

Article

# The Convex Information Bottleneck Lagrangian

Borja Rodríguez Gálvez <sup>\*,†</sup> , Ragnar Thobaben <sup>\*,†</sup>  and Mikael Skoglund <sup>\*,†</sup> 

Department of Intelligent Systems, Division of Information Science and Engineering (ISE), KTH Royal Institute of Technology, 11428 Stockholm, Sweden

\* Correspondence: borjarg@kth.se (B.R.G.); ragnart@kth.se (R.T.); skoglund@kth.se (M.S.)

† Current address: Malvinas väg 10, 100 44 Stockholm, Sweden

Received: 9 December 2019; Accepted: 8 January 2020; Published: 14 January 2020



**Abstract:** The information bottleneck (IB) problem tackles the issue of obtaining relevant compressed representations  $T$  of some random variable  $X$  for the task of predicting  $Y$ . It is defined as a constrained optimization problem that maximizes the information the representation has about the task,  $I(T; Y)$ , while ensuring that a certain level of compression  $r$  is achieved (i.e.,  $I(X; T) \leq r$ ). For practical reasons, the problem is usually solved by maximizing the IB Lagrangian (i.e.,  $\mathcal{L}_{\text{IB}}(T; \beta) = I(T; Y) - \beta I(X; T)$ ) for many values of  $\beta \in [0, 1]$ . Then, the curve of maximal  $I(T; Y)$  for a given  $I(X; T)$  is drawn and a representation with the desired predictability and compression is selected. It is known when  $Y$  is a deterministic function of  $X$ , the IB curve cannot be explored and another Lagrangian has been proposed to tackle this problem: the squared IB Lagrangian:  $\mathcal{L}_{\text{sq-IB}}(T; \beta_{\text{sq}}) = I(T; Y) - \beta_{\text{sq}} I(X; T)^2$ . In this paper, we (i) present a general family of Lagrangians which allow for the exploration of the IB curve in all scenarios; (ii) provide the exact one-to-one mapping between the Lagrange multiplier and the desired compression rate  $r$  for known IB curve shapes; and (iii) show we can approximately obtain a specific compression level with the convex IB Lagrangian for both known and unknown IB curve shapes. This eliminates the burden of solving the optimization problem for many values of the Lagrange multiplier. That is, we prove that we can solve the original constrained problem with a single optimization.

**Keywords:** information bottleneck; representation learning; mutual information; optimization

## 1. Introduction

Let  $X \in \mathcal{X}$  and  $Y \in \mathcal{Y}$  be two statistically dependent random variables with joint distribution  $p_{(X,Y)}$ . The information bottleneck (IB) [1] investigates the problem of extracting the relevant information from  $X$  for the task of predicting  $Y$ .

For this purpose, the IB defines a bottleneck variable  $T \in \mathcal{T}$  obeying the Markov chain  $Y \leftrightarrow X \leftrightarrow T$  so that  $T$  acts as a representation of  $X$ . Tishby et al. [1] define the relevant information as the information the representation keeps from  $Y$  after the compression of  $X$  (i.e.,  $I(T; Y)$ ), provided a certain level of compression (i.e.,  $I(X; T) \leq r$ ). Therefore, we select the representation which yields the value of the IB curve that best fits our requirements.

**Definition 1 (IB Functional).** Let  $X$  and  $Y$  be statistically dependent variables. Let  $\Delta$  be the set of random variables  $T$  obeying the Markov condition  $Y \leftrightarrow X \leftrightarrow T$ . Then the IB functional is

$$F_{\text{IB,max}}(r) = \max_{T \in \Delta} \{I(T; Y)\} \text{ s.t. } I(X; T) \leq r, \forall r \in [0, \infty). \quad (1)$$

**Definition 2 (IB Curve).** The IB curve is the set of points defined by the solutions of  $F_{\text{IB,max}}(r)$  for varying values of  $r \in [0, \infty)$ .

**Definition 3 (Information Plane).** *The plane is defined by the axes  $I(T; Y)$  and  $I(X; T)$ .*

This method has been successfully applied to solve different problems from a variety of domains. For example:

- **Supervised learning.** In supervised learning, we are presented with a set of  $n$  pairs of input features and task outputs instances. We seek an approximation of the conditional probability distribution between the task outputs  $Y$  and the input features  $X$ . In classification tasks (i.e., when  $Y$  is a discrete random variable), the introduction of the variable  $T$  learned through the information bottleneck principle maintained the performance of standard algorithms based on the cross-entropy loss while providing with more adversarial attacks robustness and invariance to nuisances [2–4]. Moreover, by the nature of its definition the information bottleneck appears to be closely related with a trade-off between accuracy on the observable set and generalization to new, unseen instances (see Section 2).
- **Clustering.** In clustering, we are presented with a set of  $n$  pairs of instances of a random variable  $X$  and their attributes of interest  $Y$ . We seek groups of instances (or clusters  $T$ ) such that the attributes of interest within the instances of each cluster are similar and the attributes of interest of the instances of different clusters are dissimilar. Therefore, the information bottleneck can be employed since it allows us to aim for attribute representative clusters (maximizing the similarity between instances within the clusters) and enforce a certain compression of the random variable  $X$  (ensuring a certain difference between instances of the different clusters). This has been successfully implemented, for instance, for gene expression analysis and word, document, stock pricing, or movie rating clustering [5–7].
- **Image segmentation.** In image segmentation, we want to partition an image into segments such that each pixel in a region shares some attributes. If we divide the image into very small regions  $X$  (e.g., each region is a pixel or a set of pixels defined by a grid), we can consider the problem of segmentation as that of clustering the regions  $X$  based on the region attributes  $Y$ . Hence, we can use the information bottleneck so that we seek region clusters  $T$  that are maximally informative about the attributes  $Y$  (e.g., the intensity histogram bins) and maintain a level of compression of the original regions  $X$  [8].
- **Quantization.** In quantization, we consider a random variable  $X \in \mathcal{X}$  such that  $\mathcal{X}$  is a large or continuous set. Our objective is to map  $X$  into a variable  $T \in \mathcal{T}$  such that  $\mathcal{T}$  is a smaller, countable set. If we fix the quantization set size to  $|\mathcal{T}| = \lfloor r \rfloor$  and aim at maximizing the information of the quantized variable with another random variable  $Y$  and restrict the mapping to be deterministic, then the problem is equivalent to the information bottleneck [9,10].
- **Source coding.** In source coding, we consider a data source  $\mathcal{S}$  which generates a signal  $Y \in \mathcal{Y}$ , which is later perturbed by a channel  $\mathcal{C} : \mathcal{Y} \rightarrow \mathcal{X}$  that outputs  $X$ . We seek a coding scheme that generates a code  $T \in \mathcal{T}$  from the output of the channel  $X$  which is as informative as possible about the original source signal  $Y$  and can be transmitted at a small rate  $I(X; T) \leq r$ . Therefore, this problem is equivalent to the the formulation of the information bottleneck [11].

Furthermore, it has been employed as a tool for development or explanation in other disciplines like reinforcement learning [12–14], attribution methods [15], natural language processing [16], linguistics [17] or neuroscience [18]. Moreover, it has connections with other problems such as source coding with side information (or the Wyner-Ahlsvede-Körner (WAK) problem), the rate-distortion problem or the cost-capacity problem (see Sections 3, 6 and 7 from [19]).

In practice, solving a constrained optimization problem such as the IB functional is challenging. Thus, in order to avoid the non-linear constraints from the IB functional, the IB Lagrangian is defined.

**Definition 4 (IB Lagrangian).** *Let  $X$  and  $Y$  be statistically dependent variables. Let  $\Delta$  be the set of random variables  $T$  obeying the Markov condition  $Y \leftrightarrow X \leftrightarrow T$ . Then we define the IB Lagrangian as*

$$\mathcal{L}_{IB}(T; \beta) = I(T; Y) - \beta I(X; T). \quad (2)$$

Here  $\beta \in [0, 1]$  is the Lagrange multiplier which controls the trade-off between the information of  $Y$  retained and the compression of  $X$ . Note we consider  $\beta \in [0, 1]$  because (i) for  $\beta \leq 0$  many uncompressed solutions such as  $T = X$  maximize  $\mathcal{L}_{IB}(T; \beta)$ , and (ii) for  $\beta \geq 1$  the IB Lagrangian is non-positive due to the data processing inequality (DPI) (Theorem 2.8.1 from Cover and Thomas [20]) and trivial solutions like  $T = \text{const}$  are maximizers with  $\mathcal{L}_{IB}(T; \beta) = 0$  [21].

We know the solutions of the IB Lagrangian optimization (if existent) are solutions of the IB functional by the Lagrange's sufficiency theorem (Theorem 5 in Appendix A of Courcoubetis [22]). Moreover, since the IB functional is concave (Lemma 5 of Gilad-Bachrach et al. [19]) we know they exist (Theorem 6 in Appendix A of Courcoubetis [22]).

Therefore, the problem is usually solved by maximizing the IB Lagrangian with adaptations of the Blahut-Arimoto algorithm [1], deterministic annealing approaches [23] or a bottom-up greedy agglomerative clustering [6] or its improved sequential counterpart [24]. However, when provided with high-dimensional random variables  $X$  such as images, these algorithms do not scale well and deep learning-based techniques, where the IB Lagrangian is used as the objective function, prevailed [2,25,26].

Note the IB Lagrangian optimization yields a representation  $T$  with a given performance  $(I(X; T), I(T; Y))$  for a given  $\beta$ . However, there is no one-to-one mapping between  $\beta$  and  $I(X; T)$ . Hence, we cannot directly optimize for the desired compression level  $r$  but we need to perform several optimizations for different values of  $\beta$  and select the representation with the desired performance; e.g., [2]. The Lagrange multiplier selection is important since (i) sometimes even choices of  $\beta < 1$  lead to trivial representations such that  $p_{T|X} = p_T$ , and (ii) there exist some discontinuities on the performance level w.r.t. the values of  $\beta$  [27].

Moreover, recently Kolchinsky et al. [21] showed how in deterministic scenarios (such as many classification problems where an input  $x_i$  belongs to a single particular class  $y_i$ ) the IB Lagrangian could not explore the IB curve. Particularly, they showed that multiple  $\beta$  yielded the same performance level and that a single value of  $\beta$  could result in different performance levels. To solve this issue, they introduced the squared IB Lagrangian,  $\mathcal{L}_{sq-IB}(T; \beta_{sq}) = I(T; Y) - \beta_{sq} I(X; T)^2$ , which is able to explore the IB curve in any scenario by optimizing for different values of  $\beta_{sq}$ . However, even though they realized a one-to-one mapping between  $\beta_{sq}$  and the compression level existed, they did not find such mapping. Hence, multiple optimizations of the Lagrangian were still required to find the best trade-off solution.

The main contributions of this article are:

1. We introduce a general family of Lagrangians (the convex IB Lagrangians) which are able to explore the IB curve in any scenario for which the squared IB Lagrangian [21] is a particular case of. More importantly, the analysis made for deriving this family of Lagrangians can serve as inspiration for obtaining new Lagrangian families that solve other objective functions with intrinsic trade-offs such as the IB Lagrangian.
2. We show that in deterministic scenarios (and other scenarios where the IB curve shape is known) one can use the convex IB Lagrangian to obtain a desired level of performance with a single optimization. That is, there is a one-to-one mapping between the Lagrange multiplier used for the optimization and the level of compression and informativeness obtained, and we provide the exact mapping. This eliminates the need for multiple optimizations to select a suitable representation.
3. We introduce a particular case of the convex IB Lagrangians: the shifted exponential IB Lagrangian, which allows us to approximately obtain a specific compression level in any scenario. This way, we can approximately solve the initial constrained optimization problem from Equation (1) with a single optimization.

Furthermore, we provide some insight for explaining why there are discontinuities in the performance levels w.r.t. the values of the Lagrange multipliers. In a classification setting, we connect those discontinuities with the intrinsic clusterization of the representations when optimizing the IB bottleneck objective.

The structure of the article is the following: In Section 2 we motivate the usage of the IB in supervised learning settings. Then, in Section 3 we outline the important results used about the IB curve in deterministic scenarios. Later, in Section 4 we introduce the convex IB Lagrangian and explain some of its properties like the bijective mapping between Lagrange multipliers and the compression level and the range of such multipliers. After that, we support our (proved) claims with some empirical evidence on the MNIST [28] and TREC-6 [29] datasets in Section 5. Finally, in Section 6 we discuss our claims and empirical results. A PyTorch [30] implementation of the article can be found at <https://github.com/burklight/convex-IB-Lagrangian-PyTorch>.

In the Appendices A–F we provide with the proofs of the theoretical results. Then, in Appendix G we show some alternative families of Lagrangians with similar properties. Later, in Appendix H we provide with the precise experimental setup details to reproduce the results from the paper, and further experimentation with different datasets and neural network architectures. To conclude, in Appendix I we show some guidelines on how to set the convex information bottleneck Lagrangians for practical problems.

## 2. The IB in Supervised Learning

In this section, we will first give an overview of supervised learning in order to later motivate the usage of the information bottleneck in this setting.

### 2.1. Supervised Learning Overview

In supervised learning we are given a dataset  $\mathcal{D}_n = \{(x_i, y_i)\}_{i=1}^n$  of  $n$  pairs of input features and task outputs. In this case,  $X$  and  $Y$  are the random variables of the input features and the task outputs. We assume  $x_i$  and  $y_i$  are sampled i.i.d. from the true distribution  $p_{(X,Y)} = p_{Y|X}p_X$ . The usual aim of supervised learning is to use the dataset  $\mathcal{D}_n$  to learn a particular conditional distribution  $q_{\hat{Y}|X}$  of the task outputs given the input features, parametrized by  $\theta$ , which is a good approximation of  $p_{Y|X}$ . We use  $\hat{Y}$  and  $\hat{y}$  to indicate the predicted task output random variable and its outcome. We call a supervised learning task regression when  $Y$  is continuous-valued and classification when it is discrete.

Usually, supervised learning methods employ intermediate representations of the inputs before making predictions about the outputs; e.g., hidden layers in neural networks (Chapter 5 from Bishop [31]) or transformations in a feature space through the kernel trick in kernel machines like SVMs or RVMs (Sections 7.1 and 7.2 from Bishop [31]). Let  $T$  be a possibly stochastic function of the input features  $X$  with a parametrized conditional distribution  $q_{T|X}$ , then,  $T$  obeys the Markov condition  $Y \leftrightarrow X \leftrightarrow T$ . The mapping from the representation to the predicted task outputs is defined by the parametrized conditional distribution  $q_{\hat{Y}|T}$ . Therefore, in representation-based machine learning methods, the full Markov Chain is  $Y \leftrightarrow X \leftrightarrow T \leftrightarrow \hat{Y}$ . Hence, the overall estimation of the conditional probability  $p_{Y|X}$  is given by the marginalization of the representations; i.e.,  $q_{\hat{Y}|X} = \mathbb{E}_{t \sim q_{T|X}} [q_{\hat{Y}|T=t}]$  (The notation  $q_{\hat{Y}|T=t}$  represents the probability distribution  $q_{\hat{Y}|T}(\cdot|t; \theta)$ ). For the rest of the text, we will use the same notation to represent conditional probability distributions where the conditioning argument is given).

In order to achieve the goal of having a good estimation of the conditional probability distribution  $p_{Y|X}$ , we usually define an instantaneous cost function  $\mathbf{j} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ . The value of this function  $\mathbf{j}(x, y; \theta)$  serves as a heuristic to measure the loss of our algorithm, parametrized by  $\theta$ , obtains when trying to predict the realization of the task output  $y$  with the input realization  $x$ .

