Control Decisions in the Bayesian Automatic Adaptive Quadrature

Gheorghe Adam & Sanda Adam
LIT-JINR Dubna & IFIN-HH Bucharest

ROLCG 2015, Cluj-Napoca, October 28-30, 2015

SUMMARY: Three topics in the Bayesian automatic adaptive quadrature are discussed:

- (1) The ***accuracy specifications*** controlling the error tolerance of the derived output need a ***twofold scrutiny***:
  - check of the **reliability** of the accuracy specification **at input**;
  - check for the **need of integrand dependent accuracy bounds**.
- (2) **Avoidance of overcomputing & minimization of the hidden floating point loss of precision** asks for the consideration of **three classes** of integration domain lengths endowed with specific quadrature sums: ***microscopic*** (trapezoidal rule), ***mesoscopic*** (Simpson rule), and ***macroscopic*** (quadrature sums of high algebraic degrees of precision).
- (3) ***Sensitive diagnostic tools for the Bayesian inference*** on macroscopic ranges, coming from the use of Clenshaw-Curtis quadrature, are illustrated

Overview

- **General frame of the Bayesian automatic adaptive quadrature (BAAQ)**
- Reliability of accuracy specifications
- Three classes of integration domain lengths
- Improved diagnostic tools for Bayesian inference over macroscopic range lengths
- Conclusions

We consider the BAAQ numerical solution of the Riemann integral

$$I \equiv I[f] = \int_a^b g(x)f(x)dx, \qquad -\infty < a < b < \infty,$$

under the assumption that the real valued ***integrand function*** $f(x)$ is continuous almost everywhere on $[a, b]$ such that I exists and is finite.

The ***weight function*** $g(x)$ either absorbs an analytically integrable difficult factor in the integrand (e.g., endpoint singularity or oscillatory function), or else $g(x) \equiv 1,\ \forall\, x \in [a, b]$.

The automatic adaptive quadrature (AAQ) solution of I rests on the use of an ***interpolatory quadrature sum*** to get an approximation $Q \equiv Q[f]$ to $I[f]$.

The ***meaningfulnes***s of $Q[f]$ is assessed by deriving a bound $E \equiv E[f]>0$ to the remainder $R[f] = I[f] - Q[f]$.

For a prescribed accuracy τ ***requested at input***, the approximation Q to I is assumed to end the computation provided

$$|R[f]| < E < \tau.$$

The definition of τ needs two parameters: the **absolute accuracy** ε_a and the **relative accuracy** ε_r , such that

$$\tau = \max\{\varepsilon_a,\ \varepsilon_r \cdot |I|\} \simeq \max\{\varepsilon_a,\ \varepsilon_r \cdot |Q|\}\ .$$

If the remainder boundedness condition is not satisfied, the AAQ approach to the solution attempts at decreasing the error E by the ***subdivision*** of the integration domain $[a,b]$ ***into subranges*** using ***bisection*** and the computation of a ***local pair*** $\{q,\ e > 0\}$ over each newly defined subrange $[\alpha, \beta] \subset [a, b]$.

This procedure builds a ***subrange binary tree*** the evolution of which is controlled by an associated ***priority queue***.

Local pairs $\{q_i, e_i > 0\}$ are computed over the i-th subrange of $[a, b]$ and ***global*** outputs $\{Q_N, E_N > 0\}$ are got by summing the results obtained over the N existing subranges in $[a, b]$.

After each subrange binary tree update, the termination criterion is checked until it gets fulfilled.

The existing strict mathematical bounds to $R[f]$ are of little use in the implementation of practical codes.
The derivation of a ***practical bound*** e > 0 to q rests on ***probabilistic arguments*** the validity of which is always subject to doubt.
The BAAQ advancement to the solution incorporates the rich AAQ accumulated empirical evidence into a general frame based on the ***Bayesian inference***.
Essentially, the probabilistic character of the AAQ approach is preserved. However, each step of the gradual advancement to the solution is scrutinized based on a set of ***hierarchically ordered*** criteria which enable decision taking in terms of the established diagnostics.

