SENSITIVITY DISCREPANCY FOR GEOMETRIC PARAMETERS

http://people.sc.fsu.edu/~jburkardt/presentations/sanfran_1995.pdf

....

John Burkardt *
Max Gunzburger *
Interdisciplinary Center for Applied Mathematics
Virginia Tech
Blacksburg, VA, 24061

December 23, 2010

ABSTRACT

When discretized sensitivities are used to approximate the sensitivities of a discretized variable, a discrepancy may occur, particularly when the underlying parameter represents a geometric or implicit quantity. For a model problem, the primary source of error is found to be in the boundary conditions, which in turn are affected by errors in approximate spatial derivatives. A related optimization problem shows that the discretized sensitivities provide superior approximation of the behavior of the continuous variables.

INTRODUCTION

The sensitivities, or partial derivatives of state variables with respect to problem parameters, provide a very useful analytic tool for a parameterized state system. But it may be quite difficult to deduce the appropriate sensitivity system when the state system was derived by discretization via, for instance, the finite element method, from a continuous state system. In such a case, it may be more convenient to derive the *discretized sensitivity* system, which reverses the order of application of discretization and differentiation to the continuous state equation. For implicit or geometric variables, the discretized sensitivities should not be expected to be equal to the sensitivities of the discretized state variables. Instead, an approximation error ensues, which typically diminishes as some power of the discretization parameter h. Compounding this error, the boundary conditions defining the discretized sensitivity may have to be approximated. These two sources of error can result in a computed discretized sensitivity which is an unsuitable approximation to the sensitivity of the discretized

*Work supported by the Air Force Office of Scientific Research under grant F49620-93-1-0280

Figure 1: Flow region for $\alpha = 1$.

state variables.

A typical situation where sensitivities are useful occurs in the optimization of a cost functional associated with a state solution. The computation of the cost gradient via the chain rule results in a formula involving the partial derivatives of the cost with respect to the state variables, and of the state variables with respect to the parameters (the sensitivities). Errors in estimating the sensitivities can result in poor approximation of the cost gradients, leading to a failure of the optimization itself.

THE MODEL PROBLEM

The problem to be considered is the steady two-dimensional flow of a viscous incompressible fluid through a rectangular channel of height 3 units and length 10 units. The channel is partially obstructed by a parabolic bump whose height is controlled by a parameter, α. The original problem is related to a wind tunnel simulation studied by Huddleston (1990). A typical configuration is shown in Figure 1.

The state system comprises the time independent Navier Stokes equations, the continuity equation, zero velocities at the walls, straight outflow, a prescribed inflow, and the specification of the value of the pressure at a single point. The state variables are the horizontal and vertical velocities, u and v, and the pressure p. For brevity, the state variable u may be used to stand for all three.

In the usual way, the finite element method may be applied, with a discretization parameter h, to produce u^h (and v^h and p^h), a discrete approximation to the state variables. Because the flow region is not convex, the choice of the Taylor Hood element results in an approximation error that is $O(h^2)$ for the velocities u and v, and $O(h)$ for the pressure p. First spatial velocity derivatives, such as u_y, are only approximated to order $O(h)$, a point which will shortly cause difficulties. For details on this formulation, refer to Gunzburger (1989).

Both the continuous and discrete state variables will be assumed to be smoothly differentiable functions of the parameters, which will be denoted, for instance, as $u(x, y, \alpha)$. The bump parameter α influences the problem in an implicit way. That is, instead of explicitly influencing a quantity in the state system, it determines *where* a particular boundary condition is applied. The discretization method is also implicitly affected by α. This is primarily because the position of nodes along the moving bump boundary must be adjusted. How-

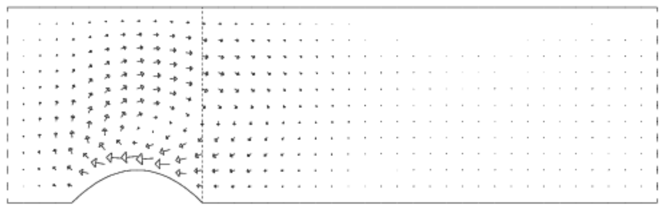

Figure 2: Discretized velocity sensitivities.

ever, for the model problem, the position of all interior nodes above the bump will also be adjusted as α varies, so as to avoid extremely distorted elements. This adjustment causes a further influence of the parameter α, which extends through all the elements above the bump.

