Optimal Average Case Strategy for Looking around a Corner

Reza Dorrigiv * Alejandro López-Ortiz † Selim Tawfik †

Abstract

A robot is free to move in a non-convex polygonal region, starting against an edge on the boundary. Ahead of the robot at one unit of distance is a corner, i.e. a reflex vertex (see Figure 1). The angle θ is unknown to the robot. The robot's task is to look at the region around the corner, which it initially cannot see.

We let $\varphi = \pi - \theta$. If $\varphi \geq \pi/2$, the robot is best off moving straight to the vertex. However, if φ is close to 0, much shorter paths exist, making this solution suboptimal. Therefore we ask: What is the best path for the robot to follow?

In this paper, we look into the problem of finding an optimal average-case strategy under a homogeneous probability distribution for φ. The average-case performance of a strategy is measured by its average cost, defined as the expected value of the strategy's competitive function. Given a value for φ, the competitive function of a strategy gives the ratio of the distance the robot travels to look around the corner (as prescribed by the strategy) to the shortest distance it must travel to do so.

We give strong evidence that an optimal average-case strategy exists and achieves an average cost of ~ 1.189.

1 Introduction

The corner exploration problem was first examined by Icking, Klein and Ma in [7]. The authors measure the performance of a strategy by its competitive factor, defined as the maximum value attained by its competitive function (lower is better). Under this measure, they show there exists an optimal strategy characterized as the solution of a certain differential equation. This strategy's competitive function is in fact identically equal to a constant $c \approx 1.21218$. Thus both its competitive factor and its average cost are equal to c. Although this strategy's competitive factor is optimal, strategies with better average costs exist.

The problem of finding an optimal average-case strategy reduces to that of finding a curve which minimizes an integral giving the average cost. The appropriate tool for handling such problems is the calculus of variations, on which we state a couple of useful theorems in the appendix. This said, formal proofs of some of our results would require complex techniques from that field. For the problem at hand, we obtain near optimal results using discretizations of the instance.

After giving discretizations of the problem, we apply some techniques from the calculus of variations to produce a strong candidate for an optimal average-case strategy. The average cost of this strategy is ~ 1.189. This is better than the average cost of the strategy described in [7], which is ~ 1.21218 as we noted earlier. However, we should state that the competitive factor of our strategy is worse: ~ 1.3136 for ours vs. ~ 1.21218 for [7].

Unfortunately, we do not have a closed-form for our strategy, so we propose a closed-form approximation whose average cost exceeds ours by less than 0.202%.

2 Competitive Strategies

2.1 Preliminaries

We need some of the early definitions and results in [7], which we reproduce here.

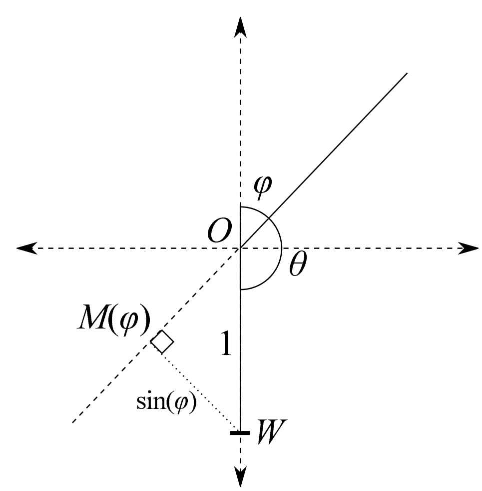

Figure 1: The robot's predicament.

We start by introducing a coordinate system with the origin located at the corner, and with the robot's starting position, W, one unit away from the origin. We let φ be the angle between the invisible wall and the pro-

*Faculty of Computer Science, Dalhousie University, Halifax, NS, B3H 4R2, Canada, email: rdorrigiv@cs.dal.ca

†David Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, N2L 3G1, Canada, email: {alopez-o,stawfik}@uwaterloo.ca

longation of the visible wall, $M(\pi)$. The distance from W to $M(\varphi)$ is denoted by $a(\varphi)$. We observe that

$$a(\varphi) = \begin{cases} \sin(\varphi) & : \quad \varphi < \frac{\pi}{2} \\ 1 & : \quad \varphi \geq \frac{\pi}{2} \end{cases} \tag{1}$$

A *strategy* is a curve which starts at the point W and finishes on the prolongation of the visible wall, $M(\pi)$. For any possible value of φ, there is a point on such a curve from which the other wall is visible, the intersection of the curve with $M(\varphi)$.

