Transversality, the maximum principle, and the approximation problem

Héctor J. Sussmann

Department of Mathematics — Rutgers University
Piscataway, NJ 08854, USA
`sussmann@math.rutgers.edu`

Conference on
Optimization, State Constraints and Geometric Control
Dipartimento di Matematica "Tullio Levi-Civita"
Università degi Studi di Padova
May 23-24, 2018

HAPPY

BIRTHDAY

FRANCO !!!!!

HAPPY

BIRTHDAY

GIOVANNI !!!!!

REMARKS FOR THE EXPERTS I

All the ideas of this talk are reall contained, at least implicitly, in the original work of Pontryagin-Boltyanskii-Gamkrelidze-Mischenko.

And they were understood quite explicitly by Jack Warga.

The purposes of this talk are

1. to clarify these old ideas and explain them in simple modern language

2. to show, how properly formulated, these ideas can be extended further, in particular to the cases where the relevant maps are set-valued.

REMARKS FOR THE EXPERTS II

Most of the ideas discussed in this talk have been presented in previous lectures and papers.

But the approach used in this talk is new.

In particular, I have been able to do away with the distinction between "transversality" and "strong transversality", thus making the new approach much simpler.

REMARKS FOR THE EXPERTS III

The work discussed here is closely related to that of Michele Palladino and Franco Rampazzo on the gap problem.

STRUCTURE OF THE TALK

1. Transversality

2. The maximum principle as a transversality theorem

3. Approximation of trajectories by trajectories for a smaller class of controls

TRANSVERSALITY

The idea of transversality is very simple.

EXAMPLE: Let $\gamma_1 : [0,1] \mapsto \mathbb{R}^2$, $\gamma_2 : [0,1] \mapsto \mathbb{R}^2$, be two continuous curves in the plane.

Then γ_1, γ_2 may:

