LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-JRNL-655515

Error Analysis in Nuclear Density Functional Theory

N. Schunck, J. D. McDonnell, J. Sarich, S. M. Wild, D. Higdon

June 5, 2014

Journal of Physics G: Nuclear and Particle Physics

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Error Analysis in Nuclear Density Functional Theory

Nicolas Schunck[1], Jordan D McDonnell[1], Jason Sarich[2], Stefan M Wild[2], and Dave Higdon[3],

[1] Physics Division, Lawrence Livermore National Laboratory, Livermore, CA 94551, USA

[2] Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA

[3] Los Alamos National Laboratory, Los Alamos, NM 87545, USA

E-mail: `schunck1@llnl.gov`

Abstract. Nuclear density functional theory (DFT) is the only microscopic, global approach to the structure of atomic nuclei. It is used in numerous applications, from determining the limits of stability to gaining a deep understanding of the formation of elements in the universe or the mechanisms that power stars and reactors. The predictive power of the theory depends on the amount of physics embedded in the energy density functional as well as on efficient ways to determine a small number of free parameters and solve the DFT equations. In this article, we discuss the various sources of uncertainties and errors encountered in DFT and possible methods to quantify these uncertainties in a rigorous manner.

PACS numbers: 21.10.-k, 21.30.Fe, 21.60.Jz, 21.65.Mn

Submitted to: *J. Phys. G: Nucl. Phys.*

1. Introduction

The past decade has seen two major developments that have had a profound impact on nuclear theory. First, progress in nuclear astrophysics has increased the need for reliable data in neutron-rich and superheavy nuclei. This data has become a critical input to reaction networks, neutron stars, and supernovae simulations [1, 2]. Large uncertainties in predictions (e.g., of the abundance of elements in the universe) have been traced to uncertainties in the properties of both ground and excited states of the nuclei involved during the formation of elements. Second, targeted programs by funding agencies have fostered the use of leadership-class computing facilities in science. In particular, the UNEDF collaboration has deployed high-performance computing (HPC) methods on various problems in nuclear theory [3].

The increased need for reliable, accurate data in systems where no experimental information is available has shined a spotlight on the predictive power of nuclear models. It has become increasingly clear that theoretical predictions must be accompanied with estimates of their error bars. Standard methods of statistics are mature by now, but often their dissemination has been hampered by their high computational cost for many problems. In this respect, the availability of ever more powerful supercomputers and concomitant spread of HPC techniques can be seen as game changers.

Among nuclear models, density functional theory (DFT) plays a unique role: it is currently the only global approach to nuclear structure that is applicable across the whole nuclear chart [4]. It thus features prominently in high-level applications of nuclear science such as nuclear fission [5], superheavy element predictions [6], fundamental symmetries [7, 8], and neutrinoless double-beta decay [9]. In all these applications, high-accuracy, high-precision DFT calculations are essential, yet relatively few attempts have been made to estimate the uncertainties underlying these calculations.

In this paper, we discuss some of the issues related to the estimation of theoretical uncertainties in nuclear DFT. In section 2, we discuss the essential features of the theory. In section 3, we analyze various sources of errors of DFT calculations, before presenting some of the techniques used to propagate these errors in model predictions in section 4. Section 5 gives a brief summary and outlook.

2. Nuclear Density Functional Theory

In the context of nuclear structure, density functional theory exists in two variants: the self-consistent mean field (SCMF) theory and the energy density functional (EDF) theory. We now discuss the main assumptions of both approaches and introduce the notation used throughout this paper.

2.1. Overview of the Theory

The first modern formulation of DFT was from the work of Hohenberg, Kohn, and Sham on electronic structure theory [10, 11]. Since then, considerable effort has been

devoted to perfecting the approach, including extensions for excited states, statistical ensembles, and time-dependent DFT; see [12, 13] for reviews. The central assumption of DFT is that a many-body system of electrons interacting via the Coulomb force can be recast into an independent-particle picture. In this picture, the energy of the interacting system is expressed as an (unknown) functional of the density of electrons $n(\boldsymbol{r})$. The DFT equations are formally identical to the Hartree-Fock (HF) equations.

Contrary to many-electron systems, the nucleus is self-bound, and the interaction is unknown a priori. Until recently the prevailing approach in nuclear structure has thus been slightly different from electronic DFT. The starting point has often been an effective nuclear force (or pseudopotential) $\hat{V}$, which is designed so that, when combined with the Hartree-Fock approximation to the many-body problem, essential properties of nuclei (e.g., saturation density, masses, radii, shell structure) are reproduced. In essence, this is similar to DFT, except that the energy density $\mathcal{H}$ is explicitly derived from the effective potential. This approach is known as the self-consistent mean field (SCMF) approach to nuclear structure [4].

Scientists have long recognized that pairing correlations are essential ingredients for explaining the structure of low-lying excited states and differences between even and odd nuclei. Given the form of the effective interaction $\hat{V}$, the most salient features of pairing are naturally incorporated in the SCMF theory by upgrading from the HF to the Hartree-Fock-Bogoliubov (HFB) approximation for the ground-state wave function. In practice, the energy density now involves an additional term, the pairing energy density, which is a functional of the pairing tensor (or pairing density) [14].

Guided by the Kohn-Sham and Hohenberg-Kohn theorems (and their generalization to atomic nuclei [15]), one may therefore decide to view the energy density $\mathcal{H}$ of the system as the fundamental unknown of the theory, inasmuch as it is a functional of the one-body density and two-body pairing tensor $\mathcal{H}[\rho, \kappa]$. This is the spirit of the energy density functional (EDF) variant of DFT. In principle, there exists a functional $\mathcal{H}_0$ that will give the exact energy of the system for the solution ρ, κ of the HFB equations and for which *all* many-body correlations will be included. Alternatively, one may decide to retain the effective interaction $\hat{V}$ as the central component of the theory and improve upon the HF and HFB approximations, typically by choosing a different ansatz for the nuclear wave function.

2.2. Model Parameters and Notations

In this paper, we view both the EDF and SCMF approaches as *models* of nuclear structure. In practice, they are thus characterized by a finite number, n_x, of model parameters $\boldsymbol{x} = (x_1, \ldots, x_{n_x})$. In nuclear DFT, model parameters are not specified by some underlying theory: they must be determined based on some experimental data. By $\boldsymbol{d} = (d_1, \ldots, d_{n_d})$ we denote the set of n_d experimental data points used to fit the model. These data points can be of different types (e.g., atomic masses, excitation energies, radii, transition strengths); see section 3.2.1.

