Reduced Basis Approximation and

A Posteriori Error Estimation for

Parametrized Partial Differential Equations

Anthony T. Patera   Gianluigi Rozza

Massachusetts Institute of Technology

Department of Mechanical Engineering

to appear in (tentative rubric)

MIT Pappalardo Graduate Monographs in Mechanical Engineering

Current Edition: V1.0 January 2007

Notice

Preface

Motivation

This book is about the evaluation of *input-output relationships* in which the *output* is evaluated as a functional of a *field* that is the solution of an *input-parametrized* partial differential equation (PDE). We will focus on applications in mechanics — heat and mass transfer, solid mechanics, acoustics, fluid dynamics — but we do not preclude other domains of inquiry within engineering (e.g., electromagnetics) or even more broadly within the quantitative disciplines (e.g., finance).

The *input-parameter* vector typically characterizes the geometric configuration, the physical or effective properties of the constitutive or phenomenological model, the boundary conditions and initial conditions, and any loads and sources. The *outputs of interest* might be the maximum system temperature, a crack stress intensity factor, structural resonant frequencies, an acoustic waveguide transmission loss, or a channel flowrate or pressure drop. Finally, the *fields* that connect the input parameters to the outputs can represent a distribution function, temperature or concentration, displacement, (acoustic) pressure, or velocity.

Our interest is in two particular contexts: the *real-time context*, and the *many-query context*. Both these contexts are crucial to computational engineering and to more widespread adoption and application of numerical methods for PDEs in engineering practice.

We first consider the *real-time context*; we can also characterize this context as "deployed" or "in the field" or "embedded." Typical activities in the real-time context fall within the broad "Assess, Predict, Act" framework. In the Assess stage we pursue parameter estimation or inverse analysis — to characterize current state; in the Predict stage we consider prognosis — to understand possible subsequent states; and in the Act stage we intervene to achieve our objectives — to influence realized future state. We now give several examples of these real-time processes;

note in all cases not all three stages of Assess, Predict, Act are invoked.

(*i*) Consider a crack in a critical structural component such as a composite-reinforced concrete support (or an aircraft engine). In the Assess stage we pursue Non-Destructive Evaluation (NDE) — say by vibration or thermal transient analysis — to determine the location and configuration of the delamination crack in the support. In the Predict stage we evaluate stress intensity factors to determine the critical load for brittle failure or the anticipated crack growth due to fatigue. And in the Act stage we modify the installation or subsequent mission profile to prolong life. In all cases, safety and economics require rapid and reliable response in the field. (See Part VIII for a detailed discussion of this particular problem.)

(*ii*) Consider an "immersed" object such as a tumor (or unexploded ordnance, or moving military target). In the Assess stage we apply parameter estimation techniques — say by acoustic or electromagnetic analysis — to determine the tumor location and geometric and physical characteristics. In the Predict stage we evaluate the potential lethality of the tumor. And in the Act stage we steer the intervention — therapy or surgery — to minimize the threat. Again, the timeliness *and* reliability of the analysis is all-important to safe and successful conclusion of the operation.

(*iii*) Consider heat treatment of a workpiece such as a turbine disk. In the Assess stage we apply inverse procedures to determine the (effective) heat transfer coefficients between the workpiece and the quenching bath. In the Predict stage we evaluate the anticipated residual stresses in the quenched workpiece for a given annealing schedule. And finally in the Act stage we apply optimal feedforward or feedback control to modify the annealing schedule in order to achieve lower residual stresses. For expensive materials, reliable quality for each workpiece "in process" is critical.

Clearly there are many other examples from other fields in engineering.

We next consider the *many-query context*, in which we make many (many) appeals to the input-output evaluation. One example of the many-query context, directly related to our discussion above, is *Robust* Assess, Predict, and Act scenarios: a parameter estimation, prognosis, and optimization framework in which we determine and subsequently accommodate *all possible model variations* — crack lengths, tumor dimensions, or heat transfer coefficients — consistent with (typically noisy or uncertain) experimental measurements of system and environmental conditions. This robust framework — a form of uncertainty quantification admittedly within a restrictive parametric context — requires extensive exploration of parameter space to determine and exploit appropriate "possibility regions"; we discuss this further in Part VIII.