COMPUTATION OF SENSITIVITIES

Derivatives such as u_α and u^h_α are respectively referred to as the *sensitivities* of the continuous and discrete state variables with respect to α. Sensitivities can be used to estimate the solution value at nearby parameter values, to assess the strength of the influence of a parameter, or to produce, via the chain rule, the derivative of a cost functional with respect to the parameters.

Once a solution to the continuous or discrete state system has been obtained, the sensitivities of that solution can be determined from a *sensitivity system*. When the parameter of interest occurs explicitly, then simple differentiation of the state system with respect to that parameter should be enough to determine the sensitivity system. However, for implicit parameters, such as those that control geometry, it may be necessary instead to carefully consider limits of difference quotients in order to determine the proper form of the sensitivity system.

Thus, once a solution u^h of the discretized state system has been computed, the sensitivity of the discretized state variable u^h_α can be found by solving the sensitivity system.

By contrast, the discretized sensitivities $\hat{u}^h_\alpha$ are found by computing the continuous sensitivity system, discretizing it, and solving it.

The definition of sensitivities and their use in the model problem and a variety of other applications is treated in Borggaard (1994), Borggaard et al. (1993), Burkardt (1995), and Burkardt and Peterson (1995a).

A typical discretized velocity sensitivity field for the bump parameter is displayed in Figure 2. The sensitivities may be regarded as the solution of the homogeneous Oseen equations. There are no source terms, and zero boundary conditions everywhere except along the bump. A suitable estimation of the bump boundary conditions is crucial to computing reliable sensitivities.

There are actually three reasonable ways to estimate a sensitivity:

- Derive the sensitivity system from the discrete state system, and solve it for the sensitivities;
- Discretize the continuous sensitivity system, computing the discretized sensitivity;
- Compute u^h at a nearby value of α and apply finite differences, producing an estimate which will be denoted by $\Delta_\alpha u^h$.

If, in fact, it is exactly the sensitivities of the discretized variable which are desired, then the first method has the advantage that it produces precisely these quantities, that is, the partial derivatives of the discrete solution u^h with respect to the parameter. And for a parameter which occurs explicitly, the formulation of the necessary sensitivity system can be surprisingly easy. In the case where the finite element method is used, this is because the differentiation operator passes under the integral sign of the finite element state system, and the test functions are unaffected by the parameter. Thus the sensitivity system for the discretized state variables is actually identical to the discretized sensitivity system.

For a geometric parameter, however, direct differentiation of the discretized state system can result in many new terms. This is because the region of integration is affected by changes in the parameter. Moreover, if the finite element method is used, the shape of individual elements, the placement of nodes, and the form of the test functions will all be affected as well. These dependencies are particularly irksome because, aside from changes at the boundary, they have little to do with the physical problem; instead, they are artefacts that arise from differentiating the discretization algorithm that is applied to the problem. The result can be a very cumbersome form for the differentiated discrete state system, and one which must be adjusted whenever the algorithm is modified.

For the second method, it is comparatively quite easy to derive the continuous sensitivity system, since no discretization operations have been applied. The same discretization algorithm can be applied to solve both the state and sensitivity systems. And if a Newton or quasi-Newton method has been used to solve the nonlinear state system, then the currently factored Newton matrix is exactly the matrix needed to solve the linear sensitivity system. Thus, the computation of discretized sensitivities is extremely cheap in terms of algorithm development and CPU time.

