

MULTI-VIEW GENERATIVE ADVERSARIAL NETWORKS

Mickaël Chen

Sorbonne Universités, UPMC Univ Paris 06, UMR 7606, LIP6, F-75005, Paris, France
mickael.chen@lip6.fr

Ludovic Denoyer

Sorbonne Universités, UPMC Univ Paris 06, UMR 7606, LIP6, F-75005, Paris, France
ludovic.denoyer@lip6.fr

ABSTRACT

Learning over multi-view data is a challenging problem with strong practical applications. Most related studies focus on the classification point of view and assume that all the views are available at any time. We consider an extension of this framework in two directions. First, based on the BiGAN model, the Multi-view BiGAN (MV-BiGAN) is able to perform density estimation from multi-view inputs. Second, it can deal with missing views and is able to update its prediction when additional views are provided. We illustrate these properties on a set of experiments over different datasets.

1 INTRODUCTION

Many concrete applications involve multiple sources of information generating different views on the same object (Cesa-Bianchi et al., 2010). If we consider human activities for example, GPS values from a mobile phone, navigation traces over the Internet, or even photos published on social networks are different views on a particular user. In multimedia applications, views can correspond to different modalities (Atrey et al., 2010) such as sounds, images, videos, sequences of previous frames, etc...

The problem of multi-view machine learning has been extensively studied during the last decade, mainly from the classification point of view. In that case, one wants to predict an output y based on multiple views acquired on an unknown object x . Different strategies have been explored but a general common idea is based on the (early or late) fusion of the different views at a particular level of a deep architecture (Wang et al., 2015; Ngiam et al., 2011; Srivastava & Salakhutdinov, 2012).

The existing literature mainly explores problems where outputs are chosen in a discrete set (e.g categorization), and where all the views are available. An extension of this problem is to consider the density estimation problem where one wants to estimate the conditional probabilities of the outputs given the available views. As noted by Mathieu et al. (2015), minimizing classical prediction losses (e.g Mean square error) will not capture the different output distribution modalities.

In this article, we propose a new model able to estimate a distribution over the possible outputs given any subset of views on a particular input. This model is based on the (Bidirectional) *Generative Adversarial Networks* (BiGAN) formalism. More precisely, we bring two main contributions: first, we propose the CV-BiGAN (*Conditional Views BiGAN* – Section 3) architecture that allows one to model a conditional distribution $P(y|\cdot)$ in an original way. Second, on top of this architecture, we build the Multi-view BiGANs (MV-BiGAN – Section 4) which is able to both **predict when only one or few views are available**, and to **update its prediction if new views are added**. We evaluate this model on different multi-views problems and different datasets (Section 5). The related work is provided in Section 6 and we propose some future research directions in Section 7.

2 BACKGROUND AND GENERAL IDEA

2.1 NOTATIONS AND TASK

Let us denote \mathcal{X} the space of objects on which different views will be acquired. Each possible input $x \in \mathcal{X}$ is associated to a target prediction $y \in \mathbb{R}^n$. A classical machine learning problem is to estimate $P(y|x)$ based on the training set. But we consider instead a multi-view problem in which different views on x are available, x being unknown. Let us denote V the number of possible views and \tilde{x}_k the k -th view over x . The description space for view k is \mathbb{R}^{n_k} where n_k is the number of features in view k . Moreover, we consider that some of the V views can be missing. The subset of available views for input x^i will be represented by an index vector $s^i \in \mathcal{S} = \{0, 1\}^V$ so that $s_k^i = 1$ if view k is available and $s_k^i = 0$ elsewhere. Note that all the V views will not be available for each input x , and the prediction model must be able to predict an output given any subset of views $s \in \{0, 1\}^V$.

In this configuration, our objective is to estimate the distributions $p(y|v(s, x))$ where $v(s, x)$ is the set of views \tilde{x}_k so that $s_k = 1$. This distribution p will be estimated based on a training set \mathcal{D} of N training examples. Each example is composed of a subset of views $s^i, v(s^i, x^i)$ associated to an output y^i , so that $\mathcal{D} = \{(y^1, s^1, v(s^1, x^1)), \dots, (y^N, s^N, v(s^N, x^N))\}$ where s^i is the index vector of views available for x^i . Note that x^i is not directly known in the training set but only observed through its associated views.