A second important example of the many-query context is multiscale (temporal, spatial) and multiphysics "multimodels," in which behavior at a larger scale must "invoke" many spatial or temporal realizations of behavior at a smaller scale. Particular illustrative cases include stress intensity factor evaluation [7] within a crack fatigue growth model [61]; calculation of spatially varying cell properties [25, 29] within homogenization theory [24] predictions for macroscale composite properties; assembly and interaction of many similar building blocks [81] in large (e.g., cardio-vascular) biological networks; or molecular dynamics computations based on quantum-derived energies/forces [36]. In all these cases, the number of input-output evaluations is often measured in the tens of thousands.

Both the real-time and many-query contexts present a significant and often unsurmountable challenge to "classical" numerical techniques such as the finite element (FE) method. These contexts are often much better served by the reduced basis approximations and associated *a posteriori* error estimation techniques described in this book. Two important notes that the reader will soon appreciate.

First, the methods in this book do not replace, but rather build upon and are measured (as regards accuracy) relative to, a "truth" finite element approximation [21, 41, 125, 144, 159] — this is not an algorithmic competition, but rather an algorithmic collaboration. Second, the methods in this book are decidedly *ill-suited* for the "single-query" or "few-query" context.

Historical Perspective

The development of the Reduced Basis (RB) method can perhaps be viewed as a response to the imperatives described above. In particular, the two contexts described represent not only computational challenges, but also computational opportunities. We identify two key opportunities that can be gainfully exploited.

(I) In the parametric setting, we restrict our attention to a typically smooth and rather low-dimensional parametrically induced manifold: the set of fields engendered as the input varies over the parameter domain; in the case of single parameter, the parametrically induced manifold is a one-dimensional filament within the infinite dimensional space which characterizes *general* solutions to the PDE. Clearly, generic approximation spaces are unnecessarily rich and hence unnecessarily expensive within the parametric framework.

(II) In the real-time or many-query contexts, in which the premium is on *marginal cost* (or perhaps asymptotic average cost) per input-output evaluation, we can accept greatly increased pre-processing or "Offline" cost — not tolerable for a single or few evaluations — in exchange for greatly decreased "Online" (or deployed) cost for each new/additional input-output evaluation. Clearly, resource allocation typical for "single-query" investigations will be far from optimal for many-query and real-time exercises.

We shall view the development of RB methods in terms of these two opportunities.

As always, it is difficult to find all initial sources of a good idea, as good ideas tend to be prompted by common stimuli and to simultaneously occur to several investigators; hence we apologize for any omissions. Initial work on the Reduced Basis approximation — Galerkin projection on approximation spaces which focus (typically through Taylor expansions or "snapshots") on the low-dimensional parametrically induced manifold of opportunity (I) — grew out of two related streams of inquiry: from the need for more effective, and perhaps also more interactive, many-query design evaluation — [51] considers linear structural examples; and from the need for more efficient parameter continuation methods — [6, 99, 100, 102, 105, 106] consider nonlinear structural analysis problems. (Several modal analysis techniques from this era [95] are also closely related to RB notions.)

The ideas present in these early somewhat domain-specific contexts were soon extended to (*i*) general finite-dimensional systems as well as certain classes of PDEs (and ODEs) [19, 50, 76, 101, 107, 119, 130, 131], and (*ii*) a variety of different reduced basis approximation spaces — in particular Taylor and Lagrange [118] and more recently Hermite [66]. The next decade(s) saw further expansion into different applications and classes of equations, such as fluid dynamics and the Navier-Stokes equations [57, 65, 66, 67, 68, 112].

However, in these early methods, the approximation spaces tended to be rather local and typically rather low-dimensional in parameter (often a single parameter). In part, this was due to the nature of the applications — parametric continuation. But it was also due to the absence of *a posteriori* error estimators and effective sampling procedures. (In fairness, several early papers [103, 104, 105] did indeed discuss *a posteriori* error estimation and even adaptive improvement of the RB space; however, the approach could not be efficiently or rigorously applied to PDEs due to the computational requirements, the residual norms employed, and the absence of any stability considerations.) It is clear that in more global, higher-dimensional parameter domains the reduced basis predictions "far" from any sample points can not *a priori* be trusted, and hence *a posteriori* error estimators are crucial to reliability (and ultimately, safe engineering interventions *in particular in the real-time context*). It is equally clear that in more global, higher-dimensional — even three-dimensional — parameter domains simple tensor-product/factorial "designs" are not practicable, and hence sophisticated sampling strategies are crucial to efficiency.

