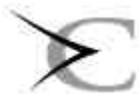


2002s-52

Input Decay: Simple and Effective Soft Variable Selection

Nicolas Chapados, Yoshua Bengio

Série Scientifique
Scientific Series



CIRANO
Centre interuniversitaire de recherche
en analyse des organisations

Montréal
Mai 2002

CIRANO

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du ministère de la Recherche, de la Science et de la Technologie, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO is a private non-profit organization incorporated under the Québec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the Ministère de la Recherche, de la Science et de la Technologie, and grants and research mandates obtained by its research teams.

Les organisations-partenaires / The Partner Organizations

- École des Hautes Études Commerciales
- École Polytechnique de Montréal
- Université Concordia
- Université de Montréal
- Université du Québec à Montréal
- Université Laval
- Université McGill
- Ministère des Finances du Québec
- MRST
- Alcan inc.
- AXA Canada
- Banque du Canada
- Banque Laurentienne du Canada
- Banque Nationale du Canada
- Banque Royale du Canada
- Bell Canada
- Bombardier
- Bourse de Montréal
- Développement des ressources humaines Canada (DRHC)
- Fédération des caisses Desjardins du Québec
- Hydro-Québec
- Industrie Canada
- Pratt & Whitney Canada Inc.
- Raymond Chabot Grant Thornton
- Ville de Montréal

© 2002 Nicolas Chapados et Yoshua Bengio. Tous droits réservés. *All rights reserved.* Reproduction partielle permise avec citation du document source, incluant la notice ©.
Short sections may be quoted without explicit permission, if full credit, including © notice, is given to the source.

Les cahiers de la série scientifique (CS) visent à rendre accessibles des résultats de recherche effectuée au CIRANO afin de susciter échanges et commentaires. Ces cahiers sont écrits dans le style des publications scientifiques. Les idées et les opinions émises sont sous l'unique responsabilité des auteurs et ne représentent pas nécessairement les positions du CIRANO ou de ses partenaires.

This paper presents research carried out at CIRANO and aims at encouraging discussion and comment. The observations and viewpoints expressed are the sole responsibility of the authors. They do not necessarily represent positions of CIRANO or its partners.

Input Decay: Simple and Effective Soft Variable Selection

Nicolas Chapados^{*} and Yoshua Bengio[†]

Résumé / Abstract:

Pour tenir compte des problèmes de sur-entraînement qui apparaissent quand il n'y a pas assez d'exemples comparativement au nombre de variables d'entrées durant l'apprentissage supervisé, les approches traditionnelles sont la pénalisation de la norme des paramètres (weight decay) et la sélection de variables vorace. Une alternative qui est apparue tout récemment est de garder toutes les variables, mais de mettre plus d'emphase sur celles qui sont le plus utiles. Nous introduisons une nouvelle méthode de régularisation, appelé "pénalisation sur la norme des entrées" (input decay), qui applique une plus grande pénalité relative sur les paramètres associés aux entrées qui contribuent le moins à la fonction apprise. Cette méthode, comme la pénalisation de la norme des paramètres (weight decay) et la sélection de variables, demande tout de même d'appliquer une sorte de sélection de modèle. Une série d'expériences comparatives avec cette nouvelle méthode ont été appliquées à deux tâches de régression, une qui était simulée et l'autre à partir d'une vrai problème financier.

To deal with the overfitting problems that occur when there are not enough examples compared to the number of input variables in supervised learning, traditional approaches are weight decay and greedy variable selection. An alternative that has recently started to attract attention is to keep all the variables but to put more emphasis on the most useful ones. We introduce a new regularization method called "input decay" that exerts more relative penalty on the parameters associated with the inputs that contribute less to the learned function. This method, like weight decay and variable selection, still requires to perform a kind of model selection. Successful comparative experiments with this new method were performed both on a simulated regression task and a real-world financial prediction task.

