$, when there is no β subscript.

The proof is given in §A.2.

The alternating projections between the encoder and decoder, which converge to a solution of the dualIB at a given value of the Lagrange multiplier β, are implemented by the following iterative algorithm[4]:

Algorithm 1 dualIB iterative algorithm

1: **while** $\left|p_\beta^{t+1}(\mathrm{y}\mid\hat{\mathrm{x}}) - p_\beta^{t}(\mathrm{y}\mid\hat{\mathrm{x}})\right| > \epsilon$ **do**
2: $\quad Z^{t+1}_{\hat{\mathbf{x}}|\mathbf{x}}(\mathrm{x};\beta) = \sum_{\hat{\mathrm{x}}} p_\beta^t(\hat{\mathrm{x}}) e^{-\beta D\left[p_\beta^t(\mathrm{y}|\hat{\mathrm{x}})\|p(\mathrm{y}|\mathrm{x})\right]}$
3: $\quad p_\beta^{t+1}(\hat{\mathrm{x}}\mid\mathrm{x}) = \frac{p_\beta^t(\hat{\mathrm{x}})}{Z^{t+1}_{\hat{\mathbf{x}}|\mathbf{x}}(\mathrm{x};\beta)} e^{-\beta D\left[p_\beta^t(\mathrm{y}|\hat{\mathrm{x}})\|p(\mathrm{y}|\mathrm{x})\right]}$
4: $\quad p_\beta^{t+1}(\hat{\mathrm{x}}) = \sum_{\mathrm{x}} p_\beta^{t+1}(\hat{\mathrm{x}}\mid\mathrm{x}) p(\mathrm{x})$
5: $\quad Z^{t+1}_{\mathrm{y}|\hat{\mathrm{x}}}(\hat{\mathrm{x}};\beta) = \sum_{\mathrm{y}} e^{\sum_{\mathrm{x}} p_\beta^{t+1}(\mathrm{x}|\hat{\mathrm{x}}) \log p(\mathrm{y}|\mathrm{x})}$
6: $\quad p_\beta^{t+1}(\mathrm{y}\mid\hat{\mathrm{x}}) = \frac{1}{Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}};\beta)} e^{\sum_{\mathrm{x}} p_\beta^{t+1}(\mathrm{x}|\hat{\mathrm{x}}) \log p(\mathrm{y}|\mathrm{x})}$
7: **return** $p_\beta(\hat{\mathrm{x}}), p_\beta(\hat{\mathrm{x}}\mid\mathrm{x}), p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})$

Here the Z's are standard normalization factors (partition functions).
As in the IB, the encoder update (row 3) and decoder update (row 6) are the core of the algorithm.
Figure 2 illustrates the properties of the dualIB solutions on the same small problem as in Figure 1.

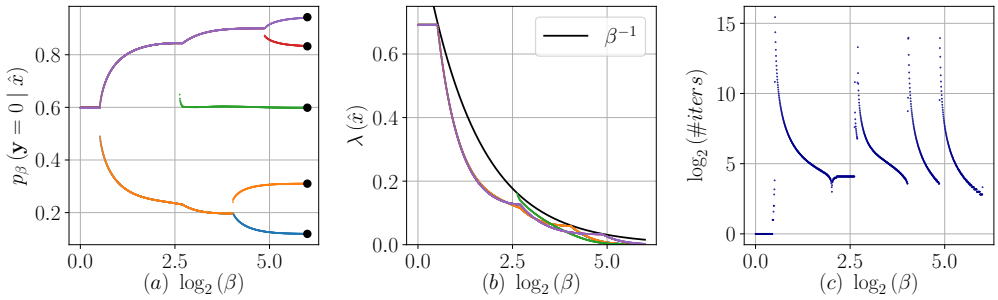

Figure 2: Same as Figure 1 for the dualIB solutions with the $C^{\mathrm{dualIB}}_{\mathrm{yy}'}$ matrix. (a) The representations bifurcate at the critical values of β. (b) Critical points appear when an eigenvalue (color) crosses the β^{-1} (black) line. (c) Near these points there is *critical slowing down* and the numbers of iterations of Algorithm 1 diverge.

2.1 The Critical Points of the dualIB

As discussed in §1.4 the skeleton of the IB optimal bound (the information curve) is constituted by the critical points in which the topology (cardinality) of the representation changes. A similar stability analysis of the dualIB equations reveals similar conditions for the critical points.

Theorem 2. *The* dualIB *critical points are detected by non-trivial solutions of the nonlinear eigenvalue problem:*

$$\left[I - \beta C^{\mathrm{dualIB}}_{\mathrm{xx}'}(\hat{\mathrm{x}},\beta)\right]\frac{\partial \log p_\beta(\mathrm{x}'\mid\hat{\mathrm{x}})}{\partial\hat{\mathrm{x}}} = 0\,, \quad \left[I - \beta C^{\mathrm{dualIB}}_{\mathrm{yy}'}(\hat{\mathrm{x}},\beta)\right]\frac{\partial \log p_\beta(\mathrm{y}'\mid\hat{\mathrm{x}})}{\partial\hat{\mathrm{x}}} = 0, \quad (11)$$

[4]Unless stated otherwise, we use the convention $\log = \ln \equiv \log_e$

with the matrices C^{dualIB} *given by:*

$$C^{\text{dualIB}}_{\text{xx}'}(\hat{\text{x}};\beta) = \sum_{\text{y},\tilde{\text{y}},\tilde{\text{x}}} p_\beta(\text{y}\mid\hat{\text{x}})p_\beta(\tilde{\text{x}}\mid\hat{\text{x}})\log\frac{p(\text{y}\mid\text{x})}{p(\text{y}\mid\tilde{\text{x}})}\cdot p_\beta(\text{x}'\mid\hat{\text{x}})p_\beta(\tilde{\text{y}}\mid\hat{\text{x}})\log\frac{p(\text{y}\mid\text{x}')}{p(\tilde{\text{y}}\mid\text{x}')}$$

$$C^{\text{dualIB}}_{\text{yy}'}(\hat{\text{x}};\beta) = \sum_{\text{x},\tilde{\text{x}},\tilde{\text{y}}} p_\beta(\text{x}\mid\hat{\text{x}})p_\beta(\tilde{\text{y}}\mid\hat{\text{x}})\log\frac{p(\text{y}\mid\text{x})}{p(\tilde{\text{y}}\mid\text{x})}\cdot p_\beta(\text{y}'\mid\hat{\text{x}})p_\beta(\tilde{\text{x}}\mid\hat{\text{x}})\log\frac{p(\text{y}'\mid\text{x})}{p(\text{y}'\mid\tilde{\text{x}})}. \tag{12}$$

The proof to *theorem* 2 is given in §A.3.

Lemma 3. *The matrices* $C^{\text{dualIB}}_{\text{xx}'}(\hat{\text{x}};\beta), C^{\text{dualIB}}_{\text{yy}'}(\hat{\text{x}};\beta)$ *have the same eigenvalues* $\{\lambda_i\}$*, with* $\lambda_1(\hat{\text{x}}) = 0$*. With binary* Y*, the critical points are obtained at* $\lambda_2(\hat{\text{x}}) = \beta^{-1}$*.*

The proof of *lemma* 3 given in section §A.3.1.