1. not intersect at all,

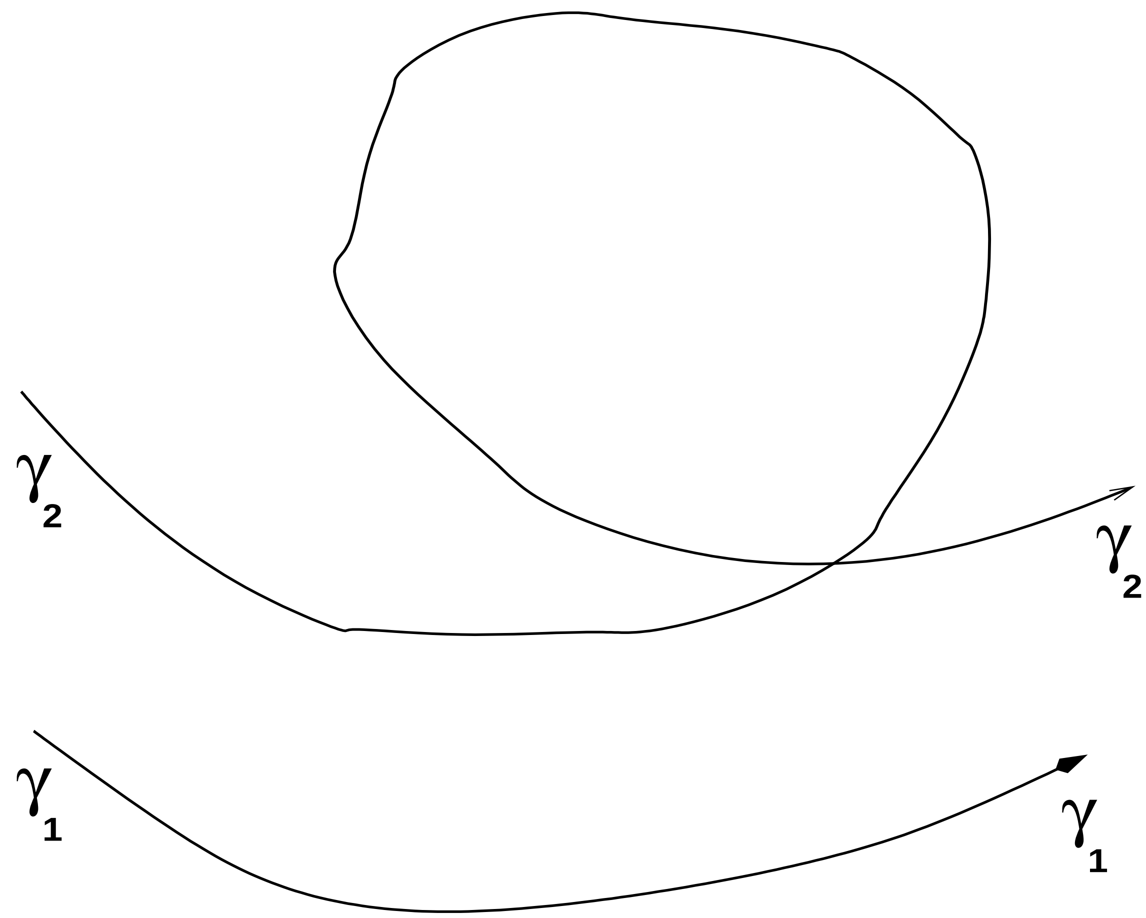

2. intersect "tangentially without crossing",

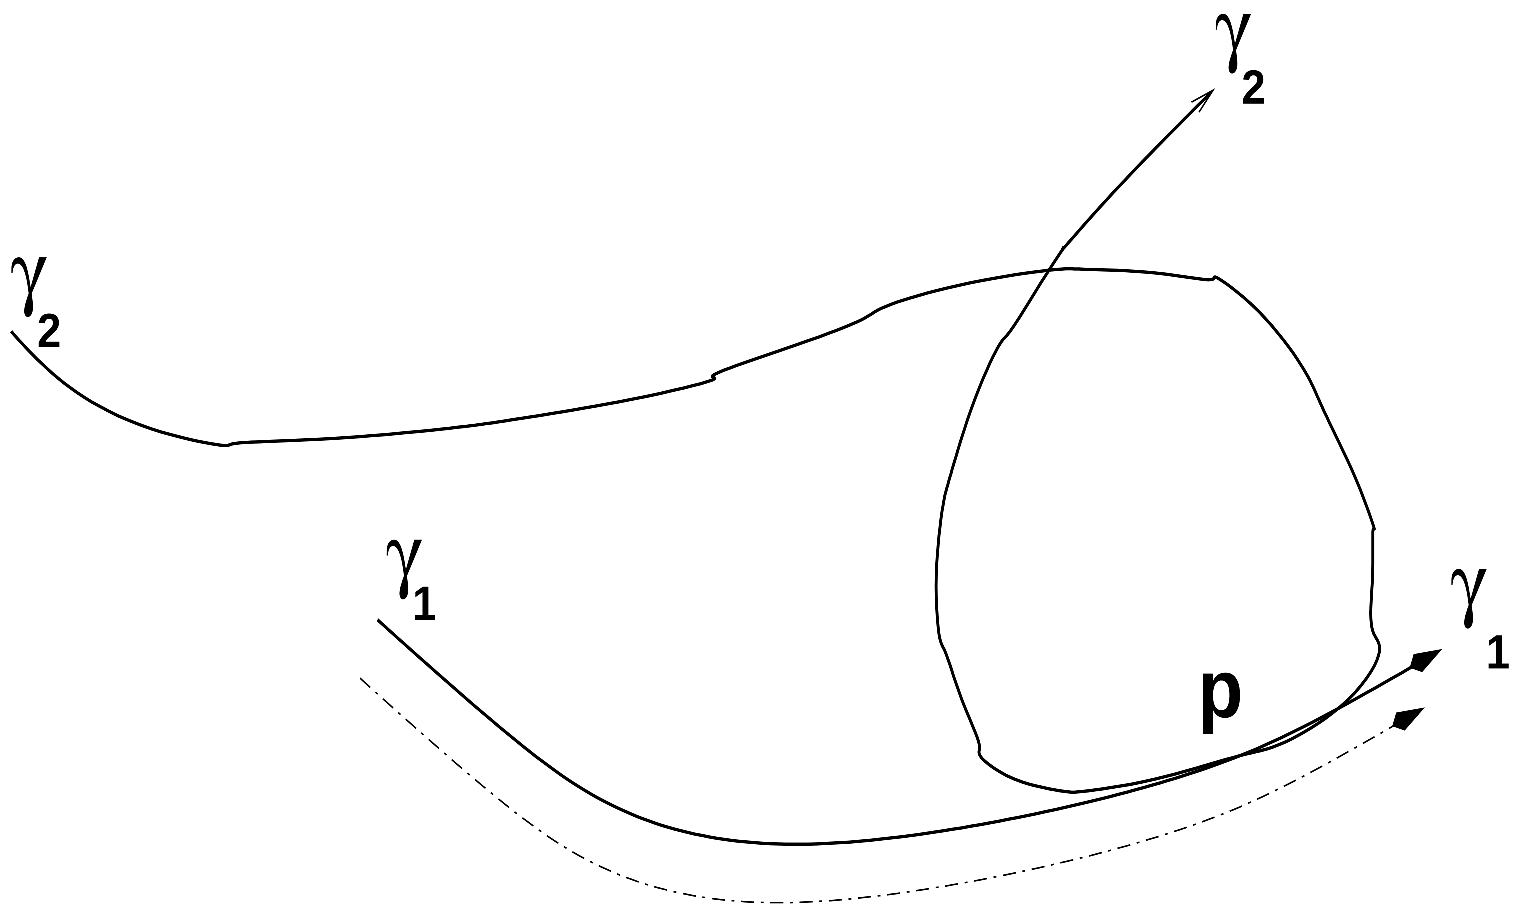

in which case it's possible to make arbitrarily small perturbations of γ_1, γ_2 that will not intersect at all;

3. intersect transversally (i.e., cross),

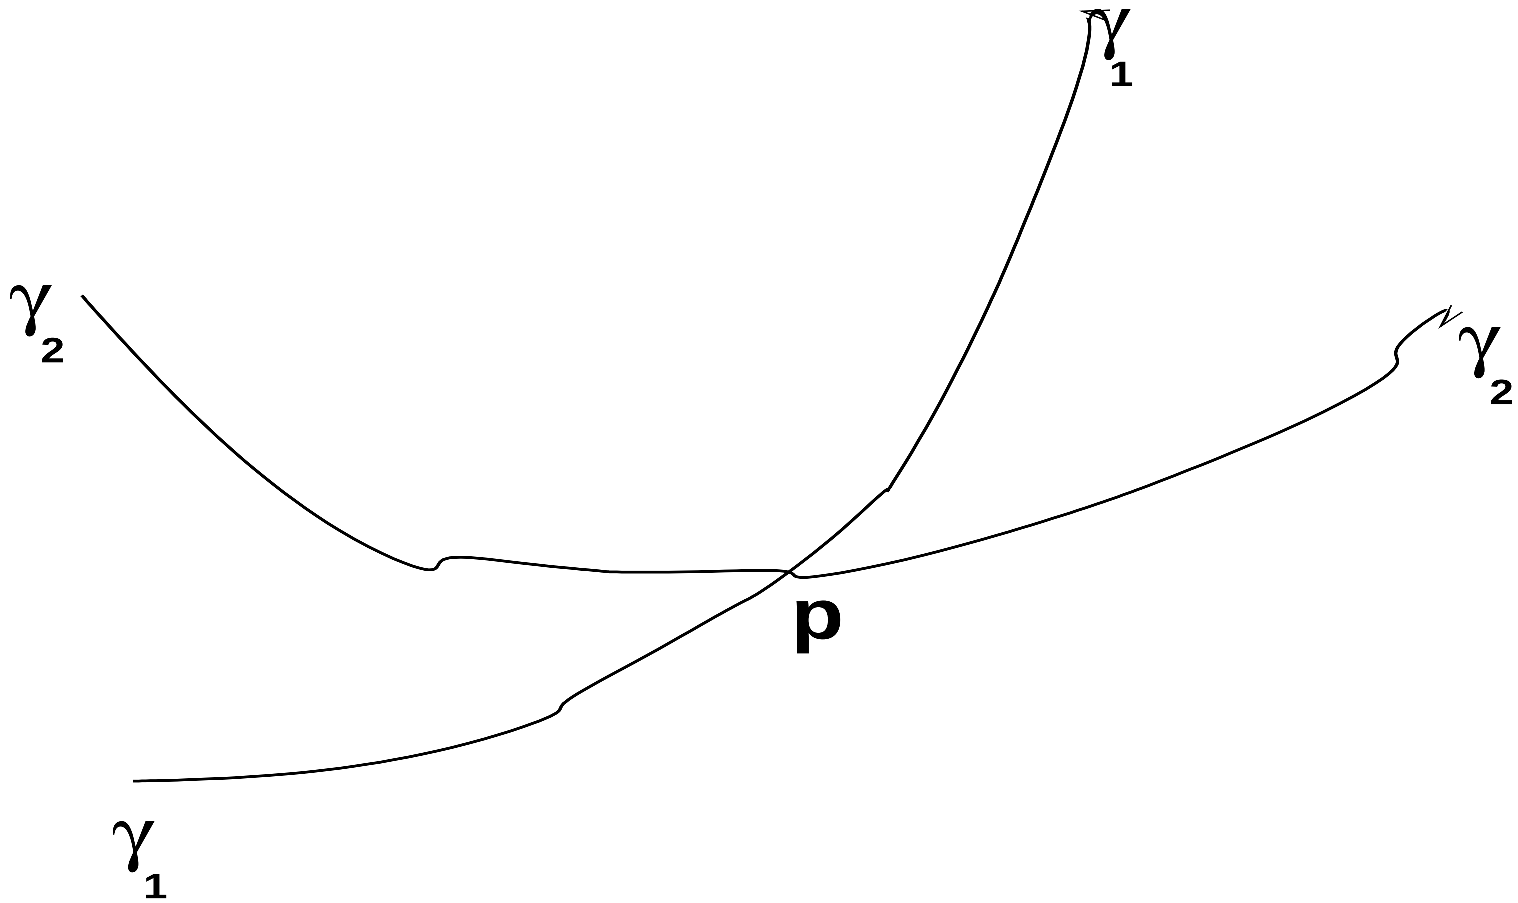

in which case all sufficiently small perturbations of γ_1, γ_2 will also intersect.

Transversality of the tangent approximations L_1, L_2 to γ_1 and γ_2 at an intersection point p is sufficient for the curves to intersect transversally:

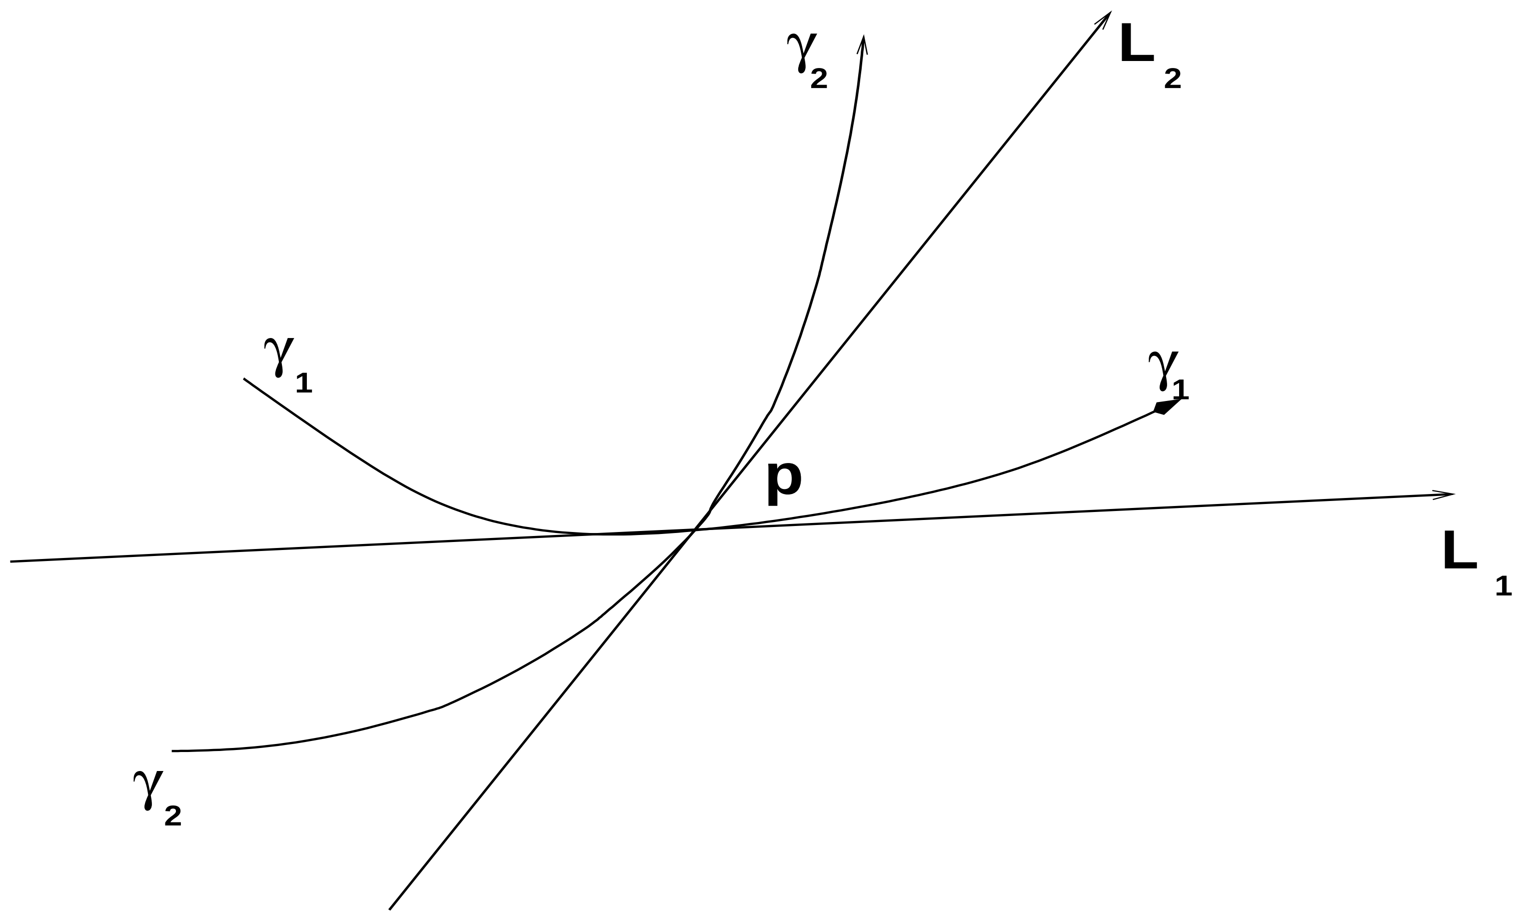

but is not necessary:

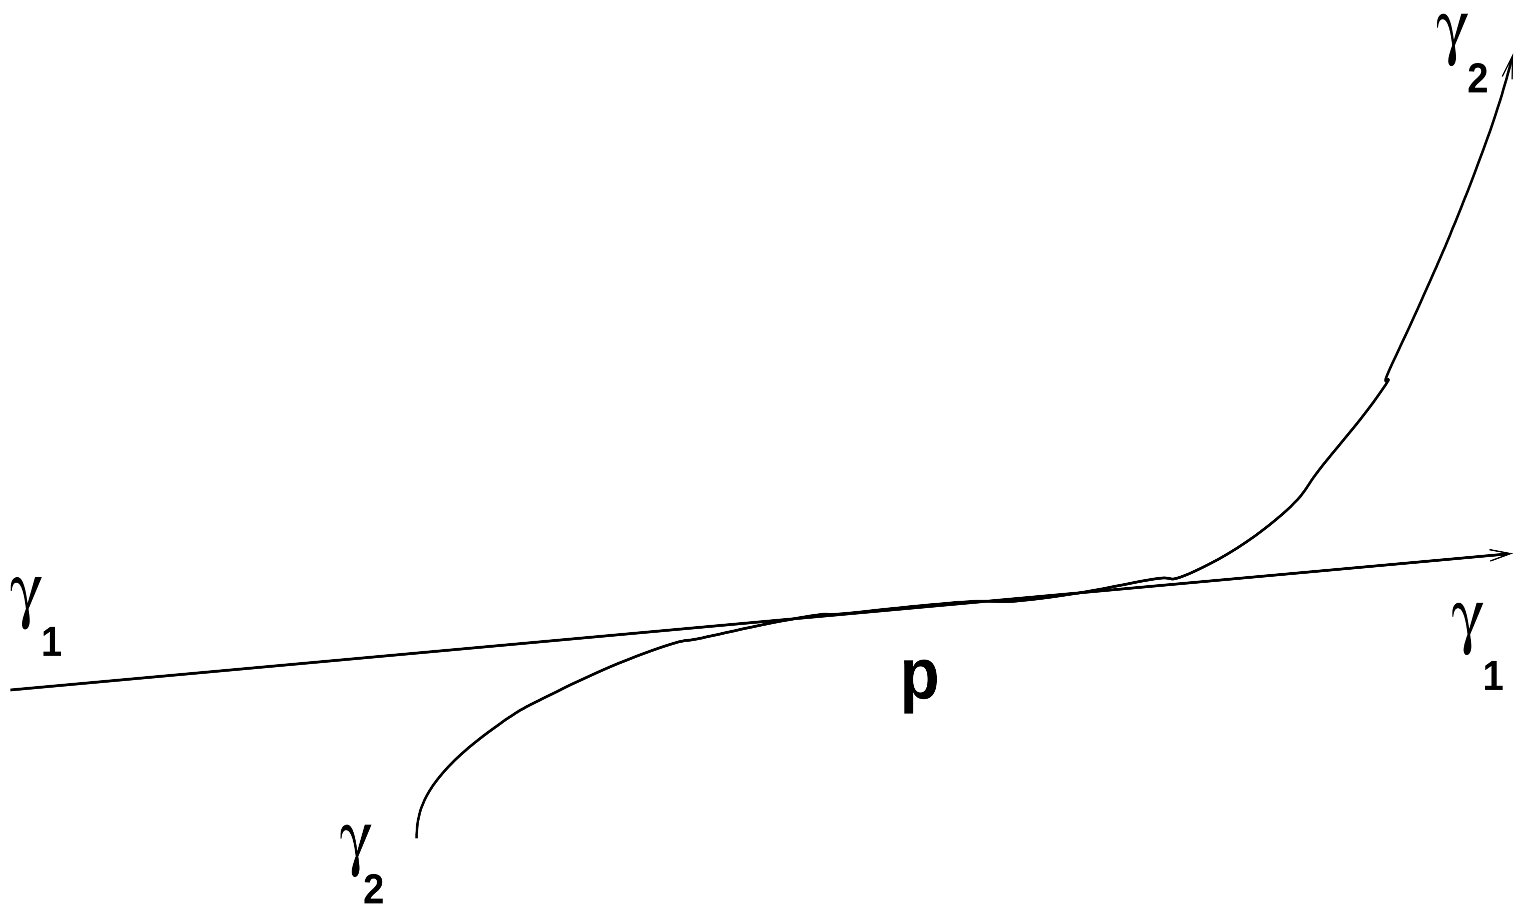

And the linear approximations that matter are approximations by tangent cones, not necessarily by tangent subspaces:

