EPJ Web of Conferences **302**, 07015 (2024) https://doi.org/10.1051/epjconf/202430207015
SNA + MC 2024

A Compact Model for Unresolved Resonance Region Cross Sections

Gavin Ridley[1,2,*], *Timothy* Burke[3], and *Benoit* Forget[1]

[1]Massachusetts Institute of Technology
[2]Aalo Atomics
[3]Los Alamos National Laboratory

Abstract. For decades, the probability table technique has been used to model the behavior of neutron cross sections in the unresolved resonance region with a stochastic approach. Early work from Hwang hinted at the possibility of a universal probability distribution to fit the URR cross section distribution. In this work we present a solution to Hwang's problem: the normal inverse Gaussian distribution fits the probabilistic behavior of the total cross section almost perfectly for the most important nuclides, both in asymptotic behavior and around the more probable regions.

1 Introduction

1.1 The Unresolved Resonance Region

Where resonances cannot be conveniently measured for a nuclide, evaluators resort to modeling unknown resonances stochastically. This method has proven successful in improving the agreement of MC codes with experiments for decades by correctly capturing self-shielding effects, but the same fundamental approach has been employed: probability tables [1]. In this work, we present what an alternative to probability tables (PT) which does not rely on tabulation but rather continuous fitting. We leverage the fact that a certain distribution fits the distribution of cross section behavior quite accurately, and describe our methods for fitting the relevant parameters to nuclear data preprocessing code.

Probability tables are created by sampling the statistics of resonance locations and widths then binning the probability density of σ_t and conditional expectations of its partials. NJOY's [2] PURR module can build probability tables for MC codes. More generally, deterministic some codes employing the subgroup approach may also employ PT, as described by [3] and implemented in their CALENDF code. Using PT enhances the accuracy of continuous energy Monte Carlo calculations to eliminate biases typically around 100-200 pcm in intermediate spectrum problems that occur when neglecting oscillation of the cross section in the URR range, as shown in [4, 5] among other works. The typical behavior of the cross section in the unresolved resonance region has been depicted in Fig. 1.

*e-mail: gavin@aalo.com

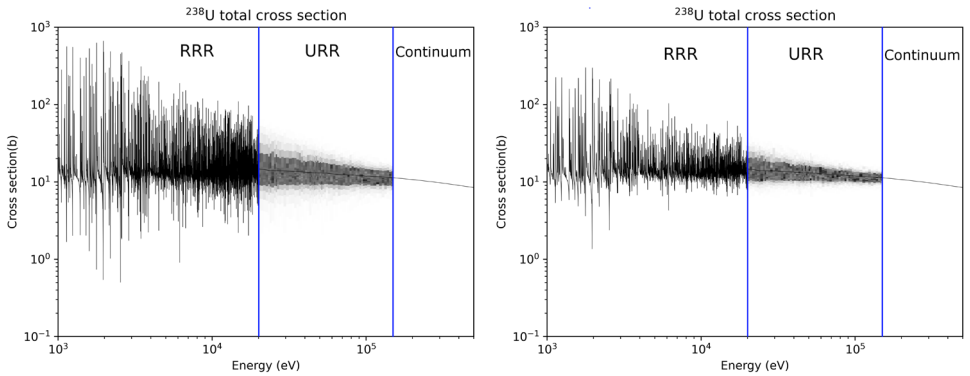

(a) 294K (b) 2500K

Figure 1: The unresolved resonance region for ^{238}U begins around 20 keV and continues to about 500 keV. The influence of temperature in the URR can be observed here; Doppler broadening of resonances results in narrowing of the total cross section distribution. Darker grey indicates a higher probability of encountering that value of σ_t.

The elastic scattering cross section that can be sampled after sampling artificial resonance locations and widths is:

$$\sigma_n(E) = \sigma_{\text{pot}} + \sum_{l=0}^{N_l-1} \sum_{j=1}^{N_j(l)} \sum_{\lambda=1}^{N_\lambda} \sigma_\lambda \left(\psi(\theta_\lambda, x_\lambda(E)) \left[\cos(2\phi_l) - \left(1 - \frac{\Gamma_{n,\lambda}}{\Gamma_\lambda} \right) \right] + \chi(\theta_\lambda, x_\lambda(E)) \sin(2\phi_l) \right) , \tag{1}$$

and we refer the reader to [4] for clarification of the terms appearing here and the other contributions to the total cross section.

