ARTICLE IN PRESS

Computer Physics Communications ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Computer Physics Communications

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

A Bayesian method for analysing relaxation spectra

L. Ciocci Brazzano [a,c,*], L.J. Pellizza [b,c], C.L. Matteo [a,c], P.A. Sorichetti [a]

[a] Facultad de Ingeniería, Universidad de Buenos Aires, Paseo Colón 850 CABA, Argentina

[b] Instituto Argentino de Radioastronomía (CONICET), C.C. 5, (1894), Villa Elisa, Pcia. de Buenos Aires, Argentina

[c] Consejo Nacional de Investigaciones Científicas y Técnicas, Argentina

ARTICLE INFO

Article history:
Received 11 March 2015
Received in revised form
22 July 2015
Accepted 25 August 2015
Available online xxxx

Keywords:
Dielectric spectroscopy
Mechanical spectroscopy
Spectral analysis
Bayesian methods
Monte Carlo

ABSTRACT

The knowledge of electrical and mechanical properties of material, relies on a precise analysis of the relaxation spectra. We explore the ability of a Bayesian method to achieve an accurate estimation of spectral parameters. We implemented a parallel-tempering Markov-chain Monte Carlo algorithm and used it to fit simulated and measured spectra. An exhaustive testing of the code shows that it presents an extremely good performance, accurately fitting complex spectra under strong noise and overlapping components. We conclude that this technique is quite suitable for relaxation spectra analysis, complementing classical methods.

1. Introduction

A precise knowledge of the electrical and mechanical properties of solids and liquids is of great importance for pure research, as well as for various applications. It is a key ingredient for the development of new materials, quality control and preventive maintenance. Mechanical [1] and electrical [2,3] relaxation spectroscopy techniques are widely used to characterise materials (see for example [4–8]).

It is well known that, as it follows from linear response theory, there is a complete analogy between the phenomenological descriptions of linear viscoelastic and dielectric relaxation processes (e.g. [9] and Chapter 1, Appendix 1.1 of [3]). In consequence, the analysis methods for dielectric spectra may be also used for the study of mechanical relaxations. In this work we will present a general algorithm for the analysis of relaxation spectra. Although the discussion of applications will focus on dielectric processes, the extension to mechanical spectra is straightforward.

Broadband dielectric spectroscopy allows the measurement of electrical properties of a material, through the analysis of the interaction between it and an electrical excitation field. In the macroscopic scale, this interaction is described by two parameters: the permittivity ε, associated with the relaxation of the molecular polarisability; and the conductivity σ, linked to free charge transportation. When a harmonic excitation field is applied, the complex permittivity spectra describe jointly the frequency dependence of both parameters [3],

$$\varepsilon(\omega, T) = \varepsilon'(\omega, T) - j\varepsilon''(\omega, T) - j\frac{\sigma}{\omega}\varepsilon_0. \quad (1)$$

For excitation frequencies lower than 10^{12} Hz, the electrical properties are dominated by relaxation processes. Relaxation spectra may be analysed through models that describe the dependence of the complex permittivity with frequency, by means of a small number of parameters that characterise the material. A precise fitting of these parameters is of great importance for material sciences, as well as for device design and applications.

At present, low cost and high quality measurements of the aforementioned properties are routinely performed both in the laboratory and in the field. Hence, the development of powerful tools for the analysis of experimental data is still of great importance to obtain reliable results. Spectral parameter estimation is usually done, both by commercial and laboratory-designed software, through fitting techniques based on the minimisation of some figure of merit applied separately to the real and imaginary parts of the spectra. One of the most popular techniques is the least squares (LS) method, which aims at minimising the distance between the measured data and the fitting model. Typical LS algorithms, such as the Levenberg–Marquardt scheme [10], implement this minimisation by starting from a user-provided initial guess for

* Corresponding author at:Facultad de Ingeniería, Universidad de Buenos Aires, Paseo Colón 850 CABA, Argentina. Tel.: +54 11 4343 0891.
E-mail address: bciocci@fi.uba.ar (L. Ciocci Brazzano).

