

Users' Guide for imfil Version 0.7

C. T. Kelley

Version of September 21, 2008
Copyright ©2008 by C. T. Kelley

Contents

| | |
|---|-----------|
| Preface | iii |
| How to Get the Software | iv |
| 1 Getting Started | 1 |
| 1.1 Computing Environment | 1 |
| 1.2 What imfil.m does | 1 |
| 1.3 Basic Usage | 3 |
| 1.3.1 A Simple Example | 4 |
| 1.3.2 Calling imfil.m and Looking at Results | 5 |
| 1.3.3 Nonlinear Least Squares | 6 |
| 1.3.4 Parallel Computing | 7 |
| 1.3.5 Revisiting the Example | 8 |
| 1.4 Setting Options | 8 |
| 1.4.1 Scaling f | 9 |
| 1.4.2 Terminating the Iteration | 9 |
| 1.4.3 Changing the Scales | 10 |
| 1.4.4 Scale-Aware Functions | 10 |
| 2 Reference for imfil.m | 11 |
| 2.1 Input | 11 |
| 2.1.1 The Initial Iterate | 11 |
| 2.1.2 The Objective Function f | 12 |
| 2.1.3 The Budget | 12 |
| 2.1.4 The Bounds | 12 |
| 2.2 Output and Troubleshooting | 12 |
| 2.3 Options | 12 |
| 2.3.1 fscale | 13 |
| 2.3.2 Nonlinear Least Squares | 14 |
| 2.3.3 Limits on the Iterations | 14 |
| 2.3.4 Parallel Computing | 14 |
| 2.3.5 The quasi-Newton Method | 14 |
| 2.3.6 Scalestart and Scaleddepth | 15 |

| | | |
|--------|-------------------------------|-----------|
| 2.3.7 | Custom Scales | 15 |
| 2.3.8 | Scale_Aware | 15 |
| 2.3.9 | Stencils | 15 |
| 2.3.10 | Vstencil | 15 |
| 2.3.11 | Random Vectors | 16 |
| 2.3.12 | Termination Options | 16 |
| 2.3.13 | The Inner Iteration | 16 |
| 2.3.14 | Verbose | 17 |
| | Bibliography | 19 |
| | Index | 21 |

Preface

This is the users' guide for the MATLAB version of implicit filtering **imfil.m**. As of September 21, 2008, the code is ready for beta-testing. You will find that parts of the code, especially some of the parts deep inside, are still poorly documented.

I assume that you have a background in optimization at the level of [9, 15]. If you do not, and simply want to use **imfil.m** as a consumer, I have tried to make that possible, but make no guarantees.

A more complete account of **imfil.m**, including a review of the important ideas from traditional optimization, details of the algorithmic decisions, and some of the theory, will wind up in [17], a book in preparation. If you want a copy of the draft of the book, let me know. The draft will go to SIAM as part of a book proposal by the end of September 2008.

Implicit filtering is a projected quasi-Newton method for bound constrained optimization problems. The gradients are computed with a finite difference and the difference increment varies as the optimization progresses. The points on the difference stencil are also used to guide the search. In this way implicit filtering is a hybrid of classical smooth optimization methods and coordinate search.

As implicit filtering has evolved since [25], it and several related approaches have moved closer together. In the current version, as reflected in this book, ideas from [1, 8, 14, 21] have found their way into the implementation in **imfil.m**.

imfil.m is a MATLAB implementation of the implicit filtering method. This version differs in significant ways from our older FORTRAN code [6], and we support only the MATLAB version now. This document is a complete reference to Version 0.7 of **imfil.m**, covering installation, testing, and its use in both serial and parallel environments. 0.7 is an update of 0.6, with some new features and the expected bug fixes.

This is a work in progress, as you will see when you run into the **boldface notes to myself** that are scattered all over the place.

C. T. Kelley
Raleigh, North Carolina
September, 2008

Latest changes:

- You can solve nonlinear least squares using the methods from [16]. We have improved the algorithm since [16] was published and more improvements are on the way.
- I have changed the details of the options structure. This will only affect you if you have been **hacking** that structure.

How to get the software

The codes live at

<http://www4.ncsu.edu/~ctk/imfil.html>

On that page you will find

- A pdf file of the users' guide [18].
- **imfil.m**
This is the main implicit filtering code. This manual is current as of September 21, 2008. This is version 0.7. This version is a significant change from 0.6. If you are using any version prior to 0.7, please upgrade to 0.7. Of course, let me know if you have any problems.
- **imfil_optset.m** handles the options.
- The codes which generate Figures 1.1 and 1.2: **pid_example_chapter_1.m**, **pidobj.m**, and **pidlsq.m**.

One can obtain MATLAB from
The MathWorks, Inc.
24 Prime Park Way
Natick, MA 01760,
(508)653-1415
Fax: (508)653-2997
Email: info@mathworks.com
WWW: <http://www.mathworks.com>

Chapter 1

Getting Started

In this chapter we give a brief description of what **imfil.m** does, and illustrate its use for a simple problem. We do not describe all the options, the details of the algorithms, or the ways to extend the code on your own. The later chapters will do that.

1.1 Computing Environment

As of September 21, 2008, you will need to get the software, put the codes in your MATLAB path, and be running a recent version of MATLAB. We have tested **imfil.m** on versions 6.5 and higher.

imfil.m uses very little memory on its own. The codes which define your problem may use much more. MATLAB will complain if it runs out of memory, which is less likely if you run version 7.5 or later.