2.2 BIDIRECTIONAL GENERATIVE ADVERSARIAL NETS (BiGAN)

We quickly remind the principle of BiGANs since our model is an extension of this technique. Generative Adversarial Networks (GAN) have been introduced by Goodfellow et al. (2014) and have demonstrated their ability to model complex distributions. They have been used to produce compelling natural images from a simple latent distribution (Radford et al., 2015; Denton et al., 2015). Exploring the latent space has uncovered interesting, meaningful patterns in the resulting outputs. However, GANs lack the ability to retrieve a latent representation given an output, missing out an opportunity to exploit the learned manifold. Bidirectional Generative Adversarial Networks (BiGANs) have been proposed by Donahue et al. (2016) and Dumoulin et al. (2016), independently, to fill that gap. BiGANs simultaneously learn both an encoder function E that models the encoding process $P_E(z|y)$ from the space \mathbb{R}^n to a latent space \mathbb{R}^Z , and a generator function G that models the mapping distribution $P_G(y|z)$ of any latent point $z \in \mathbb{R}^Z$ to a possible object $y \in \mathbb{R}^n$. From both the encoder distribution and the generator distribution, we can model two joint distributions, respectively denoted $P_E(y, z)$ and $P_G(y, z)$:

$$\begin{aligned} P_G(y, z) &= P(z)P_G(y|z) \\ P_E(y, z) &= P(y)P_E(z|y) \end{aligned} \quad (1)$$

assuming that $P(z) = \mathcal{N}(0, 1)$ and $P(y)$ can be estimated over the training set by a uniform sampling. The BiGAN framework also introduces a discriminator network D_1 whose task is to determine whether a pair (y, z) is sampled from $p_G(y, z)$ or from $p_E(y, z)$, while E and G are trained to fool D_1 , resulting in the following learning problem:

$$\min_{G, E} \max_{D_1} \mathbb{E}_{y \sim P(y), z \sim P_E(z|y)} [\log D_1(y, z)] + \mathbb{E}_{z \sim P(z), y \sim P_G(y|z)} [1 - \log D_1(y, z)] \quad (2)$$

It can be shown, by following the same steps as in Goodfellow et al. (2014), that the optimization problem described in Equation 2 minimizes the Jensen-Shanon divergence between $P_E(y, z)$ and $P_G(y, z)$, allowing the model to learn both a decoder and a generator over a training set that will model the joint distribution of (y, z) pairs. As proposed by Dumoulin et al. (2016), we consider in the following that $P_G(y|z)$ is modeled by a deterministic non-linear model G so that $G(z) = y$, and P_E as a diagonal Gaussian distribution $E(z) = (\mu(y), \sigma(y))$. G , μ and σ are estimated by using gradient-based descent techniques.

2.3 GENERAL IDEA

We propose a model based on the Generative Adversarial Networks paradigm adapted to the multi-view prediction problem. Our model is based on two different principles:

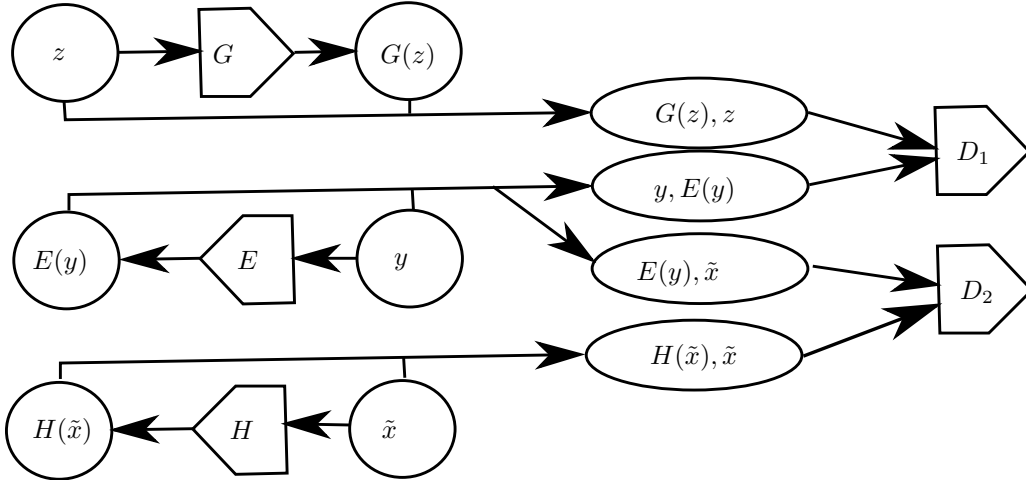


Figure 1: The CV-BiGAN Architecture. The two top levels correspond to the BiGAN model, while the third level is added to model the distribution over the latent space given the input of the CV-BiGAN. The discriminator D_2 is used to constraint $P(z|y)$ and $P(z|\tilde{x})$ to be as close as possible.

