

Adding noise to the input of a model trained with a regularized objective

Salah Rifai, Xavier Glorot, Yoshua Bengio and Pascal Vincent

April 20, 2011

Dept. IRO, Université de Montréal. Montréal (QC), H3C 3J7, Canada

Technical report 1359

Abstract

Regularization is a well studied problem in the context of neural networks. It is usually used to improve the generalization performance when the number of input samples is relatively small or heavily contaminated with noise. The regularization of a parametric model can be achieved in different manners some of which are early stopping (Morgan and Bourlard, 1990), weight decay, output smoothing that are used to avoid overfitting during the training of the considered model. From a Bayesian point of view, many regularization techniques correspond to imposing certain prior distributions on model parameters (Krogh and Hertz, 1991). Using Bishop's approximation (Bishop, 1995) of the objective function when a restricted type of noise is added to the input of a parametric function, we derive the higher order terms of the Taylor expansion and analyze the coefficients of the regularization terms induced by the noisy input. In particular we study the effect of penalizing the Hessian of the mapping function with respect to the input in terms of generalization performance. We also show how we can control independently this coefficient by explicitly penalizing the Jacobian of the mapping function on corrupted inputs.

1 Introduction

Regularization is a well studied problem in the context of neural networks. It is usually used to improve the generalization performance when the number of input samples is relatively small or heavily contaminated with noise. The regularization of a parametric model can be achieved in different manners some of which are early stopping (Morgan and Bourlard, 1990), weight decay or output smoothing, and are used to avoid overfitting during the training of the considered model. From a Bayesian point of view, many regularization techniques correspond to imposing certain prior distributions on model parameters (Krogh and Hertz, 1991).

In this paper we propose a novel approach to achieve regularization that combines noise in the input and explicit output smoothing by regularizing the L2-norm of the Jacobian's mapping function with respect to the input. Bishop (1995) has proved that the

two approaches are essentially equivalent under some assumptions using a Taylor approximation up to the second order of the noisy objective function. Using his theoretical analysis, we derive the approximation of our cost function in the weak noise limit and show the advantage of our technique from the theoretical and empirical point of view. In particular, we show that we achieve a better smoothing of the output of the considered model with a little computational overhead.

2 Definitions

For the ease of readability, most of our analysis involves only vectors and matrices except for section 6 for which it was not possible to avoid using tensors objects. Also our analysis assumes that the model's output is scalar which will prevent the use of tensors for the low order terms of the Taylor expansion. We will use the following notations:

- $\langle \cdot, \cdot \rangle$: inner product,
- \otimes : tensor product,
- $J_f(x), H_f(x), T_f^{(n)}(x)$: respectively the Jacobian, Hessian and n -th order derivative of f with respect to vector x .

We consider the following set of points:

$$\mathcal{D}_n = \left\{ z_i = (x_i, y_i) \in (\mathbb{R}^d, \mathbb{R}) \mid \forall i \in \llbracket 1; n \rrbracket \right\}$$

where the (x_i, y_i) are the (input, target) of an arbitrary dataset. In the paper we will consider a particular family of parametric models

$$\mathcal{F} = \left\{ F_\theta \in C^\infty(\mathbb{R}^d) \mid \theta \in \mathbb{R}^p \mid p \in \mathbb{N} \right\}$$

and $F_\theta(\mathbb{R}^d) \subseteq \mathbb{R}$. We define the expected cost of the true distribution $p(z)$ of our data points as being:

$$\mathcal{C}(\theta) = \int \mathcal{L}(z, \theta) p(z) dz \tag{1}$$

The expected empirical cost when using the data without noise can be expressed as:

$$\mathcal{C}^{\text{clean}}(\theta) = \int \mathcal{L}(z, \theta) \delta(z_i - z) dz = \frac{1}{n} \sum_i^n \mathcal{L}(z_i, \theta) \tag{2}$$

where δ is the Dirac function. When adding noise to the input, we will consider $p(z)$ as being a parzen density estimator:

$$p(z) = \frac{1}{n} \sum_i^n \psi(z_i - z) \tag{3}$$

with the kernels ψ centered on the points of our dataset. In the rest of this paper we will consider kernels for which the following assumptions hold:

- (a) every kernel has zero mean,
- (b) different components of a kernel are independent.