The total cross section is found by summing the elastic, absorption, and fission cross sections. Notably, NJOY does not include the optional fourth competitive (often inelastic scattering) reaction in the sum. While in principle a large number of reaction types could be included, only data for the elastic, absorption, fission, and an optional fourth channel are given in ENDF [6]. The fourth channel acts to change the distribution of the Γ_λ terms only, since it is not added in to the total cross section.

Two things are required to model the URR region: a way of representing the probability distribution of σ_t and a way of representing $\mathbb{E}[\sigma_x|\sigma_t]$, where $x \in \gamma, f$. Both models depend on temperature. Our new method consists of an efficient way of accomplishing both tasks. The conventional probability table approach uses a histogram to represent the σ_t distribution and piecewise linear interpolation for $\mathbb{E}[\sigma_x|\sigma_t]$.

Hwang once suggested that simple probability distributions may be able to fit the URR region behaviors [7]. His observation was that when plotting the probability density on semilogarithmic scale, the tails tend to each form a straight line. Hwang proposed that the hypoexponential distribution may be the right choice, and found reasonable fits for a few nuclides. In our experience, while the hypoexponential approach may yield reasonable fits for a few nuclides and exhibits the requisite linear tails in log probability space, it fails to be practically usable for modelling URR effects due to poor fits for a variety of nuclides. It was precisely this same linear-tailed log probability behavior, also exhibited by the log-log particle size distribution of wind-blown sands, that inspired Barndorff-Nielsen to formulate the hyperbolic distribution [8].

1.2 A Reasonable Form of the Distribution of σ_t For All Regimes

1.2.1 Applicability of the NIG Distribution

The hyperbolic distribution meets the criteria Hwang had proposed would be required to model this situation well, but we found that another member of Barndorff-Nielsen's family of distributions that generalized the hyperbolic stands out to model the unresolved resonance region cross section distribution particularly well: the normal inverse Gaussian (NIG) distribution [9]. It, too, has the correct asymptotically linear in log probability density tail behavior. We will expose some more rigorous asymptotic arguments in favor of this distribution in a future journal article, and can only summarize the main points here.

We have identified three key behaviors that the distribution of the total cross section may present. Each of these behaviors tend to become more or less prominent depending on the ratio of the mean resonance width to spacing. The behaviors are:

- For very widely spaced resonances, the σ_t distribution closely resembles a truncated Cauchy distribution
- For intermediate resonance spacing, the tails of the σ_t distribution have linear log-probability density behavior
- Tight resonance spacings lead to near normality of the σ_t distribution, seen in ^{235}U

The NIG distribution [9] can meet these three criteria with four free parameters. It is defined by the density function:

$$\mathcal{NIG}(x;\mathfrak{a},\mathfrak{b},\mu,\delta) = \frac{\mathfrak{a}\delta K_1\left(\mathfrak{a}\sqrt{\delta^2+(x-\mu)^2}\right)}{\pi\sqrt{\delta^2+(x-\mu)^2}}\exp\left(\delta\gamma+\mathfrak{b}(x-\mu)\right) \quad , \tag{2}$$

where $\gamma = \sqrt{\mathfrak{a}^2-\mathfrak{b}^2}$. This distribution offers the computational benefit of being sampled easily, despite having a superficially complex density function.

1.2.2 Sampling The NIG Distribution

The whole family of generalized hyperbolic distributions can be interpreted as infinite mixtures of normal random variables with variances distributed by generalized inverse Gaussian distributions [9]. Specifically, samples z which have $z \sim \mathcal{IG}(\delta(\mathfrak{a}^2-\mathfrak{b}^2)^{-1/2}, \mathfrak{a}^2-\mathfrak{b}^2)$ which are used to sample $x \sim \mathcal{N}(\mu+\mathfrak{b}z, z)$ result in $x \sim \mathcal{NIG}(\mathfrak{a},\mathfrak{b},\mu,\delta)$. Equivalently,

$$\mathcal{NIG}(x;\mathfrak{a},\mathfrak{b},\mu,\delta) = \int_0^\infty \mathcal{IG}(z;\delta/\gamma,\delta^2)\mathcal{N}(x;\mu+\beta z,z)\,\mathrm{d}z \quad . \tag{3}$$

This fact can be combined with the efficient sampling procedure for inverse Gaussian distributed variables presented in [10] to efficiently sample a NIG-distributed random variable. Algorithm 1 presents the combined IG sample and variance mixture sampling algorithm. As a rule of thumb, shallow branches that fit in the instruction cache can run efficiently on a modern GPU, so we can expect this algorithm to outperform PTs in that context.

