LA-UR-13-26477
Approved for public release; distribution is unlimited.

Title: Flux and Reaction Rate Kernel Density Estimators in OpenMC

Author(s): Burke, Timothy P.
Kiedrowski, Brian C.
Martin, William R.

Intended for: MCNP Website
Report

Issued: 2013-08-15

Flux and Reaction Rate Kernel Density Estimators in OpenMC

[1]Timothy P. Burke, [2]Brian C. Kiedrowski, and [1]William R. Martin

[1]*Department of Nuclear Engineering and Radiological Sciences, University of Michigan, Ann Arbor, MI*
[2]*X-Computational Physics Division, Los Alamos National Laboratory, Los Alamos, NM*
tpburke@umich.edu

INTRODUCTION

Current methods for obtaining flux and reaction rate profiles in a Monte Carlo simulation rely on histogram tallies that can suffer from large uncertainties when fine detail is required. Recently, another density estimation technique has been applied to Monte Carlo radiation transport simulations: the Kernel Density Estimator (KDE) [Banerjee, K.]. KDEs are capable of calculating the flux at a point without the necessity of tracking particles on a spatial mesh, making KDEs ideal tools for multi-physics problems. Furthermore, KDEs allow for one collision or track-length to contribute to the estimate of the flux at multiple points, circumventing the large statistical uncertainties associated with fine-mesh histogram tallies. However, KDEs have only been proven to be capable of estimating flux values at a series of points; no work has been done to estimate reaction rates or other quantities of interest. Furthermore, little work has been done in optimizing the KDE bandwidth for the energy-dependence of neutron transport simulations.

This summary introduces an optimal physics-based bandwidth for use in eigenvalue problems to obtain reaction rates and flux estimates at specified points using collision and track-length KDEs, as well as the implementation of this method into OpenMC [Romano, P. K.], a recently developed open-source Monte Carlo code primarily used for reactor-type problems. First, the current method of obtaining the optimal bandwidth in an eigenvalue problems is introduced and a simple thought experiment is presented to reveal the weaknesses in this approach. Next, the physics-based bandwidth is introduced and compared to the current method of obtaining and optimal bandwidth. Results from both methods of obtaining an optimal bandwidth are compared against a fine-histogram reference solution for one-group and continuous 1D problems. The results indicate that the collision and track-length KDEs are capable of accurately estimating the flux and reaction rate densities in one-group heterogeneous reactor-type problems; however, inaccuracies exist when the standard bandwidth is used in continuous-energy problems when the quantity of interest exhibits steep gradients at material interfaces. The results obtained using the physics-based bandwidth correct these innacuracies in the continuous-energy problem and provide nearly identical results to those in the one-group problem.

BACKGROUND & THEORY

KDEs are nonparametric estimators that take a collection of samples from an unknown density function and place a kernel function around each sample to obtain a smooth estimate of the underlying density function. The multivariate KDE used in this paper is composed of the product of univariate kernels and is defined by

$$\hat{f}(\mathbf{x}) = \frac{1}{N}\sum_{i=1}^{N}\prod_{l=1}^{d}\frac{1}{h_l}k\left(\frac{x_l - X_{l,i}}{h_l}\right), \tag{1}$$

where N is the number of samples, $x_l - X_{l,i}$ is the difference between the node at location $\mathbf{x}$ and the location of sample i in dimension l, h_l are the bandwidth parameters in each dimension, and k is the univariate kernel function.

Previous work has shown that the accuracy of KDEs are more severely affected by multimodal densities (e.g., fluxes in heterogeneous problems) than histogram tallies. A suggested solution is to segment the multimodal problem into several unimodal problems [L. Devroye, L. Györfi]. Banerjee's region-based bandwidth method does exactly this, and was shown to be a more accurate method for estimating quantities in geometries with heterogeneous materials than a KDE that utilizes a global bandwidth [Banerjee, K.]. The region-based bandwidth approach allows for bandwidths to be calculated on a per-region basis rather than across the entire domain, effectively transforming the larger problem into a collection of subproblems that can be handled by the KDE.

