Medical Image Analysis 35 (2017) 116–132

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Federated optimisation of kinetic analysis problems

Nicholas Dowson[a,*], Charles Baker[a,c], Paul Thomas[b,c], Jye Smith[b], Simon Puttick[d], Christopher Bell[a,e], Olivier Salvado[a,e], Stephen Rose[a,c]

[a] CSIRO, The Australian eHealth Research Centre, Level 5 UQ Health Sciences Building, Royal Brisbane and Women's Hospital, Herston, Queensland, 4029, Australia
[b] Specialised PET Services Queensland, Royal Brisbane and Women's Hospital, Herston, Brisbane, Australia
[c] School of Medicine, University of Queensland, St Lucia, Brisbane, Australia
[d] Australian Institute for Bioengineering and Nanotechnology, University of Queensland, St Lucia, Brisbane, Australia
[e] School of Information Technology and Electrical Engineering, University of Queensland, St Lucia, Brisbane, Australia

ARTICLE INFO

Article history:
Received 16 July 2015
Revised 8 June 2016
Accepted 15 June 2016
Available online 16 June 2016

Keywords:
Kinetic analysis
Optimisation
Dynamic images
Positron emission tomography
Magnetic resonance imaging

ABSTRACT

Positron Emission Tomography (PET) data is intrinsically dynamic, and kinetic analysis of dynamic PET data can substantially augment the information provided by static PET reconstructions. Yet despite the insights into disease that kinetic analysis offers, it is not used clinically and seldom used in research beyond the preclinical stage. The utility of PET kinetic analysis is hampered by several factors including spatial inconsistency within regions of homogeneous tissue and relative computational expense when fitting complex models to individual voxels. Even with sophisticated algorithms inconsistencies can arise because local optima frequently have narrow basins of convergence, are surrounded by relatively flat (uninformative) regions, have relatively low-gradient valley floors, or combinations thereof. Based on the observation that cost functions for individual voxels frequently bear some resemblance to each-other, this paper proposes the federated optimisation of the individual kinetic analysis problems within a given image. This approach shares parameters proposed during optimisation with other, similar voxels. Federated optimisation exploits the redundancy typical of large medical images to improve the optimisation residuals, computational efficiency and, to a limited extent, image consistency. This is achieved without restricting the formulation of the kinetic model, resorting to an explicit regularisation parameter, or limiting the resolution at which parameters are computed.

1. Introduction

This work is concerned with the development of a more robust approach for analysing dynamic medical images that is agnostic to model and sufficiently efficient for practical use. The approach developed here obtains rapid and robust regressions of non-linear models despite the relatively high noise typical of certain images, particularly Positron Emission Tomography (PET) images. This is achieved by exploiting the substantial redundancy within dynamic medical images, which frequently contain hundreds of thousands of time varying voxel intensities.

The analysis of dynamic Positron Emission Tomography (PET) images, and Magnetic Resonance Images (MRI), is useful for gaining an understanding of biological phenomena such as cancer. For PET, where acquisition must be performed over a period of time, the raw data is intrinsically dynamic, *i.e.* the dynamic data is obtained without altering the imaging protocol by extending scan time or otherwise. Using *kinetic analysis* techniques, dynamic data can be converted into a form that is more amenable to biological interpretation for clinicians. This is achieved by specifying a model that encodes the movement of a biological tracer through a set of states, which may be considered as metabolic states or anatomical locations (Carson, 2005; Cherry et al., 2003; Schmidt and Turkheimer, 2002; Watabe et al., 2006). The free parameters in the model can be fitted to the data using standard non-linear optimisation techniques (Levenberg, 1944; Marquardt, 1963), albeit with some customisation to clamp values to a physiologically feasible range, *e.g.* by enforcing positivity.

Yet despite the availability of dynamic data and apparent ease with which it may be analysed, dynamic information is not generally used in the clinic, and even in research it is seldom used beyond the preclinical stage. A reason for the limited usage of dynamic data could be that the consistency required for reliable interpretation of processed data is not currently attained. The high noise typical of PET images can result in relatively unstable

* Corresponding author.
E-mail address: nicholas.dowson@csiro.au (N. Dowson).