Given the parameters $\boldsymbol{x}$, the value $y_i(\boldsymbol{x})$ of an observable computed in DFT differs from the experimental value d_i by some error ϵ_i. We collect these calculated observable values and error values in $\boldsymbol{y} = \boldsymbol{y}(\boldsymbol{x})$ and $\boldsymbol{\epsilon}$, respectively. In the following, we will use Greek indices α to denote parameters, x_α, and Latin indices to denote observables, d_i or y_i.

For the sake of completeness, we mention an important conceptual difference between the EDF and SCMF approaches to nuclear structure. In the pure EDF approach, the DFT equations are always solved at the HFB level. All information about the nucleus is assumed to be encapsulated in the energy density $\mathcal{H}$. In principle, it should include both short- and long-range correlations, symmetry restoration effects, and so forth. The EDF picture is thus characterized by a single model with n_x parameters $\boldsymbol{x}$. By contrast, the model parameters of the SCMF approach are those of the effective Hamiltonian. For a given Hamiltonian, however, a hierarchy of approximations is available, which reflect the different ansatzes for the nuclear wave function. his observation has important practical consequences when it comes to model uncertainties: if the fit of the Hamiltonian parameters $\boldsymbol{x}$ is performed at a given approximation, say HF, it is not, in principle, applicable at another level. Therefore, the fit of an effective force within the SCMF approach should always be done at each level of approximation for the ground-state wave function. Clearly, this requirement increases considerably the difficulty of the task. In the rest of this paper, we will illustrate our considerations with studies of the Skyrme functional in the context of the EDF rather than the SCMF approach.

2.3. Sources of Errors

One may distinguish three main sources of uncertainties and errors in nuclear DFT:

Model Error - The most difficult source of uncertainties to quantify comes from the choice of the energy density. For example, relativistic EDFs are specified by the number and type of mesons that will be considered and whether coupling constants will be made density dependent [16, 17]. Nonrelativistic EDFs may be derived from a Skyrme-like, zero-range, two-body effective force [18] or a Gogny-like, finite-range, two-body force [19], with or without specific terms such as tensor or generalized density dependencies [20, 21]; alternatively, one may substitute all density dependencies by effective three-body forces [22] or, following the spirit of electronic DFT, build up the EDF by coupling densities and currents up to some order [22, 23, 24]. When EDFs are derived from an effective Hamiltonian, the choice of the ansatz for the ground-state and excited-state wave functions (HF, HFB, etc.) introduces additional uncertainties. Quantifying these sources of uncertainties can be done only on an empirical basis by comparing with experimental data; see section 3.1.

Fitting Bias - The EDF itself, irrespective of its origin, always contains a number of free parameters, from about a dozen for standard Skyrme or Gogny forces up to

about 30 for the generalized Skyrme forces used in mass models [20]. In fact, even for a given EDF, there is sometimes an ambiguity about the choice of parameters: for example, the term $\hbar^2/2m$, in principle a constant, has often be adjusted along the parameters of the Skyrme EDF; see Table II in [4]; in relativistic meson-exchange EDFs, the mass of some mesons is also sometimes taken as adjustable parameter [17]. The determination of these parameters requires experimental data, an optimization algorithm, and, most often, a number of hidden assumptions when solving the DFT equations (spherical symmetry, time-reversal symmetry, etc.). Clearly, the choice of the data, the performance of the optimization algorithm, and the various hypotheses made in the fit will be a source of errors. In the past few years, the nuclear theory community has expended considerable effort to quantify the resulting uncertainties [25, 26, 27]; see section 3.2.

Numerical Implementation - Given the form of the EDF and a set of coupling constants, actual calculations are performed based on a given numerical implementation. In practice, the DFT equations can be solved in multiple ways:

- In a basis formed by, for example, the eigenstates of the harmonic oscillator (HO) [28, 29]
- In coordinate space by mesh discretization and direct numerical integration [30, 31]
- On a lattice [32]
- With finite element [33] or multiwavelet resolution analysis [34].

Each of these implementations possesses inherent, unavoidable numerical errors. Basis expansions are always truncated, possibly inducing a dependence on additional parameters (e.g., the oscillator frequency of the HO). The precision of mesh or lattice calculations is also contingent on the resolution of the underlying grid. We give additional examples in section 3.3.

3. Estimating Errors in Nuclear DFT

In this section, we provide further details about the various sources of uncertainties in nuclear DFT calculations. We begin with model errors and emphasize that they are extremely difficult to completely isolate. We then present several factors impacting the fit of a given energy density functional. We close this section with a reminder about numerical errors due to the particular implementation of the DFT equations.

3.1. Model Errors

As discussed in section 2.3, model errors in DFT can only be estimated on an empirical basis by carefully comparing predictions of selected observables obtained with different functionals. For this comparison to be meaningful, one should in principle ensure that the optimization procedure used to determine the coupling constants of the functionals

was the same for all EDFs considered and that the numerical implementation is also identical. In practice, this is rarely the case.

Mass models are an example where such a comparison is sometimes possible. The original mass model from the Bruxelles-Montréal collaboration was based on a standard Skyrme force and the BCS approximation [35]. In subsequent incarnations of the model, the HFB approximation to the pairing solution was substituted for the BCS one, more realistic pairing forces were introduced, and several phenomenological corrections were added. Although the experimental data used to perform the fit also evolved over the years, in several instances the experimental data, the optimization procedure, and the DFT solver used are identical, and only the form of the functional is different. For example, in [36], the performance of a generalized Skyrme force is compared with that of a standard Skyrme force, all things being equal.

Another (unpublished) example comes from the UNEDF collaboration and the UNEDF1 protocol [25]. This original parametrization of the Skyrme EDF was performed at the HFB approximation with an approximate restoration of particle number using the seniority limit of the Lipkin-Nogami (LN) prescription. Details of the fit itself – choice of data, form of the χ^2, numerical algorithms, and so on – can be found in the publications by the UNEDF collaboration [37, 25, 26]. The same fit as UNEDF1 was repeated at the pure HFB level, that is, without the LN prescription. All calculations were performed with the HFBTHO solver and the POUNDERS algorithm [28, 38, 39]. Results are summarized in Table 1.

