Iain Smears
supervised by
Dr. Max Jensen

Hamilton-Jacobi-Bellman Equations

Durham University

Analysis

Optimal Control (and notation).

Optimal control problems model situations where a stochastic process $x(s) \in \mathbb{R}^n$ is steered by the choice of a function $\alpha(\cdot)$, called a *control*, that takes values in a compact metric space Λ.

The dynamics of $x(s)$ are specified by a stochastic differential equation:

$$dx(s) = b\big(x(s), s, \alpha(x(s), s)\big)ds + \sigma\big(x(s), s, \alpha(x(s), s)\big)dW(s) \quad \text{for } s \in (t, T], \tag{1a}$$
$$x(t) = x; \tag{1b}$$

where $W(s) \in \mathbb{R}^d$ is a Brownian motion, $x \in \overline{U}$ for $U \in \mathbb{R}^n$ open and bounded, $T > 0$ is a terminal time, $t \in [0, T)$ is a start time, $b\colon \mathbb{R}^n \times \mathbb{R} \times \Lambda \mapsto \mathbb{R}^n$ and $\sigma\colon \mathbb{R}^n \times \mathbb{R} \times \Lambda \mapsto \mathbb{R}^{n\times d}$.

The objective is to minimise the *cost functional*, defined as

$$J\big(x, t, \alpha(\cdot)\big) = \mathbb{E}^{x,t}\left[\int_t^\tau f\big(x(s), s, \alpha(x(s), s)\big)ds + g\big(x(\tau), \tau\big)\right]; \tag{2}$$

where $\mathbb{E}$ means expectation, τ is the time of first exit of $x(\cdot)$ from $O = U \times (0, T]$, $f\colon \mathbb{R}^n \times \mathbb{R} \times \Lambda \mapsto \mathbb{R}$ is the *running cost* and $g\colon \mathbb{R}^n \times \mathbb{R} \mapsto \mathbb{R}$ is the *exit cost*.

To study optimal control problems, the *value function* is introduced as

$$u\,(x, t) = \inf_{\alpha(\cdot)\in\mathcal{A}} J\big(x, t, \alpha(\cdot)\big); \tag{3}$$

where $\mathcal{A}$ is a problem-given subset of stochastic processes that take values in Λ.

A control $\alpha^*(\cdot)$ is called optimal for start data $(x, t) \in O$ if $J\,(x, t, \alpha^*(\cdot)) = u(x, t)$. The *optimal control problem* is to find such an optimal control $\alpha^*(\cdot)$.

Solving the Optimal Control Problem. Under smoothness assumptions, the value function u satisfies the HJB equation (5) and that optimal controls maximise the term $-u_t(x, t) + L^\alpha u(x, t) - f(x, t, \alpha)$. Further details are found in [3].

The HJB equation is used twice in the process of solving the optimal control problem. First it is used as a partial differential equation for u. Having found u, it is then used to find an optimal control α^* by finding, for each $(x, t) \in \bar{O}$, some $\alpha^*(x, t) \in \Lambda$ a maximiser of $-u_t(x, t) + L^\alpha u(x, t) - f(x, t, \alpha)$.

Viscosity Solutions.

Definition. ([2]) A function u is a *viscosity solution* of the HJB equation if it is continuous on O, satisfies the boundary conditions, and satisfies the following property. For $\varphi \in C^2(O)$, if $u - \varphi$ has a maximum, respectively a minimum, at $(x, t) \in O$ then

$$-\varphi_t(x, t) + \max_{\alpha\in\Lambda}\left[L^\alpha \varphi(x, t) - f(x, t, \alpha)\right] \le 0; \tag{4}$$

respectively greater than or equal to 0. A viscosity solution is not necessarily a differentiable function that satisfies the equation in the usual pointwise meaning. This is why it is said that a viscosity solution is a form of *generalised* solution to a PDE.

Challenge. The two plots A and B below display two functions that satisfy almost everywhere the equation:

$$|u_x| - 1 = 0 \quad \text{on} \quad (-1, 1); \qquad u = 0 \quad \text{at} \; -1 \text{ and } 1.$$

With the definition of the previous paragraph, can you find which one is the viscosity solution of this equation?
The answer is at the bottom left of this poster.