Clearly, we can be interested in minimizing the expectation of the instantaneous cost function over all the possible input features and task outputs, which we call the cost function. However, since we only have a finite dataset  $\mathcal{D}_n$  we have instead to minimize the empirical cost function.

**Definition 5 (Cost Function and Empirical Cost Function).** Let  $X$  and  $Y$  be the input features and task output random variables and  $x \in \mathcal{X}$  and  $y \in \mathcal{Y}$  their realizations. Let also  $j$  be the instantaneous cost function,  $\theta$  the parametrization of our learning algorithm, and  $\mathcal{D}_n = \{(x_i, y_i)\}_{i=1}^n$  the given dataset. Then, we define:

$$1. \text{ The cost function: } J(p_{(X,Y)}; \theta) = \mathbb{E}_{(x,y) \sim p_{(X,Y)}} [j(x, y; \theta)] \quad (3)$$

$$2. \text{ The empirical cost function: } \hat{J}(\mathcal{D}_n; \theta) = \frac{1}{n} \sum_{i=1}^n j(x_i, y_i; \theta) \quad (4)$$

The discrepancy between the normal and empirical cost functions is called the generalization gap or generalization error (see Section 1 of Xu and Raginsky [32], for instance) and intuitively, the smaller this gap is, the better our model generalizes; i.e., the better it will perform to new, unseen samples in terms of our cost function.

**Definition 6 (Generalization Gap).** Let  $J(p_{(X,Y)}; \theta)$  and  $\hat{J}(\mathcal{D}_n; \theta)$  be the cost and the empirical cost functions as defined in Definition 5. Then, the generalization gap is defined as

$$\text{gen}(\mathcal{D}_n; \theta) = J(p_{(X,Y)}; \theta) - \hat{J}(\mathcal{D}_n; \theta), \quad (5)$$

and it represents the error incurred when the selected distribution is the one parametrized by  $\theta$  when the rule  $\hat{J}(\mathcal{D}_n; \theta)$  is used instead of  $J(p_{(X,Y)}; \theta)$  as the function to minimize.

Ideally, we would want to minimize the cost function. Hence, we usually try to minimize the empirical cost function and the generalization gap simultaneously. The modifications to our learning algorithm which intend to reduce the generalization gap but not hurt the performance on the empirical cost function are known as regularization.

## 2.2. Why Do We Use the IB?

**Definition 7 (Representation cross-entropy cost function).** Let  $X$  and  $Y$  be two statistically dependent variables with joint distribution  $p_{(X,Y)} = p_{Y|X}p_X$ . Let also  $T$  be a random variable obeying the Markov condition  $Y \leftrightarrow X \leftrightarrow T$  and  $q_{T|X}$  and  $q_{\hat{Y}|T}$  be the encoding and decoding distributions of our model, parametrized by  $\theta$ . Finally, let  $\mathbb{C}(p_Z || q_Z) = -\mathbb{E}_{z \sim p_Z} [\log(q_Z(z))]$  be the cross entropy between two probability distributions  $p_Z$  and  $q_Z$ . Then, the cross-entropy cost function is

$$J_{\text{CE}}(p_{(X,Y)}; \theta) = \mathbb{E}_{(x,t) \sim q_{T|X}p_X} [\mathbb{C}(q_{Y|T=t} || q_{\hat{Y}|T=t})] = \mathbb{E}_{(x,y) \sim p_{(X,Y)}} [j_{\text{CE}}(x, y; \theta)], \quad (6)$$

where  $j_{\text{CE}}(x, y; \theta) = -\mathbb{E}_{t \sim q_{T|X=x}} [q_{\hat{Y}|T=t}(y|t; \theta)]$  is the instantaneous representation cross-entropy cost function and  $q_{Y|T} = \mathbb{E}_{x \sim p_X} [p_{Y|X=x}q_{T|X=x} / q_T]$  and  $q_T = \mathbb{E}_{x \sim p_X} [q_{T|X=x}]$ .

The cross-entropy is a widely used cost function in classification tasks (e.g., Teahan [8], Krizhevsky et al. [33], Shore and Gray [34]) which has many interesting properties [35]. Moreover, it is known that minimizing the  $J_{\text{CE}}(p_{(X,Y)}; \theta)$  maximizes the mutual information  $I(T; Y)$ . That is:

**Proposition 1 (Minimizing the Cross Entropy Maximizes the Mutual Information).** Let  $J_{\text{CE}}(p_{(X,Y)}; \theta)$  be the representation cross-entropy cost function as defined in Definition 7. Let also  $I(T; Y)$  be the mutual information between random variables  $T$  and  $Y$  in the setting from Definition 7. Then, minimizing  $J_{\text{CE}}(p_{(X,Y)}; \theta)$  implies maximizing  $I(T; Y)$ .

The proof of this proposition can be found in Appendix A.

**Definition 8 (Nuisance).** A nuisance is any random variable that affects the observed data  $X$  but is not informative to the task we are trying to solve. That is,  $\Xi$  is a nuisance for  $Y$  if  $Y \perp \Xi$  or  $I(\Xi, Y) = 0$ .

Similarly, we know that minimizing  $I(X; T)$  minimizes the generalization gap for restricted classes when using the cross-entropy cost function (Theorem 1 of Vera et al. [36]), and when using  $I(T; Y)$  directly as an objective to maximize (Theorem 4 of Shamir et al. [37]). Furthermore, Achille and Soatto [38] in Proposition 3.1 upper bound the information of the input representations,  $T$ , with nuisances that affect the observed data,  $\Xi$ , with  $I(X; T)$ . Therefore, minimizing  $I(X; T)$  helps generalization by not keeping useless information of  $\Xi$  in our representations.

Thus, jointly maximizing  $I(T; Y)$  and minimizing  $I(X; T)$  is a good choice both in terms of performance in the available dataset and in new, unseen data, which motivates studies on the IB.

### 3. The Information Bottleneck in Deterministic Scenarios

Kolchinsky et al. [21] showed that when  $Y$  is a deterministic function of  $X$  (i.e.,  $Y = f(X)$ ), the IB curve is piecewise linear. More precisely, it is shaped as stated in Proposition 2.

**Proposition 2 (The IB Curve is Piecewise Linear in Deterministic Scenarios).** Let  $X$  be a random variable and  $Y = f(X)$  be a deterministic function of  $X$ . Let also  $T$  be the bottleneck variable that solves the IB functional. Then the IB curve in the information plane is defined by the following equation:

$$\begin{cases} I(T; Y) = I(X; T) & \text{if } I(X; T) \in [0, I(X; Y)) \\ I(T; Y) = I(X; Y) & \text{if } I(X; T) \geq I(X; Y) \end{cases} \quad (7)$$

Furthermore, they showed that the IB curve could not be explored by optimizing the IB Lagrangian for multiple  $\beta$  because the curve was not strictly concave. That is, there was not a one-to-one relationship between  $\beta$  and the performance level.

#### Theorem 1 (In Deterministic Scenarios, the IB Curve cannot be Explored Using the IB Lagrangian).

Let  $X$  be a random variable and  $Y = f(X)$  be a deterministic function of  $X$ . Let also  $\Delta$  be the set of random variables  $T$  obeying the Markov condition  $Y \leftrightarrow X \leftrightarrow T$ . Then:

1. Any solution  $T \in \Delta$  such that  $I(X; T) \in [0, I(X; Y))$  and  $I(T; Y) = I(X; T)$  solves  $\arg \max_{T \in \Delta} \{\mathcal{L}_{\text{IB}}(T; \beta)\}$  for  $\beta = 1$ . That is, many different compression and performance levels can be achieved for  $\beta = 1$ .
2. Any solution  $T \in \Delta$  such that  $I(X; T) > I(X; Y)$  and  $I(T; Y) = I(X; Y)$  solves  $\arg \sup_{T \in \Delta} \{\mathcal{L}_{\text{IB}}(T; \beta)\}$  for  $\beta = 0$ . That is, many compression levels can be achieved with the same performance for  $\beta = 0$ .

Note we use the supremum in this case since for  $\beta = 0$  we have that  $I(X; T)$  could be infinite and then the search set from Equation (1); i.e.,  $\{T : Y \leftrightarrow X \leftrightarrow T\} \cap \{T : I(X; T) < \infty\}$  is not compact anymore.

3. Any solution  $T \in \Delta$  such that  $I(X; T) = I(T; Y) = I(X; Y)$  solves  $\arg \max_{T \in \Delta} \{\mathcal{L}_{\text{IB}}(T; \beta)\}$  for all  $\beta \in (0, 1)$ . That is, many different  $\beta$  achieve the same compression and performance level.

An alternative proof for this theorem can be found in Appendix B.

### 4. The Convex IB Lagrangian

#### 4.1. Exploring the IB Curve

Clearly, a situation like the one depicted in Theorem 1 is not desirable, since we cannot aim for different levels of compression or performance. For this reason, we generalize the effort from Kolchinsky et al. [21] and look for families of Lagrangians which are able to explore the IB curve. Inspired by the squared IB Lagrangian,  $\mathcal{L}_{\text{sq-IB}}(T; \beta_{\text{sq}}) = I(T; Y) - \beta_{\text{sq}} I(X; T)^2$ , we look at the

conditions a function of  $I(X; T)$  requires in order to be able to explore the IB curve. In this way, we realize that any monotonically increasing and strictly convex function will be able to do so, and we call the family of Lagrangians with these characteristics the convex IB Lagrangians, due to the nature of the introduced function.

**Theorem 2 (Convex IB Lagrangians).** *Let  $\Delta$  be the set of r.v.  $T$  obeying the Markov condition  $Y \leftrightarrow X \leftrightarrow T$ . Then, if  $u$  is a monotonically increasing and strictly convex function, the IB curve can always be recovered by the solutions of  $\arg \max_{T \in \Delta} \{\mathcal{L}_{IB,u}(T; \beta_u)\}$ , with*

$$\mathcal{L}_{IB,u}(T; \beta_u) = I(T; Y) - \beta_u u(I(X; T)). \quad (8)$$

*That is, for each point  $(I(X; T), I(T; Y))$  s.t.  $dI(T; Y)/dI(X; T) > 0$  there is a unique  $\beta_u$  for which maximizing  $\mathcal{L}_{IB,u}(T; \beta_u)$  achieves this solution. Furthermore,  $\beta_u$  is strictly decreasing w.r.t.  $I(X; T)$ . We call  $\mathcal{L}_{IB,u}(T; \beta_u)$  the convex IB Lagrangian.*

The proof of this theorem can be found in Appendix C. Furthermore, by exploiting the IB curve duality (Lemma 10 of Gilad-Bachrach et al. [19]) we were able to derive other families of Lagrangians which allow for the exploration of the IB curve (Appendix G).

**Remark 1.** *Clearly, we can see how if  $u$  is the identity function (i.e.,  $u(I(X; T)) = I(X; T)$ ) then we end up with the normal IB Lagrangian. However, since the identity function is not strictly convex, it cannot ensure the exploration of the IB curve.*

During the proof of this theorem we observed a relationship between the Lagrange multipliers and the solutions obtained of the normal IB Lagrangian  $\mathcal{L}_{IB}(T; \beta)$  and the convex IB Lagrangian  $\mathcal{L}_{IB,u}(T; \beta_u)$ . This relationship is formalized in the following corollary.

**Corollary 1 (IB Lagrangian and IB convex Lagrangian connection).** *Let  $\mathcal{L}_{IB}(T; \beta)$  be the IB Lagrangian and  $\mathcal{L}_{IB,u}(T; \beta_u)$  the convex IB Lagrangian. Then, maximizing  $\mathcal{L}_{IB}(T; \beta)$  and  $\mathcal{L}_{IB,u}(T; \beta_u)$  can obtain the same point in the IB curve if  $\beta_u = \beta / u'(I(X; T))$ , where  $u'$  is the derivative of  $u$ .*

This corollary allows us to better understand why the addition of  $u$  allows for the exploration of the IB curve in deterministic scenarios. If we note that for  $\beta = 1$  we can obtain any point in the increasing region of the curve, then we clearly see how evaluating  $u'$  for different values of  $I(X; T)$  define different values of  $\beta_u$  that obtain such points. Moreover, it lets us see how if for  $\beta = 0$  maximizing the IB Lagrangian could obtain any point  $(I(X; Y); I(X; T))$  with  $I(X; T) > I(X; Y)$ , then the same happens for the IB convex Lagrangian.

#### 4.2. Aiming for a Specific Compression Level

Let  $B_u$  denote the domain of Lagrange multipliers  $\beta_u$  for which we can find solutions in the IB curve with the convex IB Lagrangian. Then, the convex IB Lagrangians do not only allow us to explore the IB curve with different  $\beta_u$ . They also allow us to identify the specific  $\beta_u$  that obtains a given point  $(I(X; T), I(T; Y))$ , provided we know the IB curve in the information plane. Conversely, the convex IB Lagrangian allows finding the specific point  $(I(X; T), I(T; Y))$  that is obtained by a given  $\beta_u$ .

**Proposition 3 (Bijective Mapping between IB Curve Point and Convex IB Lagrange multiplier).** *Let the IB curve in the information plane be known; i.e.,  $I(T; Y) = f_{IB}(I(X; T))$  is known. Then there is a bijective mapping from Lagrange multipliers  $\beta_u \in B_u \setminus \{0\}$  from the convex IB Lagrangian to points in the IB curve  $(I(X; T), f_{IB}(I(X; T)))$ . Furthermore, these mappings are:*

$$\beta_u = \frac{df_{IB}(I(X; T))}{dI(X; T)} \frac{1}{u'(I(X; T))} \quad \text{and} \quad I(X; T) = (u')^{-1} \left( \frac{df_{IB}(I(X; T))}{dI(X; T)} \frac{1}{\beta_u} \right), \quad (9)$$

where  $u'$  is the derivative of  $u$  and  $(u')^{-1}$  is the inverse of  $u'$ .

This is especially interesting since in deterministic scenarios we know the shape of the IB curve (Theorem 2) and since the convex IB Lagrangians allow for the exploration of the IB curve (Theorem 2). A proof for Proposition 3 can be found in Appendix D.

**Remark 2.** Note that the definition from Tishby et al. [1]  $\beta = df_{\text{IB}}(I(X;T))/dI(X;T)$  only allows for a bijection between  $\beta$  and  $I(X;T)$  if  $f_{\text{IB}}$  is a strictly convex, and known function, and we have seen this is not the case in deterministic scenarios (Theorem 1).

A direct result derived from this proposition is that we know the domain of Lagrange multipliers,  $B_u$ , which allows for the exploration of the IB curve if the shape of the IB curve is known. Furthermore, if the shape is not known we can at least bound that range.

