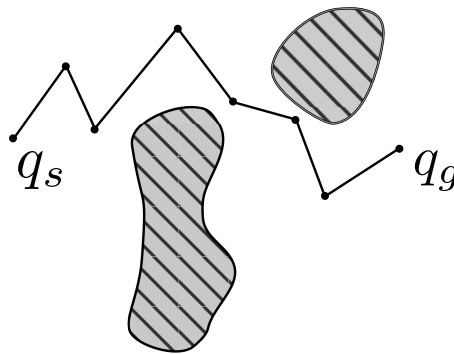


Lecture 5: Trajectory Optimization

Scribes: Esther Rolf, Matthew Matl

5.1 Overview

Last lecture, we saw that randomized planners can get us from a start configuration to a goal configuration while overcoming some of the issues associated with high dimensionality. In practice, RRT runs quickly (milliseconds to seconds in a 7 d.o.f. robot). However, there are still several drawbacks of RRT. It is only probabilistically complete, which is a rather weak assertion, as we discussed last time. RRT is also not optimal, and as we see in practice, it is *very* non-optimal. For example, consider the figure below, which shows an example path that RRT might generate in C-space.



This path in C-space will zig-zag randomly, and the corresponding motion in world-space can be very jerky and unintuitive, which is a problem when robots are working with humans. One way to address this is to use a post-processing step called *shortcutting*, which samples pairs of points on the path returned by RRT and tries to connect them with a straight line in C-space. If the connection will not cause a collision, it is added to the path and redundant vertices are removed so that the overall path is shorter. This takes a long time, can still result in paths that are in the wrong homotopy class, can still result in jerky and counterintuitive paths.

5.2 Optimal Motion Planning

The idea behind optimal motion planning (optimal randomized sampling based planners) is to leverage the benefits of randomized sampling (namely, reaching a feasible solution quickly), with asymptotic optimality. Here, asymptotic optimality means that as the number of iterations goes to infinity, the resulting trajectory converges to an optimal one.

An example of an optimal motion planning algorithm is RRT*, introduced by [Karman, Frazzoli '2010]. RRT* is based on unidirectional RRT with two modifications:

1. Parent selection. When a new point is sampled, rather than making it a child of the closest node in the graph so far, we connect it to the point that would produce the shortest path from the starting point.
2. Rewiring Step. Once a new point x_s is sampled and added to the graph, consider whether we can improve the existing graph using this new point. For every node x_v in the vicinity of the new point, check if that point can be reached faster by an alternate path using x_s . If so, make x_s the parent of x_v , and delete all redundant edges so that the graph is again a tree.

RRT* is an asymptotically optimal algorithm. However, the convergence to the optimal solution is slow, and the algorithm itself can be slow due to all the extra graph modification steps. While RRT may be more feasible for use in practice, RRT* is an important theoretical result.

5.3 Trajectory Optimization

When robots are working in an environment with people, it's important that they move in an understandable way, and in a way that humans can anticipate. Trajectory optimization can get us closer to that goal. The idea behind trajectory optimization is to start with a simple path, and let the obstacles guide you away from collision while optimizing for efficiency/smoothness or other costs. The resulting algorithm is:

- probabilistically complete
- fast
- locally optimal

The problem statement of trajectory optimization is to find a trajectory ζ which maps time to C-space configurations:

$$\zeta : [0, T] \rightarrow C$$

Here $\zeta \in \Xi$ is a particular trajectory among a set of possible trajectories Ξ . A trajectory ζ is a function of $t \in [0, T]$ which produces a C-space configuration $q \in C$. Thus, a trajectory tells the robot where to be for each point in time.

There are many possible trajectories in the set Ξ . To distinguish between them, we define a *cost functional* (also called an objective functional) \mathcal{U} , which maps trajectories onto non-negative scalars¹ :

$$\mathcal{U} : \Xi \rightarrow \mathbb{R}^+$$

Our goal is to optimize \mathcal{U} , resulting in an optimal trajectory, where optimal is defined according to the factors encoded in \mathcal{U} . Some potential factors encoded in \mathcal{U} include:


- length or efficiency of the resulting path (in C-space or world-space)
- obstacle avoidance
- smoothness
- caution/ slow-down near obstacles
- energy minimization

¹ \mathcal{U} maps functions to scalars, and thus is a *functional*.