Conditional Views BiGANs (CV-BiGAN): First, since one wants to model an output distribution based on observations, our first contribution is to propose an adaptation of BiGANs to model conditional probabilities, resulting in a model able to learn $P(y|\tilde{x})$ where \tilde{x} can be either a single view or an aggregation of multiple views. If conditional GANs have already been proposed in the literature (see Section 6) they are not adapted to our problem which require explicit mappings between input space to latent space, and from latent space to output space.

Multi-View BiGANs (MV-BiGAN): On top of the CV-BiGAN model, we build a multi-view model able to estimate the distribution of possible outputs based on any subset of views $v(s, x)$. If a natural way to extend the Conditional BiGANs for handling multi-view is to define a mapping function which map the set of views to a representation space (see Section 4.1) the resulting model has shown undesirable behaviors (see Section 5.1). Therefore, we propose to constrain the model based on the idea that adding one more view to any subset of views must decrease the uncertainty on the output distribution i.e the more views are provided, the less variance the output distribution has. This behavior is encouraged by using a Kullback-Leibler divergence (KL) regularization (see Section 4.2).

3 THE CONDITIONAL BiGAN MODEL (CV-BiGAN)

Our first objective is to extend the BiGAN formalism to handle an input space (e.g a single observed view) in addition to the output space \mathbb{R}^n . We will denote \tilde{x} the observation and y the output to predict. In other words, we wish to capture the conditional probability $P(y|\tilde{x})$ from a given training dataset. Assuming one possesses a bidirectional mapping between the input space and an associated representation space, ie. $P_E(z|y)$ and $P_G(y|z)$, one can equivalently capture $P(z|\tilde{x})$. The CV-BiGAN model keeps the encoder E and generator G defined previously but also includes an additional encoder function denoted H which goal is to map a value \tilde{x} to the latent space \mathbb{R}^Z . Applying H on any value of \tilde{x} results in a distribution $P_H(z|\tilde{x}) = \mathcal{N}(\mu_H(\tilde{x}), \sigma_H(\tilde{x}))$ so that a value of z can be sampled from this distribution. This would then allow one to recover a distribution $P(y|\tilde{x})$.

Given a pair (\tilde{x}, y) , we wish a latent representation z sampled from $P_H(z|\tilde{x})$ to be similar to one from $P_E(z|y)$. As our goal here is to learn $P(z|\tilde{x})$, we define two joint distributions between \tilde{x} and z :

$$\begin{aligned} P_H(\tilde{x}, z) &= P_H(z|\tilde{x})P(\tilde{x}) \\ P_E(\tilde{x}, z) &= \sum_y P_E(z|y)P(\tilde{x}, y) \end{aligned} \quad (3)$$




Figure 2: The MV-BiGAN additional components. In this example, we consider a case where only \tilde{x}^1 is available (top level) and a second case where both \tilde{x}^1 and \tilde{x}^3 are available. The distribution $P(z|\tilde{x}^1, \tilde{x}^3)$ is encouraged to be "included" in $P(z|\tilde{x}^1)$ by the KL constraint. The aggregation of the views is made by the ϕ_k functions that are learned conjointly with the rest of the model.

Minimizing the Jensen-Shanon divergence between these two distributions is equivalent to solving the following adversarial problem:

$$\min_{E,H} \max_{D_2} \mathbb{E}_{\tilde{x}, y \sim P(\tilde{x}, y), z \sim P_E(z|y)} [\log D_2(\tilde{x}, z)] + \mathbb{E}_{\tilde{x}, y \sim P(\tilde{x}, y), z \sim P_H(z|\tilde{x})} [1 - \log D_2(\tilde{x}, z)] \quad (4)$$

Note that when applying stochastic gradient-based descent techniques over this objective function, the probability $P(\tilde{x}, y)$ is approximated by sampling uniformly from the training set. We can sample from $P_H(\tilde{x}, z)$ and $P_E(\tilde{x}, z)$ by forwarding the pair (\tilde{x}, y) into the corresponding network.