Much of this book is devoted to (*i*) recent work on rigorous *a posteriori* error estimation and in particular error bounds for outputs of interest [87, 89, 120], and (*ii*) effective sampling strategies in particular for higher (than one) dimensional parameter domains [32, 98, 134, 149]. In fact, as we shall see, the former are a crucial ingredient in the latter — the inexpensive error bounds permit us first, to explore much larger subsets of the parameter domain in search of most representative or best "snapshots," and second, to know when we have *just enough* basis functions — and hence the simultaneous development of error estimation and sampling capabilities is not a coincidence. (We note that the greedy sampling methods described in this book are similar in objective to, but very different in approach from, more well-known Proper Orthogonal Decomposition (POD) methods [8, 23, 58, 75, 77, 97, 126, 127, 128, 142, 143, 156] typically applied in the temporal domain. However, POD economization techniques can be and have successfully been applied within the parametric RB context [31, 40, 43, 59, 86]. A brief comparative study is provided in Part I and again in Part IV.)

Early work certainly exploited the opportunity (II), but not fully. In particular, and perhaps at least partially because of the difficult nonlinear nature of the initial applications, early RB approaches did not fully decouple the underlying "truth" FEM approximation — of very high dimension $\mathcal{N}_{\mathrm{t}}$ — from the subsequent reduced basis projection and evaluation — of very low dimension N. More precisely, most often the Galerkin stiffness equations for the reduced basis system were generated by direct appeal to the high–dimensional "truth" representation: in nuts and bolts terms, pre- and post-multiplication of the "truth" stiffness system by rectangular basis matrices. As a result of this expensive projection the computational savings provided by RB treatment (relative to classical FEM "truth" evaluation) were typically rather modest [99, 118, 119] *despite* the very small size of the ultimate reduced basis stiffness system.

Much of this book is devoted to *full* decoupling of the "truth" and RB spaces through Offline-Online procedures: the complexity of the Offline stage depends on $\mathcal{N}_{\mathrm{t}}$ (the dimension of the "truth" finite element space); but the complexity of the Online stage — in which we respond to a new value of the input parameter — depends only on N (the dimension of the reduced basis space) and the parametric complexity of the operator and data. In essence, we are guaranteed the accuracy of a high-fidelity finite element model but at the very low cost of a reduced-order model. In the context of *affine parameter dependence*, in which the operator is expressible as the sum of products of parameter-dependent functions and parameter-independent operators, the Offline-Online idea is quite self-apparent and indeed has been re-invented often [15, 65, 70, 112]; however, application of the concept to *a posteriori* error estimation — note the Online complexity of both the output *and* the output error bound calculation must be independent of $\mathcal{N}_{\mathrm{t}}$ — is more involved and more recent [62, 120, 121]. In the case of *nonaffine parameter dependence* the development of Offline-Online strategies is much less transparent, and only in the last few years have effective procedures — in effect, efficient methods for approximate reduction to affine form — been established [18, 54, 135]. Clearly, Offline-Online procedures are a crucial ingredient in the real-time context.

We note that historically [50] and in this book RB methods have been built upon, and measured (as regards accuracy) relative to, *finite element* "truth" discretizations (or related spectral element approaches [81, 82, 83, 84, 111]) — the variational framework provides a very convenient setting for approximation and error estimation. However there are certainly many good reasons to consider alternative "truth" settings: a systematic finite volume framework for RB approximation and *a posteriori* error estimation is proposed and developed in [60]. We do note that boundary and integral approximations are less amenable to RB treatment or at least Offline-Online decompositions, as the inverse operator will typically not be affine in the parameter.

