NONLINEAR CONTROL MARKOV PROCESSES AND GAMES[1]

Vassili N. Kolokoltsov[2]

[1]Dynamics and control Program review, 9-12 August 2010, Arlington, VA 22226.

[2]Department of Statistics, University of Warwick, Coventry CV4 7AL, UK. Email: v.kolokoltsov@warwick.ac.uk

General Objective

The objective is to analyze a new class of stochastic games that I call nonlinear Markov games, as they arise as a (competitive) controlled version of nonlinear Markov processes (an emerging field of intensive research).
This class of games can model a variety of situation for economics and epidemics, statistical physics, and pursuit - evasion processes. The official introduction to the theme is now covered in my monograph 'Nonlinear Markov processes and kinetic equations', CUP 2010.
Two basic examples to have in mind:
(1) Pursuit - evasion ; false targets
(2) Finances - hidden internal parameters of firms in competition

Connection with usual Markov processes

Nonlinear Markov process: future depends on the past not only via its present position, but also its present distribution.
A nonlinear Markov semigroup can be considered as a nonlinear deterministic dynamic system, though on a weird state space of measures.
Thus, as the stochastic control theory is a natural extension of the deterministic control, we extend it further by turning back to deterministic control, but of measures.

Nonlinear Markov chains (discrete time)

A *nonlinear Markov semigroup* Φ^k, $k \in \mathbf{N}$, is specified by an arbitrary continuous mapping $\Phi : \Sigma_n \to \Sigma_n$, where the simplex

$$\Sigma_n = \{\mu = (\mu_1, ..., \mu_n) \in \mathbf{R}_+^n : \sum_{i=1}^n \mu_i = 1\}.$$

Stochastic representation for Φ:

$$\Phi(\mu) = \{\Phi_j(\mu)\}_{j=1}^n = \{\sum_{i=1}^n P_{ij}(\mu)\mu_i\}_{i=1}^n, \qquad (1)$$

where $P_{ij}(\mu)$ is a family of stochastic matrices ($\sum_{j=1}^d P_{ij}(\mu) = 1$ for all i), depending on μ (nonlinearity!), whose elements specify the *nonlinear transition probabilities*.

Convergence to a stationary regime

Proposition

(i) For any continuous $\Phi : \Sigma_n \to \Sigma_n$ *there exists a stationary distribution, i.e. a measure* $\mu \in \Sigma_n$ *such that* $\Phi(\mu) = \mu$*. (ii) If a representation* (1) *for* Φ *is chosen in such a way that there exists a* $j_0 \in [1, n]$*, a time* $k_0 \in \mathbf{N}$ *and a positive* δ *such that*

$$P_{ij_0}^{k_0}(\mu) \geq \delta \tag{2}$$

for all i, μ, *then* $\Phi^m(\mu)$ *converges to a stationary measure for any initial* μ.

Proof.

Statement (i) is a consequence of the Brower fixed point principle. Statement (ii) follows from the representation (given above) of the corresponding nonlinear Markov chain as a time non-homogeneous Markov process. $\square$

Nonlinear Markov semigroup (continuous time)

A *nonlinear Markov semigroup* with the finite state space $\{1, ..., n\}$ is a semigroup Φ^t, $t \geq 0$, of continuous transformations of Σ_n. As in the case of discrete time the semigroup itself does not specify a process.
Stochastic representation for Φ^t:

$$\Phi_j^t(\mu) = \sum_i \mu_i P_{ij}(t, \mu), \quad t \geq 0, \mu \in \Sigma_n, \tag{3}$$

where $P(t, \mu) = \{P_{ij}(t, \mu)\}_{i,j=1}^n$ is a family of stochastic matrices depending continuously on $t \geq 0$ and $\mu \in \Sigma_n$ (nonlinear transition probabilities).

Nonlinear Markov chain (continuous time)

Once a stochastic representation (3) for the semigroup Φ^t is chosen one can define the corresponding stochastic process started at $\mu \in \Sigma_n$ as a time nonhomogeneous Markov chain with the transition probabilities from time s to time t being

$$p_{ij}(s, t, \mu) = P_{ij}(t - s, \Phi^s(\mu)).$$

Thus, to each trajectory of a nonlinear semigroup there corresponds a *tangent Markov process*.
Stochastic representation for the semigroup depends on the stochastic representation for the generator.

