

Continuous Optimization of Hyper-Parameters

Yoshua Bengio

Département d'informatique et recherche opérationnelle

Université de Montréal

C.P. 6128 Succ. Centre-Ville, Montréal, Québec, Canada, H3C 3J7

bengioy@iro.umontreal.ca

Abstract

Many machine learning algorithms can be formulated as the minimization of a training criterion which involves a hyper-parameter. This hyper-parameter is usually chosen by trial and error with a model selection criterion. In this paper we present a methodology to optimize several hyper-parameters, based on the computation of the gradient of a model selection criterion with respect to the hyper-parameters. In the case of a quadratic training criterion, the gradient of the selection criterion with respect to the hyper-parameters is efficiently computed by back-propagating through a Cholesky decomposition. In the more general case, we show that the implicit function theorem can be used to derive a formula for the hyper-parameter gradient involving second derivatives of the training criterion.

1 Introduction

Many machine learning algorithms can be formulated as the minimization of a **training criterion** which involves a hyper-parameter, kept fixed during this minimization. For example, in the regularization framework [12], one hyper-parameter controls the strength of the penalty term: a larger penalty term reduces the “complexity” of the resulting function. A common example is *weight decay*, used with neural networks and linear regression, to penalize the 2-norm of the parameters. A **model selection criterion** can be used to select hyper-parameters, more generally to compare and choose among models which may have a different capacity. Many model selection criteria have been proposed in the past [13, 1, 7]. When there is only a single hyper-parameter one can easily explore how its value affects the model selection criterion: typically one tries a finite number of values of the hyper-parameter and picks the one which gives the lowest value of the model selection criterion.

In this paper we present a methodology to simultaneously select *many hyper-parameters using the gradient of the model selection criterion with respect to the hyper-parameters*. This methodology can be applied when some differentiability and continuity conditions of the training criterion are satisfied. The use of multiple hyper-parameters has already been proposed in the Bayesian literature: one hyper-parameter per input feature was used to control the prior on the parameters associated to that input feature [9]. In this case, the hyper-parameters can be interpreted as scale parameters for the prior distribution on the parameters, for different directions in parameter space. In Sections 2, 3 and 4, we explain how the gradient with respect to the hyper-parameters can be computed. In the conclusion, we briefly describe the results of preliminary experiments performed with the proposed methodology (described in more details in [?, 4]), and we raise some important open questions concerning the kind of “over-fitting” that can occur with the proposed methodology.

2 Objective Functions for Hyper-Parameters

The methodology presented in this paper can be applied in a variety of settings, but we concentrate here on the familiar supervised machine learning task from *i.i.d* data: we are given a set of independent data points, $D = \{z_1, \dots, z_T\}$, all generated by the same unknown distribution $P(Z)$. We want to choose a function f from a given set of functions \mathcal{F} to minimize the expectation $E_Z(Q(f, Z))$ of a given loss functional $Q(f, Z)$. In supervised learning problems, we have input/output pairs $Z = (X, Y)$, with $X \in \mathcal{X}$, $Y \in \mathcal{Y}$, and $f : \mathcal{X} \rightarrow \mathcal{Y}$. For example, we will consider the case of the quadratic loss, with real-valued vectors $\mathcal{Y} \subseteq \mathcal{R}^m$ and $Q(f, (X, Y)) = \frac{1}{2}(f(X) - Y)'(f(X) - Y)$.

2.1 Training Criteria

In its most general form, a **training criterion** C is any real-valued function of the set of empirical losses $Q(f, z_i)$:

$$C = c(Q(f, z_1), Q(f, z_2), \dots, Q(f, z_T))$$

In the cases in which we are interested, C is parameterized by a vector of real-valued hyper-parameters $\lambda = (\lambda_1, \dots, \lambda_q)$, and the function $f = F(x, \theta) \in \mathcal{F}$ is associated to the parameter vector $\theta \in \Omega \subset \mathcal{R}^s$. We can also rewrite the **training criterion** C in the more compact form $C = c(\theta, \lambda, D)$. The proposed method relies on the assumption that C is continuous and differentiable almost everywhere with respect to θ and λ . When the hyper-parameters are fixed, the learning algorithm attempts to perform the following minimization: $\theta(\lambda, D) = \operatorname{argmin}_{\theta} c(\theta, \lambda, D)$. The corresponding solution is the function $f_{\lambda, D}(x) = F(x, \theta(\lambda, D))$.

