Bringing IDEAs into Practice:
Optimization in a Minimally Invasive
Vascular Intervention Simulation System

Peter A.N. Bosman

Tanja Alderliesten

Bringing IDEAs into Practice: Optimization in a Minimally Invasive Vascular Intervention Simulation System

Peter A.N. Bosman[1]   Tanja Alderliesten[2]

Peter.Bosman@cs.uu.nl   *tanja@isi.uu.nl*

[1] Institute of Information and Computing Sciences, Utrecht University
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
[2] Image Sciences Institute, University Medical Center Utrecht
P.O. Box 85.500, 3508 GA Utrecht, The Netherlands

Abstract

For real–valued continuous optimization problems various evolutionary algorithms (EAs) have obtained promising results on benchmark problems. The Iterated Density–Estimation Evolutionary Algorithm (IDEA) is an example of one such algorithm. However, little is known about the practical benefits of these algorithms even though AI–techniques are often favored in practice because of their general applicability and good performance on complicated real–world problems. In this paper we focus on one specific practical medical application that imposes many optimization tasks. The application is the simulation of minimally invasive vascular interventions. We compare the use of a hybrid IDEA with the conjugate gradients algorithm and a problem–specific optimization algorithm and indicate that although the application of the conjugate gradients algorithm already leads to highly useful results, IDEAs yet improve on these results in the area of scalability, making a clear statement that IDEAs can indeed also be useful in practice.

1   Introduction

Our goal in this paper is to compare the use of three different optimization techniques in a practical setting: a simulation system for minimally invasive vascular interventions. To this end we have performed experiments with an implementation of the conjugate gradients algorithm [9], an evolutionary algorithm (EA) [4, 5] and a problem–specific algorithm [8]. Most real–valued EAs have only been thoroughly tested on benchmark problems. Here we want to investigate the advantages and disadvantages of the use of EAs in a real–world (medical) application.

A minimally invasive vascular intervention is a form of surgery in which the instruments of operation enter the body, specifically the vasculature, through a small incision. These instruments are then navigated towards the vascular abnormality to be treated. Although compared to traditional open vascular surgery, minimally invasive vascular interventions have some important advantages for patients, it is hard to master the required skills. Therefore, thorough training is needed. Simulation is becoming an accepted and established possibility for training. In this paper we experiment with one such simulation system [2] that focuses on the main instrument used during these interventions, the guide wire.

2 Modelling and optimization

At the heart of the simulation lies an optimization problem rooted in elementary physics. The main idea is that the preferred configuration of the guide wire is one in which the associated energy is minimal. To model this concept, a simplified representation of energy is used that consists of only two sources. An internal energy is associated with the flexibility of the guide wire and an external energy is associated with the deformation properties of the vasculature. The objective of the simulation now is to keep the guide wire in a state of minimal energy.

The simulation is based on quasi–static mechanics, which means that the velocity and acceleration aspects associated with the propagation of the guide wire are ignored. This is a valid approximation, as much of human movement is performed at moderate speed, and is therefore essentially quasi–static [7].

2.1 Guide wire, vasculature and energy model

The guide wire is modelled using a set of joint positions $(\boldsymbol{x}_0, \ldots, \boldsymbol{x}_{k-1})$, where joint $\boldsymbol{x}_0$ is the tip of the guide wire and joint $\boldsymbol{x}_{k-1}$ is the bottom. Between each two joints a straight, not bendable or compressible segment is defined with a predefined constant length (λ_i). The thickness of the guide wire is taken into account by using sample points on the outer hull of the guide wire instead of using the joint positions when calculating the interaction with the vasculature (Figure 1).
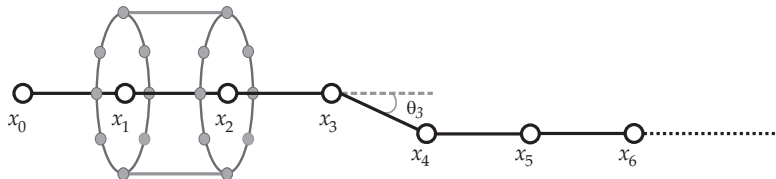

Figure 1: Representation of a guide wire. The angle (θ_i) between the two segments connected by joint i as used in equation 1 is illustrated for joint 3. For joint 1 and 2 the sample points on the outer hull of the guide wire, which allow the modelling of guide wires with different thickness, are drawn.