We let $A_S(\varphi)$ be the length of the path described by strategy S between W and the first point of intersection with $M(\varphi)$. The *competitive function*, $f_S(\varphi)$, of S is the ratio of $A_S(\varphi)$ to $a(\varphi)$ and its *competitive factor*, c_S, is the supremum of the values taken by $f_S(\varphi)$ in $(0, \pi]$.

$$f_S(\varphi) = \frac{A_S(\varphi)}{a(\varphi)}, \qquad c_S = \sup_{\varphi \in (0,\pi]} f_S(\varphi)$$

We say that S is *competitive* if $c_S < \infty$. If a strategy reaches $M(\varphi')$ for the first time, turns back and meets $M(\varphi')$ again, this part of the path may be cut off and replaced by a radial line segment, giving a better strategy. Strategies with radial line segments can then be approximated arbitrarily closely by strategies that can be described in polar coordinates.

Definition 1 ([7]) A curve $S = (\varphi, s(\varphi))$, where s is defined on $[0, \pi]$, is called a *strategy* for the corner problem if the following holds.

(i) s is a continuous function on an interval $[0, \sigma]$, where $\sigma \leq \pi$.

(ii) On the open interval $(0, \sigma)$, s is piecewise continuously differentiable and $s'(0)$ exists (possibly $\pm\infty$).

(iii) s is *rectifiable*, i.e. s has finite arc length.[1]

(iv) $s(0) = 1$.

(v) If $s(\sigma) \neq 0$ then $\sigma = \pi$.

The last property says that the strategy must end somewhere on $M(\pi)$, possibly the corner. By agreeing that $s(\varphi) = 0$ for $\sigma < \varphi \leq \pi$, we can regard $s(\varphi)$ as defined on all of $[0, \pi]$.

Lemma 1 ([7]) *Let $S = (\varphi, s(\varphi))$ be a strategy. Then S is competitive iff $|s'(0)| < \infty$. The estimation*

$$c_S \geq \sqrt{s'^2(0) + 1}$$

holds for the competitive factor.

[1]This criterion is not listed in [7], but it is assumed. Since $s'(\sigma)$ is not known to exist, it is necessary to require this.

Thus competitive strategies are piecewise continuously differentiable in $[0, \pi)$. Using the formula for arc length in polar coordinates, we have

$$A_S(\varphi) = \int_0^{\varphi} \sqrt{s^2(t) + s'^2(t)}\, dt \tag{2}$$

By the fundamental theorem of calculus, $A'_S(\varphi) = \sqrt{s^2(\varphi) + s'^2(\varphi)}$ on $[0, \pi)$, and so A_S is continuous therein. Since s is rectifiable, $A_S(\pi) < \infty$ so that A_S is bounded on $[0, \pi]$. By L'Hôpital's rule,

$$\begin{aligned} \lim_{\varphi \to 0} f_S(\varphi) &= \lim_{\varphi \to 0} \frac{\sqrt{s^2(\varphi) + s'^2(\varphi)}}{\cos(\varphi)} \\ &= \sqrt{s'^2(0) + 1} \end{aligned}$$

Defining $f_S(0) := \sqrt{s'^2(0) + 1}$, f_S is continuous on $[0, \pi)$ and bounded on $[0, \pi]$. In particular, f_S is integrable on $[0, \pi]$ (this is given by Lebesgue's criterion for Riemann-integrability, see [1] pp. 171).

2.2 The Objective

We are interested in finding a competitive strategy S which minimizes the average value taken by $f_S(\varphi)$ in the interval $[0, \pi]$. More precisely, we wish to minimize the expected value of the ratio of the distance traveled before the corner is seen over the shortest path to the line of sight:

$$\begin{aligned} E[f_S] &= \frac{1}{\pi} \int_0^{\pi} f_S(\varphi)\, d\varphi \\ &= \frac{1}{\pi} \int_0^{\pi} \int_0^{\varphi} \frac{\sqrt{s^2(t) + s'^2(t)}}{a(\varphi)}\, dt\, d\varphi \end{aligned} \tag{3}$$

which we define as the *average cost* of f_S. Since f_S is integrable on $[0, \pi]$, the above integral exists and is finite.

