

A New Method For Numerical Constrained Optimization

Ronald N. Perry

Abstract

Numerical constrained optimization is an important tool with extensive applications in computer graphics and many other diverse fields. In computer animation, for example, constrained optimization has been used for finding initial conditions, smoothing or finishing motion paths, interpolating orientations, animating flexible bodies, self-assembly of objects, and finding boundaries.

An ideal problem for constrained optimization is one that has a single measure defining the quality of a solution plus some requirements upon that solution that must not be violated. The quality measure is called the *objective function*; the requirements are called *constraints*. A constrained optimization solver maximizes (or minimizes) the objective function while satisfying the constraints. In this report, we propose a new method for constraint handling that can be applied to established optimization algorithms and which significantly improves their ability to traverse through constrained space. To make the presentation concrete, we apply the new constraint method to the Nelder and Mead polytope algorithm. The resulting technique, called SPIDER, has shown great initial promise for solving difficult (e.g., nonlinear, nondifferentiable, noisy) constrained problems.

Presented at SIGGRAPH 2001 Conference Abstracts and Applications.

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories of Cambridge, Massachusetts; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories. All rights reserved.

A New Method For Numerical Constrained Optimization

Ronald N. Perry
Mitsubishi Electric Research Laboratory

Introduction

Numerical constrained optimization is an important tool with extensive applications in computer graphics [3] and many other diverse fields [1]. In computer animation, for example, constrained optimization has been used for finding initial conditions, smoothing or finishing motion paths, interpolating orientations, animating flexible bodies, self-assembly of objects, and finding boundaries [3].

An ideal problem for constrained optimization is one that has a single measure defining the quality of a solution plus some requirements upon that solution that must not be violated. The quality measure is called the **objective function** (denoted as F below); the requirements are called **constraints** (denoted as C_i below). A constrained optimization solver maximizes (or minimizes) the objective function while satisfying the constraints. In this sketch, we propose a new method for constraint handling that can be applied to established optimization algorithms and which significantly improves their ability to traverse through constrained space. To make the presentation concrete, we apply the new constraint method to the Nelder and Mead polytope algorithm [2]. The resulting technique, called SPIDER, has shown great initial promise for solving difficult (e.g., nonlinear, nondifferentiable, noisy) constrained problems.

The Idea

Many constrained problems have optima that lie near constraint boundaries. Consequently, avoidance of constraints can hinder an algorithm's path to the answer. By allowing (and even encouraging) an optimization algorithm to move its vertices into constrained space, a more efficient and robust algorithm emerges. In the new method, constraints are partitioned into multiple levels. A constrained performance, independent of the objective function, is defined for each level. A set of rules, based on these partitioned performances, specify the ordering and movement of vertices as they straddle constraint boundaries; these rules (employing the insight stated above) have been shown to significantly aid motion along constraints toward an optimum. Note that the new approach uses no penalty function and thus does not warp the performance surface, thereby avoiding the possible ill-conditioning of the objective function typical in penalty methods.

Problem Statement

Maximize $F(x)$ for all $x \in \mathbb{R}^N$ such that $C_i(x) \geq 0$ where $F: \mathbb{R}^N \rightarrow \mathbb{R}$, $C_i: \mathbb{R}^N \rightarrow \mathbb{R}$, and $i = [1..M]$. F or any C_i may be highly nonlinear, nondifferentiable, and/or noisy. A function (e.g., F) is considered **noisy** when repeated evaluation of that function with identical input yields different values.

SPIDER Algorithm

- Pick a starting location $x_s \in \mathbb{R}^N$
- Partition F and the constraints C_i into W levels $[L_1..L_w]$ with $L_w = \{ F \}$. We define the **performance** of a location $x \in \mathbb{R}^N$ as a 2-tuple $\langle P, L \rangle$ consisting of a floating point scalar P and an integer level indicator L . P represents the "goodness" of x at level L . P and L are computed as follows:
 - $P_1 \leftarrow \sum \min(C_k(x), 0)$, for all C_k in L_1
 - If (P_1 is nonzero) $P \leftarrow P_1$, $L \leftarrow 1$, Stop
 - $P_2 \leftarrow \sum \min(C_k(x), 0)$, for all C_k in L_2
 - If (P_2 is nonzero) $P \leftarrow P_2$, $L \leftarrow 2$, Stop
 - ...
 - ...
 - $P_w \leftarrow F(x)$, $P \leftarrow P_w$, $L \leftarrow W$, Stop

The performances of two locations $x_1 (\langle P_1, L_1 \rangle)$ and $x_2 (\langle P_2, L_2 \rangle)$ are compared as follows:

