Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Reducing the Effect of Outlying Data on the Identification of Insulinaemic Pharmacokinetic Parameters with an Adapted Gauss-Newton Approach

P.D. Docherty, R.A.L. Gray, and E.J. Mansell

Department of Mechanical Engineering, University of Canterbury, Christchurch, Private Bag 4800, New Zealand
(Tel: +64 3 3642987 ex7211; e-mail: paul.docherty@canterbury.ac.nz).

Most parameter identification procedures find the parameter values that best fit some observed behavior according to the least squares criterion. However, the least squares criterion can be heavily influenced by outlying data or un-modelled effects and can thus yield poor results. Outlying data is often removed to avoid inaccurate outcomes. However, this process is complex and tedious.

This research presents an adaptation of the Gauss-Newton parameter identification method that effectively ignores the contribution from outlying data. The adapted method was compared to the original Gauss-Newton method in two modeling exercises. The first exercise was a C-peptide pharmaco-kinetic model with noisy data that did not contain outliers. The second exercise is an insulin pharmacokinetic model with data that contained outliers and un-modelled behavior.

The adapted method yielded similar results to the original methods for the C-peptide data with high correlations between identified parameters across approaches (R=0.90 and R=0.92). This was expected due to the relative lack of outliers in the C-peptide data. In contrast, the high rate of outliers and un-modelled behavior in the insulin data caused the significant differences between the parameter values identified by the approaches (R=0.47 and R=0.00). While the original method consistently found the least squares optima, the adapted method located the parameter set that fitted the majority of data points. The subtle difference between these approaches can yield large difference in identified parameter values. The adapted approach exhibits no erroneous effects in typical data, and should be considered more applicable in modeling exercises with outlier data or un-modelled effects.

Keywords: Parameter Identification, Pharmaco-kinetics, Insulin sensitivity and kinetics, Gradient descent, Outlier data

1. INTRODUCTION

Parameter identification methods are used to manipulate the values of certain model parameters such that the model accurately captures some observed behaviour (Carson and Cobelli). Most parameter identification methods optimise model parameters by minimising the least squares objective (or penalty) function (Bard 1970; Davidon 1991; Docherty *et al.* 2012; Levenberg 1944; Marquardt 1963; Steihaug1983). This means that a particular data point that is a certain distance from the modelled behaviour will have four times the influence on the objective function than a similar data point that is half the distance from the modelled behaviour.

While this approach works well in most data sets, it often leads to inaccurate parameter values when outlying data is present. In particular, outliers cause the optimal parameter set in terms of the least squares criteria to diverge significantly from the optimal parameter set defined by the 'inlying' data points. In such cases, the typical approach is to perform the inverse problem over a number of observations, determine the variance of the residuals, and then declare any points that fall outside three standard deviations from the simulated behaviour to be outliers and omit them from a second inverse problem. This process is tedious, and can lead to ambiguous outcomes or diminished operator independence.

We have previously presented an adaptation of the Gauss-Newton gradient-descent parameter identification method that reduces the contribution of outliers to the inverse problem (Gray *et al.* 2013). The method was presented using *in silico* derived data with induced outliers. In this current analysis we test the approach in model-based analyses of C-peptide kinetics and insulin kinetics. While the C-peptide data is known to be relatively free from outliers, the insulin data is known to contain both random outliers and un-modelled behaviour. Both data sets contain approximately 5% normally distributed noise.

2. METHODS

2.1 Clinical Protocol

The data used in this analysis was gathered during a dietary intervention study that measured the effect of dietary fibre in females that were deemed at risk of developing type 2 diabetes. The outcomes of the trial were presented by TeMorenga et al. (TeMorenga *et al.* 2010). Eighty-three individuals underwent the DISST at week 0, week 12 and

week 24 of the intervention. Some participants were lost to follow-up and a total of 218 DISST tests were undertaken.

Participants fasted from 10pm the night before the test and attended the clinic in the morning. Participants sat in a relaxed position for the duration of the test and had a cannula placed in their antecubital-fossa. This was used to both administer glucose and insulin boluses, and draw blood samples. A 10 g glucose bolus (50% dextrose) was administered at t = 6 minutes and an insulin bolus (actrapid) was administered at t = 16 minutes. Blood samples were taken at t = 0, 5, 10, 15, 20, 25, 30, 35, 40, and 50 minutes. The glucose levels were measured at the bedside (Enzymatic glucose hexokinase assay, Abbot Labs, Illinois USA) and the samples were then spun and frozen for batch assays of insulin and C-peptide (ELISA Immunoassay, Roche, Germany).