Observation 1 *If $S = (\varphi, s(\varphi))$ is an optimal strategy, then s is non-increasing on $[0, \pi]$. Indeed, the robot is always better off staying at the same radial distance from the corner than getting farther from it.*

3 A Discretization

As a first attempt to gain insight into the problem, we look at a discretization of it. To this end, we start by partitioning the interval $[0, \pi]$ into n equal subintervals (for simplicity, we choose n to be even), so that we get partition points $x_0, \ldots, x_n$ with $x_0 = 0$, $x_n = \pi$ and $x_{i+1} - x_i = \frac{\pi}{n}$ for $i = 0, \ldots, n-1$. Putting $\theta = \frac{\pi}{n}$, we will assume that the angle φ in the original problem can only take values in $\{k\,\theta \ : \ 1 \leq k \leq n\}$ with equal probability $\frac{1}{n}$.

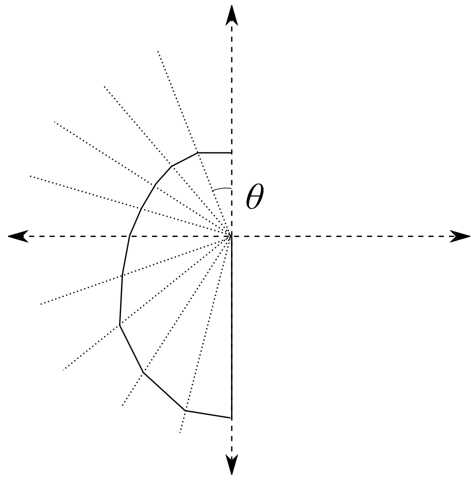

Figure 2: Robot's path in the discretization.

For non-negative values $y_0, \ldots, y_n$, let $P(y_0, \ldots, y_n)$ denote the polygonal path which joins the polar points $(x_0, y_0), \ldots, (x_n, y_n)$ in that order. The problem is then to find values for $y_0, \ldots, y_n$ which minimize the average cost incurred by the strategy $P(y_0, \ldots, y_n)$. The requirement that the robot starts one unit away from the origin is worked in by demanding that $y_0 = 1$. Let $L_i(y_0, \ldots, y_n)$ denote the i^{th} segment of the path $P(y_0, \ldots, y_n)$. The law of cosines gives us

$$\begin{aligned} |L_i(y_0, \ldots, y_n)|^2 &= y_{i-1}^2 + y_i^2 - 2\cos(\theta) y_{i-1} y_i \\ &= (y_{i-1} - \cos(\theta) y_i)^2 + \sin^2(\theta) y_i^2 \end{aligned}$$

And so $|L_i(y_0, \ldots, y_n)|$ may be expressed as the norm of a vector

$$\|(y_{i-1} - \cos(\theta) y_i, \sin(\theta) y_i)\| \quad (4)$$

Notice that the expression above is *convex* in the arguments $y_0, \ldots, y_n$. For $1 \le k \le n$, the robot travels a distance of $\sum_{i=1}^{k} |L_i(y_0, \ldots, y_n)|$ before it reaches the ray $k\theta$. Remembering (1), we see that the competitive function, f_S takes on values according to

$$f_S(k\theta) = \sum_{i=1}^{k} \frac{|L_i(y_0, \ldots, y_n)|}{\sin(k\theta)}$$

if $1 \le k \le \frac{n}{2}$ and

$$f_S(k\theta) = \sum_{i=1}^{k} |L_i(y_0, \ldots, y_n)|$$

if $\frac{n}{2} + 1 \le k \le n$. The average cost of the strategy $P(y_0, \ldots, y_n)$ is thus given by

$$\frac{1}{n} \sum_{k=1}^{\frac{n}{2}} \sum_{i=1}^{k} \frac{|L_i(y_0, \ldots, y_n)|}{\sin(k\theta)} + \frac{1}{n} \sum_{k=\frac{n}{2}+1}^{n} \sum_{i=1}^{k} |L_i(y_0, \ldots, y_n)|$$

Now, the convexity of (4) gives us the convexity of the above expression and so the problem at hand is actually a convex program. Using a numerical solver [2], we are able to find solutions for $n = 100$, plotted below:

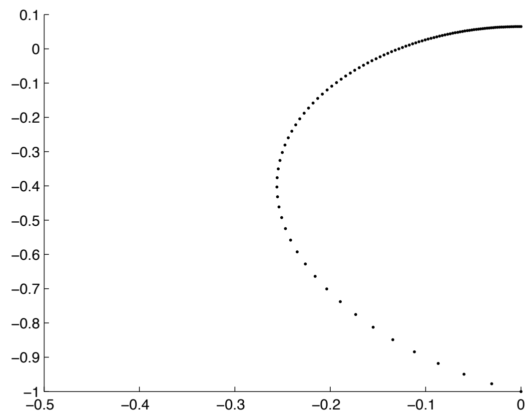

Figure 3: Optimal discrete strategy for $n = 100$.

