On ODE estimation algorithms

M.R. Osborne

Mathematical Sciences Institute
Australian National University

CTAC08
ANU, July 13–16

Outline

Problem data

Differential equation:

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}\left(t, \mathbf{x}, \beta\right), \; \mathbf{x}, \mathbf{f} \in R^m, \; \beta \in R^p.$$

"Observed" data:

$$\mathbf{y}_i = \mathcal{O}\mathbf{x}^*\left(t_i, \beta^*\right) + \varepsilon_i, \; i = 1, 2, \cdots, n,$$
$$\mathcal{O} \in R^m \to R^k, \; \mathbf{y}_i \in R^k, \; k \leq m,$$
$$\varepsilon_i \in R^k, \; \sim N\left(0, \sigma^2 I_k\right), \text{ independent}$$
$$t_i \in [0, 1], \; i = 1, 2, \cdots, n.$$

Problem

It is required to estimate “true” vector of parameters β^* by

$$\widehat{\beta}_n = \arg\min_{\beta \in R^p} \Sigma_{i=1}^{n} \left\| \mathbf{y}_i - \mathcal{O}\mathbf{x}_i \right\|^2$$

where the $\mathbf{x}_i = \mathbf{x}\left(t_i, \beta\right)$, and the values of $\mathbf{x} = \mathbf{x}\left(t, \beta\right)$ are constrained by the differential equation.
Trapezoidal rule frequently a good enough discretization of ODE:

$$\mathbf{x}_{i+1} - \mathbf{x}_i = \frac{h}{2}\left(\mathbf{f}\left(t_{i+1}, \mathbf{x}_{i+1}, \beta\right) + \mathbf{f}\left(t_i, \mathbf{x}_i, \beta\right)\right), \; i = 1, 2, \cdots, n-1.$$

This has the form go Sim

$$\mathbf{c}\left(\mathbf{x}_c\right)_i = \mathbf{c}_{ii}\left(\mathbf{x}_i\right) + \mathbf{c}_{i(i+1)}\left(\mathbf{x}_{i+1}\right) = 0,$$

where $\left(\mathbf{x}_c\right)_i = \mathbf{x}_i, \; i = 1, 2, \cdots, n, \; \mathbf{x}_c \in R^{n \times m}$. Sparsity!

Embedding

The embedding approach leads to an unconstrained optimization problem which can be solved by standard methods e.g. Gauss-Newton. It removes the differential equation constraint on the state variable $\mathbf{x}(t, \beta)$ by embedding the differential equation into a parametrised family of boundary value problems which is solved explicitly at each step to generate trial values. It imposes boundary conditions:

$$B_1\mathbf{x}(0) + B_2\mathbf{x}(1) = \mathbf{b},$$

where $B_1, B_2 \in R^m \rightarrow R^m$ are assumed known while $\mathbf{b}$ is a vector of additional parameters which is sought as part of the estimation process. The key requirement is that the resulting system has a numerically well determined solution $\mathbf{x}(t, \beta, \mathbf{b})$ for all $\beta, \mathbf{b}$ in a large enough neighborhood of $\beta^*, \mathbf{b}^* = B_1\mathbf{x}^*(0) + B_2\mathbf{x}^*(1)$.

Selection of $B_1,\ B_2$

Calculation of $\mathbf{x}\left(t, \beta, \mathbf{b}\right)$ and $\nabla_\beta \mathbf{x},\ \nabla_b \mathbf{x}$ are required at each Gauss-Newton step. Typical is sequence of linear problems

$$\frac{d}{dt}\nabla_b \mathbf{x} - \nabla_x \mathbf{f} \nabla_b \mathbf{x} = 0,$$
$$B_1 \nabla_b \mathbf{x}\left(0\right) + B_2 \nabla_b \mathbf{x}\left(1\right) = I$$

These linear systems have matrix

$$F = \begin{bmatrix} C_{11} & C_{12} & & & & \\ & C_{22} & C_{23} & & & \\ & & & \ddots & & \\ & & & & C_{(n-1)(n-1)} & C_{(n-1)n} \\ B_1 & & & & & B_2 \end{bmatrix}$$

where $C_{ij} = \nabla_x \mathbf{c}_{ij}$.

Idea: Choose $B_1,\ B_2$ so this matrix is well conditioned at $\mathbf{x}^*\left(t, \beta^*\right)$.

Computation

Begin by permuting the first block column of F to the last position. A transformation of the first $n-1$ block rows of the permuted matrix to upper triangular form by orthogonal S yields

$$S^T FP = \begin{bmatrix} R_{11} & R_{12} & 0 & \cdots & 0 & R_{1n} \\ & R_{22} & R_{23} & \cdots & 0 & R_{2n} \\ & & & \cdots & \vdots & \vdots \\ & & & & R_{(n-1)(n-1)} & R_{(n-1)n} \\ & & & & B_2 & B_1 \end{bmatrix}.$$

The last two block rows now determine the conditioning of the transformed matrix. Make a second orthogonal factorization

$$\begin{bmatrix} R_{(n-1)(n-1)} & R_{(n-1)n} \end{bmatrix} = \begin{bmatrix} U^T & 0 \end{bmatrix} \begin{bmatrix} Q_1^T \\ Q_2^T \end{bmatrix}.$$

then $\begin{bmatrix} B_2 & B_1 \end{bmatrix} = Q_2^T$ provides the desired conditions.