Generator of a nonlinear Markov semigroup

Namely, assuming the semigroup Φ^t is differentiable in t one can define the *(nonlinear) infinitesimal generator* of the semigroup Φ^t as the nonlinear operator on measures given by

$$A(\mu) = \frac{d}{dt}\Phi^t|_{t=0}(\mu).$$

The semigroup identity for Φ^t (nonlinear Chapman-Kolmogorov equation) implies that $\Phi^t(\mu)$ solves the Cauchy problem

$$\frac{d}{dt}\Phi^t(\mu) = A(\Phi^t(\mu)), \quad \Phi^0(\mu) = \mu. \tag{4}$$

The mapping A is *conditionally positive* in the sense that $\mu_i = 0$ for a $\mu \in \Sigma_n$ implies $A_i(\mu) \geq 0$ and is also *conservative* in the sense that A maps the measures from Σ_n to the space of the signed measures $\Sigma_n^0 = \{\nu \in \mathbf{R}^n : \sum_{i=1}^n \nu_i = 0\}$.

Generator: stochastic representation

We shall say that such an A has a *stochastic representation* if it is written in the form

$$A_j(\mu) = \sum_{i=1}^{n} \mu_i Q_{ij}(\mu) = (\mu Q(\mu))_j, \tag{5}$$

where $Q(\mu) = \{Q_{ij}(\mu)\}$ is a family of infinitesimally stochastic matrices (also referred to as Q-matrices or Kolmogorov's matrices) depending on $\mu \in \Sigma_n$. Thus in stochastic representation the generator has the form of a usual Markov chain generator, though additionally depending on the present distribution. The existence of a stochastic representation for the generator is not obvious, but is not difficult to get.

Exeample: replicator dynamics (RD)

The RD of the evolutionary game arising from the classical paper-rock-scissors game has the form

$$\begin{cases} \dfrac{dx}{dt} = (y - z)x \\ \dfrac{dy}{dt} = (z - x)y \\ \dfrac{dz}{dt} = (x - y)z \end{cases} \tag{6}$$

Its generator has a clear stochastic representation with

$$Q(\mu) = \begin{pmatrix} -z & 0 & z \\ x & -x & 0 \\ 0 & y & -y \end{pmatrix} \tag{7}$$

where $\mu = (x, y, z)$.

Example: simplest epidemics (1)

Let $X(t), L(t), Y(t)$ and $Z(t)$ denote respectively the numbers of susceptible, latent, infectious and removed individual at time t and the positive coefficients λ, α, μ (which may actually depend on X, L, Y, Z) reflect the rates at which susceptible individuals become infected, latent individuals become infectious and infectious individuals become removed.
Basic model, written in terms of the proportions $x = X/\sigma$, $y = Y/\sigma$, $l = L/\sigma$, $z = Z/\sigma$, where $\sigma = X + L + Y + Z$:

$$\begin{cases} \dot{x}(t) = -\sigma\lambda x(t)y(t) \\ \dot{l}(t) = \sigma\lambda x(t)y(t) - \alpha l(t) \\ \dot{y}(t) = \alpha l(t) - \mu y(t) \\ \dot{z}(t) = \mu y(t) \end{cases} \tag{8}$$

with $x(t) + y(t) + l(t) + z(t) = 1$.

Example: simplest epidemics (2)

Subject to the often made assumption that $\sigma\lambda$, α and μ are constants, the r.h.s. is an infinitesimal generator of a nonlinear Markov chain in Σ_4. This generator depends again quadratically on its variable and has an obvious stochastic representation (5) with the infinitesimal stochastic matrix

$$Q(\mu) = \begin{pmatrix} -\lambda y & \lambda y & 0 & 0 \\ 0 & -\alpha & \alpha & 0 \\ 0 & 0 & -\mu & \mu \\ 0 & 0 & 0 & 0 \end{pmatrix} \tag{9}$$

where $\mu = (x, I, y, z)$, yielding a natural probabilistic interpretation to the dynamics (8).