Note that the normal and uniform distribution have these properties. Using (a), we can write (An, 1996):

$$\forall i, \int \varepsilon_i \psi(\varepsilon) d\varepsilon = 0 \quad (4)$$

and using (b) :

$$\forall (i, j), \int \varepsilon_i \varepsilon_j \psi(\varepsilon) d\varepsilon = \sigma^2 \delta_{ij} \quad (5)$$

where σ is the variance of the distribution ϕ , and $\varepsilon = (\varepsilon_1, \dots, \varepsilon_d)$. In our analysis we will restrict ourselves to gaussian kernels:

$$\psi(z_i - z) = \mathcal{N}_{z_i, \sigma^2}(z) \quad (6)$$

Substituting (6) in (3) we can write the objective function with noisy input as being:

$$\mathcal{C}^{\text{noisy}}(\theta) = \frac{1}{n} \sum_i^n \int \mathcal{L}(z, \theta) \mathcal{N}_{z_i, \sigma^2}(z) dz \quad (7)$$

3 Penalty term induced by noisy input

Bishop (1995) already showed that tuning the parameters of a model with corrupted inputs is asymptotically equivalent to minimizing the true error function when simultaneously decreasing the level of corruption to zero as the number of corrupted inputs tends to infinity. He also showed that adding noise in the input is equivalent to minimize a different objective function that includes one or more penalty terms, and he uses a Taylor expansion to derive an analytic approximation of the noise induced penalty. Using the above assumption we can write:

$$\mathcal{C}^{\text{noisy}}(\theta) = \mathcal{C}^{\text{clean}}(\theta) + \phi(\theta) \quad (8)$$

where ϕ is the penalty term. Substituting (7) and (2) in (8) we express the penalty term as being:

$$\phi(\theta) = \frac{1}{n} \sum_i^n \left[\int \mathcal{L}(z, \theta) \mathcal{N}_{z_i, \sigma^2}(z) dz - \mathcal{L}(z_i, \theta) \right] \quad (9)$$

We define the noise vector as being $\varepsilon = z - z_i$ and omitting θ for simplicity we can write $\forall i$, the term inside the sum of (9):

$$D = \int \mathcal{L}(z_i + \varepsilon) \mathcal{N}_{0, \sigma^2}(\varepsilon) d\varepsilon - \mathcal{L}(z_i) \quad (10)$$

Now that we have identified the term to approximate, let's write the Taylor approximation of our loss function when our sample is shifted by a noise vector ε :

$$\mathcal{L}(z + \varepsilon) = \mathcal{L}(z) + \langle J_{\mathcal{L}}(z), \varepsilon \rangle + \frac{1}{2} \varepsilon^T \cdot H_{\mathcal{L}}(z) \cdot \varepsilon + o(\varepsilon^2) \quad (11)$$

To match equation (10) we multiply with (6) and integrate with respect to ε both sides of (11):

$$\begin{aligned} & \int \mathcal{L}(z + \varepsilon) \mathcal{N}_{0, \sigma^2}(\varepsilon) d\varepsilon = \\ & \int \left[\mathcal{L}(z) + \langle J_{\mathcal{L}}(z), \varepsilon \rangle + \frac{1}{2} \varepsilon^T \cdot H_{\mathcal{L}}(z) \cdot \varepsilon + o(\varepsilon) \right] \mathcal{N}_{0, \sigma^2}(\varepsilon) d\varepsilon \end{aligned} \quad (12)$$