2.2 DISST model

The DISST model defines the behaviour of glucose, insulin and C-peptide kinetics (Lotz *et al.* 2010). However, only the insulin and C-peptide models are required for this analysis. The models are defined:

$$\dot{C} = k_2 Y - (k_1 + k_3)C + U_N \tag{1}$$

$$\dot{Y} = k_1 C - k_2 Y \tag{2}$$

$$U_N = U_B(t) + U_1(t) + U_2(t) \tag{1a}$$

$$U_B(t) = \begin{cases} k_3 C_0, & 0 \le t < 5 \\ 0, & 5 \le t \end{cases} \tag{1b}$$

$$U_1(t) = \begin{cases} U_1, & 5 \le t < 6 \\ 0, & t \le 5, 6 \le t \end{cases} \tag{1c}$$

$$U_2(t) = \begin{cases} U_2(1.5 - (t-6)/45), & 6 \le t \\ 0, & t < 6 \end{cases} \tag{1d}$$

$$\dot{I} = \frac{n_I}{V_P} Q - \left(n_T + \frac{n_I}{V_P} \right) I + x_L \xi U_N + \frac{U_X}{V_P} \tag{3}$$

$$\dot{Q} = \frac{n_I}{V_Q} I - \left(n_C + \frac{n_I}{V_Q} \right) Q \tag{4}$$

where: U_N is the endogenous insulin production which comprises of the basal rate (U_B) the first phase secretion (U_1) and the second phase secretion (U_2) (pmol·L^{-1}·min^{-1}); C is the plasma C-peptide concentration (pmol·L^{-1}); Y is the interstitial C-peptide concentration (pmol·L^{-1}); I is the plasma insulin concentration (mU·L^{-1}); Q is the interstitial insulin concentration (mU·L^{-1}); V_P is the plasma insulin distribution volume (L); V_Q is the Interstitial insulin distribution volume (L); $k_{1\text{-}3}$ are the C-peptide kinetic parameters (min^{-1}); n_I is the plasma-interstitial diffusion rate (L·min^{-1}); n_T is the plasma insulin clearance rate (min^{-1}); n_C is the interstitial insulin degradation rate (min^{-1}); ξ is the pmol→mU conversion factor (1/6); and x_L is the fractional hepatic clearance.

2.3 Parameter identification methods

This analysis compares the outcomes of the adapted Gauss-Newton method with the original approach. The original application of the Gauss-Newton parameter identification method iterates toward the optimal parameter set ($\mathbf{x}_{opt}$) using the iterative process:

$$\mathbf{x}_{i+1} = \mathbf{x}_i - (\mathbf{J}^T\mathbf{J})^{-1}\mathbf{J}^T\boldsymbol{\psi} \tag{5}$$

where:

$$\mathbf{J} = \left[\frac{\delta \psi_j}{\delta x_h} \right] = \begin{bmatrix} \frac{\delta \psi_1}{\delta x_1} & \frac{\delta \psi_1}{\delta x_2} & \cdots & \frac{\delta \psi_1}{\delta x_n} \\ \frac{\delta \psi_2}{\delta x_1} & \frac{\delta \psi_2}{\delta x_2} & \cdots & \frac{\delta \psi_2}{\delta x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\delta \psi_m}{\delta x_1} & \frac{\delta \psi_m}{\delta x_2} & \cdots & \frac{\delta \psi_m}{\delta x_n} \end{bmatrix} \tag{5a}$$

$$\boldsymbol{\psi} = \left[\hat{\psi}_j \right] = \left[X(\mathbf{x}_i, t_j) - X_{M,j} \right] = \begin{bmatrix} X(\mathbf{x}_i, t_1) - X_{M,1} \\ X(\mathbf{x}_i, t_2) - X_{M,2} \\ \vdots \\ X(\mathbf{x}_i, t_m) - X_{M,m} \end{bmatrix} \tag{5b}$$

and $\mathbf{J}$ is the Jacobian, $\boldsymbol{\psi}$ is the residual vector, X is the measured property, j is the sample index (j = 1..m), k is the parameter index (h = 1..n); $X(\mathbf{x}_i,t_j)$ is the simulated value of X at $t = t_j$; and $X_{M,j}$ is the measured value of X at $t = t_j$.