By merging the two objective functions defined in Equation 2 and 4, the final learning problem for our Conditionnal BiGANs is defined as:

$$\min_{G,E,H} \max_{D_1, D_2} \mathbb{E}_{\tilde{x}, y \sim P(\tilde{x}, y), z \sim P_E(z|y)} [\log D_1(y, z)] + \mathbb{E}_{z \sim P(z), y \sim P_G(y|z)} [1 - \log D_1(y, z)] \\ + \mathbb{E}_{\tilde{x}, y \sim P(\tilde{x}, y), z \sim P_E(z|y)} [\log D_2(\tilde{x}, z)] + \mathbb{E}_{\tilde{x}, y \sim P(\tilde{x}, y), z \sim P_H(z|\tilde{x})} [1 - \log D_2(\tilde{x}, z)] \quad (5)$$

The general idea of CV-BiGAN is illustrated in Figure 1.

4 MULTI-VIEW BIGAN

4.1 AGGREGATING MULTI-VIEWS FOR CV-BIGAN

We now consider the problem of computing an output distribution conditioned by multiple different views. In that case, we can use the CV-BiGAN Model (or other conditional approaches) conjointly with a model able to aggregate the different views where A is the size of the aggregation space. Instead of considering the input \tilde{x} , we define an aggregation model Ψ . $\Psi(v(s, x))$ will be the representation of the aggregation of all the available views \tilde{x}_k^1 :

$$\Psi(v(s, x)) = \sum_{k=1}^V s_k \phi_k(\tilde{x}^k) \quad (6)$$

where ϕ_k is a function that will be learned that maps a particular view in \mathbb{R}^{n_k} to the aggregation space \mathbb{R}^A . By replacing \tilde{x} in Equation 5, one can then simultaneously learn the functions ϕ_k and the distributions P_H , P_E and P_D , resulting in a multi-view model able to deal with any subset of views.

4.2 UNCERTAINTY REDUCTION ASSUMPTION

However, the previous idea suffers from a very high instability when learning, as it is usually noted with complex GANs architectures (see Section 5). In order to stabilize our model, we propose to add a regularization based on the idea that adding new views to an existing subset of views should reduce the uncertainty over the outputs. Indeed, under the assumption that views are consistent one another, adding a new view should allow to refine the predictions and reduce the variance of the distribution of the outputs.

¹ Note that other aggregation scheme can be used like recurrent neural networks for example.

Let us consider an object x and two index vectors s and s' such that $v(x, s) \subset v(x, s')$ ie. $\forall k, s'_k \geq s_k$. Then, intuitively, $P(x|v(x, s'))$ should be "included" in $P(x|v(x, s))$. In the CV-GAN model, since $P(y|z)$ is deterministic, this can be enforced at a latent level by minimizing $KL(P(z|v(x, s')) || P(z|v(x, s)))$. By assuming those two distributions are diagonal gaussian distributions (ie. $P(z|v(x, s')) = \mathcal{N}(\mu_1, \Sigma_1)$ and $P(z|v(x, s)) = \mathcal{N}(\mu_2, \Sigma_2)$ where Σ_k are diagonal matrices with diagonal elements $\sigma_{k(i)}$), the KL divergence can be computed as in Equation 7 and differentiated.

$$KL(P(z|v(x, s')) || P(z|v(x, s))) = \frac{1}{2} \sum_{i=1}^Z \left(-1 - \log \left(\frac{\sigma_{1(i)}^2}{\sigma_{2(i)}^2} \right) + \frac{\sigma_{1(i)}^2}{\sigma_{2(i)}^2} + \frac{(\mu_{1(i)} - \mu_{2(i)})^2}{\sigma_{2(i)}^2} \right) \quad (7)$$

Note that this divergence is written on the estimation made by the function H and will act as a regularization over the latent conditional distribution.

The final objective function of the MV-BiGAN can be written as:

$$\begin{aligned} \min_{G, E, H} \max_{D_1, D_2} & \mathbb{E}_{s, x, y \sim P(s, x, y), z \sim P_E(z|y)} [\log D_1(y, z)] + \mathbb{E}_{z \sim P(z), y \sim P_G(y|z)} [1 - \log D_1(y, z)] \\ & + \mathbb{E}_{s, x, y \sim P(s, x, y), z \sim P_E(z|y)} [\log D_2(v(x, s), z)] + \mathbb{E}_{s, x, y \sim P(s, x, y), z \sim P_H(z|v(x, s))} [1 - \log D_2(v(x, s), z)] \\ & + \lambda \mathbb{E}_{x \sim P(x)} \sum_{\substack{s, s' \in \mathcal{S}_x \\ \forall k, s'_k \geq s_k}} KL(H(v(x, s')) || H(v(x, s))) \end{aligned} \quad (8)$$

where λ controls the strength of the regularization. Note that aggregation models Ψ are included into H and D_2 and can be optimized conjointly in this objective function.