**Corollary 2 (Domain of Convex IB Lagrange Multiplier with Known IB Curve Shape).** Let the IB curve in the information plane be  $I(T;Y) = f_{\text{IB}}(I(X;T))$  and let  $I_{\text{max}} = I(X;Y)$ . Let also  $I(X;T) = r_{\text{max}}$  be the minimum mutual information s.t.  $f_{\text{IB}}(r_{\text{max}}) = I_{\text{max}}$ ; i.e.,  $r_{\text{max}} = \arg \inf_r \{f_{\text{IB}}(r)\}$  s.t.  $f_{\text{IB}}(r) = I_{\text{max}}$ . Then, the range of Lagrange multipliers that allow the exploration of the IB curve with the convex IB Lagrangian is  $B_u = [\beta_{u,\text{min}}, \beta_{u,\text{max}}]$ , with

$$\beta_{u,\text{min}} = \lim_{r \rightarrow r_{\text{max}}} \left\{ \frac{f'_{\text{IB}}(r)}{u'(r)} \right\} \quad \text{and} \quad \beta_{u,\text{max}} = \lim_{r \rightarrow 0^+} \left\{ \frac{f'_{\text{IB}}(r)}{u'(r)} \right\}, \quad (10)$$

where  $f'_{\text{IB}}(r)$  and  $u'(r)$  are the derivatives of  $f_{\text{IB}}(I(X;T))$  and  $u(I(X;T))$  w.r.t.  $I(X;T)$  evaluated at  $r$  respectively. Also, note that there are some scenarios where  $r_{\text{max}} \rightarrow \infty$  (see, e.g., [39]), in these scenarios  $\beta_{u,\text{min}} = \lim_{r \rightarrow \infty} \{f'_{\text{IB}}(r)/u'(r)\} \geq 0$ .

**Corollary 3 (Domain of Convex IB Lagrange Multiplier Bound).** The range of the Lagrange multipliers that allow the exploration of the IB curve is contained by  $[0, \beta_{u,\text{top}}]$  which is also contained by  $[0, \beta_{u,\text{top}}^+]$ , where

$$\beta_{u,\text{top}} = \frac{(\inf_{\Omega_x \subset \mathcal{X}} \{\beta_0(\Omega_x)\})^{-1}}{\lim_{r \rightarrow 0^+} \{u'(r)\}}, \quad \text{and} \quad \beta_{u,\text{top}}^+ = \frac{1}{\lim_{r \rightarrow 0^+} \{u'(r)\}}, \quad (11)$$

where  $u'(r)$  is the derivative of  $u(I(X;T))$  w.r.t.  $I(X;T)$  evaluated at  $r$ ,  $\mathcal{X}$  is the set of possible realizations of  $X$  and  $\beta_0$  and  $\Omega_x$  are defined as in [27] (Note in [27] they consider the dual problem (see Appendix G), so when they refer to  $\beta^{-1}$  it translates to  $\beta$  in this article). That is,  $B_u \subseteq [0, \beta_{u,\text{top}}] \subseteq [0, \beta_{u,\text{top}}^+]$ .

Corollaries 2 and 3 allow us to reduce the range search for  $\beta$  when we want to explore the IB curve. Practically,  $\inf_{\Omega_x \subset \mathcal{X}} \{\beta_0(\Omega_x)\}$  might be difficult to calculate so Wu et al. [27] derived an algorithm to approximate it. However, we still recommend setting the numerator to 1 for simplicity. The proofs for both corollaries are found in Appendices E and F.

## 5. Experimental Support

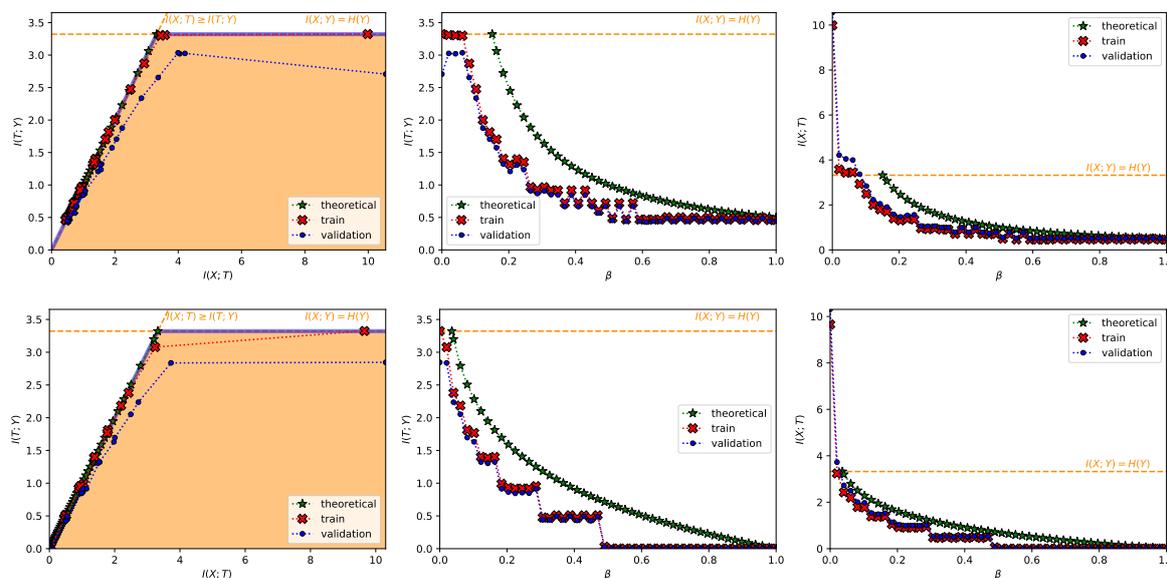
In order to showcase our claims we use the MNIST [28] and the TREC-6 [29] datasets. We modify the nonlinear-IB method [26], which is a neural network that minimizes the cross-entropy while also minimizing a differentiable kernel-based estimate of  $I(X;T)$  [40]. Then, we used this technique to maximize a lower bound on the convex IB Lagrangians by applying the functions  $u$  to the  $I(X;T)$  estimate.

The network structure is the following: first, a stochastic encoder  $T = f_{\text{enc}}(X; \theta) + W$  with  $p_W = \mathcal{N}(0, I_d)$  such that  $T \in \mathbb{R}^d$ , where  $d$  is the dimension of the bottleneck variable (Note that the encoder needs to be stochastic to (i) ensure a finite and well-defined mutual information [21,41] and (ii) make gradient-based optimization methods over the IB Lagrangian useful [41]). Second, a

deterministic decoder  $q_{\hat{Y}|T} = f_{\text{dec}}(T; \theta)$ . For the MNIST dataset both the encoder and the decoder are fully-connected networks, for a fair comparison with [26]. For the TREC-6 dataset, the encoder is a set of convolutions of word embeddings followed by a fully-connected network and the decoder is also a fully-connected network. For further details about the experiment setup, additional results for different values of  $\alpha$  and  $\eta$  and supplementary experimental results for different datasets and network architectures, please refer to Appendix H.

In Figure 1 we show our results for two particularizations of the convex IB Lagrangians:

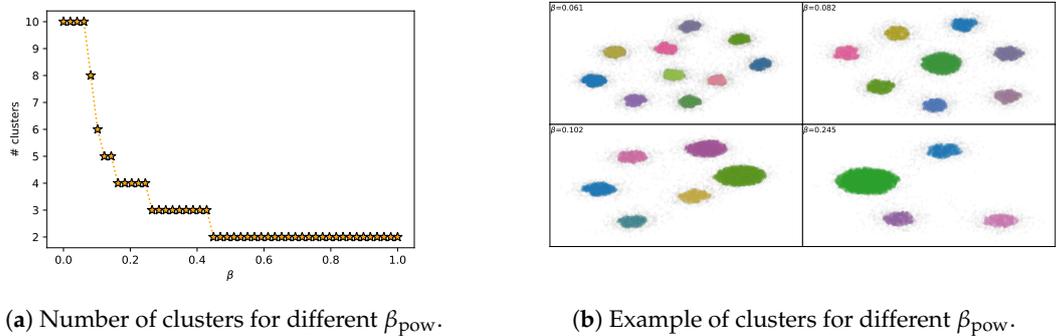
1. the power IB Lagrangians:  $\mathcal{L}_{\text{IB,pow}}(T; \beta_{\text{pow}}, \alpha) = I(T; Y) - \beta_{\text{pow}} I(X; T)^{1+\alpha}$ ,  $\alpha > 0$  (Note when  $\alpha = 1$  we have the squared IB functional from Kolchinsky et al. [21]).
2. the exponential IB Lagrangians:  $\mathcal{L}_{\text{IB,exp}}(T; \beta_{\text{exp}}, \eta) = I(T; Y) - \beta_{\text{exp}} \exp(\eta I(X; T))$ ,  $\eta > 0$ .



**Figure 1.** The top row shows the results for the power information bottleneck (IB) Lagrangian with  $\alpha = 1$ , and the bottom row for the exponential IB Lagrangian with  $\eta = 1$ , both in the MNIST dataset. In each row, from left to right it is shown (i) the information plane, where the region of possible solutions of the IB problem is shadowed in light orange and the information-theoretic limits are the dashed orange line; (ii)  $I(T; Y)$  as a function of  $\beta_u$ ; and (iii) the compression  $I(X; T)$  as a function of  $\beta_u$ . In all plots, the red crosses joined by a dotted line represent the values computed with the training set, the blue dots the values computed with the validation set and the green stars the theoretical values computed as dictated by Proposition 3. Moreover, in all plots, it is indicated  $I(X; Y) = H(Y) = \log_2(10)$  in a dashed, orange line. All values are shown in bits.

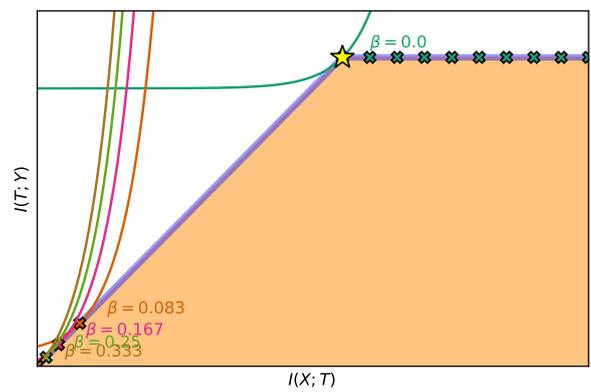
We can clearly see how both Lagrangians are able to explore the IB curve (first column from Figure 1) and how the theoretical performance trend of the Lagrangians matches the experimental results (second and third columns from Figure 1). There are small mismatches between the theoretical and experimental performance. This is because using the nonlinear-IB, as stated by Kolchinsky et al. [21], does not guarantee that we find optimal representations due to factors like (i) inaccurate estimation of  $I(X; T)$ , (ii) restrictions on the structure of  $T$ , (iii) use of an estimation of the decoder instead of the real one and (iv) the typical non-convex optimization issues that arise with gradient-based methods. The main difference comes from the discontinuities in performance for increasing  $\beta$ , which cause is still unknown (cf. Wu et al. [27]). It has been observed, however, that the bottleneck variable performs an intrinsic clusterization in classification tasks (see, for instance, [21,26,42] or Figure 2b). We observed how this clusterization matches with the quantized performance levels observed (e.g., compare Figure 2a with the top center graph in Figure 1); with

maximum performance when the number of clusters is equal to the cardinality of  $Y$  and reducing performance with a reduction of the number of clusters, which is in line with the concurrent work from Wu and Fischer [43]. We do not have a mathematical proof for the exact relationship between these two phenomena; however, we agree with Wu et al. [27] that it is an interesting matter and hope this observation serves as motivation to derive new theory.



**Figure 2.** Depiction of the clusterization behavior of the bottleneck variable for the power IB Lagrangian in the MNIST dataset with  $\alpha = 1$ . The clusters were obtained using the DBSCAN algorithm [44,45].

In practice, there are different criteria for choosing the function  $u$ . For instance, the exponential IB Lagrangian could be more desirable than the power IB Lagrangian when we want to draw the IB curve since it has a finite range of  $\beta_u$ . This is  $B_u = [(\eta \exp(\eta I_{\text{max}}))^{-1}, \eta^{-1}]$  for the exponential IB Lagrangian vs.  $B_u = [((1 + \alpha) I_{\text{max}}^\alpha)^{-1}, \infty)$  for the power IB Lagrangian. Furthermore, there is a trade-off between (i) how much the selected  $u$  function resembles a linear function in our region of interest; e.g., with  $\alpha$  or  $\eta$  close to zero, since it will suffer from similar problems as the original IB Lagrangian; and (ii) how fast it grows in our region of interest; e.g., higher values of  $\alpha$  or  $\eta$ , since it will suffer from value convergence; i.e., optimizing for separate values of  $\beta_u$  will achieve similar levels of performance (Figure 3). Please, refer to Appendix I for a more thorough explanation of these two phenomena.

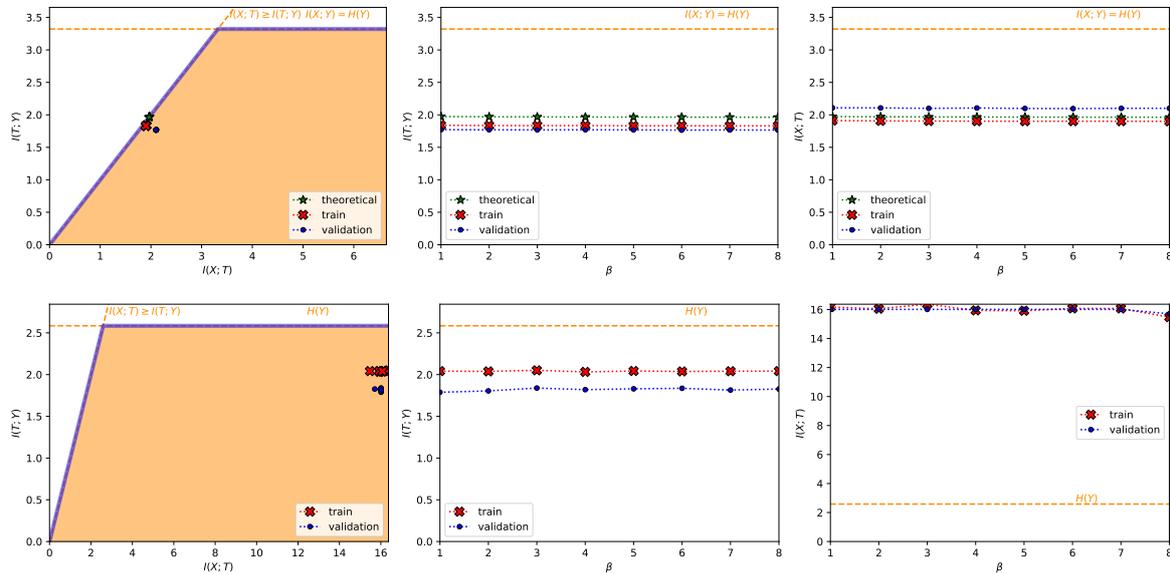


**Figure 3.** Example of value convergence with the exponential IB Lagrangian with  $\eta = 3$ . We show the intersection of the isolines of  $\mathcal{L}_{\text{IB,exp}}(T; \beta_{\text{exp}})$  for different  $\beta_{\text{exp}} \in B_{\text{exp}} \approx [1.56 \times 10^{-5}, 3^{-1}]$  using Corollary 2.

Particularly, the value convergence phenomenon can be exploited in order to approximately obtain a particular level of compression  $r^*$ , both for known and unknown IB curves (see Appendix I or the example in Figure 4). For known IB curves, we also know the achieved predictability  $I(T; Y)$  since it is the same as the level of compression  $I(X; T)$ . For this exploitation, we can employ the shifted version of the exponential IB Lagrangian (which is also a particular case of the convex IB Lagrangian):

- the shifted exponential IB Lagrangians:

$$\mathcal{L}_{IB,sh-exp}(T; \beta_{sh-exp}, \eta, r^*) = I(T; Y) - \beta_{sh-exp} \exp(\eta(I(X; T) - r^*)), \eta > 0, r^* \in [0, \infty).$$



**Figure 4.** Example of value convergence exploitation with the shifted exponential Lagrangian with  $\eta = 200$ . In the top row, for the MNIST dataset aiming for a compression level  $r^* = 2$  and in the bottom row, for the TREC-6 dataset aiming for a compression level of  $r^* = 16$ . In each row, from left to right it is shown (i) the information plane, where the region of possible solutions of the IB problem is shadowed in light orange and the information-theoretic limits are the dashed orange line; (ii)  $I(T; Y)$  as a function of  $\beta_u$ ; and (iii) the compression  $I(X; T)$  as a function of  $\beta_u$ . In all plots, the red crosses joined by a dotted line represent the values computed with the training set, the blue dots the values computed with the validation set and the green stars the theoretical values computed as dictated by Proposition 3. Moreover, in all plots, it is indicated  $H(Y)$  in a dashed, orange line. All values are shown in bits.