- “naturalness”
- information gathering
- predictability
- legibility/ intent expression

Before we define the technical optimization problem, consider the worlds in which each of these elements live. A robot’s pose, $R(q)$, in world space $W(\mathbb{R}^3)$ is a single point, q , in C-space $C(D)$, where D is robot’s number of degrees of freedom. A trajectory, ζ , is a timed path through C-space and is a single point in trajectory-space, Ξ . Ξ is a space of functions and is thus infinite-dimensional.

In our optimization problem, we will seek to find an optimal $\zeta \in \Xi$. We will be using gradient descent on \mathcal{U} .



Definition

Trajectory optimization is the process of finding an optimal trajectory ζ^* :


$$\begin{aligned} \zeta^* &= \arg \min_{\zeta \in \Xi} \mathcal{U}[\zeta] \\ \text{s.t.} \quad &\zeta(0) = q_s \\ &\zeta(T) = q_g \end{aligned}$$

One method for solving this optimization problem is gradient descent (GD), which iteratively updates the guess for an optimal ζ :

$$\zeta_{t+1} = \zeta_t - \frac{1}{\alpha} \nabla_{\zeta} \mathcal{U}(\zeta_t)$$

It is important to note that gradient descent returns a local minimizer, not a global minimizer of the cost functional. We can minimize globally for some \mathcal{U} (convex), but not for general \mathcal{U} .

Unfortunately, most objective functions we will want to use are non-convex. Because we are only guaranteed a local minimum, we might threshold an acceptable $\mathcal{U}^* = \mathcal{U}(\zeta^*)$, and if the \mathcal{U}^* returned by GD is above this threshold, resample using a different initialization ζ_0 , or introduce randomization in more sophisticated ways (e.g. Hamiltonian Monte Carlo).



5.4 Calculus Review

5.4.1 Motivation

In this section, we review the basics of single-variable, multivariate, and vector calculus. Before doing so, however, it is useful to think about how calculus can be used in trajectory optimization. Consider a trajectory $\tilde{\zeta} : [0, T] \rightarrow C$ for a configuration space with d dimensions. This trajectory is a function that maps from some time $t \in [0, T]$ to a vector q of coordinates in C-space:

$$\tilde{\zeta}(t) : t \mapsto q, q = \begin{bmatrix} \tilde{\zeta}_1(t) \\ \tilde{\zeta}_2(t) \\ \vdots \\ \tilde{\zeta}_d(t) \end{bmatrix}$$

Because of this, we say that $\tilde{\zeta}$ is a *vector-valued function* – a function operates on one or more variables and maps them to vectors. In order to analyze these sorts of functions, we need to understand *vector calculus*.

Furthermore, if we discretize time with N timesteps, then $\tilde{\zeta}$ becomes a vector of configuration vectors:

$$\tilde{\zeta} = \begin{bmatrix} q_1 \\ q_2 \\ \vdots \\ q_N \end{bmatrix}$$

Then, our cost function for trajectory optimization, $\mathcal{U} : \Xi \rightarrow \mathbb{R}^+$ maps from some discrete trajectory, which is simply a vector, to a single real number. This mapping from many variables to a single value makes \mathcal{U} a *multivariate function*, and we need *multivariable calculus* in order to analyze this function properly.

Finally, if we do not discretize trajectories, \mathcal{U} is a functional, and we need *calculus of variations*.

5.4.2 Derivative

Definition

Let $f(x) : \mathbb{R} \rightarrow \mathbb{R}$ be a function of a single real variable. The *derivative* of $f(x)$ is a function of that same variable that, when evaluated at a given input point, produces the rate of change of f at that point. The derivative of $f(x)$, which is written as $f'(x)$, can be defined as

$$f'(x) = \lim_{\epsilon \rightarrow 0} \frac{f(x + \epsilon) - f(x)}{\epsilon}$$

Example

Let $f(x) = 2x$. We can compute the derivative as follows.

$$f'(x) = \lim_{\epsilon \rightarrow 0} \frac{(2x + 2\epsilon) - 2x}{\epsilon} = 2$$

In this case, the derivative is a constant. This is not always the case.