4.3 LEARNING THE MV-BiGAN

The different functions E, G, H, D_1 and D_2 are implemented as parametric neural networks and trained by mini-batch stochastic gradient descent (see Appendix for more details concerning the architectures). We first update the discriminators networks D_1 and D_2 , then we update the generator and encoders G, E and H with gradient steps in the opposite direction.

As with most other implementation of GAN-based models, we find that using an alternative objective proposed by Goodfellow et al. (2014) for E, G and H instead leads to more stable training. The new objective consist of swapping the labels for the discriminators instead of reversing the gradient. We also find that we can update all the modules in one pass instead of taking alternate gradient steps while obtaining similar results.

Note that the MV-BiGAN model is trained based on datasets where all the V views are available for each data point. In order to generate examples where only subsets of views are available, the ideal procedure would be to consider all the possible subsets of views. Due to the number of data points that would be generated by such a procedure, we build random sequences of incremental sets of views and enforce the KL regularization over successive sets.

5 EXPERIMENTS

We evaluate our model on three different types of experiments, and on two different datasets. The first dataset we experiment on is the MNIST dataset of handwritten digits. The second dataset is the CelebA (Liu et al., 2015) dataset composed of both images of faces and corresponding attributes. The MNIST dataset is used to illustrate the ability of the MV-BiGAN to handle different subset of views, and to update its prediction when integrating new incoming views. The CelebA dataset is used to demonstrate the ability of MV-BiGAN to deal with different types (heterogeneous) of views.

5.1 MNIST, 4 VIEWS

We consider the problem where 4 different views can be available, each view corresponding to a particular quarter of the final image to predict – each view is a vector of $\mathbb{R}^{(14 \times 14)}$. The MV-BiGAN




Figure 3: Results of the MV-BiGAN on sequences of 4 different views. The first column corresponds to the provided views, while the other columns correspond to outputs sampled by the MV-BiGAN.

is used here to recover the original image. The model is trained on the MNIST training digits, and results are provided on the MNIST testing dataset.

Figure 3 illustrates the results obtained for some digits. In this figure, the first column displays the input (the subset of views), while the other columns shows predicted outputs sampled by the MV-BiGAN. An additional view is added between each row. This experiment shows that when new views are added, the diversity in the predicted outputs decreases due to the KL-constraint introduced in the model, which is the desired behavior i.e more information implied less variance. When removing the KL constraint (Figure 4), the diversity still remains important, even if many views are provided to the model. This show the importance of the KL regularization term in the MV-BiGAN objective.




Figure 4: MV-BiGAN without KL-constraint (i.e $\lambda = 0$).

5.2 MNIST, SEQUENCE OF INCOMING VIEWS

We made another set of experiments where the views correspond to images with missing values (missing values are replaced by 0.5). This can be viewed as a data imputation problem – Figure 5. Here also, the behavior of the MV-BiGAN exhibits interesting properties: the model is able to predict the desired output as long as enough information has been provided. When only non-informative views are provided, the model produces digits with a high diversity, the diversity decreasing when new information is added.

5.3 CELEBA, INTEGRATING HETEROGENEOUS INFORMATION

At last, the third experiment aims at measuring the ability of MV-BiGAN to handle heterogeneous inputs. We consider two views: (i) the attribute vector containing information about the person in the picture (hair color, sex, ...), and (ii) a incomplete face. Figure 9 illustrates the results obtained on two faces. The first line corresponds to the faces generated based on the attribute vector. One can see that the attribute information has been captured by the model: for example, the sex of the




Figure 5: MV-BiGAN with sequences of incoming views. Here, each view is a 28×28 matrix (values are between 0 and 1 with missing values replaced by 0.5).




Figure 6: Results obtained on the CelebA dataset for two examples. The first line corresponds to the images generated based on the attribute vector, the second line corresponds to images generated based on the incomplete face, the third line corresponds to the images generated based on the two partial views. The groundtruth face is given in the bottom-left corner, while the incomplete face is given in the top-left corner.

generated face is constant (only women) showing that MV-BiGAN has captured this information from the attribute vector. The second line corresponds to the faces generated when using the incomplete face as an input. One can also see that the generated outputs are "compatible" with the incomplete information provided to the model. But the attribute are not considered (for example, women and men are generated). At last, the third line corresponds to images generated based on the two partial views (attributes and incomplete face) which are close to the ground-truth image (bottom left). Note that, in this set of experiments, the convergence of the MV-BiGAN was quite difficult to obtain, and the quality of the generated faces is still not satisfying.