Advantages

- Given initial $\mathbf{x}_c^0$ can compute estimate of optimal boundary conditions.
- Readily adapted to make use of standard Gauss-Newton nonlinear least squares and BVP software.
- Availability of good BVP software important if ODE is difficult.

Disadvantages:

- Good initial conditions are important. What happens if $1 - \left\| Q_2^T \left(\mathbf{x}_c^0\right)^T Q_2\left(\mathbf{x}_c^*\right) \right\|$ is close to 1?
- The economics of solving a nonlinear BVP for every function evaluation needs watching.
- The extra parameters $\mathbf{b}$ are not directly relevant to the problem formulation.

Simultaneous

Constrained least squares formulation. Let

$$\mathbf{r}_i = \mathbf{y}_i - \mathcal{O}\mathbf{x}_i,$$

$$\Phi\left(\mathbf{x}_c\right) = \frac{1}{2n}\sum_{i=1}^{n}\left\|\mathbf{r}_i\right\|^2.$$

Then the simultaneous method formulates the estimation problem as the constrained nonlinear least squares problem

$$\widehat{\boldsymbol{\beta}}_n = \arg\min_{\boldsymbol{\beta}} \Phi\left(\mathbf{x}_c\right);\ \mathbf{c}\left(\mathbf{x}_c, \boldsymbol{\beta}\right) = 0.$$

go Pro Standard methods of sequential quadratic programming are available. Important to exploit sparsity. Note the number of constraints increases as the discretization of the ODE is refined.

Computation

Introduce Lagrangian

$$\mathcal{L}\left(\mathbf{x}_c, \beta\right) = \Phi\left(\mathbf{x}_c\right) + \sum_{i=1}^{n-1} \boldsymbol{\lambda}_i^T \mathbf{c}_i\left(\mathbf{x}_c, \beta\right).$$

Necessary conditions give $\nabla_{(x,\lambda)}\mathcal{L} = 0$. Corresponding Newton iteration is

$$\begin{bmatrix} \nabla_{xx}^2\mathcal{L} & C^T \\ C & 0 \end{bmatrix} \begin{bmatrix} \triangle\mathbf{x} \\ \triangle\boldsymbol{\lambda} \end{bmatrix} = -\begin{bmatrix} \nabla_x\mathcal{L}^T \\ \mathbf{c}_c \end{bmatrix},$$

where $C = \nabla_x \mathbf{c}_c \in R^{nm} \to R^{(n-1)m}$.

Advantages

- Completely specified given initial estimates.
- Economy - avoids the inner iterations needed to solve BVP in embedding method.

Disadvantages:

- The number of constraints grows without bound as the discretization is refined.
- So does the number of constraint second derivatives that must be computed in a Newton iteration.
- It is problematic to formulate such solution strategies as mesh refinement as the state variables are only known exactly at the solution.

Equivalence

Let $S_E(\mathbf{x}_E)$ and $S_S(\mathbf{x}_S)$ be minimum sums of squares of residuals achieved at single isolated points in a large enough ball $B\left(\mathbf{x}^*, \rho(n)\right)$ by the embedding and simultaneous methods respectively. Then $\mathbf{x}_E$ satisfies the constraints on the simultaneous method so that

$$S_S\left(\mathbf{x}_S\right) \leq S_S\left(\mathbf{x}_E\right) = S_E\left(\mathbf{x}_E\right),$$

but direct substitution gives

$$B_1\mathbf{x}_S(0) + B_2\mathbf{x}_S(1) = \mathbf{b}_S.$$

Thus $\mathbf{x}_S$ can generates comparison quantities for the embedding problem. It follows that

$$S_E\left(\mathbf{x}_E\right) \leq S_E\left(\mathbf{x}_S\right) = S_S\left(\mathbf{x}_S\right).$$

Thus $\mathbf{x}_E = \mathbf{x}_S$ provided the solutions are isolated and $B\left(\mathbf{x}^*, \rho(n)\right)$ is large enough.

Better results

The problem with this derivation of equivalence is the condition on the size of $\rho(n)$. More satisfactory would be results of the kind:

- Satisfaction of necessary conditions for either the embedding or simultaneous methods could be deduced from satisfaction of the other. It would follow that both problems have the same stationary points.
- If choice $\rho(n) \to 0$ with $n^{-\alpha}$, $\alpha < 1/2$ given consistency:

$$\mathbf{x}_E \underset{n\to\infty}{\overset{a.s.}{\longrightarrow}} \mathbf{x}^*, \ \mathbf{x}_S \underset{n\to\infty}{\overset{a.s.}{\longrightarrow}} \mathbf{x}^*,$$

  then possibility of deducing identity for large enough n by the preceding argument. Not always possible. A direct proof of consistency for the simultaneous method is lacking.