Mots clés : pénalisation sur la norme des entrées, sur-entraînement, réseau de neurones

Keywords: *input decay, overfitting, neural network*

^{*} CIRANO and Département d'informatique et recherche opérationnelle, Université de Montréal, Montréal, Québec, Canada, H3C 3J7. Email: chapados@iro.umontreal.ca

[†] CIRANO and Département d'informatique et recherche opérationnelle, Université de Montréal, Montréal, Québec, Canada, H3C 3J7. Tel: +1 (514) 343-6804, email: bengioy@iro.umontreal.ca

1 Introduction

In a large number of applications of machine learning algorithms, we face an implementation dilemma: a great number of input features is often available to solve the problem, but the limited size of the training set makes it seemingly impossible to use them all without running the risk of severely overfitting the data. This dilemma provides the rationale justifying classical variable selection procedures, such as stepwise selection [5] or branch-and-bound. These methods attempt to select the “good variables”, those yielding good generalization performance, to the exclusion of the others. We can alternatively define “good variables” to be those that are part of the generative model of the data; however, some variables that are part of the generative model may not, by themselves, be predictive enough to justify additional parameters in the model. In this paper, we argue that *penalized parameters* might allow to take such variables into account, albeit to a lesser extent. This enables to account for the fact that, in many situations, the distinction between “good variables” and “bad variables” is not nearly so clear cut. Some variables are certainly clearly useful; others are less so, *but they are not totally useless*.

Instead of reducing capacity by selecting a particular subset of variables, one can use regularization methods to reduce the capacity of the model. The most classical example is the weight decay [7] or “ridge” regression, which penalizes the squared norm of the parameter vector. However, weight decay penalizes all the input variables in the same way. More recently, several methods have been proposed to penalize input variables in different ways, depending on how “useful” they are. Examples of this class of algorithms are the adaptive ridge estimation procedure [6], the LASSO [11], and instances of hyper-parameter tuning such as those done in [8, 9] or in [1]. In this paper, we introduce a new approach to regularization for performing a “soft” selection of the variables, which we call *input decay*. It is well-suited to neural networks as well as classical linear regression, and is extremely easy to implement. Furthermore, contrarily to the combinatorial variable selection methods, it is computationally very cheap, requiring only a modest amount of effort over that normally required for an ordinary neural network training.

In section 2 we introduce notation and formalize and justify the proposed penalty. In section 3 we describe simulations in which we compare the proposed penalty method with more classical approaches, in a controlled setup where we can easily measure performance. In section 4, we present an application of the proposed method to a neural network regression problem occurring in financial decision-making.

2 Input Decay

Input decay is a method for performing “soft” variable selection during the regular training of a linear regression or non-linear neural network. The basic idea is that the training criterion penalizes the network connections coming from the inputs that have a less important role in determining the value of the output prediction.

Input decay works by adding a regularization term to the cost function used for training the network; the same principle can be applied to linear regression but we shall describe the general case of multi-layer perceptrons (MLPs). For a network trained to minimize the mean-squared error on a length- N training set $\{x_i, y_i\}$, the cost function incorporating input decay is

$$C = \frac{1}{2N} \sum_{i=1}^N (f(x_i; \theta) - y_i)^2 + C_{\text{ID}}(\theta), \quad (1)$$

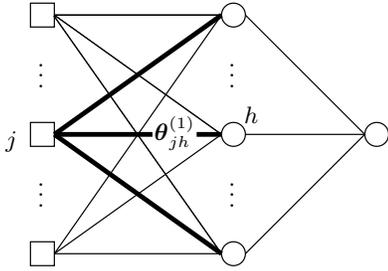


FIGURE 1. Network weights affected by the input decay term $C_{\text{ID}}^{(j)}(\boldsymbol{\theta})$, for an input j in a one-hidden-layer MLP (thick lines).

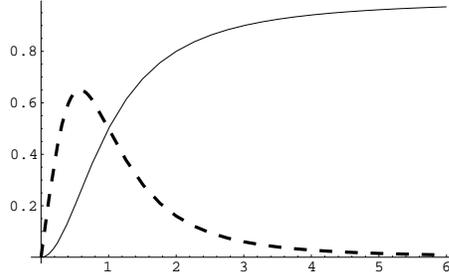


FIGURE 2. The input decay penalization function $x^2/(\eta + x^2)$ (solid) and its first derivative (dashed), for $\eta = 1$.

