HEC MONTRÉAL

Institut
d'économie appliquée

Forecasting chaotic systems: The role of local Lyapunov exponents

Dominique GUÉGAN
Justin LEROUX

Cahier de recherche n° IEA-07-12

3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7
www.hec.ca/iea

Forecasting chaotic systems: The role of local Lyapunov exponents

Dominique Guégan and Justin Leroux

December 14th 2007

Abstract

We propose a novel methodology for forecasting chaotic systems which is based on the nearest-neighbor predictor and improves upon it by incorporating local Lyapunov exponents to correct for its inevitable bias. Using simulated data, we show that gains in prediction accuracy can be substantial. The general intuition behind the proposed method can readily be applied to other non-parametric predictors.

1 Introduction

When taking a deterministic approach to predicting the future of a system, the main premise is that future states can be fully inferred from the current state. Hence, deterministic systems should in principle be easy to predict. Yet, some systems can be difficult to forecast accurately: such chaotic systems are extremely sensitive to initial conditions, so that a slight deviation from a trajectory in the state space can lead to dramatic changes in future behavior. We propose a novel methodology for forecasting deterministic systems which can then be extended to chaotic time series. For illustrative purposes, we describe how our methodology can be used to improve upon the nearest-neighbor predictor, but the same intuition can be applied to any non-parametric predictor (such as methods based on kernels, radial functions, neural nets, wavelets, etc.; see [1] and [2]) as it corrects for their inevitable bias by incorporating additional information on the local chaoticity of the system via the so-called local Lyapunov exponents (LLE).

The nearest-neighbor predictor has proved to be a simple yet useful tool for forecasting chaotic systems (see [3]). In the case of a one-neighbor predictor, it takes the observation in the past which most resembles today's state and returns that observation's successor as a predictor of tomorrow state. The rationale behind this nearest-neighbor predictor is quite simple: given that the system is assumed to be deterministic and ergodic, one obtains a sensible prediction of the variable's future by looking back at its evolution from a similar, past situation. For predictions more than one step ahead, the procedure is iterated by successively merging the predicted values with the observed data.

The nearest-neighbor predictor performs reasonably well in the short run but is not satisfactory for even medium-run predictions ([4], [5]). The generally accepted intuition being that the two trajectories (of the current state and of its neighbor) will have separated significantly by then, and the nearest neighbor's medium-run future will have little to do with the future we are trying to predict. Intuitively, this failure to perform well in the medium run arises mainly from the fact that short-run predictions are not accurate enough to withstand the complex dynamics of the system and to remain accurate after being iterated over a period of time of significant length. We argue that this lack of accuracy is inherent to the prediction method itself because the nearest neighbor on which predictions are based can never exactly coincide with today's state (or else the underlying process, being deterministic, would also be periodic and, thus, could not be chaotic).

We aim to correct the above shortcoming by incorporating information carried by the system's LLEs into the prediction. The LLE (see [6], [7]) represents the local dispersion rate of the system at a given point: a positive value meaning that two nearby points in the state space tend to grow apart over time, while a negative value indicates that nearby points will come closer together in the near future (but may diverge later on). In other words, the LLE is a measure of local chaoticity of a system. Typically, even a "globally chaotic" system is made up of both "chaotic regions" where the LLE is positive and more stable regions where it is negative. We illustrate this fact, which has been suggested in [8], more systematically in a companion paper.

By definition, the LLE tells us precisely by how much the distance between the current state and its nearest neighbor will expand (or contract) over time, so that we can easily obtain the distance between the nearest-neighbor predictor (i.e., the neighbor's successor) and the future we are trying to predict (tomorrow's state). Thus, we know exactly by how much to correct the prediction of the nearest-neighbor predictor.[1] To the best of our knowledge, no work has yet been done in this direction.

The rest of the paper is organized as follows. In Section 2, we develop our methodology by first pointing out why the nearest-neighbor predictor information on is biased and then suggesting how to correct this bias using information carried by the system's LLEs. In Section 3, we present simulations carried out on known chaotic systems to illustrate the extent of the (large) potential accuracy gains our methodology generates. Finally, Section 4 concludes by discussing the significance of the approach we propose and by pointing to directions in future work in order to refine it.