Consistency of maximum likelihood estimates

If BVP is solved exactly then ODE estimation by the embedding method becomes a conventional maximum likelihood estimation problem.

$$\widehat{\boldsymbol{\beta}}_n = \arg\max_{\boldsymbol{\beta}} \mathbf{L}_n(\mathbf{y}, \boldsymbol{\beta}) = \arg\max_{\boldsymbol{\beta}} \sum_{i=1}^{n} L(\mathbf{y}_i, \mathbf{x}(t_i, \boldsymbol{\beta})).$$

Assume t_i equispaced, then

$$\frac{1}{n}\nabla_{\boldsymbol{\beta}} \mathbf{L}(\mathbf{y}, \boldsymbol{\beta}) \underset{n\to\infty}{\overset{a.s.}{\to}} \int_0^1 \mathcal{E}^* \{\nabla_{\boldsymbol{\beta}} L(\mathbf{y}, \mathbf{x}(t, \boldsymbol{\beta}))\}\, dt.$$

This gives a limiting form of the necessary conditions and it follows from a standard identity that $\boldsymbol{\beta} = \boldsymbol{\beta}^*$ is a solution. To show $\widehat{\boldsymbol{\beta}}_n \underset{n\to\infty}{\overset{a.s.}{\to}} \boldsymbol{\beta}^*$ can use the Kantorovich form of Newton's method. The idea is to apply this to $\nabla_{\boldsymbol{\beta}} \mathbf{L}_n(\mathbf{y}, \boldsymbol{\beta}) = 0$ starting from $\boldsymbol{\beta}^*$. Use Theorem to deduce from small residuals that $\widehat{\boldsymbol{\beta}}_n$ is close to $\boldsymbol{\beta}^*$.

Kantorovich Theorem

Let $\mathcal{J}_n = \frac{1}{n}\nabla^2_{\beta\beta}\mathbf{L}$. If the following conditions are satisfied in a ball $S_\varrho = \{\beta; \|\beta - \beta_0\| < \varrho\}$:

(i) $\|\mathcal{J}_n(\mathbf{u}) - \mathcal{J}_n(\mathbf{v})\| \leq K_1 \|\mathbf{u} - \mathbf{v}\|, \ \forall \mathbf{u}, \mathbf{v} \in S_\varrho,$

(ii) $\left\|\mathcal{J}_n(\beta_0)^{-1}\right\| = K_2,$

(iii) $\left\|\mathcal{J}_n(\beta_0)^{-1}\frac{1}{n}\nabla_x \mathbf{L}_n(\mathbf{y};\beta_0)^T\right\| = K_3$, and

(iv) $\xi = K_1K_2K_3 < \frac{1}{2},$

then the Newton iteration converges to a point $\widehat{\beta} \in S_\varrho$ satisfying the estimating equation, and $\widehat{\beta}$ is the only root in S_ϱ. The step to the solution $\widehat{\mathbf{x}}$ is bounded by

$$\left\|\widehat{\beta} - \beta_0\right\|_2 < 2K_3 < \varrho.$$

Considering discretization error

The embedding consistency result can be extended to two important cases:

1. when each differential equation discretization grid $\mathbf{K}_n$ corresponds to the observation grid $\mathbf{T}_n$; and
2. when the discretization is made on a fixed grid $t_j \in \mathbf{K}$ independent of $\mathbf{T}_n,\ n \to \infty$.

In the first case the condition that the differential equation be integrated satisfactorily ensures that the maximum mesh spacing $\Delta t \to 0,\ n \to \infty$. In the second case Δt is fixed and finite. This means that truncation error effects persist in the solution of the discretized problem as $|T_n| \to \infty$.

Types of results

1: If $\Delta t \rightarrow 0$ consistency follows using a similar argument. The idea is to start the iteration for each n at the exact integration solution $\widehat{\beta}_n$, $\widehat{\mathbf{b}}_n$ and use knowledge of the discretization error to show $K_3 = O\left(\Delta t\right)^2$ so this start is close to $\beta^n_\Delta, \mathbf{b}^n_\Delta$. Consistency now follows from the consistency for exact integration.
2: If Δt fixed, small enough, then best result possible starting from the corresponding exact integration solution is

$$\lim_{n\rightarrow\infty}\begin{bmatrix}\beta^n_\Delta\\ \mathbf{b}^n_\Delta\end{bmatrix}\subset S\left(\begin{bmatrix}\beta^*\\ \mathbf{b}^*\end{bmatrix}, O\left(\Delta t^2\right)\right),\ n\rightarrow\infty.$$

It uses $K_3 = O\left(\Delta t\right)^2$, $\forall n$large enough. Now n is number of observations.

Convergence rate results

The Gauss-Newton method for nonlinear least squares is typically the method of choice in the embedding method. A key result is that the convergence rate approaches second order asymptotically if the discretization error tends to zero as $|T_n| \rightarrow \infty$. If Δt fixed, small enough the rate reduces to fast first order.

The Bock iteration is the method of choice in the simultaneous method. Here the Newton iteration is modified by setting the constraint second derivatives to 0. This iteration has similar convergence rate behaviour provided the error terms are normally distributed. This is a stronger condition than that on the embedding method which requires only that the errors be independent and have bounded variance.