For this Lagrangian, the optimization procedure converges to representations with approximately the desired compression level  $r^*$  if the hyperparameter  $\eta$  is set to a large value.

In Figure 4 we show the results of aiming for a compression level of  $r^* = 2$  bits in the MNIST dataset and of  $r^* = 16$  bits in the TREC-6 dataset, both with  $\eta = 200$ . We can see how for different values of  $\beta_{sh-exp}$  we can obtain the same desired compression level, which makes this method stable to variations in the Lagrange multiplier selection.

To sum up, in order to achieve a desired level of performance with the convex IB Lagrangian as an objective one should:

- In a deterministic or close to a deterministic setting (see  $\epsilon$ -deterministic definition in Kolchinsky et al. [21]): Use the adequate  $\beta_u$  for that performance using Proposition 3. Then if the performance is lower than desired, i.e., we are placed in the wrong performance plateau, gradually reduce the value of  $\beta_u$  until reaching the previous performance plateau. Alternatively, exploit the value convergence phenomenon with, for instance, the shifted exponential IB Lagrangian.
- In a stochastic setting: exploit the value convergence phenomenon with, for instance, the shifted exponential IB Lagrangian. Alternatively, draw the IB curve with multiple values of  $\beta_u$  on the range defined by Corollary 3 and select the representations that best fit their interests.

## 6. Conclusions

The information bottleneck is a widely used and studied technique. However, it is known that the IB Lagrangian cannot be used to achieve varying levels of performance in deterministic scenarios. Moreover, in order to achieve a particular level of performance, multiple optimizations with different Lagrange multipliers must be done to draw the IB curve and select the best traded-off representation.

In this article we introduced a general family of Lagrangians which allow to (i) achieve varying levels of performance in any scenario, and (ii) pinpoint a specific Lagrange multiplier  $\beta_u$  to optimize for a specific performance level in known IB curve scenarios; e.g., deterministic. Furthermore, we showed the  $\beta_u$  domain when the IB curve is known and a  $\beta_u$  domain bound for exploring the IB curve when it is unknown. This way we can reduce and/or avoid multiple optimizations and, hence, reduce the computational effort for finding well traded-off representations. Moreover, (iii) when the IB curve is not known, we saw how we can exploit the value convergence issue of the convex IB Lagrangian to approximately obtain a specific compression level for both known and unknown IB curve shapes. Finally, (iv) we provided some insight into the discontinuities on the performance levels w.r.t. the Lagrange multipliers by connecting those with the intrinsic clusterization of the bottleneck variable.

**Author Contributions:** Conceptualization, B.R.G. and R.T.; formal analysis, B.R.G.; funding acquisition, M.S.; methodology, B.R.G. and R.T.; resources, M.S.; software, B.R.G.; supervision, R.T. and M.S.; visualization, B.R.G.; writing—original draft, B.R.G.; writing—review and editing, B.R.G., R.T. and M.S. All authors have read and agreed to the published version of the manuscript.

**Funding:** This work was supported in part by the Swedish Research Council.

**Acknowledgments:** We want to thank the anonymous reviewers for their insightful comments.

**Conflicts of Interest:** The authors declare no conflict of interest.

## Appendix A. Proof of Proposition 1

**Proof.** We can easily prove this statement by finding  $I(T; Y)$  is lower bounded by the  $\gamma J_{\text{CE}}(p_{(X,Y)}; \theta) + C$  where  $\gamma < 0$  and  $C$  does not depend on  $T$ . This way maximizing such lower bound would be equivalent to minimizing  $J_{\text{CE}}(p_{(X,Y)}; \theta)$  and, moreover, it would imply maximizing  $I(T; Y)$ .

We can find such an expression as follows:

$$I(T; Y) = \mathbb{E}_{(y,t) \sim q_{Y|T} q_T} \left[ \log \left( \frac{q_{Y|T=t}(y|t; \theta)}{p_Y(y)} \right) \right] = H(Y) + \mathbb{E}_{(y,t) \sim q_{Y|T} q_T} \left[ \log(q_{Y|T=t}(y|t; \theta)) \right] \quad (\text{A1})$$

$$= H(Y) + \mathbb{E}_{t \sim q_T} \left[ D_{\text{KL}} \left( q_{Y|T=t} \| q_{\hat{Y}|T=t} \right) \right] + \mathbb{E}_{(y,t) \sim q_{Y|T} q_T} \left[ \log(q_{\hat{Y}|T}(y|t; \theta)) \right] \quad (\text{A2})$$

$$\geq H(Y) + \mathbb{E}_{(x,y,t) \sim q_{Y|T} q_T p_X} \left[ \log(q_{\hat{Y}|T=t}(y|t; \theta)) \right] = H(Y) - \mathbb{E}_{(x,t) \sim q_{T|X} p_X} \left[ \mathbb{C}(q_{Y|T=t} \| q_{\hat{Y}|T=t}) \right] \quad (\text{A3})$$

$$= H(Y) - J_{\text{CE}}(p_{(X,Y)}; \theta). \quad (\text{A4})$$

Here, in Equation (A1) we just used the definition of the mutual information between two random variables, and then we decoupled it using the definition of the entropy of a variable (Note we used  $H(\cdot)$  which is usually employed for discrete variables. However, in this setting  $H(\cdot)$  could also refer to the differential entropy  $h(\cdot)$  of a continuous random variable since we employed the general definition using the expectation). Then, in Equation (A2) we only multiplied and divided by  $q_{\hat{Y}|T}$  inside the logarithm and employed the definition of the Kullback–Leibler divergence. Finally, in Equation (A3) we first used the fact the Kullback–Leibler divergence is always positive (Theorem 2.6.3 from Cover and Thomas [20]) and then the properties of the Markov chain  $T \leftrightarrow X \leftrightarrow Y$ .

Therefore, since  $H(Y)$  does not depend on  $T$  and we have a negative multiplicative term on  $J_{\text{CE}}(p_{(X,Y)}; \theta)$  the proposition is proved.  $\square$

### Appendix B. Alternative Proof of Theorem 1

**Proof.** We will prove all the enumerated statements sequentially, since the third one requires from the two first ones to be proved.

1. Proposition 2 states that the IB curve in the information plane follows the equation  $I(T; Y) = I(X; T)$  if  $I(X; T) \in [0, I(X; Y))$ . Then, since  $\beta = dI(T; Y)/dI(X; T)$  [1], we know  $\beta = 1$  in all these points. Therefore, for  $\beta = 1$  all points  $(I(X; T), I(X; T))$  such that  $I(X; T) \in [0, I(X; Y))$  are solutions of optimizing the IB Lagrangian.
2. Similarly, Proposition 2 states that the IB curve follows the equation  $I(T; Y) = I(X; Y)$  if  $I(X; T) \geq I(X; Y)$ . Then, since  $\beta = dI(T; Y)/dI(X; T)$  [1], we know  $\beta = 0$  in all points such that  $I(X; T) > I(X; Y)$ . We cannot ensure it at  $I(X; T) = I(X; Y)$  since  $\beta = 1$  for  $I(X; T) = \lim_{\epsilon \rightarrow 0^+} \{I(X; Y) - \epsilon\}$ .
3. Finally, in order to prove the last statement we will first prove that if  $\beta \in (0, 1)$  achieves a solution, it is  $(I(X; Y), I(X; Y))$ . Then, we will prove that if the solution  $(I(X; Y), I(X; Y))$  exists, this can be yield by any  $\beta \in (0, 1)$ . Hence, the solution  $(I(X; Y), I(X; Y))$  is achieved  $\forall \beta \in (0, 1)$  and it is the only solution achievable.
  - (a) Since the IB curve is concave we know  $\beta$  is non-increasing in  $I(X; T) \in \mathbb{R}^+$ . We also know  $\beta = 1$  at the points in the IB curve where  $I(X; T) \leq \lim_{\epsilon \rightarrow 0^+} \{I(X; Y) - \epsilon\}$  and  $\beta = 0$  at the points in the IB curve where  $I(X; T) \geq \lim_{\epsilon \rightarrow 0^+} \{I(X; Y) + \epsilon\}$ . Hence, if we achieve a solution with  $\beta \in (0, 1)$ , this solution is  $I(X; T) = I(T; Y) = I(X; Y)$ .
  - (b) We can upper bound the IB Lagrangian by

$$\mathcal{L}_{IB}(T; \beta) = I(T; Y) - \beta I(X; T) \leq (1 - \beta)I(T; Y) \leq (1 - \beta)I(X; Y), \tag{A5}$$

where the first and second inequalities use the DPI (Theorem 2.8.1 from Cover and Thomas [20]).

Then, we can consider the point of the IB curve  $(I(X; Y), I(X; Y))$ . Since the function is concave a tangent line to  $(I(X; Y), I(X; Y))$  exists such that all other points in the curve lie below this line. Let  $\beta$  be the slope of this curve (which we know it is from Tishby et al. [1]). Then,

$$I(X; Y) - \beta I(X; Y) = (1 - \beta)I(X; Y) \geq F_{IB, \max}(r) - \beta r, \quad \forall r \in [0, \infty). \tag{A6}$$

As we see, by the upper bound on the IB Lagrangian from Equation (A5), if the point  $(I(X; Y), I(X; Y))$  exists, any  $\beta$  can be the slope of the tangent line to  $(I(X; Y), I(X; Y))$  that ensures concavity.

□

### Appendix C. Proof of Theorem 2

**Proof.** We start the proof by remembering the optimization problem at hand (Definition 1):

$$F_{IB, \max}(r) = \max_{T \in \Delta} \{I(T; Y)\} \text{ s.t. } I(X; T) \leq r \tag{A7}$$

We can modify the optimization problem by

$$\max_{T \in \Delta} \{I(T; Y)\} \text{ s.t. } u(I(X; T)) \leq u(r) \tag{A8}$$

iff  $u$  is a monotonically non-decreasing function since otherwise  $u(I(X; T)) \leq u(r)$  would not hold necessarily. Now, let us assume  $\exists T^* \in \Delta$  and  $\beta_u^*$  s.t.  $T^*$  maximizes  $\mathcal{L}_{IB, u}(T; \beta_u^*)$  over all  $T \in \Delta$ , and  $I(X; T^*) \leq r$ . Then, we can operate as follows:

$$\max_{\substack{T \in \Delta \\ u(I(X;T)) \leq u(r)}} \{I(T; Y)\} = \max_{\substack{T \in \Delta \\ u(I(X;T)) \leq u(r)}} \{I(T; Y) - \beta_u^*(u(I(X; T)) - u(r) + \xi)\} \tag{A9}$$

$$\leq \max_{T \in \Delta} \{I(T; Y) - \beta_u^*(u(I(X; T)) - u(r) + \xi)\} \tag{A10}$$

$$= I(T^*; Y) - \beta_u^*(u(I(X; T^*) - u(r) + \xi) = I(T^*; Y). \tag{A11}$$

Here, the equality from Equation (A9) comes from the fact that since  $I(X; T) \leq r$ , then  $\exists \xi \geq 0$  s.t.  $u(I(X; T)) - u(r) + \xi = 0$ . Then, the inequality from Equation (A10) holds since we have expanded the optimization search space. Finally, in Equation (A11) we use that  $T^*$  maximizes  $\mathcal{L}_{IB,u}(T; \beta_u^*)$  and that  $I(X; T^*) \leq r$ .

Now, we can exploit that  $u(r)$  and  $\xi$  do not depend on  $T$  and drop them in the maximization in Equation (A10). We can then realize we are maximizing over  $\mathcal{L}_{IB,u}(T; \beta_u^*)$ ; i.e.,

$$\max_{\substack{T \in \Delta \\ u(I(X;T)) \leq u(r)}} \{I(T; Y)\} \leq \max_{T \in \Delta} \{I(T; Y) - \beta_u^*(u(I(X; T)) - u(r) + \xi)\} \tag{A12}$$

$$= \max_{T \in \Delta} \{I(T; Y) - \beta_u^*(I(X; T))\} = \max_{T \in \Delta} \{\mathcal{L}_{IB,u}(T; \beta_u^*)\}. \tag{A13}$$

Therefore, since  $I(T^*; Y)$  satisfies both the maximization with  $T^* \in \Delta$  and the constraint  $I(X; T^*) \leq r$ , maximizing  $\mathcal{L}_{IB,u}(T; \beta_u^*)$  obtains  $F_{IB,max}(r)$ .

Now, we know if such  $\beta_u^*$  exists, then the solution of the Lagrangian will be a solution for  $F_{IB,max}(r)$ . Then, if we consider Theorem 6 from the Appendix of Courcoubetis [22] and consider the maximization problem instead of the minimization problem, we know if both  $I(T; Y)$  and  $-u(I(X; T))$  are concave functions, then a set of Lagrange multipliers  $S_u^*$  exists with these conditions. We can make this consideration because  $f$  is concave if  $-f$  is convex and  $\max\{f\} = \min\{-f\}$ . We know  $I(T; Y)$  is a concave function of  $T$  for  $T \in \Delta$  (Lemma 5 of Gilad-Bachrach et al. [19]) and  $I(X; T)$  is convex w.r.t.  $T$  given  $p_X$  is fixed (Theorem 2.7.4 of Cover and Thomas [20]). Thus, if we want  $-u(I(X; T))$  to be concave we need  $u$  to be a convex function.

Finally, we will look at the conditions of  $u$  so that for every point  $(I(X; T), I(T; Y))$  in the IB curve, there exists a unique  $\beta_u^*$  s.t.  $\mathcal{L}_{IB,u}(T; \beta_u^*)$  is maximized. That is, the conditions of  $u$  s.t.  $|S_u^*| = 1$ . For this purpose we will look at the solutions of the Lagrangian optimization:

$$\frac{d\mathcal{L}_{IB,u}(T; \beta_u)}{dT} = \frac{d(I(T; Y) - \beta_u u(I(X; T)))}{dT} = \frac{dI(T; Y)}{dT} - \beta_u \frac{du(I(X; T))}{dI(X; T)} \frac{dI(X; T)}{dT} = 0 \tag{A14}$$

Now, if we integrate both sides of Equation (A14) over all  $T \in \Delta$  we obtain

$$\beta_u = \frac{dI(T; Y)}{dI(X; T)} \left( \frac{du(I(X; T))}{dI(X; T)} \right)^{-1} = \frac{\beta}{u'(I(X; T))}, \tag{A15}$$

where  $\beta$  is the Lagrange multiplier from the IB Lagrangian [1] and  $u'(I(X; T))$  is  $\frac{du(I(X; T))}{dI(X; T)}$ . Also, if we want to avoid indeterminations of  $\beta_u$  we need  $u'(I(X; T))$  not to be 0. Since we already imposed  $u$  to be monotonically non-decreasing, we can solve this issue by strengthening this condition. That is, we will require  $u$  to be monotonically increasing.