Overview

- General frame of the Bayesian automatic adaptive quadrature
- **Reliability of accuracy specifications**
- Three classes of integration domain lengths
- Improved diagnostic tools for Bayesian inference over macroscopic range lengths
- Conclusions

Input Reliability Check

- Let $\{\varepsilon_a^{(i)}, \varepsilon_r^{(i)}\}$ denote the input provided values for the accuracy parameters.
- The input reliability check aims at setting up *reliable* values $\{\varepsilon_a^{(r)}, \varepsilon_r^{(r)}\}$ to be used within the BAAQ.
- $\varepsilon_a^{(i)}$ is mapped onto a non-negative value $\varepsilon_a^{(r)}$,

$$\varepsilon_a^{(r)} = \max\{\varepsilon_a^{(i)}, 0.0\}.$$

- $\varepsilon_r^{(i)}$ is mapped onto an inner value $\varepsilon_r^{(r)}$ satisfying

$$\varepsilon_r^{(r)} = \min\{rceil(), \max\{\varepsilon_r^{(i)}, rfloor()\}\};$$

$rceil() = 2^{-8}$; $rfloor() = epmach()/rceil()$, $epmach() = 2^{-52}$ denote two empirically defined *environment functions*.

Integrand Dependent Accuracy Bounds

- For the derivation of integrand dependent accuracy bounds, we compute $Q_N = Q_N[f]$ and $T_N = Q_N[|f|]$.
- At the first attempt to solve the integral:

If the result $Q_1 = 0.0$ was obtained, then we set $E_1 = 0.0$ and decide on the *end of computation*.

- Otherwise, $Q_N \neq 0.0$ will hold, hence we may define

$$\rho_N = epmach() \cdot (T_N / |Q_N|).$$

The condition $\rho_N > rceil()$ diagnoses the occurrence of catastrophic cancellation by subtraction, hence the *end of computation* with a message of useless output.

Integrand Dependent Accuracy Bounds

- Otherwise, the ***termination of computations is checked*** for ***integrand dependent accuracy bounds at output***

$\{\varepsilon_a{}^{(o)}, \varepsilon_r{}^{(o)}\}$

$$|I - Q_N| < E_N < \max\{\varepsilon_a{}^{(o)}, \varepsilon_r{}^{(o)}|Q|\}.$$

The output accuracy parameters are obtained from the validated input $\{\varepsilon_a{}^{(r)}, \varepsilon_r{}^{(r)}\}$ as follows:

$$\varepsilon_a{}^{(o)} = \min\{\varepsilon_a{}^{(r)}, \max\{|Q_N|, 1.0\} \cdot rceil()\}.$$

$$\varepsilon_r{}^{(o)} = \max\{\varepsilon_r{}^{(r)}, \rho_N\}.$$

Overview

- General frame of the Bayesian automatic adaptive quadrature
- Reliability of accuracy specifications
- **Three classes of integration domain lengths**
- Improved diagnostic tools for Bayesian inference over macroscopic range lengths
- Conclusions

Algebraic and Floating Point Degrees of Precision

- The ***algebraic degree of precision***, d, is an invariant feature of a quadrature sum over the field $\mathbb{R}$ of the real numbers: its value remains *constant* irrespective of the *extent* and the *localization* of the current integration domain over the real axis.
- Under floating point computations, the characterization of an interpolatory quadrature sum is made by its ***floating point degree of precision***, d_{fp}.

Given the integration domain $[a, b]$ $(a \neq b)$, the value of d_{fp} is determined by the magnitude of the parameter

$$\lambda = |L| / \max\{1.0, X\} \qquad (\lambda > 0.0),$$

where

$$L = b - a \quad (L \neq 0.0); \quad X = \max\{|a|, |b|\} \quad (X > 0.0).$$

The quantity λ defines the *floating point scale length* of $[a, b]$.