This plot suggests that an optimal continuous strategy exists. The solutions for $n = 1000$ and $n = 5000$ further support this hypothesis. Note that the solution does not reach the corner as is the case with the strategy found in [7]. Instead, it continually approaches the corner until it reaches the ray π, at a distance of ~ 0.067 from the corner.

Observation 2 *This discretization method suggests that if $H = (\varphi, h(\varphi))$ is an optimal continuous strategy, then h is differentiable at π and $h'(\pi) = 0$. To see this, consider y_n in relation to y_{n-1}: once the robot has reached the ray $(n-1)\,\theta$, it is best off heading to the ray $n\,\theta = \pi$ in the shortest possible path, which is the line segment perpendicular to $n\,\theta$. Thus $y_n = \cos(\theta)\, y_{n-1}$ and we expect the difference quotient*

$$\frac{y_n - y_{n-1}}{\theta} = \frac{(\cos(\theta) - 1)}{\theta}\ y_{n-1}$$

to converge to $h'(\pi)$ as θ approaches 0, if it converges at all. Since

$$\lim_{\theta \to 0} \frac{(\cos(\theta) - 1)}{\theta} = 0$$

[2] We used `CVX`, a package for specifying and solving convex programs [6].

and $0 \leq y_{n-1} \leq 1$*, we would have* $h'(\pi) = 0$*. It is actually possible to prove this rigorously with the techniques of the calculus of variations. An argument along similar lines, by considering* y_n *and* y_{n-1} *in relation to* y_{n-2}*, gives* $h''(\pi) = \infty$*.*

4 A Continuous Solution

Our discretization has given us good evidence that an optimal continuous strategy exists. We now explore this possibility with more appropriate tools coming from the calculus of variations. First however We start by considering yet another discretization which will reinforce our evidence. We also derive an expression for the average cost E in the form of a single integral.

Next, we apply the techniques of the calculus of variations to the problem of minimizing E. Namely, we solve the Euler-Lagrange equation associated with E. In general, proving that a given solution to the Euler-Lagrange equation is a minimizer for a variational problem is difficult and attempting to do so would lead us too far deep into the theory of the calculus of variations. Instead, we will be content with the near perfect match we observe between our candidate solution and the discrete optimal average-case solutions.

4.1 The Average Cost as a Single Integral

In view of (3), the expression for the average cost is

$$\begin{aligned} E[f_S] &= \frac{1}{\pi} \int_0^\pi f_S(\varphi)\, d\varphi \\ &= \frac{1}{\pi} \int_0^\pi \int_0^\varphi \frac{\sqrt{s(t)^2 + s'^2(t)}}{a(\varphi)}\, dt\, d\varphi \qquad (5) \end{aligned}$$

Before continuing, we define a function b:

$$b(\varphi) = \begin{cases} \ln\left(\frac{1-\cos(\varphi)}{\sin(\varphi)}\right) & : \quad \varphi < \frac{\pi}{2} \\ \varphi - \frac{\pi}{2} & : \quad \varphi \geq \frac{\pi}{2} \end{cases} \qquad (6)$$

Note that $b(\varphi)$ is continuous and that $b'(\varphi) = \frac{1}{a(\varphi)}$ on $(0, \pi]$, i.e. b is an antiderivative for the reciprocal of a. Using Fubini's theorem (see [3] pp. 64) we may change the order of integration in (5):

$$\frac{1}{\pi} \int_0^\pi \int_t^\pi \frac{\sqrt{s(t)^2 + s'^2(t)}}{a(\varphi)}\, d\varphi\, dt$$

Finally, invoking the fundamental theorem of calculus, this gives

$$E[f_S] = \int_0^\pi \frac{\sqrt{s(t)^2 + s'^2(t)}\left(\frac{\pi}{2} - b(t)\right)}{\pi}\, dt \qquad (7)$$