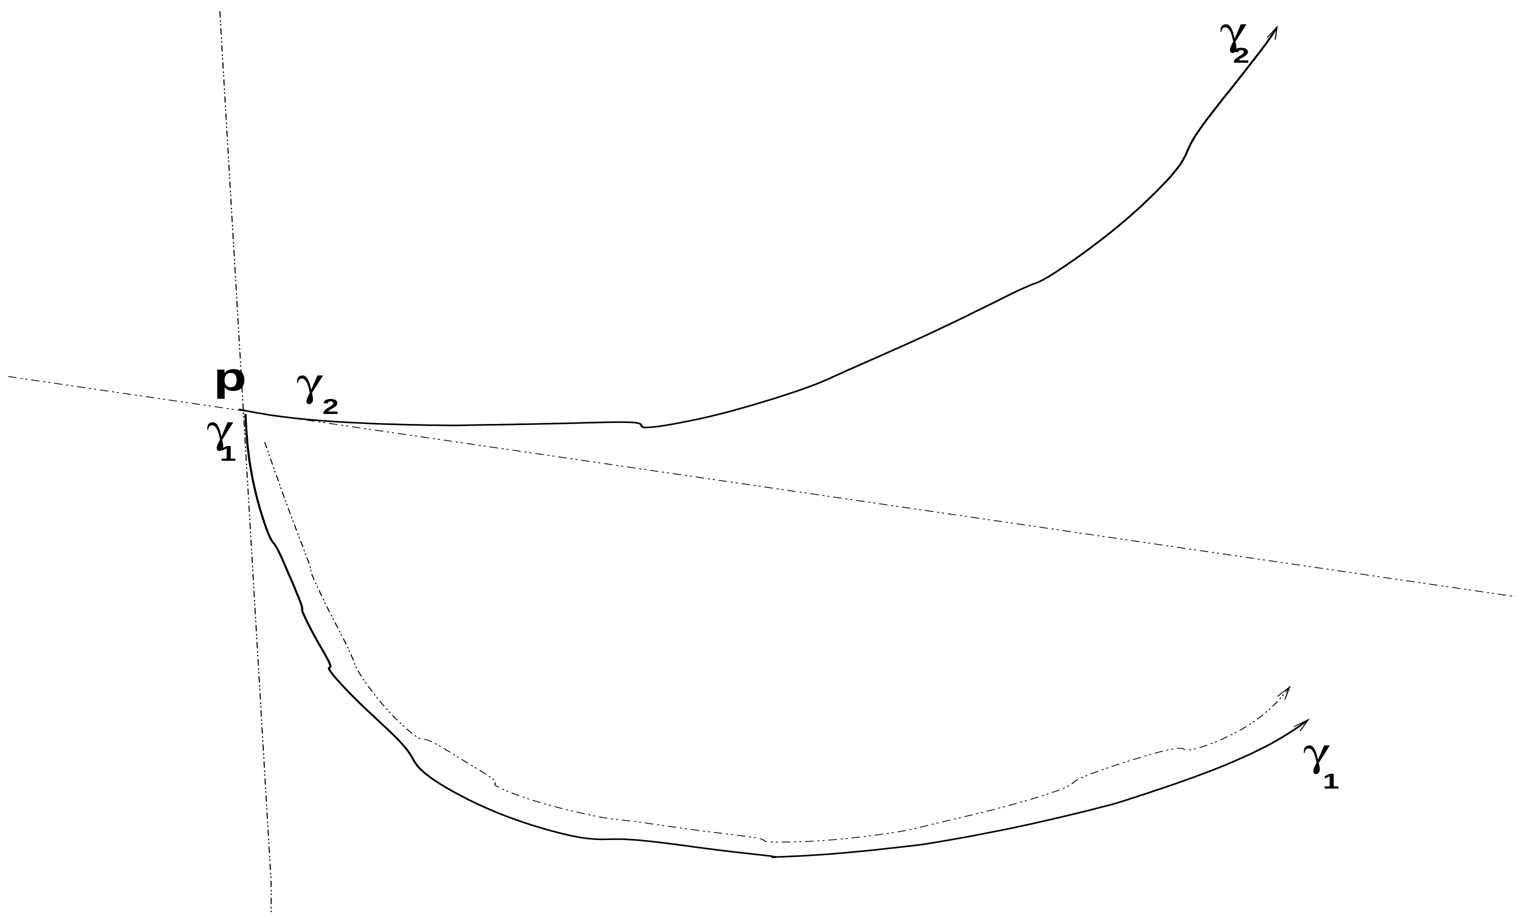

The situation is totally different for two curves in three-space:

If $\gamma_1 : [0,1] \mapsto \mathbb{R}^3$, $\gamma_2 : [0,1] \mapsto \mathbb{R}^3$ *are two continuous curves in* $\mathbb{R}^3$, *then for every positive* ε *there exist curves* $\tilde{\gamma}_1 : [0,1] \mapsto \mathbb{R}^3$, $\tilde{\gamma}_2 : [0,1] \mapsto \mathbb{R}^3$, *such that*

i. $\|\tilde{\gamma}_1 - \gamma_1\|_{sup} < \varepsilon$,

ii. $\|\tilde{\gamma}_2 - \gamma_2\|_{sup} < \varepsilon$,

and

iii. $\tilde{\gamma}_1$ *and* $\tilde{\gamma}_2$ *do not meet at all,*

that is,

iii'. $\tilde{\gamma}_1(t_1) \neq \tilde{\gamma}_2(t_2)$ *for all* $(t_1, t_2) \in [0,1] \times [0,1]$.

We consider pointed continuous maps (PCMs), that is, triples

$$(S, f, p)$$

where:

i. S is a topological space,

ii. f is a continuous map from S to some other topological space T,

iii. p is a point of S.

If S, f, p, T are as above, then we say that f, or the PCM (S, f, p), have target T.

DEFINITION: Let (S_1, f_1, p_1), (S_2, f_2, p_2), be PCMs with target T. Assume that T is a metric space. We say that (S_1, f_1, p_1) and (S_2, f_2, p_2) meet transversally if

(TR) *For every pair (N_1, N_2) consisting of neighborhoods N_j of p_j in S_j, there exists a positive real number ε such that, if $g_j : N_j \mapsto T$ are arbitrary continuous maps such that*

$$\text{dist}(g_j(x), f_j(x)) \leq \varepsilon \qquad \text{whenever } x \in N_j\,,\ j = 1, 2$$

(that is, g_1, g_2 are "ε-perturbations" of f_1, f_2 on N_1, N_2) it follows that g_1 and g_2 meet, that is, there exist $q_j \in N_j$ for which $g_1(q_1) = g_2(q_2)$.

An obvious necessary condition for (S_1, f_1, p_1) and (S_2, f_2, p_2)

to meet transversally is

$$f_1(p_1) = f_2(p_2)\,.$$

REMARK: Rather than assuming that T is a metric space, it would suffice to assume that T has a uniform structure, i.e., a structure that makes it possible to talk about two maps into T being "uniformly cloee".

For example, T could be a topological vector space. In that case, instead of talking about "ε-perturbations" of a map μ into T we would talk about "V-small perturbations", where V is a neighborhood of 0 in T: a map $\nu : S \mapsto T$ is a V-small perturbation of a map $\mu :\mapsto T$ if $\nu(s) - \mu(s) \in V$ for all $s \in S$.

If S is compact, then T can be an arbitrary toplogical space, because the space $C^0(S,T)$ of continuous maps has a natural topology (the compact-open topology).

LINEAR PCMs

A linear PCM is a PCM of the form $(D, L, 0)$, where

i. D is a convex cone,

ii. L is a linear map with target a real vector space T.

A REMARK ON THE DEFINITION OF "CONVEX CONE": A convex cone is a nonempty subset D of a real vector space V, which is closed under addition and multipication by nonnegative scalars. In particular, 0 always belongs to D.

The space V has a natural topology $\mathcal{T}_V$, namely, the one in which a subset Ω of V is open if and only if $\Omega \cap W$ is open in W for every finite-dimensional subspace W of V. (Also: (i) $\mathcal{T}_V$ is the inductive limit of the topologies of the finite-dimensional subspaces of V; (ii) $\mathcal{T}_V$ is the strongest topology that makes all the inclusion maps $W \ni w \mapsto w \in V$ continuous, for all finite-dimensional subspaces W of V.)
So D has a natural topology as well.

TRANSVERSALITY OF LINEAR PCMs

THEOREM Let T be a finite-dimensional real vector space, and let $(D_1, L_1, 0)$, $(D_2, L_2, 0)$, be linear PCMs with target T. Let $C_i = L_i D_i$ for $i = 1, 2$. Then $(D_1, L_1, 0)$ and $(D_2, L_2, 0)$ meet transversally if and only if

$$C_1 - C_2 = T\,. \qquad (1)$$

REMARK; The transversality condition (1) is equivalent to the following nonseparation condition:

(NS) There does not exist a nonzero lineal functional $\lambda : T \mapsto \mathbb{R}$ such that

$$\langle \lambda, c_1 \rangle \leq 0 \leq \langle \lambda, c_2 \rangle \qquad \text{for all } c_1 \in C_1\,,\ c_2 \in C_2\,. \qquad (2)$$

This theorem is not true for infinite-dimensional targets.

EXAMPLE: Let T be an infinite-dimensional Hilbert space.

Then, if $D_1 = T$, $L_1 = \mathrm{id}_T$, and $D_2 = \{0\}$, $L_2 = 0$, we have $C_1 = T$, $C_2 = \{0\}$, so the linear "transversality condition" $C_1 - C_2 = T$ is satisfied.