Features of the Floating Point Degree of Precision

- *Gliding integration range* [0,1] *on the real axis.*

The following plot gives outputs for the family of 1024 integration ranges

$$\{[j\alpha, j\alpha + \beta], \alpha = \beta = 1; j = 0, 1, ..., 1023\}$$

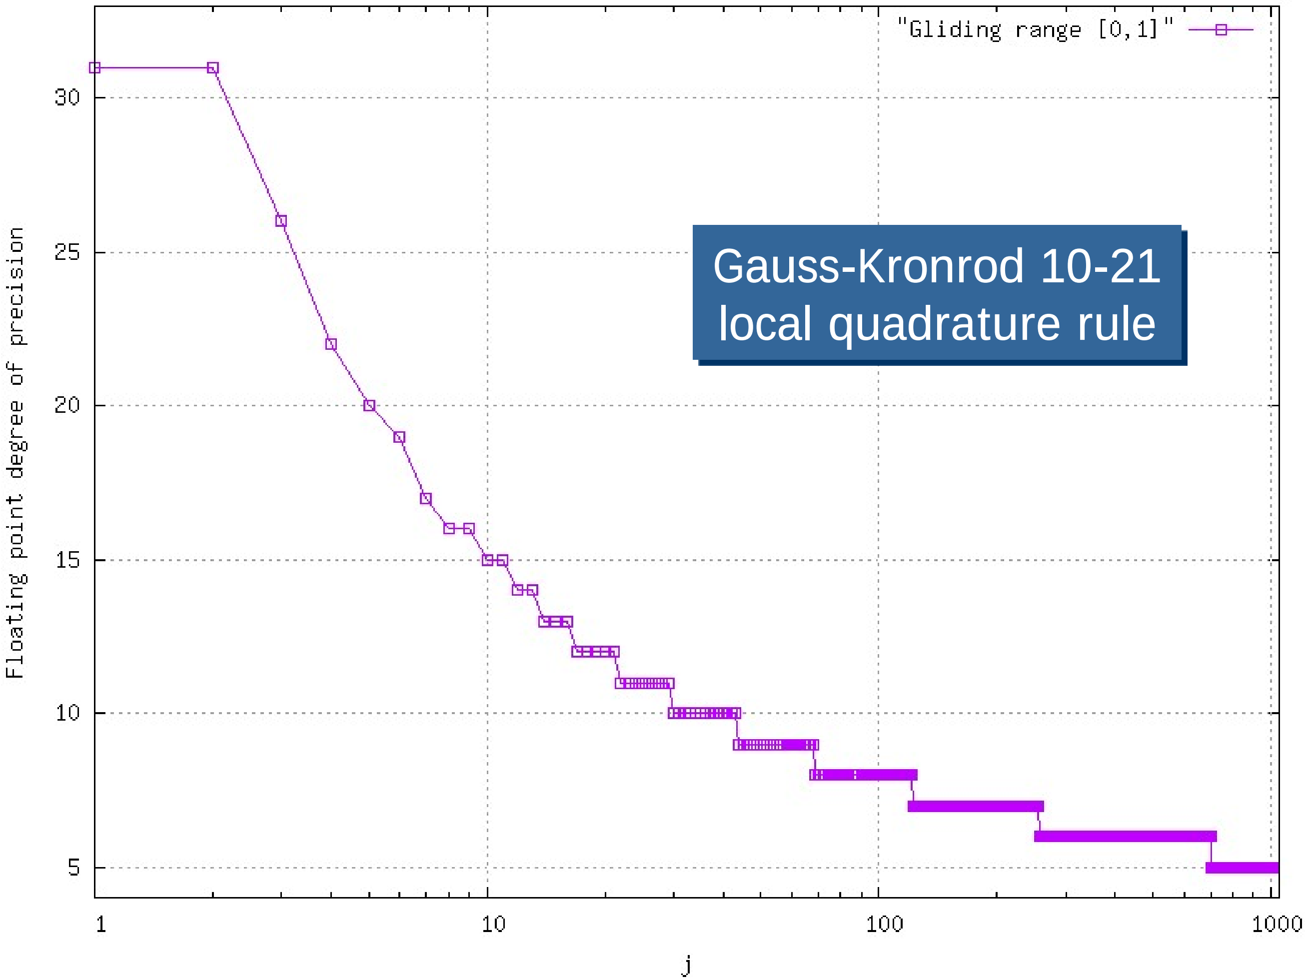

Variation of the floating point degree of precision of the GK 10-21 local quadrature rule over the gliding range [0, 1] versus its distance j from the origin. It is shown that $d_{fp} = d = 31$ at low j values ($j = 0, 1, 2$), then d_{fp} abruptly decreases at larger but small enough j, to show slower decreasing rates under the displacement of [0,1] far away from the origin, reaching a bottom value $d_{fp} = 5$ at $701 \le j \le 1023$.

The Inverse Problem

- Find the family of the integration ranges $[\alpha, \beta]$ over which the floating point degree of precision cannot exceed a prescribed value d.
- Possibilities at hand: $d \gg 1$ (the standard assumption of the AAQ and previous BAAQ implementations), $d = 4$ (the, perhaps composite, Simpson rule), $d = 2$ (the, perhaps composite, trapezoidal rule).

Each of these three cases corresponds to specific integration domain lengths, which are separated from each other by two empirically chosen thresholds, τ_μ and τ_m , defined below. They separate ***three classes of integration domain lengths*** corresponding to various quadrature sums at hand.