where $f(\cdot; \boldsymbol{\theta})$ is the function computed by the MLP and $C_{\text{ID}}(\boldsymbol{\theta})$ is the input decay term. The fundamental idea behind input decay is to impose a penalty on the squared-norm of the weights linking a particular network input to all the hidden units. Let $\boldsymbol{\theta}^{(i)}$ be the parameters on the i -th MLP layer, and $\theta_{jh}^{(1)}$ the first-layer network weight linking input j to hidden unit h ; the squared-norm of the weights from input j is:

$$C_{\text{ID}}^{(j)}(\boldsymbol{\theta}) = \sum_{h=1}^H \left(\theta_{jh}^{(1)}\right)^2, \quad (2)$$

where H is the number of hidden units in the network. The weights that are part of $C_{\text{ID}}^{(j)}(\boldsymbol{\theta})$ are illustrated in figure 1. The complete contribution $C_{\text{ID}}(\boldsymbol{\theta})$ to the cost function is obtained by a non-linear combination of the $C_{\text{ID}}^{(j)}$:

$$C_{\text{ID}}(\boldsymbol{\theta}) = \phi \sum_j \frac{C_{\text{ID}}^{(j)}}{\eta + C_{\text{ID}}^{(j)}(\boldsymbol{\theta})}, \quad (3)$$

where the hyper-parameter ϕ governs the relative important of the input decay term in the overall cost function. The behavior of the function $x^2/(\eta + x^2)$ is shown in figure 2. Intuitively, this function acts as follows: if the weights emanating from input j are small, the network must absorb a high marginal cost (locally quadratic) in order to increase the weights; the net effect, in this case, is to bring those weights closer to zero. On the other hand, if the weights associated with that input have become large enough, the penalty incurred by the network turns into a constant independent of the value of the weights; those are then free to be adjusted as appropriate. The hyper-parameter η acts as a threshold that determines the point beyond which the penalty becomes constant.

Input decay is similar to the *weight elimination* procedure [12] sometimes applied for training neural networks, with the difference that input decay applies in a collective way to the weights associated with a given input.

3 Experiments with Simulated Data

To ascertain the effectiveness of the input decay regularizer in principle, we performed experiments with generated data in a “difficult” linear regression setting. We compared the results obtained with the input decay method to stepwise (forward) variable selection and to the benchmark ordinary least-squares (OLS) regressor that uses all the variables. Experiments with MLPs on real data are described in section 4.

3.1 Experimental Setting

Data Generation We made use of the experimental framework described by Breiman [2], which consists in a linear regression problem. The generating model is

$$y^i = \boldsymbol{\beta}' \mathbf{x}^i + \epsilon^i, \quad (4)$$

where $\epsilon^i \sim \mathcal{N}(0, 1)$, and $\boldsymbol{\beta}$ and \mathbf{x}^i are length- M vectors (the prime denotes the transpose operation). The number of variables is fixed to $M = 30$. We describe below how the coefficients vector $\boldsymbol{\beta}$ (fixed during an experiment) is chosen. The input vector \mathbf{x}^i is drawn from a multivariate normal distribution of mean zero, and covariance matrix whose (i, j) -th element is $\rho^{|i-j|}$. Our experiments focused on the moderate-correlation ($\rho = 0.5$) and high-correlation ($\rho = 0.9$) situations.

In keeping with [2], we used three sets of $\boldsymbol{\beta} = (\beta_1, \dots, \beta_M)'$ coefficients. Within each set, we have $K = 3$ clusters of non-zero coefficients, centered at $c_1 = 5$, $c_2 = 15$, and $c_3 = 25$. Given the cluster centers $\{c_k\}$, the coefficients $\{\beta_m\}$ are computed as

$$\beta_m = C \sum_{k=1}^K [(h - |m - c_k|)^+]^2, \quad m = 1, \dots, 30, \quad (5)$$

where $(x)^+ \stackrel{\text{def}}{=} \max(0, x)$; $h = 1, 3, 5$, respectively for the first, second and third set of coefficients; and C is a constant chosen such that the coefficient of determination R^2 of the regression on the generated data is ≈ 0.75 .

Experiment Details For each of $\rho \in \{0.5, 0.9\}$ and $h \in \{1, 3, 5\}$, we generated 20 repetitions of 10000-element datasets. Within each dataset, the first 60 elements are used for training, the input values for the following 4940 are used in conjunction with the ADJ algorithm described below for selecting hyper-parameters, and the remaining 5000 are used for final testing.