1.2 What imfil.m does

Implicit filtering solves **bound constrained optimization** problems :

$$\min_{x \in \Omega} f(x), \quad (1.1)$$

by which we mean that the goal is to minimize the **objective function** f subject to the condition that $x \in R^N$ is in the **feasible region** (or **nominal design space**)

$$\Omega = \{x \in R^N \mid L_i \leq (x)_i \leq U_i\}, \quad (1.2)$$

which is a **hyper-rectangle** in R^N . **imfil.m** is a MATLAB implementation of implicit filtering.

Implicit filtering, and the other methods that are derived from coordinate search, are best used in cases where f is either not smooth, not everywhere defined, discontinuous, or when derivatives of f are too costly to obtain. The motivating examples for the construction of implicit filtering were problems in which f was a

smooth function corrupted by low-amplitude, high-frequency noise, which was not defined (*i. e.* the code for computing f failed) at many points in the nominal design space Ω .

In the classical nonlinear programming problem [11, 13, 22] f is a smooth (*e. g.* twice Lipschitz continuously differentiable) function and Ω can be described by smooth inequality constraints, *i. e.*

$$\Omega_G = \{x \in R^N \mid g_i(x) \leq 0, 1 \leq i \leq M\}. \quad (1.3)$$

There are several good gradient-based methods and codes for solving this classical problem [3–5, 7, 12, 24]. Sampling methods such as implicit filtering are not among them.

Implicit filtering is a **sampling method**. By this we mean that the optimization is controlled only by evaluating f at a cluster of points in Ω . That evaluation determines the next cluster.

Implicit filtering’s samples are arranged on a stencil, and it is important to understand how that stencil is built. Given a current iterate x_c The default behavior is to sample the $2N$ points

$$x_c \pm hv_i \quad 1 \leq i \leq N,$$

where

$$v_i = (L_i - U_i)e_i,$$

e_i is the unit vector in the i th coordinate direction, and h , the **scale** varies as the optimization progresses. The sequence of scales

$$\{2^{-n}\}_{n=1}^{\text{scaledepth}}.$$

The algorithmic parameter **scaledepth** can be changed from the default of 7 with the **imfil_optset** command. The optimization will terminate when the sequence of scales has been exhausted.

imfil.m uses the values of f on the stencil in several ways, one of which is to construct a difference gradient and use that in a Quasi-Newton method. **imfil.m** reports results after each quasi-Newton iteration is complete. When the supply of scales has been exhausted, the optimization terminates.

imfil.m scales the bounds by changing variables so that $L_i = 0$ and $U_i = 1$ for all i . Scaling helps **imfil.m** take steps of relatively equal size in all the variables. You do not have to scale the variables. **imfil.m** does that for you.

Implicit filtering is able to respond to the function’s failure to return a value. When this happens, we say that a **hidden constraint** has been violated. **imfil.m** treats a point in Ω for which f has no value as missing data, and will proceed without the value. Your implementation of f must communicate a failure to **imfil.m**, and an example below.

The most common way to terminate a sampling algorithm is to assign a **budget** of function evaluations to the optimization, and to stop the computation when that budget is exceeded. When the function may fail, keeping track of the budget becomes a bit subtle, and your code for f must help **imfil.m** keep track of the

budget. One thing to consider, for example, is that sometimes a failed point is significantly cheaper to detect than a complete call to f .

So, at a minimum, you must give **imfil.m** an initial iterate, the objective function, the budget, and the bounds. **imfil.m** will return the optimal point x , and (optionally) a history of the iteration. You can use the history to evaluate the performance of the algorithm or see if you've gotten stuck.

1.3 Basic Usage

You must write a MATLAB code for f , which will take as its input $x \in R^N$ and return

- a value $f_{out} = f(x)$,
- a flag $ifail$ to signal a failed evaluation ($ifail = 0$ unless the evaluation fails, if the evaluation fails set $ifail = 1$ and $f_{out} = NaN$), and
- $icount$, an estimate of the cost.

So, the call to f would look like

```
[fout,ifail,icount]=f(x)
```

In the most simple case, $icount$ will be the number of calls to f . If some failed points cost more than others, or you want $icount$ to reflect some other measure of cost, such as wall clock time, you may do that as well. The budget will be computed in terms of $icount$.

You should put the bounds in a $N \times 2$ array, with L in the first column and U in the second.

Finally, you must specify a budget. **imfil.m** will examine a cumulative cost estimate (which uses $icount$) and terminate the optimization when the budget is exceeded. **imfil.m** will not interrupt an iteration in the middle, so you should expect a modest overshoot in the cost of the optimization. **imfil.m** will also terminate when the list of scales has been exhausted. § 2.3.12 describes other ways to terminate the iteration.

A complete call would look like

```
x=imfil(x0,f,budget,bounds);
```

or

```
[x,history]=imfil(x0,f,budget,bounds);
```

if you want the history of the iteration. Remember that if your objective function is a MATLAB .m file, say `myfun.m`, you'll have to put single quotes around the name of the function. Then the call would look like

```
x=imfil(x0,'myfun',budget,bounds);
```

`myfun.m` would have to be either in your MATLAB path or in the current directory.

The *history* array is an $IT \times 5$ array, where IT is the number of quasi-Newton iterations used in the optimization. For now we will concentrate on the first two columns, which contain the cumulative value of *icount* and the value of f at the end of the iteration.

1.3.1 A Simple Example

We will use an example from [2, 15]. In this example $N = 2$. The goal is to identify the damping c and spring constant k of a linear spring by minimizing the difference of a numerical prediction and measured data. The experimental scenario is that the spring-mass system will be set into motion by an initial displacement from equilibrium and measurements of displacements will be taken at equally spaced increments in time.