Let $f(x) = x^2$. Omitting all steps for brevity, we can show that

$$f'(x) = 2x$$

In this case, the derivative is a function of x .

5.4.3 Partial Derivative

Definition

A *partial derivative* is an extension of the concept of a derivative to a multivariate function. Let $f(x_1, x_2, \dots, x_n) : \mathbb{R}^n \rightarrow \mathbb{R}$ be a function of n real variables. Partial derivatives of f are always defined with respect to only one of f 's variables. Specifically, the partial derivative of f with respect to x_i is a multivariate function of f 's variables that, when evaluated at a given input point, produces the rate of change of f in the direction of x_i .

In order to calculate the partial derivative of f with respect to x_i , we assume that all variables but x_i are held constant. This reduces f to a function of only x_i , and we can calculate the partial derivative, $\frac{\partial f}{\partial x_i}$, as the simple derivative of this new function using a limit:

$$\frac{\partial f}{\partial x_i}(x_1, x_2, \dots, x_n) = \lim_{\epsilon \rightarrow 0} \frac{f(x_1, x_2, \dots, x_{i-1}, x_i + \epsilon, x_{i+1}, \dots, x_n) - f(x_1, x_2, \dots, x_n)}{\epsilon}$$

Example

Let $f(x, y) = 2x + y + xy$. We can calculate the derivative of f with respect to x as follows:

$$\frac{\partial f}{\partial x} = \lim_{\epsilon \rightarrow 0} \frac{(2(x + \epsilon) + y + (x + \epsilon)y) - (2x + y + xy)}{\epsilon} = \lim_{\epsilon \rightarrow 0} \frac{2\epsilon + y\epsilon}{\epsilon} = 2 + y$$

Similarly,

$$\frac{\partial f}{\partial y} = 1 + x$$

5.4.4 Gradient

Definition

A gradient is a generalization of the concept of partial derivatives. Specifically, let $f(x_1, x_2, \dots, x_n)$ be a function of n real variables. Then, the gradient is simply a vector whose elements are the n partial derivatives of f with respect to each of its variables:

$$\nabla_{x_1, x_2, \dots, x_n} f = \begin{bmatrix} \frac{\partial f}{\partial x_1} \\ \frac{\partial f}{\partial x_2} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{bmatrix}$$

The gradient is a function that maps from f 's variables to a vector. Thus, we can describe the gradient as a multivariate vector-valued function. When the gradient is evaluated at a given input point, it produces a vector in \mathbb{R}^n which points in the direction of the greatest rate of ascent in f and whose magnitude specifies that rate of ascent.

Example

Let $f(x, y) = 2x + y + xy$. We computed the partial derivatives of this function above, so we can simply say that

$$\nabla_{x,y} f = \begin{bmatrix} 2 + y \\ 1 + x \end{bmatrix}$$

5.4.5 Ordinary Derivatives of Vector-Valued Functions

Definition

A vector-valued function is a function whose range is a set of vectors. For example, let

$$f(x) = \begin{bmatrix} f_1(x) \\ f_2(x) \\ \vdots \\ f_n(x) \end{bmatrix}$$

be a single-variable vector-valued function. This function maps a single variable into a n -dimensional vector. Because f is a function of a single variable, we can take its ordinary derivative by simply differen-

tiating each of the functions f_1, f_2, \dots, f_n that compose f 's matrix:

$$f'(x) = \begin{bmatrix} \frac{df_1}{dx} \\ \frac{df_2}{dx} \\ \vdots \\ \frac{df_n}{dx} \end{bmatrix}$$

This ordinary derivative is a vector of functions of a single variable, and when it is evaluated at a given input point, it produces a vector.

Example

Let

$$f(x) = \begin{bmatrix} x^2 + 2 \\ 2x \end{bmatrix}$$

Then, we can compute

$$f'(x) = \begin{bmatrix} 2x \\ 2 \end{bmatrix}$$

This is a vector of functions of x , and we can evaluate it at a particular point to get a vector.