[1] Note that correction is required even in regions of the state space where the system is stable (i.e., where nearby trajectories come closer together, corresponding to a negative value of the LLE).

2 Methodology

Consider a one-dimensional series of T observations from a chaotic system, $(x_1, ...x_T)$, whose future values we are trying to forecast. Recall that a chaotic system is characterized by the existence of an attractor in a d-dimensional phase space (see [9]), where $d > 1$ is the embedding dimension.[2] A possible embedding method involves building a d-dimensional orbit, (X_t), with $X_t = (x_t, x_{t-\tau}, ..., x_{t-(d-1)\tau})$. For the sake of exposition, we shall assume $\tau = 1$ in the remainder of the paper.

By definition, the local Lyapunov exponent (or LLE) of a dynamical system which characterizes the rate of separation of infinitesimally close points of an orbit. Quantitatively, two neighboring points in phase space with initial separation δX_0 are separated, t periods later, by the distance:

$$|\delta X| \approx |\delta X_0| e^{\lambda_0 t},$$

where $|\cdot|$ represents the modulus of the considered vectors and λ_0 is the local Lyapunov exponent of the system in the vicinity of the initial points. Typically, this local rate of divergence (or convergence, if $\lambda_0 < 0$) depends on the orientation of the initial vector δX_0. Thus, strictly speaking, a whole spectrum of local Lyapunov exponents exists, one per dimension of the state space. A dynamic system is considered to be (locally) chaotic if $\lambda_0 > 0$, and (locally) stable if $\lambda_0 < 0$. (see, e.g., [8])

Our goal is to exploit the local information carried by the LLEs to improve upon existing methods of reconstruction and prediction. We propose a methodology which builds upon the classical nearest-neighbor predictor, which we now recall. Consider an orbit $(X_1, ..., X_T)$ whose one-step-ahead future, X_{T+1}, we are trying to predict. The nearest-neighbor predictor returns $\hat{X}_{T+1} = X_{i+1}$, where X_i is the element of the orbit with minimal distance to X_T. Because the dynamic system at hand is aperiodic (or else, forecasting would not be an issue), the nearest-neighbor predictor is inevitably biased. Indeed, because $|X_T - X_i| > 0$, it must also be the case that:

$$|\hat{X}_{T+1} - X_{i+1}| \approx |X_T - X_i| e^{\lambda_i} > 0, \tag{1}$$

where λ_i can be approximated in practice by the following expression:

$$\hat{\lambda}_i = \ln \frac{|X_{i+1} - X_{j+1}|}{|X_j - X_j|} \qquad \text{with } X_j = \arg\min_{t \neq i,T} |X_t - X_i| \tag{2}$$

It follows from Expression (1) that knowing the distance between the predictee and the nearest neighbor as well as the LLE at the nearest neighbor allows us to predict the distance of the predictee's image to the neighbor's image. Note that this is true regardless of the sign of λ_i; i.e., regardless of whether the system is locally chaotic or locally stable. Moreover, because the orbit considered

[2] The choice of the embedding dimension has been the object of much work (see [10] for a survey) and is beyond the scope of this note.

results from the embedding of a one-dimensional series, we also know all but the first coordinate of $X_{T+1} = (x_{T+1}, x_T, ..., x_{T-d+2})$. Hence, X_{T+1} lies at the intersection of the sphere of radius $|X_T - X_i|e^{\hat{\lambda}_i}$ centered on X_T and the line defined by $\{(z, x_T, ..., x_{T-d+2})|z \in \mathbb{R}\}$ which, in the Euclidean space, amounts to solving the following polynomial for $z \in \mathbb{R}$:

$$(z - x_{i+1})^2 + (x_T - x_i)^2 + ... + (x_{T-d+2} - x_{i-d+2})^2 - |X_T - X_i|e^{\hat{\lambda}_i} = 0 \quad (3)$$

Typically, two candidates emerge, $\hat{x}^-_{T+1}$ and $\hat{x}^+_{T+1}$, respectively underestimating and overestimating the true value of observation x_{T+1}(see Figure 1 in Appendix)[3].