One possible drawback of using discretized sensitivities has already been noted: if the user desires the sensitivity of the discretized variables, then the discretized sensitivities are only approximately equal to those quantities, and the first method might be preferred. On the other hand, if the quantity of interest is the sensitivity of the continuous variables, (the physically meaningful quantity), then both the discretized sensitivities and the sensitivities of the discretized variables may be regarded as approximations to this quantity, and the choice between them should be made on other grounds.

In any case, the approximating power of the discretized sensitivities depends on the value of the discretization parameter h. If the value of h is appropriate for solution of the state system, but turns out to be too coarse for estimating the desired (discrete or continuous) sensitivities, a reduction of h, even by half, can be very costly. If the model problem is treated via finite elements, and a banded Gaussian solver is used, for instance, then halving h increases the work of factoring the system matrix by a multiplier of 16.

Finally, consider the third method, of finite differences. This has the great advantage that almost no special programming is required beyond that which computes the state solution itself. Instead, each parameter is in turn slightly perturbed, and the state solution is recomputed. A difference quotient then estimates the partial derivative. Moreover, if a given value of $\Delta\alpha$ does not seem to produce a suitable approximation, the computation can easily be repeated with a smaller value, at the same cost, until roundoff effects come into play.

On the other hand, each computed derivative comes at the cost of a full solution of the state system; in some cases, a single solution can be enormously expensive, and the cost of computing a new solution for each parameter to be investigated may be insupportable. There are also some subtle difficulties with this approach, particularly for a finite element formulation with moving nodes. If a particular node (x_i, y_i) is not moving, then we can compute the value of the finite difference approximation to the sensitivity of the discretized variables as:

$$\Delta_\alpha u^h(x_i, y_i, \alpha) = \frac{u^h(x_i, y_i, \alpha + \Delta\alpha) - u_i^h(x_i, y_i, \alpha)}{\Delta\alpha} \tag{1}$$

Unfortunately, this formula cannot be applied if the i-th node moves with α. In the model problem, all nodes above the bump may move, though only in the vertical direction. Then the correct formula to apply at node i with vertical coefficient $y_i(\alpha)$ is:

$$\begin{aligned} &\Delta_\alpha u^h(x_i, y_i(\alpha), \alpha) \approx \\ &\frac{u_i^h(x_i, y_i(\alpha + \Delta\alpha), \alpha + \Delta\alpha) - u_i^h(x_i, y_i(\alpha), \alpha)}{\Delta\alpha} \\ &-\frac{\partial u_i^h}{\partial y_i}\frac{\partial y_i}{\partial \alpha} \end{aligned} \tag{2}$$

Both the discretized sensitivities and the finite difference approach require the values of $\dfrac{\partial u_i^h}{\partial y_i}$ and $\dfrac{\partial y_i}{\partial \alpha}$. These should be easily computable from the discretization of u^h and from the behavior prescribed for the nodes $(x_i, y_i(\alpha))$ as α varies. For the Taylor Hood elements, u_y^h is not continuously defined across element interfaces, and so, for a finite difference approach, this term must be approximated, perhaps by the average of its values in each of the elements that impinge on the given node.

THE BOUNDARY CONDITION

Since the bump parameter does not occur explicitly in the continuous state system, simple differentiation will not produce the desired continuous sensitivity system. The equations will be derived following the approach of Burns et al.

(1991). Consider, then, a fixed point (x, y), and compare the values of u for α and for the perturbed parameter value $\alpha + \Delta\alpha$.

If (x, y) is strictly inside the flow region for a given value of α, then, for small perturbations $\Delta\alpha$, the perturbed flow region will still strictly contain the point (x, y). Therefore, both $u(x, y, \alpha)$ and $u(x, y, \alpha + \Delta\alpha)$ will satisfy the state system, and the limit of their difference quotient, u_α, satisfies the sensitivity equations formed by taking the partial derivative with respect to α of the state equations that apply inside the flow region. Assuming sufficient continuous differentiability, the orders of differentiation may be interchanged to produce the usual sensitivity system for the continuous problem.