We compare side-by-side the performance of the following models:

- *Classical OLS regression.* This involves no variable selection at all.
- *Input decay*, as described above, but for linear regression. Since we don't know a priori what are good values for the hyper-parameters ϕ and η , we train the models for every combination of $\phi \in \{0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1.0, 3.0, 10.0\}$ and $\eta \in \{0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1.0\}$, in addition to the “degenerate models” $\{\phi = 0, \eta = 100\}$ and $\{\phi = 100, \eta = 10^{-5}\}$ (large and small capacity models). We then select one model using the ADJ algorithm described below.
- *Forward variable selection.* We add variables one at a time (allowing up to 30 variables), and we choose the subset of variables giving the best estimated generalization error using 12-fold cross validation on the training set.

3.2 Review of the ADJ Model Selection Algorithm

The ADJ model selection algorithm, introduced by Shuurmans [10], is based on the idea of exploiting the natural geometry of the distribution of input vectors to achieve a re-ranking of competing models on the basis of those that can be “trusted” the most. The algorithm only needs access to *unlabeled* examples drawn from the input distribution, in order to estimate this distribution.

The full justification and geometric intuition behind ADJ are given in [10]; we give here an operational description. First, ADJ needs a partial order \prec on hypothesis

classes, ordering the classes by their complexity. For input decay, given hypothesis classes H_1 and H_2 , we define the order to be:

$$H_1 \prec H_2 \iff (\phi_1 \geq \phi_2 \wedge \eta_1 \leq \eta_2) \wedge (\phi_1 \neq \phi_2 \vee \eta_1 \neq \eta_2), \quad (6)$$

where ϕ_1 and η_1 are the hyper-parameters associated with H_1 , and correspondingly for H_2 . Intuitively, this order corresponds to the learning capacity allowed by the hyper-parameters. (We note that a larger input decay parameter ϕ *reduces* the capacity; in contrast, a smaller threshold η increases it.)

Next, the *expected distance* between two hypotheses h_1 and h_2 is defined as $d(h_1, h_2) = (\frac{1}{2} \int_X (h_1(x) - h_2(x))^2 dP_X)^{1/2}$, where the integral is over the input distribution space; similarly, the *empirical distance* on the training set $\{(x_i, y_i)\}$ of length N is $\hat{d}(h_1, h_2) = (\frac{1}{2N} \sum_i (h_1(x_i) - h_2(x_i))^2)^{1/2}$ (note that the computation of this empirical distance between two hypotheses does not use the targets y_i from the training set.)

Given a set of hypotheses $\{h_j^*\}$ (obtained, as usual, by minimizing the empirical error on the training set), each having obtained an empirical RMSE on the training set of $\widehat{\text{RMSE}}_j = (\frac{1}{2N} \sum_i (h_j^*(x_i) - y_i)^2)^{1/2}$, the ADJ algorithm re-ranks the $\widehat{\text{RMSE}}_j$ as follows:

- ① It finds the largest observed ratio of expected distance to empirical distance for the “smaller” hypotheses in the partial order \prec :

$$r_j = \max_{k, H_k \prec H_j} d(h_k^*, h_j^*) / \hat{d}(h_k^*, h_j^*). \quad (7)$$

- ② It adjusts the empirical $\widehat{\text{RMSE}}_j$ by this ratio: $\widehat{\text{RMSE}}'_j = r_j \widehat{\text{RMSE}}_j$.

The model ultimately selected by the algorithm is the one having the smallest adjusted $\widehat{\text{RMSE}}'$.

The computation of expected distance $d(h_k^*, h_j^*)$ between two hypotheses requires a model of the input distribution P_X . This model can be estimated by having only access to unlabeled data drawn from the input distribution; either kernel estimators or Monte Carlo methods can be used for this purpose. In our experiments, we used 4940 (unlabeled) vectors (separate from either the training or the test set) drawn from the input distribution to form a Monte Carlo estimate of the expected distance.

3.3 Results

The results of the experiments are summarized in tables 1 and 2. The first table gives the mean-squared errors obtained by each method, averaged over the 20 generated training and test sets. Standard errors under a Student t_{19} distribution are also given. From this table, we note that input decay model selected by ADJ performs much better than either a standard OLS regression using all the variables or a regression after stepwise variable selection. This is true for both moderately ($\rho = 0.5$) and highly correlated ($\rho = 0.9$) input variables, and for all coefficient vectors ($h = 1, 3, 5$).