Discrete nonlinear Markov games and controlled processes

U, V are metric spaces of the control parameters,
$g(u, v, \mu)$, $u \in U$, $v \in V$, $\mu \in \Sigma_n$ is a continuous transition cost function,
$\nu(u, v, \mu)$ is a transition law $\nu(u, v, \mu)$.
One-step game (with sequential moves) is specified by the Bellman (or Shapley) operator

$$(BS)(\mu) = \min_u \max_v [g(u, v, \mu) + S(\nu(u, v, \mu))] \qquad (10)$$

for a given final cost function S on Σ_n.
The multi-step game solution is given by the iterations $B^k S$.
Long horizon problem: behavior of the optimal cost $B^k S(\mu)$ as $k \to \infty$.

Discrete nonlinear Markov games: long horizon (1)

Assume a stochastic representation for transitions is chosen:

$$\nu_j(u, v, \mu) = \sum_{i=1}^{n} \mu_i P_{ij}(u, v, \mu),$$

and

$$g(u, v, \mu) = \sum_{i,j=1}^{n} \mu_i P_{ij}(u, v, \mu) g_{ij}$$

with certain real coefficients g_{ij} (averages over the random transitions), then

$$(BS)(\mu) = \min_u \max_v \Big[\sum_{i,j=1}^{n} \mu_i P_{ij}(u, v, \mu) g_{ij} + S\left(\sum_{i=1}^{n} \mu_i P_{i.}(u, v, \mu) \right) \Big]. \tag{11}$$

Discrete nonlinear Markov games: long horizon (2)

Proposition

If the mapping ν is a contraction uniformly in u, v, i.e. if

$$\|\nu(u,v,\mu^1) - \nu(u,v,\mu^2)\| \leq \delta\|\mu^1 - \mu^2\|, \quad \delta < 1, \qquad (12)$$

and g is Lipschitz:

$$\|g(u,v,\mu^1) - g(u,v,\mu^2)\| \leq C\|\mu^1 - \mu^2\|, \qquad (13)$$

then there exists a unique $\lambda \in \mathbf{R}$ and a Lipschitz continuous function S on Σ_n such that $B(S) = \lambda + S$ and

$$\lim_{m\to\infty} \frac{B^m g}{m} = \lambda, \quad g \in C(\Sigma_n). \qquad (14)$$

Other results: turnpike theorems.

Discrete nonlinear Markov games: connection with usual Markov games

We can now identify the (not so obvious) place of the usual stochastic control theory in this nonlinear setting. Namely, assume P_{ij} above do not depend on μ. But even then the set of the linear functions $S(\mu) = \sum_{i=1}^n s_i \mu^i$ on measures (identified with the set of vectors $S = (s_1, ..., s_n)$) is not invariant under B. Hence we are not automatically reduced to the usual stochastic control setting, but to a game with incomplete information, where the states are probability laws on $\{1, ..., n\}$, i.e. when choosing a move the players do not know the position precisely, but only its distribution. Only if we allow only Dirac measures μ as a state space (i.e. no uncertainty on the state), the Bellman operator would be reduced to the usual one of the stochastic game theory:

$$(\bar{B}S)_i = \min_u \max_v \sum^n P_{ij}(u, v)(g_{ij} + S_j). \tag{15}$$

Continuous state spaces (general nonlinear Markov semigroups)

General kinetic equation in the weak form:

$$\frac{d}{dt}(f,\mu_t) = (L_{\mu_t} f, \mu_t), \quad \mu_t \in \mathcal{P}(\mathbf{R}^d), \quad \mu_0 = \mu, \tag{16}$$

(that should hold, say, for all $f \in C_c^2(\mathbf{R}^d)$), where

$$L_\mu f(x) = \frac{1}{2}(G(x,\mu)\nabla,\nabla)f(x) + (b(x,\mu),\nabla f(x))$$

$$+ \int (f(x+y) - f(x) - (\nabla f(x), y)\mathbf{1}_{B_1}(y))\nu(x,\mu,dy). \tag{17}$$

They play indispensable role in the theory of interacting particles (mean field approximation) and exhaust all positivity preserving evolutions on measures subject to certain mild regularity assumptions. They include Vlasov, Boltzmann, Smoluchovski, Landau-Fokker-Planck equations, McKean diffusions and many other models.

