5th International Conference on Industrial Engineering and Industrial Management
XV Congreso de Ingeniería de Organización
Cartagena, 7 a 9 de Septiembre de 2011

Teaching the mean-field approximation

Luis R. Izquierdo [1], Segismundo S. Izquierdo [2] José I. Santos[1], José M. Galán[1] , Ricardo del Olmo[1]

[1] Department of Civil Engineering, University of Burgos, 09001 Spain; InSiSoc, luis@izquierdo.name
[2] Department of Industrial Organization, University of Valladolid, 47011 Spain; InSiSoc

Keywords: Mean field, approximation, teaching

1. Introduction

Many University subjects that we teach in engineering degrees (e.g. *Control Theory*, *Operations Research*, *Game Theory*, *Physics*, *Fluid Dynamics*, *Modelling and Control of Queuing and Production Systems*, *Stochastic Processes*, and *Network Theory*) make extensive use of the so-called mean-field analysis. In these subjects, the mean-field analysis often appears as an approximation of a *discrete time stochastic* process where a) the inherent *stochasticity* of the original process is replaced with *determinism*, and b) the time *discreteness* of the original process is replaced with time *continuity*. Thus, the mean-field approximation is presented as a *continuous* time *differential* equation that can approximate the dynamics of the *discrete* time *stochastic* process under investigation. In this paper we present a teaching methodology that we have found useful for introducing students to the mean-field analysis, and we provide some accompanying teaching material –in the form of computer models– that other academics may want to use in their own lectures. Our approach is twofolde:

- First, we provide the intuition underlying the mean-field approximation, and we discuss the type of models for which the approximation can be a useful tool. For this, we use a purpose-built computer model that allows students to grasp the reasoning behind the mean-field analysis. The model, which is freely available under the GNU General Public License, can be run as an applet at http://luis.izqui.org/models/gardendesert/ using any Internet browser.

- Then, as an illustration of the possible applications of the theory, and with the aim of fostering students' comprehension, we discuss specific examples of discrete stochastic processes that can be usefully approximated using the mean field. These examples are implemented as computer models with which students can play and, in this way, see for themselves the conditions under which the mean-field approximation works best.

This paper focuses on discrete stochastic processes that can be easily formulated as Markov chains, since this is the most common case in engineering contexts. Thus, first of all, we include in section 2 a brief summary of some of the basic concepts of Markov chains and, in

particular, of the difference between transient and asymptotic dynamics in that context. Then, in section 3 we describe the first stage of the methodology, i.e. we explain the intuition that underlies the mean-field approximation and we informally discuss the conditions under which the mean-field approximation can be useful to understand the transient dynamics of Markov models.

We consider this part of the methodology especially important since, most often, comprehending the intuition behind the mean-field analysis enables students to significantly enhance their understanding of the original stochastic processes, without getting lost in technical details. The second stage of the methodology is presented in the full paper, and includes an illustrative example of how to analyse discrete stochastic processes combining the theory of Markov chains with the mean-field approximation.

2. Markov Chains. Transient and asymptotic dynamics

Before embarking on the explanation of the mean-field analysis, some teachers may need to refresh their students' memory about Markov chains. For the sake of completeness, this section briefly revisits the most relevant concepts for our purposes. The main objective of this preliminary stage of the methodology is to set the scope of applicability of the mean-field approximation, namely the study of the transient behaviour of systems with large transition matrices (or with an infinite state space).

Markov chains are discrete stochastic processes in which, in every time-step n, a vector of state variables X_n can condense all the past information of the process that is relevant for its future. In a way, the values of such variables depict a "complete and accurate snapshot" of the current situation of the system. The picture is sufficient to characterise the future stochastically, in the sense that it contains all the necessary information to calculate the probability that each possible outcome will actually take place at any given time-step in the future. Thus, the value of the vector of state variables X_n is often called "the state of the system in time-step n". The current state of the system determines the probability distribution over the possible states of the system for every future time-step. Slightly more formally, given some initial state $X_0 = i$, we can calculate the probability distribution of the state of the system after n time steps: $\mathrm{P}(X_n = \cdot \mid X_0 = i)$.

The study of Markov processes is usually divided into transient and asymptotic behaviour. The transient behaviour is characterised by the probability distribution of the state vector X_n in a given time-step n. The asymptotic behaviour is characterised by the limit of the distribution of X_n as n goes to infinity, when this limit exists.