Table 1. Comparison of UNEDF1 parametrization with and without the Lipkin-Nogami approximate particle number restoration. ρ_c is in fm^{-3}; E^{NM}/A, K^{NM}, $a_{\rm sym}^{NM}$, and $L_{\rm sym}^{NM}$ are in MeV; $1/M_s^*$ is dimensionless; $C_t^{\rho\Delta\rho}$ and $C_t^{\rho\nabla J}$, $t = 0, 1$ are in MeV fm^5; and V_0^n and V_0^p are in MeV fm^3.

Parameters	UNEDF1	UNEDF1-HFB
ρ_c	0.158707	0.156247
E^{NM}/A	-15.80000	-15.800000
K^{NM}	220.00000	244.839379
$a_{\rm sym}^{NM}$	28.986789	28.668204
$L_{\rm sym}^{NM}$	40.004790	40.109081
$1/M_s^*$	0.992423	1.067970
$C_0^{\rho\Delta\rho}$	-45.135131	-45.599763
$C_1^{\rho\Delta\rho}$	-145.382168	-143.935074
V_0^n	-186.065399	-234.380010
V_0^p	-206.579594	-260.437001
$C_0^{\rho\nabla J}$	-74.026333	-73.946388
$C_1^{\rho\nabla J}$	-35.658261	-51.912548

As anticipated, the major difference between the two parameterizations is in the pairing strengths, which are much larger in magnitude when the LN prescription is dropped. The two other notable changes are a sizable increase of the incompressibility K^{NM} and a sizable decrease of the scalar effective mass. The r.m.s. deviations for

the observables used in the fit are listed in Table 2. Overall, the original UNEDF1 model, based on the HFB+LN approximation, performs better than the UNEDF1-HFB model. Since the data used in the fit – the χ^2 objective function, the DFT solver and the optimization algorithms – and even the number of actual parameters, are exactly the same in both cases, it is tempting to interpret the differences as *model differences*. However, this is not exactly true: the optimization is always initialized from a given point, and one cannot dismiss a dependence of the final result on the starting point, see discussion in [40] in this Focus Issue. In fact, there is no guarantee that the minimum obtained in any of these optimizations is the absolute minimum of the 12-dimensional hypersurface. Even in this near-ideal calibration, therefore, it is almost impossible to completely disentangle *model uncertainties* and *fitting bias*. In the next section, we discuss these fitting biases in more detail.

Table 2. Root-mean-square deviations for each observable in the UNEDF1 optimization protocol compared for UNEDF1 and UNEDF1-HFB. All r.m.s. values are in MeV, except the ones for proton radii, which are in fm.

R.m.s.	UNEDF1	UNEDF1-HFB
Deformed masses	0.721	0.776
Spherical masses	1.461	1.836
Proton radii	0.016	0.022
OES neutrons	0.023	0.051
OES protons	0.080	0.075
Fission isomer	0.208	0.558

3.2. Fit of Model Parameters

In practice, parameter estimation comes down to solving an optimization problem. Given a probability distribution function for the errors $\boldsymbol{\epsilon}$, a common method of determining model parameter values is to maximize the associated likelihood function. For specific choices of the distribution, this is equivalent to a least-square fitting, where a χ^2 function of the squared errors $\boldsymbol{\epsilon}^2$ is minimized. In this section, we give a brief example of how the form of the χ^2 function may affect the least-square minimization.

3.2.1. Sensitivity on Experimental Data. An important feature of DFT is that it is a global approach to nuclear structure, which is meant to describe a variety of nuclear properties, including ground and excited states, and collective motion [4]. In order for a functional (or effective interaction) to be truly predictive, each parameter thus must be carefully constrained. Doing so could be challenging, however, because certain parameters may be especially sensitive to specific experimental data and much less sensitive to other data. This fact was pointed out in the first systematic studies of uncertainties in DFT by means of a singular value decomposition of the model parameters [41]. The authors found that nuclear binding energies were sensitive only

to 3 of the 8 parameters of a standard Skyrme functional. Similar conclusions were obtained when looking at single-particle energies in doubly magic spherical nuclei [27].

In view of these results, the strategy of nuclear mass models, where the single source of experimental data is atomic masses, may appear too restrictive: while it will provide excellent agreement with masses, it will do so at the price of having many parameters ill-constrained. The predictive power of such an (overconstrained) model for observables that are not masses should be questioned.

In a least-square fit, the usual way to incorporate different data types is to take a composite χ^2 objective function of the form

$$\chi^2(\boldsymbol{x}) = \frac{1}{n_d - n_x} \sum_{t=1}^{n_T} \sum_{j=1}^{n_t} \left(\frac{y_{tj}(\boldsymbol{x}) - d_{tj}}{\sigma_t} \right)^2 , \tag{1}$$

where n_T denotes the number of different data types, n_t the number of data points for type t, $n_d = \sum_t n_t$ the total number of data points over all types, and n_x the number of model parameters. The calculated value of data point number j of type t is denoted by y_{tj}, with d_{tj} the corresponding experimental value. Because there are different types of data, relative distances must be properly normalized by the quantity σ_t (which, in the general case, could also vary within a data type: σ_{tj}). If we viewed all the data $d_i \equiv d_{tj}$ as independent random variables following a normal distribution centered about the model, $\mathcal{N}(y_i, \sigma_i)$, then the quantity defined by Eq. (1) would follow the actual χ^2 probability distribution function. This is the reason that σ_t is interpreted as an estimate of the error on the parameter of type t. Note that experimental errors on the data $\boldsymbol{d}$ are not considered here.

3.2.2. Objective Function. As shown in Eq. (1), the objective function explicitly depends on an estimate of the choice for the initial errors σ_t of data type t, and it was shown earlier that this dependence could be significant [42]. In this section, we further illustrate this point in the realistic setting of the UNEDF1 optimization protocol. As a reference, we take the UNEDF1-HFB parametrization of Table 1 and look in particular at the standard deviation of the odd-even staggering (OES) and of the fission isomer excitation energy. These two particular data types indeed appear to be the main drivers of the UNEDF1 parametrization; see fig. 3 in [25].

Table 3 shows the r.m.s. deviations of each observable in the UNEDF1 optimization protocol when the standard deviation of the OES data varies from 0.025 MeV to 0.100 MeV by a step of 0.025 MeV. For reference, the UNEDF1-HFB solution, which was obtained with $\sigma_{OES} = 0.050$ MeV, is also shown. We observe overall large variations, especially in the reproduction of the masses and the fission isomer excitation energy. Interestingly, the UNEDF1-HFB obtained with the original weights probably gives the best compromise between all different types of observables.