Nonlinear Markov process: definition, approach via SDE

A resolving semigroup $U_t : \mu \mapsto \mu_t$ of the Cauchy problem for this equation specifies a so called *generalized or nonlinear Markov process* $X(t)$, whose distribution μ_t at time t can be determined by the formula $U_{t-s}\mu_s$ from its distribution μ_s at any previous moment s.

We exploit the idea of nonlinear integrators combined with a certain coupling of Lévy processes in order to push forward the probabilistic construction in a way that allows the natural Lipschitz continuous dependence of the coefficients G, b, ν on x, μ. Thus obtained extension of the standard SDEs with Lévy noise represents a probabilistic counterpart of the celebrated extension of the Monge mass transformation problem to the generalized Kantorovich one.

Wasserstein-Kantorovich metrics

$$W_p(\nu_1, \nu_2) = \left(\inf_{\nu} \int |y_1 - y_2|^p \nu(dy_1 dy_2) \right)^{1/p}, \qquad (18)$$

where inf is taken over the class of probability measures ν on $\mathbf{R}^{2d}$ that couple ν_1 and ν_2.
The Wasserstein distances between the distributions in the Skorohod space $D([0, T], \mathbf{R}^d)$:

$$W_{p,T}(X_1, X_2) = \inf \left(\mathbf{E} \sup_{t \leq T} |X_1(t) - X_2(t)|^p \right)^{1/p}, \qquad (19)$$

where inf is over the couplings of the random paths X_1, X_2.
To compare Lévy measures, we extend these distances to unbounded measures with a finite pth moment.

Basic well-posedness: setting

$$L_\mu f(x) = \frac{1}{2}(G(x,\mu)\nabla,\nabla)f(x) + (b(x,\mu),\nabla f(x))$$

$$+ \int (f(x+z) - f(x) - (\nabla f(x), z))\nu(x,\mu;dz) \qquad (20)$$

with $\nu(x,\mu;.) \in \mathcal{M}_2(\mathbf{R}^d)$ (has a finite second moment). Let $Y_\tau(z,\mu)$ be a family of Lévy processes depending measurably on the points z and probability measures μ in $\mathbf{R}^d$ and specified by their generators

$$L[z,\mu]f(x) = \frac{1}{2}(G(z,\mu)\nabla,\nabla)f(x) + (b(z,\mu),\nabla f(x))$$

$$+ \int (f(x+y) - f(x) - (\nabla f(x), y))\nu(z,\mu;dy) \qquad (21)$$

where $\nu(z,\mu) \in \mathcal{M}_2(\mathbf{R}^d)$.

Position dependent SDE with a nonlinear noise

Our approach to solving (16) is via the solution to the following *nonlinear distribution dependent stochastic equation* with *nonlinear Lévy type integrators*:

$$X(t) = X + \int_0^t dY_s(X(s), \mathcal{L}(X(s))), \ \mathcal{L}(X) = \mu, \tag{22}$$

with a given initial distribution μ and a random variable X independent of $Y_\tau(z, \mu)$.
Euler-Ito approximation:

$$X_\mu^\tau(t) = X_\mu^\tau(l\tau) + Y_{t-l\tau}^l(X_\mu^\tau(l\tau), \mathcal{L}(X_\mu^\tau(l\tau))), \tag{23}$$

$\mathcal{L}(X_\mu^\tau(0)) = \mu$, where $l\tau < t \leq (l+1)\tau$, $l = 0, 1, 2, \ldots$, and $Y_\tau^l(x, \mu)$ is a collection (depending on l) of independent families of the Lévy processes $Y_\tau(x, \mu)$ introduced above.