As in the IB, at the critical points, β_c^{dualIB}, the partial derivatives of the encoder and decoder with respect to β, $\partial\log p_\beta(\text{x}\mid\hat{\text{x}})/\partial\beta$, $\partial\log p_\beta(\text{y}\mid\hat{\text{x}})/\partial\beta$, have multiple (at least two) values. This results in discontinuities (cusps) in the encoder and decoder mutual information values as functions of β along the optimal line, with an undefined second derivative.

2.2 The Information Plane of the dualIB

The *Information-Plane*, $I_{\text{X}} = I(\hat{\text{X}};\text{X})$ vs. $I_{\text{Y}} = I(\hat{\text{X}};\text{Y})$, is the standard visualization of the compression-prediction trade off of the IB. It can be defined for *any* encoder once the decoder is the Bayes optimal (equation 3-(iii)), for which the I_{Y} is the actual information of the representation on the desired label (Tishby et al., 1999).

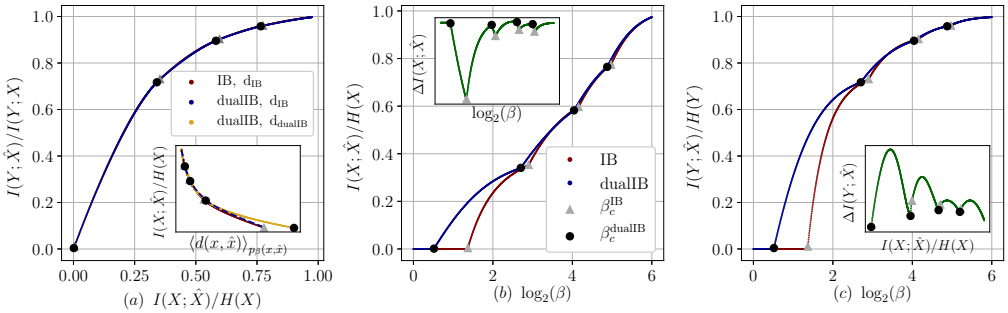

Figure 3: The IB's and dualIB's *Information Plane*. (a) I_{Y} vs. I_{X} for the two algorithms. The black dots are the dualIB critical points, β_c^{dualIB}, and the grey triangles are the IB critical points, β_c^{IB}. The corresponding distortion functions are shown in the inset. (b) The functions $I_{\text{X}}^{\text{IB}}(\beta)$ and $I_{\text{X}}^{\text{dualIB}}(\beta)$. Both curves are monotonic and concave between the critical points. The inset indicates the relative difference between the curves, where the alternating order of the critical points is clearly observed. (c) Similarly, $I_{\text{Y}}^{\text{IB}}(\beta)$ and $I_{\text{Y}}^{\text{dualIB}}(\beta)$ are monotonic and piece-wise concave. The relative discrepancy between the information curves is clearly minimized at the dualIB critical points (inset). The functions approach each other for large β.

Comparing the dualIB information curve to the IB curve shows the quality of this approximation. Figure 3(a) depicts this comparison. While we know that $I_{\text{Y}}^{\text{IB}}(\beta)$ is always higher, the two curves are almost indistinguishable. To better understand the relationship between these two curves, we look at the values of I_{X} and I_{Y} as functions of the corresponding β (Figure 3(b),(c)). The important role of the critical points is revealed as the corresponding cusps along these curves. As we argue below, the IB information values are strictly below those of the dualIB, but the distance between them is minimized precisely at the dual critical points.

Lemma 4. $I_{\text{X}}(\beta)$ *and* $I_{\text{Y}}(\beta)$*, along the optimal lines, are non-decreasing piece-wise concave, functions of* β*. When their second derivative (with respect to* β*) is defined, it is strictly negative.*

Lemma 5. *For any sub-optimal information curve ($I_{\mathrm{X}}, I_{\mathrm{Y}}$), $I_{\mathrm{X}}^{\mathrm{IB}}(\beta) \leq I_{\mathrm{X}}(\beta)$ and $I_{\mathrm{Y}}^{\mathrm{IB}}(\beta) \leq I_{\mathrm{Y}}(\beta)$, for all values of β.*

Proofs of the above are given in §A.4.

The information plane properties are summarized by the following theorem.

Theorem 6. *(i) The critical points of the two algorithms alternate: for each critical point, $\beta_c^{\mathrm{dualIB}} \leq \beta_c^{\mathrm{IB}}$. (ii) The distance between the two information curves is minimized precisely at the* dualIB *critical points $\beta_c^{\mathrm{dualIB}}$. (iii) The two curves approach each other as $\beta \to \infty$.*

Proof. The proof follows from *lemmas* 4 and 5, together with the critical points analysis above, and is only sketched here. As the encoder and decoder at the critical points, β_c^{IB} and $\beta_c^{\mathrm{dualIB}}$, have different left and right derivatives, they form cusps in the curves of the mutual information (I_{X} and I_{Y}) as functions of β. These cusps can only be consistent with the optimality of the IB curves if $\beta_c^{\mathrm{dualIB}} < \beta_c^{\mathrm{IB}}$ (this is true for any sub-optimal distortion), otherwise the curves intersect.

Moreover, at the dualIB critical points, the distance between the curves is minimized due to the strict concavity of the functions segments between the critical points. As the critical points imply discontinuity in the derivative, this results in a "jump" in the information values. Therefore, at any $\beta_c^{\mathrm{dualIB}}$ the distance between the curves has a (local) minimum. This is depicted in Figure 3, comparing $I_{\mathrm{X}}(\beta)$ and $I_{\mathrm{Y}}(\beta)$ and their differences for the two algorithms.

The two curves approach each other for large β since the two distortion functions become close in the low distortion limit (as long as $p(\mathrm{y} \mid \mathrm{x})$ is bounded away from 0). □

3 THE dualIB FOR EXPONENTIAL FAMILIES

Distributions of exponential families form the elegant theoretical core of parametric statistics and often emerge as maximum entropy (Jaynes, 1957) or stochastic equilibrium distributions, subject to observed constraints. They also form the class of parametric distributions for which exact, finite dimensional and additive, Minimal Sufficient Statistics exist (Kullback, 1959). One of the key properties of the dualIB is that it preserves the exponential form of the original data distribution, $p(\mathrm{x}, \mathrm{y})$ for all values of β.

Assuming that the data distribution is of the form:

$$p(\mathrm{y} \mid \mathrm{x}) = e^{-\sum_{r=0}^{d} \lambda^r(y) A_r(x)} = e^{-\boldsymbol{\lambda}(\mathrm{y})\cdot\boldsymbol{A}(\mathrm{x}) - \log Z_{\mathbf{y}|\mathbf{x}}(\mathrm{x})} = \prod_{r=0}^{d} e^{-\lambda^r(\mathrm{y}) A_r(\mathrm{x})} \,, \tag{13}$$

where $A_r(\mathrm{x})$ are d linearly independent functions of the input x and $\lambda^r(\mathrm{y})$ are functions of the label y, or the parameters of this exponential family. The $\boldsymbol{\lambda}(\mathrm{y})$ can also be considered Lagrange multipliers associated with the constraints conditional expectations $\langle A_r(\mathrm{x})\rangle_{p(\mathrm{x}|\mathrm{y})}$ in entropy maximization. The normalization factors, $Z_{\mathbf{y}|\mathbf{x}}(\mathrm{x})$, are written, for brevity, as $\lambda^0_{\mathbf{x}} \equiv \log(\sum_{\mathrm{y}} \prod_{r=1}^{d} e^{-\lambda^r(\mathrm{y}) A_r(\mathrm{x})})$ with $A_0(\mathrm{x}) \equiv 1$. We do not constrain the marginal $p(\mathrm{x})$.