One difficulty lies in determining when the nearest-neighbor predictor overestimates or underestimates the true value to be predicted. Being able to discriminate accurately between $\hat{x}^-_{T+1}$ and $\hat{x}^+_{T+1}$ may lead to significant improvements upon the nearest-neighbor predictor.

3 Simulations

We illustrate our point by simulating two well-known chaotic systems: the Lorenz system (see [12]) and the logistic map (see [13]). The Lorenz system is characterized by the following system of differential equations:

$$\begin{cases} \frac{dx}{dt} = \sigma(y - x) \\ \frac{dy}{dt} = x(R - z) - y \\ \frac{dz}{dt} = xy - bz \end{cases}$$

We simulated this system with values $\sigma = 16$, $R = 45.92$ and $b = 4$, initial values $x_0 = -10$, $y_0 = -10$ and $z_0 = 30$, and a step size of 0.01, integrated every fifth step. Taking the last 4,000 of our 5,000 observations to ensure that we are working within the attractor and considering the values on the x-coordinate as its own series, we successively predicted the last 1,000 in-sample observations. Each prediction was carried out with the full (and true) information set leading up to it, each time using the best of the two candidates, $\hat{x}^-_{T+1}$ and $\hat{x}^+_{T+1}$(measured in distance to the—known—successor). We obtain results which are always better than with the nearest-neighbor predictor and a mean-squared error which is half that of the nearest-neighbor predictor.

The logistic map is defined by:

$$x_{t+1} = 4x_t(1 - x_t),$$

Keeping the last 4500 of 5000 iterations, embedded in dimension 2, we predicted the in-sample 5001st observation of 1000 simulated trajectories with initial values drawn from U(0,1). We again obtain results which are always better than with the nearest-neighbor predictor and much smaller mean squared errors: of the order of 10^{-11} as opposed to 10^{-7} with the nearest-neighbor predictor; i.e. the best-candidate predictor is one hundred times more accurate!

[3] The situation whereby Expression (3) has no real solution would only arise if λ_i had been greatly underestimated, which has never occured to us in practice using Expression (2).

4 Concluding comments

The above preliminary analysis goes to show that there is great potential in improving upon the accuracy of the nearest-neighbor predictor by incorporating the information contained in local Lyapunov exponents as in Expression (1). Moreover, such increased precision in short-run prediction shall translate into accuracy gains for medium-run predictions, which is currently unsatisfactory with existing techniques. The general intuition behind the proposed method can readily be applied to other non-parametric predictors.

Several aspects of the implementation are still to be refined, and will be the object of future work. For instance, consistently discriminating between the two candidates,$\hat{x}_{T+1}^{-}$ and $\hat{x}_{T+1}^{+}$, can prove to be a difficult task due to the inherent chaotic nature of the systems at hand. As a first guess, one can select the candidate which maximizes the colinearity between the $X_{i+1} - X_i$ vector and the vector $\hat{X}_{T+1} - X_T$ (with $\hat{X}_{T+1}$ standing for $\hat{X}_{T+1}^{-}$ or $\hat{X}_{T+1}^{+}$). With the simulation of the Lorenz system described above, we achieve 97.5% accuracy, . Nonetheless, even considering only the observations where this so-called "LLE-corrected nearest-neighbor predictor" chooses the wrong candidates, errors are of the same order of magnitude as obtained with the classical nearest-neighbor estimator for the Lorenz system. This suggest that our LLE-corrected predictor still performs... On the other hand, we obtain only 66% accuracy with the logistic map. This is quite intuitive as the Lorenz system is "much less chaotic" than the logistic map (in the sense that its LLEs are "less often" positive and typically smaller than those of the logistic map; we elaborate on such distinctions of chaoticity in a companion paper) and, hence, is better behaved. Thus, with our rule of thumb, we achieve accuracy gains which are close to those obtained with best-candidate predictor on the Lorenz system. However, in the case of the logistic map (and highly chaotic systems, in general) our selector still needs refining. In another companion paper, we propose specific methods to improve upon the above rule of thumb to discriminate between candidates and, ultimately, yield better prediction results.