Scope and Roadmap

We begin with General Preliminaries. The purpose of the General Preliminaries is to recall — in a form relevant to the subsequent development — the background material on which the rest of the book shall rest. We discuss basic elements of functional analysis: Hilbert spaces (real and complex); product spaces; bases; inner products and norms; the Cauchy-Schwarz inequality; linear bounded forms and dual spaces; the Riez representation theorem; and finally bilinear forms. We review the fundamental properties associated with bilinear forms — the coercivity and inf-sup stability conditions [9] and the continuity condition — and introduce associated eigenproblems of computational relevance. And finally we introduce the basic smoothness hypotheses, function spaces, and norms associated with variational formulation and approximation of second order partial differential equations. In all cases we consider both the standard definition as well as the (in most cases, rather self-evident) extension to the parametric context of particular interest in this book.

Following the General Preliminaries each subsequent Part of this book addresses a different class of problems. In each case we first identify the abstract formulation of the problem and then develop particular instantiations (corresponding to particular equations/physical phenomena) and associated specific examples. We then proceed to the formulation and analysis: reduced basis approximation; optimal sampling procedures; *a priori* theory; rigorous *a posteriori* error estimation; and Offline-Online computational strategies. Finally, in each case we provide software in the form of MATLAB® `.m` files that, given an appropriate "truth" finite element model provided by the reader implements the methods developed.

In Part I of this book we treat the particularly simple case of Parametrically Coercive and Compliant Affine Linear Elliptic Problems. This class of problems — in which each term of the parametric development is independently symmetric positive semidefinite, and for which furthermore the load/source functional and output functional coincide — permits a simple exposition of the key ideas of the book: RB spaces and suitably orthogonalized bases; Galerkin projection and optimality; greedy quasi-exhaustive sampling procedures; the role of parametric smoothness in convergence; rigorous and relatively sharp *a posteriori* error estimation; and Offline-Online computational strategies. We illustrate this Part of the book with thermal conduction and linear elasticity examples that involves $O(10\text{–}20)$ independent parameters [137]. (The reader should guard against disappointment: it is really only for this simple class of problems, for which lower bounds for stability constants can be explicitly and readily extracted, that we can entertain so many parameters.)

In Part II of this book we consider the more general case of Coercive Affine Linear Elliptic Problems. We no longer require *parametric* coercivity; we now permit non-symmetric bilinear forms a; and we now consider arbitrary linear (bounded) output functionals — and perhaps multiple outputs. At this stage we can also treat, and we hence we introduce and exercise, the general class of piecewise-affine geometric and coefficient parametric variations consistent with the requirement of affine parameter dependence. Physical instantiations include general heat conduction (the Poisson equation) problems; forced-convection heat transfer (the convection-diffusion equation) problems; and general linear elasticity problems. In relation to Part I, the key new methodological elements are the development of (*i*) primal-dual (with adjoint) approximation [115] for RB [120], (*ii*) an *a posteriori* theory of general (non-compliant, and hence "non-energy") outputs [5, 22, 110] for RB [120, 139], and (*iii*) an efficient Offline-Online computational procedure for the construction of a lower bound for the general coercivity constant [62] required by our *a posteriori* estimators.

In Part III of this book we consider the most general case of *Non-Coercive* Affine Linear Elliptic Problems (we address the particular stability issues [26, 27] associated with Saddle Problems, in particular the Stokes equations of incompressible flow [133, 135, 136], in Part VI). Physical instantiations include the ubiquitous Helmholtz equation relevant to time-harmonic acoustics, elasticity, and electromagnetics. In this Part we also introduce a special formulation (perforce non-coercive) for *quadratic* output functionals [61] — important in such applications as acoustics and linear elastic fracture theory. In relation to Part II, the key new methodological elements are the development of (*i*) *discretely* stable primal-dual RB approximations [89], (*ii*) an efficient Offline-Online computational procedure [62, 139] for the construction of a lower bound for the general *inf-sup* constant [9] required by our *a posteriori* estimators. (The latter, in essence a lower bound for a *singular value*, demands considerable Offline resources and is certainly a limiting factor in the treatment of higher dimensional parameter spaces: an opportunity for further work.)