4.1.1 Another Discretization

The above derivation suggests we try to approximate the average cost E by a Riemann sum. Suppose $S = (\varphi, s(\varphi))$ is an optimal strategy. Choosing an $n > 0$ sufficiently large, we let $\theta = \frac{\pi}{n}$ and $\varphi_i = i\ \theta$. We approximate $s'(\varphi_i)$ by the Newton quotient

$$\frac{s(\varphi_{i+1}) - s(\varphi_i)}{\theta}$$

Putting $y_i = s(\varphi_i)$, we make an approximation for $E[f_S]$:

$$\frac{1}{\pi} \sum_{i=0}^{n-1} \sqrt{y_i^2 + \left(\frac{y_{i+1} - y_i}{\theta}\right)^2} \left(\frac{\pi}{2} - b(\varphi_i)\right) \theta \qquad (8)$$

Naturally, we look for values $y_0, \ldots, y_n$ which will minimize (8). Fortunately, this problem is once again convex and in fact computationally easier than the one before. Using a numerical solver, we are able to increase n to 100000, the result is shown below together with our previous discretization:

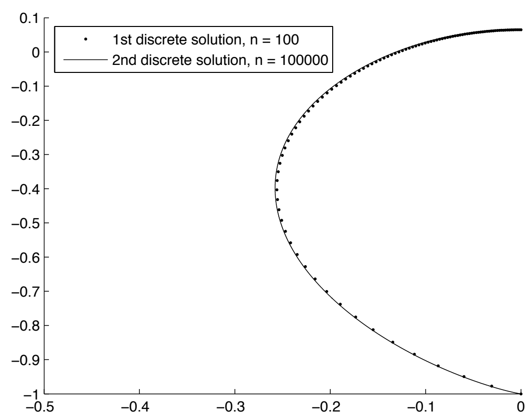

Figure 4: Optimal solutions for two types of discretizations

Notice the two solutions almost agree everywhere they are both defined. The value of (8) in this case is $\sim$ 1.1892.

4.2 An Optimal Strategy

If we regard E as a functional acting on functions defined in $[0, \pi]$ and piecewise continuously differentiable in $[0, \pi)$, our goal is to find such a function h, subject to $h(0) = 1$ and $h(x) \geq 0$ for $x \in [a, b]$, that minimizes E. In general, proving the existence of a continuous minimizer for variational problems such as this one is not an easy task. Instead, we will take the corroborating results from the first and second discretization as proof one exists, and that its average cost should be close to 1.1892. The discretizations suggest we look for such a minimizer with $h(\pi) > 0$. Before proceeding, we verify that we may apply the techniques of the calculus of variations:

Lemma 2 *Suppose h minimizes* (7) *and $h(\pi) > 0$. Then h is C^2 on $(0, \pi)$.*

Proof. As we have observed, h is non-increasing on $[0, \pi]$, thus $h(t) > 0$ for all $t \in [0, \pi]$. If we define

$$F(x, y, z) = \frac{1}{\pi}\sqrt{y^2 + z^2}\left(\frac{\pi}{2} - b(x)\right)$$

then

$$E(h) = \int_0^\pi F(t, h(t), h'(t))\, dt$$

Now we have

$$\frac{\partial^2 F}{\partial z^2}(t, h(t), h'(t)) = \frac{1}{2\pi}\frac{(\pi - 2b(t))^2}{(h(t)^2 + h'(t)^2)^{\frac{3}{2}}}$$

which remains non-zero in $(0, \pi)$. By the criterion for the regularity of minimizers, h is C^2 in $(0, \pi)$. □

As shown in the appendix, the above lemma ensures h satisfies the Euler-Lagrange equation associated with (7) in $(0, \pi)$. Although too long to produce here, this equation may be expressed in the explicit form

$$h'' = G(t, h, h')$$

so long as $h(t) > 0$ and $0 < t < \pi$. Thus if we know the values for, say, $h(\frac{\pi}{2})$ and $h'(\frac{\pi}{2})$, then we can determine h by the Existence-Uniqueness Theorem for ordinary differential equations. With this mind, we use the discretization results for $n = 100000$ to estimate values for $h(\frac{\pi}{2})$ and $h'(\frac{\pi}{2})$. The resulting solution to the Euler-Lagrange equation associated with (7) is plotted below alongside the optimal discrete solution for $n = 100$.