2 Results

2.1 Generation of a Finely Resolved URR Library

We generated a highly resolved library of σ_t distributions for every energy point of every nucleus exhibiting unresolved resonance behavior in ENDF/B-VIII.0 [11]. Two hundred bin

Input : $\mathfrak{a}, \mathfrak{b}, \mu, \delta$
$\gamma \leftarrow \sqrt{\mathfrak{a}^2 - \mathfrak{b}^2}$;
Sample $w_0 \sim \mathcal{N}(0, 1)$;
$w \leftarrow \delta w_0^2/\gamma$;
$c \leftarrow (2\delta\gamma)^{-1}$;
$x \leftarrow \delta/\gamma + c\left(w - \sqrt{w(4\delta^2 + w)}\right)$;
Sample $u \sim$ Unif([0, 1]);
if $u \geq \delta/(\gamma(\delta/\gamma + x))$ **then**
  $z \leftarrow \delta^2/(x\gamma^2)$;
else
  $z \leftarrow x$;
end
return sample from $\mathcal{N}(\mu + \beta z, z)$;

Algorithm 1: Algorithm to sample a variable distributed as $\mathcal{NIG}(\mathfrak{a}, \mathfrak{b}, \mu, \delta)$.

histograms were used to tally the total cross section probability density and conditional expectations of the partial cross section given the total. In contrast, default OpenMC nuclear data processing options [12] use NJOY to build a twenty bin histogram, with sixty-four independent resonance ladders. While NJOY samples the ladder at ten thousand uniformly sampled energy points, we sampled each ladder at one hundred thousand random points using an in-house code. Our code still employed sixty-four independent ladders for each energy point of each nuclide.

Creating a new program devoted to highly resolved unresolved resonance PT generation proved necessary. On an AMD Ryzen 7950X CPU, NJOY sampled the full set of PTs for ^{238}U over six different temperatures in 1593 s. Assuming optimal scaling [1], NJOY would take 15,930 s to build the highly sampled PTs that our code generated. In contrast, our code took 0.1 s to sample the resonance locations and sort the energy sample grid of each ladder for ^{238}U, with 2.5 s spent reconstructing the cross section on 100,000 points for each of the six temperatures. The six temperatures are 250, 294, 600, 900, 1200, and 2500 kelvin. To achieve this high performance, we reconstruct the cross section using a CUDA kernel called from Python.

2.2 Fitting the σ_t Model

To fit the densities, we minimize the Hellinger distance as approximated by numerical integration over our finely divided histograms. The Hellinger distance can be understood as proportional to the root mean square difference between the square roots of two density functions. The use of Hellinger distance over standard least squares minimization between the density functions provides a desirable balance between minimizing error in the most probable regions and properly modeling the behavior in the tails of the probability distribution. In contrast, we found minimizing the Kullback-Leibler divergence excessively prioritized matching the values in the tails of the distribution, and naive least-squares minimization between the density function difference excessively prioritized matching the densities around the distribution's mode.

The nonlinear optimization problem presented by minimizing the Hellinger distance was solved with the Nelder-Mead method wrapped by `scipy` [14]. We found that applying bounds to the NIG parameters improved the reliability of the optimizer substantially, named

[1] [13] points out that the energy grid sorting algorithm of NJOY is $O(n^2)$, so the actual number would be higher.