But $(D_1, L_1, 0)$ does not meet $(D_2, L_2, 0)$ transversally.

Reason: Using a continuous retraction of the unit ball $\mathbb{B}$ of T onto the unit sphere $\partial\mathbb{B}$ (which exists if T is infinite-dimensional) one can construct, for any positive ε, a sequence $B_1, B_2, \ldots$ of pairwise disjoint balls in T that converge to zero, and retractions $\rho_j : B_j \mapsto \partial B_j$, thus obtaining a continuous map $M_\varepsilon : T \mapsto T$ which is an ε-perturbation of id_T, and a sequence of points p_j that are not in the image of M_ε. And, for large enough j, these points are ε-perturbations of 0.

PROOF THAT THE CONDITION $C_1 - C_2 = T$ IS NECESSARY FOR TRANSVERSALITY:

Assume $(D_1, L_1, 0)$ and $(D_2, L_2, 0)$ meet transversally. Then in particular if $v \in T$ is sufficiently small the cones $C_1 + v$ and C_2 must intersect.

So there exist $c_1 \in C_1$, $c_2 \in C_2$, such that $c_1 + v = c_2$.

Then $v = c_1 - c_2$. So $v \in C_1 - C_2$.

Hence the convex cone $C_1 - C_2$ contains a neighborhood of 0 in T.

So $C_1 - C_2 = T$.

Q.E.D.

The proof that the condition $C_1 - C_2 = T$ is sufficient for $(D_1, L_1, 0)$ and $(D_2, L_2, 0)$ to meet transversally is not very hard, but it needs some work.

Furthermore, the proof yields a somewhat stronger conclusion:

(*) If $C_1 - C_2 = T$, then there exist finitely spanned subcones $\tilde{D}_1$. $\tilde{D}_2$ of D_1, D_2 such that the PCMs $(\tilde{D}_1, \tilde{L}_1, 0)$ and $(\tilde{D}_2, \tilde{L}_2, 0)$ (where $\tilde{L}_j$ is the restriction of L_j to $\tilde{D}_j$) meet transversally with a linear rate.

FINITELY SPANNED: A convex subcone $\tilde{D}$ of a convex cone D is finitely spanned it there exists a finite subset F of D such that $\tilde{D}$ is the convex cone spanned by F.

LINEAR RATE: There exists a positive constant K such that, for all sufficiently small positive δ, if $N_j(\delta)$ is the δ-neighborhood of 0 in $\tilde{C}_j$, then

(#) If $\varepsilon = K\delta$, then, if g_1, g_2 are continuous ε-perturbations of $\tilde{L}_1, \tilde{L}_2$ on $N_1(\delta)$, $N_2(\delta)$, then g_1 and g_2 meet (that is, there exist $q_j \in N_j(\delta)$ such that $g_1(q_1) = g_2(q_2)$).

PARTIAL LINEARIZATIONS

Let (S, f, p) be a PCM with finite-dimensional target T. A partial linearization of (S, f, p) is a pair $\Big((D, L, 0), \mu\Big)$, such that

i. $(D, L, 0)$ is a linear PCM,

ii. μ is a continuous map from some neighborhood Dom(μ) of 0 in D into S,

iii. $\mu(0) = p$,

iii. the map $f_\mu \stackrel{\text{def}}{=} f \circ \mu$ satisfies

$$f_\mu(x)) - f_\mu(0) = Lx + o(\|x\|) \qquad \text{as } x \to 0\,,\ x \in \text{Dom}(\mu)\,. \tag{3}$$

i.e.,

$$f(\mu(x)) - f(p) = Lx + o(\|x\|) \qquad \text{as } x \to 0\,,\ x \in \text{Dom}(\mu)\,. \tag{4}$$

PARTIAL LINEARIZATIONS

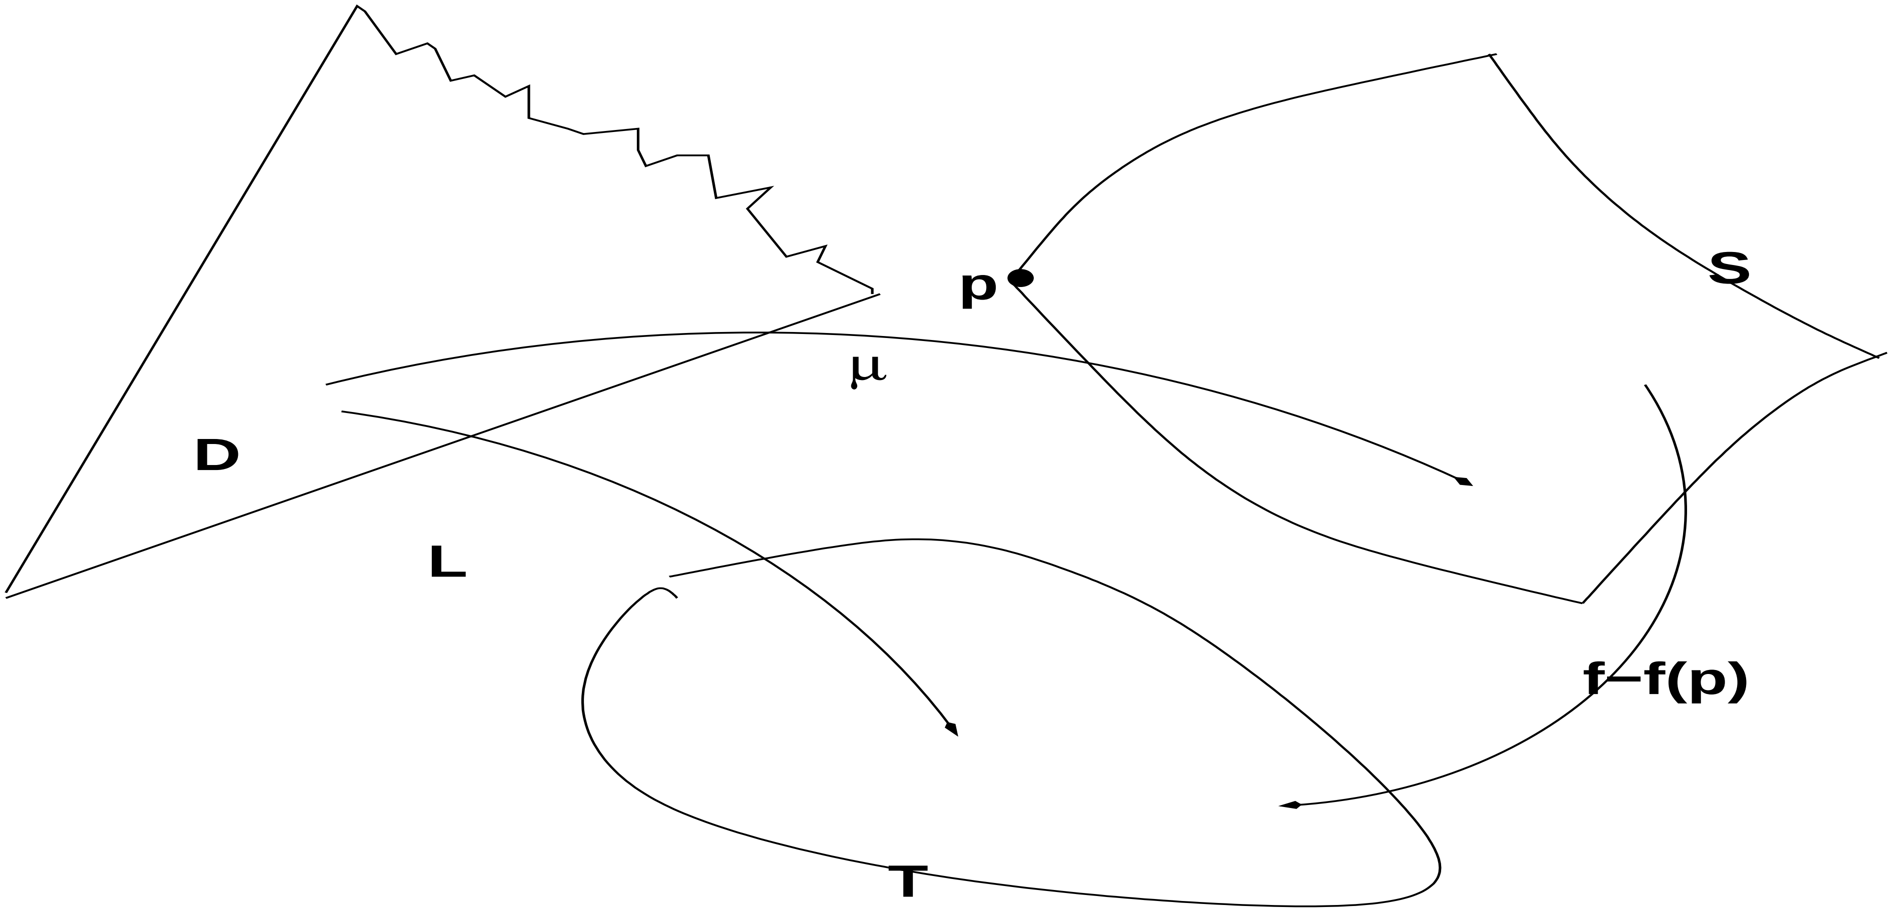

APPROXIMATING CONES

If S is a subset of a finite-dimensional real vector space V, an approximating cone (a.k.a. "Boltyanskii approximating cone") to S at a point p if S is a convex cone C in T which is the image of a partial linearization $((D, L, 0), \mu)$ of the identity map id_S.