The important fact about the exponential form is that all the mutual information, $I(\mathrm{X};\mathrm{Y})$, is fully captured by the d conditional expectations, $\langle A_r(\mathrm{x})\rangle_{p(\mathrm{x}|\mathrm{y})}$, since these are the (minimal) sufficient statistics for the parameters. This means that all the relevant information (in the training sample) is captured solely by d-dimensional empirical expectations. This can lead to a huge reduction in computational complexity (from $\dim(\mathrm{X})$ to d).

We next show that for the dualIB, this dimension reduction is preserved or improved along the dual information curve, for all values of β.

Theorem 7. *For data from an exponential family, equation 13, the optimal decoder of the* dualIB*, at a given β, is given by:*

$$p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) = e^{-\sum_{r=1}^{d} \lambda^r(\mathrm{y}) A_{r,\beta}(\hat{\mathrm{x}}) - \lambda^0_\beta(\hat{\mathrm{x}})} \,, \quad \lambda^0_\beta(\hat{\mathrm{x}}) = \log\Big(\sum_{\mathrm{y}} e^{-\sum_{r=1}^{d} \lambda^r(\mathrm{y}) A_{r,\beta}(\hat{\mathrm{x}})}\Big) \,, \tag{14}$$

and the respective encoder:

$$p_\beta(\hat{\mathrm{x}} \mid \mathrm{x}) = \frac{p_\beta(\hat{\mathrm{x}}) e^{\beta \lambda_\beta^0(\hat{\mathrm{x}})}}{Z_{\hat{\mathbf{x}}|\mathbf{x}}(\mathrm{x};\beta)} e^{-\beta \sum_{r=1}^{d} \lambda_\beta^r(\hat{\mathrm{x}})[A_r(\mathrm{x}) - A_{r,\beta}(\hat{\mathrm{x}})]} \,, \tag{15}$$

with the constraints and multipliers expectations,

$$A_{r,\beta}(\hat{\mathrm{x}}) \equiv \sum_{\mathrm{x}} p_\beta(\mathrm{x} \mid \hat{\mathrm{x}}) A_r(\mathrm{x}) \,,\; \lambda_\beta^r(\hat{\mathrm{x}}) \equiv \sum_{\mathrm{y}} p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) \lambda^r(\mathrm{y}) \,,\; 1 \le r \le d \,. \tag{16}$$

The complete derivations for this section are given in §A.5.

This defines a simplified iterative algorithm to solve the dualExpIB problem, since we can replace the decoders' update rule at each iteration in the dualIB algorithm (Algorithm 1, row 6) with the simplified expression given in equation 14. The dualExpIB algorithm is more efficient because decoders' update amounts to estimating the d dimensional *constraints expectation* $A_{r,\beta}(\hat{\mathrm{x}})$.

One would expect that as we decrease β the dimensionality of the *constraints expectation*, $\boldsymbol{A}_\beta(\hat{\mathrm{x}})$, should reduce as well. This is the most natural implication of the dimensionality reduction of $\hat{\mathrm{x}}$. Figure 4(c) indeed shows that this is the case. $\boldsymbol{A}_\beta(\hat{\mathrm{x}})$ follows the same bifurcation pattern as $p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})$.

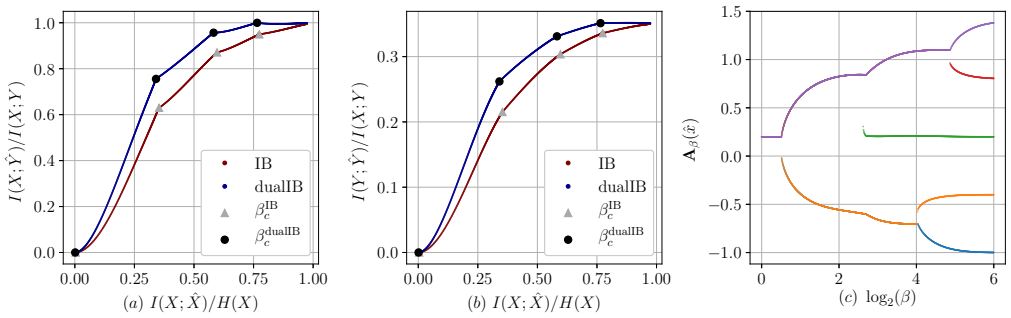

Figure 4: The two left plots present the *Information Planes* with respect to $\hat{\mathrm{Y}}$, in both the dualIB appears above the IB. (a) $I(\mathrm{X};\hat{\mathrm{Y}})$ vs. $I(\mathrm{X};\hat{\mathrm{X}})$. (b) $I(\mathrm{Y};\hat{\mathrm{Y}})$ vs. $I(\mathrm{X};\hat{\mathrm{X}})$. (c) The *constraint expectation*, $\boldsymbol{A}_\beta(\hat{\mathrm{x}})$ as a function of β for a problem with a single constraint ($d = 1$). The relevant dimension of the representation decreases at the critical points.

4 Conclusion

We presented a new, dual formulation of the Information Bottleneck framework, based on switching the arguments in the original IB distortion function. This simple change has several interesting consequences: (i) it provides a good approximation to the original IB while keeping the algorithm in the low relevant dimension of the original data. This can significantly reduce the complexity of finding good IB representations; (ii) it optimizes the information between the representation and the *predicted label* rather than the desired label as in the original IB. This can improve the generalization error when trained on small samples since the predicted label is the one used in practice. (iii) It preserves the exponential form of data from exponential families, while reducing the dimensionality of the compressed representations. This important property was known to be satisfied by the Gaussian IB (Chechik et al., 2005) but not known for other distributions. The Gaussian case is self-dual in that sense. Generalizing this property to other exponential families was an open problem for many years. (iv) The exponential form of the optimal encoder-decoder pairs allows for the application to distributions with special symmetries, which can be naturally expressed in this form.

The topology of the reduced representations is determined by the critical points, where the cardinality of the representation changes. The critical points form the skeleton of the optimal solutions and most of the computation time is spent near the critical points. We analyzed the critical points of the dualIB and their analytic relations to those of the original IB.

This paves the way for completely new applications of the IB to representation learning, in particular deep learning of data with low internal dimensionality or special symmetries.

REFERENCES

Lawrence D. Brown. Fundamentals of statistical exponential families with applications in statistical decision theory. *Lecture Notes-Monograph Series*, 9:i–279, 1986. ISSN 07492170. URL http://www.jstor.org/stable/4355554.

Gal Chechik, Amir Globerson, Naftali Tishby, and Yair Weiss. Information bottleneck for gaussian variables. *J. Mach. Learn. Res.*, 6:165–188, December 2005. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=1046920.1046926.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, New York, NY, USA, 2006. ISBN 0471241954.

Domenico Felice and Nihat Ay. Divergence functions in information geometry. In Frank Nielsen and Frédéric Barbaresco (eds.), *Geometric Science of Information - 4th International Conference, GSI 2019, Toulouse, France, August 27-29, 2019, Proceedings*, volume 11712 of *Lecture Notes in Computer Science*, pp. 433–442. Springer, 2019. ISBN 978-3-030-26979-1. doi: 10.1007/978-3-030-26980-7_45. URL https://doi.org/10.1007/978-3-030-26980-7_45.

