Available online at www.sciencedirect.com
ScienceDirect
Comput. Methods Appl. Mech. Engrg. 298 (2016) 121–158
Computer methods in applied mechanics and engineering
www.elsevier.com/locate/cma

A fast, certified and "tuning free" two-field reduced basis method for the metamodelling of affinely-parametrised elasticity problems

K.C. Hoang[a], P. Kerfriden[a,*], S.P.A. Bordas[a,b]

[a] *School of Engineering, Cardiff University, The Parade, CF24 3AA Cardiff, United Kingdom*
[b] *Faculté des Sciences, de la Technologie et de la Communication, University of Luxembourg, 6 rue Richard Coudenhove-Kalergi, L-1359, Luxembourg*

Received 24 November 2014; received in revised form 25 August 2015; accepted 27 August 2015

Abstract

This paper proposes a new reduced basis algorithm for the metamodelling of parametrised elliptic problems. The developments rely on the Constitutive Relation Error (CRE), and the construction of separate reduced order models for the primal variable (displacement) and flux (stress) fields. A two-field greedy sampling strategy is proposed to construct these two fields simultaneously and in an efficient manner: at each iteration, one of the two fields is enriched by increasing the dimension of its reduced space in such a way that the CRE is minimised. This sampling strategy is then used as a basis to construct goal-oriented reduced order modelling. The resulting algorithm is certified and "tuning-free": the only requirement from the engineer is the level of accuracy that is desired for each of the outputs of the surrogate. It is also shown to be significantly more efficient in terms of computational expense than competing methodologies.

Keywords: Two-field reduced basis method (TF-RBM); Model order reduction; Constitutive relation error; Goal-oriented greedy sampling; *a posteriori* error estimation

1. Introduction

Model order reduction is an increasingly popular family of metamodelling techniques for parametrised boundary value problems (BVP) solved using numerical methods. As opposed to response surface methodologies, the output of the computation is not interpolated directly over the parameter domain. Instead, one constructs an approximation of the BVP that can be solved efficiently, and from which the quantities of interest can be post-processed. The applicability of reduced order modelling requires a certain smoothness of the solution to the original BVP over the parameter domain.

Reduced order modelling (ROM) can be performed in various ways (*e.g. a priori* reduction approach [1], Proper Generalised Decomposition [2–4], operator interpolation in attractive manifolds [5], machine-learning-based

* Corresponding author.
E-mail address: pierre.kerfriden@gmail.com (P. Kerfriden).

http://dx.doi.org/10.1016/j.cma.2015.08.016

interpolations in attractive manifolds [6], classical mode synthesis, ...), but we will focus our discussion on the popular case of projection-based ROM (*e.g.* [7–9]). In this context, the reduced model is obtained by the projection of the original boundary value problem in a space of small dimension. Three ingredients are required for the metamodel to be efficient: (i) a reliable way to construct the projection space, (ii) an efficient (if possible optimum) projection of the solution to the BVP in this space and (iii) a method to decompose the numerical complexity of tasks (i) and (ii) in an "offline/online" manner. The latter point means that the expensive operations should be performed in advanced ("offline"), whilst the solution of the reduced model itself ("online") should remain computationally inexpensive.

The most popular way to extract the projection subspace is probably the Proper Orthogonal Decomposition (POD) [10–12]. The idea is to solve a finite number of realisations of the parametrised BVP to be reduced, and perform a spectral analysis of the space spanned by the corresponding solutions (*i.e.* the "snapshot"). This procedure delivers a hierarchy of subspaces in which the snapshots are best approximated in an average sense (*i.e.* an average of the distance between the original snapshot vectors and an optimal representation of these vectors in the reduced space is minimised). The BVP in then projected into one of these subspaces using a Galerkin [7,8,13,14], Petrov–Galerkin [15] or residual minimising formulation [16,17]. This projected BVP can be solved "online" at reduced numerical costs. The success of this approach lies in the fact that it is very general with regard to the nature of the original BVP. In particular, it remains the main candidate for the projection-based ROM of nonlinear problems [15,16,18–20]. However, it suffers from two drawbacks: (i) it is a priori difficult to choose the location of the snapshots, although some authors have recently addressed this issue, and (ii) the optimality of this procedure is established on the average over the parameter domain, and consequently its accuracy may be sensitive to outliers.

The reduced basis method (RBM) [21] is a powerful alternative to POD-based methodologies. It overcomes the two limitations mentioned previously by construction. The main idea is to construct the reduced spaces in such a way that the maximum error of projection of the solution over the entire parameter domain is minimised. As this problem is not solvable using direct algebraic tools, its solution is approximated by means of a greedy algorithm [22,23]. Typically, one-dimensional enrichments of the reduced space are looked for by identifying the point of the parameter domain that displays the largest projection error. The solution corresponding to this point is computed and its projection error is reduced by adding some of its components to the reduced basis. The main difficulty of the approach is to evaluate the projection error in an affordable and efficient manner over the entire parameter domain. This can be done efficiently in the case of affinely parametrised linear elliptic or parabolic problems, as reliable and inexpensive error bounds are available. This is the setting in which the RBM has been the most successful so far, but extensions to linear hyperbolic problems or more complex settings have been proposed, for instance using direct search [24,25], global optimisation [26] or gradient-based optimisation [27].

In the context of affinely parametrised elliptic PDEs [21,28], the reduced basis method proceeds as follows. Given a certain projection subspace, one evaluates an upper bound of the projection error (measured in "energy norm") over an exhaustive sampling of the parameter domain. The solution of the BVP is computed at this point and added to the basis of the previous reduced space after an orthonormalisation. Efficient upper bounds have been developed over the years, and the most widely used is probably the *residual-based a posteriori* error bounds based on the Successive Constraint Method (SCM) [29,30]. Once the projection subspace is sufficiently large (*i.e.* the maximum error of projection over the parameter domain is sufficiently small), this "offline" greedy algorithm is stopped. The metamodel can then be used online by performing a simple Galerkin projection of the original BVP in the reduced space.

In this paper, we propose an alternative to this procedure, based on our previous work on the Constitutive Relation Error (CRE) [14] and goal-oriented reduced basis sampling [24]. The first idea is to create two separate ROMs: one for the primal field of the elliptic BVP and an other one for the flux field. The primal ROM is required to deliver continuous solutions that satisfy the Dirichlet boundary conditions, whilst the flux ROM satisfies the balance equation of the flux and the Neumann boundary conditions. In this context, the Prager–Synge theorem gives us a relationship between the projection errors of the two fields and a certain distance between the primal and flux reduced solutions (*i.e.* the CRE [31]). A remarkable fact is that this distance can be computed in an inexpensive manner, and gives us (a) a joint indication of the accuracy of the two ROMs (b) an upper bound for the projection error in terms of primal variable and (c) an upper bound for the projection error in terms of flux. In view of remark (b), the CRE could be used as an alternative to a *residual-based a posteriori* error bound in the traditional RBM algorithm, as argued in the perspectives of our previous work [14]. However, we will propose a different, more elegant approach, the two-field reduced basis (TF-RBM) (see Fig. 1). The CRE will be used as an indication of accuracy of the two ROMs over the entire parameter domain, following remark (a). The training phase of the proposed reduced basis method