References

[1] Shintani M and Linton O. Nonparametric neural network estimation of Lyapunov exponents and a direct test for chaos, 2004; 120: 1-33.

[2] Guégan D and Mercier L. Stochastic and chaotic dynamics in high-frequency financial data. Signal Processing and Prediction, Prochazka Eds, 1998; 365-372.

[3] Farmer JD and Sidorowich JJ. Predicting chaotic time series. Physical Review Letters, 1987; 59: 845 – 848.

[4] Ziehmann C, Smith LA and Kurths J. Localized Lyapunov exponents and the prediction of predictability. Physics Letters A, 2000; 271: 237-251.

[5] Guégan D. Les Chaos en Finance: Approche Statistique. Economica Eds.; 2003, 420 pp.

[6] Abarbanel HDI. Local and global Lyapunov exponents on a strange attractor. Nonlinear Modeling and Forecasting, SFI Studies in the Science of Complexity, Casdagli M and Eubank S Eds, Addison-Wesley, 1992; Proc. Vol. XII: 229-247.

[7] Wolff RCL. Local Lyapunov exponents: looking closely at chaos. Journal of the Royal Statistical Society B, 1992; 54: 353 – 371.

[8] Bailey B. Local Lyapunov exponents: predictability depends on where you are. Nonlinear Dynamics and Economics, Kirman et al. Eds, 1997.

[9] Eckmann JP and Ruelle D. Ergodic theory of chaos and strange attractors. Review of Modern Physics, 1985; 57: 615-656.

[10] Takens F. Estimation of dimension and order of time series. Progress in Nonlinear Differential Equations and their Applications, 1996; 19: 405-422.

[11] Takens F. Detecting strange attractors in turbulence. Lecture Notes in Mathematics, 1981; 898: 366-381.

[12] Lorenz EN. Deterministic non-periodic flow. Journal of Atmospheric Science, 1963; 20: 130-141.

[13] May RM. Simple mathematical with complicated dynamics. Nature, 1976; 261: 459-467.