Equation (4) implies that all odd-moments of the approximation are null and in conjunction with (5) we can now simplify (12) into:

$$\begin{aligned} & \int \mathcal{L}(z + \varepsilon) \mathcal{N}_{0, \sigma^2}(\varepsilon) d\varepsilon \\ & = \mathcal{L}(z) \underbrace{\int \mathcal{N}_{0, \sigma^2}(\varepsilon) d\varepsilon}_1 + \underbrace{\int \langle J_{\mathcal{L}}(z), \varepsilon \rangle \mathcal{N}_{0, \sigma^2}(\varepsilon) d\varepsilon}_0 \\ & \quad + \frac{1}{2} \int \varepsilon^T \cdot H_{\mathcal{L}}(z) \cdot \varepsilon \mathcal{N}_{0, \sigma^2}(\varepsilon) d\varepsilon + R \end{aligned} \quad (13)$$

and with some algebra we can finally write:

$$\int \mathcal{L}(z + \varepsilon) \mathcal{N}_{0, \sigma^2}(\varepsilon) d\varepsilon - \mathcal{L}(z) \approx \frac{\sigma^2}{2} \text{Tr}(H_{\mathcal{L}}(z)) \quad (14)$$

by substituting $z = z_i$ and summing over all the elements of \mathcal{D}_n :

$$\phi(\theta) \approx \frac{\sigma^2}{2n} \sum_i^n \text{Tr}(H_{\mathcal{L}}(z_i)) \quad (15)$$

Hence training with corrupted inputs is approximately equivalent to minimize the following objective:

$$\boxed{\mathcal{C}^{\text{noise}}(\theta) \approx \mathcal{C}^{\text{clean}}(\theta) + \lambda \text{Tr}(H_{\bar{\mathcal{C}}}(\theta))} \quad (16)$$

This relation holds for any objective \mathcal{C} with the above two assumptions.

All the above results are already well established in the literature of noise injection (Bishop, 1995; An, 1996). Our contribution is to show that, for the second order Taylor expansion, adding noise to the input on a well chosen regularized objective is equivalent to add

a L_2 -norm on the Hessian of the output function of the considered model respectively to its input, which is not the case when adding noise to a non-regularized objective. In the following sections we will consider the MSE as the objective function to tune the parameters of our model but this choice does not affect the generality of our analysis.

4 Non regularized objective

We define the following error function:

$$\mathcal{L}^{\text{clean}}(z_i, \theta) = \| F(x_i, \theta) - y_i \|^2$$

and the Hessian of the cost function being the average over the Hessian of all individual errors:

$$H_{\mathcal{L}^{\text{clean}}}(\theta) = \frac{1}{n} \sum_i^n H_{\mathcal{L}^{\text{clean}}}(z_i, \theta)$$

Assuming without loss of generality that F is a scalar function¹, and that the noise is added to the input x , we will only consider the Hessian of the loss function with respect to x :

$$\begin{aligned} H_{\mathcal{L}^{\text{clean}}}(x) &= \frac{\partial}{\partial x} \left[\frac{\partial}{\partial x} (\| F(x, \theta) - y \|^2) \right] \\ &= 2 \frac{\partial}{\partial x} \left[\frac{\partial F(x, \theta)}{\partial x} (F(x, \theta) - y) \right] \\ &= 2 \left[\frac{\partial^2 F(x, \theta)}{(\partial x)^2} (F(x, \theta) - y) + \left\langle \frac{\partial F(x, \theta)}{\partial x} \right\rangle^2 \right] \\ H_{\mathcal{L}^{\text{clean}}}(x) &= 2H_F(x) (F(x, \theta) - y) + 2 \langle J_F(x)^T, J_F(x) \rangle \end{aligned}$$

A standard result of linear algebra is the relation between the Frobenius norm and the trace operator:

$$\text{Tr}(\langle A^T, A \rangle) = \|A\|_F^2$$

By taking the trace of the above results, we get:

$$\text{Tr}(H_{\mathcal{L}^{\text{clean}}}(x)) = 2(F(x, \theta) - y) \text{Tr}(H_F(x)) + 2\|J_F(x)\|_F^2$$

and plugging this in (16) gives us the following second order approximation of the noisy objective:

$$\boxed{\mathcal{L}^{\text{noise}}(z_i, \theta) \approx \| F(x_i, \theta) - y_i \|^2 + 2\lambda \left((F(x, \theta) - y) \text{Tr}(H_F(x)) + \|J_F(x)\|_F^2 \right)} \quad (17)$$

¹in the case of multiple regression such as a reconstruction function, only the orders of the tensors involved in the approximation of \mathcal{L} will change.