Ran Gilad-bachrach, Amir Navot, and Naftali Tishby. An information theoretic tradeoff between complexity and accuracy. In *In Proceedings of the COLT*, pp. 595–609. Springer, 2003.

E. T. Jaynes. Information theory and statistical mechanics. *Phys. Rev.*, 106:620–630, May 1957. doi: 10.1103/PhysRev.106.620. URL https://link.aps.org/doi/10.1103/PhysRev.106.620.

S. Kullback. *Information Theory and Statistics*. Wiley, New York, 1959.

Amichai Painsky and Gregory W. Wornell. Bregman Divergence Bounds and the Universality of the Logarithmic Loss. *arXiv e-prints*, art. arXiv:1810.07014, Oct 2018.

Albert E. Parker, Tomáš Gedeon, and Alexander G. Dimitrov. Annealing and the rate distortion problem. In S. Becker, S. Thrun, and K. Obermayer (eds.), *Advances in Neural Information Processing Systems 15*, pp. 993–976. MIT Press, 2003. URL http://papers.nips.cc/paper/2264-annealing-and-the-rate-distortion-problem.pdf.

Ohad Shamir, Sivan Sabato, and Naftali Tishby. Learning and generalization with the information bottleneck. *Theor. Comput. Sci.*, 411:2696–2711, 2010.

Noam Slonim, Nir Friedman, and Naftali Tishby. Multivariate information bottleneck. *Neural Computation*, 18(8):1739–1789, 2006. doi: 10.1162/neco.2006.18.8.1739. URL https://doi.org/10.1162/neco.2006.18.8.1739. PMID: 16771652.

Naftali Tishby, Fernando C. Pereira, and William Bialek. The information bottleneck method. In *Proc. of the 37-th Annual Allerton Conference on Communication, Control and Computing*, pp. 368–377, 1999.

J. R. Tredicce, G. L. Lippi, Paul Mandel, B. Charasse, A. Chevalier, and B. Picqu. Critical slowing down at a bifurcation. *American Journal of Physics*, 72(6):799–809, 2004. doi: 10.1119/1.1688783. URL https://doi.org/10.1119/1.1688783.

G. Tusnady and I. Csiszar. Information geometry and alternating minimization procedures. *Statistics & Decisions: Supplement Issues*, 1:205–237, 1984.

Noga Zaslavsky and Naftali Tishby. Deterministic annealing and the evolution of optimal information bottleneck representations. *Preprint*, 2019.

A APPENDIX

A.1 DUALIB MATHEMATICAL FORMULATION

As presented in §1.5 the dualIB is solved with respect to the full Markov chain (equation 6) in which we introduce the new variable, $\hat{\mathrm{Y}}$, the *predicted label*. Thus, in analogy to the IB we want to write the optimization problem in term of $\hat{\mathrm{Y}}$.

Developing the expected distortion we find:

$$\begin{aligned}\left\langle d_{\mathrm{dualIB}}(\mathrm{x},\hat{\mathrm{x}})\right\rangle_{p_\beta(\mathrm{x},\hat{\mathrm{x}})} &= \sum_{\mathrm{x},\hat{\mathrm{x}}} p_\beta(\mathrm{x},\hat{\mathrm{x}}) \sum_{\hat{\mathrm{y}}} p_\beta(\hat{\mathrm{y}}\mid\hat{\mathrm{x}}) \log\frac{p_\beta(\hat{\mathrm{y}}\mid\hat{\mathrm{x}})}{p(\hat{\mathrm{y}}\mid\mathrm{x})} \\ &= \sum_{\hat{\mathrm{x}},\hat{\mathrm{y}}} p_\beta(\hat{\mathrm{x}}) p_\beta(\hat{\mathrm{y}}\mid\hat{\mathrm{x}}) \log\frac{p_\beta(\hat{\mathrm{y}}\mid\hat{\mathrm{x}})}{p_\beta(\hat{\mathrm{y}})} - \sum_{\mathrm{x},\hat{\mathrm{y}}} p(\mathrm{x}) p_\beta(\hat{\mathrm{y}}\mid\mathrm{x}) \log\frac{p_\beta(\hat{\mathrm{y}}\mid\mathrm{x})}{p_\beta(\hat{\mathrm{y}})} \\ &+ \sum_{\mathrm{x},\hat{\mathrm{y}}} p(\mathrm{x}) p_\beta(\hat{\mathrm{y}}\mid\mathrm{x}) \log\frac{p_\beta(\hat{\mathrm{y}}\mid\mathrm{x})}{p(\hat{\mathrm{y}}\mid\mathrm{x})} \\ &= I(\hat{\mathrm{X}};\hat{\mathrm{Y}}) - I(\mathrm{X};\hat{\mathrm{Y}}) + \left\langle D[p_\beta(\hat{\mathrm{y}}\mid\mathrm{x})\|p(\hat{\mathrm{y}}\mid\mathrm{x})]\right\rangle_{p(\mathrm{x})}.\end{aligned}$$

Allowing the dual optimization problem to be written as:

$$\mathcal{F}^*[p(\hat{\mathrm{x}}\mid\mathrm{x});p(\hat{\mathrm{x}});p(\mathrm{y}\mid\hat{\mathrm{x}})] = I(\mathrm{X};\hat{\mathrm{X}}) - \beta\Big\{I(\mathrm{X};\hat{\mathrm{Y}}) - I(\hat{\mathrm{X}};\hat{\mathrm{Y}}) - \left\langle D[p_\beta(\hat{\mathrm{y}}\mid\mathrm{x})\|p(\hat{\mathrm{y}}\mid\mathrm{x})]\right\rangle_{p(\mathrm{x})}\Big\}.$$

A.2 THE DUALIB SOLUTIONS

To prove *theorem* 1 we want to obtain the normalized distributions minimizing the dualIB rate-distortion problem.

Proof. (i) Given that the problem is formulated as a rate-distortion problem the encoder's update rule must be the known minimizer of the distortion function (Cover & Thomas, 2006). Thus the IB encoder with the dual distortion is plugged in. (ii) For the decoder, by considering a small perturbation in the distortion $d_{\mathrm{dualIB}}(\mathrm{x},\hat{\mathrm{x}})$, with $\alpha(\hat{\mathrm{x}})$ the normalization Lagrange multiplier, we obtain:

$$\begin{aligned}\delta d_{\mathrm{dualIB}}(\mathrm{x},\hat{\mathrm{x}}) &= \delta\left(\sum_{\mathrm{y}} p_\beta(\mathrm{y}\mid\hat{\mathrm{x}}) \log\frac{p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})}{p(\mathrm{y}\mid\mathrm{x})} + \alpha(\hat{\mathrm{x}})\left(\sum_{\mathrm{y}} p_\beta(\mathrm{y}\mid\hat{\mathrm{x}}) - 1\right)\right) \\ \frac{\delta d_{\mathrm{dualIB}}(\mathrm{x},\hat{\mathrm{x}})}{\delta p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})} &= \log\frac{p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})}{p(\mathrm{y}\mid\mathrm{x})} + 1 + \alpha(\hat{\mathrm{x}}).\end{aligned}$$