6 RELATED WORK

Multi-view and Representation Learning: Many application fields naturally deal with multi-view data with true advantages. For example, in the multimedia domain, dealing with a bunch of views is usual (Atrey et al., 2010): text, audio, images (different framings from videos) are starting points of these views. Besides, multimedia learning tasks from multi-views led to a large amount of fusion-based ad-hoc approaches and experimental results. The success of multi-view supervised learning approaches in the multimedia community seems to rely on the ability of the systems to deal with the complementary of the information carried by each modality. Comparable studies are of importance

in many domains, such as bioinformatics (Sokolov & Ben-Hur, 2011), speech recognition (Arora & Livescu, 2012; Koço et al., 2012), signal-based multimodal integration (Wu et al., 1999), gesture recognition (Wu et al., 2013), etc.

Moreover, multi-view learning has been theoretically studied mainly under the semi-supervised setting, but only with two facing views (Chapelle et al., 2006; Sun, 2013; Sun & Taylor, 2014; Johnson & Zhang, 2015). In parallel, ensemble-based learning approaches have been theoretically studied, in the supervised setting: many interesting results should concern multi-view learning, as long as the ensemble is built upon many views (Rokach, 2010; Zhang & Zhang, 2011). From the representation learning point of view, recent models are based on the incorporation of some "fusion" layers in the deep neural network architecture as in (Ngiam et al., 2011) or (Srivastava & Salakhutdinov, 2012) for example. Some other interesting models include the multiview perceptron (Zhu et al., 2014).

Estimating Complex Distributions: While deep learning has shown great results in many classification task for a decade, training deep generative models still remains a challenge. Deep Boltzmann Machines (Salakhutdinov & Hinton, 2009) are un-directed graphical models organized in a succession of layers of hidden variables. In a multi-view setting, they are able to deal with missing views and have been used to capture the joint distribution in bi-modal text and image data (Srivastava & Salakhutdinov, 2012; Sohn et al., 2014). Another trend started with denoising autoencoder (Vincent et al., 2008), which aims to reconstruct a data from a noisy input have been proved to possess some desirable properties for data generation (Bengio et al., 2013). The model have been generalized under the name Generative Stochastic Networks by replacing the noise function C with a mapping to a latent space (Thibodeau-Laufer et al., 2014). Pulling away from the mixing problems encountered in previous approaches, Variational Autoencoders (Kingma & Welling, 2013) attempts to map the input distribution to a latent distribution which is easy to sample from. The model is trained by optimizing a variational bound on the likelihood, using stochastic gradient descent methods. The Kullback-Leibler regularizer on the latent Gaussian representations used in our model is reminiscent of the one introduced in the variational lower bound used by the VAE.

The BiGAN model (Donahue et al., 2016; Dumoulin et al., 2016) that serves as a basis for our work is an extension of the Generative Adversarial Nets (Goodfellow et al., 2014). A GAN extension that captures conditional probabilities (CGAN) has been proposed in (Mirza & Osindero, 2014). However, as noted by (Mathieu et al., 2015) and (Pathak et al., 2016), they display very unstable behavior. More specifically, CGAN have been able to generate image of faces conditioned on an attribute vector (Gauthier, 2014), but fail to model image distribution conditioned on a part of the image or on previous frames. In both CGAN and CVBiGAN, the generation process uses random noise to be able to generate a diversity of outputs from the same input. However, in a CGAN, the generator concatenate an independent random vector to the input while CV-BiGAN learns a stochastic latent representation of the input. Also, some of the difficulties of CGAN in handling images as both inputs \tilde{x} and outputs \tilde{y} stem from the fact that CGAN's discriminator directly compares \tilde{x} and y . In CV-BiGAN, neither discriminators has access to both \tilde{x} and y but only to a latent representation z and either \tilde{x} or y .