We obtain a L_2 -norm on the gradient of the mapping function F added to our error function whereas the Hessian term is not constrained to be positive and it is not sure that its terms are going to cancel-out. E.g. if prediction overshoots or undershoots on average, then penalty may encourage very large Hessian trace inducing high curvature which would potentially harm a stochastic gradient descent and converge to a poor local minimum.

5 Our regularized objective

As the results of the previous section suggest, adding noise to the input of the objective function yields an undesirable term that might interfere with the goal of smoothing the output of our function. We propose here to overcome this difficulty by adding noise only to the input of the objective function's Jacobian J_F , doing so will avoid the unpredictable effect of additional unwanted terms. We define the error function we previously used to which we add a regularization term that is the L_2 -norm of the Jacobian of F with respect to x :

$$\mathcal{L}^{\text{reg+noise}}(z_i, \theta) = \| F(x_i, \theta) - y_i \|^2 + \lambda \| J_F(\tilde{x}_i) \|_F^2$$

We now calculate the Hessian of $\tilde{\mathcal{L}}$ by omitting the first term since the noise is added only to the input of the regularization term, and using the approximation derived in (16):

$$\left\| \frac{\partial F(\tilde{x}_i, \theta)}{\partial x} \right\|_F^2 \approx \left\| \frac{\partial F(x_i, \theta)}{\partial x} \right\|_F^2 + \varepsilon^2 \text{Tr} \left(H_{\| \frac{\partial F}{\partial x} \|^2}(x) \right)$$

We can now calculate the Hessian of the regularization term as being:

$$\begin{aligned} H_{\| \frac{\partial F}{\partial x} \|^2}(x) &= 2 \frac{\partial}{\partial x} \left[\frac{\partial^2 F(x, \theta)}{(\partial x)^2} \frac{\partial F(x, \theta)}{\partial x} \right] \\ &= 2 \left[\frac{\partial^3 F(x, \theta)}{(\partial x)^3} \frac{\partial F(x, \theta)}{\partial x} + \left\langle \frac{\partial^2 F(x, \theta)}{(\partial x)^2} \right\rangle^2 \right] \\ H_{\| \frac{\partial F}{\partial x} \|^2}(x) &= 2 \frac{\partial^3 F(x, \theta)}{(\partial x)^3} J_F(x) + 2 \langle H_F(x)^T, H_F(x) \rangle \\ \text{Tr} \left(H_{\| \frac{\partial F}{\partial x} \|^2}(x) \right) &= 2 \text{Tr} \left(\frac{\partial^3 F(x, \theta)}{(\partial x)^3} J_F(x) \right) + 2 \| H_F(x) \|_F^2 \end{aligned} \quad (18)$$

which gives us the approximation of our regularized objective:

$$\boxed{\mathcal{L}^{\text{reg+noise}}(z_i, \theta) = \| F(x_i, \theta) - y_i \|^2 + \lambda \| J_F(x) \|_F^2 + 2\lambda\sigma^2 \| H_F(x) \|_F^2 + R} \quad (19)$$

We have shown that adding noise to a well chosen regularized objective clearly penalize the L_2 -norm on the Hessian of the considered model F (without ever calculating it) using

a second order Taylor approximation of the noisy objective under two necessary assumptions on the noise distribution. In statistics regularizing the norm of the derivatives of the model to be tuned is often referred as roughness penalty (Green, 1993) and is used in the context of cubic splines (De Boor, 1998).