The situation is analogous when we vary the standard deviation of the fission isomer excitation energy, shown in table 4. Here again, there are relatively large variations of both the deformed masses and the excitation energies. Unfortunately, the two variations

Table 3. Root-mean-square deviations for each observable in the UNEDF1 optimization protocol compared for UNEDF1-HFB for a few different values of the standard deviation σ_{OES} (in MeV) of the OES data. All r.m.s. values are in MeV, except the ones for proton radii, which are in fm.

	$\sigma_{OES} = 0.025$	UNEDF1-HFB	$\sigma_{OES} = 0.075$	$\sigma_{OES} = 0.100$
Deformed masses	0.944	0.776	2.596	0.806
Spherical masses	2.427	1.836	2.669	1.718
Proton radii	0.022	0.022	0.022	0.022
OES neutrons	0.012	0.051	0.065	0.080
OES protons	0.043	0.074	0.075	0.072
Fission isomer	0.809	0.558	0.535	0.530

are anticorrelated: improving the agreement for deformed masses requires an increase in σ_{FI}, which degrades the quality of reproduction of the excitation energies. However, the effect is nonlinear: masses are degraded by 36% when dividing σ_{FI} by a factor of 2, but improved by a mere 6% when multiplying them by a factor of 2. The sweet spot of the optimization is probably somewhere between 0.25 and 0.50 MeV.

Table 4. Same as Table 3 for variations of the standard deviation σ_{FI} (in MeV) of the fission isomer excitation energy.

	$w = 0.25$	UNEDF1-HFB	$w = 0.75$	$w = 1.00$
Deformed masses	1.057	0.776	0.748	0.730
Spherical masses	1.808	1.836	1.879	1.893
Proton radii	0.023	0.022	0.021	0.021
OES neutrons	0.057	0.051	0.044	0.042
OES protons	0.079	0.074	0.073	0.072
Fission isomer	0.279	0.558	0.794	0.903

The relatively large variability of optimization results under a change of the standard deviations σ_t for each data type is a significant source of model uncertainties. At this point, there is no magic recipe that would completely remove them. Bayesian approaches could possibly be used to provide alternative estimates of these errors.

3.3. Numerical Implementation

Of all the possible sources of uncertainties in nuclear DFT computations, numerical errors stemming from the particular implementation of DFT equations in a computer code are the easiest to quantify. The vast majority of DFT solvers are based on the expansion of HF(B) wave functions on a basis. Of particular interest is the basis made of the eigenstates of the harmonic oscillator: the HO is a decent approximation of the nuclear mean-field, at least for deeply bound states; basis functions are analytical; and, most important, it is the only example where there is an analytical separation between center of mass and relative motion in a many-body system.

When solving the HFB equations in the HO basis, several approximations can be imposed on the form of the solutions. In the case of the nonrelativistic Skyrme pseudopotential, there exist three published, open-source versions of DFT solvers assuming spherical symmetry [24], axial and time-reversal symmetry [28], or no particular symmetry at all [29]. These three solvers have been carefully benchmarked against one another and give essentially identical results. A similar package of relativistic DFT solvers has recently been published [43].

Nuclear DFT calculations in the HO basis are subject to truncation errors. Because of the finite size of the basis, results become dependent on the oscillator frequency ω_0 or, equivalently, the oscillator length b_0. When computing deformed nuclear states, it is convenient to also deform, or "stretch," the basis states to accelerate convergence: basis functions are then characterized either by the frequencies $\omega = (\omega_\perp, \omega_z)$ (cylindrical coordinates, 2D) or $\omega = (\omega_x, \omega_y, \omega_z)$ (Cartesian coordinates, 3D) or by some spherical-equivalent frequency ω_0 and one (or several) deformation parameters. See, for example, the discussion in [28].

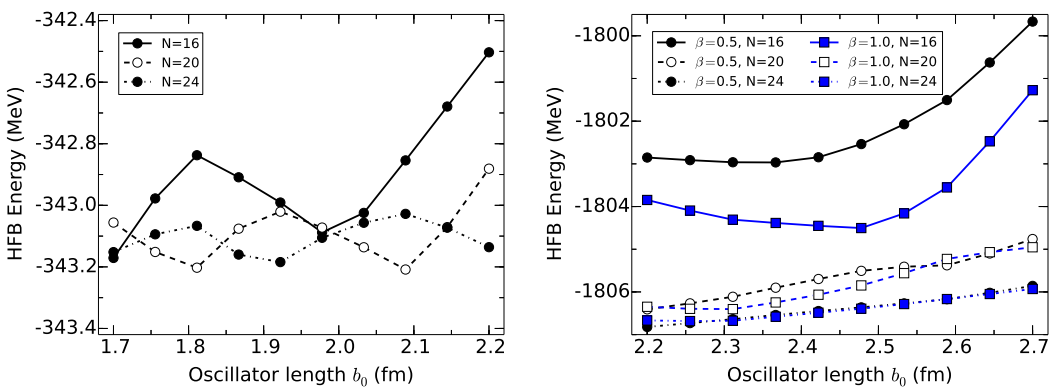

Figure 1. Comparison between the pace of convergence of DFT calculations using HO expansions. Both figures are obtained at the HFB approximation with the SkM* functional and a surface-volume pairing. Left: Total energy of ^{40}Ca as function of the oscillator length b (in fm) for different number of oscillator shells $N_{\rm shell}$. Right: Same for the configuration defined by $\langle \hat{Q}_{20} \rangle = 200$ b and $\langle \hat{Q}_{40} \rangle = 50$ b^2 in ^{240}Pu. Stretched HO bases with different deformations $\beta = 0.5$ and $\beta = 1.0$ are used.

This model dependence is illustrated in two extreme cases in fig. 1. In the left panel, we show the convergence of a simple HF calculation in the ground state of ^{40}Ca; in the right panel, we show the convergence of a deformed HFB calculation defined by $\langle \hat{Q}_{20} \rangle = 200$ b and $\langle \hat{Q}_{40} \rangle = 50$ b^2 along the fission path of ^{240}Pu. As expected, the sensitivity on basis parameters is much more pronounced in heavier nuclei and for very deformed configurations: across the range in values for b_0, $N_{\rm shell}$ and the basis deformation β shown in the figure, the total energy varies by about 0.7 MeV in ^{40}Ca and about 7 MeV in ^{240}Pu. Including pairing correlations also requires larger bases, since higher-lying states may become occupied. In mass $A > 200$, one may estimate that ground-state calculations are subject to an error greater than 1 MeV even for

large HO bases with about $N = 20$ full shells; this error may reach up to 3-4 MeV in very deformed configurations near the scission point [44]. Early attempts were made to quantify this numerical truncation error without performing a full-scale calculation [45]. Because of the nonlinearity and density dependencies of standard EDF, it is not clear whether extrapolation techniques developed in the *ab initio* community might also be applicable [46, 47].