Most of the original state boundary conditions may also be transformed by implicit differentiation with respect to the parameter, to yield boundary conditions for the discretized sensitivities. However, this approach cannot be used for the boundary conditions $u = 0$ and $v = 0$ along the bump, where the parameter sets the location of the bump boundary. To determine the form of the corresponding boundary condition for the sensitivity system at a node $(x_i, y_i(\alpha))$ that is moving with the boundary, consider the fact that the *total* derivative of u with respect to α at the node must be zero, since the value of u there is always 0. But this implies that at any such node:

$$0 = \frac{Du}{D\alpha} = \frac{\partial u}{\partial \alpha} + \frac{\partial u}{\partial y}\frac{\partial y_i}{\partial \alpha}, \tag{3}$$

which gives us the appropriate sensitivity boundary condition to apply there:

$$u_\alpha(x_i, y_i(\alpha), \alpha) = -u_y(x_i, y_i(\alpha), \alpha)\frac{\partial y_i}{\partial \alpha}. \tag{4}$$

This boundary condition completes the specification of the continuous sensitivity system that defines the quantities u_α.

Once the continuous sensitivity system has been defined, applying the discretization operation yields the corresponding discretized sensitivity system. However, the boundary conditions along the bump cause a problem, because they are given in terms of the spatial derivatives of the continuous solution u. But that solution is unavailable; only an estimate can be made from the discretization u^h. Thus, the boundary conditions are approximated with the term $-u^h_y \dfrac{\partial y_i}{\partial \alpha}$.

There remains the question of how to approximate u^h_y given u^h. The estimated derivative might be supplied by the discretization itself, as in finite elements, or by the use of finite differences applied to pointwise values of u^h. It is important, if possible, to choose an approximation for the derivative that does not significantly increase the error.

Three simple estimates are available to us:

- **FD2**, the two point finite difference formula;
- **FD3**, the three point finite difference formula;

Table 1: SENSITIVITY DISCREPANCY FOR U.

Method	$\alpha = 0.25$	$\alpha = 0.5$	$\alpha = 1.0$
	h = 1.000		
FD2	0.183	0.242	0.398
FD3	0.0698	0.138	0.343
FE	0.0802	0.162	0.433
	h = 0.500		
FD2	0.139	0.215	0.394
FD3	0.0515	0.105	0.234
FE	0.0651	0.139	0.320
	h = 0.250		
FD2	0.0852	0.151	0.308
FD3	0.0207	0.0534	0.157
FE	0.0259	0.0673	0.206
	h = 0.125		
FD2	0.0477	0.0906	0.207
FD3	0.00675	0.0192	0.0663
FE	0.00854	0.0255	0.0868
	h = 0.0625		
FD2	0.0253	0.0498	0.119
FD3	0.00213	0.00611	0.0225
FE	0.00299	0.00793	0.0322

- **FE**, the approximation supplied by the discretization method, in our case, the finite element method, which requires the averaging of nearby values.

The finite difference formulas use the value of u^h at the node where the boundary condition is to be applied, and at one or two nodes immediately above it.

To evaluate the performance of these approximations, each was used to define the boundary condition for the bump sensitivity parameter, the discretized bump sensitivities were solved for, and compared to the estimate produced by finite differences applied directly to the solution values, as in Equation 2. The sensitivity discrepancy was then recorded, that is, the maximum difference between $\hat{u}^h_\alpha$ and $\Delta_\alpha u^h$ over all the nodes. Several values of the bump parameter α and the discretization parameter h were considered. The results are presented in Table 1.

From the tables, it seems that for any method and value of h, the sensitivity discrepancy increases roughly linearly with α. This is plausible, since, as the bump rises, the flow velocity above the bump must increase. Moreover, larger bumps correspond to more disordered and nonlinear flow.