Hooke's law [3] is used as a basis for modelling the flexibility of the guide wire (internal bending energy) as well as for modelling the deformation of the vasculature (external vessel–wall energy). The bending energy $(U_{\mathrm{b}}(i))$ per joint i and the total bending energy (U_{Tb}) for the guide wire are defined as follows:

$$U_{\mathrm{b}}(i) = \frac{1}{2} c_i \theta_i^2, \ U_{\mathrm{Tb}} = \sum_{i=1}^{k-2} U_{\mathrm{b}}(i) \tag{1}$$

where θ_i denotes the angle between two segments connected by joint i and c_i is a spring constant related to the flexibility of the joint. The dependency between the flexibility of the guide wire and the segment length has been described elsewhere [2, 8] and can be formulated as $c_i = \frac{EI}{\lambda_i}$. Here, EI denotes the flexural rigidity of the guide wire, which is defined as the product of elasticity (E) and the second moment of area (I).

The external vessel–wall energy per joint ($U_{\mathrm{vw}}(i)$) and the external vessel–wall energy for the complete guide wire (U_{Tvw}) are given by:

$$U_{\mathrm{vw}}(i) = \frac{1}{2}ld_i^2, \quad U_{\mathrm{Tvw}} = \sum_{i=0}^{k-1} U_{\mathrm{vw}}(i) \tag{2}$$

The elasticity constant of the vessel wall is denoted by l and the deformation of the vessel wall by d_i. Note that the total energy to be minimized in the simulation is just the sum of U_{Tb} and U_{Tvw}.

2.2 Processing of actions

To simulate the behavior of the guide wire as it is propagated inside the vasculature a continuous propagation action is discretized into a sequence of small, forced motions. To simulate the outcome of each small forced translation of size ξ_{internal}, first each joint position is moved in the propagation direction by a length of ξ_{internal}, i.e. as if there were no vasculature. Subsequently, the influence of the vasculature must be taken into account. To this end, the total energy must be minimized to return the simulation into a state of physical equilibrium. To this end, an optimization problem must thus be solved. Therefore, an optimization algorithm is called for, which is the subject we turn to next.

3 Optimization methods

In this section we discuss three different optimization methods for finding a configuration of joint positions $\boldsymbol{x}_0, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_{k-3}$ that correspond to a minimal value of $U_{\mathrm{Tb}} + U_{\mathrm{Tvw}}$. The bottom two joint positions may not be varied because they define the propagation direction, which is fixed.

The dimensionality of the optimization problem can be reduced using a two–dimensional parameterization (instead of the three–dimensional joint positions). Such a parameterization is possible because the distance between two neighbouring joints always remains constant; each subsequent joint must lie on the surface of a sphere, which can be parameterized using only two dimensions. The actual mapping can be constructed in various ways (see for instance [10]).

It is important to realize that the goal is not to globally minimize the value of $U_{\mathrm{Tb}} + U_{\mathrm{Tvw}}$. The reason for this is that previous actions determine the path followed in the vasculature so far. We do not want the guide wire to switch from one branch of the vasculature to another one during optimization, because this is a physical impossibility. Local search methods therefore seem a reasonable approach, especially if ξ_{internal} is not too large and the nearest local optimum is the one that would have been obtained if a continuous motion had been used.