6 Higher order terms of the Taylor expansion

In this section we are interested in the higher order terms of the cost approximation, we find it convenient to use the following formalism:

if $T_{\mathcal{L}}^n(z, \varepsilon)$ denotes the n-th order derivative of \mathcal{L} with respect to z , then:

$$T_{\mathcal{L}}^n(z, \varepsilon) = \frac{1}{n!} \sum_{i_1, \dots, i_n} \varepsilon_{i_1}, \dots, \varepsilon_{i_n} T_{i_1, \dots, i_n}^n(z)$$

where T^n is a tensor of order n and

$$T_{i_1, \dots, i_n}^n(z) = \frac{\partial^n \mathcal{L}(z)}{\partial z_{i_1}, \dots, \partial z_{i_n}}$$

using this formalism we can write the fourth order derivative as being:

$$T_{\mathcal{L}}^4(z, \varepsilon) = \frac{1}{24} \sum_{i, j, k, l} \varepsilon_i \varepsilon_j \varepsilon_k \varepsilon_l T_{ijkl}^4(z)$$

Using the two assumptions made on the noise distribution, we know that the third order derivative of the approximation is zero. As for the fourth order derivative, using the second assumption of the noise distribution we know that only the terms that are on the diagonal of the T^4 will be non-zero, we can then write:

$$\int T_{\mathcal{L}}^4(z, \varepsilon) \mathcal{N}_{0, \sigma^2}(\varepsilon) d\varepsilon = \frac{\sigma^4}{4!} \sum_i T_{iii}^4(z)$$

Using the above result we can approximate our cost function in the noisy input setting more finely, for this purpose we will use the results obtained above for the Hessian and differentiate them again twice with respect to x .

$$\begin{aligned} T_{\mathcal{L}}^4(x) &= \frac{\partial^2 (H_{\mathcal{L}}(x))}{(\partial x)^2} \\ &= \frac{\partial^2}{(\partial x)^2} [2H_F(x) (F(x, \theta) - y) + 2 \langle J_F(x)^T, J_F(x) \rangle] \\ &= 2 \frac{\partial}{\partial x} [T_F^3(x) (F(x, \theta) - y) + 3H_F(x) J_F(x)] \\ &= 2 [T_F^4(x) (F(x, \theta) - y) + 4T_F^3(x) J_F(x) + 3 \langle H_F(x)^T, H_F(x) \rangle] \end{aligned}$$

hence,

$$\text{Tr}(T_{\mathcal{L}}^4(z)) = 6 \|H_F(x)\|_F^2 + \text{Tr}(R)$$

where $R = 2T_F^4(x)(F(x, \theta) - y) + 8T_F^3(x)J_F(x)$.

7 Comparison

7.1 Noise added to the input of the objective function

Now that we have a higher order approximation in the case where noise is added to the input of the function, we can compare the magnitude of the coefficients that penalize the Hessian H_F , note that in this case the Hessian term appears in the fourth order of the Taylor expansion of the cost function, whereas we need only a second order approximation in the case where we add a regularization term evaluated on a corrupted input. We can write the approximation of the cost function without regularization at the fourth order as being:

$$\mathcal{L}^{\text{noise}}(z) \approx \mathcal{L}^{\text{clean}}(z) + \frac{\sigma^2}{2!} \text{Tr}(H_{\mathcal{L}^{\text{clean}}}(z)) + \frac{\sigma^4}{4!} \text{Tr}(T_{\mathcal{L}^{\text{clean}}}^4(z))$$

$$\mathcal{L}^{\text{noise}}(z) = \|F(x_i, \theta) - y_i\|^2 + \overline{\sigma^2 \|J_F(x)\|_F^2} + \frac{\sigma^4}{4} \overline{\|H_F(x)\|_F^2} + \overline{R} \quad (20)$$

where the number i of overline denotes the terms induced by the noise obtained at the i -th order of the Taylor expansion.