Three Classes of Integration Domain Lengths

- ***Microscopic ranges*** [using (composite) ***trapezoidal rule*** ($d = 2$)], are characterized by the threshold condition

$$0 < \min(X, \rho) \le \tau_\mu = 2^{-22}\,.$$

- ***Mesoscopic ranges*** [using (composite) ***Simpson rule*** ($d = 4$)], are characterized by the threshold condition

$$\tau_\mu = 2^{-22} < \min(X, \rho) \le \tau_m = 2^{-8}\,.$$

- ***Macroscopic ranges*** [using ***quadrature sums of high algebraic degrees of precision***], are characterized by the threshold condition

$$\min(X, \rho) \le \tau_m = 2^{-8}\,.$$

== $\boldsymbol{\tau_\mu = 2^{-22}}$ corresponds to $d = 3$

== $\boldsymbol{\tau_m = 2^{-8}}$ corresponds to $d = 8$; it results in **negligible round off** over the macroscopic domain lengths.

Given a microscopic (or mesoscopic) integration range $[\alpha, \beta] \subset \mathbb{R}$, the ***minimization of the round-off errors*** within the trapezoidal rule (or Simpson rule respectively) is secured as follows.
The integration range $[\alpha, \beta]$ is mapped onto [0, 1] by the substitution (floating point representations and floating point operations with the involved quantities are assumed) $x = \alpha + hy$, $h = \beta - \alpha$
and the current Riemann integral over $[\alpha, \beta]$ is transformed accordingly to get

$$I[\varphi] = h \int_0^1 g(\alpha + hy) \cdot \varphi(y) dy, \qquad \varphi(y) = f(\alpha + hy)$$

This step associates the ***unavoidable round-off cancellation error*** coming from the computation of the integration domain length h.
Besides the minimum number of integrand evaluations asked by the corresponding quadrature rule, additional requested integrand evaluations are performed at suitable newly added ***machine number*** reduced abscissas inside [0, 1] in terms of which all the newly involved subtraction operations in the resulting composite rules are done ***exactly***.

Overview

- General frame of the Bayesian automatic adaptive quadrature
- Reliability of accuracy specifications
- Three classes of integration domain lengths
- **Improved diagnostic tools for Bayesian inference over macroscopic range lengths**
- Conclusions