DFT calculations performed directly in coordinate space are often considered more precise than basis calculations. We show in fig. 2 a comparison between the pace of convergence of a typical HFB calculation in two different numerical implementations: the direct numerical integration of DFT equations in coordinate space of the code HFBRAD [30] and the expansion on spherical HO basis of HOSPHE [24]. Calculations were performed on ^{120}Sn with the SkM* functional and a simple surface-volume pairing with a cut-off of $E_{\rm cut} = 60$ MeV. Since the Coulomb potential is treated differently in the two codes, it was neglected here for the sake of comparison. The rate of convergence is approximately exponential in the HO basis, while it is roughly Gaussian in coordinate space. Note that in order to achieve the same precision given by a mesh size of $h = 0.05$ fm (routinely achieved with HFBRAD), nearly 60 *full* HO shells are necessary.

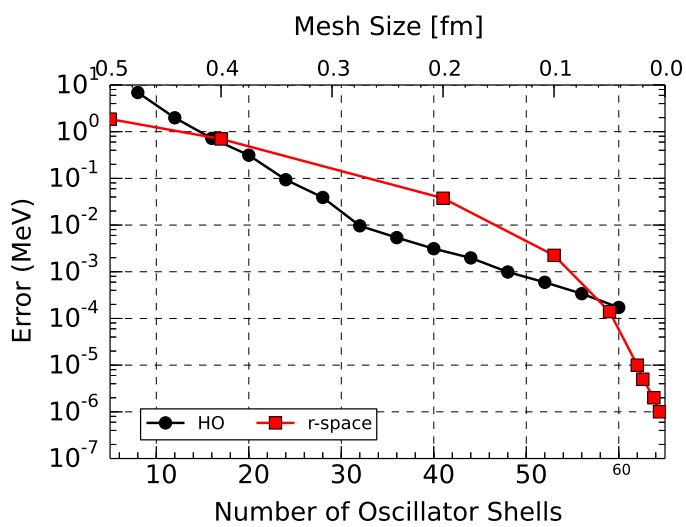

Figure 2. Comparison between the pace of convergence of a DFT calculation in coordinate-space, red squares, and configuration space (HO basis), black circles. Results were obtained by setting both direct and exchange terms of the Coulomb potentials to 0. The HO basis results are optimized with respect to the oscillator frequency. Coordinate space calculations were performed with HFBRAD in a box of 20 fm [30], HO calculations with HOSPHE [24].

The cost of directly integrating of DFT equations in coordinate space grows quickly if spherical symmetry is broken. There exist a couple of axial DFT solvers in coordinate space [31], but multicore architectures are essential for reasonable run times. To our knowledge, there is no full 3D solver in coordinate space. For such arbitrary geometries,

other representations are more promising:

- The lattice representation of the DFT solvers developed by the Bruxelles-Bordeaux-CEA collaboration is combined with the imaginary time method to solve HFB equations [32]. Its main appeal is that the numerical precision is essentially independent of the underlying geometry of the nucleus.
- In a similar spirit, multiresolution wavelet expansions of HFB wave functions guarantee, by construction, arbitrary precision for observables [34]. Originally developed for quantum chemistry applications, this technique has been recently applied to nuclei and has shown great potential for complex problems such as fission, highly excited nuclei, or nuclear reactions.
- A path hitherto neglected in nuclear structure is finite element analysis. This technique is widely employed in engineering but has only one application (published) in the context of the relativistic mean field [33].

4. Uncertainty Quantification in Nuclear DFT

Most of the discussion presented in section 3, especially in subsections 3.1-3.2, was centered on identifying sources of uncertainties in computations of nuclear properties and sampling their impact on calculations. For example, comparing predictions of different energy densities (e.g., Skyrme and Gogny) gives insights into the magnitude of model uncertainties but does not provide a rigorous metric. The purpose of uncertainty quantification in nuclear DFT is to systematically estimate error bars in calculations by deploying a variety of statistical and computational techniques.

4.1. Confidence Intervals and Error Propagation

The UNEDF collaboration was probably the first to popularize the use of traditional covariance and sensitivity methods to estimate fitting uncertainties [3]. Starting from a given parameterization $\boldsymbol{x}$ of the model, uncertainties were estimated by using confidence intervals [26]. We emphasize that several assumptions are made in practice: all the errors $y_{tj}(\boldsymbol{x}) - d_{tj}$ are independent of one another, each distributed according to a normal distribution (possibly with different variances σ_t^2). Under these conditions, confidence intervals can be computed using the central limit theorem result that the estimate $\bar{\boldsymbol{x}}$ is asymptotically normal, with mean at the true value $\boldsymbol{x}$, and $n_x \times n_x$ covariance $C(\boldsymbol{x})$. We note that the optimization process of an EDF yields only the central value $\bar{\boldsymbol{x}}$. In order to compute the covariance matrix, several additional approximations may be invoked.

The most popular of those is that the model outputs $\boldsymbol{y}(\boldsymbol{x})$ are effectively linear with respect to local variations of model parameters $\boldsymbol{x}$ around the optimized solution $\bar{\boldsymbol{x}}$. That is, we can write [48]

$$\boldsymbol{y}(\boldsymbol{x}) \approx \boldsymbol{y}(\bar{\boldsymbol{x}}) + \boldsymbol{J}(\bar{\boldsymbol{x}})(\boldsymbol{x} - \bar{\boldsymbol{x}}) \tag{2}$$

where $\boldsymbol{J}$ is the sensitivity matrix at $\bar{\boldsymbol{x}}$, $J_{\alpha i}(\bar{\boldsymbol{x}}) = \left(\dfrac{\partial y_i(\bar{\boldsymbol{x}})}{\partial x_\alpha} \right)$. Under this assumption, one can show that the covariance matrix reduces to

$$C = \sigma^2 \left(\boldsymbol{J}(\bar{\boldsymbol{x}}) \boldsymbol{J}(\bar{\boldsymbol{x}})^T \right)^{-1}. \tag{3}$$

Alternatively, one can use the Hessian of the χ^2 to estimate the covariance matrix for $\bar{\boldsymbol{x}}$

$$C = H^{-1}, \text{ where } H_{\alpha\beta}(\bar{\boldsymbol{x}}) = \left(\frac{n_d - n_x}{2} \frac{\partial^2 \chi^2}{\partial x_\alpha \partial x_\beta}(\bar{\boldsymbol{x}}) \right).$$

Uncertainties on $\bar{\boldsymbol{x}}$ induce uncertainty on the model prediction so that the variance of a computed observable $y_j(\bar{\boldsymbol{x}})$ is simply given by

$$\mathrm{Var}(y_j(\bar{\boldsymbol{x}})) = \sum_{\alpha\beta} J_{\alpha j} C_{\alpha\beta} J_{\beta j}. \tag{4}$$

This covariance technique has been recently applied to estimate the information content of the electric dipole strength [49], the correlation between electric dipole polarizability and neutron skin [50], the uncertainties on the weak charge form factor [51], and the neutron skin of neutron-rich nuclei [52].