 - If ($L_1 == L_2$) if ($P_1 > P_2$) x_1 is better else x_2 is better, Stop
 - If ($L_1 > L_2$) x_1 is better else x_2 is better
- Determine the performance of the starting location x_s to form the starting vertex $v_s = \langle x_s, P_s, L_s \rangle$
- Pick the scale of each dimension to form the **size of space** vector $s \in \mathbb{R}^N$
- Generate an initial set of $N+1$ vertices $[v_1..v_{N+1}]$ using v_s and s . For example (where $N = 2$ and vector indices begin at 1),

$$v_1, v_2, v_3 \leftarrow v_s$$

$$v_2.x[1] \leftarrow v_2.x[1] + \text{random}() * s[1]$$

$$v_3.x[2] \leftarrow v_3.x[2] + \text{random}() * s[2]$$
- Determine the performance of each vertex $[v_1..v_{N+1}]$
- Sort the vertices $[v_1..v_{N+1}]$ from worst to best. Label the overall best vertex B_{all} and the overall worst vector W_{all} .
- For each vertex v in the set $[v_1..v_{N+1}]$ from worst to best:
 - Determine the centroid c of $[v_1..v_{N+1}]$, excluding v and other vertices at a lower level than v . If there are insufficient vertices to compute a valid centroid (e.g., < 2),

incrementally include other vertices at lower levels until there are a sufficient number.

- If v is not the best vertex:
 - $x_t \leftarrow c + (c - v.x) * \text{expansionFactor}$ (e.g., 1.1)
 - $\langle P_t, L_t \rangle \leftarrow \text{ComputePerf}(x_t)$
 - Form trial vertex v_t : $v_t.x \leftarrow x_t$, $v_t.P \leftarrow P_t$, $v_t.L \leftarrow L_t$
 - Either accept v_t or reject v_t (where accept **means** replace v with v_t) using the following criteria: if $v_t.L == v.L$, accept v_t if $v_t.P > v.P$; if $v_t.L > v.L$, accept v_t if $\text{Perf}(v_t) > \text{Perf}(B_{all})$; if $v_t.L < v.L$, accept v_t if $\text{Perf}(v_t) > \text{Perf}(W_{all})$
- If v is the best vertex:
 - $x_t \leftarrow c - (c - v.x) * \text{expansionFactor}$ (e.g., 1.1)
 - $\langle P_t, L_t \rangle \leftarrow \text{ComputePerf}(x_t)$
 - Form trial vertex v_t : $v_t.x \leftarrow x_t$, $v_t.P \leftarrow P_t$, $v_t.L \leftarrow L_t$
 - Accept v_t if $\text{Perf}(v_t) > \text{Perf}(v)$
- Relabel B_{all} and W_{all} if either has changed
- If $\text{Perf}(B_{all})$ has not improved, shrink the vertices at the same level as B_{all} toward B_{all} , and flip (as well as shrink) vertices at lower levels over B_{all} . The later rule helps to move legs (vertices) across a constraint boundary towards feasibility.
- If the number of successive shrinks exceeds some threshold, rebuild the vertex set using B_{all} as v_s and the current size in each dimension as s (see step 5)
- If there are more iterations to perform, repeat from step 7

Notes and Extensions


- One effective classification of constraints is to place simple limits on x that are independent of F into level 1, all other constraints into level 2, and the objective function F into level 3. Many different strategies for classification are being explored.
- SPIDER permits dynamic classification of constraints at anytime. This classification can be specified by a user (through observation of how SPIDER is moving) or by a classification algorithm which, for example, categorizes constraints into "active" and "inactive" levels.
- Other non-polytope optimization methods [2] can be modified to use the dynamic partitioned performances and corresponding rules introduced above, thereby improving their ability to traverse constrained space.

Results

Although we have tested SPIDER on some difficult problems with great initial success, much remains to prove the viability of this technique. In the talk we will demonstrate, through a series of contour plots, how SPIDER traverses constrained space and why the partitioned performances, along with the specified rules, enable an optimization algorithm to better navigate difficult terrain.

Summary

Partitioned performances permit many optimization algorithms, including SPIDER, to better traverse constrained space. The dynamic classification of constraints, either by a user or a classification algorithm, further enhances navigation through difficult terrain. Finally, the SPIDER algorithm, including the centroid computation, leg flipping when shrinking, and cycling through all legs before resorting, is a robust method for solving difficult constrained problems.



References

- R. Fletcher, Practical Methods of Optimization, Second Edition (John Wiley, Chichester and New York, 1987).
- Press, W., Flannery, B., Teukolsky, S., and Vetterling, W., Numerical Recipes, Cambridge University Press, New York, NY, 1986.
- Goldsmith, Jeff and Alan H. Barr, "Applying constrained optimization to computer graphics," SMPTE Journal, Vol. 102, No. 10, October 1993, pp. 910-912.