The table also makes clear that, for any fixed method and value of α, the discrepancy decreases, again in a roughly linear fashion, as h is decreased. This is an important computational confirmation of the asymptotic convergence of $\hat{u}^h_\alpha$ and u^h_α to the continuous limit, u_α. Moreover, the estimate of $O(h)$ error in $\hat{u}^h_\alpha$ correlates nicely with its dependence on the boundary condition, which was approximated to within $O(h)$. This suggests that the use of higher order elements, a local refinement of the mesh near the bump, or some other technique, would be required to improve the performance of the discretized sensitivities as approximants to u^h_α.

Note that the two point finite difference method is significantly less accurate than the others, although its error maintains the same overall rate behavior as the other methods when h is decreased or α is increased. Since our discretized solution is quadratic, and methods **FD3** and **FE** are exact for quadratics, there will be no reward for pursuing a higher order finite difference method.

From the previous remark, it may be wondered why the methods **FD3** and **FE** differ at all. But this is because the three points sampled by the **FD3** method may actually lie in two different elements, across which u^h is not continuously differentiable.

IMPROVED SPATIAL DERIVATIVES

Note that values of the derivative u^h_y are required at two different parts of the computation. First, these values are needed to define the boundary condition along the bump. Secondly, these values are needed to adjust the finite difference estimates, as in Equation 2.

The estimate of the error between the sensitivities u^h_α and the discretized sensitivities $\hat{u}^h_\alpha$ uses the continuous sensitivities u_α as an intermediate quantity:

$$\|\hat{u}^h_\alpha - u^h_\alpha\| \leq \|\hat{u}^h_\alpha - u_\alpha\| + \|u_\alpha - u^h_\alpha\| \tag{5}$$

Under the pessimistic view that the right hand side is not much bigger than the left, the error is made up of two parts:

- the discretization error of approximating u_α by u^h_α,
- the error committed approximating u_α by differentiating the discretized state variable u^h.

The first portion of the error actually includes an additional term, namely the error caused by approximating the boundary condition for the discretized sensitivities, an error which we know is roughly $O(h)$. The relative importance of this error may be judged by making at attempt to reduce it. The procedure will be to compute the flow solution u^h on a grid with a relatively fine mesh parameter $h = 0.0625$ and save the nodal values of u^h_y, which is taken to be a good approximation to u_y. Then the problem will be solved on a series of coarse meshes as before, except that the saved data will be used when values of u^h_y are needed for discretized sensitivity boundary conditions. These improved values will also be used when applying Equation 2. The results for this method, designated **FE+**, are shown in Table 2 and represent a dramatic improvement over the previous finite element results.

Table 2: IMPROVED SPATIAL DERIVATIVES.

Method	$\alpha = 0.25$	$\alpha = 0.5$	$\alpha = 1.0$
	h = 1.000		
FE	0.0802	0.162	0.433
FE+	0.0306	0.107	0.484
	h = 0.500		
FE	0.0651	0.139	0.320
FE+	0.0120	0.0304	0.144
	h = 0.250		
FE	0.0259	0.0673	0.206
FE+	0.00440	0.00433	0.0254
	h = 0.125		
FE	0.00854	0.0255	0.0868
FE+	0.00178	0.00133	0.00363

Moreover, as h decreases, the error is now dropping faster than $O(h)$. This suggests that the original sensitivity boundary value estimate which used the spatial derivatives of the discrete variable is in fact a poor approximant of u_y, and that this error accounts for a significant portion of the discrepancy.

APPROXIMATING A GRADIENT

To see how sensitivities are needed in a real computation, and how discrepancies in approximated sensitivities can affect (and even doom) that computation, consider a case discussed in Burkardt and Peterson (1995b). The problem is similar in geometry to the model problem, but an inflow parameter λ is added, and there are now three bump parameters α_1, α_2, and α_3. A cost functional $J(u^h)$ is defined, and a flow field is sought which minimizes J. Since the only free variables are the parameters, the problem may be recast as the minimization of the functional

$$\mathcal{J}(\lambda, \alpha_1, \alpha_2, \alpha_3) = J(u^h(\lambda, \alpha_1, \alpha_2, \alpha_3)). \tag{6}$$

In order to carry out a minimization, it is necessary to compute the partial derivatives, such as:

$$\frac{\partial \mathcal{J}}{\partial \lambda} = \frac{dJ}{du^h}\frac{\partial u^h}{\partial \lambda} \tag{7}$$

The value of $\dfrac{dJ}{du^h}$ should be easy to derive from the formula for J, and the second factor is just u^h_λ, the sensitivity of the discrete variables.