The motion of an unforced harmonic oscillator satisfies the initial value problem

$$u'' + cu' + ku = 0; u(0) = u_0, u'(0) = 0, \quad (1.4)$$

on the interval $[0, T]$. In (1.4) $u' = du/dt$.

We let $x = (c, k)^T$ be the vector of unknown parameters and, when the dependence on the parameters needs to be explicit, we will write $u(t : x)$ instead of $u(t)$ for the solution of (1.4). If the displacement is sampled at $\{t_j\}_{j=1}^M$, where $t_j = (j-1)T/(M-1)$, and the observations for u are $\{u_j\}_{j=1}^M$, then the objective function is

$$f(x) = \frac{1}{2} \sum_{j=1}^M |u(t_j : x) - u_j|^2. \quad (1.5)$$

This is a **nonlinear least squares** problem, but we will ignore the nonlinear least squares structure for the present. We show how to solve the nonlinear least squares problem in § 1.3.3 and 1.3.5.

We will use MATLAB's `ode15s` [23] to solve (1.4), and use the solution from `ode15s` to compute f . The first step in using `ode15s` is to convert (1.4) to a first order system for

$$y = \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} u \\ u' \end{pmatrix}.$$

The resulting first order system is

$$y' = F(y) = \begin{pmatrix} v \\ -cv - ku \end{pmatrix}, \quad (1.6)$$

with initial data $y(0) = (0, 0)^T$.

When one uses an integrator like `ode15s`, one is asked to provide a local truncation error tolerance. This tolerance controls the accuracy of the integration, and thereby the resolution in f .

The examples for this chapter are in a subdirectory of software collection, and this first example uses the files

- `pid_example_chapter_1.m`, which is the driver for all the examples in this chapter, and
- `pidobj.m`, which calls the least squares formulation `pidlsq.m` to compute f using `ode15s`.

We sample at $M = 101$ equally spaced points $\{i/100\}$ for $0 \leq i \leq 100$ and configure the problem so that the solution is $(c, k) = (1, 1)$. The objective function is `pidobj.m`, and the first line is

```
function [f,ifail,icount]=pidobj(x)
```

Functions you write must **AT LEAST FOR NOW** conform to this paradigm. Even if your function never fails and always does the same amount of work when called, `imfil.m` needs *ifail* and *icount*. **This will get fixed soon.**

`pidobj.m` has a failure mode. If either c or k is negative, the spring is not physical. The code traps this and returns without calling the integrator. You could fix this yourself, by making sure that the lower bounds you give to `imfil.m` are all nonnegative. We put this in as an example so you can see how to do it yourself.

Here is a sketch of what happens when `pidobj.m` receives x . You can look at `pidlsq.m` and `pidobj.m` for the details. The first thing to do is test for a negative component of x . The MATLAB looks like

```
ifail=0; icount=1;
%
% Call the integrator only if x is physically reasonable, ie if
% x(1) and x(2) are nonnegative. Otherwise, report a failure.
%
if min(x) < 0
    ifail=1; icount=0; f=NaN;
else
% call the integrator and do the work
end
```

We omit the details of the call to the integrator, and the way `pidobj.m` uses global variables. The curious reader can look at the source in the software collection.

1.3.2 Calling `imfil.m` and Looking at Results

We can now show how the simplest call would work. The plots in the upper row of Figure 1.1 reflect the a case with an intentionally poor choice of bounds

$$L = \begin{pmatrix} 2 \\ 0 \end{pmatrix} \text{ and } U = \begin{pmatrix} 20 \\ 5 \end{pmatrix},$$

which exclude the solution. We gave the optimization a budget of 100 calls to the integrator and an artificially low upper limit of five scales $\{2^{-n}\}_{n=1}^5$. We changed the set of scales from the default set $\{2^{-n}\}_{n=1}^7$ by using the `imfil.optset` command to change . The MATLAB commands were

```

bounds=[2 20; 0 5];
x0=[5,5]'; budget= 100;
options=imfil_optset('scaleddepth',5,options);
[x,histout]=imfil(x0,'pidobj',budget,bounds,options);

```

As you can see from the plot on the upper left of Figure 1.1, the iteration terminated before the budget had been exhausted. We can return to the default set of scales by either reinitializing the `options` structure

```
options=imfil_optset;
```

and calling `imfil.m` again. The picture on the upper right reflects the results of this change. Now the optimization slightly exceeds the budget, and the final value of the objective function is much lower.

One can query the `histout` array to see how far the optimization got in the list of scales. All the plots were made with the first two columns of the `histout` array. The plots at the top of Figure 1.1 were made with the command

```
plot(histout(:,1),histout(:,2),'-');
```

The first two columns of the `histout` array are the function values and the cumulative cost, measured in this case by calls to `ode15s`. When we look at the plots we see that the optimization has made very little progress after 75 or so calls to `ode15s`. You may modify the example code to add more scales and increase the budget, but the value of the function will decrease only a little, if at all. The reason for this is that we have resolved the optimal point as far as the resolution in the integrator will allow.

