Gradient Descent Methods

Enrique Zuazua

UAM-Madrid & DeustoTech-Bilbao & LJLL-Sorbonne Univ.-Paris
enrique.zuazua@deusto.es
http://paginaspersonales.deusto.es/enrique.zuazua/

ConFlex-UAM Workshop, February 2019

Classical steepest descent

$J : H \to \mathbf{R}$. Two main assumptions:

$$< \nabla J(u) - \nabla J(v), u - v > \geq \alpha |u - v|^2, \quad |\nabla J(u) - \nabla J(v)|^2 \leq M |u - v|^2.$$

Then, for

$$u_{k+1} = u_k - \rho \nabla J(u_k),$$

we have

$$|u_k - u^*| \leq (1 - 2\rho\alpha + \rho^2 M)^{k/2} |u_1 - u^*|.$$

Convergence is guaranteed for $0 < \rho < 1$ small enough.

Compare with the continuous marching gradient system

$$u'(\tau) = -\nabla J(u(\tau)).$$

LaSalle's invariance principle

Taking the scalar product in equation $u'(t) = -\nabla J(u(t))$ with $\nabla J(u(t))$ we deduce that

$$dJ(u(t)/dt = -|\nabla J(u(t))|^2.$$

Thus, for the gradient system, $J(u(t))$ constitutes a Lyapunov function whose value diminishes along trajectories.

Assume that J is bounded below. This is typically the case when searching the minimizers of J under the standard coercivity and continuity assumptions.

Then, necessarily, $J(u(t))$ has a limit l as $t \to \infty$.

Furthermore, when J is coercive, this necessarily means that the trajectory $\{u(t)\}_{t\geq 0}$ is bounded. In the finite-dimensional context this means that the trajectory is precompact. In the infinite-dimensional case this requires further analysis of the dynamical properties of the evolution system under consideration.

Let us then define the ω-limit set. Given the initial datum u_0 of the solution of the gradient system, $\omega(u_0)$ is the set of accumulation points of the trajectory as $t \to \infty$. Obviously $J(z) = I$ for all $z \in \omega(u_0)$. On the other hand, if we denote by $z(t)$ the trajectory of the same gradient system starting at z at time $t = 0$, by the semigroup property, we also deduce that $J(z(t)) = I$ for all $t \geq 0$. This implies, in particular, that z is a critical point of J: $J(z) = 0$. In case J has an unique minimizer, as it happens when J is strictly convex, then z is this minimizer.
Taking into account that the accumulation point is unique, we deduce that $\omega(u_0) = \{z\}$. This implies that the whole trajectory $u(t)$ converges to z.

As we mentioned above, in the infinite-dimensional case, the boundedness of trajectories does not necessarily imply that they are relatively compact. The compactness of trajectories is normally achieved by imposing further monotonicity properties.

Indeed, when J is convex, distances diminish along trajectories. Indeed, if u and v are two trajectories of the same system then $|u(t) - v(t)|$ diminishes as time evolves.

According to this it is sufficient to prove convergence towards equilibrium for a dense set of initial data. This dense set is chosen normally to ensure compactness through the compactness of the embedding into the phase space, and the boundedness of the trajectories in that subspace.

The Direct Method of the Calculus of Variations (DMCV)

Consider a continuous, convex and coercive functional $J : H \to \mathbb{R}$ in a Hilbert space H. Then, the functional achieves its minimum in at least one point:

$$\exists h \in H : \; J(h) = \min_{g \in H} J(g). \tag{1}$$

This can be easily proved in a systematic manner by means of the DMCV:
Step 1. Define the infimimum

$$I = \inf_{g \in H} J(g)$$

that, by the coercivity of J, necessarily satisfies $I > -\infty$.

Step 2. Consider the minimizing sequence

$$(g_n)_{n\in\mathbb{N}} \subset H : \; J(g_n) \searrow I. \tag{2}$$

By the coercivity of the functional J we deduce that $(g_n)_{n\in\mathbb{N}}$ is bounded in H.

Step 3. H being a Hilbert space, there exists a weakly convergent subsequence $(g_n)_{n\in\mathbb{N}}$

$$g_n \rightharpoonup g \text{ en } H. \tag{3}$$

Step 4. J being continuous in H and convex it is lowe semicontinuous with respect to the weak topology. Therefore,

$$J(g) \leqslant \varliminf_{n\to\infty} J(g_n). \tag{4}$$

LaSalle's invariance principle Taking the scalar product in equation $u'(t) = -\nabla J(u(t))$ with $\nabla J(u(t))$ we deduce that

$$dJ(u(t)/dt = -|\nabla J(u(t))|^2.$$

Thus, for the gradient system, $J(u(t))$ constitutes a Lyapunov function whose value diminishes along trajectories.

Assume that J is bounded below. This is typically the case when searching the minimizers of J under the standard coercivity and continuity assumptions.

Then, necessarily, $J(u(t))$ has a limit I as $t \to \infty$.

Furthermore, when J is coercive, this necessarily means that the trajectory $\{u(t)\}_{t\geq 0}$ is bounded. In the finite-dimensional context this means that the trajectory is precompact. In the infinite-dimensional case this requires further analysis of the dynamical properties of the evolution system under consideration.

Let us then define the ω-limit set. Given the initial datum u_0 of the solution of the gradient system, $\omega(u_0)$ is the set of accumulation points of the trajectory as $t \to \infty$. Obviously $J(z) = I$ for all $z \in \omega(u_0)$. On the other hand, if we denote by $z(t)$ the trajectory of the same gradient system starting at z at time $t = 0$, by the semigroup property, we also deduce that $J(z(t)) = I$ for all $t \geq 0$. This implies, in particular, that z is a critical point of J: $J(z) = 0$. In case J has an unique minimizer, as it happens when J is strictly convex, then z is this minimizer.

Taking into account that the accumulation point is unique, we deduce that $\omega(u_0) = \{z\}$. This implies that the whole trajectory $u(t)$ converges to z.

As we mentioned above, in the infinite-dimensional case, the boundedness of trajectories does not necessarily imply that they are relatively compact. The compactness of trajectories is normally achieved by imposing further monotonicity properties.

Indeed, when J is convex, distances diminish along trajectories. Indeed, if u and v are two trajectories of the same system then $|u(t) - v(t)|$ diminishes as time evolves.

According to this it is sufficient to prove convergence towards equilibrium for a dense set of initial data. This dense set is chosen normally to ensure compactness through the compactness of the embedding into the phase space, and the boundedness of the trajectories in that subspace.

We deduce that $J(g) \leqslant I$ which, by the definition of infimum, implies that $J(g) = I$, which shows that the minimum is achieved.