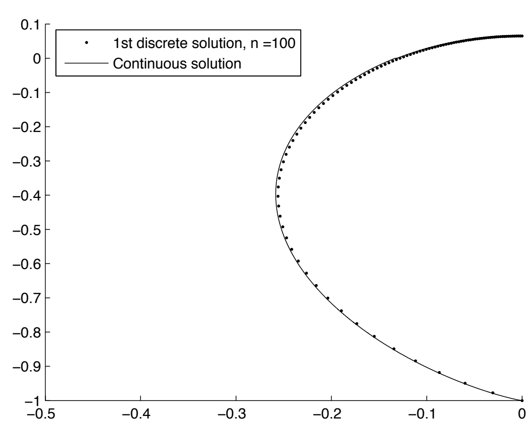

Figure 5: Optimal discrete solution together with optimal continuous solution

The very close match between the continuous solution and the discrete one gives us strong evidence that we have indeed found an optimal strategy. The average cost of this strategy is ~ 1.189. Its competitive factor is ~ 1.3136.

5 Conclusion

Our results give us strong evidence that we have found an optimal average case strategy. It achieves an average cost of ~ 1.189. Although we do not have a closed-form expression for this strategy, a good approximation is given by

$$\frac{1 + 7\,\theta}{35 + 21\,\theta + 22\,\theta^2}$$

which has an average cost of ~ 1.1914, which exceeds the presumed optimal average cost by less than 0.202%.

References

[1] T. M. Apostol. *Mathematical Analysis*. Addison Wesley, 1974.

[2] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[3] J. C. Burkill. *The Lesbesgue Integral*. Cambridge University Press, 1951.

[4] J. C. Butcher. *Numerical Methods for Ordinary Differential Equations*. Wiley, 2008.

[5] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. `http://stanford.edu/~boyd/graph_dcp.html`.

[6] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. `http://cvxr.com/cvx`, Apr. 2011.

[7] C. Icking, R. Klein, and L. Ma. How to look around a corner. In *Proceedings of the 5th Canadian Conference on Computational Geometry*, pages 443–448, Waterloo, Ontario, 1993.

[8] J. Jost and X. Li-Jost. *Calculus of Variations*. Cambridge University Press, 1998.

[9] M. Mesterton-Gibbons. *A Primer on the Calculus of Variations and Optimal Control Theory*. American Mathematical Society, 2009.

6 Appendix: The Calculus of Variations

Our analysis requires us to minimize expressions of the form

$$J[r] = \int_a^b F\left(x, r(x), r'(x)\right)\, dx \qquad (9)$$

where $F : (a,b)\times\mathbb{R}^2 \to [0,\infty)$ is a given function of class C^2 in $(a,b)\times\Omega$ where Ω is an open region in $\mathbb{R}^2$. J is the *objective function*, otherwise known as a *functional.* The goal is to find a minimizing function $r : [a,b] \to \mathbb{R}$ over the class of *admissible* functions. For us, these are the functions that are piecewise continuously differentiable in (a,b) with $(r(x), r'(x)) \in \Omega$ for all x where $r'(x)$ is defined, and moreover satisfy some boundary condition $r(a) = \alpha$ for some $\alpha \in \mathbb{R}$.

Suppose for now that r is such a minimizer. Assume furthermore that r is in fact C^2 in (a,b). It can be shown that

$$\frac{\partial F}{\partial r}(x, r(x), r'(x)) - \frac{d}{dx}\left(\frac{\partial F}{\partial r'}(x, r(x), r'(x))\right) = 0 \quad (10)$$

holds everywhere in (a,b) (see [8] for example). Equation (10) is known as the *Euler-Lagrange equation.* It is a second order differential equation in r and must be satisfied by any r which minimizes J subject to the conditions we have imposed. It is important to note that simply satisfying (10) is not enough to guarantee that r is a minimizer: the condition is necessary, but not sufficient.

In the above, we have assumed minimizers for (9) are of class C^2 in (a,b). The following result gives us conditions on F under which this is justified.

Proposition 3 *Regularity of Minimizers Let r be a piecewise continuously differentiable minimizer for the above problem, and let F be as in* (9)*. If $\frac{\partial^2 F}{\partial r'^2}(x, r(x), r'(x))$ does not vanish anywhere in (a,b), then r is C^2 in (a,b).*

A proof of this statement goes along the same lines as Theorem 1.2.3 and 1.2.3 in [8]. Note that because r is piecewise continuously differentiable, the above implies it is C^1 in $[a,b]$.