5.4.6 Jacobians of Multivariate Vector-Valued Functions

Definition

The Jacobian matrix is a generalization of the gradient to multivariate vector-valued functions. Let $f(x_1, x_2, \dots, x_n) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a multivariate vector-valued function:

$$f(x_1, x_2, \dots, x_n) = \begin{bmatrix} f_1(x_1, x_2, \dots, x_n) \\ f_2(x_1, x_2, \dots, x_n) \\ \vdots \\ f_m(x_1, x_2, \dots, x_n) \end{bmatrix}$$

Then, the Jacobian of f is defined as

$$\frac{df}{d(x_1, x_2, \dots, x_n)} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \dots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \dots & \frac{\partial f_m}{\partial x_n} \end{bmatrix}$$

This matrix is often represented by **J**.

Usage in Kinematics

Let $\phi_{EE}(q)$ be a forward kinematics function that maps a configuration q from d -dimensional C-space to a location of a robot's end effector in \mathbb{R}^3 :

$$\phi_{EE}(q) = \begin{bmatrix} x(q) \\ y(q) \\ z(q) \end{bmatrix}$$

The Jacobian of ϕ gives us a 3 by d matrix whose rows are the gradients of x , y , and z with respect to the configuration q . This is very useful, as it gives us a way to map a rate of change in configuration space to a rate of change in the real world.

The Jacobian is also useful in inverse kinematics with the psuedo-inverse Jacobian and Jacobian transpose methods.

Properties

If we have a vector

$$v = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$$

then if, for some matrix \mathbf{A} , we can write $f(x_1, x_2, \dots, x_n) = \mathbf{A}v$ then $\mathbf{J} = \mathbf{A}$.

Example

Let

$$f(x, y) = \begin{bmatrix} 2x + y \\ y \end{bmatrix}$$

Then, we can calculate

$$\frac{df}{d(x, y)} = \begin{bmatrix} 2 & 1 \\ 0 & 1 \end{bmatrix}$$

We could compute this directly or by noting that

$$f(x, y) = \begin{bmatrix} 2 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

and using the first property of the Jacobian listed above.

Lecture 6: Trajectory Optimization II

*Scribes: Huazhe Xu, Xinyu Liu***Today's Outline**

- Problem Statement
- Functional Gradient Descent
- CHOMP: Covariant Hamiltonian Optimization for Motion Planning
- Next Lecture: Non-Euclidean Inner Product

6.1 Problem Statement

- Trajectory (function) $\zeta : [0, T] \rightarrow \mathbb{C}$
- Cost (functional) $U : \Xi \rightarrow \mathbb{R}^+$
- Optimization Problem:

$$\zeta^* = \arg \min_{\zeta \in \Xi} U(\zeta)$$

subject to


$$\zeta(0) = q_s,$$

$$\zeta(T) = q_g$$

- Update Equation (functional gradient descent):

$$\zeta_{i+1} \leftarrow \zeta_i - \frac{1}{\alpha} \nabla_{\zeta} U(\zeta_i)$$

 $\nabla_{\zeta} U(\zeta_i)$: also a function of time



- Ξ is a Hilbert space, a complete vector space with an inner product.
- Inner product:
For this lecture, we assume a particular Hilbert space specified by the Euclidean inner product. Given two trajectories $\xi_1, \xi_2 \in \Xi$, the Euclidean inner product is

$$\langle \xi_1, \xi_2 \rangle = \int_0^T \xi_1(t)^T \xi_2(t) dt$$

In discrete time, $\xi = [q_1, \dots, q_N]^T$, thus $\langle \xi_1, \xi_2 \rangle = \xi_1^T \xi_2$.

Properties of inner products:

Symmetry:

$$\langle \xi_1, \xi_2 \rangle = \langle \xi_2, \xi_1 \rangle$$

Positive definite:

$$\forall \xi, \langle \xi, \xi \rangle \geq 0; \langle \xi, \xi \rangle = 0 \iff \xi = 0$$

($\xi = 0$ is the zero trajectory that always maps time to the zero configuration.)

Linearity in the first argument:

$$\langle \xi_1 + \xi_2, \xi_3 \rangle = \langle \xi_1, \xi_3 \rangle + \langle \xi_2, \xi_3 \rangle$$

(The same holds true for the second argument by *symmetry*.)

6.2 Functional Gradient Descent

We use calculus of variation in computing the derivatives of a functional.