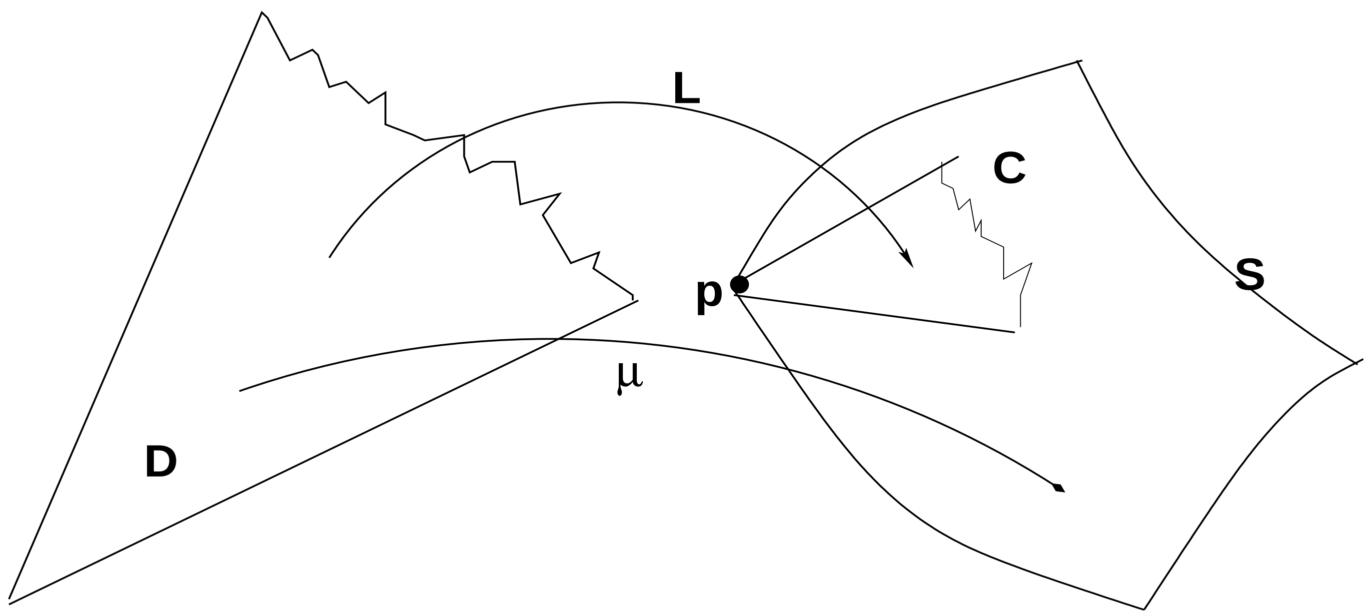

THEOREM. If

(i.) (S_1, f_1, p_1) and (S_2, f_2, p_2) are PCMs with the same finite-dimensional target T,

(ii.) $(D_1, L_1, 0)$, $(D_2, L_2, 0)$ are partial linearizations of (S_1, f_1, p_1) and (S_2, f_2, p_2),

(iii.) $(D_1, L_1, 0)$ and $(D_2, L_2, 0)$ meet transversally,

then

(&) (S_1, f_1, p_1) and (S_2, f_2, p_2) meet transversally.

PROOF: Trivial. Q.E.D.

THE MAXIMUM PRINCIPLE (MP) I: THE DATA

I. We are given

I.a: a state space Ω,

I.b: a control set U,

I.c: a time interval $[a, b]$,

I.d: a class of admissible controls $\mathcal{U}$, consisting of functions $\eta : [a, b] \mapsto U$,

I.e: a controlled dynamical law, i.e., a differential equation

$$\dot{x} = f(x, u, t)\,, \tag{5}$$

I.f: an initial condition, i.e., a point $x_* \in \Omega$,

I.g: a reference trajectory-control pair, i.e., a pair (ξ_*, η_*), such that $\eta_* \in \mathcal{U}$, $\xi_* \in W^{1,1}([a, b]; \Omega)$, and the curve ξ_* satisfies $\xi_*(a) = x_*$ and

$$\dot{\xi}_*(t) = f(\xi_*(t), \eta_*(t), t) \quad \text{for a.e. } t \in [a, b]\,,$$

I.h: a subset S of Ω.

THE MAXIMUM PRINCIPLE (MP) II: THE MAPS

One then defines:

II.a: the (possibly set-valued) control-to-trajectory maps

$$\mathcal{C}\text{–}\mathcal{TR}_{x_*} : \mathcal{U} \mapsto W^{1,1}([a,b],\Omega)\,,$$

that assign to each control $\eta \in \mathcal{U}$ the trajectory (or the set of trajectories) ξ for η with initial condition $\xi(a) = x$, i.e., the solution (or set of solutions) ξ of

$$\begin{aligned} \dot{\xi}(t) &= f(\xi(t), \eta(t), t) \quad \text{a.e}\,, \\ \xi(a) &= x_*\,. \end{aligned}$$

II.b: the endpoint map $\mathcal{E} : W^{1,1}([a,b],\Omega) \mapsto \Omega$, defined by

$$\mathcal{E}(\xi) = \xi(b) \qquad \text{for}\,, \ \xi \in W^{1,1}([a,b],\Omega)\,.$$

II.c the (also possibly set-valued) control-to-terminal-point maps $\mathcal{C}\text{–}\mathcal{TP}_{x_*} : \mathcal{U} \mapsto \Omega$ defined by

$$\mathcal{C}\text{–}\mathcal{TP}_{x_*} = \mathcal{E} \circ \mathcal{C}\text{–}\mathcal{TR}_{x_*}\,.$$

THE MAXIMUM PRINCIPLE (MP) III: THE CONCLUSION

The (MP) gives a sufficient condition for the following transversal intersection property:

(TI) $\boxed{\text{The PCM } (\mathcal{U}, \mathcal{C}\text{–}\mathcal{TP}_{x_*}, \xi_*(b)) \text{ meets } (S, \mathrm{id}_S, \xi_*(b)) \text{ transversally.}}$

Actually, (MP) gives a sufficient condition for a stronger property, namely, finite-dimensional transversal intersection

(FDTI) $\boxed{\text{There exists a continuous map } V \text{ from the unit ball } \mathbb{B}^n \text{ to } \mathcal{U} \text{ such that } V(0) = \xi_*(b) \text{ such that } (\mathbb{B}^n, \mathcal{C}\text{–}\mathcal{TP} \circ V, 0) \text{ and } (S, \mathrm{id}_S, \xi_*(b)) \text{ transversally.}}$

THE MAXIMUM PRINCIPLE (MP) IV: THE SUFFICIENT CONDITION

The sufficient condition has the form

(SC) $\boxed{\text{There does not exist a nontrivial Hamiltonian-maximizing adjoint vector } \lambda \text{ along } (\xi_*, \eta_*) \text{ that satisfies the transversality condition } -\lambda(b) \in C^\perp, \text{ where } C \text{ is an approximating cone of } S \text{ at } \xi_*(b) \text{ and } C^\perp \text{ is the polar cone of } C.}$

(Naturally, this requires having a topology on $\mathcal{U}$. This will be discussed later.)