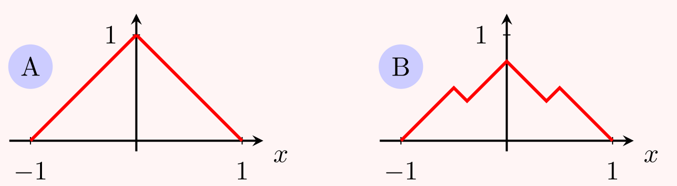

Relevance to HJB. Under reasonable assumptions, the viscosity solution of a HJB equation is unique. Often, for HJB equations originating from optimal control problems, *the value function $u(x, t)$ of equation (3) is the unique viscosity solution*, even if $u(x, t)$ is not smooth. This means that the notion of viscosity solution is the one that is relevant for solving an optimal control problem. For more details, see [3].

Answer: A

Partial Differential Equations

Hamilton-Jacobi-Bellman (HJB) Equations.

With the notation of the paragraph on optimal control, the HJB equation is

$$-u_t(x, t) + \max_{\alpha\in\Lambda}\left[L^\alpha u(x, t) - f(x, t, \alpha)\right] = 0 \quad \text{on} \quad O, \tag{5}$$

with the linear elliptic operators L^α defined by

$$L^\alpha u(x, t) = -\sum_{i,j}^{n} a_{ij}(x, t, \alpha) u_{x_i x_j}(x, t) - \sum_{i}^{n} b_i(x, t, \alpha) u_{x_i}(x, t);$$

and $a \in \mathbb{R}^{n\times n}$ defined by $a = 1/2\sigma\sigma^T$.

It is a backward parabolic partial differential equation, often considered with terminal condition $u(x, T) = g(x, T)$, $x \in U$ and lateral conditions $u(x, t) = g(x, t)$, $(x, t) \in \partial U \times [0, T)$, with g given by (2).

Theorem ([4]). *Under some regularity assumptions, if L^α is uniformly elliptic for all $(x, t, \alpha) \in \mathbb{R}^n \times \mathbb{R} \times \Lambda$, i.e. $\sum_{i,j=1}^n a_{ij}(x, t, \alpha)\xi_i\xi_j \ge c\|\xi\|_2^2$, for all $\xi \in \mathbb{R}^n$, then there exists a unique solution in $C^2(O) \cap C(\bar{O})$ to the HJB equation (5).*

However the cases where a is not uniformly elliptic are important; they include for instance all *deterministic* optimal control problems, i.e. where $\sigma \equiv 0$. *For these cases, a twice differentiable solution cannot generally be expected*, and the appropriate notion of solution is that of a viscosity solution.