- Euler-Lagrange Equation:

If

$$\langle \xi_1, \xi_2 \rangle = \int_0^T \xi_1(t)^T \xi_2(t) dt$$

(i.e. Ξ is a Hilbert space with the Euclidean inner product)
and

$$U[\tilde{\zeta}] = \int_0^T F(t, \tilde{\zeta}(t), \tilde{\zeta}'(t)) dt$$

then

$$\nabla_{\tilde{\zeta}} U(t) = \frac{\partial F}{\partial \tilde{\zeta}(t)}(t) - \frac{d}{dt} \frac{\partial F}{\partial \tilde{\zeta}'(t)}(t)$$

Note that $\nabla_{\tilde{\zeta}} U(t) \in \Xi$


- Example:

Consider the example where you minimize the squared norm of velocity in trajectory subject to starting at q_s and ending at q_g :


$$U[\tilde{\zeta}] = \frac{1}{2} \int_0^T \|\tilde{\zeta}'(t)\|^2 dt$$

The optimal trajectory has

Shape: straight line. Intuitively, in the same amount time, a trajectory traversing a longer path needs a faster velocity, thus has higher cost.



Timing: constant velocity. Intuitively, in discrete time with T time steps, $U[\xi^*] < U[\xi]$



$$F(t, \zeta(t), \zeta'(t)) = \frac{1}{2} \|\zeta'(t)\|^2$$

Apply Euler-Lagrange equation,

$$\nabla_{\zeta} U(t) = 0 - \frac{d}{dt} \zeta'(t) = -\zeta''(t)$$

Since U is quadratic/convex, to find ζ that minimizes $\text{cost } U$, we set the gradient to 0, and solve for ζ^* , which is global minimum.

$$\zeta''(t) = 0$$

$$\zeta'(t) = a$$

shows that optimal trajectory has constant velocity.

$$\zeta(t) = at + b$$

shows that optimal trajectory is a straight line.

Then to solve for a and b in ζ^* , we use constraints $\zeta(0) = q_s$ and $\zeta(T) = q_g$.

- Proof:

First order Taylor series expansion (relates f to f')

$f : \mathbb{R} \rightarrow \mathbb{R}$,


$$f(x + \epsilon) \approx f(x) + \epsilon f'(x)$$

$$f'(x) = \lim_{\epsilon \rightarrow 0} \frac{f(x + \epsilon) - f(x)}{\epsilon}$$

$U : \Xi \rightarrow \mathbb{R}^+$,

$$U[\zeta + \epsilon\eta] \approx U[\zeta] + \epsilon \langle \nabla_{\zeta} U, \eta \rangle$$

smooth disturbance $\eta \in \Xi$ s.t. $\eta(0) = \eta(T) = 0$



arbitrary small $\epsilon \in \mathbb{R}$

$$\langle \nabla_{\zeta} U, \eta \rangle = \lim_{\epsilon \rightarrow 0} \frac{U[\zeta + \epsilon\eta] - U[\zeta]}{\epsilon} \quad (1)$$

$$\langle \nabla_{\zeta} U, \eta \rangle = \int_0^T \nabla_{\zeta} U(t)^T \eta(t) dt \quad (2)$$

We are going to massage equation (1) to equation (2) and term match to find $\nabla_{\zeta} U$.

Let $\phi(\epsilon) = U[\zeta + \epsilon\eta]$,

$$\langle \nabla_{\zeta} U, \eta \rangle = \lim_{\epsilon \rightarrow 0} \frac{\phi(\epsilon) - \phi(0)}{\epsilon}$$

$$= \left. \frac{d\phi}{d\epsilon} \right|_{\epsilon=0}$$

$$= \left. \frac{d}{d\epsilon} \int_0^T F[t, \zeta(t) + \epsilon\eta(t), \zeta'(t) + \epsilon\eta'(t)] dt \right|_{\epsilon=0}$$

Exchange differentiation with integration,

$$= \int_0^T \left. \frac{d}{d\epsilon} F[t, \zeta(t) + \epsilon\eta(t), \zeta'(t) + \epsilon\eta'(t)] dt \right|_{\epsilon=0}$$