An example of training criterion with hyper-parameters is the following:

$$C = \sum_{(x_i, y_i) \in D} w_i(\lambda) (F(x_i, \theta) - y_i)^2 + \theta' A(\lambda) \theta \quad (1)$$

where the hyper-parameters provide different quadratic penalties to different parameters (with the matrix A), and different weights to different training patterns (with $w_i(\lambda)$), (as in [4, ?]).

2.2 Model Selection Criteria

The **model selection criterion** E is a criterion that is used to select hyper-parameters or more generally to choose one model among several models. Ideally, it should be the expected generalization error (for a fixed λ), but $P(Z)$ is unknown, so many alternatives have been proposed, which are either approximations, bounds, or empirical estimates. Most model selection criteria have been proposed for selecting a single hyper-parameter that controls the “complexity” of the class of functions in which the learning algorithms finds a solution, e.g. the minimum description length principle [10], structural risk minimization [13], the Akaike Information Criterion [1], or the generalized cross-validation criterion [7]. Another type of criteria are those based on held-out data, such as the cross-validation estimates of generalization error. These are almost unbiased estimates of generalization error [13] obtained by testing f on data not used to choose f within \mathcal{F} . For example, the K -fold cross-validation estimate uses K partitions of D , $(S_1^1, S_2^1), \dots$ and (S_1^K, S_2^K) , with $S_1^i \cup S_2^i = D$ and $S_1^i \cap S_2^i = \emptyset$:

$$E_{cv}(\lambda, D) = \frac{1}{K} \sum_i \frac{1}{|S_2^i|} \sum_{z_t \in S_2^i} Q(f_{\lambda, S_1^i}, z_t).$$

When f is fixed, the empirical risk $\frac{1}{T} \sum_t Q(f, z_t)$ is an unbiased estimate of the generalization error of f (but it becomes an optimistic estimate when f is chosen within \mathcal{F} to minimize the empirical risk). Similarly, when λ is fixed, the cross-validation criterion is an almost unbiased estimate (when K approaches $|D|$) of the generalization error of $f_{\lambda, D}$. When λ is chosen to minimize the cross-validation criterion, this minimum value also becomes an optimistic estimate. Likewise, when there is a greater diversity of functions $f_{\lambda, D}$ that can be obtained for different values of λ , there is more risk of **overfitting the hyper-parameters**. In this sense, the use of hyper-parameters proposed in this paper can be very different from the common use in which a hyper-parameter helps to control overfitting. Instead, a blind use of the extra freedom brought by many hyper-parameters could deteriorate generalization.

3 Optimizing Hyper-Parameters for a Quadratic Training Criterion

In this section we analyze the simpler case in which the training criterion C is a quadratic polynomial of the parameters θ . The dependence on the hyper-parameters λ can be of higher order, as long as it is continuous and differentiable almost everywhere:

$$C = a(\lambda) + b(\lambda)' \theta + \frac{1}{2} \theta' H(\lambda) \theta \quad (2)$$

where $\theta, b \in \mathcal{R}^s$, $a \in \mathcal{R}$, and $H \in \mathcal{R}^{s \times s}$. For a minimum of (2) to exist requires that H be positive definite. It can be obtained by solving the linear system

$$\frac{\partial C}{\partial \theta} = b + H\theta = 0 \quad (3)$$

which yields the solution

$$\theta(\lambda) = -H^{-1}(\lambda)b(\lambda). \quad (4)$$

The gradient of the model selection criterion E with respect to λ is

$$\left. \frac{\partial E}{\partial \lambda} \right|_{\lambda} = \left. \frac{\partial E}{\partial \theta} \right|_{\theta, \lambda} \left. \frac{\partial \theta}{\partial \lambda} \right|_{\lambda} + \left. \frac{\partial E}{\partial \lambda} \right|_{\theta, \lambda}$$