The adapted method is intended to ignore the contribution from outlying data. The combined Jacobian terms $(\mathbf{J}^T\mathbf{J})^{-1}\mathbf{J}^T$) result in a matrix that defines the optimal direction for the reduction of residuals at each data point for in each parameter. This is multiplied by the residual matrix to determine the relative weighting that should be given to each data point. The adapted method works by modulating the effect of the residual matrix according to the degree of residuals. Hence, the adapted method reduces the effect of outliers on $\boldsymbol{\psi}$ in (5), but not the contribution of $\boldsymbol{\psi}$ on $\mathbf{J}$ (5a). This is done by substituting $\boldsymbol{\psi}$ with $\hat{\boldsymbol{\psi}}$:

$$\mathbf{x}_{i+1} = \mathbf{x}_i - (\mathbf{J}^T\mathbf{J})^{-1}\mathbf{J}^T\hat{\boldsymbol{\psi}} \tag{6}$$

where:

$$\hat{\boldsymbol{\psi}} = \left[\hat{\psi}_j \right] = \psi_j e^{\frac{-|\psi_j|}{\beta|\tilde{\boldsymbol{\psi}}|}}, \tag{6a}$$

$|\tilde{\boldsymbol{\psi}}|$ is the median of the absolute values of the residuals and β is a scaling factor that determines the width of the peak as a function of $|\tilde{\boldsymbol{\psi}}|$.

In this analysis, β = 2. This value provides maximal objective function contributions at ψ = ±2 standard deviations of the residual distribution as shown in Fig. 1. In contrast, the typical Gauss-Newton optimises the least-squares residuals. This effectively means that the objective contribution increases at ψ^2as ψ increases.

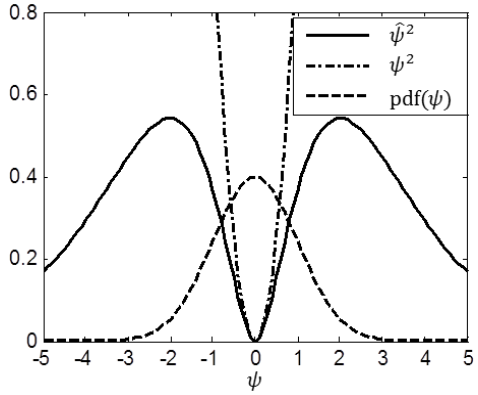

Fig. 1. Objective contributions from (5) and (6).

Initially, both Gauss Newton approaches were used to determine the contributions to U_N. The identified parameter set was $\mathbf{x}=[U_1, U_2]^T$. The remaining parameters from (1) and (2) were determined *a-priori* via the methods of Van Cauter *et al.* (Lotz *et al.* 2010; Van Cauter *et al.* 1992). This generated two U_N profiles and two sets of residuals for each DISST trial. The U_N profiles were then used in the identification of the insulinaemic pharmaco-kinetic parameters $\mathbf{x}=[n_T, x_L]^T$. The remaining parameters in (3) and (4) were determined *a-priori* via the methods of Lotz *et al.* (Lotz *et al.* 2010).

2.4 Evaluation

The two approaches were assessed via the nature of model residuals they produce. However, the adapted approach minimises $\|\hat{\boldsymbol{\psi}}\|_2$ and the typical approach minimises $\|\boldsymbol{\psi}\|_2$. Hence, a direct numerical comparison cannot be made. To highlight the differences in approach behaviour, summary statistics of the absolute ψ values and the residuals as a function of time ($\psi(t)$) will be presented.

3. RESULTS

Both approaches converged to the expected behaviour in each case tested. No parameter sets diverged, and there were no evident cases wherein the approaches led to an incorrect local minima. The endogenous insulin parameter values were relatively well correlated across identification methods (R=0.90 for U_1 and R=0.97 for U_2, respectively). However, the insulin pharmaco-kinetic parameters were not so well correlated (R=0.47 for n_T and R=0.00 for x_L). The correlation for x_L improved to R=0.44 when the Spearman correlation was used, thus implying some outlying data obscured the overall trends.