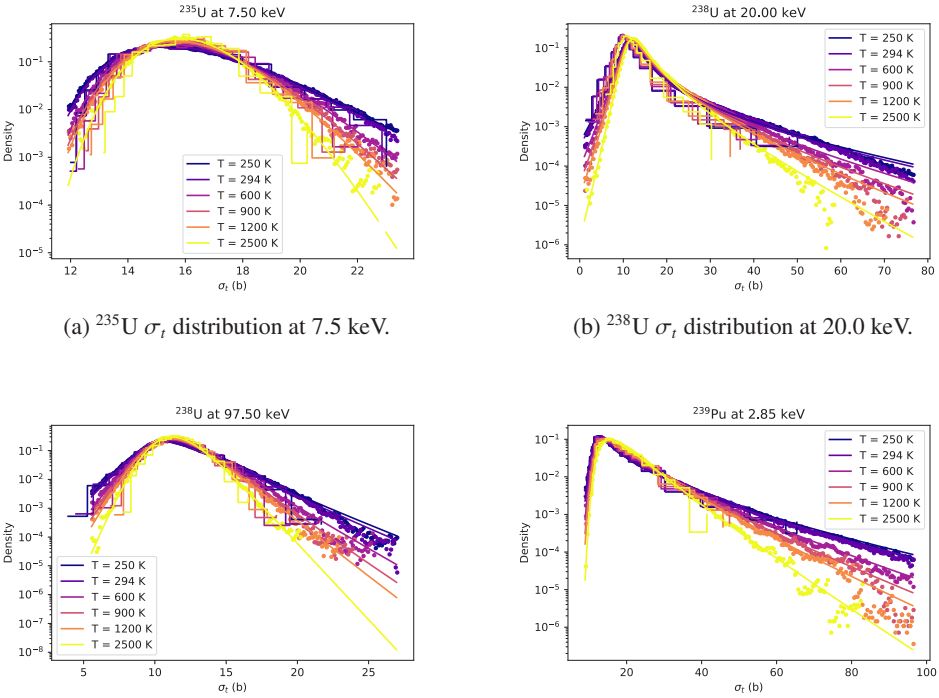

(a) ^{235}U σ_t distribution at 7.5 keV.

(b) ^{238}U σ_t distribution at 20.0 keV.

(c) ^{238}U σ_t distribution at 97.5 keV.

(d) ^{239}Pu σ_t distribution at 2.85 keV.

Figure 2: Comparison of pointwise σ_t probability density estimates from our simulation (points), fits to NIG distributions (smooth lines), and NJOY PURR results (stairsteps).

that $\alpha > 0$, $\beta > 0$, $\delta^2 > 0$, and $\mu > \min_i \sigma_{t,i}$, where $\sigma_{t,i}$ represents the realizations of the total microscopic cross section. The bound $\beta > 0$ strictly speaking is not required of the NIG distribution; it is instead based on our observation that the σ_t is exclusively skewed to large values and never otherwise.

Still, the nonlinear optimizer crashed on many nuclides due to NaN being generated. These were found to result from the multiplication of underflowed zeros and overflowed infinities respectively generated by the K_1 term of Eq. 2 and the exponential term. The code must branch to handle this asymptotic case separately. We chose the criterion for the asymptotic branch as any argument to the exponential function over seven hundred.

A select few results for nuclides of importance to reactor simulation are presented in Fig. 2 and 3. The plots show some representative results of obvious importance for ^{235}U, ^{238}U, ^{239}Pu, and ^{240}Pu. ^{107}Ag has been included due to its use in control rods and natural abundance around 52%. ^{137}Cs has been presented due to its prominence in radiological source terms. We include some results for ^{55}Mn due to it constituting all natural manganese; we have observed that in both our code and in NJOY, the traditionally used single-level Breit-Wigner approach yields a large fraction of negative cross sections. Notably, in Fig. 3c we have not zeroed out the negative values in order to show the fit of our model to the raw data; negative samples are treated in the same manner as NJOY.