In a computation involving Equation 7, the sensitivities of the discrete variables were approximated by the discretized sensitivities. For a particular value of h, the resulting gradient field was so inconsistent with the cost functional that the optimization had to be halted. At the termination point, a step in the computed direction of descent would actually produce an increase in the cost functional.

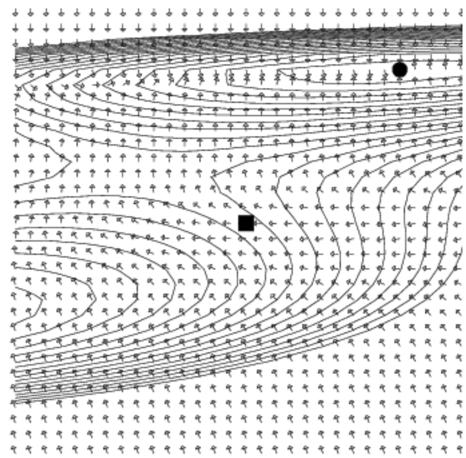

Figure 3: Approximated gradients, H = 0.5.

Gradients approximated by discretized sensitivities.
The global minimizer is at the solid disk.
The optimization halts at the square.

This situation is illustrated in Figure 3, which shows contours of the cost functional $\mathcal{J}$ evaluated on a two-dimensional "slice" through its four-dimensional parameter space. The approximate gradient field computed from the discretized sensitivities is superimposed. The inconsistency between the two sets of data is pictorially clear; there are many points where the approximate gradients are not perpendicular to the contours. The point where the optimization halted is shown, and it is clear that the gradient data is especially inconsistent with the contour values there.

But the gradients and contour values are presumably inconsistent because the value of h is large, which implies that the discretized sensitivities and the sensitivities of the discretized variables are both far from the continuous sensitivities, and hence from each other. One simple remedy is to reduce h. Figure 4 shows the same slice of parameter space when $h = 0.25$. The same contour levels are plotted, and it may be seen that they have shifted dramatically, while the gradients have changed much less. The two sets of data now appear to be consistent over a larger neighborhood of the minimizer. When the optimization algorithm is applied to this data, the global minimizer is correctly computed.

This graphical investigation raises an interesting point. Which quantity is really "inconsistent"? Refining the mesh one more time, as in Figure 5, makes it clear that the contour levels were very corrupt at $h = 0.5$, while the gradient field was quite close to being correct even at that coarse value of h. This suggests that the discrepancy between the discretized sensitivities and the sensitivities of the discretized variables may be due to the fact that the discretized sensitivities

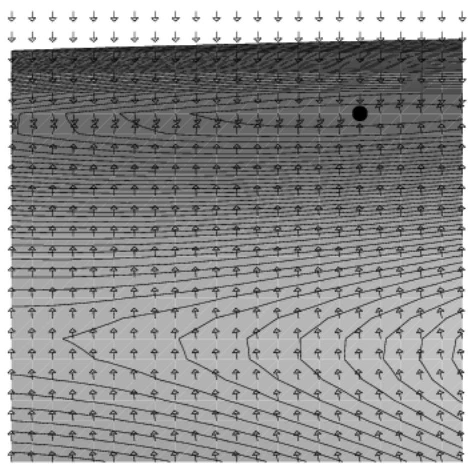

Figure 4: Approximated gradients, H = 0.25.

do a better job of approximating the continuous sensitivities.