The plots on the bottom of Figure 1.1 are from an optimization where the global minimum is within the bounds. Here we set

```
bounds=[0 20; 0 5];
```

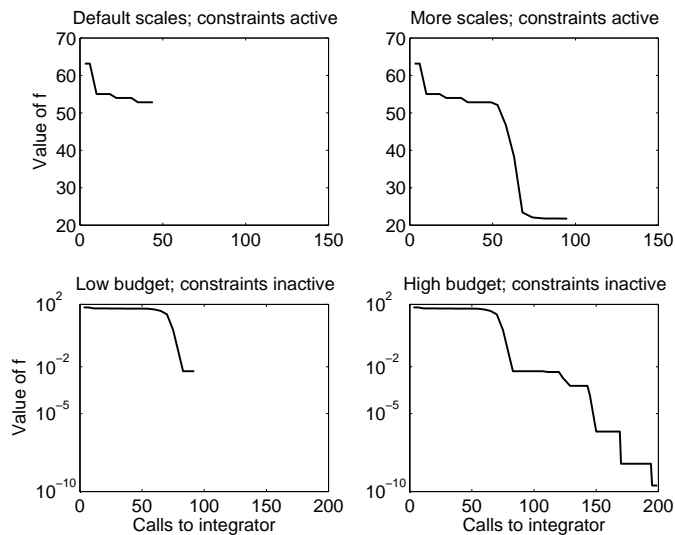
The lower left plot in Figure 1.1 shows the progress of the optimization with a budget of 100 and the default set of scales. In this case the budget and the number of scales are insufficient to fully resolve the optimal value. The lower right plot shows the results with a budget of 200 and 20 scales. The results are very different, which returns us to the issue of termination. How do we know when to stop the optimization? How can we tell if the budget is too small? These are open research questions at this time (2008). In § 2.3.12 we explore the options in `imfil.m` for termination.

1.3.3 Nonlinear Least Squares

Many problems, such as the example in this section, are best formulated as nonlinear least squares problems, where F returns a vector of residuals in R^M and the function to be minimized is

$$f(x) = F(x)^T F(x)/2. \quad (1.7)$$

You can tell `imfil.m` that your problem is a nonlinear least squares problem by setting the `least_squares` option to 1 with the command

Figure 1.1. *Optimization History: Parameter ID*

```
options=imfil_optset('least_squares',1,options);
```

If you do this you need to write your function so that $F \in R^M$ is returned. **imfil.m** will construct $f(x) = F(x)^T F(x)/2$ for you. The optimization method is also tuned to a nonlinear least squares computation, and the underlying method is a damped finite-difference Gauss-Newton iteration.

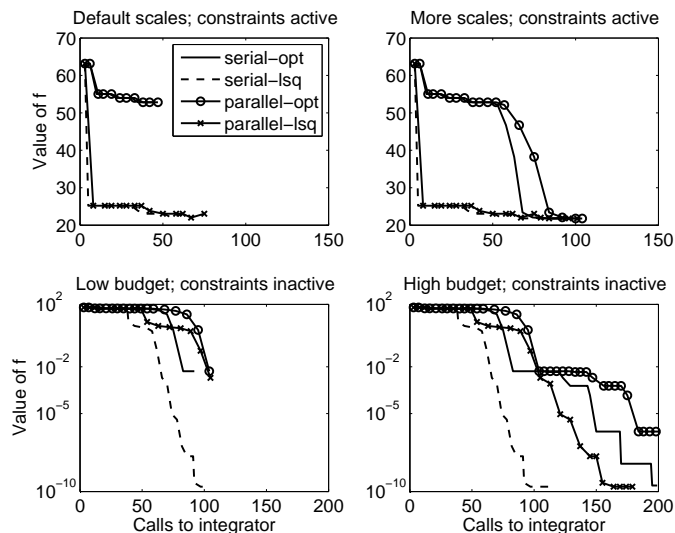
1.3.4 Parallel Computing

The `parallel` option tells **imfil.m** that f can be called with multiple arguments, and will return a matrix whose columns are the values of f , $ifail$, and $icount$. So if x is an $N \times P$ array of M arguments to f and `parallel` is set to 1, a call to $f(x)$ will return a $3 \times P$ matrix of values and flags. It is your responsibility to write f to do the parallel evaluation in an efficient way. Our example of a parallel call in § 1.3.5 only shows how **imfil.m** responds to

```
options=imfil_optset('parallel',1,options);
```

If you are solving a nonlinear least squares problem, where a call to f returns an $M \times 1$ column vector, your parallel function should return an $N \times P$ array of residual values as well as the row vectors of $iflag$ and $icount$. The parallel algorithm is not the same as the serial method because all the line search possibilities are examined at the same time (see § 2.3.4 for the details). One implication of this is that more function evaluations will be used even if the final results is the same as in the serial case and the total runtime is significantly less. One should interpret graphs like Figure 1.1 with care when one does the function evaluations in parallel. The default is `parallel = 0`.

Figure 1.2. *Optimization History: Parameter ID Revisited*



The latest versions of MATLAB support some parallelism, and there are some very useful resources in the **MATLAB Central File Exchange**. **MULTICORE** is a package that lets you use multiple cores with MATLAB. Each core runs its own copy of MATLAB. The package moves data between cores with file I/O, an approach with can slow down the computation if function calls are very inexpensive. The MATLAB parallel toolbox provides the `parfor` loop. Similarly to **MULTICORE** the `parfor` loop runs a copy of MATLAB on each core. Neither approach supports fine-grained parallelism, but both work well for very expensive function evaluations.

1.3.5 Revisiting the Example

We will close this section by reexamining the results in Figure 1.1 by comparing the serial optimization results in that figure with a serial least squares computation and parallel results. Clearly the least squares formulation is better because of the rapid convergence of the Gauss-Newton iteration for this small-residual problem. As one can see from Figure 1.2, the serial and parallel iteration histories differ by over 50%, which is not surprising since $N = 2$ and the parallel line search queries three or more possibilities at once. This is an extreme example of the difference between the parallel and serial algorithms.