where we have denoted by $\left. \frac{\partial y}{\partial x} \right|_{x, v}$ the partial derivative of y with respect to x when y is seen as a function of x and v (i.e, keeping x and v fixed, so that even if v is a function x , that dependency is not accounted in that derivative). Later on in this text we will drop the $|_{x, v}$ notation when all the dependencies are taken into account (i.e, this is $\left. \frac{\partial y}{\partial x} \right|_x$). For example, in the case of the cross-validation criteria, $\left. \frac{\partial E_{cv}}{\partial \lambda} \right|_{\theta, \lambda} = 0$, and

$$\frac{\partial E_{cv}}{\partial \theta} = \frac{1}{K} \sum_i \frac{1}{|S_2^i|} \sum_{z_t \in S_2^i} \frac{\partial Q(\theta, z_t)}{\partial \theta}.$$

In the quadratic case, the influence of λ on θ is spelled out by (4), yielding

$$\frac{\partial \theta_i}{\partial \lambda} = - \sum_j \frac{\partial H_{i,j}^{-1}}{\partial \lambda} b_j - \sum_j H_{i,j}^{-1} \frac{\partial b_j}{\partial \lambda} \quad (5)$$

Although the second sum can be readily computed, $\frac{\partial H_{i,j}^{-1}}{\partial \lambda}$ in the first sum is more challenging: we consider several methods below. One solution is based on the computation of gradients through the inverse of a matrix. A general but inefficient solution is the following: $\frac{\partial H_{i,j}^{-1}}{\partial \lambda} = \sum_{k,l} \frac{\partial H_{i,j}^{-1}}{\partial H_{k,l}} \frac{\partial H_{k,l}}{\partial \lambda}$ where

$$\frac{\partial H_{i,j}^{-1}}{\partial H_{k,l}} = -H_{i,j}^{-1} H_{l,k}^{-1} + I_{i \neq l, j \neq k} H_{i,j}^{-1} \text{minor}(H, j, i)_{l', k'}^{-1}, \quad (6)$$

where $\text{minor}(H, j, i)$ denotes the “minor matrix”, obtained by removing the j -th row and the i -th column from H , and the indices (l', k') in the above equation refer to the position within a minor matrix that corresponds to the position (l, k) in H (note $l' \neq i$ and $k' \neq j$). Unfortunately, the computation of this gradient requires $O(s^5)$ multiply-add operations for an $s \times s$ matrix, which is much more than the inversion of H ($O(s^3)$). A better solution is based on the following equality: $HH^{-1} = I$, where I is the $s \times s$ identity matrix. This implies, by differentiating with respect to λ : $\frac{\partial H}{\partial \lambda} H^{-1} + H \frac{\partial H^{-1}}{\partial \lambda} = 0$. Isolating $\frac{\partial H^{-1}}{\partial \lambda}$, we get

$$\frac{\partial H^{-1}}{\partial \lambda} = -H^{-1} \frac{\partial H}{\partial \lambda} H^{-1} \quad (7)$$

which requires only about $2s^3$ multiply-add operations.

An even better solution (which was suggested by Léon Bottou) is to return to equation (3), which can be solved in about $s^3/3$ multiply-add operations (when $\theta \in \mathcal{R}^s$). The idea is to *back-propagate* gradients through each of the operations performed to solve the linear system. The objective is to compute the gradient of E with respect to H and b through the effect of H and b on θ , in order to finally compute $\frac{\partial E}{\partial \lambda}$, as illustrated in figure 1. The back-propagation costs the same as the linear system solution, i.e., $s^3/3$ operations, so this is the approach that we have kept for our implementation. Since H is the Hessian matrix, it is positive definite and symmetric, and (3) can be solved through the Cholesky decomposition of H (assuming H is full rank, which is likely if the hyper-parameters provide some sort of weight decay). The Cholesky decomposition of a symmetric positive definite matrix H gives $H = LL'$ where L is a lower diagonal matrix (with zeros above the diagonal). It is computed in time $O(s^3)$ as follows:

$$\begin{aligned} & \text{for } i = 1, \dots, s \\ & \quad L_{i,i} = \sqrt{H_{i,i} - \sum_{k=1}^{i-1} L_{i,k}^2} \\ & \quad \text{for } j = i + 1, \dots, s \\ & \quad \quad L_{j,i} = (H_{i,j} - \sum_{k=1}^{i-1} L_{i,k} L_{j,k}) / L_{i,i} \end{aligned}$$




Figure 1: Illustration of forward paths (full lines) and gradient paths (dashed) for computation the model selection criterion E and its derivative with respect to the hyper-parameters (λ), when using the method based on the Cholesky decomposition and back-substitution to solve for the parameters (θ).

Once the Cholesky decomposition is achieved, the linear system $LL'\theta = -b$ can be easily solved, in two back-substitution steps: first solve $Lu = -b$, then solve $L'\theta = u$. First step, iterating once forward through the rows of L : for $i = 1, \dots, s$, $u_i = (-b_i - \sum_{k=1}^{i-1} L_{i,k}u_k)/L_{i,i}$. Second step, iterating once backward through the rows of L : for $i = s, \dots, 1$, $\theta_i = (u_i - \sum_{k=i+1}^s L_{k,i}\theta_k)/L_{i,i}$.

The computation of the gradient of θ with respect to the elements of H and b proceed in exactly the reverse order. We start by back-propagating through the back-substitution steps, and then through the Cholesky decomposition. Note that for another application it has already been proposed to differentiate the Cholesky algorithm [11]. See [3] for the details of the gradient computation.

3.1 Weight Decays for Linear Regression

In this subsection, we consider as an example the case of multiple weight decays for linear regression, with K-fold cross-validation as the model selection criterion. The hyper-parameter λ_j will be a weight decay associated to the j -th input variable. The training criterion for the k -th partition is

$$C_k = \frac{1}{|S_1^k|} \sum_{(x_t, y_t) \in S_1^k} \frac{1}{2} (\Theta x_t - y_t)' (\Theta x_t - y_t) + \frac{1}{2} \sum_j \lambda_j \sum_i \Theta_{i,j}^2 \quad (8)$$

The objective is to penalize separately each of the input variables (as in [9]), a kind of “soft variable selection” (see [?] for more discussion and experiments with this setup). The training criterion is quadratic, as in (2), with coefficients

$$a = \frac{1}{2} \sum_t y_t' y_t, \quad b_{(ij)} = - \sum_t y_t x_{t,j}, \quad H_{(ij), (i'j')} = \delta_{i,i'} \sum_t x_{t,j} x_{t,j'} + \delta_{i,i'} \delta_{j,j'} \lambda_j,$$

where $\delta_{i,j} = 1$ when $i = j$ and 0 otherwise, and (ij) is an index corresponding to indices (i, j) in the weight matrix Θ , e.g., $(ij) = (i-1) \times s + j$. From the above definition of the coefficients of C , we obtain their partial derivatives with respect to λ :

$$\frac{\partial b}{\partial \lambda} = 0, \quad \frac{\partial H_{(ij), (i'j')}}{\partial \lambda_k} = \delta_{i,i'} \delta_{j,j'} \delta_{j,k}.$$

Note that here H is block-diagonal, with m identical blocks of size $(n+1)$, so the Cholesky decomposition (and similarly back-propagating through it) can be performed in about $(s/m)^3/3$ operations rather than $s^3/3$, where m is the number of outputs (the dimension of the output variable).