Hence, minimizing the expected distortion becomes:

$$\begin{aligned}0 &= \sum_{\mathrm{x}} p_\beta(\mathrm{x}\mid\hat{\mathrm{x}})\left[\log\frac{p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})}{p(\mathrm{y}\mid\mathrm{x})} + 1\right] + \alpha(\hat{\mathrm{x}}) \\ &= \log p_\beta(\mathrm{y}\mid\hat{\mathrm{x}}) - \sum_{\mathrm{x}} p_\beta(\mathrm{x}\mid\hat{\mathrm{x}}) \log p(\mathrm{y}\mid\mathrm{x}) + 1 + \alpha(\hat{\mathrm{x}}),\end{aligned}$$

which yields Algorithm 1, row 6. □

Considering the dualIB encoder-decoder, Algorithm 1, we find that $\left\langle d_{\mathrm{dualIB}}(\mathrm{x},\hat{\mathrm{x}})\right\rangle_{p_\beta(\mathrm{x},\hat{\mathrm{x}})}$ reduces to the expectation of the decoder's log partition function:

$$\begin{aligned}\left\langle d_{\mathrm{dualIB}}(\mathrm{x},\hat{\mathrm{x}})\right\rangle_{p_\beta(\mathrm{x},\hat{\mathrm{x}})} &= \sum_{\mathrm{x},\hat{\mathrm{x}}} p_\beta(\mathrm{x},\hat{\mathrm{x}}) \sum_{\mathrm{y}} p_\beta(\mathrm{y}\mid\hat{\mathrm{x}}) \log\frac{p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})}{p(\mathrm{y}\mid\mathrm{x})} \\ &= -\left\langle\log Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}};\beta)\right\rangle_{p_\beta(\hat{\mathrm{x}})} + \sum_{\hat{\mathrm{x}},\mathrm{y}} p_\beta(\hat{\mathrm{x}})\left[\sum_{\mathrm{x}'} p_\beta(\mathrm{x}'\mid\hat{\mathrm{x}}) \log p(\mathrm{y}\mid\mathrm{x}') - \sum_{\mathrm{x}} p_\beta(\mathrm{x}\mid\hat{\mathrm{x}}) \log p(\mathrm{y}\mid\mathrm{x})\right] \\ &= -\left\langle\log Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}};\beta)\right\rangle_{p_\beta(\hat{\mathrm{x}})}.\end{aligned}$$

A.3 STABILITY ANALYSIS

Here we provide the detailed stability analysis allowing the definition of the matrices $C^{\mathrm{dualIB}}_{xx'}, C^{\mathrm{dualIB}}_{yy'}$ (equation 12) and which allows us to claim that they obey the same rules as the C matrices (equation 5) in equation 4. Considering a variation in $\hat{\mathrm{x}}$ we get:

$$\begin{aligned}
\frac{\partial \log p_\beta(\mathrm{x} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} =& \beta \sum_{\mathrm{y}} p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) \left(\log \frac{p(\mathrm{y} \mid \mathrm{x})}{p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})} - 1 \right) \frac{\partial \log p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} \\
=& \beta \sum_{\mathrm{y}} p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) \left[\log p(\mathrm{y} \mid \mathrm{x}) - \sum_{\tilde{\mathrm{x}}} p_\beta(\tilde{\mathrm{x}} \mid \hat{\mathrm{x}}) \log p(\mathrm{y} \mid \tilde{\mathrm{x}}) \right] \frac{\partial \log p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} \\
&+ \beta \sum_{\mathrm{y}} \log Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}}; \beta) \frac{\partial p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} \\
=& \beta \sum_{\mathrm{y}, \tilde{\mathrm{x}}} p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{x}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y} \mid \mathrm{x})}{p(\mathrm{y} \mid \tilde{\mathrm{x}})} \frac{\partial \log p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}}, \qquad (17)
\end{aligned}$$

$$\begin{aligned}
\frac{\partial \log p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} =& - \frac{1}{Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}}; \beta)} \frac{\partial Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}}; \beta)}{\partial \hat{\mathrm{x}}} + \sum_{\mathrm{x}} p_\beta(\mathrm{x} \mid \hat{\mathrm{x}}) \log p(\mathrm{y} \mid \mathrm{x}) \frac{\partial \log p_\beta(\mathrm{x} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} \\
=& - \sum_{\tilde{\mathrm{y}}} p_\beta(\tilde{\mathrm{y}} \mid \hat{\mathrm{x}}) \sum_{\mathrm{x}} p_\beta(\mathrm{x} \mid \hat{\mathrm{x}}) \log p(\tilde{\mathrm{y}} \mid \mathrm{x}) \frac{\partial \log p_\beta(\mathrm{x} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} \\
&+ \sum_{\mathrm{x}} p_\beta(\mathrm{x} \mid \hat{\mathrm{x}}) \log p(\mathrm{y} \mid \mathrm{x}) \frac{\partial \log p_\beta(\mathrm{x} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} \\
=& \sum_{\mathrm{x}, \tilde{\mathrm{y}}} p_\beta(\mathrm{x} \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{y}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y} \mid \mathrm{x})}{p(\tilde{\mathrm{y}} \mid \mathrm{x})} \frac{\partial \log p_\beta(\mathrm{x} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}}. \qquad (18)
\end{aligned}$$

Substituting equation 18 into equation 17 and vice versa one obtains:

$$\begin{aligned}
\frac{\partial \log p_\beta(\mathrm{x} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} =& \beta \sum_{\mathrm{x}', \mathrm{y}, \tilde{\mathrm{y}}, \tilde{\mathrm{x}}} p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{x}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y} \mid \mathrm{x})}{p(\mathrm{y} \mid \tilde{\mathrm{x}})} \\
& \cdot p_\beta(\mathrm{x}' \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{y}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y} \mid \mathrm{x}')}{p(\tilde{\mathrm{y}} \mid \mathrm{x}')} \frac{\partial \log p_\beta(\mathrm{x}' \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} \\
\frac{\partial \log p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} =& \beta \sum_{\mathrm{x}, \mathrm{y}', \tilde{\mathrm{x}}, \tilde{\mathrm{y}}} p_\beta(\mathrm{x} \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{y}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y} \mid \mathrm{x})}{p(\tilde{\mathrm{y}} \mid \mathrm{x})} \\
& \cdot p_\beta(\mathrm{y}' \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{x}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y}' \mid \mathrm{x})}{p(\mathrm{y}' \mid \tilde{\mathrm{x}})} \frac{\partial \log p_\beta(\mathrm{y}' \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}}.
\end{aligned}$$

We now define the C^{dualIB} matrices as follows:

$$\begin{aligned}
C^{\mathrm{dualIB}}_{\mathrm{xx}'}(\hat{\mathrm{x}}; \beta) =& \sum_{\mathrm{y}, \tilde{\mathrm{y}}, \tilde{\mathrm{x}}} p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{x}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y} \mid \mathrm{x})}{p(\mathrm{y} \mid \tilde{\mathrm{x}})} \cdot p_\beta(\mathrm{x}' \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{y}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y} \mid \mathrm{x}')}{p(\tilde{\mathrm{y}} \mid \mathrm{x}')} \\
C^{\mathrm{dualIB}}_{\mathrm{yy}'}(\hat{\mathrm{x}}; \beta) =& \sum_{\mathrm{x}, \tilde{\mathrm{x}}, \tilde{\mathrm{y}}} p_\beta(\mathrm{x} \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{y}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y} \mid \mathrm{x})}{p(\tilde{\mathrm{y}} \mid \mathrm{x})} \cdot p_\beta(\mathrm{y}' \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{x}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y}' \mid \mathrm{x})}{p(\mathrm{y}' \mid \tilde{\mathrm{x}})}.
\end{aligned}$$