We would like  $\beta_u$  to be continuous, this way there would be a unique  $\beta_u$  for each value of  $I(X; T)$ . We know  $\beta$  is a non-increasing function of  $I(X; T)$  (Lemma 6 of Gilad-Bachrach et al. [19]). Hence, if we want  $\beta_u$  to be a strictly decreasing function of  $I(X; T)$ , we will require  $u'$  to be a strictly increasing function of  $I(X; T)$ . Therefore, we will require  $u$  to be a strictly convex function.

Thus, if  $u$  is a strictly convex and monotonically increasing function, for each point  $(I(X;T), I(T;Y))$  in the IB curve s.t.  $dI(T;Y)/dI(X;T) > 0$  there is a unique  $\beta_u$  for which maximizing  $\mathcal{L}_{IB,u}(T; \beta_u)$  achieves this solution.  $\square$

### Appendix D. Proof of Proposition 3

**Proof.** In Theorem 2 we showed how each point of the IB curve  $(I(X;T), I(T;Y))$  can be found with a unique  $\beta_u$  maximizing  $\mathcal{L}_{IB,u}(T; \beta_u)$ . Therefore, since we also proved  $\mathcal{L}_{IB,u}(T; \beta_u)$  is strictly concave w.r.t.  $T$  we can find the values of  $\beta_u$  that maximize the Lagrangian for fixed  $I(X;T)$ .

First, we look at the solutions of the Lagrangian maximization:

$$\frac{d\mathcal{L}_{IB,u}(T; \beta_u)}{dT} = \frac{d(f_{IB}(I(X;T)) - \beta_u u(I(X;T)))}{dT} = \frac{df_{IB}(I(X;T))}{dT} - \beta_u \frac{du(I(X;T))}{dI(X;T)} \frac{dI(X;T)}{dT} = 0. \tag{A16}$$

Then as before we can integrate at both sides for all  $T \in \Delta$  and solve for  $\beta_u$ :

$$\beta_u = \frac{df_{IB}(I(X;T))}{dI(X;T)} \frac{1}{u'(I(X;T))}. \tag{A17}$$

Moreover, since  $u$  is a strictly convex function it's derivative  $u'$  is strictly increasing. Hence,  $u'$  is an invertible function (since a strictly increasing function is bijective and a function is invertible iff it is bijective by definition). Now, if we consider  $\beta_u > 0$  to be known and  $I(X;T)$  to be the unknown we can solve for  $I(X;T)$  and get:

$$I(X;T) = (u')^{-1} \left( \frac{df_{IB}(I(X;T))}{dI(X;T)} \frac{1}{\beta_u} \right). \tag{A18}$$

Note we require  $\beta_u$  not to be 0 so the mapping is defined.  $\square$

### Appendix E. Proof of Corollary 2

**Proof.** We will start the proof by proving the following useful Lemma.

**Lemma A1.** Let  $\mathcal{L}_{IB,u}(T; \beta_u)$  be a convex IB Lagrangian, then  $\sup_{T \in \Delta} \{\mathcal{L}_{IB,u}(T; 0)\} = I(X;Y)$ .

**Proof.** Since  $\mathcal{L}_{IB,u}(T; 0) = I(T;Y)$ , maximizing this Lagrangian is directly maximizing  $I(T;Y)$ . We know  $I(T;Y)$  is a concave function of  $T$  for  $T \in \Delta$  (Theorem 2.7.4 from Cover and Thomas [20]); hence it has a supremum. We also know  $I(T;Y) \leq I(X;Y)$ . Moreover, we know  $I(X;Y)$  can be achieved if, for example,  $Y$  is a deterministic function of  $T$  (since then the Markov Chain  $X \leftrightarrow T \leftrightarrow Y$  is formed). Thus,  $\sup_{T \in \Delta} \{\mathcal{L}_{IB,u}(T; 0)\} = I(X;Y)$ .  $\square$

For  $\beta_u = 0$  we know maximizing  $\mathcal{L}_{IB,u}(T; 0)$  we can obtain the point in the IB curve  $(r_{\max}, I_{\max})$  (Lemma A1). Moreover, we know that for every point  $(I(X;T), f_{IB}(I(X;T)))$  such that  $df_{IB}(I(X;T))/dI(X;T) > 0$ ,  $\exists! \beta_u$  s.t.  $\max\{\mathcal{L}_{IB,u}(T; \beta_u)\}$  achieves that point (Theorem 2). Thus,  $\exists! \beta_{u,\min}$  s.t.  $\lim_{r \rightarrow r_{\max}^-} (r, f_{IB}(r))$  is achieved. From Proposition 3 we know this  $\beta_{u,\min}$  is given by

$$\beta_{u,\min} = \lim_{r \rightarrow r_{\max}^-} \left\{ \frac{f'_{IB}(r)}{u'(r)} \right\}. \tag{A19}$$

Since we know  $f_{IB}(I(X;T))$  is a concave non-decreasing function in  $(0, r_{\max})$  (Lemma 5 of Gilad-Bachrach et al. [19]) we know it is continuous in this interval. In addition we know  $\beta_u$  is strictly decreasing w.r.t.  $I(X;T)$  (Theorem 2). Furthermore, by definition of  $r_{\max}$  and knowing  $I(T;Y) \leq I(X;Y)$  we know  $f'_{IB}(r) = 0, \forall r > r_{\max}$ . Therefore, we cannot ensure the exploration of the IB curve for  $\beta'_u$  s.t.  $0 < \beta'_u < \beta_{u,\min}$ .

Then, since  $u$  is a strictly increasing function in  $(0, r_{\max})$ ,  $u'$  is positive in that interval. Hence, taking into account  $\beta_u$  is strictly decreasing we can find a maximum  $\beta_u$  when  $I(X;T)$  approaches to 0. That is,

$$\beta_{u,\max} = \lim_{r \rightarrow 0^+} \left\{ \frac{f'_{\text{IB}}(r)}{u'(r)} \right\}, \tag{A20}$$

□

### Appendix F. Proof of Corollary 3

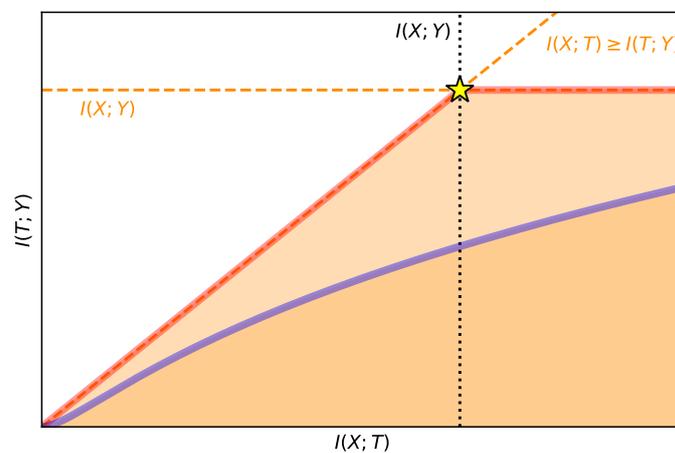
**Proof.** If we use Corollary 2, it is straightforward to see that  $\beta_u \subseteq [L_-, L_+]$  if  $\beta_{u,\min} \geq L_-$  and  $\beta_{u,\max} \leq L_+$  for all IB curves  $f_{\text{IB}}$  and functions  $u$ . Therefore, we look at a domain bound dependent on the function choice. That is, if we can find  $\beta_{\min} \leq f'_{\text{IB}}(r)$  and  $\beta_{\max} \geq f'_{\text{IB}}(r)$  for all IB curves and all values of  $r$ , then

$$B_u \subseteq \left[ \frac{\beta_{\min}}{\lim_{r \rightarrow r_{\max}^-} \{u'(r)\}}, \frac{\beta_{\max}}{\lim_{r \rightarrow 0^+} \{u'(r)\}} \right]. \tag{A21}$$

The region for all possible IB curves regardless of the relationship between  $X$  and  $Y$  is depicted in Figure A1. The hard limits are imposed by the DPI (Theorem 2.8.1 from Cover and Thomas [20]) and the fact that the mutual information is non-negative (Corollary with Equation 2.90 for discrete and first Corollary of Theorem 8.6.1 for continuous random variables from Cover and Thomas [20]). Hence, a minimum and maximum values of  $f'_{\text{IB}}$  are given by the minimum and maximum values of the slope of the Pareto frontier. Which means

$$B_u \subseteq \left[ 0, \frac{1}{\lim_{r \rightarrow 0^+} \{u'(r)\}} \right]. \tag{A22}$$

Note  $0 / (\lim_{r \rightarrow r_{\max}^-} \{u'(r)\}) = 0$  since  $u$  is monotonically increasing and, thus,  $u'$  will never be 0.



**Figure A1.** Graphical representation of the IB curve in the information plane. Dashed lines in orange represent tight bounds confining the region (in light orange) of possible IB curves (delimited by the red line, also known as the Pareto frontier). Black dotted lines are informative values. In blue we show an example of a possible IB curve confining a region (in darker orange) of an IB curve that does not achieve the Pareto frontier. Finally, the yellow star represents the point where the representation keeps the same information about the input and the output.

Then, we can tighten the bound using the results from Wu et al. [27], where, in Theorem 2, they showed the slope of the Pareto frontier could be bounded in the origin by  $f'_{\text{IB}} \leq (\inf_{\Omega_x \subset \mathcal{X}} \{\beta_0(\Omega_x)\})^{-1}$ . Finally, we know that in deterministic classification tasks  $\inf_{\Omega_x \subset \mathcal{X}} \{\beta_0(\Omega_x)\} = 1$ , which aligns with Kolchinsky et al. [21] and what we can observe from Figure A1. Therefore,

$$B_u \subseteq \left[ 0, \frac{(\inf_{\Omega_x \subset \mathcal{X}} \{\beta_0(\Omega_x)\})^{-1}}{\lim_{r \rightarrow 0^+} \{u'(r)\}} \right] \subseteq \left[ 0, \frac{1}{\lim_{r \rightarrow 0^+} \{u'(r)\}} \right]. \tag{A23}$$

□

### Appendix G. Other Lagrangian Families

We can use the same ideas we used for the convex IB Lagrangian to formulate new families of Lagrangians that allow the exploration of the IB curve. For that, we will use the duality of the IB curve (Lemma 10 of [19]). That is:

**Definition A1 (IB Dual Functional).** Let  $X$  and  $Y$  be statistically dependent variables. Let also  $\Delta$  be the set of random variables  $T$  obeying the Markov condition  $Y \leftrightarrow X \leftrightarrow T$ . Then the IB dual functional is

$$F_{IB,\min}(i) = \min_{T \in \Delta} \{I(X; T)\} \text{ s.t. } I(T; Y) \geq i, \forall i \in [0, I(X; Y)]. \tag{A24}$$

**Theorem A1 (IB Curve Duality).** Let the IB curve be defined by the solutions of  $F_{IB,\max}(r)$  for varying  $r \in [0, \infty)$ . Then,

$$\forall r \exists i \text{ s.t. } (r, F_{IB,\max}(r)) = (F_{IB,\min}(i), i) \tag{A25}$$

and

$$\forall i \exists r \text{ s.t. } (F_{IB,\min}(i), i) = (r, F_{IB,\max}(r)). \tag{A26}$$

From this definition, it follows that minimizing the dual IB Lagrangian,  $\mathcal{L}_{IB,\text{dual}}(T; \beta_{\text{dual}}) = I(X; T) - \beta_{\text{dual}}I(T; Y)$ , for  $\beta_{\text{dual}} = \beta^{-1}$  is equivalent to maximizing the IB Lagrangian. In fact, the original Lagrangian for solving the problem was defined this way [1]. We decided to use the maximization version because the domain of useful  $\beta$  is bounded while it is not for  $\beta_{\text{dual}}$ .

Following the same reasoning as we did in the proof of Theorem 2, we can ensure the IB curve can be explored if:

1. We minimize the concave IB Lagrangian  $\mathcal{L}_{IB,v}(T; \beta_v) = I(X; T) - \beta_v v(I(T; Y))$ .
2. We maximize the dual concave IB Lagrangian  $\mathcal{L}_{IB,v,\text{dual}}(T; \beta_{v,\text{dual}}) = v(I(T; Y)) - \beta_{v,\text{dual}}I(X; T)$ .
3. We minimize the dual convex IB Lagrangian  $\mathcal{L}_{IB,u,\text{dual}}(T; \beta_{u,\text{dual}}) = u(I(X; T)) - \beta_{u,\text{dual}}I(T; Y)$ .

Here,  $u$  is a monotonically increasing strictly convex function,  $v$  is a monotonically increasing strictly concave function, and  $\beta_v, \beta_{v,\text{dual}}, \beta_{u,\text{dual}}$  are the Lagrange multipliers of the families of Lagrangians defined above.

In a similar manner, one could obtain relationships between the Lagrange multipliers of the IB Lagrangian and the convex IB Lagrangian with these Lagrangian families. For instance, the convex IB Lagrangian  $\mathcal{L}_{IB,u}(T; \beta_u)$  is related with the concave IB Lagrangian  $\mathcal{L}_{IB,v}(T; \beta_v)$  as defined by Proposition A1.

**Proposition A1 (Relationship between the convex and concave IB Lagrangians).** Consider the convex and concave IB Lagrangians  $\mathcal{L}_{IB,u}(T; \beta_u)$ ,  $\mathcal{L}_{IB,v}(T; \beta_v)$ . Let the IB curve defined as in Definition 2 be  $f_{IB}$ . Then, if we fix the functions  $u$  and  $v$  we can obtain the same point in the IB curve  $(r, f_{IB}(r))$  with both Lagrangians when

$$\beta_v^{-1} = f'_{IB}(r) v' \left( f_{IB} \left( (u')^{-1} \left( \frac{f'_{IB}(r)}{\beta_u} \right) \right) \right), \tag{A27}$$

or equivalently,

$$\beta_u^{-1} = \frac{1}{f'_{IB}(r)} u' \left( f_{IB}^{-1} \left( (v')^{-1} \left( \frac{\beta_v^{-1}}{f'_{IB}(r)} \right) \right) \right). \tag{A28}$$

**Proof.** If we proceed like we did in the proof of Proposition 3 we can find the mapping between  $I(X;T)$  and  $\beta_u$  and between  $I(T;Y)$  and  $\beta_v$ . That is,

$$I(X;T) = (u')^{-1} \left( \frac{df_{\text{IB}}(I(X;T))}{dI(X;T)} \frac{1}{\beta_u} \right) \text{ and } I(T;Y) = (v')^{-1} \left( \left( \frac{df_{\text{IB}}(I(X;T))}{dI(X;T)} \right)^{-1} \frac{1}{\beta_v} \right). \quad (\text{A29})$$

Then, if we recall that  $I(T;Y) = f_{\text{IB}}(I(X;T))$ , we can directly obtain that

$$f_{\text{IB}} \left( (u')^{-1} \left( \frac{df_{\text{IB}}(I(X;T))}{dI(X;T)} \frac{1}{\beta_u} \right) \right) = (v')^{-1} \left( \left( \frac{df_{\text{IB}}(I(X;T))}{dI(X;T)} \right)^{-1} \frac{1}{\beta_v} \right). \quad (\text{A30})$$

Then, if we solve Equation (A30) with a fixed point ( $I(X;T) = r, I(T;Y) = f_{\text{IB}}(r)$ ) for  $\beta_v$  we obtain Equation (A27), and if we solve it for  $\beta_u$  we obtain Equation (A28).  $\square$

Also, one could find a range of values for these Lagrangians to allow for the IB curve exploration and define a bijective mapping between their Lagrange multipliers and the IB curve. However, (i) as mentioned in Section 2.2,  $I(T;Y)$  is particularly interesting to maximize without transformations because of its meaning. Moreover, (ii) like  $\beta_{\text{dual}}$ , the domain of useful  $\beta_v$  and  $\beta_{u,\text{dual}}$  is not upper bounded. These two reasons make these other Lagrangians less preferable. We only include them here for completeness. Nonetheless, we encourage the curious reader to explore these families of Lagrangians too. For example, a possible interesting research would be investigating if some particularization of the concave IB Lagrangian suffers from an issue like value convergence that can be exploited for approximately obtaining any predictability level  $I(T;Y) = i^*$  for many values of  $\beta_v$ .