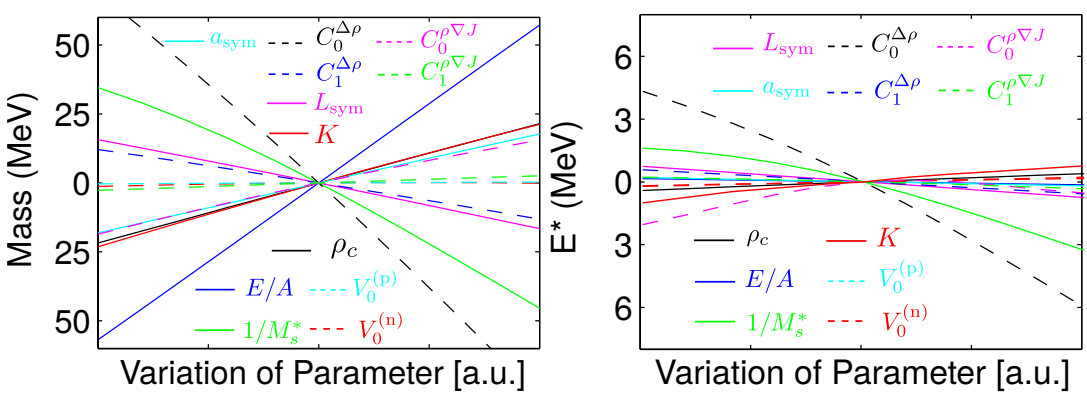

Figure 3. Variation of the mass (left panel) and the excitation energy of the fission isomer (right panel) of ^{240}Pu as a function of each of the 12 parameters of the UNEDF1 functional. All parameters have been normalized to the interval of variation listed in Table II of [25], column marked "Bounds."

The two approximations of independent errors and of linearity near the solution seem the strongest. In fact, researchers pointed out early on that such approximations were not justified [42, 27]. In fig. 3, we show the variations of the mass of the ^{240}Pu nucleus and of the excitation energy of its fission isomer as a function of each of the twelve parameters of the UNEDF1 functional. Each parameter has been normalized between 0 and 1 based on the "reasonable interval" of variation listed in [25]. We see that the mass behaves nearly linearly across the parameter range; the excitation energy of the fission isomer, however, shows some marked deviations, in particular as a function of the scalar effective mass M_s^*, the isoscalar surface, and the spin-orbit terms $C_0^{\rho\Delta\rho}$ and $C_0^{\rho\nabla J}$. These deviations are more pronounced as we go away from the local solution given by UNEDF1.

4.2. Application of Bayesian Statistics in Nuclear DFT

The Bayesian approach is based on Bayes' theorem for conditional probabilities: the probability that a continuous alternative A lies between A and $A + dA$ given B and C is given by

$$p(A|BC)dA = \frac{p(B|AC)p(A|C)dA}{\int p(B|AC)p(A|C)dA}. \tag{5}$$

Based on this theorem, the Bayesian approach consists of treating the model parameters A as genuine random variables depending on some data B and some other circumstances C. The goal is to find the probability distribution function (p.d.f.) of the random variable A, that is, the probability of having A between A and $A + dA$ given B and C. From the Bayesian perspective, uncertainty regarding the fixed, but unknown, model parameters is described with probability. In practice, Bayes' theorem is also often given as

$$\begin{aligned} &\text{posterior} \propto \text{likelihood} \times \text{prior}, \qquad \text{or} \\ &\pi(\boldsymbol{x}|d) \quad \propto \mathcal{L}(\boldsymbol{x}, d) \times p(\boldsymbol{x}). \end{aligned}$$

In the context of nuclear DFT, this equation should be interpreted as follows. We start with a prior density $p(\boldsymbol{x})$ for the model parameters $\boldsymbol{x}$. By default, one may assume a uniform distribution between some sensible, physically-motivated intervals so that $p(\boldsymbol{x}) \propto I[\boldsymbol{x} \in C]$, where C denotes the n_x-dimensional prior rectangle, $n_x = 12$ for the UNEDF1 protocol. Based on a set of experimental data $\boldsymbol{d}$, we then define our likelihood function $\mathcal{L}(\boldsymbol{x}, d)$ based on some $\chi^2(\boldsymbol{x}, \boldsymbol{d})$ function, since, for normally distributed random variables, we know that $\mathcal{L}(\boldsymbol{x}, \boldsymbol{d}) \propto p(\boldsymbol{d}|\boldsymbol{x}) \propto e^{-\chi^2(\boldsymbol{x},\boldsymbol{d})}$. The posterior distribution will be estimated by sampling from the posterior density of the parameters, $\pi(\boldsymbol{x}|d)$.

The Bayesian approach offers numerous advantages. It provides a full probabilistic description of the model parameters, allowing very general dependence between model parameters, from which one may deduce the mean, standard deviation, and covariance matrix, if desired. It can easily incorporate the impact of new data: a posterior distribution obtained from a given set of data can serve as a prior distribution if the dataset is extended (or modified). On the other hand, the computational cost of building a full p.d.f. in the context of nuclear DFT can be significant. In the case of the UNEDF parameterizations, each χ^2 function involves on the order of 100 deformed HFB calculations, each taking on the order of 5-10 minutes. The parameter space has dimension 12: in such a space, Markov chain Monte Carlo techniques, which are often used to build the posterior distribution, could easily require dozens of thousands of iterations before convergence. Various techniques can be deployed to mitigate this cost, such as the construction of metamodels (or response functions, or emulators) for the χ^2 function.