CONCLUSION

Discretized sensitivities provide a simple and straightforward way of approximating the parameter sensitivity of a continuous or discretized state variable. Discretized sensitivities and sensitivities of discretized state variables should both be regarded as approximations to the sensitivity of the continuous state variables, and only secondarily as approximations to each other.

The discretized sensitivities can be significantly easier to compute, because only the original state equation is differentiated, not the discretized version. On the other hand, the sensitivities of the discretized variable have the advantage that they are consistent with the discretized solution u^h.

Discretized sensitivities may not be accurately computed if their boundary conditions or other data depend on knowledge of the continuous solution; the discrete solution u^h may be used in place of the unknown data involving u, but this can cause inaccuracies in the computation of the entire discretized sensitivity field, of a magnitude depending on the approximating error of u^h or its spatial derivatives.

Typically, the discretized sensitivities will be approximated to a lower order than the state solution itself. Therefore, a particular value of the discretization parameter h might be suitable for computing an approximate state solution but not for the discretized sensitivities. When accurate discretized sensitivities are needed, it may be advisable to employ a discretization scheme for the state variable with a higher than usual approximating power to guarantee satisfactory computation. In the case of the model problem, this would mean moving from the Taylor Hood element to a higher degree element.

For the model problem, the computed value of the discretized sensitivities is essentially determined entirely from the data specified on the bump boundary. Surprisingly, the major source of the discrepancy between the discretized sensitivities and the sensitivities of the discretized variables seems to arise when the boundary condition data involving u_y is approximated by the discretized solution u_y^h.

Finally, if discretized sensitivities are used as approximations to the sensitivities of the discretized variables, then a certain degree of approximation error may be expected. In such a case, it is important that data be monitored for consistency. For instance, gradient vectors computed via discretized sensitivities might be allowed to deviate no more than some tolerance angle from a descent direction of a cost function evaluated on a discretized variable. In this way, unacceptably bad computations can be terminated promptly, and corrective measures, such as reducing the mesh parameter h, recomputing with finite differences, or reverting to a higher order element, can be tried.

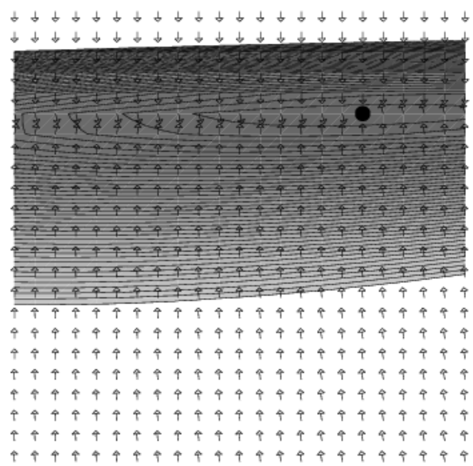

Figure 5: Approximated gradients, H = 0.125.

References

[1] J Boland and R Nicolaides, *The Stability of Finite Elements Under Divergence Constraints*, **SIAM Journal on Numerical Analysis**, Volume 20, Number 4, pages 722-731, August 1983.

[2] Jeff Borggaard, *The Sensitivity Equation Method for Optimal Design*, PhD Thesis, Virginia Polytechnic Institute and State University, Department of Mathematics, 1994.

[3] Jeff Borggaard, John Burns, Gene Cliff and Max Gunzburger, *Sensitivity Calculations for a 2D, Inviscid Supersonic Forebody Problem*, in **Identification and Control of Systems Governed by Partial Differential Equations**, H T Banks, R Fabiano, K Ito, editors, SIAM Publications, pages 14-24, 1993.

[4] John Burkardt, *Sensitivity Analyses and Computational Shape Optimization of Incompressible Flows*, PhD Thesis, Virginia Polytechnic Institute and State University, Department of Mathematics, 1995.

[5] John Burkardt and Janet Peterson, *Control of Steady Incompressible 2D Channel Flow*, **Flow Control**, The IMA Volumes in Mathematics and Its Applications, Volume 68, Springer Verlag, New York, 1995.