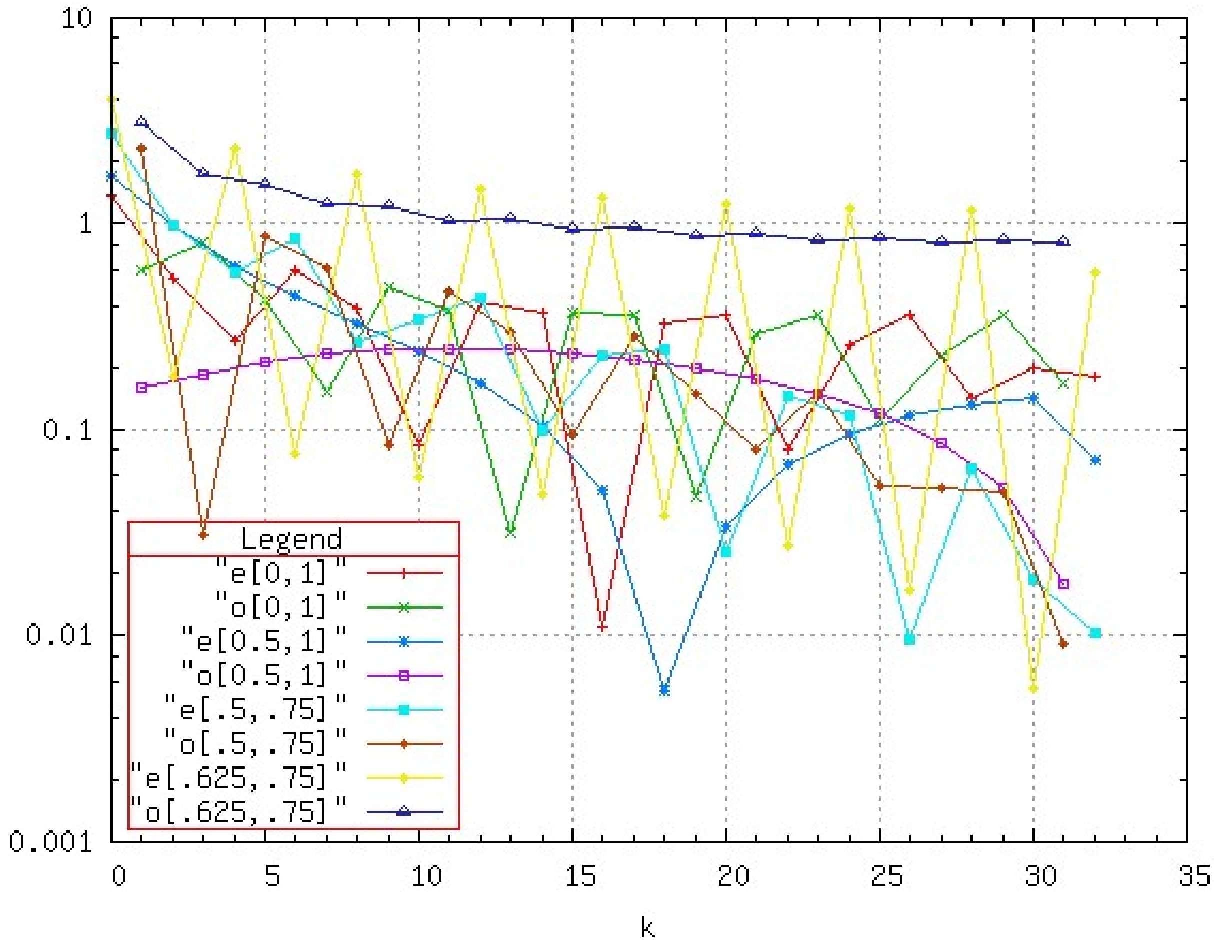

Ill-integrand behavior illustrated in the irregular variation of the Chebyshev expansion coefficients for the integrand $f_1(x) = |x^2 + 2x - 2|^{-1/2}: [0,1] \to \mathbb{R}$ which shows an inner singularity at $x_s = \sqrt{3} - 1$ over the specified subranges. The file notations start with the specification of the rank of the Chebyshev subset: 'e' (for even) and 'o' (for odd).