The median bias of U_1 between the typical and adapted method was -6.1% (IQR -31.2% to 1.4%). For U_2, the bias was 3.73% (IQR -1.73% to 10.5%). The bias in n_T was -23.8% (IQR -40% to -2.3%) and the bias in x_L was 16.1% (IQR -7.8% to 52.6%).

Summary statistics of the residual data are presented in Table 1. Note that the C-peptide model residuals were lower below the median residual (ψ_{50}) for the adapted approach. In contrast, the adapted approach yielded higher C-peptide residuals above the median value. The adapted approach yielded lower insulin residuals for all metrics shown, but exhibited higher residuals at the extreme upper limit (not shown).

Table 1. Summary statistics of model residuals

Model	Approach	Residuals (percentiles) $[\psi_{25}, \boldsymbol{\psi_{50}}, \psi_{75}, \psi_{95}, \psi_{99}]$
C-peptide	*Eqn. 6*	[9.39, **39.00**, 98.20, 273.91, 604.23]
[pmol.L^{-1}]	*Eqn. 5*	[12.16, **38.97**, 80.38, 184.89, 294.97]
Insulin	*Eqn. 6*	[1.11, **3.48**, 11.23, 107.48, 345.09]
[mU.L^{-1}]	*Eqn. 5*	[4.21, **12.60**, 30.91, 139.93, 419.22]

Fig. 3 shows the distribution of the residuals about the measured points. Note that the C-peptide residuals were relatively well centred about zero with a seemingly normal distribution. In contrast, the insulin residuals were sporadic, non-normal and show definite point dependence. Fig. 4 shows the least squares and the adapted objective surfaces for the responses shown in Fig. 2.

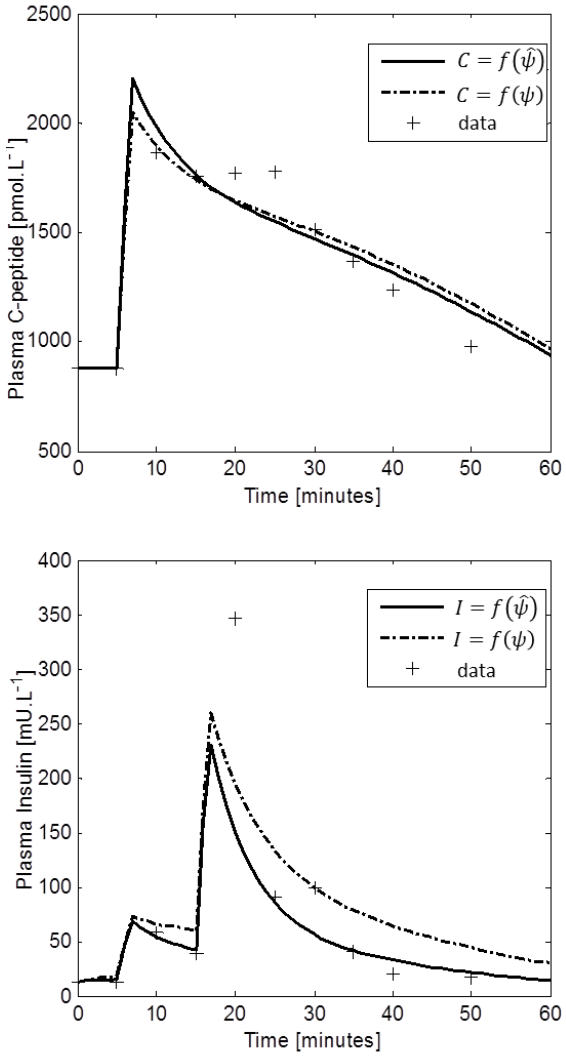

Fig. 2. Plasma C-peptide and insulin simulations for a typical patient response to the DISST test with conspicuous outliers in the insulin data.