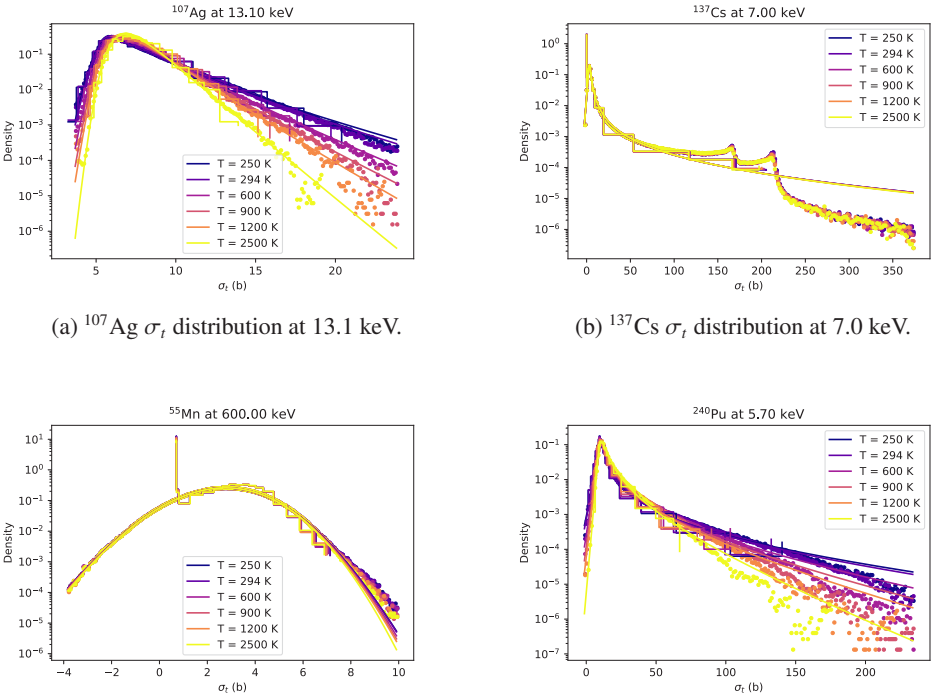

(a) ^{107}Ag σ_t distribution at 13.1 keV.

(b) ^{137}Cs σ_t distribution at 7.0 keV.

(c) ^{55}Mn σ_t distribution at 600 keV.

(d) ^{240}Pu σ_t distribution at 5.70 keV.

Figure 3: Comparison of pointwise σ_t probability density estimates from our simulation (points), fits to NIG distributions (smooth lines), and NJOY PURR results (stairsteps).

2.3 Fitting the Partial Cross Section Model

As we will discuss in the upcoming associated journal paper, the conditional expectations of the partial cross section components suffice to obtain properly self-shielded fluxes in a Monte Carlo simulation. Consequently, we must obtain a curve fit that outputs the expected partial cross section of interest given a value of σ_t. In order to yield a clear benefit over the traditional PTs approach, the model should require a fairly small amount of data yet approximate the cross section curve accurately.

Ideally, a physically motivated expression for $\mathbb{E}[\sigma_x|\sigma_t]$ could be obtained. Our preliminary investigation, under some crude assumptions, suggested that the absorption and fission partial cross sections might be approximated accurately using a rational polynomial expression. Regardless, rational polynomial functions have universal approximation properties to analytic and more generally Lipschitz-continuous functions [15], so further justification of this approach is unnecessary.

Rational polynomial fitting is a nontrivial research problem itself, and algorithms specialized for it typically must be used as discussed in [16]. The recently presented adaptive Antoulas-Anderson (AAA) algorithm [17] is, in our opinion, the most straightforward and powerful means of rational polynomial fitting at present. The strength of this algorithm orig-

inates from its representation of rational functions in rational barycentric form,

$$r(z) = \frac{n(z)}{d(z)} = \frac{\sum_{j=1}^{m} \frac{w_j f_j}{z - z_j}}{\sum_{j=1}^{m} \frac{w_j}{z - z_j}} \quad , \tag{4}$$

where $n(z)$ and $d(z)$ are polynomials, and z_j are evaluation points to evaluate function values f_j at. The coefficients of the polynomials $n(z)$ and $d(z)$ need not be computed. The work [18] discusses the numerical advantages that the barycentric rational form poses. For a certain choice of w_j, the classical Lagrange interpolant is obtained. For the choice $w_j = \{1, -1, 1, -1, \ldots\}$, a smooth rational interpolant is obtained [19].

Our initial attempts at fitting to the conditional expectations of the partial cross sections using general rational polynomials of the direct rational form $n(z)/d(z)$ were clear failures. The coefficients in the numerator and denominator can have enormous differences in their influence on the fit quality: for example high order monomial coefficients clearly vary the fit more than the lower order parts. Using a Chebyshev basis for each polynomial might have remediated this issue, but a nonlinear fitting approach using the barycentric rational form offers superior numerical characteristics and interpretability.

An initial guess to a nonlinear fitting problem like this can accelerate computations appreciably, especially considering the problem of curve fitting at every energy point, at each temperature, for each nuclide. Barycentric rational forms furnish a straightforward approach to this. The number of summation terms m is decided ahead of time in building the library; $m = 5$ appears to be more than enough for the problem at hand. Next, five points z_j have to be chosen. Because the data at hand comes from conditional expectations within each bin of a histogram of the total cross section, the points z_j can be distributed evenly in σ_t probability.