1.4 Setting Options

You can set several algorithmic parameters with the `imfil_optset` options command. Many of these are rarely needed or intended for the specialist. We will discuss only the most useful and important in this section. You have already seen

the `scaledepth` option. We will explain the details for all the options in § 2.3.

Before setting the options, you need to get the default options structure from the `imfil_optset` command by calling that command with no arguments.

```
options=imfil_optset;
```

You need only do this once; additional calls to `imfil_optset` will update the the options structure you've already created. For example, if you want to change `fscale` to 1.0 and `scaledepth` to 10, you could call `imfil_optset` three times:

```
options=imfil_optset;
options=imfil_optset('fscale',1.0,options);
options=imfil_optset('scaledepth',10,options);
```

prior to the call to `imfil.m`. You can also put all the calls to `optset` on a single line

```
options=imfil_optset('fscale',1.0,'scaledepth',10);
```

If you want to change an existing set of options, you would add the name of the options structure to the `imfil_optset` command. For example, to change `scaledepth` from 10 to 8, the call would be

```
options=imfil_optset('scaledepth',8,options);
```

1.4.1 Scaling f

If the values of $|f|$ are very small or very large, the quality of the difference gradient which `imfil.m` uses in its search can be poor. `imfil.m` attempts to solve this problem by **scaling** the function by dividing it by the size of a “typical value”. Unless you tell `imfil.m` otherwise, this value is 1.2 times the absolute value of the value at the initial iterate.

You can change this by setting the `fscale` option. Setting `fscale` to a negative value will tell `imfil.m` to use $|fscale| \times |f(x_0)|$ as the typical value for f . Setting `fscale` to a positive value will tell `imfil.m` to use $fscale$ as the typical value. If you blunder and set $fscale = 0$, `imfil.m` will restore the default. See § 2.3.1 for more details on `fscale` and its role in `imfil.m`.

Scaling f to order 1 means that we can compare the variation in f to the scale and make a termination decision. The story on that will appear in § 2.3.12. **It is not in there yet. Stay tuned.**

1.4.2 Terminating the Iteration

There are two iterations which require termination parameters. The **inner iteration** is the quasi-Newton iteration for each value of h . The **outer iteration** is the implicit filtering iteration. In this section we will explain the default termination criteria and list some other ways to terminate these iterations. The details are in § 2.3.12.

The inner iteration will terminate

- if the value of f at the current point is smaller than the values elsewhere on the finite difference stencil, a condition we will call **stencil failure**,
- if the internal termination criteria of the quasi-Newton iteration are satisfied.

One can tune both of these criteria, and a user interested in doing that should look at [17] to understand the details.

The outer iteration, by default, terminates when either

- a budget of calls to the function has been exceeded or
- the list of scales has been exhausted.

The budget and the list of scales are input arguments to **imfil.m** and this mode of termination usually works well. One can do more, and set several options to terminate the iteration when

- the function value has been decreased to a desired target or
- the variation in the function on the stencil is sufficiently small.

See § 2.3.12 for the details.

1.4.3 Changing the Scales

imfil.m uses a stencil that is build from the bounds. If your current point is x_c , **imfil.m**'s default behavior is to sample the $2N$ points

$$x_c \pm h(L_i - U_i)e_i \quad 1 \leq i \leq N,$$

where e_i is the unit vector in the i th coordinate direction, and h , the **scale** varies as the optimization progresses. The sequence of scale is

$$\{2^{-n}\}_{n=\text{scalestart}}^{\text{scaleddepth}}.$$

`scalestart` and `scaleddepth` can be changed with the options command. The defaults are `scalestart = 1` and `scaleddepth = 7`.

imfil.m has a few options for messing with the stencil and future versions will let you do all kinds of stuff.

1.4.4 Scale-Aware Functions

Your function may be able to adjust its own accuracy or resolution. In this case we will say that your function is **scale-aware**. One example of this possibility is if the tolerance in a solver can be reduced as the scale is reduced. If your function has this capability, you may enable communication between **imfil.m** and the function call by adding the scale as an extra argument to f , making the call look like

```
[fout,ifail,icount]=f(x,h)
```

You must tell **imfil.m** that f is taking the extra argument by setting the `scale_aware` option to 1, the default is 0.

Put an example in here!

Chapter 2

Reference for `imfil.m`

The full calling sequence for `imfil.m` is

```
[x,histout]=imfil(x0,f,budget,bounds,options);
```

The last argument `options` is optional. If the `options` argument is omitted, `imfil.m` will use the defaults.

2.1 Input

The input data are

- $x0 \in R^N$: the initial iterate,
- f : the function to be minimized,
- *budget*, the maximum number of function evaluations allowed to the optimization,
- the bounds array `bounds`, and
- the `options` structure.

We will discuss all but the `options` argument in this section. We will explain the `options` at length in § 2.3.

2.1.1 The Initial Iterate

`imfil.m` requires a feasible initial iterate. This means that $x0$ must satisfy the bound constraints, *i. e.*

$$bounds(i,1) \leq x0(j) \leq bounds(i,2),$$

for all j , and that $f(x0)$ must be defined, *i. e.* f will return a value for $x0$ with `ifail` = 0.