Table 2 formally confirms these observations by tabulating the p -values obtained under the t_{19} distribution for the MSE differences between input decay and the other two methods. All p -values are highly statistically significant. In addition, the column ‘# significant’ lists the number of times, out of 20 repetitions, that input decay was found to be significantly better than the other method, using paired t -tests on the test set results.

TABLE 1. Test MSE obtained by the regression with input decay, versus a standard OLS regression and one using stepwise variable selection. The results are averaged over 20 different training and test sets.

ρ	h	Input Decay		OLS		Stepwise Selection	
		Avg. MSE	Std Err.	Avg. MSE	Std Err.	Avg. MSE	Std Err.
0.5	1	0.866	(0.027)	1.046	(0.050)	1.085	(0.087)
0.5	3	0.813	(0.027)	1.046	(0.050)	1.354	(0.101)
0.5	5	0.740	(0.019)	1.046	(0.050)	1.104	(0.040)
0.9	1	0.730	(0.028)	1.045	(0.050)	0.848	(0.029)
0.9	3	0.693	(0.023)	1.045	(0.050)	0.810	(0.026)
0.9	5	0.700	(0.020)	1.045	(0.050)	0.840	(0.041)

TABLE 2. MSE differences between regression with input decay and the other two methods, averaged over 20 generated datasets. The p -values result from t -tests over sequences of 20 differences. The ‘# Significant’ columns list the number of experiments in which input decay was found significantly better than the competing method (computed with a paired t -test on the test sets for the 20 datasets.)

ρ	h	Input Decay vs. OLS			Input Decay vs. Stepwise Selection		
		MSE diff.	p -value	# Significant	MSE diff.	p -value	# Significant
0.5	1	-0.180	< 0.001 *	19/20	-0.219	0.030 *	12/20
0.5	3	-0.233	< 0.001 *	18/20	-0.541	< 0.001 *	15/20
0.5	5	-0.306	< 0.001 *	19/20	-0.364	< 0.001 *	19/20
0.9	1	-0.315	< 0.001 *	20/20	-0.117	0.008 *	15/20
0.9	3	-0.352	< 0.001 *	20/20	-0.118	0.001 *	16/20
0.9	5	-0.348	< 0.001 *	20/20	-0.143	0.004 *	14/20

We conclude from these results that linear regression with input decay, given a reasonable model of the input distribution and a good model selection algorithm such as the ADJ algorithm, performs significantly better than either the benchmark OLS regression or stepwise variable selection.

4 Experiments with an Asset-Allocation Problem

We also applied input decay to a real-world asset-allocation problem. Our experiments consisted in allocating among the 14 sectors (sub-indices) of the Toronto Stock Exchange TSE 300 index. Our input variables consisted in technical indicators related to each asset, including moving averages (at several depths) of asset returns and estimated asset volatilities; a total of 75 input variables were used.

We used standard MLPs to make the asset allocation decisions. They were trained to make forecasts of future asset performance, those forecasts serving as input to a fixed decision system. In all cases, both ordinary weight decay and the input decay regularizer were incorporated into the cost functions used for training the MLPs. The complete experimental setup, including details on our investment framework that allocates assets according to a value-at-risk (VaR) constraint, is explained in [3, 4].

We compared the performance of various MLP topologies, varying the weight decay level, the input decay level (ϕ in eq. (3)) and the number of hidden units. In all cases, the input decay threshold η was kept fixed at 1.0. The performance criterion was

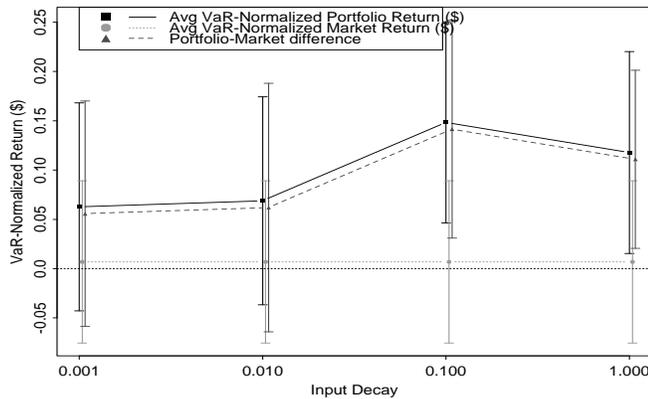


FIGURE 3. Effect of Input Decay on the financial performance obtained by an MLP in an asset-allocation task (solid). The (constant) benchmark market performance is given (dotted), along with the MLP-market difference (dashed). The error bars represent 95% confidence intervals.