To find the points, the cumulative distribution function approximated by summing histogram bins can be found, calling this function $C(\sigma_t)$. Then, the z_j are $z_j = C^{-1}(j(\max \hat{\sigma}_t - \min \hat{\sigma}_t)/(m-1) + \min \hat{\sigma}_t)$. The interpolation points thus cluster around the most likely values of σ_t. The starting guess values for the partial cross section, f_j, are the conditional expectations of partial cross sections from each histogram bin. The weights are set to Berrut's set, $\{1, -1, 1, \ldots\}$ [19].

At this point, another layer of complexity can be added. The weights and interpolation points can be shared among all the temperatures data must be tabulated at, while the interpolation values f_j can vary with respect to temperature. We find fits by solving:

$$\min_{\vec{\sigma}_t, \vec{w}, \vec{\sigma}_x(T_t)} \sum_{t=1}^{n_t} \sum_{i}^{n_\sigma} \tilde{P}(\sigma_{t,i})^2 \left(\mathbb{E}\left[\sigma_x(T_t, \sigma_{t,i}) | \sigma_{t,i} \right] - b(\vec{\sigma}_t, \vec{w}, \vec{\sigma}_x(T_t), \sigma_{t,i}) \right)^2 \tag{5}$$

where $\sigma_x(T_t, \sigma_{t,i})$ is the random partial cross section on the histogram bin i. $b(\vec{\sigma}_t, \vec{w}, \vec{\sigma}_x(T_t), x)$ is the barycentric rational approximation formed by using interpolation points $\vec{\sigma}_t$, weights $\vec{w}$, and interpolation values $\sigma_x(T_t)$, evaluated at the point x. The outer sum in Eq. 5 is over temperature points T_t, and the inner sum is over $\sigma_{t,i}$, the midpoints of each histogram bin. Notably, the sum is weighted by the square of the probability density function estimated on the histogram. Squaring is not *per se* necessary, but we found the resulting fits to be most aesthetically appealing if using this choice. The black-box Nelder-Mead optimizer from `scipy` provided reasonable results in a reasonable amount of compute time. The use of the barycentric rational form appears to transform the problem to a space of coordinates which is reasonably well-conditioned and lacking in local optima, as our experience using black-box solvers to solve Eq. 5 has been generally positive. The same cannot be said of directly using a rational polynomial form.

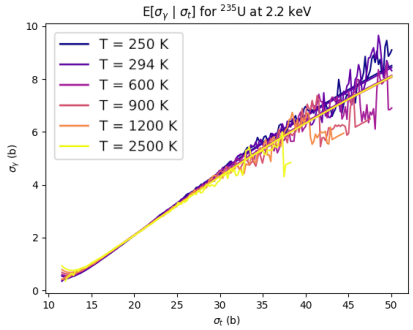

(a) ^{235}U conditional expectation of σ_γ given σ_t function at 2.2 keV.

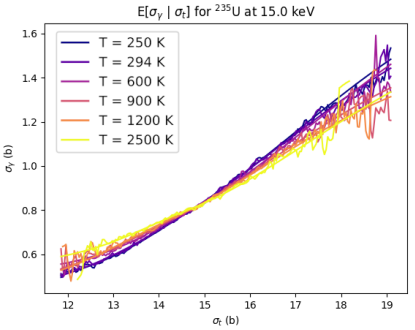

(b) ^{235}U conditional expectation of σ_γ given σ_t function at 15 keV.

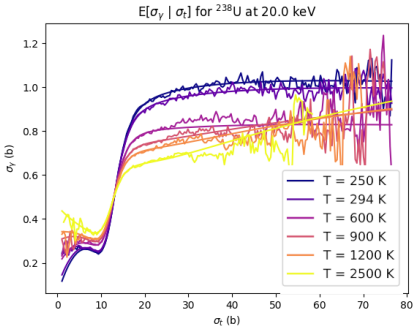

(c) ^{238}U conditional expectation of σ_γ given σ_t function at 20 keV.