5. Conclusions

In this paper, we have discussed the various sources of errors and uncertainties in nuclear density functional theory. In particular, we have distinguished between model errors, fitting errors, and numerical implementation errors. Implementation errors are, in principle, the easiest to control, although they can become significant in specific applications such as fission, where systems become extremely elongated, or neutron-rich nuclei near or beyond the drip lines, where the coupling to continuum becomes significant. Considerable work was recently devoted to estimating and propagating fitting errors, mostly through covariance techniques. Model errors are unavoidable in the theoretical description of any quantum many-body problem, in particular in the nuclear physics case where the interaction is not known. They are by far the most difficult uncertainties to estimate, and preliminary indications are that they can be very large. We have also emphasized that Bayesian techniques may represent a promising path toward a more rigorous quantification of uncertainties in DFT. These methods are computationally costly but can be deployed in a variety of settings.

Acknowledgment

This work was partly performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. It was supported by the SciDAC activity within the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research under contract number DE-AC02-06CH11357. Computational resources were provided through an INCITE award "Computational Nuclear Structure" by the National Center for Computational Sciences (NCCS) and National Institute for Computational Sciences (NICS) at Oak Ridge National Laboratory, through an award by the Livermore Computing Resource Center at Lawrence Livermore National Laboratory, and through an award by the Laboratory Computing Resource Center at Argonne National Laboratory.

References

[1] F. Käppeler, R. Gallino, S. Bisterzo, and W. Aoki. The s process: Nuclear physics, stellar models, and observations. *Rev. Mod. Phys.*, 83(1):157, 2011.

[2] K. Langanke and G. Martínez-Pinedo. Nuclear weak-interaction processes in stars. *Rev. Mod. Phys.*, 75(3):819, 2003.

[3] S. Bogner, A. Bulgac, J. Carlson, J. Engel, G. Fann, R. J. Furnstahl, S. Gandolfi, G. Hagen, M. Horoi, C. Johnson, et al. Computational nuclear quantum many-body problem: The UNEDF project. *Comput. Phys. Comm.*, 184(10):2235, 2013.

[4] M. Bender, P.-H. Heenen, and P.-G. Reinhard. Self-consistent mean-field models for nuclear structure. *Rev. Mod. Phys.*, 75(1):121, 2003.

[5] W. Younes and D. Gogny. Nuclear scission and quantum localization. *Phys. Rev. Lett.*, 107(13):132501, 2011.

[6] S. Hofmann and G. Münzenberg. The discovery of the heaviest elements. *Rev. Mod. Phys.*, 72(3):733, 2000.

[7] J. Dobaczewski and P. Olbratowski. Solution of the Skyrme–Hartree–Fock–Bogolyubov equations in the Cartesian deformed harmonic-oscillator basis. (V) HFODD(v2.08k). *Comput. Phys. Comm.*, 167(3):214–216, 2005.

[8] S. Ban, J. Dobaczewski, J. Engel, and A. Shukla. Fully self-consistent calculations of nuclear Schiff moments. *Phys. Rev. C*, 82(1):015501, 2010.

[9] T. R. Rodríguez and G. Martínez-Pinedo. Energy density functional study of nuclear matrix elements for neutrinoless $\beta\beta$ decay. *Phys. Rev. Lett.*, 105(25):252503, 2010.

[10] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Phys. Rev.*, 136(3B):B864, 1964.

[11] W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, 140(4A):A1133, 1965.

[12] R. Eschrig. *Fundamentals of Density Functional Theory.* Teubner, Leipzig, 1996.

[13] R.G. Parr and W. Yang. *Density Functional Theory of Atoms and Molecules.* Oxford University Press, Oxford, 1989.

[14] E. Perlińska, S. Rohoziński, J. Dobaczewski, and W. Nazarewicz. Local density approximation for proton-neutron pairing correlations: Formalism. *Phys. Rev. C*, 69(1):014316, 2004.

[15] J. Messud, M. Bender, and E. Suraud. Density functional theory and Kohn-Sham scheme for self-bound systems. *Phys. Rev. C*, 80(5):054314, 2009.

[16] T. Nikšić, D. Vretenar, and P. Ring. Relativistic nuclear energy density functionals: Mean-field and beyond. *Prog. Part. Nucl. Phys.*, 66(3):519, 2011.

[17] P. Ring. Relativistic mean field theory in finite nuclei. *Prog. Part. Nucl. Phys.*, 37:193, 1996.

[18] D. Vautherin and D. M. Brink. Hartree-Fock calculations with Skyrme's interaction, I: Spherical nuclei. *Phys. Rev. C*, 5(3):626, 1972.

[19] J. Dechargé and D. Gogny. Hartree-Fock-Bogolyubov calculations with the D1 effective interaction on spherical nuclei. *Phys. Rev. C*, 21(4):1568, 1980.

[20] N. Chamel, S. Goriely, and J. M. Pearson. Further explorations of Skyrme-Hartree-Fock-Bogoliubov mass formulas, XI: Stabilizing neutron stars against a ferromagnetic collapse. *Phys. Rev. C*, 80(6):065804, 2009.

[21] T. Lesinski, M. Bender, K. Bennaceur, T. Duguet, and J. Meyer. Tensor part of the Skyrme energy density functional: Spherical nuclei. *Phys. Rev. C*, 76(1):014312, 2007.

[22] F. Raimondi, K. Bennaceur, and J. Dobaczewski. Nonlocal energy density functionals for low-energy nuclear structure. *J. Phys. G: Nucl. Part. Phys.*, 41(5):055112, 2014.

[23] F. Raimondi, B. G. Carlsson, and J. Dobaczewski. Effective pseudopotential for energy density functionals with higher-order derivatives. *Phys. Rev. C*, 83(5):054311, 2011.

[24] B.G. Carlsson, J. Dobaczewski, J. Toivanen, and P. Veselý. Solution of self-consistent equations for the N3LO nuclear energy density functional in spherical symmetry. the program hosphe (v1.02). *Comput. Phys. Comm.*, 181(9):1641, 2010.

[25] M. Kortelainen, J. McDonnell, W. Nazarewicz, P.-G. Reinhard, J. Sarich, N. Schunck, M. V. Stoitsov, and S. M. Wild. Nuclear energy density optimization: Large deformations. *Phys. Rev. C*, 85(2):024304, 2012.

[26] M. Kortelainen, T. Lesinski, J. Moré, W. Nazarewicz, J. Sarich, N. Schunck, M. V. Stoitsov, and S. Wild. Nuclear energy density optimization. *Phys. Rev. C*, 82(2):024313, 2010.