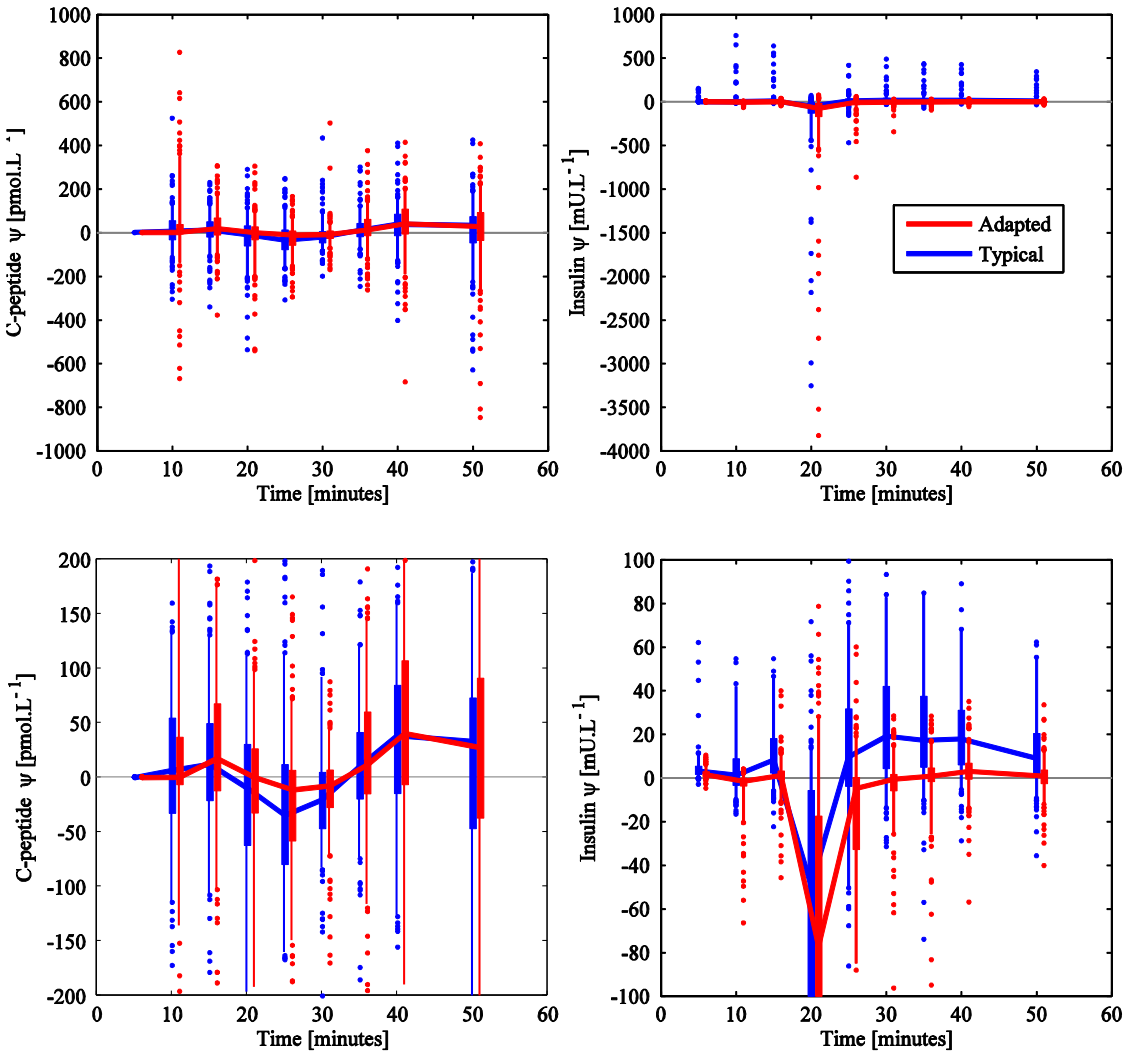

Fig. 3. Residual plots for C-peptide and insulin. The bottom plots are cropped to show the general behaviour. The thick errorbars show the interquartile range, thin errorbars show the 5th to 95th percentile range and the dots show the outlying points.

4. DISCUSSION

The proposed adaptation to the Gauss-Newton parameter identification method yielded different results to the typical approach. The adaptation allows the identified parameter set to diverge away from the mathematical optima defined by the least squares objective criterion and locate a minima that is most representative of the least-squares optima for the majority of data points. In cases wherein there was no outlying data, this makes virtually no difference (Fig. 2-top). However, when there are outlying data points, the outcomes defined by the original and the adapted approaches become distinct (Fig. 2-bottom, Fig. 4).