Using the above definition we have an equivalence to equation 4 in the form of:

$$\left[I - \beta C^{\mathrm{dualIB}}_{\mathrm{xx}'}(\hat{\mathrm{x}}, \beta) \right] \frac{\partial \log p_\beta(\mathrm{x}' \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} = 0 \,, \quad \left[I - \beta C^{\mathrm{dualIB}}_{\mathrm{yy}'}(\hat{\mathrm{x}}, \beta) \right] \frac{\partial \log p_\beta(\mathrm{y}' \mid \hat{\mathrm{x}})}{\partial \hat{\mathrm{x}}} = 0.$$

Note that for the binary case, the matrices may be simplified to:

$$\begin{aligned}
C^{\mathrm{dualIB}}_{\mathrm{xx}'}(\hat{\mathrm{x}}; \beta) =& \sum_{\mathrm{y}, \tilde{\mathrm{x}}} p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{x}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y} \mid \mathrm{x})}{p(\mathrm{y} \mid \tilde{\mathrm{x}})} \cdot p_\beta(\mathrm{x}' \mid \hat{\mathrm{x}}) (1 - p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})) \log \frac{p(\mathrm{y} \mid \mathrm{x}')}{1 - p(\mathrm{y} \mid \mathrm{x}')} \\
C^{\mathrm{dualIB}}_{\mathrm{yy}'}(\hat{\mathrm{x}}; \beta) =& \sum_{\mathrm{x}, \tilde{\mathrm{x}}} p_\beta(\mathrm{x} \mid \hat{\mathrm{x}}) (1 - p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})) \log \frac{p(\mathrm{y} \mid \mathrm{x})}{1 - p(\mathrm{y} \mid \mathrm{x})} \cdot p_\beta(\mathrm{y}' \mid \hat{\mathrm{x}}) p_\beta(\tilde{\mathrm{x}} \mid \hat{\mathrm{x}}) \log \frac{p(\mathrm{y}' \mid \mathrm{x})}{p(\mathrm{y}' \mid \tilde{\mathrm{x}})}.
\end{aligned}$$

Repeating the above steps considering a small change in β we see that the above obey the same, up to an additive constant, non-linear eigenvalue problem as the $\hat{\mathrm{x}}$ derivatives. Therefore, we conclude that they appear at the same critical points as the variation w.r.t $\hat{\mathrm{x}}$ (the solutions to the non-linear eigenvalue problem depend solely on the derivative coefficients).

A.3.1 Proof of Lemma 3

We show that the C^{dualIB} matrices share the same eigenvalues with $\lambda_1(\hat{\mathrm{x}}) = 0$.

Proof. The matrices, $C^{\mathrm{dualIB}}_{\mathrm{xx}'}(\hat{\mathrm{x}};\beta)$, $C^{\mathrm{dualIB}}_{\mathrm{yy}'}(\hat{\mathrm{x}};\beta)$, are given by:

$$C^{\mathrm{dualIB}}_{\mathrm{xx}'}(\hat{\mathrm{x}};\beta) = A_{\mathrm{xy}}(\hat{\mathrm{x}};\beta)B_{\mathrm{yx}'}(\hat{\mathrm{x}};\beta)\ ,\ C^{\mathrm{dualIB}}_{\mathrm{yy}'}(\hat{\mathrm{x}};\beta) = B_{\mathrm{yx}}(\hat{\mathrm{x}};\beta)A_{\mathrm{xy}'}(\hat{\mathrm{x}};\beta),$$

with:

$$A_{\mathrm{xy}}(\hat{\mathrm{x}};\beta) = p_\beta(\mathrm{y}\mid\hat{\mathrm{x}})\sum_{\tilde{\mathrm{x}}} p_\beta(\tilde{\mathrm{x}}\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}\mid\mathrm{x})}{p(\mathrm{y}\mid\tilde{\mathrm{x}})}\ ,\ B_{\mathrm{yx}}(\hat{\mathrm{x}};\beta) = p_\beta(\mathrm{x}\mid\hat{\mathrm{x}})\sum_{\tilde{\mathrm{y}}} p_\beta(\tilde{\mathrm{y}}\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}\mid\mathrm{x})}{p(\tilde{\mathrm{y}}\mid\mathrm{x})}.$$

Given that the matrices are obtained by the multiplication of the same matrices, it follows that they have the same eigenvalues $\{\lambda_i(\hat{\mathrm{x}};\beta)\}$.

To prove that $\lambda_1(\hat{\mathrm{x}};\beta) = 0$ we show that $\det(C^{\mathrm{dualIB}}_{\mathrm{yy}'}) = 0$. We present the exact calculation for a binary label $\mathrm{Y}\in\{\mathrm{y}_0,\mathrm{y}_1\}$ (the argument for general Y follows by encoding the label as a sequence of bits and discussing the first bit only, as a binary case.):

$$\begin{aligned}
\det(C^{\mathrm{dualIB}}_{\mathrm{yy}'}(\hat{\mathrm{x}};\beta)) =& \sum_{\mathrm{x},\tilde{\mathrm{x}}} p_\beta(\mathrm{x}\mid\hat{\mathrm{x}})p_\beta(\mathrm{y}_1\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_0\mid\mathrm{x})}{p(\mathrm{y}_1\mid\mathrm{x})}\cdot p_\beta(\mathrm{y}_0\mid\hat{\mathrm{x}})p_\beta(\tilde{\mathrm{x}}\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_0\mid\mathrm{x})}{p(\mathrm{y}_0\mid\tilde{\mathrm{x}})}\\
&\cdot\sum_{\mathrm{x}',\tilde{\mathrm{x}}',} p_\beta(\mathrm{x}'\mid\hat{\mathrm{x}})p_\beta(\mathrm{y}_0\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_1\mid\mathrm{x}')}{p(\mathrm{y}_0\mid\mathrm{x}')}\cdot p_\beta(\mathrm{y}_1\mid\hat{\mathrm{x}})p_\beta\big(\tilde{\mathrm{x}}'\mid\hat{\mathrm{x}}\big)\log\frac{p(\mathrm{y}_1\mid\mathrm{x}')}{p\big(\mathrm{y}_1\mid\tilde{\mathrm{x}}'\big)}\\
&-\sum_{\mathrm{x},\tilde{\mathrm{x}}} p_\beta(\mathrm{x}\mid\hat{\mathrm{x}})p_\beta(\mathrm{y}_0\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_1\mid\mathrm{x})}{p(\mathrm{y}_0\mid\mathrm{x})}\cdot p_\beta(\mathrm{y}_0\mid\hat{\mathrm{x}})p_\beta(\tilde{\mathrm{x}}\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_0\mid\mathrm{x})}{p(\mathrm{y}_0\mid\tilde{\mathrm{x}})}\\
&\cdot\sum_{\mathrm{x}',\tilde{\mathrm{x}}'} p_\beta(\mathrm{x}'\mid\hat{\mathrm{x}})p_\beta(\mathrm{y}_1\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_0\mid\mathrm{x}')}{p(\mathrm{y}_1\mid\mathrm{x}')}\cdot p_\beta(\mathrm{y}_1\mid\hat{\mathrm{x}})p_\beta\big(\tilde{\mathrm{x}}'\mid\hat{\mathrm{x}}\big)\log\frac{p(\mathrm{y}_1\mid\mathrm{x}')}{p\big(\mathrm{y}_1\mid\tilde{\mathrm{x}}'\big)}\\
=&\sum_{\mathrm{x},\mathrm{x}',\tilde{\mathrm{x}},\tilde{\mathrm{x}}'} p_\beta(\mathrm{x}\mid\hat{\mathrm{x}})p_\beta(\mathrm{x}'\mid\hat{\mathrm{x}})p^2_\beta(\mathrm{y}_0\mid\hat{\mathrm{x}})p^2_\beta(\mathrm{y}_1\mid\hat{\mathrm{x}})p_\beta(\tilde{\mathrm{x}}\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_0\mid\mathrm{x})}{p(\mathrm{y}_0\mid\tilde{\mathrm{x}})}p_\beta\big(\tilde{\mathrm{x}}'\mid\hat{\mathrm{x}}\big)\log\frac{p(\mathrm{y}_1\mid\mathrm{x}')}{p\big(\mathrm{y}_1\mid\tilde{\mathrm{x}}'\big)}\\
&\cdot\left[\log\frac{p(\mathrm{y}_0\mid\mathrm{x})}{p(\mathrm{y}_1\mid\mathrm{x})}\log\frac{p(\mathrm{y}_1\mid\mathrm{x}')}{p(\mathrm{y}_0\mid\mathrm{x}')}-\log\frac{p(\mathrm{y}_0\mid\mathrm{x})}{p(\mathrm{y}_1\mid\mathrm{x})}\log\frac{p(\mathrm{y}_1\mid\mathrm{x}')}{p(\mathrm{y}_0\mid\mathrm{x}')}\right]=0.
\end{aligned}$$