http://dx.doi.org/10.1016/j.cpc.2015.08.033
0010-4655/

Please cite this article in press as: L. Ciocci Brazzano, et al., A Bayesian method for analysing relaxation spectra, Computer Physics Communications (2015), http://dx.doi.org/10.1016/j.cpc.2015.08.033

the parameter values, and stepping along the gradient of the figure of merit in the parameter space until they reach the minimum. This method is the usual choice in most cases, because it is simple and gives good results in typical problems. Moreover, it is available in most software packages. However, it is less efficient when confronted to more complex problems, such as spectra with many overlapping relaxations (see [11] for a comprehensive discussion) or in the presence of noise. Particularly, its convergence is not guaranteed for an arbitrary initial guess, as it is prone to get stuck in local minima, and the convergence speed for multiparameter models is slow. These disadvantages usually imply a lot of time and effort dedicated to the search of an adequate initial guess (which may be difficult to find when many relaxations are present), in order for the method to converge to the true, global minimum.

An alternative to classical LS schemes, the interval analysis (IA) method, was recently put forth by several authors [11–14]. This method is based on a sequential partition of the parameter space, which allows the establishment of lower and upper bounds to the best-fit parameters. Therefore, it provides an interval for each parameter, in which the best-fit value is guaranteed to lie. This method avoids many of the shortcomings of the LS. However, the IA method is mathematically complex, and relies on the key assumption of known bounds for the measurement uncertainties. Finding proper error bounds for the method to give sensible results might be a difficult task. This key assumption also makes difficult to compare the results of the IA method with either LS or any other that makes a different assumption for the statistics of the measurement uncertainties.

The approach to parameter estimation problems presented in this work has become popular in many scientific fields during the last decade. It relies on Bayesian statistical methods (e.g., [15,16] and references therein) to provide a powerful and accurate fitting technique for complex, multiparameter problems. The method, called Markov-chain Monte Carlo (MCMC), is based on the sampling of the posterior probability density function (PDF) of the parameters, given the data. The best-fit parameters are those that maximise this PDF. MCMC avoids by construction most of the problems of the LS method. It can be applied to models of arbitrary complexity, the only requirement being that the posterior PDF can be computed either analytically or numerically. It works even for non-deterministic (stochastic) models that include random components other than those arising from measurement uncertainties (e.g., [17]). It can deal naturally with parameter bounds and symmetries, and performs a fast sampling of the parameter space, spending more time in high probability regions and avoiding uninteresting, low probability ones. The convergence speed can be automatically controlled. A simple extension of the method, known as parallel tempering MCMC (PT-MCMC) allows also to avoid local maxima, ensuring the convergence for any reasonable initial guess. The quality of the fit is directly and statistically linked to the measurement accuracy. Indeed, the complete posterior PDF of the parameters is given as a result, which allows the estimation of their uncertainties (or confidence intervals) and correlations. Finally, MCMC provides a model comparison technique that naturally implements the Occam's razor rule.

Given these advantages, it is worth to explore the application of MCMC to the problem of parameter estimation in relaxation spectra. This work develops such an application, and is organised as follows. In Section 2 we describe an implementation of the PT-MCMC algorithm to fit Havriliak–Negami models to relaxation spectra. In Sections 3 and 4 we investigate the performance of the algorithm on simulated and real (measured) data, respectively. In Section 5 we present our conclusions.

2. A Bayesian method to analyse relaxation spectra

Relaxation spectra can be described as a sum of several relaxation processes taking place in the material. The most general model used to describe them is that of Havriliak–Negami (hereafter HN, [18]),

$$\varepsilon_{\mathrm{HN}} = \varepsilon_\infty + \sum_{i=1}^{M} \frac{\Delta\varepsilon_i}{(1+(j\omega\tau_i)^{\alpha_i})^{\beta_i}}. \tag{2}$$

Here M is the number of relaxation processes, and ε_∞ is the high frequency limit of the real permittivity. For each relaxation process i, τ_i is its characteristic relaxation time (related to the maximum loss frequency), $\Delta\varepsilon_i$ measures the relaxation strength, and α_i and β_i describe the broadening and the shape of the relaxation peak. We will denote these parameters collectively by $\vec{\theta}$. In this paper, we will focus on the problem of estimating their best values, assuming that the number of relaxation processes that describe the spectrum is fixed, and that each of them can be described by a HN model. The problem of determining the true number of relaxation processes and the model that best describes each of them (either Debye, HN, etc.), which are different aspects of a model comparison problem, will be treated elsewhere.