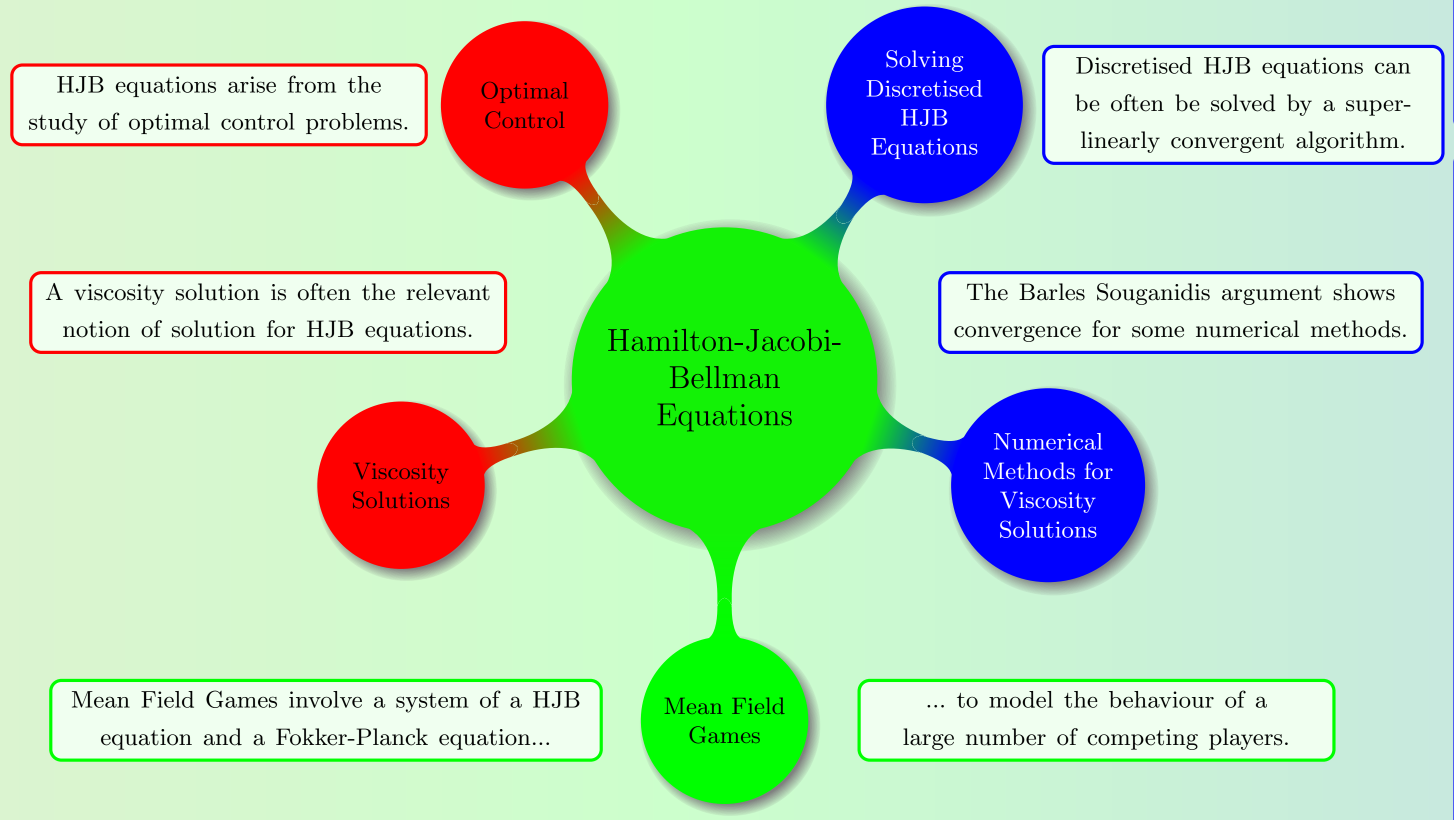

Mean Field Games.

Mean field games were introduced by P.-L. Lions and J.-M. Lasry in 2006 ([5]) to model a large system of "players" each undergoing an individual optimal control problem. The density of players at position x at time t is $m(x, t)$.

The position of each player evolves via the stochastic differential equation (1) with $\sigma := \sqrt{2\varepsilon}$, $\varepsilon > 0$ constant, and the players aim to optimise their cost functionals $J(x, \alpha) = \mathbb{E}\int_0^T f\,(x(s), \alpha) + V\,[m(\cdot, s)]\,(x(s))ds$, where $V\,[m(\cdot, s)]\,(x)$ represents the costs of the interactions between players.

With the ideas from the section on optimal control, the players determine their strategies through an optimal control $\alpha^*(x, t)$ determined by the HJB equation

$$-u_t(x, t) - \varepsilon\Delta u(x, t) + \max_{\alpha\in\Lambda}\left[-b(x, \alpha)\cdot\nabla u(x, t) - f(x, \alpha)\right] = V\,[m(\cdot, t)]\,(x); \tag{6a}$$

and the distribution of players evolves through the Fokker-Planck equation

$$m_t(x, t) - \varepsilon\Delta m(x, t) + \nabla\cdot\Big(b\big(x, \alpha^*(x, t)\big)m(x, t)\Big) = 0. \tag{6b}$$

Equations (6a) and (6b) constitute the mean field game equations, a nonlinearly coupled system of parabolic equations evolving in opposite time directions.

Numerical Analysis

Solving Discretised HJB Equations.

Many numerical schemes for the HJB equation can be represented by a sequence of equations of the form

$$F_h\,[u_h] = \max_{\alpha\in\Lambda}\left[A_h^\alpha u_h - d^\alpha\right] = 0, \quad u_h \in \mathbb{R}^{N_h}; \tag{7}$$

where N_h is the number of degrees of freedom of the scheme, $F\colon \mathbb{R}^{N_h} \mapsto \mathbb{R}^{N_h}$, A_h^α is an N_h-by-N_h matrix and $d^\alpha \in \mathbb{R}^{N_h}$. The maximum is understood here in the component-wise sense: for $x, y \in \mathbb{R}^{N_h}$, $(\max(x, y))_i = \max(x_i, y_i)$.