RB-like snapshot ideas (typically enhanced by sophisticated POD sampling variants) are also common in certain Reduced Order Model (ROM) approaches in the temporal domain [14, 38, 39, 93, 114, 129, 143, 151, 152]; more recently greedy sampling approaches have also been considerated in [20]. However, combined "parameter + time" approaches — essentially the marriage of ROM in time with

RB in parameter, and sometimes referred to as PROM (Parametric ROM) — are relatively uncommon [31, 40, 43, 59, 49, 86, 141]. In Part IV of this book we explore the "parameter + time" paradigm in the important context of Affine Linear Parabolic Problems such as the heat or diffusion equation and the passive scalar convection-diffusion equation (also the Black-Sholes equation of derivative theory [117]). In particular, in Part IV we extend to parabolic PDEs the primal-dual approximations, greedy sampling strategies, *a posteriori* error estimation concepts, and Offline-Online computational strategies developed in Part II for elliptic PDEs — with particular focus on the accommodation of an "evolution" parameter t [56, 60, 132]. Two qualifications: we restrict attention to discrete-time parabolic equations corresponding to simple Euler backward (or Crank-Nicolson) discretization of the original continuous PDE; and except briefly (where we permit a weaker Garding inequality) we only consider parabolic equations associated with *coercive* spatial operators.

In Part V of the book we consider the extension, in both the elliptic and parabolic cases, to nonaffine problems. The strategy is ostensibly simple: reduce the nonaffine operator and data to approximate affine form, and then apply the methods developed for affine operators in Parts II, III and IV. However, this reduction must be done efficiently in order to avoid a proliferation of parametric functions and a corresponding degradation of Online response time. The approach we describe here is based on the Empirical Interpolation Method (EIM) [18]. We first describe the Empirical Interpolation Method for efficient approximation of fields which depend (smoothly) on parameters: a collateral RB space for the offending nonaffine coefficient functions; an interpolation system that avoids costly ($\mathcal{N}_t$-dependent) projections; and several (from less rigorous/simple to completely rigorous/quite cumbersome) *a posteriori* error estimators. We then apply the EIM within the context of RB treatment of elliptic and parabolic PDEs with nonaffine coefficient functions [54, 135, 145]; the resulting approximations preserve the usual Offline-Online efficiency — the complexity of the Online stage is independent of $\mathcal{N}_{\rm t}$.

In Part VI of the book we treat several elliptic problems with polynomial nonlinearities. Here our coverage is admittedly somewhat more anecdotal: we can not uphold our standards of rigor (in *a posteriori* error estimation) or efficiency (in Online requirements) for general nonlinear problems; we thus consider several representative examples that illustrate essential points. First [149], we consider an elliptic problem with stabilizing cubic nonlinearity [34] . This problem illustrates both the possibility and difficulty of efficient RB Galerkin approximation of (lowish-order) polynomial nonlinearities, and the availability in some very special circumstances of a very simple nonlinear *a posteriori* error theory. Second, we proceed to the (quadratically nonlinear) Navier-Stokes equations [27, 52, 57] of incompressible fluid flow; for simplicity we consider here only a single parameter, the Reynolds number. For the Navier-Stokes equations (and for nonlinear equations more generally) we can not appeal to any simple monotonicity arguments; our focus is thus on the computational (quantitative) realization of the general Brezzi-Rappaz-Raviart ("BRR") *a posteriori* theory [28, 34] — and development of associated sampling procedures — within the reduced basis Offline-Online context [98, 147, 148]. (We also address here the constuction of div-stable [26, 27] RB (Navier)-Stokes approximations [122, 136].) Third and finally, we consider symmetric eigenproblems associated with (say) the Laplacian [10] or linear elasticity operator: we present formulations for one or two lowest eigenvalues [87] and for the first "many" eigenvalues (as relevant in quantum chemistry [35, 36]). Here, implicitly, the interpretation of the BRR theory is unfortunately less compelling due to the (guaranteed!) multiplicity of often nearby solutions.