3.1 First order analytical solution

An explicit analytical solution in the first order approximation has been derived to minimize the energy per joint [8]. Instead of working with the joint positions, in this approach a vector indicating the difference in displacement of two neighboring joints is used. The solution indicates how a single joint should be displaced such

that the length of the segment between this joint and its predecessor is preserved. An algorithm has been formulated that applies this analytical solution to each joint in sequence, starting at the tip and moving towards the bottom of the guide wire. As a result of displacing joint $\boldsymbol{x}_i$ however, all joints $\boldsymbol{x}_j$, $j \in \{i-1, i-2, \dots, 0\}$ must also be displaced to preserve the segment lengths, possibly moving them away from their optimal position. Therefore, to ensure convergence, this process is repeated several times.

3.2 Conjugate gradients

The optimization problem to be solved is inherently real–valued and continuous. Therefore, classical well–known gradient–based optimization methods can be applied. In this paper we have performed experiments with the well–known conjugate gradients algorithm [6]. This algorithm makes efficient use of gradient information to find a local minimum.

3.3 IDℰAs and GLIDE

The field of evolutionary computation offers algorithms that are even more generally applicable for real–valued continuous optimization. They are less dependent on smooth gradients for instance. A specific EA that has recently been shown to give good results on a variety of benchmark problems is the Iterated Density–Estimation Evolutionary Algorithm (IDℰA) [4]. The main difference with traditional EAs is that in IDℰAs a probabilistic model is learned on the basis of the selected solutions. The probabilistic model can capture various structural properties (e.g. regions of interest and correlations between variables) of the optimization problem at hand. By drawing new solutions from the probabilistic model these structural properties can be exploited to obtain more efficient optimization. Especially if the probabilistic model is a low–complexity, highly–generalizing probability distribution such as the normal distribution, IDℰAs tend to be good global optimizers. To speed up convergence it has been shown that the hybridization of such algorithms with local optimizers can be beneficial. For instance the Gradient–Leveraged IDℰA (GLIDE–EA) in which the conjugate gradients algorithm is applied to a small percentage of solutions in the population has been shown to give superior results. In this paper we performed experiments with GLIDE based on the normal distribution in which each variable is taken to be independent of all the other variables. This means that for each variable the mean and standard deviation of a one–dimensional normal distribution is estimated from the selected solutions and that a new solution is constructed by sampling one value per variable from the associated one–dimensional normal distribution. This approach was previously observed to give good results [4].

Since EAs are by construction capable of escaping local optima to a certain extent, additional precautions need to be taken to prevent physically incorrect behavior of the simulation. Therefore, we have added a penalty term to the total energy to be minimized. The penalty term increases exponentially if the guide wire is displaced by more than a predefined distance.

4 Experiments

The accuracy of the simulation has been determined by comparing physical experiments in which the guide wire was propagated over a distance of 175mm with simulation results. A highly realistic elastrat normal abdominal aorta phantom was used [1]. The vessels were extracted from a CT data set of this phantom by means of segmentation. A CT data set however is a discretized representation of the real world. Hence, the distance of a joint to the vessel wall as required in Equation 2 can only be computed in a discretized fashion. This may pose a problem for algorithms based on continuous gradients such as conjugate gradients. The first order analytical approximation (FOA) algorithm internally also works with gradient information. Therefore, this algorithm is also hampered. Preliminary experiments have shown however that using trilinear interpolation provides a smooth enough representation for these algorithms to work.

Simulation results are evaluated based on a root–mean–square (RMS) error measure. The RMS error is computed from the distances between the joint positions in the simulated guide wire and the corresponding points in a reference that has been extracted from the physical experiments. For a guide wire with k joints the error measure is defined as RMS $= \sqrt{\frac{1}{k}\sum_{i=0}^{k-1}(d_i^{\mathrm{SR}})^2}$, where $d_i^{\mathrm{SR}} = \|\boldsymbol{x}_i^{\mathrm{S}} - \boldsymbol{x}_i^{\mathrm{R}}\|$, $\boldsymbol{x}_i^{\mathrm{S}}$ denotes the simulated joint position and $\boldsymbol{x}_i^{\mathrm{R}}$ the corresponding reference point. All experiments were run on a 1.7 GHz Pentium with 512 Mb of RAM.