For each $v \in \mathbb{R}^{N_h}$, there exists $A(v) \in \mathbb{R}^{N_h\times N_h}$ and $d(v) \in \mathbb{R}^{N_h}$ such that $F(v) = A(v)v - d(v)$. In particular $A(v)$ is composed as a combination of the rows of the matrices A_h^α for which the maximum in (7) is attained.

Newton's method can be effective at solving equations with differentiable functions. However the function in (7) is not in general differentiable, because $x \mapsto |x|$ is a particular case of the general discrete HJB equation.

However, it is possible to use a so-called semismooth Newton method, based on a weaker notion of derivative. The end result of this analysis is the following theorem.

Theorem. ([6]) *If Λ is compact and A_h^α, d^α depend continuously on α, under some further assumptions, the semismooth Newton iterates defined by*

$$A(v_k)\,(v_{k+1} - v_k) = F_h(v_k), \quad k \in \mathbb{N}, \quad v_0 \text{ chosen},$$

will converge superlinearly to a solution u_h provided the initial guess v_0 is close enough to u_h.

Convergence of Numerical Methods to Viscosity Solutions.

Because of the definition of viscosity solutions, classical convergence arguments, such as those based on truncation errors, are not valid for showing convergence of numerical methods to viscosity solutions.

An important strategy for proving convergence, from Barles and Souganidis, can be used for monotone methods. To summarise this strategy informally, consider a finite difference scheme, written abstractly as $F_h\,[u_h]\,(x_i, t_k) = 0$ for all grid points (x_i, t_k).

Assume that F_h is monotone: it has the property that if u, v are functions and that $u - v$ has a maximum at a point (x_i, t_k) of the grid, then $F_h\,[u]\,(x_i, t_k) \ge F_h\,[v]\,(x_i, t_k)$.

Now suppose that for $\varphi \in C^2(O)$, $\lim_{h\to 0} u_h - \varphi$ has a maximum at a point $(x, t) \in O$. Then it is possible to show that for h small, $u_h - \varphi$ has a maximum at a point (x_i, t_k) of the grid that is near (x, t). Then by the monotonocity property,

$$0 = F_h\,[u_h]\,(x_i, t_k) \ge F_h\,[\varphi]\,(x_i, t_k). \tag{8}$$

So by taking the limit as $h \to 0$, if the scheme is consistent,

$$-\varphi_t + \max_{\alpha\in\Lambda}\left[L^\alpha\varphi(x, t) - f(x, t, \alpha)\right] \le 0.$$

Thus $\lim_{h\to 0} u_h$ satisfies the first requirement of a viscosity solution, as given in equation (4). A similar argument shows that $\lim_{h\to 0} u_h$ satisfies the second requirement and is thus a viscosity solution. *So the approximations u_h converge to the viscosity solution of the equation.* A specific example for which this argument applies is the Kushner-Dupuis scheme; more details can be found in [1] and [3].

References

[1] G. Barles and P. E. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Anal.*, 4(3):271–283, 1991.

[2] M. Crandall, H. Ishii, and P.-L. Lions. User's guide to viscosity solutions of second order partial differential equations. *Bull. Amer. Math. Soc. (N.S.)*, 27(1):1–67, 1992.

[3] W. Fleming and H. Soner. *Controlled Markov processes and viscosity solutions*, volume 25 of *Stochastic Modelling and Applied Probability*. Springer, New York, second edition, 2006.

[4] N. V. Krylov. *Nonlinear elliptic and parabolic equations of the second order*, volume 7 of *Mathematics and its Applications (Soviet Series)*. D. Reidel Publishing Co., Dordrecht, 1987. Translated from the Russian by P. L. Buzytsky.

[5] J.-M. Lasry and P.-L. Lions. Mean field games. *Jpn. J. Math.*, 2(1):229–260, 2007.

[6] I. Smears. Hamilton-Jacobi-Bellman equations and mean field games. *Fourth Year MMath Project Report*, 2011.