a financial measure (the average return per period, normalized by the value-at-risk incurred) rather than a more conventional mean-squared error criterion.

Extensive statistical analysis of the results is presented elsewhere [3, 4]. We performed analyses of variance (ANOVA) to single out the effects of specific factors on the overall performance. In summary, we obtained the following results (we reserve the term “significant” to denote statistical significance at least at the 0.05 level):

- The effect of input decay is always significant (an example appears in figure 3).
- Weight decay is never significant.
- No higher-order interaction between the above factors is significant (as assessed by the ANOVA).

5 Conclusion

We introduced a new penalty-based method for soft variable selection that is very well-suited to multi-layer neural networks and classical linear regression settings. We showed this method to be successful on difficult simulated regression tasks and a real-world financial application. Moreover, it is exceedingly easy to implement, and, compared with combinatorial variable selection, is quite cheap computationally.

References

- [1] Y. Bengio. Gradient-based optimization of hyper-parameters. *Neural Computation*, 12(8):1889–1900, 2000.
- [2] Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1994.
- [3] N. Chapados. Optimization criteria for learning algorithms in portfolio selection. Master’s thesis, Université de Montréal, Montréal, Canada, January 2000.
- [4] N. Chapados and Y. Bengio. Cost functions and model combination for var-based asset allocation using neural networks. *IEEE Transactions on Neural Networks*, 2001. To appear.
- [5] N.R. Draper and H. Smith. *Applied Regression Analysis*. John Wiley and Sons, 1981.
- [6] Y. Grandvalet. Least absolute shrinkage is equivalent to quadratic penalization. In L. Niklasson, M. Boden, and T. Ziemskie, editors, *ICANN’98*, volume 1 of *Perspectives in Neural Computing*, pages 201–206. Springer, 1998.

- [7] G.E. Hinton. Learning translation invariant in massively parallel networks. In J.W. de Bakker et al., editor, *Proc. of PARLE Conference on Parallel Architectures and Languages Europe*, pages 1–13, Berlin, 1987. Springer.
- [8] 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 *Network: Computation in Neural Systems*, v. 6, pp. 469–505.
- [9] R.M. Neal. Assessing relevance determination methods using delve. In C.M. Bishop, editor, *Neural Networks and Machine Learning*, pages 97–129. Springer-Verlag, 1998.
- [10] D. Shuurmans. A new metric-basic approach to model selection. In *Proc. of the 14th National Conference on Artificial Intelligence*, Providence, RI, July 1997.
- [11] R.J. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society B*, 58:267–288, 1995.
- [12] A. S. Weigend, D. E. Rumelhart, and B. A. Huberman. Back-propagation, weight-elimination and time series prediction. In D.S Touretzky, J. L. Elman, T. J. Sejnowski, and G. E. Hinton, editors, *Connectionist Models: Proceedings of the 1990 Summer School*, San Mateo, CA, 1991. Morgan Kaufmann.

Liste des publications au CIRANO*

Série Scientifique / *Scientific Series* (ISSN 1198-8177)