2.1.2 The Objective Function f

If the `parallel` option is off (= 0), the calling sequence for f should be

```
[fout,ifail,icount]=f(x);
```

or

```
[fout,ifail,icount]=f(x,h);
```

if your function is **scale-aware**, *i. e.* can use the `scale` to manage its own internal control of accuracy.

If your function is scale-aware, set the `scale_aware` option to 1. .

Put an example of a scale-aware f somewhere.

If $f(x)$ successfully returns a value, `fout` should be that value, the failure flag `ifail` should be 0, and `icount` should be an estimate of the cost. `imfil.m` uses `icount` when comparing the cost of the optimization to the `budget` and to build the first column of the `histout` array, and you have the flexibility to assign non-integer values to `icount`. If, for example, a function call fails after performing half of the normal work, you might set `icount = .5`.

`ifail = 1` is the signal that the function cannot return a value, *i. e.* a **hidden constraint** has been violated. You may elect to return a NaN as the value when this happens, but that is not required. `imfil.m` will eliminate failed points from the stencil when computing the stencil gradient.

If the `parallel` option is on (= 1), then `imfil.m` will send an array of input arguments to f . These arrays vary in size. `imfil.m` will send the elements of the stencil that do not violate the bound constraints to f before it computes the stencil gradient. During the line search `imfil.m` will send every point that could be queried in the line search to f all at once; the default being the three points $\{x + \lambda d\}$ for $\lambda = 1, 1/2, 1/4, 1/8$. Your parallel function must be able to accept an $N \times M$ array of M arguments to f , and return three $M \times 1$ arrays of values for `fout`, `ifail`, and `icount`. It is your job to construct your function to use what parallelism you have efficiently.

2.1.3 The Budget

The optimization will terminate when the cumulative cost (as measured by `icost`) exceeds the `budget`. A budget that is too small will force premature termination (as will a list of scales that is too short). The optimization is likely to finish over budget because `imfil.m` does not stop the outer (optimization) loop in mid-stream.

2.1.4 The Bounds

2.2 Output and Troubleshooting

2.3 Options

You can change most of `imfil.m`'s algorithmic parameters with the `imfil_optset` command. One way to do this is to begin with a call with no arguments.

```
options=imfil_optset;
```

The output of this call is a MATLAB structure with the default options for **imfil.m**. You need only do this once and then use **imfil_optset** to update the options structure you've created. So, if you want to change **scaledepth** to 20 and use the SR1 quasi-Newton update, you could call **imfil_optset** three times

```
options=imfil_optset;
options=imfil_optset('quasi',2,options);
options=imfil_optset('scaledepth',20,options);
```

prior to the call to **imfil.m**. You can also put all the calls to **optset** on a single line

```
options=imfil_optset('quasi',2,'scaledepth',20);
```

If you wish to use the **options** structure, you add that as a final argument to **imfil.m** when you call it. So your call would look like

```
[x,history]=imfil(x0,f,budget,bounds,options);
```

instead of

```
[x,history]=imfil(x0,f,budget,bounds);
```

2.3.1 fscale

If the values of $|f|$ are very small or very large, the quality of the difference gradient which **imfil.m** uses in its search can be poor. **imfil.m** attempts to solve this problem by **scaling** the function by dividing it by the size of a “typical value”, which we call *imfil_scale*.

The default is

$$imfil_scale = 1.2|f(x_0)|,$$

which is usually fine.

If *imfil_scale* is too large, the quasi-Newton iteration within **imfil.m** may terminate too soon, and you may fail to exhaust the information in the current scale. This can lead to poor results, or even complete stagnation (*i. e.* x_0 is never changed).

If *imfil_scale* is too small, the quasi-Newton steps may be too large, and the line search may fail. In this case **imfil.m** becomes a form of coordinate search, and the performance will suffer.

You can change this by setting the **fscale** option. Setting **fscale** to a negative value will tell **imfil.m** to use

$$imfil_scale = |fscale||f(x_0)|,$$

so $fscale = -1.2$ is the default. If $fscale > 0$ then

$$imfil_scale = fscale.$$

$fscale = 0$ is not a sensible value; if you blunder and set $fscale = 0$, **imfil.m** will restore the default.

2.3.2 Nonlinear Least Squares

`/code/` will also solve nonlinear least squares problems where

$$f(x) = F(x)^T F(x)/2.$$

You tell the code that you have a nonlinear least squares problem by setting `least_squares` option to 1 with the command

```
options=imfil_optset('least_squares',1,options);
```

And write your function so that F is returned. `imfil.m` will compute f for you.

As of September 21, 2008 the solver is a damped Gauss-Newton iteration [15].

2.3.3 Limits on the Iterations

`imfil.m` has an outer iteration over the scales, an inner iteration in the finite-difference quasi-Newton loop, and an iteration within the line search. All of these iterations have limits you can set. You set the limit on the number of scales with `scalestart` and `scaledepth` (see § 2.3.6).

`maxit` is the upper limit on the number of quasi-Newton (or Gauss-Newton) iterations. The default is 50.

The line search will reduce the step at most `maxitarm` times before returning a failure. The default is 3. The line search is limited in this way for good reason. If you don't find something useful after three reductions, you're not likely to do better with more effort.

2.3.4 Parallel Computing

The `parallel` option tells `imfil.m` that f can be called with multiple arguments, and will return a matrix whose columns are the values of f , `ifail`, and `icount`. So if x is an $N \times P$ array of M arguments to f and `parallel` is set to 1, a call to $f(x)$ will return a $3 \times P$ matrix of values and flags.