After a sufficiently long period of time, a Markov process is likely to be evolving very closely to what its asymptotic behaviour (stochastically) predicts; as a result, many studies of Markov processes focus on their asymptotic dynamics. Techniques to characterise the asymptotic behaviour of Markov chains can be found in introductory textbooks. However, "a sufficiently long period of time" may be too long, i.e. it may be unattainable in practical terms (e.g. it may require years of computation) or, simply, it may be significantly longer than the time scope we are actually interested in.

Thus, let us now turn to the transient behaviour. Consider a Markov chain with s possible states. In simple words, we are after a vector $a^{(n)} = [a_1^{(n)}, \ldots, a_s^{(n)}]$ containing the probability of finding the process in each of the s possible states in time-step n. The i^{th} element of $a^{(n)}$ is $a_i^{(n)} = \mathrm{P}(X_n = i)$, and it denotes the probability that the system is in state i at time-step n. To

calculate $a^{(n)}$ we must first introduce the transition matrix $\boldsymbol{P}$. Let $p_{i,j}$ denote the probability that the system moves from state i to state j in one time-step: $p_{i,j} = \mathrm{P}(X_{n+1} = j \mid X_n = i)$. The probabilities $p_{i,j}$ are called transition probabilities and they are often arranged in a matrix, namely the transition matrix $\boldsymbol{P}$. This matrix $\boldsymbol{P}$ characterises the stochastic process.

If the transition matrix $\boldsymbol{P}$ of a Markov chain is known and tractable, computing the transient behaviour of the system is straightforward: the probability distribution of the state of the system in time-step n is characterised by $a^{(n)} = a^{(0)} \cdot \boldsymbol{P}^n$. (Naturally, in general the distribution of the state of the system in time-step n depends on the initial conditions $a^{(0)}$.)

If the state transition matrix is either unknown or intractable, the situation is not so simple. In some cases, given some initial conditions $a^{(0)}$, computer simulations can approximate the distribution of the state of the system in time-step n, and the statistical accuracy of these approximations can be assessed (Izquierdo et al., 2009). Unfortunately, the nature of the parameter space of many models means that obtaining good approximations for their transient behaviour in the general case (i.e. for any parameterisation and all initial conditions) is often too demanding in computational terms. When this is the case, can we still say something about the expected behaviour of the Markov process before it gets close to its asymptotic behaviour? Here we explain that we sometimes can, by using the mean-field analysis.

3. The mean-field approximation:

We kick off our lectures on the mean-field approximation by showing our students a purpose-built computer model. We use this simulated example to show them –in a graphical way– the conditions under which the mean field can be useful to understand the behaviour of a Markov process. The story goes as follows: Consider one agent situated at a point in a bounded 2-dimensional space. In each time-step this agent takes one step in one of the four cardinal directions. The agent's movement is stochastic and slightly biased towards the North: a northward step is twice more likely than a step in any of the other three cardinal directions. Suppose that we place several agents like this one in the middle of a small area, such as a garden, and several others in the middle of a very large area, such as a desert (see Fig. 1).

If we observe the behaviour of the agents in the garden before they leave it, it will probably look quite random, with the agents leaving the garden through all its different sides. In contrast, if we observe the behaviour of the agents in the desert, they will seem quite determined to advance northwards, and with high probability they will all be leaving the desert through its northern border. When the number of steps required to cover some distance is very large, the randomness of the agents' movement averages out, and the agents will seem to be following their expected move northwards. This is the basic idea of the mean-field approximation (Benveniste et al., 1990).

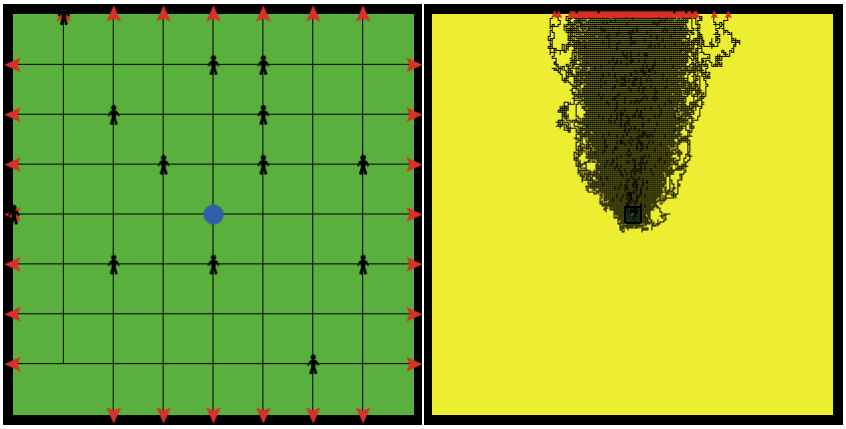

Fig. 1. The applet shows two independent areas: a small (8x8) green garden on the left, and a large (200x200) yellow desert on the right. For each area, an iteration of the model implies (a) the creation of an agent in the middle of the area, and (b) that each agent within the area advances one step: upwards (with probability 0.4), leftwards (with probability 0.2), rightwards (with probability 0.2) or downwards (with probability 0.2). Agents that reach the border of their corresponding area cannot return. Applet available at http://luis.izqui.org/models/gardendesert/