In Part VII we consider nonpolynomial nonlinearities (in the operator and also output functional) for both elliptic and parabolic PDEs. Our focus is on the application/extension of the Empirical Interpolation Method to this important class of problems [54]: in effect, we expand the nonlinearity in a collateral reduced basis expansion, the coefficients of which are then obtained by interpolation relative to the reduced basis approximation of the field variable. We are therefore able to maintain, albeit with some effort, our usual Offline-Online efficiency — Online evaluation of the output is independent of $\mathcal{N}_{\rm t}$. (For alternative approaches to nonlinearities in the ROM context, see [16, 39, 113].) Unfortunately, for this difficult class of problems we can not provide (efficient) rigorous *a posteriori* error estimators. (It it thus perhaps not surprising that initial work in RB methods [99, 102], focused on highly nonlinear problems, did not attempt complete Offline-Online decoupling or rigorous error estimation.) The trade-off between rigor and model complexity is inevitable; we hope the reader finds the methods of Part VI useful despite the lower standards of certainty.

In Part VIII we depart from our usual format and instead consider a real-time and many-query application of reduced basis approximation and *a posteriori* error estimation: robust parameter estimation for systems described by elliptic and parabolic PDEs — from outputs we wish to deduce inputs [55]. (Other applications of RB, in particular to optimization and control, can be found in [109, 123].) Our focus is on the rigorous incorporation of experimental error and numerical (RB) error bounds in the specification of "possibility regions": regions of the parameter domain consistent with available (noisy) experimental data. (For well-posed or "identifiable" [17] systems the possibility region will shrink to the unique value of the unknown parameter(s) as the experimental error and reduced basis error tend to zero. However, many interesting "systems" — which should be understood to comprise the model, the experimental measurements, the unknown inputs, and the selected outputs — are not identifiable.) In practice, except for special problems, these possibility regions can not be constructed except by truly exhaustive and exhausting (even at "reduced basis speed") calculation; we thus also consider various efficient procedures for approximating these possibility regions. We consider an example of transient thermal conduction inverse analysis.

Finally, in Part IX, we discuss briefly two topics on the research frontier. First, we shall consider the "reduced basis element method" [81, 82, 83, 84]: a marriage of reduced basis and domain decomposition concepts that permits much greater parametric complexity and also provides a framework for the integration of multiple models. Second, (at least linear) hyperbolic problems are also ripe for further development: although there are many issues related to smoothness, stability, and locality, there are also important proofs-of-concept [60, 111] in both the first order and second order contexts which demonstrate that RB approximation and *a posteriori* error estimation can be gainfully applied to hyperbolic equations. For both topics, we briefly review the current status and identify outstanding challenges.

Intended Audience

We have in mind four audiences. The first audience is professional researchers, faculty, and graduate students in the area of numerical methods for PDEs: *developers* (and analysts) of numerical methods. We hope that the formulations and theory summarized in our research monograph will provide a good foundation for further developments in reduced basis methodology and analysis. (We also hope that the book might be appropriate as a secondary source in a graduate course on numerical methods for PDEs: the RB framework is a very good laboratory in which to understand, exercise, and observe many basic aspects of computational approaches to PDEs.)

The second audience is computational engineers — professionals or graduate students for which application of computational methods for PDEs plays an essential role: advanced *users* of numerical methods. It is for this audience (and educators, see below) that we hope the software component of the book will prove useful: a rapid and easy way to apply the reduced basis approximation, greedy sampling, *a posteriori* error estimation, and Offline-Online approaches described in this book to problems of interest in the research, development, design/optimization, and "Assess-Predict-Act" contexts. We do provide one word of caution: the software we provide is blackbox (actually, only "somewhat" blackbox since the formulations and theory must be understood to properly invoke and connect the modules) once the "truth" finite element approximation is appropriately specified; however, the requisite finite element ingredients can not always be generated by a (third-party) FE program without access to the source code. Hence we implicitly assume that the computational engineer is willing and able to take screwdriver (though hopefully not jackhammer) in hand.

We note that FE packages oriented towards, or based upon, domain decomposition for definition of problem geometry and coefficients are particularly well suited to the reduced basis approach. Example of such packages are the MATLAB PDE Toolbox® or COMSOL Multiphysics™. In this case, it is possible to create the finite element inputs to the RB software without modification of, or *even access to*, the source code — the available assembly and export features suffice. We shall indicate on several occasions (with the MATLAB PDE Toolbox® as our vehicle) the simple and clean interface between a domain-decomposition "cognizant" finite element package and our own reduced basis software.