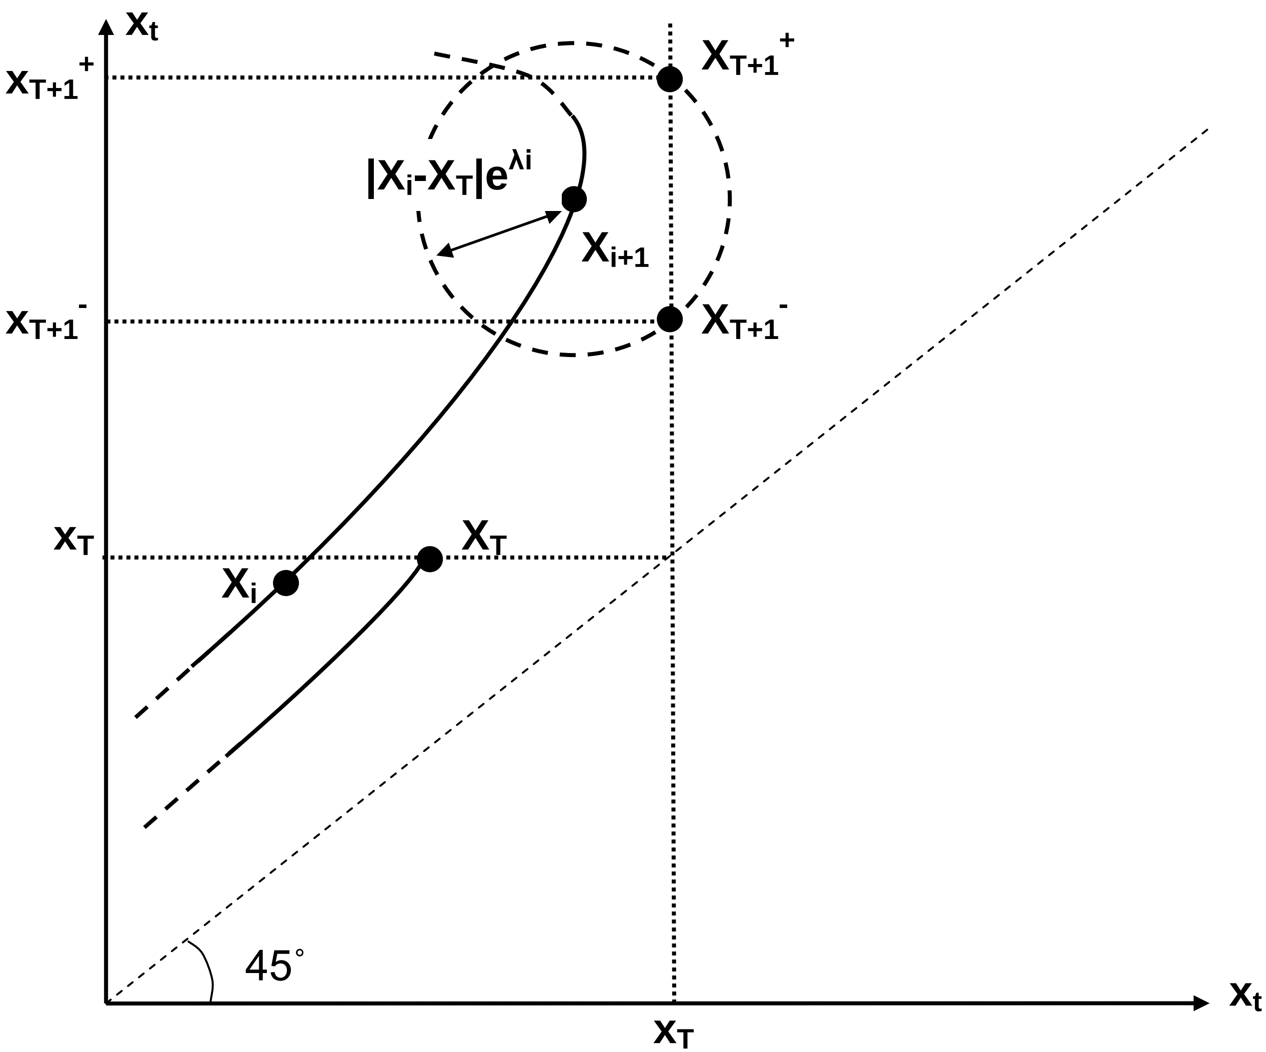

x_t
x_{T+1}^+
X_{T+1}^+
|X_i-X_T|e^{λi}
X_{i+1}
x_{T+1}^-
X_{T+1}^-
x_T
X_T
X_i
45°
x_t
x_T

Liste des cahiers de recherche publiés par les professeurs 2006-2007

Institut d'économie appliquée

IEA-06-01 DOSTIE, BENOIT ET LÉGER PIERRE THOMAS. « Self-selection in migration and returns to unobservable skills », 88 pages

IEA-06-02 JÉRÉMY LAURENT-LUCCHETTI AND ANDREW LEACH. « Induced innovation in a decentralized model of climate change », 34 pages.

IEA-06-03 BENOIT DOSTIE, RAJSHRI JAYARAMAN AND MATHIEU TRÉPANIER. « The Returns to Computer Use Revisited, Again », 27 pages.

IEA-06-04 MICHEL NORMANDIN. « The Effects of Monetary-Policy Shocks on Real Wages: A Multi-Country Investigation », 38 pages.

IEA-06-05 MICHEL NORMANDIN. « Fiscal Policies, External Deficits, and Budget Deficits », 50 pages.

IEA-06-06 J. DAVID CUMMINS, GEORGES DIONNE, ROBERT GAGNÉ AND ADBELHAKIM NOUIRA. « Efficiency of Insurance Firms with Endogenous Risk Management and Financial Intermediation Activities », 41 pages.

IEA-06-07 LUC BAUWENS AND JEROEN V.K. ROMBOUTS. « Bayesian Inference for the Mixed Conditional Heteroskedasticity Model », 25 pages.