Given that the determinant is 0 implies that $\lambda_1(\hat{\mathrm{x}}) = 0$. □

For a binary problem we can describe the non-zero eigenvalue using $\lambda_2(\hat{\mathrm{x}}) = \mathrm{Tr}(C^{\mathrm{dualIB}}_{\mathrm{yy}'}(\hat{\mathrm{x}};\beta))$. That is:

$$\begin{aligned}
\lambda_2(\hat{\mathrm{x}}) =& \sum_{\mathrm{x},\tilde{\mathrm{x}}} p_\beta(\mathrm{x}\mid\hat{\mathrm{x}})p_\beta(\mathrm{y}_1\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_0\mid\mathrm{x})}{p(\mathrm{y}_1\mid\mathrm{x})}\cdot p_\beta(\mathrm{y}_0\mid\hat{\mathrm{x}})p_\beta(\tilde{\mathrm{x}}\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_0\mid\mathrm{x})}{p(\mathrm{y}_0\mid\tilde{\mathrm{x}})}\\
&+\sum_{\mathrm{x},\tilde{\mathrm{x}}} p_\beta(\mathrm{x}\mid\hat{\mathrm{x}})p_\beta(\mathrm{y}_0\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_1\mid\mathrm{x})}{p(\mathrm{y}_0\mid\mathrm{x})}\cdot p_\beta(\mathrm{y}_1\mid\hat{\mathrm{x}})p_\beta(\tilde{\mathrm{x}}\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_1\mid\mathrm{x})}{p(\mathrm{y}_1\mid\tilde{\mathrm{x}})}\\
=&p_\beta(\mathrm{y}_1\mid\hat{\mathrm{x}})p_\beta(\mathrm{y}_0\mid\hat{\mathrm{x}})\sum_{\mathrm{x},\tilde{\mathrm{x}}} p_\beta(\mathrm{x}\mid\hat{\mathrm{x}})p_\beta(\tilde{\mathrm{x}}\mid\hat{\mathrm{x}})\log\frac{p(\mathrm{y}_0\mid\mathrm{x})}{p(\mathrm{y}_1\mid\mathrm{x})}\left[\log\frac{p(\mathrm{y}_0\mid\mathrm{x})}{p(\mathrm{y}_0\mid\tilde{\mathrm{x}})}-\log\frac{p(\mathrm{y}_1\mid\mathrm{x})}{p(\mathrm{y}_1\mid\tilde{\mathrm{x}})}\right].
\end{aligned}$$

A.4 INFORMATION PLANE ANALYSIS

We rely on known results for the rate-distortion problem and the information plane:

Lemma 8. *$I(\mathrm{X};\hat{\mathrm{X}})$ is a non-increasing convex function of the distortion $\langle d(\mathrm{x},\hat{\mathrm{x}})\rangle_{p_\beta(\mathrm{x},\hat{\mathrm{x}})}$ with a slope of $-\beta$.*

We emphasis that this is a general result of rate-distortion thus holds for the dualIB as well.

Lemma 9. *For a fixed encoder $p_\beta(\hat{\mathrm{x}} \mid \mathrm{x})$ and the Bayes optimal decoder $p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})$:*

$$\langle d_{\mathrm{IB}}(\mathrm{x},\hat{\mathrm{x}})\rangle_{p_\beta(\mathrm{x},\hat{\mathrm{x}})} = I(\mathrm{X};\mathrm{Y}) - I(\hat{\mathrm{X}};\mathrm{Y}).$$

Thus, the information curve, I_Y vs. I_X, is a non-decreasing concave function with a positive slope, β^{-1}. The concavity implies that β increases along the curve.

(Cover & Thomas, 2006; Gilad-bachrach et al., 2003).

A.4.1 PROOF OF LEMMA 4

In the following section we provide a proof to *lemma* 4, for the IB and dualIB problems.

Proof. We want to analyze the behavior of $I_\mathrm{X}(\beta)$, $I_\mathrm{Y}(\beta)$, that is the change in each term as a function of the corresponding β. From *lemma* 9, the concavity of the information curve, we can deduce that both are non-decreasing functions of β. As the two β derivatives are proportional it's enough to discuss the first one.

Next, we focus on their behavior between two critical points. That is, where the cardinality of $\hat{\mathrm{X}}$ is fixed (clusters are "static"). For "static" clusters, the β derivative of I_X, along the optimal line is given by:

$$\begin{aligned}
\frac{\partial I(\mathrm{X};\hat{\mathrm{X}})}{\partial\beta} &= -\frac{\partial}{\partial\beta}\left[\sum_{\mathrm{x},\hat{\mathrm{x}}} p_\beta(\mathrm{x},\hat{\mathrm{x}})\big(\log Z_{\hat{\mathbf{x}}|\mathbf{x}}(\mathrm{x};\beta) + \beta d(\mathrm{x},\hat{\mathrm{x}})\big)\right] \\
&= -\beta\left\langle d(\mathrm{x},\hat{\mathrm{x}})\frac{\partial \log p_\beta(\hat{\mathrm{x}} \mid \mathrm{x})}{\partial\beta}\right\rangle_{p_\beta(\mathrm{x},\hat{\mathrm{x}})} \\
&\approx \beta\left\langle d(\mathrm{x},\hat{\mathrm{x}})\left[\frac{\partial \log Z_{\hat{\mathbf{x}}|\mathbf{x}}(\mathrm{x};\beta)}{\partial\beta} + d(\mathrm{x},\hat{\mathrm{x}})\right]\right\rangle_{p_\beta(\mathrm{x},\hat{\mathrm{x}})} \\
&\approx \beta\left\langle \underbrace{\left\langle d^2(\mathrm{x},\hat{\mathrm{x}})\right\rangle_{p_\beta(\hat{\mathrm{x}}|\mathrm{x})} - \left\langle d(\mathrm{x},\hat{\mathrm{x}})\right\rangle^2_{p_\beta(\hat{\mathrm{x}}|\mathrm{x})}}_{\mathrm{Var}(d(\mathrm{x}))}\right\rangle_{p(\mathrm{x})} .
\end{aligned}$$

This first of all reassures that the function is non-decreasing as $\mathrm{Var}(d(\mathrm{x})) \geq 0$.