The third audience is university engineering educators (and ultimately, as *end users*, students). The application of finite element simulations for visualization, assessment of classical engineering models and approximations, and parameter estimation and design/optimization — both in class and as part of homework assignments and projects — has remained quite limited. Of course, complex user interfaces are part of the problem. But even more fundamental is the relatively slow response time of even very good codes: for an in-class demonstration, one minute or even 10 seconds for typically just a *single* parameter value is an eternity; and even for homework assignments, large operation counts and storage requirements can quickly obscure the pedagogical point. Clearly, this context can benefit from a real-time and many-query (many-student) perspective: in particular, we hope that, with the software we provide, educators can "automatically" and quickly develop *very fast* — and, thanks to the *a posteriori* error estimators, *reliable and physically relevant* — Online modules for visualization and input-output evaluation of complex problems. However, we do again caution that the professor — or able-bodied Teaching Assistant — must have access to the necessary finite element infrastructure in order to develop the "truth" prerequisities.

Our fourth audience is very broad: we hope that our text will become a "coffee-table" book. In this age of technology, all informed citizens should be acutely interested in reduced basis approximation and *a posteriori* error estimation for parametrized partial differential equations. Perhaps families can organize group readings so that both young and old can appreciate the content and implications. Or even better, perhaps each member of the family can purchase his or her own copy of the book to keep — or to give as thoughtful gifts.

Acknowledgements

This book discusses research that is not our own — and that preceded our own research by in some cases several decades — and also research that is our own. But "our own" is a misnomer, as our research project has been a collaborative effort with many colleagues. We give here thanks, we hope without any unintentional omissions.

First and foremost, we would like to acknowledge the many and important contributions of our longstanding co-conspirator, Yvon Maday. Our collective effort on reduced basis concepts (and on other ideas in earlier years) has been enjoyable and productive.

We would also like to acknowledge our very fruitful collaborations with Claude Le Bris, Analisa Buffa, Eric Cancès, Jan Hesthaven, Jaime Peraire, Christophe Prud'homme, Alfio Quarteroni, Einar Rønquist, Gabriel Turinici, and Karen Willcox. In many cases our joint activities have taken the form of co-supervision of talented students — a particularly enjoyable aspect of academia.

We would specifically like to thank the many MIT Masters and PhD students and post-docs, visiting students and post-docs, and also Singapore-MIT Alliance students, who have contributed to our reduced basis effort over the years: Shidrati Ali, ANG Wei Sin, Sebastien Boyaval, Luca Dedè, Simone Deparis, Revanth Garlapati, Ginger, Martin Grepl, Thomas Leurent, Alf Emil Løvgren, HUYNH Dinh Bao Phuong, Luc Machiels, NGUYEN Ngoc Cuong, Ivan Oliveira, George Pau, Jerónimo Rodríguez, Dimitrios Rovas, Sugata Sen, TAN Alex Yong Kwang, and Karen Veroy. A long alphabetical list of this variety perforce approaches anonymity, and hides the singular contributions of many of these very talented individuals. We would also like to acknowledge the contributions of Roberto Milani and Annalisa Quaini, Masters students at Politecnico di Milano.

We can not thank enough Ms Debra Blanchard for all of her efforts in managing our research projects and collaborations, and also, and in particular, for her many direct contributions to this book: for organizing our sometimes chaotic work flow; for performing all the mathematical typesetting with exacting standards; for preparing all the artwork and figures; and, occasionally, for spotting a mathematical error or two. Ms Blanchard is also responsible for the design and implementation of the website with which this book is associated.

ATP wishes to thank the Chief Technology Officer at his Lincoln office, Ms Anddie Chan, for designing, maintaining, and operating a computer system with all the latest information and document processing technology. This book owes much to Ms Chan's dedication, professionalism, and Victorian work ethic.

We would like to recognize that much of "our own" research reflected in these pages was supported by the US Defense Advanced Research Projects Agency, by the US Air Force Office of Scientific Research, and by the Singapore MIT-Alliance. We are also grateful to the "MIT Pappalardo Graduate Monographs in Mechanical Engineering" for supporting the development of this book and associated website and software.

ATP & GR, January 2007.