- 2002s-52 Input Decay: Simple and Effective Soft Variable Selection / N. Chapados et Y. Bengio
- 2002s-51 On Out-of-Sample Statistics for Time-Series / F. Gingras, Y. Bengio et C. Nadeau
- 2002s-50 Forecasting Non-Stationary Volatility with Hyper-Parameters / Y. Bengio et C. Dugas
- 2002s-49 Cost Functions and Model Combination for VaR-based Asset Allocation using Neural Networks / N. Chapados et Y. Bengio
- 2002s-48 Experiments on the Application of IOHMMs to Model Financial Returns Series / Y. Bengio, V.-P. Lauzon et R. Ducharme
- 2002s-47 Valorisation d'Options par Optimisation du Sharpe Ratio / Y. Bengio, R. Ducharme, O. Bardou et N. Chapados
- 2002s-46 Incorporating Second-Order Functional Knowledge for Better Option Pricing / C. Dugas, Y. Bengio, F. Bélisle, C. Nadeau et R. Garcia
- 2002s-45 Étude du biais dans le Prix des Options / C. Dugas et Y. Bengio
- 2002s-44 Régularisation du Prix des Options : Stacking / O. Bardou et Y. Bengio
- 2002s-43 Monotonicity and Bounds for Cost Shares under the Path Serial Rule / Michel Truchon et Cyril Tétédo
- 2002s-42 Maximal Decompositions of Cost Games into Specific and Joint Costs / Michel Moreaux et Michel Truchon
- 2002s-41 Maximum Likelihood and the Bootstrap for Nonlinear Dynamic Models / Sílvia Gonçalves, Halbert White
- 2002s-40 Selective Penalization Of Polluters: An Inf-Convolution Approach / Ngo Van Long et Antoine Soubeyran
- 2002s-39 On the Mediation Role of Feelings of Self-Determination in the Workplace: Further Evidence and Generalization / Marc R. Blais et Nathalie M. Brière
- 2002s-38 The Interaction Between Global Task Motivation and the Motivational Function of Events on Self-Regulation: Is Sauce for the Goose, Sauce for the Gander? / Marc R. Blais et Ursula Hess
- 2002s-37 Static Versus Dynamic Structural Models of Depression: The Case of the CES-D / Andrea S. Riddle, Marc R. Blais et Ursula Hess
- 2002s-36 A Multi-Group Investigation of the CES-D's Measurement Structure Across Adolescents, Young Adults and Middle-Aged Adults / Andrea S. Riddle, Marc R. Blais et Ursula Hess
- 2002s-35 Comparative Advantage, Learning, and Sectoral Wage Determination / Robert Gibbons, Lawrence F. Katz, Thomas Lemieux et Daniel Parent

* Consultez la liste complète des publications du CIRANO et les publications elles-mêmes sur notre site Internet :

- 2002s-34 European Economic Integration and the Labour Compact, 1850-1913 / Michael Huberman et Wayne Lewchuk
- 2002s-33 Which Volatility Model for Option Valuation? / Peter Christoffersen et Kris Jacobs
- 2002s-32 Production Technology, Information Technology, and Vertical Integration under Asymmetric Information / Gamal Atallah
- 2002s-31 Dynamique Motivationnelle de l'Épuisement et du Bien-être chez des Enseignants Africains / Manon Levesque, Marc R. Blais, Ursula Hess
- 2002s-30 Motivation, Comportements Organisationnels Discrétionnaires et Bien-être en Milieu Africain : Quand le Devoir Oblige / Manon Levesque, Marc R. Blais et Ursula Hess
- 2002s-29 Tax Incentives and Fertility in Canada: Permanent vs. Transitory Effects / Daniel Parent et Ling Wang
- 2002s-28 The Causal Effect of High School Employment on Educational Attainment in Canada / Daniel Parent
- 2002s-27 Employer-Supported Training in Canada and Its Impact on Mobility and Wages / Daniel Parent
- 2002s-26 Restructuring and Economic Performance: The Experience of the Tunisian Economy / Sofiane Ghali and Pierre Mohnen
- 2002s-25 What Type of Enterprise Forges Close Links With Universities and Government Labs? Evidence From CIS 2 / Pierre Mohnen et Cathy Hoareau
- 2002s-24 Environmental Performance of Canadian Pulp and Paper Plants : Why Some Do Well and Others Do Not ? / Julie Doonan, Paul Lanoie et Benoit Laplante
- 2002s-23 A Rule-driven Approach for Defining the Behavior of Negotiating Software Agents / Morad Benyoucef, Hakim Alj, Kim Levy et Rudolf K. Keller
- 2002s-22 Occupational Gender Segregation and Women's Wages in Canada: An Historical Perspective / Nicole M. Fortin et Michael Huberman
- 2002s-21 Information Content of Volatility Forecasts at Medium-term Horizons / John W. Galbraith et Turgut Kisinbay
- 2002s-20 Earnings Dispersion, Risk Aversion and Education / Christian Belzil et Jörgen Hansen
- 2002s-19 Unobserved Ability and the Return to Schooling / Christian Belzil et Jörgen Hansen
- 2002s-18 Auditing Policies and Information Systems in Principal-Agent Analysis / Marie-Cécile Fagart et Bernard Sinclair-Desgagné
- 2002s-17 The Choice of Instruments for Environmental Policy: Liability or Regulation? / Marcel Boyer, Donatella Porrini
- 2002s-16 Asymmetric Information and Product Differentiation / Marcel Boyer, Philippe Mahenc et Michel Moreaux
- 2002s-15 Entry Preventing Locations Under Incomplete Information / Marcel Boyer, Philippe Mahenc et Michel Moreaux
- 2002s-14 On the Relationship Between Financial Status and Investment in Technological Flexibility / Marcel Boyer, Armel Jacques et Michel Moreaux
- 2002s-13 Modeling the Choice Between Regulation and Liability in Terms of Social Welfare / Marcel Boyer et Donatella Porrini