[27] M. Kortelainen, J. Dobaczewski, K. Mizuyama, and J. Toivanen. Dependence of single-particle energies on coupling constants of the nuclear energy density functional. *Phys. Rev. C*, 77(6):064307, 2008.

[28] M.V. Stoitsov, N. Schunck, M. Kortelainen, N. Michel, H. Nam, E. Olsen, J. Sarich, and S. Wild. Axially deformed solution of the Skyrme–Hartree–Fock–Bogolyubov equations using the transformed harmonic oscillator basis (II) HFBTHO v2.00d: A new version of the program. *Comput. Phys. Comm.*, 184(6):1592, 2013.

[29] N. Schunck, J. Dobaczewski, J. McDonnell, W. Satuła, J.A. Sheikh, A. Staszczak, M. Stoitsov, and P. Toivanen. Solution of the Skyrme–Hartree–Fock–Bogolyubov equations in the Cartesian deformed harmonic-oscillator basis. (VII) HFODD (v2.49t): A new version of the program.

Comput. Phys. Comm., 183(1):166, 2012.

[30] K. Bennaceur and J. Dobaczewski. Coordinate-space solution of the Skyrme–Hartree–Fock–Bogolyubov equations within spherical symmetry: The program HFBRAD (v1.00). *Comput. Phys. Comm.*, 168(2):96, 2005.

[31] J. Pei, M. Stoitsov, G. Fann, W. Nazarewicz, N. Schunck, and F. Xu. Deformed coordinate-space Hartree-Fock-Bogoliubov approach to weakly bound nuclei and large deformations. *Phys. Rev. C*, 78(6):064306, 2008.

[32] P. Bonche, H. Flocard, and P. H. Heenen. Solution of the Skyrme HF + BCS equation on a 3D mesh. *Comput. Phys. Comm.*, 171(1):49, 2005.

[33] W. Pöschl, D. Vretenar, A. Rummel, and P. Ring. Application of finite element methods in relativistic mean-field theory: spherical nucleus. *Comput. Phys. Comm.*, 101(1–2):75, 1997.

[34] G. I. Fann, J. Pei, R. J. Harrison, J. Jia, J. Hill, M. Ou, W. Nazarewicz, W. A. Shelton, and N. Schunck. Fast multiresolution methods for density functional theory in nuclear physics. *J. Phys.: Conf. Ser.*, 180(1):012080, 2009.

[35] F. Tondeur, S. Goriely, J. M. Pearson, and M. Onsi. Towards a Hartree-Fock mass formula. *Phys. Rev. C*, 62(2):024308, 2000.

[36] S. Goriely, N. Chamel, and J. M. Pearson. Hartree-Fock-Bogoliubov nuclear mass model with 0.50 MeV accuracy based on standard forms of Skyrme and pairing functionals. *Phys. Rev. C*, 88(6):061302, 2013.

[37] M. Kortelainen, J. McDonnell, W. Nazarewicz, E. Olsen, P.-G. Reinhard, J. Sarich, N. Schunck, S. M. Wild, D. Davesne, J. Erler, and A. Pastore. Nuclear energy density optimization: Shell structure. *Phys. Rev. C*, 89(5):054314, 2014.

[38] T. Munson, J. Sarich, S. M. Wild, S. Benson, and L. Curfman McInnes. TAO 2.0 users manual. Tech. Memo. ANL/MCS-TM-322, Argonne Nat. Lab., Argonne, IL, 2012.

[39] S. M. Wild. Solving derivative-free nonlinear least squares with POUNDERS. Preprint ANL/MCS-P5120-0414, Argonne Nat. Lab., Argonne, IL, April 2014.

[40] J. Sarich, S. Wild, and N. Schunck. Derivative-free optimization methods for parameter estimation in computational nuclear physics. 2014.

[41] G. Bertsch, B. Sabbey, and M. Uusnäkki. Fitting theories of nuclear binding energies. *Phys. Rev. C*, 71(5):054311, 2005.

[42] J. Toivanen, J. Dobaczewski, M. Kortelainen, and K. Mizuyama. Error analysis of nuclear mass fits. *Phys. Rev. C*, 78(3):034306, 2008.

[43] T. Nikšić, N. Paar, D. Vretenar, and P. Ring. DIRHB–A relativistic self-consistent mean-field framework for atomic nuclei. *Computer Physics Communications*, 185(6):1808, 2014.

[44] N. Schunck. Density functional theory approach to nuclear fission. *Acta Phys. Pol. B*, 44:263, 2013.

[45] S. Hilaire and M. Girod. Large-scale mean-field calculations from proton to neutron drip lines using the D1S Gogny force. *Eur. Phys. J. A*, 33(2):237, 2007.

[46] R. J. Furnstahl, G. Hagen, and T. Papenbrock. Corrections to nuclear energies and radii in finite oscillator spaces. *Phys. Rev. C*, 86(3):031301, 2012.

[47] S. A. Coon, M. I. Avetian, M. K. G. Kruse, U. van Kolck, P. Maris, and J. P. Vary. Convergence properties of ab initio calculations of light nuclei in a harmonic oscillator basis. *Phys. Rev. C*, 86(5):054002, 2012.

[48] J. Dobaczewski, W. Nazarewicz, and P.-G. Reinhard. Error estimates of theoretical models: A guide. *J. Phys. G: Nucl. Part. Phys.*, 41:074001, 2014.

[49] P.-G. Reinhard, J. Piekarewicz, W. Nazarewicz, B. K. Agrawal, N. Paar, and X. Roca-Maza. Information content of the weak-charge form factor. *Phys. Rev. C*, 88(3):034325, 2013.

[50] J. Piekarewicz, B. K. Agrawal, G. Colò, W. Nazarewicz, N. Paar, P.-G. Reinhard, X. Roca-Maza, and D. Vretenar. Electric dipole polarizability and the neutron skin. *Phys. Rev. C*, 85(4):041302, 2012.

[51] P.-G. Reinhard and W. Nazarewicz. Information content of the low-energy electric dipole strength:

Correlation analysis. *Phys. Rev. C*, 87(1):014324, 2013.

[52] M. Kortelainen, J. Erler, W. Nazarewicz, N. Birge, Y. Gao, and E. Olsen. Neutron-skin uncertainties of Skyrme energy density functionals. *Phys. Rev. C*, 88(3):031305, 2013.