EXAMPLE 1: LOCAL CONTROLLABILITY

Assume S is just the single point $\xi_*(b)$. Then transversality of $(\mathcal{U}, \mathcal{C}\text{–}\mathcal{TP}_{x*}, \xi_*(b))$ to $\xi_*(b)$ implies:

(LC) $\boxed{\text{The reachable set from } \xi_*(a) \text{ over } [a,b] \text{ contains a full neighborhood of } \xi_*(b),}$

i.e., local controllability along ξ_*.

REASON: Every suffiently small perturbation of $\xi_*(b)$ must meet the reachable set. That is, the reachable set must contain every point in some ε-neighborhood of $\xi_*(b)$.

EXAMPLE 2: OPTIMALITY

Assume we want to minimize an integral $\int_a^b L(\xi(t), \eta(t), t)dt$ among all trajectory-control pairs (ξ, η) such that $\xi(b) \in S$.

In this case we apply the MP to the augmented control system

$$\begin{aligned} \dot{\xi}(t) &= f(\xi(t), \eta(t), t)\,, \\ \dot{\xi}_0(t) &= L(\xi(t), \eta(t), t)\,, \end{aligned}$$

and the terminal set

$$\widehat{S} = (-\infty, c_*] \times S\,,$$

where c_* is the cost along the reference trajectory, that is,

$$c_* = \int_a^b L(\xi_*(t), \eta_*(t), t)dt\,.$$

Transversality of the control-to-terminal-point map $\widehat{\mathcal{C}\text{–}\mathcal{TP}}_{x_*,0}$ of the agumented system to the set $\widehat{S}$ implies that $\widehat{\mathcal{C}\text{–}\mathcal{TP}}_{x_*,0}$ meets every sufficiently small perturbation of the identity map of $\widehat{S}$. In particular, we can consider the map

$$(k, s) \mapsto (k - \varepsilon, s)$$

for a small positive ε. Then $\widehat{\mathcal{C}\text{–}\mathcal{TP}}_{x_*,0}$ meets this map, so there exists a point $(c, q) \in \widehat{\mathcal{C}\text{–}\mathcal{TP}}_{x_*,0}(\mathcal{U})$ such that $c = k - \varepsilon$, $q = s$ for some $(k, s) \in \widehat{S}$.

Then $k \le c_*$, so $c < c_*$. So q is reachable from $\xi_*(a)$ over $[a, b]$ with cost $c < c_*$, and $q \in S$. So (ξ_*, η_*) is not optimal.

Hence the MP gives a sufficient condition for non-optimality, which is of course equivalent to a necessary condition for optimality.

The necessary condition for optimality has the form:

(NCO) There exist a nontrivial Hamiltonian-maximizing adjoint vector $(\hat{\lambda}_0, \lambda)$ along $(\xi_{0,*}, \xi_*, \eta_*)$ that satisfies the transversality condition

$$-\hat{\lambda}(b) \in C^{\perp}\,, \ -\hat{\lambda}_0 \geq 0\,,$$

If we write $\lambda_0 = -\hat{\lambda}_0$ (so the control theory Hamiltonian H beecomes

$$H = \langle \lambda, f(x,u,t) \rangle - \lambda_0 L(x,u,t)$$

(i.e., "$H =$momentum times velocity minus Lagrangian", as in Physics), then the transversality condition for λ_0 takes the familiar form

$$\lambda_0 \geq 0\,.$$

So we see that in both cases, local controllabilty and optimal control, the transversality version of the MP implies, but is not equivalent to, the usual versions, namely,

a. a sufficient condition for local controllability along a trajectory,

b. a necessary condition for optimality.

We now explore one of the stronger consequences of the MP.

THE APPROXIMATION PROBLEM

PROBLEM: Given a dense subset $\mathcal{U}_0$ of the set $\mathcal{U}$ of admissible controls, we want to know which trajectories ξ, corresponding to controls $\eta \in \mathcal{U}$, can be approximated by trajectories corresponding to a control in a subset $\mathcal{U}_0$ of $\mathcal{U}$ **with the same endpoints as** ξ.

Obviously, for this question to make sense we need, to begin with, an f-adequate topology on $\mathcal{U}$, i.e., a topology on such that

(C)

1. The "initial condition and control to trajecory map", i.e. the map

$$\Omega \times [a,b] \times \mathcal{U} \ni (x,t,\eta) \mapsto \xi_{x,t,\eta} \in C^0([a,b],\Omega) \tag{6}$$

(where $\xi_{x,\eta}$ is the trajectory for the control η such that $\xi(t)=x$) is continuous.

2. The map

$$\mathcal{U} \times \mathcal{U} \times [a,b] \ni (\eta,\zeta,t) \mapsto \zeta \#_t \eta \in \mathcal{U}\,, \tag{7}$$

where

$$(\zeta \#_t \eta)(s) = \begin{cases} \eta(s) & \text{if} \quad s < t \\ \zeta(s) & \text{if} \quad s \geq t \end{cases},$$

is continuous.

With any such topology,

1. the "packets of needle variations" used is the classical proof of the Maximum Principle, which are maps
$$\mathbb{R}^m_+(\varepsilon_*) \ni (\varepsilon_1, \ldots, \varepsilon_m) \mapsto V(\varepsilon_1, \ldots, \varepsilon_m) \in \mathcal{U}$$
(where $\mathbb{R}^m_+(\varepsilon_*) = \{\, (\varepsilon_1, \ldots, \varepsilon_m) \in \mathbb{R}^n_+ : \varepsilon_1 + \ldots + \varepsilon_m \leq \varepsilon_* \,\}$), are continuous,

2. the "parameter-to-trajectory" maps $\mathcal{C}\text{–}\mathcal{TR} \circ V$ are continuous,

3. the "parameter-to-terminal-point" maps $\mathcal{C}\text{–}\mathcal{TP} \circ V$ are continuous.

A natural conjecture (F. Rampazzo, M. Palladino) is:

(RP) If $\mathcal{U}_0$ is dense in $\mathcal{U}$ in an appropriate way, then for every trajectory-control pair (ξ_*, η_*) that satisfies (SC), there exist controls $\eta_n \in \mathcal{U}_0$ that converge to η_* aad are such that the correesponding trajectories ξ_n with initial condition $\xi_n(a) = \xi_*(a)$ satisfy $\xi_n(b) = \xi_*(b)$.

(SC) There does not exist a nontrivial Hamiltonian-maximizing adjoint vector λ along (ξ_*, η_*) that satisfies the transversality condition $-\lambda(b) \in C^\perp$, where C is an approximating cone of S at $\xi_*(b)$ and $C^\perp$ is the polar cone of C.

A natural conjecture is:

(RP) $\boxed{\text{If } \mathcal{U}_0 \text{ is dense in } \mathcal{U} \text{ in an appropriate way, then ...}}$

If $\mathcal{U}_0$ is dense in $\mathcal{U}$ in an appropriate way, then for every trajectory-control pair (ξ_*, η_*) that satisfies (SC), there exist controls $\eta_n \in \mathcal{U}_0$ that converge to η_* aad are such that the corresponding trajectories ξ_n with initial condition $\xi_n(a) = \xi_*(a)$ satisfy $\xi_n(b) = \xi_*(b)$.

(SC) There does not exist a nontrivial Hamiltonian-maximizing adjoint vector λ along (ξ_*, η_*) that satisfies the transversality condition $-\lambda(b) \in C^\perp$, where C is an approximating cone of S at $\xi_*(b)$ and $C^\perp$ is the polar cone of C.

The conjecture is true, if the word “dense” is interpreted in an appropriate way”

Just “dense” is not enough.

EXAMPLE: Consider the system

$$\dot{x} = u, \qquad -1 \le u \le 1\,,$$

on $\mathbb{R}$. Let $\mathcal{U}$ be the class of all measurable functions $\eta : [0,1] \mapsto [-1,1]$. Then 0 is reachable from 0 over the interval $[0,1]$, using any control $\eta : [0,1] \mapsto [-1,1]$ that satisfies

$$\int_0^1 \eta(t)dt = 0\,. \tag{8}$$

Let $\mathcal{U}_0$ be the class of all $\eta \in \mathcal{U}$ that do not satisfy (8). Then $\mathcal{U}_0$ is obviously dense in $\mathcal{U}$, but no trajectory for a control in $\mathcal{U}_0$ will go from 0 to 0 over $[0,1]$.