If you are solving a nonlinear least squares problem, where a call to f returns an $M \times 1$ column vector, your parallel function should return an $N \times P$ array of residual values as well as the row vectors of `iflag` and `icount`. The parallel algorithm is not the same as the serial method because all the line search possibilities are examined at the same time.

The default is `parallel = 0`.

Code fragments will go here.

2.3.5 The quasi-Newton Method

The `quasi` option lets you choose the quasi-Newton method. The options are BFGS (`quasi = 1`), SR1 (`quasi = 2`), and no model Hessian (`quasi = 0`). The default is `quasi = 1`, the BFGS update.

2.3.6 Scalestart and Scaledepth

imfil.m samples f on a stencil centered at the current point. The size of that stencil varies at the optimization progresses. The default shape of the stencil is a central difference stencil with $2N$ points. The range of sizes can be controlled by the **scalestart** and **scaledepth** option.

If the directions in the stencil are vectors $\{v_i\}_{i=1}^m$, **imfil.m** will sample f at the points

$$x_c + h(L_i - U_i)v_i$$

for $1 \leq i \leq m$. The default vectors are the $2N$ unit vectors in the positive and negative coordinate directions. The **scale** h varies as the optimization progresses. The sequence of scale is

$$\{2^{-n}\}_{n=\text{scalestart}}^{\text{scaledepth}}.$$

scaledepth can be changed with the **imfil.optset** command. The defaults are **scalestart** = 1 and **scaledepth** = 7. If you see stagnation in the iteration, reducing **scaledepth** will save some effort, but be aware of the risk of early termination.

2.3.7 Custom Scales

The default scales are $\{2^{-n}\}_{n=1}^7$. The iteration will terminate when the supply of scales has been exhausted.

Custom scales on the way soon.

2.3.8 Scale_Aware

The **scale_aware** option tells **imfil.m** that your function is scale-aware. If **scale_aware** is set to 1, **imfil.m** will use the scale as a second input argument to f .

2.3.9 Stencils

As for September 21, 2008, **imfil.m** offers three stencils. You can change from the default centered difference stencil with the **stencil** option. The choices are a one-sided difference stencil, which uses the positive coordinate e_i if $x_c + he_i$ satisfies the bound constraints, and $-e_i$ otherwise, and the **positive basis stencil** [19,20] which uses the $N + 1$ points $\{e_i\}_{i=1}^N$ and

$$v_{N+1} = -\sum_{i=1}^N e_i.$$

The **stencil** options are 0 for the default, 1 for the one-sided stencil, and 2 for the positive basis stencil.

2.3.10 Vstencil

If you don't like any of the build-in stencils, you may create your own by setting the **vstencil** option to a matrix with your directions in the columns.

To do that, create a matrix VS with your directions in the columns, and then

```
options=imfil_optset('vstencil',VS,options).
```

2.3.11 Random Vectors

You can augment the stencil with k random vectors by setting the `random_stencil` option to k . The theory from [1, 10] will apply if $k > 1$.

The default is $k = 0$ (no random vectors) because we have seen better performance overall with the basic centered difference stencil. One reason for this is that more vectors make it likely that stencil failure will not happen, and then the iteration spends too much time in the line search.

If you suspect that the optimal point is on a constraint boundary, especially a hidden constraint boundary, then using random vectors is a useful idea.

2.3.12 Termination Options

Most problems can be solved with the default termination criteria for the optimization (or outer) iteration. These are exhausting the list of scales or exceeding the budget. Sometimes, however, you may know things that can help `imfil.m` do its job better.

The quasi-Newton (or inner) iteration has its own termination criteria. Changing the termination criteria for the inner iteration requires detailed knowledge of how the quasi-Newton loops are implemented, and we refer the reader to [17] for the details.

The Outer Iteration

Target

You may set a `target` value for the optimization. The optimization will terminate once f is below the target. The default value is -10^8 which means that `target` will play no role in the optimization.

Function_Error

If you know how accurate your function is, you may want to terminate once the variation of the function is smaller than your estimate for the error in the function. Setting `function_error` to your estimate of the absolute error in f will terminate the optimization when the maximum absolute difference of function values on the stencil is smaller than `function_error`. To turn this optional termination test on set the `function_error` option to your estimate of the error. The default is -1 which means the option is off.

2.3.13 The Inner Iteration

You have a good deal of control over the inner iteration, but should take care before messing about with these options.

If your problem is a nonlinear least squares problem, then the inner iteration is (for now) limited to damped Gauss-Newton. For general optimization problems you may set the `quasi` option to 0 (steepest descent, *i. e.* the model Hessian is the identity matrix), 1 (BFGS) or 2 (SR1).

Because **imfil.m** is intended for small problems, **imfil.m** maintains an approximation to the full model Hessian and does not use a sparse or limited-memory [15] formulation of the quasi-Newton methods.

2.3.14 Verbose

The `verbose` option lets you watch **imfil.m** at work. If you set `verbose = 1`, you will see the rows `histout` array appear on the screen as they are computed. The default is `verbose = 0`, which tells **imfil.m** to print only the most serious warnings on the screen.

This is a useful option when troubleshooting, as it is easy to see problems with the line search or stagnation when `verbose = 1`, and then stop the optimization in mid-stream to fix the problems.