## Appendix H. Experimental Setup Details and Further Experiments

In order to generate empirical support for our claims, we performed several experiments on different datasets with different neural network architectures and different ways of calculating the information bottleneck.

### Appendix H.1. Information Bottleneck Calculations

The information bottleneck is calculated modifying either the nonlinear-IB [26]. This method of calculating the information bottleneck is a neural network that minimizes the cross-entropy while also minimizing an upper bound estimate of the mutual information  $I_\theta \approx I(X;T)$ . The nonlinear-IB relies on a kernel-based estimate of this mutual information [40]. We modify this calculation method by applying the function  $u$  to the  $I(X;T)$  estimate.

For the nonlinear-IB calculations, we estimated the gradients of both  $I_\theta(X;T)$  and the cross-entropy with the same mini-batch. Moreover, we did not learn the covariance of the mixture of Gaussians used for the kernel density estimation of  $I_\theta(X;T)$  and we set it to  $(\exp(-1))^2$ .

In both methods, and for all the experiments, we assumed a Gaussian stochastic encoder  $T = f_{\text{enc}}(X; \theta) + W$  with  $p_W = \mathcal{N}(0, I_d)$ , where  $d$  are the number of dimensions of the representations. We trained the neural networks with the Adam optimization algorithm [46] with a learning rate of  $10^{-4}$  and a 0.6 decay rate every 10 epochs. We used a batch size of 128 samples and all the weights were initialized according to the method described by Glorot and Bengio [47] using a Gaussian distribution.

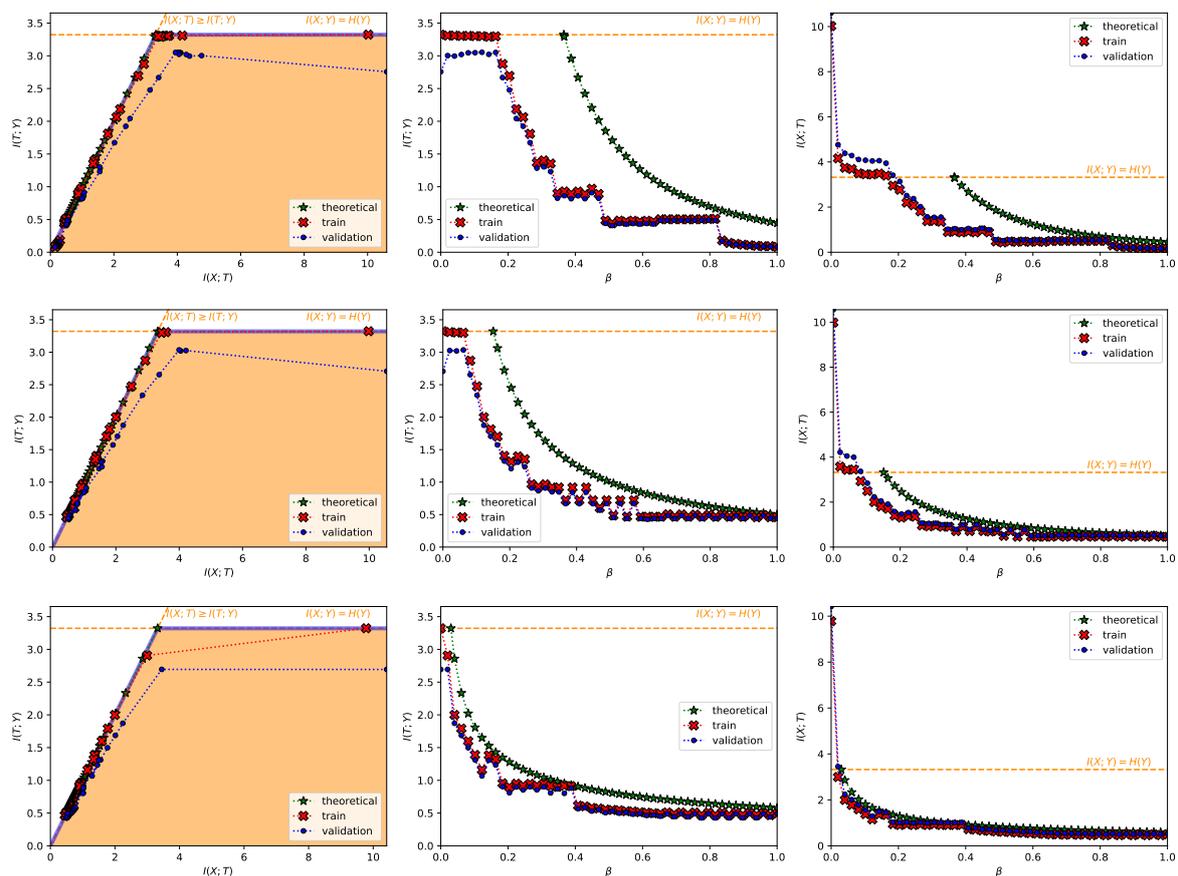
Then, we used the DBSCAN algorithm [44,45] for clustering. Particularly, we used the scikit-learn [48] implementation with  $\epsilon = 0.3$  and  $\text{min\_samples} = 50$ .

The reader can find the PyTorch [30] implementation in the following link: <https://github.com/burklight/convex-IB-Lagrangian-PyTorch>.

Appendix H.2. The Experiments

We performed experiments in four different datasets:

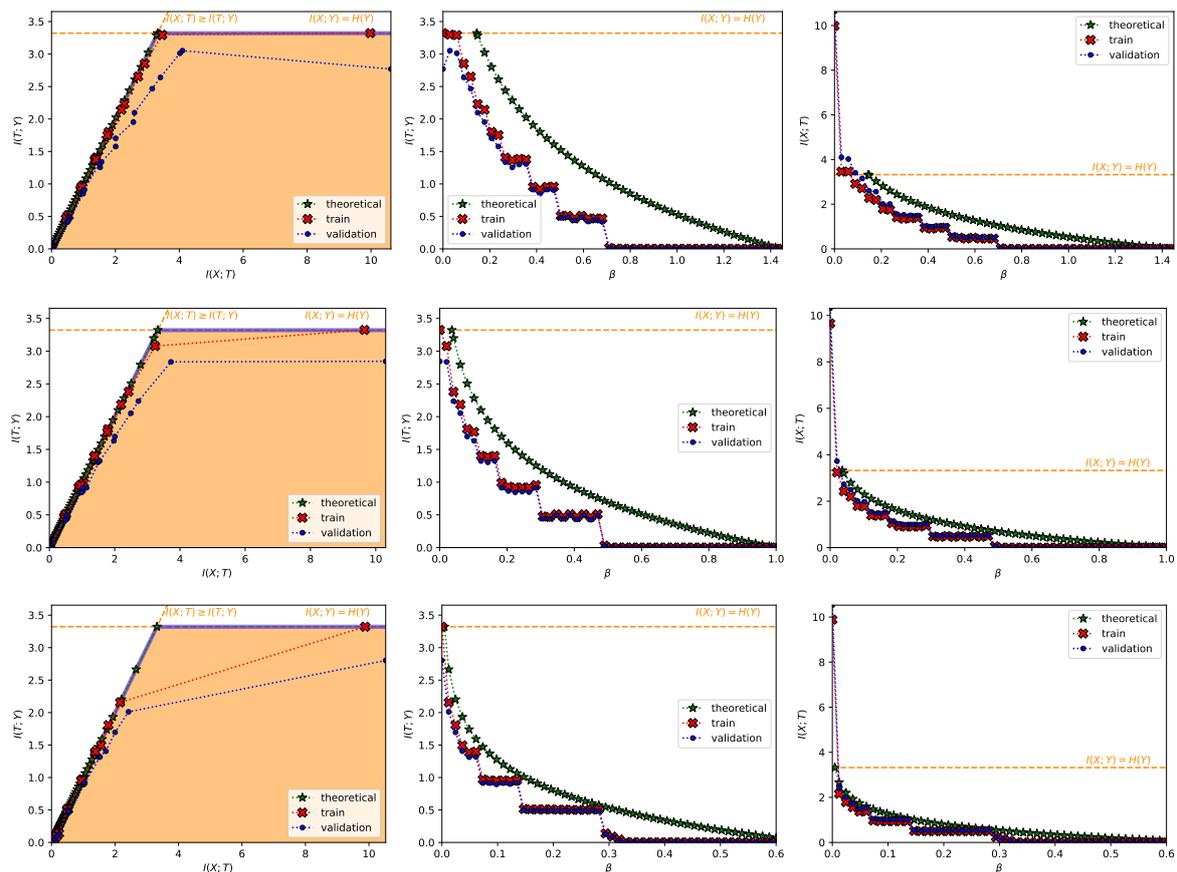
- A Classification Task on the MNIST Dataset [28] (Figures 1, 2, and A2–A4 and top row from Figure 3). This dataset contains 60,000 training samples and 10,000 testing samples of hand-written digits. The samples are 28x28 pixels and are labeled from 0 to 9; i.e.,  $\mathcal{X} = \mathbb{R}^{784}$  and  $\mathcal{Y} = \{0, 1, \dots, 9\}$ . The data is pre-processed so that the input has zero mean and unit variance. This is a deterministic setting, hence the experiment is designed to showcase how the convex IB Lagrangians allow us to explore the IB curve in a setting where the normal IB Lagrangian cannot and the relationship between the performance plateaus and the clusterization phenomena. Furthermore, it intends to showcase the behavior of the power and exponential Lagrangians with different parameters of  $\alpha$  and  $\eta$ . Finally, it wants to demonstrate how the value convergence can be employed to approximately obtain a specific compression value. In this experiment, the encoder  $f_{\text{enc}}$  is a three fully-connected layer encoder with 800 ReLU units on the first two layers and two linear units on the last layer ( $T \in \mathbb{R}^2$ ), and the decoder  $f_{\text{dec}}$  is a fully-connected 800 ReLU unit layers followed by an output layer with 10 softmax units. The convex IB Lagrangian was calculated using the nonlinear-IB.



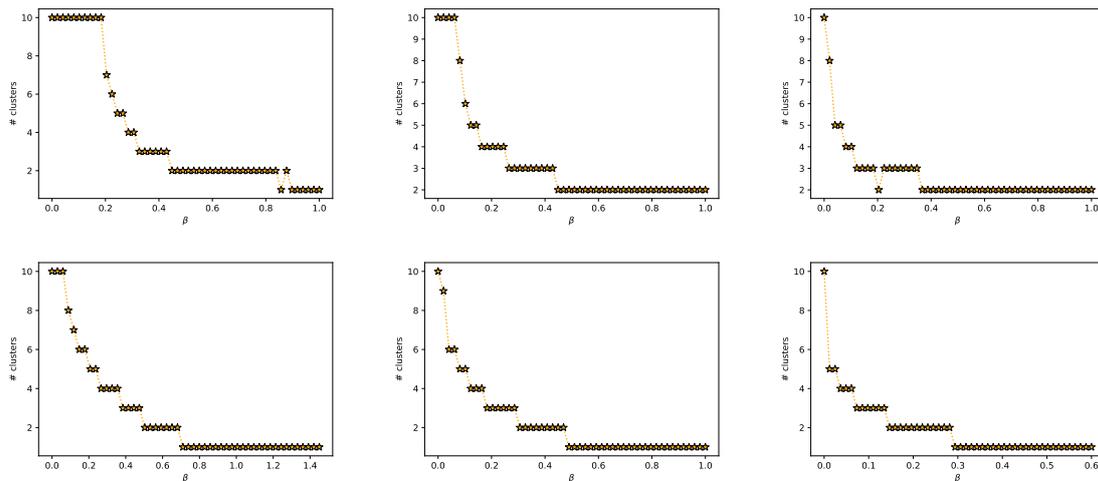
**Figure A2.** Results for the power IB Lagrangian in the MNIST dataset with  $\alpha = \{0.5, 1, 2\}$ , from top to bottom. In each row, from left to right it is shown (i) the information plane, where the region of possible solutions of the IB problem is shadowed in light orange and the information-theoretic limits are the dashed orange line; (ii)  $I(T; Y)$  as a function of  $\beta_u$ ; and (iii) the compression  $I(X; T)$  as a function of  $\beta_u$ . In all plots, the red crosses joined by a dotted line represent the values computed with the training set, the blue dots the values computed with the validation set and the green stars the theoretical values computed as dictated by Proposition 3. Moreover, in all plots, it is indicated  $I(X; Y) = H(Y) = \log_2(10)$  in a dashed, orange line. All values are shown in bits.

In Figure A2 we show how the IB curve can be explored with different values of  $\alpha$  for the power IB Lagrangian and in Figure A3 for different values of  $\eta$  and the exponential IB Lagrangian.

Finally, in Figure A4 we show the clusterization for the same values of  $\alpha$  and  $\eta$  as in Figure A2 and A3. In this way the connection between the performance discontinuities and the clusterization is more evident. Furthermore, we can also observe how the exponential IB Lagrangian maintains better the theoretical performance than the power IB Lagrangian (see Appendix I for an explanation of why).

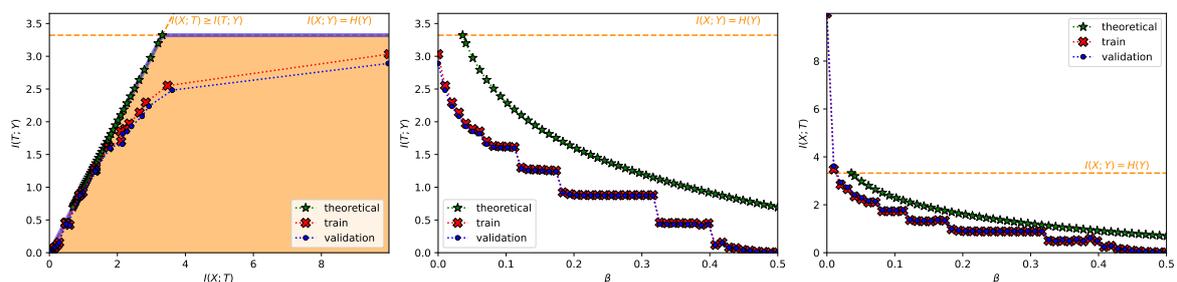


**Figure A3.** Results for the exponential IB Lagrangian in the MNIST dataset with  $\eta = \{\log(2), 1, 1.5\}$ , from top to bottom. In each row, from left to right it is shown (i) the information plane, where the region of possible solutions of the IB problem is shadowed in light orange and the information-theoretic limits are the dashed orange line; (ii)  $I(T; Y)$  as a function of  $\beta_u$ ; and (iii) the compression  $I(X; T)$  as a function of  $\beta_u$ . In all plots, the red crosses joined by a dotted line represent the values computed with the training set, the blue dots the values computed with the validation set and the green stars the theoretical values computed as dictated by Proposition 3. Moreover, in all plots, it is indicated  $I(X; Y) = H(Y) = \log_2(10)$  in a dashed, orange line. All values are shown in bits.