But: the trajectory-control pair $(0,0)$ satisfies (SC). (Reason: The Hamiltonian $H(x,\lambda,u)$ is λu. For this to be maximimzed at $u=0$ we need $\lambda = 0$. So (SC) holds.)

CONCLUSION: we need something stronger that "density".

In the previous example, the set $\mathcal{U}_0$ is the complement of $\mathcal{B}$, where $\mathcal{B} = \{\eta \in \mathcal{U} : \int_0^1 \eta(s)ds = 0\}$.

So $\mathcal{U}_0$ is the complement of the "bad" set $\mathcal{B}$, where $\mathcal{B}$ is defined by one condition, namely, $\int_0^1 \eta(s)ds = 0$.

So $\mathcal{B}$ is a subset of $\mathcal{U}$ of codimension one.

This means that in general a continuous curve $\gamma : [0,1] \mapsto \mathcal{U}$ cannot be approximated by curves $\gamma_0 : [0,1] \mapsto \mathcal{U}_0$.

In particular, it is easy to construct continuous curves $\gamma : [-1,1] \mapsto \mathcal{U}$ such that the control $\gamma(s)$ steers 0 to s in time 1. (Just let $\gamma(s)$ be the constant control with value s.)

If we could approximate one of these curves by a $\mathcal{U}_0$-valued curve γ_0, this would give us a continuous curve $[-1,1] \ni s \mapsto \mathcal{C}\text{–}\mathcal{TP}_0(\gamma_0(s))$ of terminal points and, by continuity, one of these terminal points would have to be 0.

But we cannot approximate γ by $\mathcal{U}_0$-valued curves.

This suggests an idea:

The conjecture should be true if $\mathcal{U}$-valued curves, or, more generally, continuous maps from finite-dimensional balls $\mathbb{B}^\nu$ to $\mathcal{U}$, can be approximated by $\mathcal{U}_0$-valued maps.

So the extra hypothesis on $\mathcal{U}_0$ must be:

(EH) $\boxed{\mathcal{U}_0 \text{ is the complement of a subset of sufficiently high codimension.}}$

Let us make this precise.

DEFINITION: Assume that

1. X is a topological space,
2. X_0 is a subset of X,
3. ν is a nonnegative integer.

We say that X_0 is <u>ν-dense in X</u> if

(*) $C^0(\mathbb{B}^\nu, X_0)$ is dense in $C^0(\mathbb{B}^\nu, X)$,

i.e., if

(**) every continuous map $\theta : \mathbb{B}^\nu \mapsto X$ is a uniform limit of continuous maps $\theta_\alpha : \mathbb{B}^\nu \mapsto X_0$. ■

REMARK: If M is a smooth manifold and S is a smooth submanifold of M, then the complement of S is ν-dense if $\text{codim}(S) > \nu$.

EXAMPLES:

1. X_0 is 0-dense iff X_0 is dense.

2. X_0 is 1-dense in X if every curve in X can be appproximated by curves in X_0.

3. If $X = \mathbb{R}^2$, and X_0 is the complement of a line, then X_0 is dense but is not 1-dense.

4. On the other hand, if $X = \mathbb{R}^2$ and X_0 is the complement of a point then X_0 is 1-dense but not 2-dense.

DEFINITION (Jack Warga): Assume that

1. X is a topological space,

2. X_0 is a subset of X,

We say that X_0 is an <u>abundant subset of X</u> is X_0 is ν-dense in X for every nonnegative integer ν. ■

A sufficient condition for a subset X_0 to be abundant in X is:

(R) The identity map id_X is a limit of continuous maps $\Phi_\alpha : X \mapsto X_0$, uniformly on compact sets.

The precise meaning of the convergence condition in (R), for a net $\{\Phi_\alpha\}_{\alpha \in A}$, is: *For every compact subset K of X and every neighborhood Ω in $X \times X$ of the set $\{(x, x) : x \in K\}$, there exists α_* such that $\{(x, \Phi_\alpha(x)) : x \in K\} \subseteq \Omega$ whenever $\alpha_* \preceq \alpha$.*

THIS

WORKS

EXAMPLE: RELAXED CONTROLS

Assume that

1. U is a compact metrizable space.
2. $\mathcal{U}_0$ is the class of all measurable maps $\eta : [a, b] \mapsto U$.
3. $f(x, u, t)$ is (jointly) continuous with respect to x, u for each fiexed t,
4. $f(x, u, t)$ is measurable with respect to t for each fixed x, u,
5. f satisfies Carathéodory-Lipschitz bounds

$$\|f(x, u, t)\| \leq C_K(t)\,, \tag{9}$$

$$\|f(x, u, t) - f(y, u, t)\| \leq C_K(t)\|x - y\| \tag{10}$$

for all $(x, y, u, t) \in K \times K \times U \times [a, b]$, for every compact subset K of Ω, where $C_K \in L^1([a, b], \mathbb{R})$.

THE RELAXED CONTROL VALUES

We let $\mathbf{P}(U)$ be the set of all Borel probablity measures on U, and for $\mu \in \mathbf{P}(U)$, we define

$$f(x, \mu, t) = \int_U f(x, u, t) d\mu(u) \,.$$

Then $\mathbf{P}(u)$ is a weak*-closed, bounded subset of $C^0(U, \mathbb{R})^\perp$, the dual of the Banach space $C^0(U, \mathbb{R})$.

So $\mathbf{P}(u)$, equipped with the weak* topology, is compact and metrizable.

THE RELAXED CONTROLS

We let $\mathcal{U}$ be the set of all bounded, measurable functions from $[a,b]$ to $\mathbf{P}(U)$. (The members of $\mathcal{U}$ are the "relaxed controls".)

Then $\mathcal{U}$ has a natural topology $\mathcal{T}_{\mathcal{U}}$, namely, the weakest topology on $\mathcal{U}$ that makes all the maps

$$\mathcal{U} \ni \eta \mapsto \int_a^b \Big[\int_U \varphi(u,t) d\eta(t)(u) \Big] dt$$

continuous, for all functions $\varphi \in L^1([a,b], C^0(U,\mathbb{R}))$, i.e., all functions

$$U \times [a,b] \ni (u,t) \mapsto \varphi(u,t) \in \mathbb{R}$$

that are continuous with respect to u for each t, measurable with respect to t for each u, and such that

$$\int_a^b \|\varphi(\cdot,t)\|_{sup} dt < \infty\,.$$

(Basically, $\mathcal{U}$ is a subset of the unit ball of $L^\infty([a,b], C^0(U,,\mathbb{R})^\perp)$, the dual of $L^1([a,b], C^0(U,\mathbb{R}))$. And $\mathcal{T}_f$ is the weak* topology arising from this duality.)

The topology $\mathcal{T}_{\mathcal{U}}$ is an f-adequate control topology.

And we get:

THEOREM: $\mathcal{U}_0$ is an abundant subset of $\mathcal{U}$.

SKETCH OF THE PROOF: Take your favorite way of approximating relaxed controls by ordinary ones, and verify that it yields a sequence $\{\Phi_k\}_{k=1}^{\infty}$ of continuous maps $\Phi_k : \mathcal{U} \mapsto \mathcal{U}_0$ such that $\Phi_k(\eta) \to \eta$ uniformly as $k \to \infty$.

Then, using the Maximum Principle, we get:

THEOREM: If a trajectory-control pair (ξ_*, η_*) for a relaxed control η_* is such that there does not exist a nontrivial Hamiltonian-maximizing adjoint vector that satisfies the transversality condition $-\lambda(b) \in C^\perp$ (where C is an approximating cone to S), then ξ_* can be approximated by trajectories η_k ccrresponding to ordinary controls η_k in such a way that $\xi_k(b) \in S$.

PROOF: The hypothesis implies that there exists a control vaiation $V : \mathbb{R}^{\nu}_{+}(\varepsilon_*) \mapsto \mathcal{U}$ such that the map $\mathcal{C}\text{–}\mathcal{TP} \circ V$ meets S transversally.

Let V_k be continuous maps from $\mathbb{R}^{\nu}_{+}(\varepsilon_*)$ to $\mathcal{U}_0$ that converge to V. Then for sufficienctly large k the maps $\mathcal{C}\text{–}\mathcal{TP} \circ V_k$ are small perturbations of $\mathcal{C}\text{–}\mathcal{TP} \circ V$, so they meet S.