Bibliography

- [1] C. AUDET AND J. E. DENNIS, *Mesh adaptive direct search algorithms for constrained optimization*, Tech. Rep. TR04-02, Department of Computational and Applied Mathematics, Rice University, 2004.
- [2] H. T. BANKS AND H. T. TRAN, *Mathematical and experimental modeling of physical processes*. Department of Mathematics, North Carolina State University, unpublished lecture notes for Mathematics 573-4, 1997.
- [3] J. T. BETTS, *Practical Methods for Nonlinear Control Using Nonlinear Programming*, no. 3 in *Advances in Design and Control*, SIAM, Philadelphia, 2000.
- [4] J. T. BETTS, M. J. CARTER, AND W. P. HUFFMAN, *Software for nonlinear optimization*, Tech. Rep. MEA-LR-083 R1, Mathematics and Engineering Analysis Library Report, Boeing Information and Support Services, June 6 1997.
- [5] R. BYRD, J. C. GILBERT, AND J. NOCEDAL, *A trust region method based on interior point techniques for nonlinear programming*, *Mathematical Programming A*, 89 (2000), pp. 149–185.
- [6] T. D. CHOI, O. J. ESLINGER, P. GILMORE, A. PATRICK, C. T. KELLEY, AND J. M. GABLONSKY, *IFFCO: Implicit Filtering for Constrained Optimization, Version 2*, Tech. Rep. CRSC-TR99-23, North Carolina State University, Center for Research in Scientific Computation, July 1999.
- [7] A. R. CONN, N. I. M. GOULD, AND P. L. TOINT, *LANCELOT: A Fortran Package for Large-Scale Nonlinear Optimization (Release A)*, no. 17 in *Springer Series in Computational Mathematics*, Springer Verlag, Heidelberg, Berlin, New York, 1992.
- [8] A. R. CONN, K. SCHEINBERG, AND L. N. VICENTE, *Introduction to derivative-free optimization*. unpublished manuscript, 2007.
- [9] J. E. DENNIS AND R. B. SCHNABEL, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, no. 16 in *Classics in Applied Mathematics*, SIAM, Philadelphia, 1996.

-
- [10] D. E. FINKEL AND C. T. KELLEY, *Convergence analysis of sampling methods for perturbed Lipschitz functions*. To appear in Pacific J. Opt., 2008.
- [11] R. FLETCHER, *Practical methods of optimization*, John Wiley and Sons, New York, 1987.
- [12] P. E. GILL, W. MURRAY, AND M. A. SAUNDERS, *SNOPT: An SQP algorithm for large-scale constrained optimization*, SIAM Review, 47 (2005), pp. 99–131.
- [13] P. E. GILL, W. MURRAY, AND M. H. WRIGHT, *Practical Optimization*, Academic Press, London, 1981.
- [14] P. D. HOUGH, T. G. KOLDA, AND V. J. TORCZON, *Asynchronous parallel pattern search for nonlinear optimization*, SIAM J. Sci. Comput., 23 (2001), pp. 134–156.
- [15] C. T. KELLEY, *Iterative Methods for Optimization*, no. 18 in Frontiers in Applied Mathematics, SIAM, Philadelphia, 1999.
- [16] ———, *Implicit filtering and nonlinear least squares problems*, in System Modeling and Optimization XX, E. W. Sachs and R. Tichatschke, eds., Dordrecht, 2003, Kluwer Academic Publishers, pp. 71–90.
- [17] ———, *Implicit Filtering*. Unfinished manuscript, 2008.
- [18] ———, *Users' Guide for imfl Version 0.7*. Software manual, work in progress, 2008.
- [19] T. G. KOLDA, R. M. LEWIS, AND V. J. TORCZON, *Optimization by direct search: New perspectives on some classical and modern methods*, SIAM Review, 45 (2003), pp. 385–482.
- [20] R. M. LEWIS AND V. TORCZON, *Rank ordering and positive bases in pattern search algorithms*, Tech. Rep. 96-71, Institute for Computer Applications in Science and Engineering, December 1996.
- [21] ———, *Pattern search algorithms for linearly constrained minimization*, SIAM J. Optim., 10 (2000), pp. 917–941.
- [22] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer, New York, 1999.
- [23] L. F. SHAMPINE AND M. W. REICHELDT, *The MATLAB ODE suite*, SIAM J. Sci. Comput., 18 (1997), pp. 1–22.
- [24] R. J. VANDERBEI, *LOQO: An interior point code for quadratic programming*, Optimization Methods and Software, 11 (1999), pp. 451–484.
- [25] T. A. WINSLOW, R. J. TREW, P. GILMORE, AND C. T. KELLEY, *Doping profiles for optimum class B performance of GaAs mesfet amplifiers*, in Proceedings IEEE/Cornell Conference on Advanced Concepts in High Speed Devices and Circuits, IEEE, 1991, pp. 188–197.

Index

bounds, 11
budget, 12
fscale, 7, 13
imfil_optset, 5, 7
options, 11
parallel, 12, 13
scaledepth, 5, 9, 14
scalestart, 9, 14
target, 15

Bound constrained optimization, 1
Budget, 2, 11

Feasible region, 1
function_error, 15

Hidden constraint, 2
Hyper-rectangle, 1

Initial Iterate, 11
Iteration
 Inner, 15
 inner, 8
 Outer, 15
 outer, 8

MULTICORE, 8

Nominal design space, 1
Nonlinear least squares, 4

Objective function, 1
ode15s, 4
Optimization landscape, 2
Options
 imfil_fscale, 13
 imfil_optset, 12
 scale_aware, 12, 14

vstencil, 15

Positive Basis
 stencil, 15

Sampling methods, 2
Scale, 2, 9, 14
Scale awareness of f , 9, 12
scale_aware, 9
Scaling, 2, 7, 13
Stencil failure, 8