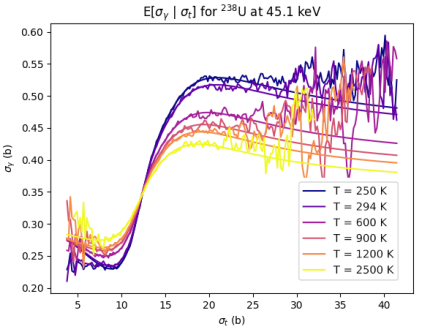

(d) ^{238}U conditional expectation of σ_γ given σ_t function at 45.1 keV.

Figure 4: The $\mathbb{E}[\sigma_\gamma|\sigma_t]$ curves for ^{235}U and ^{238}U at two energies. The jagged lines are estimates from our simulation, and the smooth lines are fits obtained by solving Eq. 5.

After obtaining a fit, the parameters of this model offer excellent interpretability. The interpolation points are where the conditional expectation curve is evaluated, and the interpolation values again are easy to interpret. On the other hand, interpreting the weights is a bit more opaque. At least, the majority of fit parameters have a clear physical meaning.

Because σ_t is defined as the sum of a capture, elastic scattering, and fission component, we only need to store the fits to the capture and fission components. In sampling a value of σ_t and its associated partials, the elastic scattering component can be found by subtracting the capture and fission parts from the total.

Figure 4 shows the absorption partial cross section conditional expectation as a function of the total cross section at two different energies for ^{235}U and ^{238}U. With only a five term barycentric rational approximation, we can see excellent agreement with the unresolved resonance cross section simulation from our simulation. Similarly, for the fission component an excellent match can also be observed in Fig. 5.

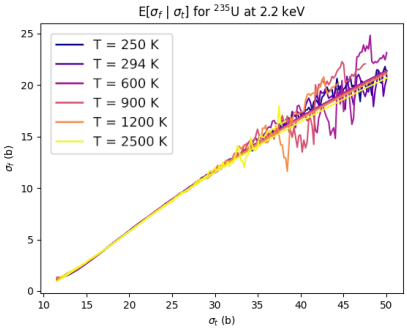

(a) ^{235}U conditional expectation of σ_f given σ_t function at 2.2 keV.

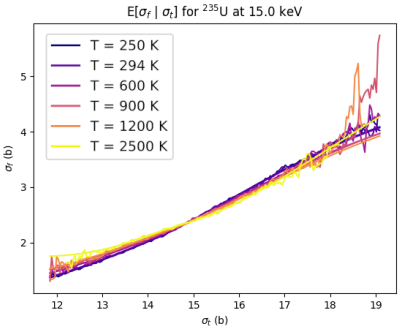

(b) ^{235}U conditional expectation of σ_f given σ_t function at 15 keV.

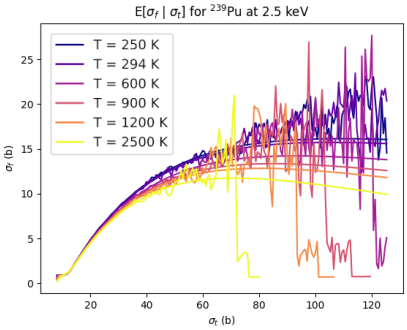

(c) ^{239}Pu conditional expectation of σ_f given σ_t function at 2.5 keV.

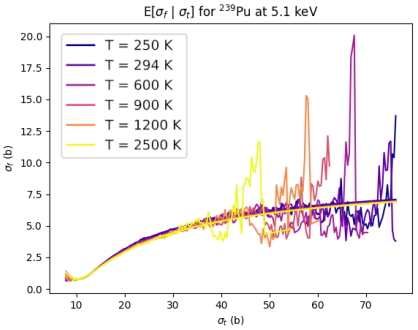

(d) ^{239}Pu conditional expectation of σ_f given σ_t function at 5.1 keV.

Figure 5: The $\mathbb{E}[\sigma_f|\sigma_t]$ curves for some key nuclides. Again, the jagged lines are estimates from the simulation, and the smooth lines are fits obtained by solving Eq. 5.

3 Conclusion

We have presented a new means of modeling cross section behavior in the unresolved resonance region that offers a few improvements. Firstly, the method consists of a continuous fit to the cross section distribution that rivals or exceeds the accuracy of the typically employed twenty bin histograms. The method also leverages the structure of the $\mathbb{E}[\sigma_\gamma|\sigma_t]$ fitting problem to express this function with a low number of degrees of freedom; barycentric rational fits allow shared interpolation points and weights across all temperatures.