4.1 First–order analytical approximation algorithm

We have performed experiments with various values for ξ_{internal} and for λ_i. A selection of results is presented in Figure 3. For multiple combinations of parameter settings the first–order analytical approximation (FOA) algorithm supplies a guide wire configuration with an RMS error of around 1mm. The use of segment lengths larger than 3mm resulted in erroneous results. In the derivation of the solution that is adopted by the FOA algorithm some assumptions have been made that enforce constraints on (combinations of) parameter values, limiting the range of parameter values for which this algorithm supplies correct results. Hence, this algorithm is mostly suited to obtain results using small steps and small segment lengths. This brings the algorithm closer to the non–discretized real–world setting, but renders it less practical.

Note that the use of a smaller segment length does not imply that a smaller RMS error is obtained. An important reason for this may be that the actual minimum of the optimization problem does not correspond to the minimum of physical reality, caused by the simplified model of energy and kinetics. Therefore, an error around 1mm is likely to be close to the optimally obtainable result using this model. This remark holds of course for all three algorithms. The best result is achieved for a segment length of 3mm, combined with an internal stepsize of 0.3mm. The required time to obtain this result was 92.04 seconds. This can be reduced to 20.36 seconds by using less global iterations of the algorithm. However, the RMS error then increases to 1.232mm.
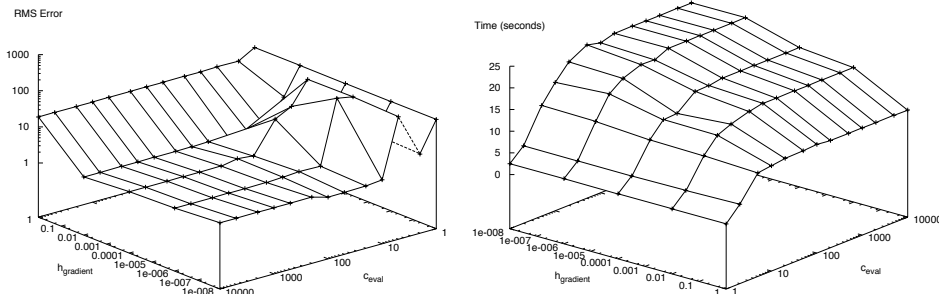
Figure 2: The RMS error (left) and the required time in seconds (right) as a function of two algorithmic parameters for conjugate gradients.

4.2 Conjugate gradients

Based on the result of the FOA algorithm, we performed initial experiments with the conjugate gradients algorithm in which the segment length was set to 3mm. We varied the values for the stepsize with which to approximate the gradient, denoted h_{gradient}, and the maximum allowed number of evaluations. The maximum allowed number of evaluations was set to $c_{\mathrm{eval}}k$ where k is the number of joints that the guide wire currently consists of. If convergence occurred before the maximum number of evaluations was reached, termination was enforced also.

As h_{gradient} decreases and $c_{\mathrm{eval}}k$ increases, the algorithm is allowed more precision and more time respectively to find a local optimum. It is to be expected that precision comes at the cost of time. Indeed this relation can be seen in Figure 2 where lower errors are associated with larger execution times. Also the influence of h_{gradient} and $c_{\mathrm{eval}}k$ can be seen clearly. The trade–off between precision and time can be made even more clear in a scatterplot, which is presented in Figure 3. The objective is to obtain an RMS error of 0, which corresponds to a simulation that perfectly mimics the physical experiment. Generally however, the lower the error, the larger the required time. Clearly, some algorithm settings lead to inferior results (i.e. larger runtimes as well as higher errors). The best trade–off results are found on the Pareto–front (i.e. the maximal set $\mathcal{P}_{F}$ of solutions for which there exists no solution that is superior to any of the solutions in $\mathcal{P}_{F}$). Based on these results, the value of $h_{\mathrm{gradient}} = 0.000001$ was found to be appropriate for the problem at hand. Moreover, it was observed that the algorithm often converges quite rapidly, for which reason we set c_{eval} to the maximum tested value of 10000.