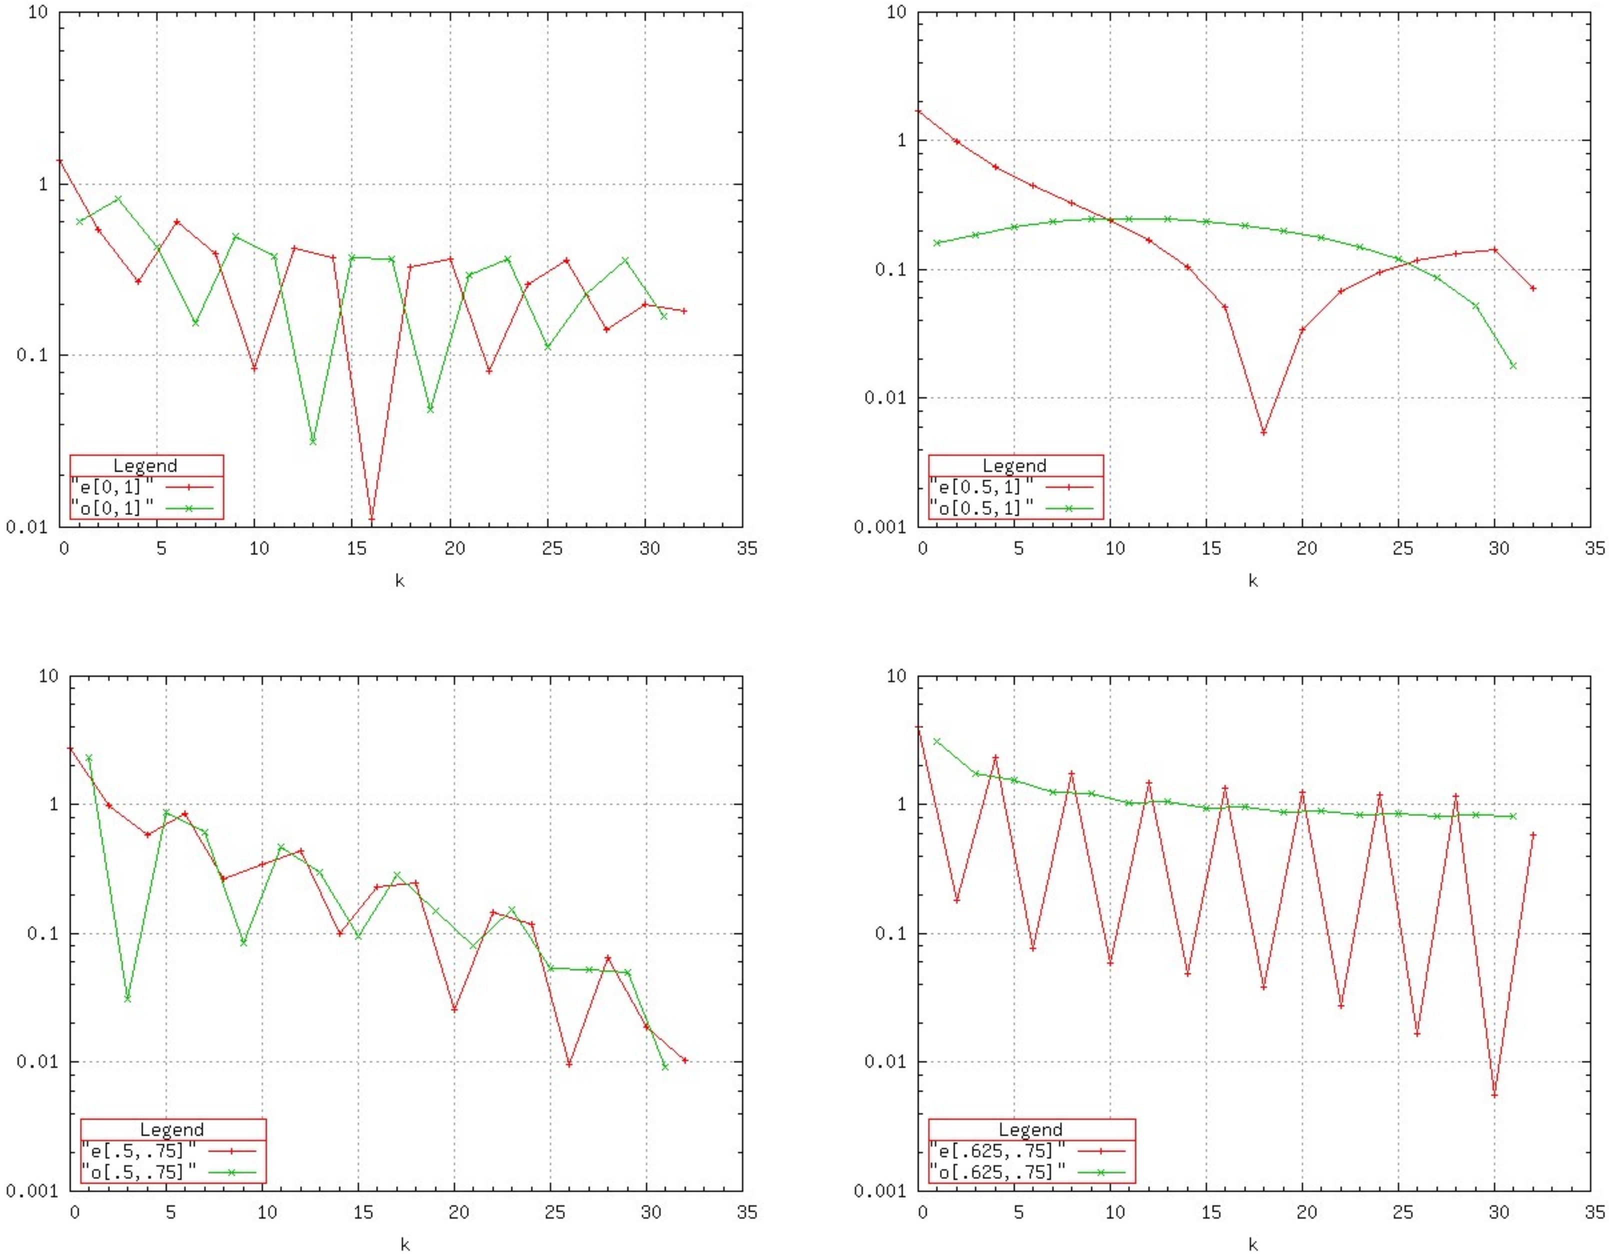
Legend
"e[0,1]"
"o[0,1]"
k
Legend
"e[0.5,1]"
"o[0.5,1]"
k
Legend
"e[.5,.75]"
"o[.5,.75]"
k
Legend
"e[.625,.75]"
"o[.625,.75]"
k

***Typical patterns of variation* of the absolute magnitudes of the Chebyshev expansion coefficients within the *even and odd rank subsets* versus the coefficient labels**