In noisy data that does not contain outliers, the adapted method provides no benefit, but also does not introduce any deleterious outcomes. When the data contains known outliers, the adaptation provides significant benefit. For example, in this study, the insulin data contained both outliers and un-modelled behaviour and the adaptation considerably changed the identified parameter values. In contrast, the C-peptide data, while noisy, did not contain significant un-modelled behaviour or outliers. Thus, the proposed method did not introduce much variation. These outcomes were observed in the correlations and relative bias seen in the parameters of the C-peptide model and the insulin model.

Fig. 3 shows that the original least-squares approach located the t = 20 minutes insulin data more accurately than the adapted approach. However, this data point is affected by incomplete mixing of insulin at the depot site, and is thus, an unmodelled phenomenon. Hence, the behaviour appeared as a consistent outlier in the measured data and the ideal parameter values that describe the overall patient behaviour would be ideally located without influence from this data. By ignoring this data point, the adapted method allowed greater adherence to the remaining insulin data. This can be seen in the comparatively unbiased insulin residuals for the t = 5, 10, 15, 25, 30, 35, 40, and 50 minutes data. The other point to note in Fig. 3 is the increased magnitude of outliers for the adapted method. This indicates that the adapted method effectively recognised and ignored these points in favour of capturing the behaviour defined by the remaining majority of data points.