7.2 Noise added to the input of the Jacobian of the objective function

In this case we just need to approximate the cost function up to the second order of the Taylor expansion:

$$\mathcal{L}^{\text{reg+noise}}(z_i, \theta) = \|F(x_i, \theta) - y_i\|^2 + \lambda \|J_F(x)\|_F^2 + \overline{2\lambda\sigma^2 \|H_F(x)\|_F^2} + \overline{R} \quad (21)$$

8 Experimental results

We have tried several experiments in order to benchmark the effect of regularization and noise combined, for this task we used the well known MNIST (LeCun *et al.*, 1998), MNIST binarised and the USPS database. Surprisingly, we were able to achieve results close to those obtained with unsupervised pretraining (Erhan *et al.*, 2010). MNIST is composed of 70k handwritten digits represented by a vector of pixels. It is divided in 50k for the training set, 10k for each of the validation and test set, the range of the features were rescaled to be within $[0, 1]$. MNIST-Binary is divided exactly the same way as

MNIST, the only difference is that the intensity of the pixels superior to $\frac{255}{2}$ were set to 1 and the others to 0. USPS dataset consists of a training set with 7291 images and a test set with 2007 images, the validation set was formed using the last 2000 images of the training set. The model $F(x)$, with parameters $\theta = \{W^{(1)}, \dots, W^{(n)}, b^{(1)}, \dots, b^{(n)}\}$, we considered to solve the classification task was a standard neural network with one or more hidden layers, and a hyperbolic tangent non-linearity in between the layers. For example, with one hidden layer we have:

$$F(x) = \sigma(W^{(2)} \tanh(W^{(1)} \cdot x + b^{(1)}) + b^{(2)})$$

where $\sigma(\cdot)$ is the logistic sigmoid function. In this setting we can write the Frobenius norm of the Jacobian of F with respect to x as being:

$$\|J_F(x)\|^2 = \sum_{i,j} J_{ij}^2$$

and with some calculus we get:

$$J_{ij} = F_i(x)(1 - F_i(x)) \sum_l W_{il}^{(2)} W_{lj}^{(1)} \left(1 - \tanh^2 \left(\sum_m W_{lm}^{(1)} x_m + b_l^{(1)} \right) \right)$$

where $F_i(x)$ is i -th output of the network. For the results in table (1), we used a number of hidden units ranging from 400 to 1000 per layer, the best results were obtained with two hidden layers on MNIST, and one hidden layer on MNIST-BINARY and USPS. The parameters of the model were learned through stochastic gradient descent with a learning rate ranging from 0.1 to 0.001. We also investigated the use of *Rectifying* units (i.e. $\max(0, x)$) (Nair and Hinton, 2010; Glorot *et al.*, 2010) as non-linearity in the hidden layer, surprisingly they seemed to benefit less from the added noise to the input than from the regularization term alone, they achieved a test classification performance of 4.8% on the USPS dataset equaling the best performance of the hyperbolic tangent activation with both regularization and added noise to the input. The best results were obtained with a gaussian isotropic noise with a standard deviation of 0.1 around training samples.

Figure 1 shows the histograms of activation values on the MNIST test set of our best MLPs with and without Jacobian regularization. The activations of the regularized model are more densely distributed at the saturation and linear regime.

9 Discussion

9.1 Constraining the solution space

When optimizing a non-convex function with an infinite amount of local minimum, it is not clear which of them yields a reasonable generalization performance, the concept of overfitting clearly illustrate this point. The proposed regularizer tries to avoid this scheme by flattening the mapping function over the training examples inducing a local