4 Optimizing Hyper-Parameters for a Non-Quadratic Criterion

If the training criterion C is not quadratic in terms of the parameters θ , it will in general be necessary to apply an iterative numerical optimization algorithm to minimize the training criterion. In this section we consider what happens after this minimization is performed, i.e., at a value of θ where $\frac{\partial C}{\partial \theta}$ is approximately zero and $\frac{\partial^2 C}{\partial \theta^2}$ is positive definite (otherwise we would not be at a minimum of C). We will use the implicit function theorem to obtain the derivative of θ with respect to λ at this point. Under appropriate conditions of continuity and differentiability, we have

$$G(\theta, \lambda) = 0 \rightarrow \frac{\partial \theta}{\partial \lambda} = -\left(\frac{\partial G}{\partial \theta}\right)^{-1} \frac{\partial G}{\partial \lambda}$$

for a vector-valued function G of the vectors θ and λ . In particular we consider here $G(\theta, \lambda) = \frac{\partial C}{\partial \theta} = 0$, so we obtain a **general formula for the gradient of the parameters with respect to the hyper-parameters**:

$$\frac{\partial \theta}{\partial \lambda} = -\left(\frac{\partial^2 C}{\partial \theta^2}\right)^{-1} \frac{\partial^2 C}{\partial \lambda \partial \theta} = -H^{-1} \frac{\partial^2 C}{\partial \lambda \partial \theta}. \quad (9)$$

Consider more closely the case of a neural network with one hidden layer and squared loss:

$$\mathcal{F}^{mlp} = \{f : \mathcal{R}^n \rightarrow \mathcal{R}^m \mid f(\mathbf{x}) = V\tilde{\mathbf{h}}, \mathbf{h} = \tanh(W\tilde{\mathbf{x}}), \mathbf{x} \in \mathcal{R}^n, \Theta \in \mathcal{R}^{m \times (n+1)}\}.$$

For example, if we want to use hyper-parameters for penalizing the use of inputs, we have a criterion similar to (8), and the cross-derivatives are easy to compute:

$$\frac{\partial^2 C}{\partial W_{i,j} \partial \lambda_k} = \delta_{k,j} W_{i,j}.$$

The Hessian and its inverse require more work, but can be done respectively in at most $O(s^2)$ and $O(s^3)$ operations. See for example [6] for the exact computation of the Hessian for multi-layer neural networks. See [2, 8] for a diagonal approximation which can be computed and inverted in $O(s)$ operations.

5 Summary of Experiments and Conclusions

In this paper, we have presented a new methodology for simultaneously optimizing several hyper-parameters, based on the computation of the gradient of a model selection criterion with respect to the hyper-parameters, taking into account the influence of the hyper-parameters on the parameters. We have considered both the simpler case of a training criterion that is quadratic with respect to the parameters and the more general non-quadratic case. We have shown a particularly efficient procedure in the quadratic case that is based on back-propagating gradients through the Cholesky decomposition and back-substitutions. This was an improvement: we have arrived at this $s^3/3$ operations procedure after studying first an $O(s^5)$ procedure and then a procedure taking about $(2s^3)$ operations for computing the gradients taking into account the influence of λ on θ . In the particular case of input weight decays for linear regression, the computation can even be reduced to about $(s/m)^3/3$ operations when there are m outputs.

We have performed preliminary experiments with the proposed methodology in several simple cases, using conjugate gradients to optimize the hyper-parameters. The application to linear regression with weight decays for each input is described in [?]. The hyper-parameter optimization algorithm is used to perform a soft selection of the input variables. A large weight decay on one of the inputs effectively forces the corresponding weights to very small values. Comparisons are made in [?] with ordinary regression as well as with stepwise regression methods and the adaptive ridge or LASSO, suggesting that the proposed method gives better results when there are many true non-zero regression coefficients and the correlation between the inputs is large.