IEA-06-08 LUC BAUWENS ARIE PREMINGER AND JEROEN V.K. ROMBOUTS. « Regime Switching Garch Models », 25 pages.

IEA-06-09 JEROEN V.K. ROMBOUTS AND TAOUFIK BOUEZMARNI. « Nonparametric Density Estimation for Positive Time Series », 32 pages.

IEA-06-10 JUSTIN LEROUX. « Cooperative production under diminishing marginal returns: Interpreting fixed-path methods », 25 pages.

IEA-06-11 JUSTIN LEROUX. « Profit sharing in unique Nash equilibrium characterization in the two-agent case », 16 pages.

IEA-06-12 ROBERT GAGNÉ, SIMON VAN NORDEN ET BRUNO VERSAEVEL. « Testing Optimal Punishment Mechnanisms under Price Regulation: the Case of the Retail Market for Gasoline », 27 pages.

IEA-06-13 JUSTIN LEROUX. « A discussion of the consistency axiom in cost-allocation problems », 11 pages.

IEA-06-14 MAURICE N. MARCHON. « Perspectives économiques canadiennes dans un contexte international », 29 pages.

IEA-06-15 BENOIT DOSTIE. « Wages, Productivity and Aging », 30 pages.

IEA-06-16 TAOUFIK BOUEZMARNI; JEROEN V.K. ROMBOUTS. « Density and Hazard Rate Estimation for Censored and α-mixing Data Using Gamma Kernels », 22 pages.

IEA-06-17 BENOIT DOSTIE et DAVID SAHN. «Labor Market Dynamics in Romania during a Period of Economic Liberalization», 37 pages

IEA-06-18 DAFNA KARIV, TERESA V. MENZIES, GABRIELLE A. BRENNER ET LOUIS-JACQUES FILION «Transnational Networking and Business Success: Ethnic entrepreneurs in Canada», 36 pages.

IEA-07-01 MOEZ BENNOURI, ROBERT CLARK, and JACQUES ROBERT. «Consumer Search and Information Intermediaries», 32 pages.

IEA-07-02 JEAN-FRANÇOIS ANGERS, DENISE DESJARDINS, GEORGES DIONNE, BENOIT DOSTIE and FRANÇOIS GUERTIN. «Poisson Models with Employer-Employee Unobserved Heterogeneity : An Aplicatin to Absence Data», 25 pages.

IEA-07-03 GEORGES DIONNE, ROBERT GAGNÉ AND ABDELHAKIM NOUIRA. «Determinants of Insurers' Performance in Risk Pooling, Risk Management, and Financial Intermediation Activities», 36 pages.

IEA-07-04 STEFAN AMBEC et PAUL LANOIE. «When and why does it pay to be green?», 40 pages.

IEA-07-05 CHRISTOS KOULOVATIANOS, LEONARD J. MIRMAN et MARC SANTUGINI. «Optimal growth and uncertainty : learning», 36 pages.

IEA-07-06 PAUL LANOIE, JÉRÉMY LAURENT-LUCCHETTI, NICK JOHNSTONE, STEFAN AMBEC. «Environmental Policy, Innovation and Performance : New Insights on the Porter Hypothesis», 40 pages.

IEA-07-07 PAUL LANOIE, DANIEL LLERENA. «Des billets verts pour des entreprises agricoles vertes?», 35 pages.

IEA-07-08 TAOUFIK BOUEZMARNI, JEROEN V.K. ROMBOUTS. « Semiparametric Multivariate Density Estimation for Positive Data Using Copulas», 29 pages.

IEA-07-09 LUC BAUWENS, ARIE PREMINGER, JEROEN V.K. ROMBOUTS. «Theory and Inference for a Markov Switching Garch Model», 26 pages.

IEA-07-10 TAOUFIK BOUEZMARNI, JEROEN V.K. ROMBOUTS. «Nonparametric Density Estimation for Multivariate Bounded Data», 31 pages.

IEA-07-11 GEORGES DIONNE and BENOIT DOSTIE. «Estimating the effect of a change in insurance pricing regime on accidents with endogenous mobility», 24 pages.