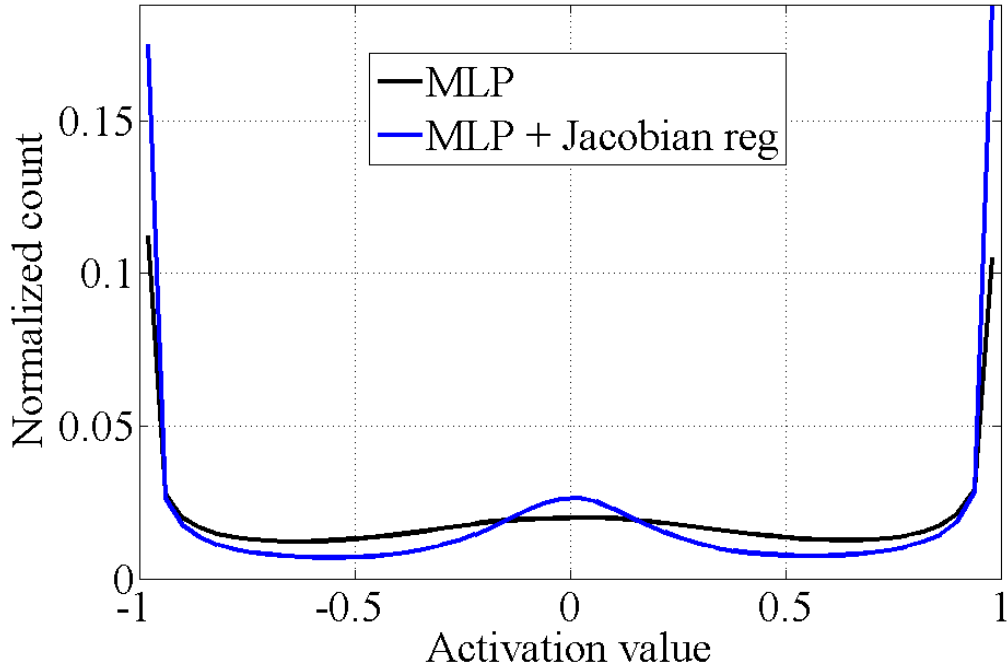


Figure 1: Normalized histograms of the activation values on the MNIST test set for the best MLPs with and without Jacobian regularization. The activations of the regularized model are more densely distributed at the saturation and linear regime.

invariance in the mapping space to infinitesimal variance in the input space. Figure 2 shows that when the input is corrupted, the models learned with the regularization term are more robust to noise on the input. Geometrically, the added regularization imposes to the model to be a *Lipschitz* function or a contracting map around the training examples imposing the constraint $F(x + \varepsilon) \approx F(x)$.

9.2 Smoothing away from the training points

Penalizing only the Jacobian of the model F with respect to the input x gives only a guarantee of flatness for an infinitesimal move away of the training example. To illustrate this point one can imagine that on the tip of a dirac function the norm of the jacobian is null and infinite just around it. Although this situation is not possible in the context of neural networks because of their smooth activation function. Given enough capacity we could converge to this solution if we do not add additional constraints. One of them would be to adjust the locality of the flatness as a hyper-parameter of the model. It requires to compute the higher order terms of the mapping function in order to regularize the magnitude of their norms. As it was discussed before, it is computationally expensive

Table 1: Test error

Model	MNIST	MNIST-BINARY	USPS
MLP	1.82%	2.01%	5.74%
MLP + L_2	1.66%	2.20%	5.6%
MLP + Noise	1.33%	1.7%	4.85%
MLP + Jacob	1.31%	1.65%	4.85%
MLP + N + J	1.19%	1.51%	4.6%

to explicitly calculate the norm of the high order derivatives because their number of components increases exponentially, instead our approach proposes an approximation of the Hessian term that allows you to simultaneously control the magnitude of both the Jacobian and hessian norms independently.

9.3 The other terms induced by the noise

As we have seen in the equation 17, the added noise does not yield only in a penalty on the norm of the successive derivatives of the mapping function and it is somehow unclear how these terms behave during gradient descent since they are not constrained to be positive. In a supervised setting, it is empirically feasible to drive those terms to zero because of the small dimensionality of the target points, whereas in a multi-dimensional regression task such as the reconstruction objective of an auto-encoder it is often impossible to achieve a “near” zero minimization of the cost with a first order optimization such as a stochastic gradient descent. The reader should note that the approximation of the noisy cost is valid when the number of corrupted inputs tends to infinity, though in practise this is never the case. It would be interesting to do an estimate of the difference between the terms induced by the noise and the real values of the term in function of the number of corrupted samples.

10 Conclusion

We have shown how to obtain a better generalization performance using a regularization term that adds a marginal computational overhead compared to the traditional approach. Using a Taylor expansion of the cost function, we also showed that by adding noise to the input of the regularization term we are able to penalize with a greater magnitude the norm of the higher order derivatives of the model avoiding the need to explicitly calculate them, which would be obviously computationally prohibitive. Initial results suggests that different parametric models clearly benefit from this approach in terms of predicting out-of-sample points. It would be interesting to investigate how this regularization approach would behave when used with non-parametric models such as gaussian-mixtures.