[6] John Burkardt and Janet Peterson, *The Role of the Cost Functional in a Shape Optimization*, **Proceedings of the Fourth Bozeman Conference on Computation and Control**, Birkhäuser, to appear 1995.

[7] John Burns, Eugene Cliff, Max Gunzburger, *Equations for the Sensitivities for the 2D, Inviscid, Supersonic Forebody Problem*, ICAM Report 91-07-05, Interdisciplinary Center for Applied Mathematics, Virginia Polytechnic Institute and State University, Blacksburg, VA, July 1991.

[8] Hayri Cabuk and Vijay Modi, *Optimum Plane Diffusers in Laminar Flow*, **Journal of Fluid Mechanics**, Volume 237, pages 373-393, April 1992.

[9] R G Carter, *On the Global Convergence of Trust Region Algorithms Using Inexact Gradient Information*, **SIAM Journal on Numerical Analysis**, Volume 28, Number 1, pages 251-265, February 1991.

[10] Carl DeBoor, **A Practical Guide to Splines,** Springer Verlag, New York, 1978.

[11] J E Dennis and Robert Schnabel, **Numerical Methods for Unconstrained Optimization and Nonlinear Equations**, Prentice Hall, Englewood Cliffs, 1983.

[12] Jack Dongarra and Eric Grosse, *Distribution of Mathematical Software Via Electronic Mail*, **Communications of the ACM**, Volume 30, pages 403-407, 1987.

[13] David Gay, *Algorithm 611, Subroutines for Unconstrained Minimization Using a Model/Trust Region Approach,* **ACM Transactions on Mathematical Software**, Volume 9, Number 4, pages 503-524, December 1983.

[14] Vivette Girault and Pierre-Arnaud Raviart, **Finite Element Methods for Navier-Stokes Equations**, Springer Verlag, Berlin, 1986.

[15] Ph Guillaume and M Masmoudi, *Computation of High Order Derivatives in Optimal Shape Design*, **Numerische Mathematik**, Volume 67, pages 231-250, 1994.

[16] Max Gunzburger, **Finite Element Methods for Viscous Incompressible Flows**, Academic Press, San Diego, 1989.

[17] Max Gunzburger and Janet Peterson, *On Conforming Finite Element Methods for the Inhomogeneous Stationary Navier-Stokes Equations*, **Numerische Mathematik**, Volume 42, pages 173-194, 1983.

[18] J Haslinger and P Neittaanmäki, **Finite Element Approximation for Optimal Shape Design**, John Wiley and Sons, Chichester, 1988.

[19] D Huddleston, *Development of a Free-Jet Forebody Simulator Design Optimization Method*, AEDC-TR-90-22, Arnold Engineering Development Center, Arnold AFB, TN, December 1990.

[20] Ohannes Karakashian, *On a Galerkin-Lagrange Multiplier Method for the Stationary Navier-Stokes Equations*, **SIAM Journal on Numerical Analysis**, Volume 19, Number 5, pages 909-923, October 1982.

[21] Hongchul Kim, **Analysis and Finite Element Approximation of an Optimal Shape Control Problem for the Steady-State Navier-Stokes Equations**, PhD Thesis, Virginia Polytechnic Institute and State University, Department of Mathematics, 1993.

[22] O A Ladyzhenskaya, **The Mathematical Theory of Viscous Incompressible Flow**, Second Edition, Gordon and Breach, New York, 1969.

[23] L Landau and E Lifshitz, **Fluid Mechanics**, Pergamon Press, Oxford, 1987.

[24] Werner Rheinboldt and John Burkardt, *A Locally Parameterized Continuation Process*, **ACM Transactions on Mathematical Software**, Volume 9, Number 2, pages 215-235, July 1983.

[25] Werner Rheinboldt and James Ortega, **Iterative Solution of Nonlinear Equations in Several Variables**, Academic Press, New York, 1970.

[26] Roger Temam, **Navier-Stokes Equations**, Revised Edition, North-Holland, Amsterdam, 1979.