**Figure A4.** Depiction of the clusterization behavior of the bottleneck variable in the MNIST dataset. In the first row, from left to right, the power IB Lagrangian with different values of  $\alpha = \{0.5, 1, 2\}$ . In the second row, from left to right, the exponential IB Lagrangian with different values of  $\eta = \{\log(2), 1, 1.5\}$ .

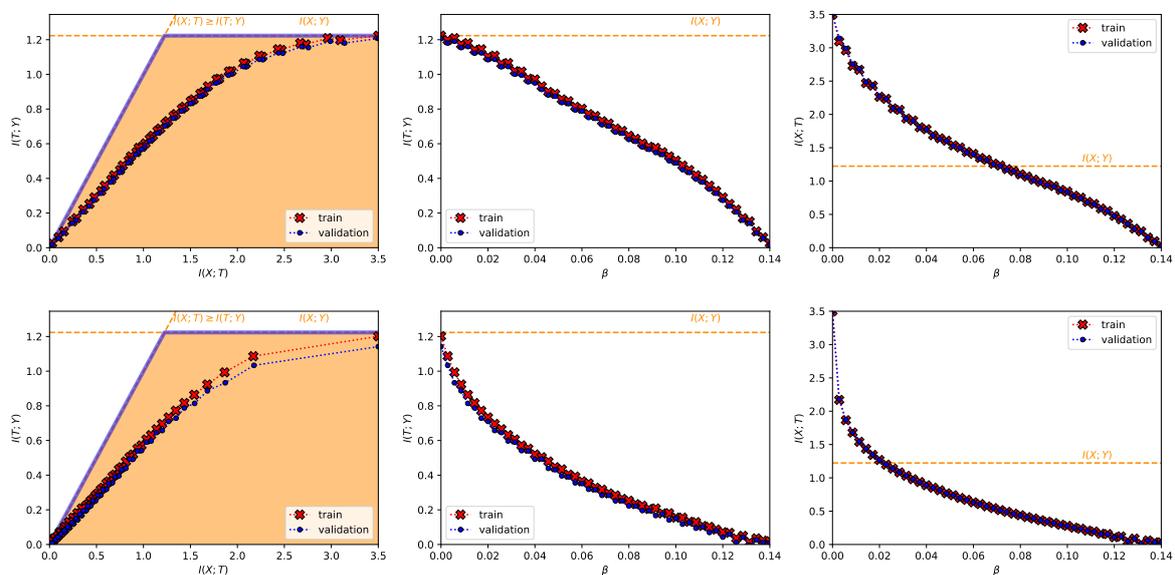
- A Classification Task on the Fashion-MNIST Dataset [49] (Figure A5). As MNSIT, this dataset contains 60,000 training and 10,000 testing samples of 28x28 pixel images labeled from 0 to 9 and constitutes a deterministic setting. The difference is that this dataset contains fashion products instead of hand-written digits and it represents a harder classification task [49]. The data is also pre-processed so that the input has zero mean and unit variance. For this experiment, the encoder  $f_{enc}$  is composed of a two-layer convolutional neural network (CNN) with 32 filters on the first layer and 128 filters on the second with kernels of size 5 and stride 2. This CNN is followed by two fully-connected layers of 128 linear units ( $T \in \mathbb{R}^{128}$ ). After the first convolution and the first fully-connected layer, a ReLU activation is employed. The decoder  $f_{dec}$  is a fully-connected 128 ReLU unit layer followed by an output layer with 10 softmax units. The convex IB Lagrangian was calculated using the nonlinear-IB. Therefore, this experiment intends to showcase how the convex IB Lagrangian can explore the IB curve for different neural network architectures and harder datasets.



**Figure A5.** Results for the exponential IB Lagrangian in the Fashion MNIST dataset with  $\eta = 1$ . From left to right it is shown (i) the information plane, where the region of possible solutions of the IB problem is shadowed in light orange and the information-theoretic limits are the dashed orange line; (ii)  $I(T; Y)$  as a function of  $\beta_u$ ; and (iii) the compression  $I(X; T)$  as a function of  $\beta_u$ . In all plots, the red crosses joined by a dotted line represent the values computed with the training set and the blue dots the values computed with the validation set. Moreover, in all plots, it is indicated  $I(X; Y) = H(Y) = \log_2(10)$ . All values are shown in bits.

- A Regression Task on the California Housing Dataset [50] (Figure A6). This dataset contains 20,640 samples of 8 real number input variables like the longitude and latitude of the house (i.e.,

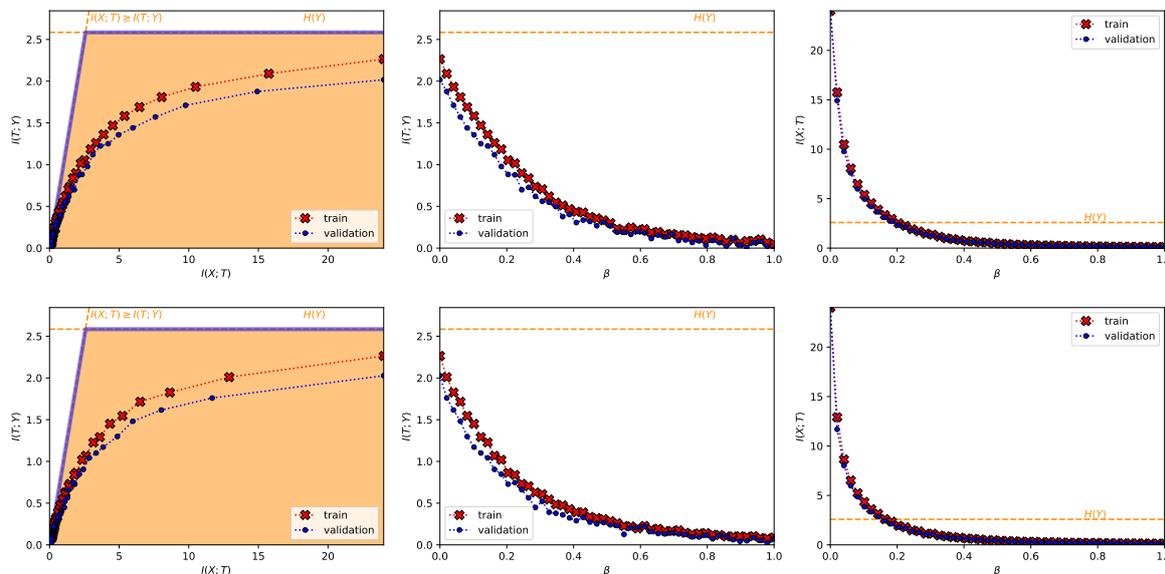
$X \in \mathbb{R}^8$ ) and a task output real variable representing the price of the house (i.e.,  $Y \in \mathbb{R}$ ). We used the log-transformed house price as the target variable and dropped the 992 samples in which the house price was equal or greater than \$500,000 so that the output distribution was closer to a Gaussian as they did in [26]. The input variables were processed so that they had zero mean and unit variance and we randomly split the samples into a 70% training and 30% test dataset. As in [40], for regression tasks we approximate  $H(Y)$  with the entropy of a Gaussian with variance  $\text{Var}(Y)$  and  $H(Y|T)$  with the entropy of a Gaussian with variance equal to the mean-squared error (MSE). This leads to the estimate  $I(T;Y) \approx 0.5 \log(\text{Var}(Y)/\text{MSE})$ . The encoder  $f_{\text{enc}}$  is a three fully-connected layer encoder with 128 ReLU units on the first two layers and 2 linear units on the last layer ( $T \in \mathbb{R}^2$ ), and the decoder  $f_{\text{dec}}$  is a fully-connected 128 ReLU unit layers followed by an output layer with 1 linear unit. The convex IB Lagrangian was calculated using the nonlinear-IB. Hence, this experiment was designed to showcase the convex IB Lagrangian can explore the IB curve in stochastic scenarios for regression tasks.



**Figure A6.** The top row shows the results for the normal IB Lagrangian, and the bottom row for the exponential IB Lagrangian with  $\eta = 1$ , both in the California housing dataset. In each row, from left to right it is shown (i) the information plane, where the region of possible solutions of the IB problem is shadowed in light orange and the information-theoretic limits are the dashed orange line; (ii)  $I(T;Y)$  as a function of  $\beta_u$ ; and (iii) the compression  $I(X;T)$  as a function of  $\beta_u$ . In all plots, the red crosses joined by a dotted line represent the values computed with the training set and the blue dots the values computed with the validation set. Moreover, in all plots, it is indicated  $I(X;Y)$  as the empirical value obtained maximizing  $I(T;Y)$  without compression limitations as in [26]. All values are shown in bits.

- A Classification Task on the TREC-6 Dataset [29] (Figure A7 and bottom row from Figure 3). This dataset is the six-class version of the TREC [51] dataset. It contains 5452 training and 500 test samples of text questions. Each question is labeled within six different semantic categories based on what the answer is; namely: Abbreviation, description and abstract concepts, entities, human beings, locations, and numeric values. This dataset does not constitute a deterministic setting since there are examples that could belong to more than one class and there are examples which are wrongly labeled (e.g., “What is a fear of parasites?” could belong both to the description and abstract concept category, however it is labeled into the entity category), and hence  $H(Y|X) > 0$ . Following Ben Trevett’s tutorial on Sentiment Analysis [52] the encoder  $f_{\text{enc}}$  is composed by a 6 billion token pre-trained 100-dimensional Glove word embedding [53], followed by a concatenation of three convolutions with kernel sizes 2–4 respectively, and finalized with a fully-connected 128 linear unit layer ( $T \in \mathbb{R}^{128}$ ). The decoder  $f_{\text{dec}}$  is a single fully-connected 6

softmax unit layer. The convex IB Lagrangian was calculated using the nonlinear-IB. Thus, this experiment intends to show an example where the classification task does not convey a deterministic scenario, that the convex IB Lagrangian can recover the IB curve in complex stochastic tasks with complex neural network architectures and that the value convergence can be employed to obtain a specific compression value even in stochastic settings where the IB curve is unknown.



**Figure A7.** The top row shows the results for the normal IB Lagrangian, and the bottom row for the power IB Lagrangian with  $\alpha = 0.1$ , both in the TREC-6 dataset. In each row, from left to right it is shown (i) the information plane, where the region of possible solutions of the IB problem is shadowed in light orange and the information-theoretic limits are the dashed orange line; (ii)  $I(T; Y)$  as a function of  $\beta_u$ ; and (iii) the compression  $I(X; T)$  as a function of  $\beta_u$ . In all plots, the red crosses joined by a dotted line represent the values computed with the training set and the blue dots the values computed with the validation set. Moreover, in all plots, it is indicated  $H(Y) = \log_2(6)$ . All values are shown in bits.

### Appendix I. Guidelines for Selecting A Proper Function in the Convex IB Lagrangian

When choosing the right  $u$  function, it is important to find the right balance between avoiding value convergence and aiming for strong convexity. Practically, this balance is found by looking at how much faster  $u$  grows w.r.t. the identity function.

When the aim is not to draw the IB curve but to find a specific level of performance, we can exploit the value convergence phenomenon in order to design a stable performance targeted  $u$  function.

#### Appendix I.1. Avoiding Value Convergence

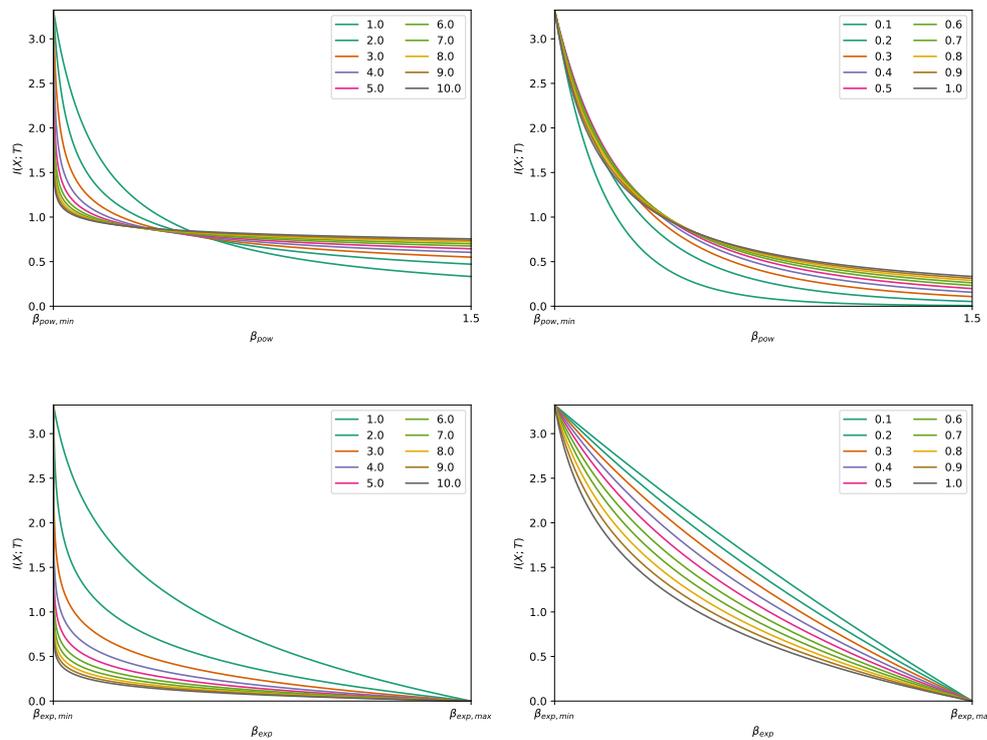
In order to explain this issue we are going to use the example of classification on MNIST [28], where  $I(X; Y) = H(Y) = \log_2(10)$ , and again the power and exponential IB Lagrangians.

If we use Proposition 3 on both Lagrangians we obtain the bijective mapping between their Lagrange multipliers and a certain level of compression in the classification setting:

1. Power IB Lagrangian:  $\beta_{\text{pow}} = ((1 + \alpha)I(X; T)^\alpha)^{-1}$  and  $I(X; T) = ((1 + \alpha)\beta_{\text{pow}})^{-\frac{1}{\alpha}}$ .
2. Exponential IB Lagrangian:  $\beta_{\text{exp}} = (\eta \exp(\eta I(X; T)))^{-1}$  and  $I(X; T) = -\log(\eta\beta_{\text{exp}})/\eta$ .

Hence, we can simply plot the curves of  $I(X; T)$  vs.  $\beta_u$  for different hyperparameters  $\alpha$  and  $\eta$  (see Figure A8). In this way, we can observe how increasing the growth of the function (e.g., increasing  $\alpha$  or

$\eta$  in this case) too much provokes that many different values of  $\beta_u$  converge to very similar values of  $I(X; T)$ . This is an issue both for drawing the curve (for obvious reasons) and for aiming for a specific performance level. Due to the nature of the estimation of the IB Lagrangian, the theoretical and practical value of  $\beta_u$  that yields a specific  $I(X; T)$  may vary slightly (see Figure 1). Then if we select a function with too high growth, a small change in  $\beta_u$  can result in a big change in the performance obtained.



**Figure A8.** Theoretical bijection between  $I(X; T)$  and different  $\alpha$  from  $\beta_{u, min}$  to 1.5 in the power IB Lagrangian (top), and different  $\eta$  in the domain  $B_u$  in the exponential IB Lagrangian (bottom).