7 CONCLUSION AND PERSPECTIVES

We have proposed the CV-BiGAN model for estimating conditional densities, and its extension MV-BiGAN to handle multi-view inputs. The MV-BiGAN model is able to both handle subsets of views, but also to update its prediction when new views are added. It is based on the idea that the uncertainty of the prediction must decrease when additional information is provided, this idea being handled through a KL constraint in the latent space. This work opens different research directions: the first one concerns the architecture of the model itself since the convergence of MV-BiGAN is still difficult to obtain and has a particularly high training cost. Another direction would be to see if this family of model could be used on data streams for anytime prediction.

ACKNOWLEDGMENTS

This work was supported by the French project LIVES ANR-15-CE23-0026-03.

REFERENCES

- R. Arora and K. Livescu. Kernel cca for multi-view acoustic feature learning using articulatory measurements. In *MLSP*, 2012.
- Pradeep K. Atrey, M. Anwar Hossain, Abdulmotaleb El Saddik, and Mohan S. Kankanhalli. Multimodal fusion for multimedia analysis: A survey. *Multimedia Syst.*, 16(6):345–379, November 2010. ISSN 0942-4962.
- Yoshua Bengio, Li Yao, Guillaume Alain, and Pascal Vincent. Generalized denoising auto-encoders as generative models. In *Advances in Neural Information Processing Systems*, pp. 899–907, 2013.
- Nicolò Cesa-Bianchi, David R. Hardoon, and Gayle Leen. Guest editorial: Learning from multiple sources. *Machine Learning*, 79(1):1–3, 2010.
- O. Chapelle, B. Schölkopf, and A. Zien. *Semi-supervised Learning*. MIT Press, Cambridge, MA, 2006.
- Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in neural information processing systems*, pp. 1486–1494, 2015.
- Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016.
- Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Alex Lamb, Martin Arjovsky, Olivier Mastropietro, and Aaron Courville. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016.
- Jon Gauthier. Conditional generative adversarial nets for convolutional face generation. *Class Project for Stanford CS231N: Convolutional Neural Networks for Visual Recognition, Winter semester*, 2014, 2014.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.
- R. Johnson and T. Zhang. Semi-supervised learning with multi-view embedding: theory and application with convolutional neural networks. *CoRR*, abs/1504.01255v1, 2015.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, number 2014, 2013.
- Sokol Koço, Cécile Capponi, and Frédéric Béchet. Applying multiview learning algorithms to human-human conversation classification. In *INTERSPEECH*, 2012.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015.
- Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multi-modal deep learning. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, 2011.
- Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. *arXiv preprint arXiv:1604.07379*, 2016.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- L. Rokach. Ensemble-based classifiers. *Artif. Intell. Rev.*, 33(1-2):1–39, 2010.

- Ruslan Salakhutdinov and Geoffrey E Hinton. Deep boltzmann machines. In *AISTATS*, volume 1, pp. 3, 2009.
- Kihyuk Sohn, Wenling Shang, and Honglak Lee. Improved multimodal deep learning with variation of information. In *Advances in Neural Information Processing Systems*, pp. 2141–2149, 2014.
- Artem Sokolov and Asa Ben-Hur. Multi-view prediction of protein function. In *ACM-BCB*, pp. 135–142, 2011.
- Nitish Srivastava and Ruslan R Salakhutdinov. Multimodal learning with deep boltzmann machines. In *Advances in neural information processing systems*, pp. 2222–2230, 2012.
- Shiliang Sun. A survey of multi-view machine learning. *Neural Comput. Appl.*, 23(7-8):2031–2038, 2013.
- Shiliang Sun and John-Shawe Taylor. PAC-Bayes analysis of multi-view learning. *CoRR*, abs/1406.5614, 2014.
- Eric Thibodeau-Laufer, Guillaume Alain, and Jason Yosinski. Deep generative stochastic networks trainable by backprop. 2014.
- Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pp. 1096–1103. ACM, 2008.
- Weiran Wang, Raman Arora, Karen Livescu, and Jeff Bilmes. On deep multi-view representation learning. In *Proc. of the 32st Int. Conf. Machine Learning (ICML 2015)*, pp. 1083–1092, 2015.
- Jiaxiang Wu, Jian Cheng, Chaoyang Zhao, and Hanqing Lu. Fusing multi-modal features for gesture recognition. In *ICMI*, pp. 453–460, 2013.
- L. Wu, S.L. Oviatt, and P.R. Cohen. Multimodal integration: a statistical view. *MM*, 1(4):334–341, 1999.
- Jianchun Zhang and Daoqiang Zhang. A novel ensemble construction method for multi-view data using random cross-view correlation between within-class examples. *Pattern Recogn.*, 44(6): 1162–1171, 2011.
- Zhenyao Zhu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Multi-view perceptron: a deep model for learning face identity and view representations. In *Advances in Neural Information Processing Systems*, pp. 217–225, 2014.