The piece-wise concavity follows from the fact that when the number of clusters is fixed (between the critical points) - increasing β decreases the clusters conditional entropy $H(\hat{\mathrm{X}} \mid \mathrm{X})$, as the encoder becomes more deterministic. The mutual information is bounded by $H(\hat{\mathrm{X}})$ and it's β derivative decreases. Further, between the critical points there are no sign changes in the second β derivative. □

A.4.2 PROOF OF LEMMA 5

Proof. The information curve has a positive slope, β^{-1}, with β increasing along it, *lemma* 9. That is, given a value of β, there exists a pair $I^{\mathrm{IB}}_\mathrm{Y}(\beta), I^{\mathrm{IB}}_\mathrm{X}(\beta)$ such that $\partial I^{\mathrm{IB}}_\mathrm{Y}(\beta)/\partial I^{\mathrm{IB}}_\mathrm{X}(\beta) = \beta^{-1}$. Now, consider a sub-optimal information curve, $I^*_\mathrm{Y}, I^*_\mathrm{X}$. There exist values β', β'' such that:

$$I^*_\mathrm{X}(\beta') = I^{\mathrm{IB}}_\mathrm{X}(\beta)\ ,\ \ I^*_\mathrm{Y}(\beta'') = I^{\mathrm{IB}}_\mathrm{Y}(\beta).$$

The optimality of the IB implies that sub-optimal curves lie below it; i.e, the IB slope is steeper:

$$\beta^{-1} > \beta'^{-1}\ ,\ \ \beta^{-1} > \beta''^{-1}.$$

Thus, given that β increases along the information curve, it holds that:

$$I_{\mathrm{X}}^{*}(\beta) > I_{\mathrm{X}}^{*}(\beta') = I_{\mathrm{X}}^{\mathrm{IB}}(\beta)\ ,\ \ I_{\mathrm{Y}}^{*}(\beta) > I_{\mathrm{Y}}^{*}(\beta'') = I_{\mathrm{Y}}^{\mathrm{IB}}(\beta).$$

□

A.5 DERIVATION OF THE DUALEXPIB

We provide elaborate derivations to *theorem* 7; that is, we obtain the dualIB optimal encoder-decoder under the exponential assumption over the data. We use the notations defined in §3.

- The *decoder*, equation 14.
  Substituting the exponential assumption into the dualIB $\log$-decoder yields:

$$\begin{aligned}\log p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) &= \sum_{\mathrm{x}} p_\beta(\mathrm{x} \mid \hat{\mathrm{x}}) \log p(\mathrm{y} \mid \mathrm{x}) - \log Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}};\beta) \\ &= -\sum_{\mathrm{x}} \sum_{r=0}^{d} p_\beta(\mathrm{x} \mid \hat{\mathrm{x}}) \lambda^r(\mathrm{y}) A_r(\mathrm{x}) - \log Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}};\beta) \\ &= -\sum_{r=1}^{d} \lambda^r(\mathrm{y}) A_{r,\beta}(\hat{\mathrm{x}}) - \left\langle \lambda_{\mathbf{x}}^0 \right\rangle_{p_\beta(\mathrm{x}|\hat{\mathrm{x}})} - \log Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}};\beta).\end{aligned}$$

  Taking a closer look at the normalization term:

$$Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}};\beta) = \sum_{\mathrm{y}} e^{\sum_{\mathrm{x}} p_\beta(\mathrm{x}|\hat{\mathrm{x}}) \log p(\mathrm{y}|\mathrm{x})} = e^{-\left\langle \lambda_{\mathbf{x}}^0 \right\rangle_{p_\beta(\mathrm{x}|\hat{\mathrm{x}})}} \sum_{\mathrm{y}} e^{-\sum_{r=1}^{d} \lambda^r(\mathrm{y}) A_{r,\beta}(\hat{\mathrm{x}})}$$

$$\log Z_{\mathbf{y}|\hat{\mathbf{x}}}(\hat{\mathrm{x}};\beta) = -\left\langle \lambda_{\mathbf{x}}^0 \right\rangle_{p_\beta(\mathrm{x}|\hat{\mathrm{x}})} + \log\left(\sum_{\mathrm{y}} e^{-\sum_{r=1}^{d} \lambda^r(\mathrm{y}) A_{r,\beta}(\hat{\mathrm{x}})} \right).$$

  From which it follows that $\lambda_\beta^0(\hat{\mathrm{x}})$ is given by:

$$\lambda_\beta^0(\hat{\mathrm{x}}) = \log\left(\sum_{\mathrm{y}} e^{-\sum_{r=1}^{d} \lambda^r(\mathrm{y}) A_{r,\beta}(\hat{\mathrm{x}})} \right),$$

  and we can conclude that the $\mathrm{dualExpIB}$ decoder takes the form:

$$\log p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) = -\sum_{r=0}^{d} \lambda^r(\mathrm{y}) A_{r,\beta}(\hat{\mathrm{x}}).$$

- The *encoder*, equation 15.
  The core of the encoder is the dual distortion function which may now be written as:

$$\begin{aligned}d_{\mathrm{dualIB}}(\mathrm{x}, \hat{\mathrm{x}}) &= \sum_{\mathrm{y}} p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) \log \frac{p_\beta(\mathrm{y} \mid \hat{\mathrm{x}})}{p(\mathrm{y} \mid \mathrm{x})} \\ &= \sum_{\mathrm{y}} p_\beta(\mathrm{y} \mid \hat{\mathrm{x}}) \left[\left(\lambda_{\mathbf{x}}^0 - \lambda_\beta^0(\hat{\mathbf{x}}) \right) + \sum_{r=1}^{d} \lambda^r(\mathrm{y}) (A_r(\mathrm{x}) - A_{r,\beta}(\hat{\mathrm{x}})) \right] \\ &= \lambda_{\mathbf{x}}^0 - \lambda_\beta^0(\hat{\mathrm{x}}) + \sum_{r=1}^{d} \lambda_\beta^r(\hat{\mathrm{x}}) (A_r(\mathrm{x}) - A_{r,\beta}(\hat{\mathrm{x}})),\end{aligned}$$

  substituting this into the encoder's definition we obtain:

$$\begin{aligned}p_\beta(\hat{\mathrm{x}} \mid \mathrm{x}) &= \frac{p_\beta(\hat{\mathrm{x}})}{Z_{\hat{\mathbf{x}}|\mathbf{x}}(\mathrm{x};\beta)} e^{-\beta\left[\lambda_{\mathbf{x}}^0 - \lambda_\beta^0(\hat{\mathrm{x}}) + \sum_{r=1}^{d} \lambda_\beta^r(\hat{\mathrm{x}})[A_r(\mathrm{x}) - A_{r,\beta}(\hat{\mathrm{x}})]\right]} \\ &= \frac{p_\beta(\hat{\mathrm{x}}) e^{\beta \lambda_\beta^0(\hat{\mathrm{x}})}}{Z_{\hat{\mathbf{x}}|\mathbf{x}}(\mathrm{x};\beta)} e^{-\beta \sum_{r=1}^{d} \lambda_\beta^r(\hat{\mathrm{x}})[A_r(\mathrm{x}) - A_{r,\beta}(\hat{\mathrm{x}})]}.\end{aligned}$$