Another type of application of the proposed method has been explored, in the context of a “real-world” problem of non-stationary time-series prediction [5]. In this case, an extension of the cross-validation criterion to sequential data which may be non-stationary is used. Because of this non-stationarity, recent data may sometimes be more relevant to current predictions than older data. The training criterion is a sum of weighted errors for the past examples, and these weights are given by a parametrized function of time (as the $w_i(\lambda)$ in eq. 1). The parameters of that function are two hyper-parameters that control when a transition in the unknown generating process would have occurred and how strong that change was or should be trusted. In these experiments, the weight given to past data points is a sigmoid function of the time: the threshold and the slope of the sigmoid are the hyper-parameters, representing respectively the time of a strong transition and the strength of that transition. Optimizing these hyper-parameters, we obtained statistically significant improvements in predicting one-month ahead future volatility of

Canadian stocks. The comparisons were made against several linear, constant, GARCH and ARMA models of the volatility. The experiments were performed on monthly return data from 473 Canadian stocks from 1976 to 1996. The measure of performance is the average out-of-sample squared error in predicting the squared returns. Single-sided significance tests were performed taking into account the auto-covariance in the temporal series of errors and the covariance of the errors between the compared models. When comparing the prediction of the first moment (expected return), no model significantly improved on the historical average of stock returns (constant model). When comparing the prediction of the second moment (expected squared returns), the method based on hyper-parameters optimization beat all the other methods, with a p-value of 1% or less.

What remains to be done? first, more experiments, in particular with the non-quadratic case (e.g., MLPs), and with model selection criteria other than cross-validation (which has large variance). Second, there are important theoretical questions that remain unanswered concerning the amount of overfitting that can be brought when too many hyper-parameters are optimized. As we have outlined in the introduction, the situation with hyper-parameters may be compared with the situation of parameters. However, whereas the form of the training criterion as a sum of independent errors allows to define the capacity for a class of functions and relate it to the difference between generalization error and training error, it does not appear clearly to us how a similar analysis could be performed for hyper-parameters.

Acknowledgments

The author would like to thank Léon Bottou, Pascal Vincent, François Blanchette, and François Gingras, as well as the NSERC Canadian funding agency.

References

- [1] H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, AC-19(6):716–728, 1974.
- [2] S. Becker and Y. LeCun. Improving the convergence of back-propagation learning with second order methods. In D. Touretzky, G. Hinton, and T. Sejnowski, editors, *Proceedings of the 1988 Connectionist Models Summer School*, pages 29–37, Pittsburg 1988, 1989. Morgan Kaufmann, San Mateo.
- [3] Y. Bengio. Continuous optimization of hyper-parameters. Technical Report 1144, Département d’informatique et recherche opérationnelle, Université de Montréal, 1999.
- [4] Yoshua Bengio and Charles Dugas. Learning simple non-stationarities with hyper-parameters. Technical Report 1145, Département d’informatique et recherche opérationnelle, Université de Montréal, 1999.
- [5] Yoshua Bengio and Charles Dugas. Learning simple non-stationarities with hyper-parameters. *submitted to Machine Learning*, 1999.
- [6] Christopher Bishop. Exact calculation of the Hessian matrix for the multi-layer perceptron. *Neural Computation*, 4(4):494–501, 1992.
- [7] P. Craven and G. Wahba. Smoothing noisy data with spline functions. *Numerical Mathematics*, 31:377–403, 1979.
- [8] Y. LeCun, J.S. Denker, and S.A. Solla. Optimal brain damage. In D.S. Touretzky, editor, *Advances in Neural Information Processing Systems 2*, pages 598–605, Denver, CO, 1990. Morgan Kaufmann, San Mateo.
- [9] D. MacKay and R. Neal. Automatic relevance determination, 1994. Unpublished report. See also MacKay D., 1995, Probable Networks and Plausible Predictions – A Review of Practical Bayesian Methods for Supervised Neural Networks, in *Neutwork: Computation in Neural Systems*, v. 6, pp. 469–505.
- [10] J. Rissanen. *Stochastic Complexity in Statistical Inquiry*. World Scientific, Singapore, 1990.
- [11] S.P. Smith. Differentiation of the cholesky algorithm. *Journal of Computational and Graphical Statistics*, 4:134–147, 1995.
- [12] A.N. Tikhonov and V.Y. Arsenin. *Solutions of Ill-posed Problems*. W.H. Winston, Washington D.C., 1977.

- [13] V.N. Vapnik. *Estimation of Dependences Based on Empirical Data*. Springer-Verlag, Berlin, 1982.