From a data compression point of view, comparing to a twenty bin histogram, the new method uses 20% as many free parameters. With fitting of the temperature effects, the method could use as low as 5% as much space. This may reduce the penalty of refining the cross section grid in the unresolved region considerably. Moreover, because the data access pattern to sample from the σ_t distribution is fixed each time, as opposed to the table search required by PTs, memory divergence on GPU-based MC codes can be reduced. This may ameliorate the GPU performance penalty of PTs observed in [20]. With less free parameters, covariance matrices on the fit parameters can be more easily be obtained, too.

We used independent fits of the σ_t distribution parameters at each temperature, but based on our empirical observations, it seems to be the case that the $\mathfrak{a}$ and $\mathfrak{b}$ parameters in the NIG distribution vary linearly with temperature. This would straightforwardly enable continuous temperature dependence of cross sections in the URR, a clear benefit for multiphysics applications. Future work will present comparisons on criticality benchmarks and performance in GPU-based MC code.

References

[1] L.B. Levitt, Nuclear Science and Engineering **49**, 450 (1972)
[2] R. Macfarlane, D.W. Muir, R.M. Boicourt, I.I.I. Kahler, J.L. Conlin, Tech. Rep. LA-UR-17-20093, Los Alamos National Lab. (LANL), Los Alamos, NM (United States) (2017)
[3] J.C. Sublet, P. Ribon, Journal of Nuclear Science and Technology **39**, 856 (2002)
[4] J. Walsh, B. Forget, K. Smith, B. Kiedrowski, F. Brown, *Direct, On-the-fly Calculation of Unresolved Resonance Region Cross Sections in Monte Carlo Simulations*, in *Proc. Joint Int. Conf. on Mathematics and Computation, Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method* (2015)
[5] A. Jimenez carrascosa, E. Fridman, N. Garcia herranz, F. Alvarez valarde, P. Romojaro, F. Bostelmann, Tech. rep., Oak Ridge National Lab. (ORNL), Oak Ridge, TN (United States) (2019)
[6] A. Trkov, D.A. Brown, Tech. Rep. BNL-203218-2018-INRE, Brookhaven National Lab. (BNL), Upton, NY (United States) (2018)
[7] R.N. Hwang, Annals of Nuclear Energy **9**, 31 (1982)
[8] O. Barndorff-Nielsen, **5**, 151 (1978), publisher: [Board of the Foundation of the Scandinavian Journal of Statistics, Wiley]
[9] O.E. Barndorff-Nielsen, Scandinavian Journal of Statistics **24**, 1 (1997), 4616433
[10] J.R. Michael, W.R. Schucany, R.W. Haas, The American Statistician **30**, 88 (1976)
[11] D.A. Brown, M.B. Chadwick, R. Capote, A.C. Kahler, A. Trkov, M.W. Herman, A.A. Sonzogni, Y. Danon, A.D. Carlson, M. Dunn et al., Nuclear Data Sheets **148**, 1 (2018)
[12] P. Romano, S. Harper, EPJ Web of Conferences **146**, 06011 (2017)
[13] J. Yu, H. Huo, W. Li, G. Yu, K. Wang, *2014 22nd International Conference on Nuclear Engineering* (American Society of Mechanical Engineers Digital Collection, 2014)
[14] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright et al., Nature Methods **17**, 261 (2020)
[15] V. Peiris, N. Sharon, N. Sukhorukova, J. Ugon, Applied Mathematics and Computation **389**, 125560 (2021)
[16] J.M. Hokanson, *Multivariate rational approximation using a stabilized sanathanan-koerner iteration*, 2009.10803 [cs, math], http://arxiv.org/abs/2009.10803
[17] Y. Nakatsukasa, O. Sète, L.N. Trefethen, SIAM Journal on Scientific Computing **40**, A1494 (2018), 1612.00337 [math]
[18] J.P. Berrut, L.N. Trefethen, SIAM Review **46**, 501 (2004), publisher: Society for Industrial and Applied Mathematics
[19] M.S. Floater, K. Hormann, Numerische Mathematik **107**, 315 (2007)
[20] J. Tramm, P. Romano, J. Doerfert, A. Lund, P. Shriwise, A. Siegel, G. Ridley, A. Pastrello, *Toward Portable GPU Acceleration of the OpenMC Monte Carlo Particle Transport Code*, in *PHYSOR2022* (2022)