Appendix I.2. Aiming for Strong Convexity

**Definition A2 ( $\mu$ -Strong Convexity).** If a function  $f(r)$  is twice continuous differentiable and its domain is confined in the real line, then it is  $\mu$ -strong convex if  $f''(r) \geq \mu \geq 0 \forall r$ .

Experimentally, we observed when the growth of our function  $u(r)$  is small in the domain of interest  $r > 0$  the convex IB Lagrangian does not perform well (see first row of Figures A2 and A3). Later we realized that this was closely related to the strength of the convexity of our function.

In Theorem 2 we imposed the function  $u$  to be strictly convex to enforce having a unique  $\beta_u$  for each value of  $I(X; T)$ . Hence, since in practice we are not exactly computing the Lagrangian but an estimation of it (e.g., with the nonlinear IB [26]) we require strong convexity in order to be able to explore the IB curve.

We now look at the second derivative of the power and exponential function:  $u''(r) = (1 + \alpha)\alpha r^{\alpha-1}$  and  $u''(r) = \eta^2 \exp(\eta r)$  respectively. Here we see how both functions are inherently 0-strong convex for  $r > 0$  and  $\alpha, \eta > 0$ . However, values of  $\alpha < 1$  and  $\eta < 1$  could lead to low  $\mu$ -strong convexity in certain domains of  $r$ . Particularly, the case of  $\alpha < 1$  is dangerous because the function approaches 0-strong convexity as  $r$  increases, so the power IB Lagrangian performs poorly when low  $\alpha$  are used to find high performances.

### Appendix I.3. Exploiting Value Convergence

When the aim is not to draw or explore the IB curve, but to obtain a specific level of performance, the power of exponential IB Lagrangians aforementioned might not be the best choice due to the problems with value convergence or non-strong convexity. However, we can exploit the former in order to design a performance targeted  $u$  function.

For instance, if we look at Figure A8 we can see how a modification of the exponential IB Lagrangian could result in such a function. More precisely, a shifted exponential  $u(r) = \exp(\eta(r - r^*))$ , with  $\eta > 0$  sufficiently large, converges to the compression level  $r^*$ . We can see this more clearly if we consider the shifted exponential IB Lagrangian  $\mathcal{L}_{\text{IB,sh-exp}}(T; \beta_{\text{sh-exp}}, \eta, r^*) = I(T; Y) - \beta_{\text{sh-exp}} \exp(\eta(I(X; T) - r^*))$ , since then the application of Proposition 3 results on  $I(X; T) = -\log(\eta \beta_{\text{sh-exp}} / f'_{\text{IB}}(I(X; T))) / \eta + r^*$ , where  $f'_{\text{IB}}(I(X; T))$  is the derivative of  $f_{\text{IB}}$  evaluated at  $I(X; T)$ . We know  $f'_{\text{IB}} = 1$  in deterministic scenarios (Theorem 2) and that  $f'_{\text{IB}} < 1$  otherwise (see, e.g., [27]). Then, for large enough  $\eta$ ,  $I(X; T) \approx r^*$  regardless of the value of  $f'_{\text{IB}}$ .

For instance, if we consider a deterministic scenario like the MNIST dataset [28] with  $I(X; Y) = H(Y) = \log_2(10)$ , for  $\eta = 200$  and  $r^* = 2$  the range of the Lagrange multipliers that allow the exploration of the IB curve, according to Corollary 2, is  $\beta_{\text{sh-exp}} \in [7.54 \times 10^{-178}, 2.61 \times 10^{171}]$ . Furthermore,  $I(X; T)$  is close to 2 for many values of  $\beta_{\text{sh-exp}}$ . For instance,  $I(X; T) = 1.974$  for  $\beta_{\text{sh-exp}} = 1$  and  $I(X; T) = 1.963$  for  $\beta_{\text{sh-exp}} = 8$ . This ensures a stability in the performance level obtained so that small changes in the choice of  $\beta_{\text{sh-exp}}$  do not result in significant changes on the performance (e.g., see top row from Figure 4).

If we now consider a stochastic scenario like the TREC-6 dataset [29] with  $H(Y) = \log_2(6)$ , for  $\eta = 200$  and  $r^* = 16$  the range of the Lagrange multipliers that allow the IB curve, according to Corollary 3, is  $\beta_{\text{sh-exp}} \in [0, 2.76(\inf_{\Omega_x \subset \mathcal{X}} \{\beta_0(\Omega_x)\})^{-1} \times 10^{1287}]$ , where  $\beta_0$  and  $\Omega_x$  are defined as in [27]. Then, unless  $(\inf_{\Omega_x \subset \mathcal{X}} \{\beta_0(\Omega_x)\})^{-1}$  is of the order of  $10^{-1287}$ , the range of possible betas is wide. Moreover,  $I(X; T)$  is close to 16 for many values of  $\beta_{\text{sh-exp}}$ . For example,  $I(X; T) = 15.939$  if  $f'_{\text{IB}} = 0.001$  at that point and  $I(X; T) = 15.973$  if  $f'_{\text{IB}} = 0.9$  for  $\beta_{\text{sh-exp}} = 1$ ; and  $I(X; T) = 15.929$  if  $f'_{\text{IB}} = 0.001$  at that point and  $I(X; T) = 15.963$  if  $f'_{\text{IB}} = 0.9$  for  $\beta_{\text{sh-exp}} = 8$ . Hence, as in the deterministic scenario, the performance level obtained is stable with changes in the choice of  $\beta_{\text{sh-exp}}$  (e.g., see bottom row from Figure 4).

## References

1. Tishby, N.; Pereira, F.C.; Bialek, W. The information bottleneck method. *arXiv* **2000**, arXiv:physics/0004057.
2. Alemi, A.A.; Fischer, I.; Dillon, J.V.; Murphy, K. Deep variational information bottleneck. *arXiv* **2016**, arXiv:1612.00410.
3. Peng, X.B.; Kanazawa, A.; Toyer, S.; Abbeel, P.; Levine, S. Variational Discriminator Bottleneck: Improving Imitation Learning, Inverse RL, and GANs by Constraining Information Flow. In Proceedings of the International Conference on Learning Representations (ICLR), New Orleans, LA, USA, 6–9 May 2019.
4. Achille, A.; Soatto, S. Information dropout: Learning optimal representations through noisy computation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2018**, *40*, 2897–2905. [[CrossRef](#)] [[PubMed](#)]
5. Slonim, N.; Tishby, N. Document clustering using word clusters via the information bottleneck method. In Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval, Athens, Greece, 24–28 July 2000.
6. Slonim, N.; Tishby, N. Agglomerative information bottleneck. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2000.
7. Slonim, N.; Atwal, G.S.; Tkačik, G.; Bialek, W. Information-based clustering. *Proc. Natl. Acad. Sci. USA* **2005**, *102*, 18297–18302. [[CrossRef](#)] [[PubMed](#)]
8. Teahan, W.J. Text classification and segmentation using minimum cross-entropy. In *Content-Based Multimedia Information Access*; LE CENTRE DE HAUTES ETUDES INTERNATIONALES D'INFORMATIQUE DOCUMENTAIRE: Paris, France, 2000; pp. 943–961.
9. Strouse, D.; Schwab, D.J. The deterministic information bottleneck. *Neur. Comput.* **2017**, *29*, 1611–1630. [[CrossRef](#)]

10. Nazer, B.; Ordentlich, O.; Polyanskiy, Y. Information-distilling quantizers. In Proceedings of the 2017 IEEE International Symposium on Information Theory (ISIT), Aachen, Germany, 25–30 June, 2017; pp. 96–100.
11. Hassanpour, S.; Wübben, D.; Dekorsy, A. On the equivalence of double maxima and KL-means for information bottleneck-based source coding. In Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC), Barcelona, Spain, 15–18 April 2018; pp. 1–6.
12. Goyal, A.; Islam, R.; Strouse, D.; Ahmed, Z.; Botvinick, M.; Larochelle, H.; Levine, S.; Bengio, Y. Infobot: Transfer and exploration via the information bottleneck. *arXiv* **2019**, arXiv:1901.10902.
13. Yingjun, P.; Xinwen, H. Learning Representations in Reinforcement Learning: An Information Bottleneck Approach. *arXiv* **2019**, arXiv:cs.LG/1911.05695.
14. Sharma, A.; Gu, S.; Levine, S.; Kumar, V.; Hausman, K. Dynamics-Aware Unsupervised Skill Discovery. In Proceeding of the International Conference on Learning Representations (ICLR), Addis Ababa, Ethiopia, 26–30 April 2020.
15. Schulz, K.; Sixt, L.; Tombari, F.; Landgraf, T. Restricting the Flow: Information Bottlenecks for Attribution. In Proceedings of the International Conference on Learning Representations (ICLR), Addis Ababa, Ethiopia, 26–30 April 2020.
16. Li, X.L.; Eisner, J. Specializing Word Embeddings (for Parsing) by Information Bottleneck. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Hong Kong, China, 3–7 November 2019; pp. 2744–2754.
17. Zaslavsky, N.; Kemp, C.; Regier, T.; Tishby, N. Efficient compression in color naming and its evolution. *Proc. Natl. Acad. Sci. USA* **2018**, *115*, 7937–7942. [[CrossRef](#)]
18. Chalk, M.; Marre, O.; Tkačik, G. Toward a unified theory of efficient, predictive, and sparse coding. *Proc. Natl. Acad. Sci. USA* **2018**, *115*, 186–191. [[CrossRef](#)]
19. Gilad-Bachrach, R.; Navot, A.; Tishby, N. An information theoretic tradeoff between complexity and accuracy. In *Learning Theory and Kernel Machines*; Springer: Berlin, Germany, 2003; pp. 595–609.
20. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*; John Wiley & Sons: Hoboken, NJ, USA, 2012.
21. Kolchinsky, A.; Tracey, B.D.; Van Kuyk, S. Caveats for information bottleneck in deterministic scenarios. In Proceedings of the International Conference on Learning Representations (ICLR), New Orleans, LA, USA, 6–9 May 2019.
22. Courcoubetis, C. *Pricing Communication Networks Economics, Technology and Modelling*; Wiley Online Library: Hoboken, NJ, USA, 2003.
23. Tishby, N.; Slonim, N. Data clustering by markovian relaxation and the information bottleneck method. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2001; pp. 640–646.
24. Slonim, N.; Friedman, N.; Tishby, N. Unsupervised document classification using sequential information maximization. In Proceedings of the 25th annual international ACM SIGIR Conference on Research and Development in Information Retrieval, Tampere, Finland, 11–15 August 2002.
25. Chalk, M.; Marre, O.; Tkacik, G. Relevant sparse codes with variational information bottleneck. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2016; pp. 1957–1965.
26. Kolchinsky, A.; Tracey, B.D.; Wolpert, D.H. Nonlinear information bottleneck. *Entropy* **2019**, *21*, 1181. [[CrossRef](#)]
27. Wu, T.; Fischer, I.; Chuang, I.; Tegmark, M. Learnability for the Information Bottleneck. In Proceedings of the International Conference on Learning Representations (ICLR), New Orleans, LA, USA, 6–9 May 2019.
28. LeCun, Y.; Bottou, L.; Bengio, Y.; Haffner, P. Gradient-based learning applied to document recognition. In Proceedings of the 1998 IEEE International Frequency Control Symposium, Pasadena, CA, USA, 27–29 May 1998.
29. Li, X.; Roth, D. Learning question classifiers. In *Proceedings of the 19th international conference on Computational linguistics—Volume 1*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2002; pp. 1–7.
30. Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; Lerer, A. Automatic differentiation in pytorch. In Proceedings of the NIPS Autodiff Workshop, Long Beach, CA, USA, 9 December 2017.
31. Bishop, C.M. *Pattern Recognition and Machine Learning*; Springer Science+ Business Media: Berlin, Germany, 2006.

32. Xu, A.; Raginsky, M. Information-theoretic analysis of generalization capability of learning algorithms. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2017; pp. 2524–2533.
33. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2012; pp. 1097–1105.
34. Shore, J.E.; Gray, R.M. Minimum cross-entropy pattern classification and cluster analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **1982**, *1*, 11–17. [[CrossRef](#)] [[PubMed](#)]
35. Shore, J.; Johnson, R. Properties of cross-entropy minimization. *IEEE Trans. Pattern Anal. Mach. Intell.* **1981**, *27*, 472–482. [[CrossRef](#)]
36. Vera, M.; Piantanida, P.; Vega, L.R. The role of the information bottleneck in representation learning. In Proceedings of the 2018 IEEE International Symposium on Information Theory (ISIT), Vail, CO, USA, 17–22 June 2018; pp. 1580–1584.
37. Shamir, O.; Sabato, S.; Tishby, N. Learning and generalization with the information bottleneck. *Theor. Comput. Sci.* **2010**, *411*, 2696–2711. [[CrossRef](#)]
38. Achille, A.; Soatto, S. Emergence of invariance and disentanglement in deep representations. *J. Mach. Learn. Res.* **2018**, *19*, 1947–1980.
39. Du Pin Calmon, F.; Polyanskiy, Y.; Wu, Y. Strong data processing inequalities for input constrained additive noise channels. *IEEE Trans. Inf. Theory* **2017**, *64*, 1879–1892. [[CrossRef](#)]
40. Kolchinsky, A.; Tracey, B. Estimating mixture entropy with pairwise distances. *Entropy* **2017**, *19*, 361. [[CrossRef](#)]
41. Amjad, R.A.; Geiger, B.C. Learning representations for neural network-based classification using the information bottleneck principle. *IEEE Trans. Pattern Anal. Mach. Intell.* **2019**, *1*. [[CrossRef](#)]
42. Alemi, A.A.; Fischer, I.; Dillon, J.V. Uncertainty in the variational information bottleneck. *arXiv* **2018**, arXiv:1807.00906.
43. Wu, T.; Fischer, I. Phase Transitions for the Information Bottleneck in Representation Learning. In Proceedings of the International Conference on Learning Representations (ICLR), Addis Ababa, Ethiopia, 26–30 April 2020.
44. Ester, M.; Kriegel, H.P.; Sander, J.; Xu, X. A density-based algorithm for discovering clusters in large spatial databases with noise. In Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, Menlo Park, CA, USA, 2–4 August 1996; pp. 226–231.
45. Schubert, E.; Sander, J.; Ester, M.; Kriegel, H.P.; Xu, X. DBSCAN revisited, revisited: Why and how you should (still) use DBSCAN. *ACM Trans. Database Syst. TODS* **2017**, *42*, 19. [[CrossRef](#)]
46. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.
47. Glorot, X.; Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, Sardinia, Italy, 13–15 May 2010; pp. 249–256.
48. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
49. Xiao, H.; Rasul, K.; Vollgraf, R. Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms. *arXiv* **2017**, arXiv:1708.07747.
50. Pace, R.K.; Barry, R. Sparse spatial autoregressions. *Stat. Probab. Lett.* **1997**, *33*, 291–297. [[CrossRef](#)]
51. Voorhees, E.M.; Tice, D.M. Building a question answering test collection. In Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Athens, Greece, 24–28 July 2000.
52. Trevett, Ben. Tutorial on Sentiment Analysis: 5—Multi-class Sentiment Analysis. April 2019. Available online: <https://github.com/bentrevett/pytorch-sentiment-analysis> (accessed on 14 January 2020).
53. Pennington, J.; Socher, R.; Manning, C. Glove: Global vectors for word representation. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 25–29 October 2014; pp. 1532–1543.