Change of variables, denote $x(\epsilon) = \zeta(t) + \epsilon\eta(t)$ and $y(\epsilon) = \zeta'(t) + \epsilon\eta'(t)$, then apply chain rule,

$$= \int_0^T \left(\frac{\partial F[t, x(\epsilon), y(\epsilon)]}{\partial x} \right)^T \frac{dx}{d\epsilon} + \left(\frac{\partial F[t, x(\epsilon), y(\epsilon)]}{\partial y} \right)^T \frac{dy}{d\epsilon} dt \Big|_{\epsilon=0}$$

$$= \int_0^T \left(\frac{\partial F[t, x(\epsilon), y(\epsilon)]}{\partial x} \right)^T \eta(t) + \left(\frac{\partial F[t, x(\epsilon), y(\epsilon)]}{\partial y} \right)^T \eta'(t) dt \Big|_{\epsilon=0}$$

$$= \int_0^T \frac{\partial F}{\partial \zeta(t)} [t, \zeta(t) + \epsilon\eta(t), \zeta'(t) + \epsilon\eta'(t)]^T \eta(t) + \frac{\partial F}{\partial \zeta'(t)} [t, \zeta(t) + \epsilon\eta(t), \zeta'(t) + \epsilon\eta'(t)]^T \eta'(t) dt \Big|_{\epsilon=0}$$

Evaluate at $\epsilon = 0$,

$$= \int_0^T \left(\frac{\partial F[t, \zeta(t), \zeta'(t)]}{\partial \zeta(t)} \right)^T \eta(t) + \left(\frac{\partial F[t, \zeta(t), \zeta'(t)]}{\partial \zeta'(t)} \right)^T \eta'(t) dt$$

Write in compact form,

$$= \int_0^T \left(\frac{\partial F}{\partial \zeta(t)} \right)^T \eta(t) + \left(\frac{\partial F}{\partial \zeta'(t)} \right)^T \eta'(t) dt$$

Apply integration by parts to solve $\int_0^T \left(\frac{\partial F}{\partial \zeta'(t)} \right)^T \eta'(t) dt$,

$$\int_0^T \left(\frac{\partial F}{\partial \zeta'(t)} \right)^T \eta'(t) dt$$

$$= \left(\frac{\partial F}{\partial \zeta'(t)} \right)^T \eta(t) \Big|_0^T - \int_0^T \frac{d}{dt} \left(\frac{\partial F}{\partial \zeta'(t)} \right)^T \eta(t) dt$$

By definition, $\eta(0) = \eta(T) = 0$,

$$= - \int_0^T \frac{d}{dt} \left(\frac{\partial F}{\partial \zeta'(t)} \right)^T \eta(t) dt$$

Thus,

$$\langle \nabla_{\zeta} U, \eta \rangle = \int_0^T \left(\frac{\partial F}{\partial \zeta(t)} - \frac{d}{dt} \frac{\partial F}{\partial \zeta'(t)} \right)^T \eta(t) dt = \int_0^T \nabla_{\zeta} U(t)^T \eta(t) dt \text{ for every } \eta$$

Therefore,

$$\nabla_{\zeta} U(t) = \frac{\partial F}{\partial \zeta(t)}(t) - \frac{d}{dt} \frac{\partial F}{\partial \zeta'(t)}(t)$$

□

6.3 CHOMP: Covariant Hamiltonian Optimization for Motion Planning

CHOMP instantiates functional gradient descent for cost

$$U[\xi] = U_{smooth}[\xi] + \lambda U_{obs}[\xi]$$

Smoothness cost is defined as in our example


$$U_{smooth}[\xi] = \frac{1}{2} \int_0^T \|\xi'(t)\|^2 dt$$

Obstacle cost is defined as

$$U_{obs}[\xi] = \int_t \int_u c(\phi_u(\xi(t))) \cdot \left\| \frac{d}{dt} \phi_u(\xi(t)) \right\| du dt$$

Understanding $U_{obs}[\xi]$:

- Define a cost function in W , $c : W \rightarrow \mathbb{R}$ that uses a signed distance field to compute distance to the closes obstacle, and returns a higher cost the closer the point is.
- Then for each time point along the trajectory (thus the integral over time), look at the configuration $\xi(t)$.
- For each body point on the robot u (thus the integral over body points), apply for forward kinematics mapping ϕ_u to get the xyz locations of the points when the robot is in configuration $\xi(t)$.
- For each body point location, compute the cost c .



The second term in the integral (the norm of the velocity) is there to create a path integral formulation.