- 2002s-12 Observation, Flexibilité et Structures Technologiques des Industries / Marcel Boyer, Armel Jacques et Michel Moreaux
- 2002s-11 Idiosyncratic Consumption Risk and the Cross-Section of Asset Returns / Kris Jacobs et Kevin Q. Wang
- 2002s-10 The Demand for the Arts / Louis Lévy-Garboua et Claude Montmarquette
- 2002s-09 Relative Wealth, Status Seeking, and Catching Up / Ngo Van Long, Koji Shimomura
- 2002s-08 The Rate of Risk Aversion May Be Lower Than You Think / Kris Jacobs
- 2002s-07 A Structural Analysis of the Correlated Random Coefficient Wage Regression Model / Christian Belzil et Jörgen Hansen
- 2002s-06 Information Asymmetry, Insurance, and the Decision to Hospitalize / Åke Blomqvist et Pierre Thomas Léger
- 2002s-05 Coping with Stressful Decisions: Individual Differences, Appraisals and Choice / Ann-Renée Blais
- 2002s-04 A New Proof Of The Maximum Principle / Ngo Van Long et Koji Shimomura
- 2002s-03 Macro Surprises And Short-Term Behaviour In Bond Futures / Eugene Durenard et David Veredas
- 2002s-02 Financial Asset Returns, Market Timing, and Volatility Dynamics / Peter F. Christoffersen et Francis X. Diebold
- 2002s-01 An Empirical Analysis of Water Supply Contracts / Serge Garcia et Alban Thomas
- 2001s-71 A Theoretical Comparison Between Integrated and Realized Volatilities Modeling / Nour Meddahi
- 2001s-70 An Eigenfunction Approach for Volatility Modeling / Nour Meddahi
- 2001s-69 Dynamic Prevention in Short Term Insurance Contracts / M. Martin Boyer et Karine Gobert
- 2001s-68 Serial Cost Sharing in Multidimensional Contexts / Cyril Tétéjédo et Michel Truchon
- 2001s-67 Learning from Strike / Fabienne Tournadre et Marie-Claire Villeval
- 2001s-66 Incentives in Common Agency / Bernard Sinclair-Desgagné
- 2001s-65 Detecting Multiple Breaks in Financial Market Volatility Dynamics / Elena Andreou et Eric Ghysels
- 2001s-64 Real Options, Preemption, and the Dynamics of Industry Investments / Marcel Boyer, Pierre Lasserre, Thomas Mariotti et Michel Moreaux
- 2001s-63 Dropout, School Performance and Working while in School: An Econometric Model with Heterogeneous Groups / Marcel Dagenais, Claude Montmarquette et Nathalie Viennot-Briot
- 2001s-62 Derivatives Do Affect Mutual Funds Returns : How and When? / Charles Cao, Eric Ghysels et Frank Hatheway
- 2001s-61 Conditional Quantiles of Volatility in Equity Index and Foreign Exchange Data / John W. Galbraith, Serguei Zernov and Victoria Zinde-Walsh
- 2001s-60 The Public-Private Sector Risk-Sharing in the French Insurance "Cat. Nat. System" / Nathalie de Marcellis-Warin et Erwann Michel-Kerjan