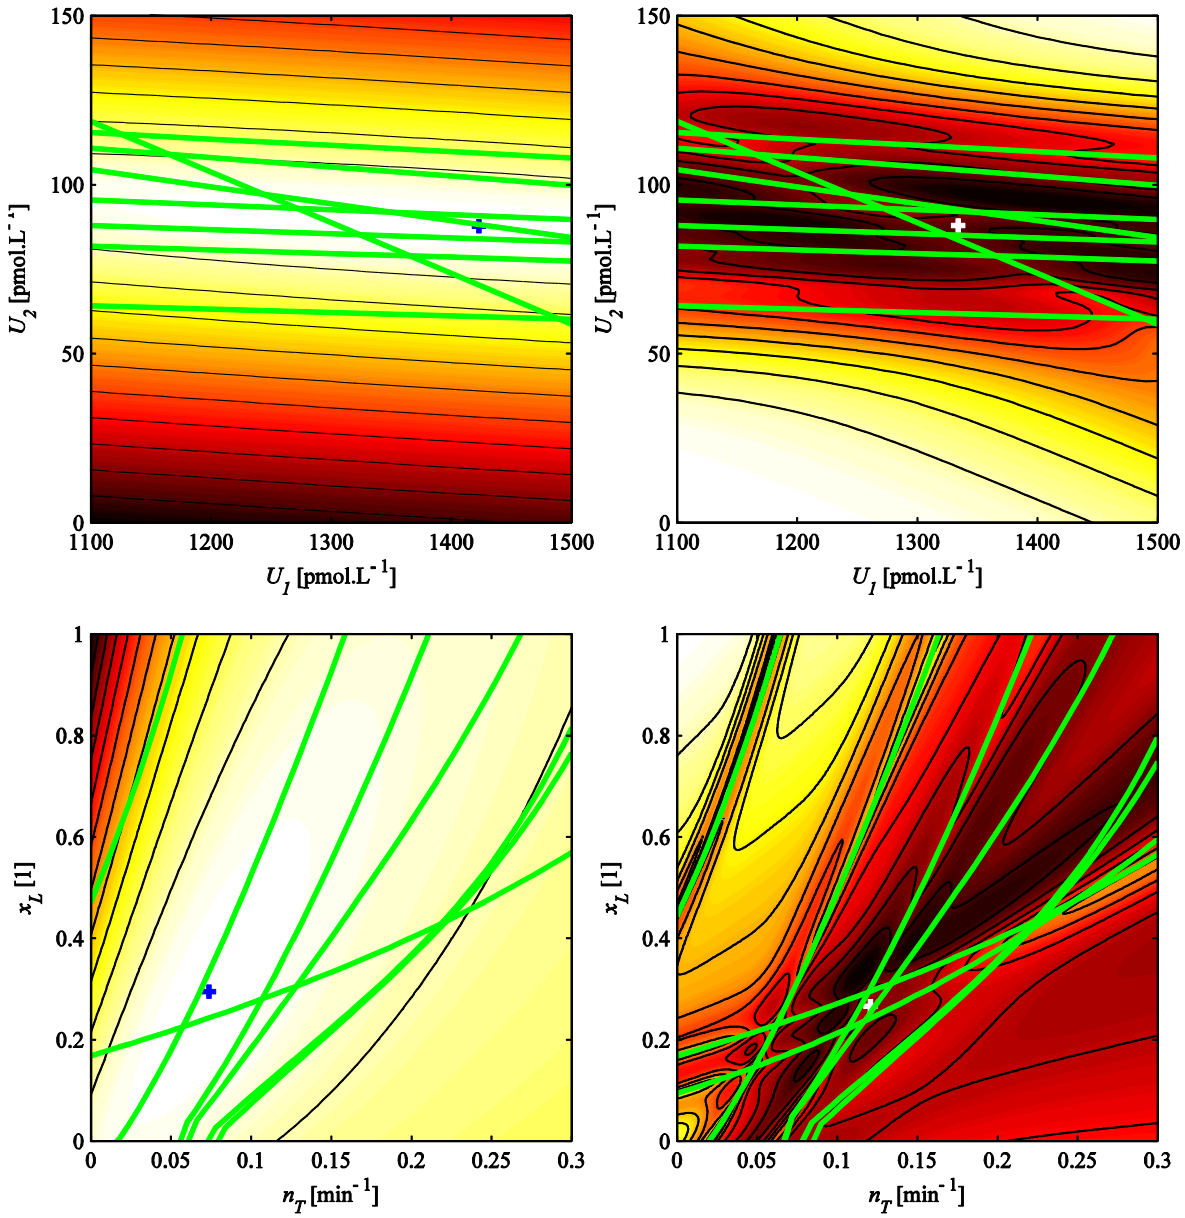

Fig. 4. Objective surfaces generated by the data shown in Fig. 2 with the residual vectors ($\boldsymbol{\psi}$ and $\hat{\boldsymbol{\psi}}$) of (5) and (6) for the typical least squares approach (left) and the adapted approach (right) for the C-peptide model (top) and the insulin model (bottom). The green lines show the parameter combinations that would achieve zero error for a particular data point. Note the different location of the parameter optimas (+) for the two approaches, and that the parameter values determined by the adapted method do not appear to be located on a minima of the objective surface.

The C-peptide residuals indicate very little systemic difference between the two approaches. This is because there is a relative lack of outliers in the C-peptide data compared to the insulin data. However, there were a few outliers and, similarly to the insulin case, the adapted method ignored these points in favour of capturing the behaviour defined by the other data points. Thus, the outlier residuals appeared further from the zero residual line for the adapted method. Table 1 shows that the adapted method yielded lower C-peptide residuals below the median residual and higher residuals above the median. This further indicates that the method gave more influence to data points close to the simulated profile, and limited the effect of those that were distant from the simulation.

The adapted method works by modulating the relative contribution of the residuals to the iteration direction and magnitude during convergence. It does not modulate the Jacobian which determines the direction that convergence should step in with respect to each residual. Hence, the objective surfaces shown in Fig. 4 right are not the surfaces used to determine the ideal direction of descent. The least-squares surfaces on the left are used by the adapted method in determining the direction of descent (the $(\mathbf{J}^T\mathbf{J})^{-1}\mathbf{J}^T$ term). Thus, it is not necessarily a negative or altogether unexpected outcome that the located optima from the adapted method does not occur at a minima on the objective surface. The insulinaemic parameter optima located by the adapted approach is located close to the intersections of the zero error

lines of 4 data points and takes influence from 3 others (Fig. 4 bottom-right). This is consistent with Fig. 2 that shows the simulation fitting closely to seven of the nine data points. The simulated model responses shown in Fig. 2 validate the overall approach and show that the behaviour determined by the adapted approach was closer to the majority of data points than the typical least squares optimisation.

Setting the contribution of the outliers as a function of the median residual allows the method to progressively reduce the influence of outliers as the method converges to a solution. In doing so, the importance of the initial parameter estimate is reduced. If this were not the case, and initial conditions were chosen such that only a small portion of the data were fitted effectively by the simulation, it would not be possible for the adapted method to converge toward the majority of the data. However, modulating the influence of the outliers by the median residual ensures that the method is robust to the choice of initial parameter estimates. In particular, The contribution from outliers is only limited when the parameter set converges close to the optima as defined by the majority of data points.