The Bayes theorem (e.g., [15,16,19]) for the problem of estimating the parameter values of a HN model may be written as

$$f_{\mathrm{post}}(\vec{\theta}|D,I) = \frac{f_{\mathrm{pri}}(\vec{\theta}|I) f_{\mathrm{L}}(D|\vec{\theta},I)}{\int f_{\mathrm{pri}}(\vec{\theta}|I) f_{\mathrm{L}}(D|\vec{\theta},I)\mathrm{d}^{4M+1}\theta}, \tag{3}$$

where f_{post} is the posterior PDF of the parameters, given the data D and any previous information I, f_{pri} the prior PDF of the parameters given I, and f_{L} the probability (*likelihood*) of the data given $\vec{\theta}$ and I. The prior (posterior) PDF is a representation of the state of knowledge about $\vec{\theta}$ before (after) performing the experiment. Therefore, the best fit to the spectrum is given by the value $\vec{\theta}_{\max}$ that maximises the posterior PDF.

To obtain $\vec{\theta}_{\max}$, a prior PDF must be adopted. The choice of the prior is an important task of the computation, as it may affect the final result in a significant way. We assume that the prior can take the form

$$f_{\mathrm{pri}}(\vec{\theta}|I) = f_{\varepsilon_\infty}(\varepsilon_\infty) \prod_i f_{\Delta\varepsilon_i}(\Delta\varepsilon_i) f_{\tau_i}(\tau_i) f_{\alpha_i}(\alpha_i) f_{\beta_i}(\beta_i), \tag{4}$$

which means that our previous knowledge of any of the parameters is independent of the others. The parameters α_i and β_i are bound to the interval (0, 1]. If we make no further assumption, a reasonable choice would then be a uniform prior $f_{\alpha_i}(\alpha_i) = f_{\beta_i}(\beta_i) = 1$ in (0, 1] (and null otherwise). On the other hand, if we assume that a sharp ($\alpha = 1$), symmetric ($\beta = 1$) peak is more probable than a broad, asymmetric one, an alternative would be to choose monotonic, increasing functions for f_{α_i} and f_{β_i}. In this paper we concentrate in the case of a known number of relaxation processes, usually individualised as steps (peaks) in the plot of the real (imaginary) parts of the spectrum. As both ε_∞ and the strengths $\Delta\varepsilon_i$ represent the heights of the different steps, they can be coarsely constrained to be in $[0, \varepsilon_{\max}]$. Here $\varepsilon_{\max}$ is an upper bound to the real part of the data, which can be determined simply by inspecting the plot. Once again, with no further assumption, the reasonable choice is a uniform prior $f_{\varepsilon_\infty} = f_{\Delta\epsilon_i} = \varepsilon_{\max}^{-1}$ in $[0, \varepsilon_{\max}]$. The parameters τ_i present a different problem. As they are related to the reciprocal of the peak frequencies, they can vary a priori by several orders of magnitude. Therefore, the best choice in this case is the Jeffreys prior $f_{\tau_i} \propto \tau_i^{-1}$ in $[\omega_{\max}^{-1}, \omega_{\min}^{-1}]$, where $\omega_{\min,\max}$ are the frequency limits of the spectrum (e.g., [16]). This prior has equal probability per time decade, and allows the exploration of the complete time domain of the experiment, which may be computationally

Please cite this article in press as: L. Ciocci Brazzano, et al., A Bayesian method for analysing relaxation spectra, Computer Physics Communications (2015), http://dx.doi.org/10.1016/j.cpc.2015.08.033