http://dx.doi.org/10.1016/j.media.2016.06.019
1361-8415/

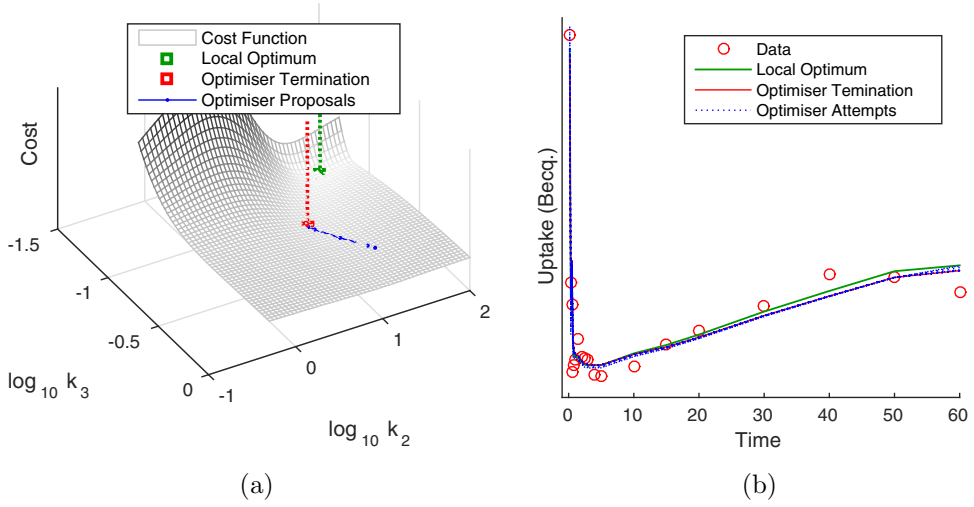

Fig. 1. (a) Cost function surface for example time activity curve after optimisation is complete, with point of termination shown (red square and pointer). Levenberg-Marquardt optimisation makes a series of suggested parameters based on local gradient (blue lines with dots), but misses the local optimum (green square), due the to local orientation of surface. **(b)** The time activity curve (red circles) has a number of similar quality solutions corresponding to points in (a), despite a relatively wide variation in parameters, which is indicative of the flatness of the cost function in many areas in (a). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

solutions, because local optima frequently have narrow basins of convergence, are surrounded by flat and uninformative regions, have low gradient valley floors or combinations thereof. An example of this is shown in Fig. 1. The net result is that many voxels have a basin of convergence that is difficult to find without sampling wide swathes of parameter-space, which can be too inefficient for practice.

1.1. Previous work

The literature contains many alternative methods that seek to robustly estimate kinetic parameters including graphical analysis, methods that rely on basis functions and direct fitting of the model via non-linear optimisation (Watabe et al., 2006).

Graphical analysis approaches, such as Gjedde-Patlak analysis (Gjedde, 1981; Patlak et al., 1983; Patlak and Blasberg, 1985) (for irreversible models), Logan analysis (Logan et al., 1990) (for reversible models), relative equilibrium approaches (Zhou et al., 2009) (to more robustly handle noise), or combinations thereof (Zhou et al., 2010), are typically rapid and robust as they involve performing linear regressions on selected time-varying functions derived from the tissues and blood. However, such graphical approaches require lumping multiple parameters that may be of interest individually (Ikoma et al., 2008). In certain applications *e.g.* modelling tracers for the first time (Huang et al., 1991), this can be overly restrictive.

A more flexible but more costly approach (computationally) is to partly linearise the model by integrating both sides of the model equation, such as applied to the Kety-Schmidt model for cerebral blood flow (Kety and Schmidt, 1948) in Huang et al. (1982). Feng et al. (1993) extended the method in an approach called Generalised Linear Least Squares (GLLS) to account for the increased bias imposed by fitting errors in earlier time-frames. This was extended in Chen et al. (1998) to multiple compartment models and to handle spillover into image derived blood input functions. These approaches are now sufficiently robust to noise that they have been applied to SPECT (Wen et al., 2009). The main challenge for GLLS and its brethren is the difficulty in enforcing parameter constraints (Zeng et al., 2012).

Blomqvist (1984) took the approach of evaluating individual integrals as functions of the non-linear parameters (the basis) and computing the linear combination of the basis functions to obtain the remaining parameters. van den Hoff et al. (1993) combined this with methods to deal with delay and signal dispersion which are non-commutative (Meyer, 1989). Boellaard et al. (2005) found the utilisation of basis functions to be more accurate than GLLS but more costly. Hong and Fryer (2010) reduce computational cost by replacing each integral with a closed form convolution of exponentials. Kadrmas and Oktay (2013) reformulate slightly to guarantee individual compartments are distinguishable and use the strategy of splitting the non-linear and linear portions of the problem and use non-negative least squares (Lawson et al., 1995) (NNLS) to enforce positivity. Smith et al. (2010) formulate the blood input function as a sum of exponential curves, for a rapidly computed analytic sum of convolved exponentials formulation, because the formulation allows new fits to be generated in closed form. This formulation also requires no branching making it a good candidate for GPU implementation which has many advantages in terms of speed.

Gunn et al. (2002) fully exploit the idea of treating TACs as linear combinations of a set of basis functions by utilising developments in sparse modelling (Chen et al., 2001). Gunn et al. decompose TACs into a sparse set of entries from a dictionary of basis vectors consisting of exponentials convolved with the blood input function. This approach combines the advantages of globally optimal solutions using quadratic programming techniques while allowing model complexity to vary appropriately with the data. Solution accuracy is (slightly) limited by dictionary resolution, but required dictionary resolution is independent of the dimension of the parameters.

Various approaches enable integration with image reconstruction, as proposed by Carson et al. (1986) for the Kety-Schmidt model, allowing the data to be fitted on a temporally continuous basis. Basis functions formulations allow more complex models to be incorporated into reconstruction (Reader et al., 2006), also using alternative basis selections, such as exponential splines (Verhaeghe et al., 2008). More recently, motion correction has been incorporated as well (Jiao et al., 2012; Pedemonte et al., 2011). Integration with image reconstruction requires list-mode data to be available, are relatively costly and can restrict the model formulation.

Of all the proposed approaches, it is acknowledged that non-linear least squares approaches generally retain the greatest