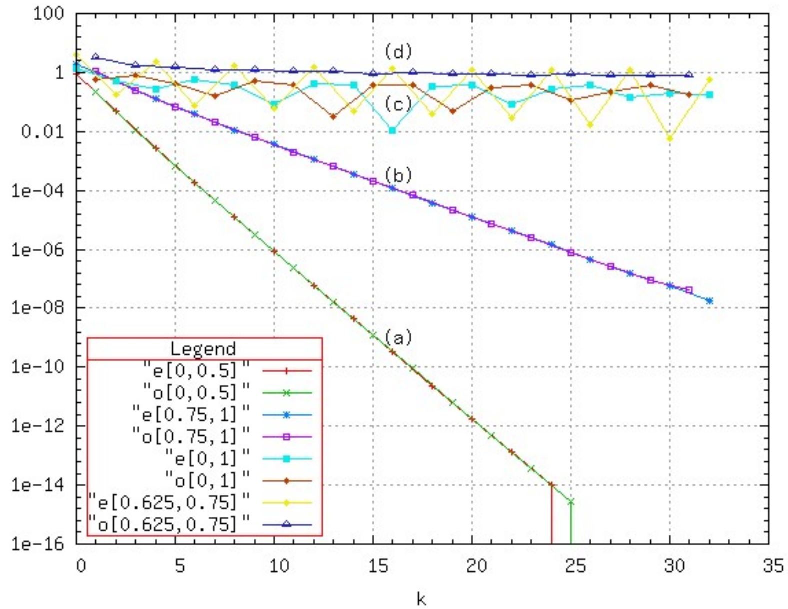

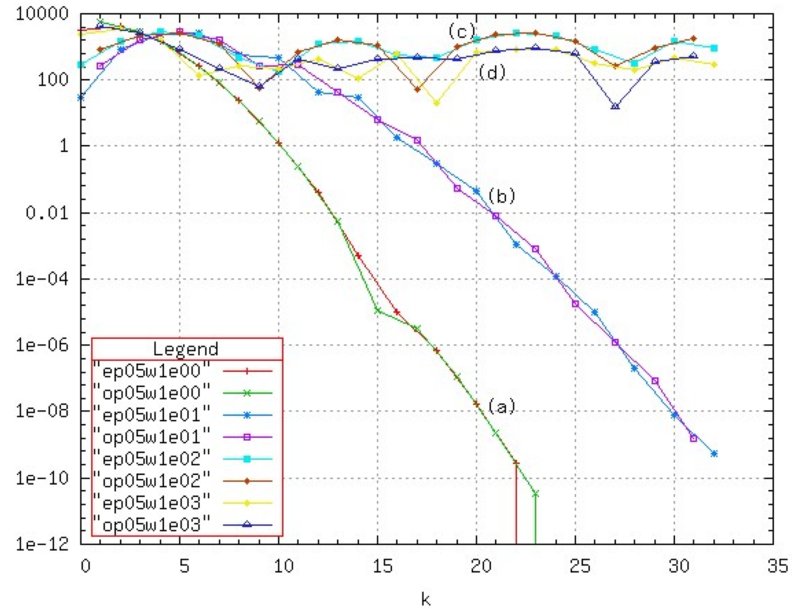

The data on the ***left figure*** were derived for the integrand $f_1(x) = |x^2 + 2x - 2|^{-1/2}: [0,1] \to \mathbb{R}$ which shows ***an inner singularity*** at $x_s = \sqrt{3} - 1$ over the specified subranges.
The data on the ***right figure*** were derived for the family of integrand functions $f_2(x) = e^{p(x-x_0)} \sin(\omega x) : [-1,1] \to \mathbb{R}$ in terms of the variable parameters p, x_0, and ω at $p = 5$ (marked as 'p05' in the file names), at fixed $x_0 = -1$ (not marked), and at the specified four ω values.
The file notations start with the specification of the rank of the Chebyshev subset: 'e' (for even) and 'o' (for odd).
Three typical integrand conditioning diagnostics are apparent:
(1) Cases (a): *well-conditioned, fast converging.*
(2) Cases (b): *well-conditioned, hopefully converging.*
(3) Cases (c) and (d): *ill-conditioned – integrand profile analysis requested to set precise diagnostic.*