Using these settings we performed experiments with three different values for the segment length and various different values for ξ_{internal}. For similar settings of the conjugate gradients algorithm for which the FOA algorithm succeeded, the conjugate gradients algorithm was found to be able to obtain similar RMS results, but required more time. Therefore, we chose to experiment with values such that we can explore the part of the simulation settings that can speedup the results of the FOA algorithm and to obtain results where the FOA algorithm failed. The results are shown in Table 1. Clearly, conjugate gradients then is much faster. Moreover, although the RMS error degrades, the results are still highly useful.

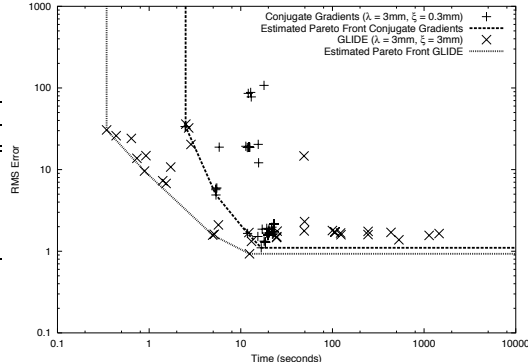| $\lambda_i/\xi_{int.}$	$\lambda_i\ (mm)$		
	1	2	3
10	0.996	1.066	0.983
20	1.217	0.995	1.075
40	1.141	1.117	1.097
80	1.276	1.120	1.140
160	1.215	1.130	1.151

Figure 3: Results of the FOA algorithm (left) and a scatterplot of results for conjugate gradients and GLIDE for different simulation settings.

4.3 GLIDE

Since the conjugate gradients algorithm is used within GLIDE, the only room for improvement is for those results where the conjugate gradients algorithm obtained results with a much larger RMS error than around 1mm. Therefore we aimed at performing experiments similar to the ones performed with conjugate gradients.

First, we performed experiments in which we varied the population size and the value of c_{eval}. However, this time we immediately performed experiments with a much larger internal stepsize. Therefore, the results as shown in Figure 3 should not lead the reader to conclude superior performance of GLIDE compared to conjugate gradients. The results do however indicate a similar trade–off between required time and obtainable RMS errors. Based on these experiments we experimented further with population sizes of 20 and 50 and $c_{\mathrm{eval}} \in \{50, 10000\}$. The results are tabulated in Table 1. From these results we can conclude that GLIDE is indeed the better optimizer since lower RMS errors can be found if c_{eval} is set large enough. However, this greatly comes at the expense of time. Since the final objective is to build a real–time simulator, these settings are not preferable. However, by choosing for instance a population size of 20 combined with a c_{eval} of 50, the EA is still able to come up with better results than the conjugate gradients algorithm in most cases, but it takes only marginally more time to run it. Moreover, the EA can still find good solutions for extreme cases. In these cases, only a few joints are required because the segment length is relatively large. Therefore, as the guide wire is propagated further, the scalability of the EA is much better as the optimization problem to solve contains less variables.

5 Discussion and conclusions

We have investigated the applicability of a hybrid real–valued IDEA to a practical problem: the simulation of minimally invasive vascular interventions. We compared the approach with the conjugate gradients algorithm and a problem–specific