APPENDIX

MODEL ARCHITECTURE AND HYPERPARAMETERS

Operation	Kernel	Strides	Padding	Feature maps	BN	Nonlinearity
Convolution	4×4	2×2	1×1	64	\times	Leaky ReLU
Convolution	4×4	2×2	1×1	128	\checkmark	Leaky ReLU
Convolution	4×4	2×2	1×1	256	\checkmark	Leaky ReLU
Convolution	4×4	2×2	1×1	512	\checkmark	Leaky ReLU
Convolution	4×4	1×1		output size	\times	Linear
Fractionally-strided convolution	4×4	1×1		512	\checkmark	ReLU
Fractionally-strided convolution	4×4	2×2	1×1	256	\checkmark	ReLU
Fractionally-strided convolution	4×4	2×2	1×1	128	\checkmark	ReLU
Fractionally-strided convolution	4×4	2×2	1×1	64	\checkmark	ReLU
Fractionally-strided convolution	4×4	2×2	1×1	3	\times	Tanh

Table 1: Convolution architectures used in our experiments on the CelebA dataset. The top part is used for encoding images to the aggregation space. The bottom part is used in G to generate images from a vector z .

All models are optimized using Adam with hyperparameters $\alpha = 2 \cdot 10^{-5}$, $\beta_1 = 0.5$, $\beta_2 = 10^{-3}$ and a learning rate of $2 \cdot 10^{-5}$. The slope of leaky rectified linear units is set to 0.2. Latent representations $(\mu, \log(\sigma^2))$ are of size 128.

MNIST EXPERIMENTS

The generator function G has three fully connected layers. The second and the third are followed by parameter free batch normalizations. All layers uses a rectified linear unit (ReLU) non linearity except the output layer that has a sigmoid.

The aggregation model Ψ is the sum of mapping functions ϕ_k that map each input views to the aggregation space. Each ϕ_k is a linear transformation of the inputs. The encoding functions E and H are both neural networks that include an aggregation model and adds two hidden fully connected layers. They takes a set of views as inputs, and outputs a pairs of vectors $(\mu, \log(\sigma^2))$. The first layer is followed by a ReLU non linearity. The second layer is followed by a parameter free batch normalization and a ReLU. The output layer has a tanh layer for μ and a negative exponential linear unit (nELU) for $\log \sigma^2$ to help stabilizing σ at the beginning of the training procedure.

The discriminator D_1 has three fully connected layers with parameter free batch normalization at the third layer, Leaky ReLU non linearity at the first three layers and a sigmoid for the output layer. The vector z is concatenated to the representation at the second layer.

The discriminator D_2 is similar to E and H except it uses Leaky ReLUs instead of ReLUs and a sigmoid at the output level. z is concatenated directly to the aggregation vector $\Psi(v(x, s))$.

All hidden layers and the aggregation space are of size 1500. λ is set to $1 \cdot 10^{-5}$. Minibatch size is set to 128. The models have been trained for 300 epochs.

CELEBA EXPERIMENTS

The generator function G is a network of fractionally-strided convolution layers (often called deconvolution) discribed in table 7.

The mapping function ϕ_k for image views is a convolutional network (Table 7). For attribute vectors, it is a simple linear transformation. E and H builds a neural networks with one hidden layer on top of the aggregation space. The hidden layer has a parameter free batch normalization followed by a ReLU. The output layer is the same as the encoders used in the MNIST experiments.

The discriminator D_1 has two fully connected layers with a Leaky ReLU after the first one, and a sigmoid at the output, build on top of a fractionally-strided convolutions network. z is concatenated at the output level of the convolutional network.

As in the MNIST experiments, the discriminator D_2 is similar to E and H except it uses Leaky ReLUs instead of ReLUs and a sigmoid at the output level. z is concatenated directly to the aggregation vector $\Psi(v(x, s))$.

Aggregation space is of size 1000. λ is set to $1 \cdot 10^{-3}$, and minibatch size is 16. The model has been trained for 15 epochs.

ADDITIONAL RESULTS




Figure 7: Additional results on the MNIST in four parts task.




Figure 8: Additional results on the MNIST in streams task.




Figure 9: Additional results on the CelebA dataset.