Loosely speaking, a requisite for the mean-field approximation to be useful is that any two consecutive states of the process are "very close" to each other –according to some measure of distance in the space of states– so the process needs to take a large number of time-steps in order to advance a significant distance in that space. Simply put, time proximity must imply space proximity. Consequently, we need to work with some measure of distance between different states. In most engineering applications (and in the simulated example in particular) this does not represent a problem, as the state variables of the stochastic process –the agents' position in our example– take real values within a closed and bounded subset of *d*-dimensional Euclidean space $\mathbf{R}^d$. Thus, we assume this is the case for the rest of the explanation.

Going back to our northward-biased random walker, suppose that the length of his step is γ, and let the 2-dimensional vector X_n^γ be his position vector in the grid at time *n*. Note that the law of motion of an agent with step size γ is $\Delta X_n^\gamma = \gamma \cdot Y_n$, where Y_n is a random vector which may take the value [0, 1] (northwards) with probability 0.4 or the values [1, 0], [-1, 0] or [0, -1] with probability 0.2 each. Note also that the expected move of X_n^γ is $E(\Delta X_n^\gamma) = \gamma \cdot E(Y_n) = \gamma$ [0, 0.2] (i.e. northwards). The mean-field approximation relates the actual (stochastic) move of X_n^γ with its (deterministic) expected move, for low values of the step size γ.

The question then is: under what conditions can one legitimately expect that the original *stochastic* process will tend to move roughly in the local direction of its *expected* (deterministic) move? This question is answered in the full paper, and is also illustrated with an example from the field of *Evolutionary Game Theory*.

Acknowledgements

The authors gratefully acknowledge financial support from the Spanish JCyL (VA006B09, GR251/2009), MICINN (SICOSSYS: TIN2008-06464-C03-02; CONSOLIDER-INGENIO 2010: CSD2010-00034; DPI2010-16920) and L.R.I. from the Spanish Ministry of Education (JC2009-00263).

References

Benaim, M., Weibull, J.W., 2003. Deterministic approximation of stochastic evolution in games. Econometrica 71**,** 873-903.

Benveniste, A., Priouret, P., Métivier, M., 1990. Adaptive algorithms and stochastic approximations. Springer-Verlag New York, Inc..

Binmore, K., Samuelson, L., 1997. Muddling Through: Noisy Equilibrium Selection. J. Econ. Theory 74**,** 235-265.

Edwards, M., Huet, S., Goreaud, F., Deffuant, G., 2003. Comparing an individual-based model of behaviour diffusion with its mean field aggregate approximation. J. Artif. Soc. Soc. Simulat. 6**,** 9.

Galan, J.M., Izquierdo, L.R., 2005. Appearances can be deceiving: Lessons learned re-implementing Axelrod's 'evolutionary approach to norms'. J. Artif. Soc. Soc. Simulat. 8.

Izquierdo, L.R., Izquierdo, S.S., Galán, J.M., Santos, J.I., 2009. Techniques to Understand Computer Simulations: Markov Chain Analysis. J. Artif. Soc. Soc. Simulat. 12.

Izquierdo, L.R., Izquierdo, S.S., Gotts, N.M., Polhill, J.G., 2007. Transient and asymptotic dynamics of reinforcement learning in games. Game. Econ. Behav. 61**,** 259-276.

Izquierdo, S.S., Izquierdo, L.R., Vega-Redondo, F., 2010. The option to leave: Conditional dissociation in the evolution of cooperation. Journal of Theoretical Biology 267**,** 76-84.

Kulkarni, V.G., 1995. Modeling and Analysis of Stochastic Systems. Chapman & Hall/CRC.

Kushner, H.J., Yin, G.G., 2003. Stochastic Approximation Algorithms and Applications. Springer-Verlag. New York.

Norman, M.F., 1972. Markov Processes and Learning Models. Academic Press, New York.