λ_i	ξ_{int}	Conjugate gradients		GLIDE 20–50		GLIDE 20–10000		GLIDE 50–50		GLIDE 50–10000	
		RMS	Time	RMS	Time	RMS	Time	RMS	Time	RMS	Time
3	3	1.789	3.703	1.697	12.344	1.715	106.328	0.930	12.360	1.381	531.140
3	6	1.678	3.093	1.289	6.281	1.678	78.562	1.619	6.313	1.652	256.360
3	8	1.453	2.718	1.808	5.000	1.618	81.704	4.043	5.016	1.666	270.438
3	12	1.571	2.046	1.701	3.469	2.416	6.344	1.273	3.500	1.628	126.187
3	14	8.313	0.718	1.118	2.984	1.787	55.328	4.753	3.015	1.395	179.578
3	21	15.782	0.437	12.424	2.172	14.360	91.187	14.308	2.187	12.991	273.234
9	3	1.923	0.532	1.894	1.344	1.970	12.688	1.552	1.375	2.037	70.907
9	6	3.058	0.312	2.892	0.688	2.275	8.765	2.987	0.703	2.036	35.313
9	8	1.845	0.266	1.923	0.547	2.433	6.547	2.496	0.546	2.424	36.329
9	12	2.492	0.203	2.773	0.391	1.654	37.031	2.707	0.391	2.417	21.328
9	14	2.338	0.219	2.656	0.328	2.397	3.672	4.055	0.344	2.235	19.875
9	21	7.028	0.109	6.602	0.235	2.458	3.156	6.648	0.250	2.344	12.438
18	3	2.430	0.188	2.333	0.344	3.375	6.953	3.464	0.359	3.383	25.703
18	6	3.133	0.110	123.961	0.172	3.265	3.406	38.017	0.187	3.049	14.297
18	8	2.766	0.094	2.855	0.140	3.013	2.594	2.911	0.140	3.037	12.265
18	12	5.834	0.063	2.987	0.110	3.329	1.969	5.541	0.110	3.039	6.625
18	14	2.892	0.078	2.855	0.094	2.996	1.281	4.049	0.094	3.018	6.234
18	21	2.995	0.062	3.432	0.063	3.070	1.141	2.986	0.078	3.085	5.797

Table 1: Additional results for conjugate gradients and GLIDE. The RMS error in mm and the required time in seconds are shown. The index for GLIDE represents: population size–c_{eval}.

optimization algorithm. The results indicate that although the application of the conjugate gradients algorithm already leads to highly useful results, IDℰAs can yet improve on these results in the area of scalability. This important observation indicates that bringing IDℰAs into practice is indeed a good idea.

References

[1] Shelley medical imaging technologies. http://www.simutec.com/Shelley/models/AbdominalAorta.html

[2] T. Alderliesten, M. K. Konings, and W. J. Niessen. Simulation of minimally invasive vascular interventions for training purposes. *Computer Aided Surgery*, 2004. Accepted for publication

[3] G. Arfken. *Mathematical methods for physicists.* Academic Press, Inc., San Diego, California, 1985

[4] P. A. N. Bosman. *Design and Application of Iterated Density–Estimation Evolutionary Algorithms.* PhD thesis, Universiteit Utrecht, Utrecht, 2003

[5] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learing.* Addison Wesley, Reading, Massachusetts, 1989

[6] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *J. of Research of the National Bureau of Standards*, 6(49):409–436, 1952

[7] N. Hogan and J. M. Winters. Principles underlying movement organization: upper limb and single joint. In J. M. Winters and S. L-Y. Woo, editors, *Multiple Muscle Systems: Biomechanics and Movement Organization*, pages 182–194. Springer–Verlag, New York, 1990

[8] M. K. Konings, E. B. van de Kraats, T. Alderliesten, and W. J. Niessen. Analytical guide wire motion algorithm for simulation of endovascular interventions. *Medical & Biological Engineering & Computing*, 41:689–700, 2003

[9] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipies In C: The Art Of Scientific Computing.* Cambridge University Press, Cambridge, Massachusetts, 1992

[10] F. Wattenberg. Spherical coordinates. Department of Mathematics, Montana State University, Bozeman, MT 59717. http://www.math.montana.edu/frankw/ccp/multiworld/multipleIVP/spherical/body.htm, 1997