Basic well-posedness: formulation

Theorem

Assume

$$\|\sqrt{G(x,\mu)} - \sqrt{G(z,\eta)}\| + |b(x,\mu) - b(z,\eta)|$$

$$+ W_2(\nu(x,\mu;.),\nu(z,\eta;.)) \le \kappa(|x-z| + W_2(\mu,\eta)), \quad (24)$$

$$\sup_{x,\mu}\left(\sqrt{G(x,\mu)} + |b(x,\mu)| + \int |y|^2\nu(x,\mu,dy)\right) < \infty. \quad (25)$$

Then for any $\mu \in \mathcal{P}(\mathbf{R}^d) \cap \mathcal{M}_2(\mathbf{R}^d)$ the approximations $X_\mu^{\tau_k}$, $\tau_k = 2^{-k}$, converge to a process $X_\mu(t)$ in W^2_{2,t_0} and the resolving operators $U_t : \mu \mapsto \mu_t$ of the Cauchy problem (16) form a nonlinear Markov semigroup. If $L[z,\mu]$ do not depend explicitly on μ the operators $T_t f(x) = \mathbf{E} f(X_x(t))$ form a conservative Feller semigroup preserving the space of Lipschitz continuous functions.

Basic well-posedness: example

(1) $\nu(x;.) = \sum_{n=1}^{\infty} \nu_n(x;.)$, $\nu_n(x,.)$ are probability measures with

$$W_2(\nu_i(x;.), \nu_i(z;.)) \leq a_i |x - z|$$

and the series $\sum a_i^2$ converges.
It is well known that the optimal coupling of probability measures (Kantorovich problem) can not always be realized via a mass transportation (a solution to the Monge problem), thus leading to the examples when the construction of the process via standard stochastic calculus would not work.
(2) common star shape of the measures $\nu(x;.)$:

$$\nu(x;dy) = \nu(x,s,dr)\,\omega(ds), \quad r = |y|, s = y/r, \qquad (26)$$

with a certain measure ω on S^{d-1} and a family of measures $\nu(x,s,dr)$ on $\mathbf{R}_+$. This allows to reduce the general coupling problem to a much more easily handled one-dimensional one.

Controlled nonlinear processes and games (1)

Consider a single control variable u and assume that μ only is observable, so that the control is based on μ. This leads to the following infinite-dimensional HJB equation

$$\frac{\partial S}{\partial t} + \max_u \left(L_{\mu,u} \frac{\delta S}{\delta \mu} + g_u, \mu \right) = 0. \tag{27}$$

If the Cauchy problem for the corresponding kinetic equation $\dot{\mu} = L^{\star}_{\mu,u}\mu$ is well posed (say, above theorem applies) uniformly for controls u from a compact set, with a solution denoted by $\mu^t(\mu, u)$ this can be resolved via discrete approximations

$$S_k(t-s) = B^k S(t), \quad k = (t-s)/\tau,$$

$$BS(\mu) = \max_u [S(\mu^{\tau}(\mu, u) + (g_u, \mu)].$$

Controlled nonlinear processes and games (2)

Convergence proof (yielding a Lipshitz continuous function for a Lipshitz continuous initial one) is the same as in book [1], Section 3.2, yielding a resolving operator $R_s(S)$ for the inverse Cauchy problem (27) as a *linear operator in the* max-*plus algebra*, i.e. satisfying the condition

$$R_s(a_1 \otimes S_1 \oplus a_2 \otimes S_2) = a_1 \otimes R_s(S_1) \oplus a_2 \otimes R_s(S_2)$$

with $\oplus = \max$, $\otimes = +$. This linearity allows for effective numeric schemes.
Extensions to a competitive control case (games) is settled via the approach with generalized dynamic systems as presented in Section 11.4 of book [2].

Selected bibliography

Books:

[1] V.N. Kolokoltsov, V.P. Maslov. Idempotent analysis and its applications. Kluwer Publishing House, 1997.

[2] V.N. Kolokoltsov, O.A. Malafeyev. Understanding Game Theory (Monograph). World Scientific, 2010.

[3] V. N. Kolokoltsov. Nonlinear Markov processes and kinetic equations. Cambridge University Press, August 2010.

Papers and talks:

[4] V. N. Kolokoltsov. Nonlinear Markov Semigroups and Interacting Lévy Type Processes. Journ. Stat. Physics **126:3** (2007), 585-642.

[5] V.N. Kolokoltsov. The Lévy-Khintchine type operators with variable Lipschitz continuous coefficients generate linear or nonlinear Markov processes and semigroups. arXiv:0911.5688v1 [Math.PR] (2009). To appear in PTRF.

[6] V.N. Kolokoltsov. Nonlinear Markov games. Proceedings of the 16th ILAS Conference, June 2010.