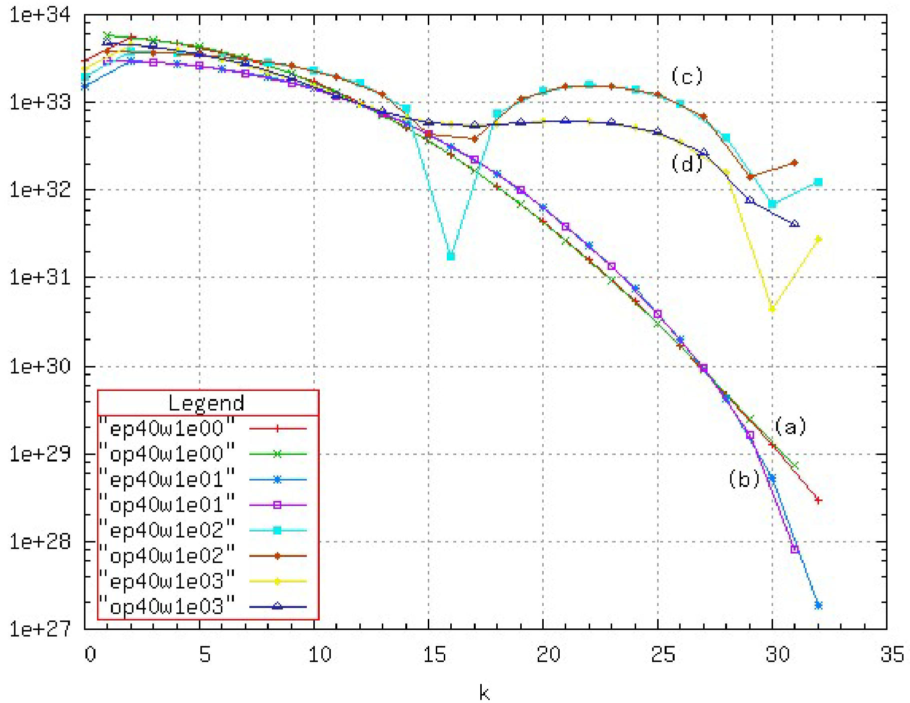

The data were derived for the family of integrand functions $f_2(x) = e^{p(x-x_0)} \sin(\omega x) : [-1, 1] \to \mathbb{R}$ in terms of the variable parameters p, x_0, and ω at $p = 40$ (marked as 'p40' in the file names), at fixed $x_0 = -1$ (not marked), and at the specified four ω values.

The file notations start with the specification of the rank of the Chebyshev subset: 'e' (for even) and 'o' (for odd).

The same three typical integrand conditioning diagnostics are apparent:

(1) Cases (a): *well-conditioned, fast converging*.

(2) Cases (b): *well-conditioned, hopefully converging*.

(3) Cases (c) and (d): *ill-conditioned – integrand profile analysis requested to set precise diagnostic*.

Overview

- General frame of the Bayesian automatic adaptive quadrature
- Reliability of accuracy specifications
- Three classes of integration domain lengths
- Improved diagnostic tools for Bayesian inference over macroscopic range lengths
- **Conclusions**

• The definition of the ***admissible output error level*** in the computed approximation Q to the given integral I involves two stages:

(i) *The user defined accuracy specifications*, through the requested absolute accuracy ε_a and the requested relative accuracy ε_r , need specific *validation checks* enforcing values within reliability bounds.

(ii) The accumulation of knowledge about the solution enables the *derivation of integrand dependent accuracy bounds* both for ε_a and ε_r, which secure the elimination of the over-computing with fake fatal round-off error diagnostic in the case of sign changing integrands.

• The numerical solution of the Riemann integrals within a BAAQ approach which *avoids the overcomputing* and secures the *minimization of the direct and hidden floating point loss of precision* is possible provided the manifold of the nonvanishing integrand domain lengths is split into three submanifolds of distinct extension ranges endowed with specific quadrature sums: ***microscopic*** *-- trapezoidal rule*, ***mesoscopic*** *-- Simpson rule*, and ***macroscopic*** *-- quadrature sums of high algebraic degrees of precision*. Many fine details are established.

• Over *macroscopic integration ranges*, the Clenshaw-Curtis (CC) quadrature provides *fast and sensitive diagnostic tools* which promote it as the best BAAQ candidate to the provision of the principal quadrature sum.

Three early Bayesian inference concerning the integrand conditioning and the expected output quality can be made:

(i) *well-conditioned integrand*, typical for an easy (or hopefully converging) integral within the standard automatic adaptive quadrature approach;

(ii) *heavily oscillatory integrand* asking for the scrutiny of the possible redefinition of the attainable output accuracy within the BAAQ approach;

(iii) *highly probable integrand ill-conditioning* asking for the activation of the integrand profile analysis procedure for the inference of precise conditioning diagnostics.

Thank you for your attention !