The adapted method can be modulated by changing the value of β in (6a). β has the role within the approach of altering how much of the measured data is captured within the objective contribution defined by (6a). Hence, a greater β value would allow a greater influence from outlying data. Conversely, a smaller value would allow a lesser influence from outlier data. Setting β = infinity yields the same behaviour as the original least squares approach. Fig. 1 assumes that the residuals are normally distributed about the model simulation, which is not the case when outliers are present. However, as the adapted method is intended to ignore outliers, this apparent contradiction of the assumptions that drive the method actually yields the intended outcome.

The method could be relatively simply applied to parameter identification methodologies as the additional term in (6) is not difficult to evaluate. Furthermore, it is often known *a-priori* whether data is likely to contain outliers. In particular, prior to performing the inverse problem, the raw data is often plotted for a few cases. It is often then apparent whether the data contains outliers. Hence, whether there is a need for the removal of outlying data via the adapted method could be known prior to a full run of the inverse problem and an analysis of the residuals yielded.

The proposed method eliminates the need for manual removal of outlier data. Removal of outlier data is typically a tedious task that is most often undertaken by manual location of data points that fall more than three standard deviations of residual data from the simulated model and removing the points for a second run of inverse problem. However this process is very costly in terms of operator time. It also runs the risk of reducing operator independence as discretion may be exercised when selecting or applying the rules regarding outlier omission.

In this analysis we have tested the adapted method for both noisy data, that includes outliers and un-modelled effects and data that was only noisy. Thus, the potential benefit of the method and the limit of the potential benefit could be determined. It was shown that the adapted method captures the observed behaviour better than the original method in data that contains un-modelled effects or outlying data, and provides no hindrance when data is simply noisy.

5. CONCLUSIONS

We have presented a method for the identification of model parameters that automatically modulates the objective surface such that outlier data is ignored. The method enables a relatively simple amendment to a well-known and understood parameter identification approach that is operator independent and can be tailored to various and distinct modelling situations.

REFERENCES

Bard Y. (1970) Comparison of Gradient Methods for the Solution of Nonlinear Parameter Estimation Problems, *SIAM Journal on Numerical Analysis,* 7(1), 157-186

Carson E.R., Cobelli C.: *Modelling methodology for physiology and medicine*. San Diego, Academic Press, 2001

Davidon W. (1991) Variable Metric Method for Minimization, *SIAM Journal on Optimization,* 1(1), 1-17

Docherty P.D., Chase J.G., David T. (2012) Characterisation of the iterative integral parameter identification method, *Medical and Biological Engineering and Computing,* 50(2), 127-134

Gray R.A.L., Docherty P.D., Fisk L.M. (2013) A modified approach to objective surface generation within the Gauss-Newton parameter identification to ignore outlier data points., *Med. Biol. Eng. Comput.,* In review(

Levenberg K. (1944) A method for the solution of certain non-linear problems in least squares, *Quarterly of Applied mathematics,* 2(164-168

Lotz T.F., Chase J.G., McAuley K.A., Shaw G.M., Docherty P.D., Berkeley J.E., Williams S.M., Hann C.E., Mann J.I. (2010) Design and clinical pilot testing of the model based Dynamic Insulin Sensitivity and Secretion Test (DISST), *J Diabetes Sci Technol,* 4(6), 1195-1201

Marquardt D.W. (1963) An Algorithm for Least-Squares Estimation of Nonlinear Parameters, *SIAM Journal on Applied Mathematics,* 11(2), 431-441

Steihaug T. (1983) The Conjugate Gradient Method and Trust Regions in Large Scale Optimization, *SIAM Journal on Numerical Analysis,* 20(3), 626-637

TeMorenga L., Williams S.M., Brown R., Mann J.I. (2010) Effect of a relatively high protein, high fiber diet on body composition and metabolic risk factors in overweight women, *Eur J Clin Nutr,* 64(11), 1323-1331

Van Cauter E., Mestrez F., Sturis J., Polonsky K.S. (1992) Estimation of insulin secretion rates from C-peptide levels. Comparison of individual and standard kinetic parameters for C-peptide clearance, *Diabetes,* 41(3), 368-377