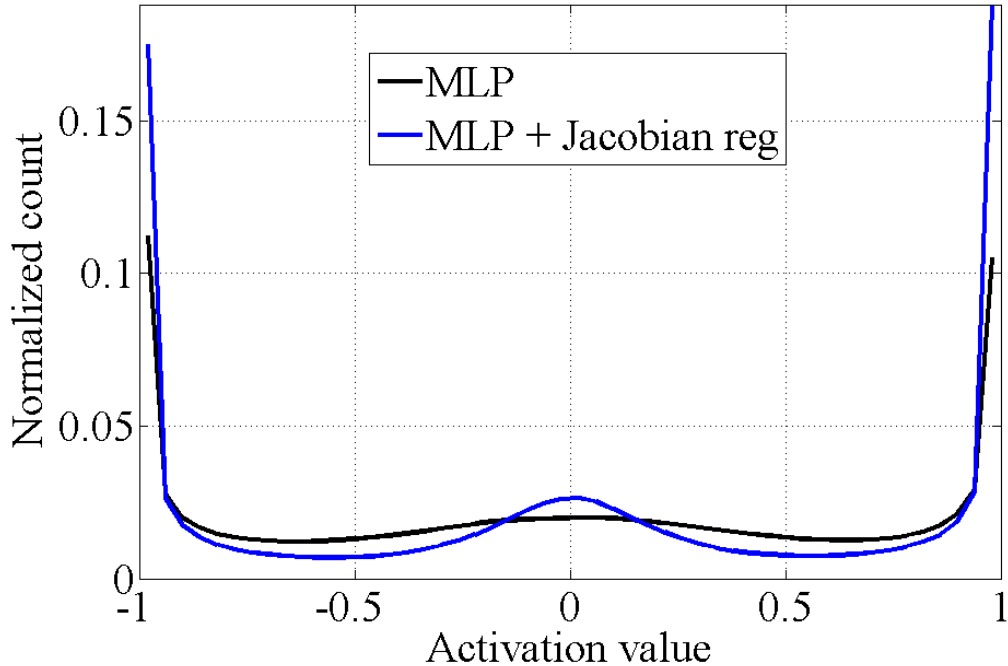


Figure 2: Robustness of the generalization error with respect to a gaussian corruption noise added to the input, the model trained with the combination of input noise and Jacobian regularization is more robust.

Bibliography

- An, G. (1996). The effects of adding noise during backpropagation training on a generalization performance. *Neural Comput.*, **8**, 643–674.
- Bishop, C. M. (1995). Training with noise is equivalent to Tikhonov regularization. *Neural Computation*, **7**(1), 108–116.
- De Boor, C. (1998). Calculation of the smoothing spline with weighted roughness measure. In *Mathematical Models and Methods in Applied Sciences*, page 2001.
- Erhan, D., Bengio, Y., Courville, A., Manzagol, P.-A., Vincent, P., and Bengio, S. (2010). Why does unsupervised pre-training help deep learning? *JMLR*, **11**, 625–660.
- Glorot, X., Bordes, A., and Bengio, Y. (2010). Deep sparse rectifier neural networks. Deep Learning and Unsupervised Feature Learning Workshop — NIPS ’10.
- Green, Bernard W Silverman, B. W. S. (1993). *Nonparametric Regression and Generalized Linear Models*. CRC Press.

- Krogh, A. and Hertz, J. A. (1991). A simple weight decay can improve generalization. In *NIPS*, pages 950–957.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, **86**(11), 2278–2324.
- Morgan, N. and Bourlard, H. (1990). Generalization and parameter estimation in feedforward nets: some experiments. In *NIPS'89*, pages 413–416, Denver, CO. Morgan Kaufmann.
- Nair, V. and Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *ICML '10*.