KDEs can be used to estimate reaction rates at a point in a straightforward manner. The collision KDE takes an estimate of the flux from a collision point and distributes that estimate over a finite space based upon the kernel being used. Thus, estimates of the reaction rate at a point can be obtained by multiplying every collision's contribution to the flux by the cross section of the reaction of interest at that point. This results in a collision rate estimator:

$$\hat{f}(\mathbf{x}) = \frac{1}{Nh}\sum_{i=1}^{N}\sum_{c=1}^{c_i}\frac{w_{i,c}\Sigma_r(\mathbf{x},E)}{\Sigma_t(\mathbf{X}_{i,c},E)}k\left(\frac{\mathbf{x}-\mathbf{X}_{i,c}}{h}\right), \tag{2}$$

where N is the number of histories, h is the product of the bandwidths in each dimension, c_i is the number of collisions in history i, Σ_r is the cross section for the reaction of interest, Σ_t is the total cross section, $w_{i,c}$ is the weight of particle i prior to collision c, E is the energy of the particle that caused the collision, $\mathbf{x}$ is the node location where the result is being tallied, $\mathbf{X}_{i,c}$ is the location of collision c in particle history i, and k is the multivariate KDE kernel defined as a product of univariate kernels as seen in Eq. (1).

Recently, an integral track-length KDE was presented that showed favorable properties when compared to the sub-track-length KDE approach [Dunn, K. L., Wilson, P.P.H.]. The physics-based bandwidth can also be applied to the integral-track KDE (further refered to as the track-length KDE in this paper). The track-length KDE can be used to determine reaction rates using

$$\hat{f}(\mathbf{x}) = \frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{c_i}w_{i,c}\int_0^{d_{i,c}}\Sigma_r(\mathbf{x},E)K_s dS, \tag{3}$$

where K_s is a 3D kernel function integrated over path length S that is defined in terms of the

starting point of the track $(X_o, Y_o.Z_o)$ and its unit direction vector (u, v, w):

$$K_s = \frac{1}{h_x} K\left(\frac{x - X_o - uS}{h_x}\right) \frac{1}{h_y} K\left(\frac{y - Y_o - vS}{h_y}\right) \frac{1}{h_x} K\left(\frac{z - Z_o - wS}{h_z}\right). \tag{4}$$

This integral is computed using a 4-point Gaussian quadrature scheme.

The optimal bandwidth for general KDEs has been discussed in depth by Silverman and is defined as the bandwidth that minimizes the sum of the integrated square bias and the integrated variance, known as the mean integrated square error (MISE) [SILVERMAN, B.W]. The approximate MISE is defined as

$$MISE = \frac{1}{4} h^4 k_2^2 \int f''(x)^2 dx + n^{-1} h^{-1} \int K(t)^2 dt, \tag{5}$$

where $f(x)$ is the density distribution being estimated. As seen in 5, a reduction in the bias creates an increase in the variance, and vice versa. This creates a complication in conducting figure of merit (FOM) comparisons between histogram tallies and the collision and track-length KDEs since the collision and track-length estimators attempt to minimize the MISE and the histogram estimator makes no attempt to do so. Thus, a figure of merit comparison between the KDE tallies and the histogram tallies would produce a better FOM for the histogram tallies due to their increased bias and consequently reduced variance.

Routines for obtaining the optimal bandwidth have been discussed by Banerjee in significant detail [Banerjee, K.]. The ideal bandwidth can be found by minimizing the MISE in Eq. (5) [E. PARZEN]

$$h_{opt} = k_2^{-2/5} \left(\int K(t)^2 dt\right)^{1/5} \left(\int f''(x)^2 dx\right)^{-1/5} n^{-1/5}. \tag{6}$$

Since the optimal bandwidth is dependent upon the distribution being used, a standard family of distributions is used to estimate quantities based on $f(x)$. The most commonly used trick for obtaining an optimal bandwidth is to use a normal distribution to estimate $f(x)$ and use the moments of the estimated distribution as parameters in the optimal bandwidth. For neutron transport problems, optimal bandwidths are currently computed for every region by

$$h_l = \left(\frac{4}{(2+d)N_c}\right)^{1/(4+d)} \sigma_l, \tag{7}$$

where d is the number of dimensions used in the KDE tally, N_c is the number of collisions that occur in a specific region, and σ_l is the standard deviation of the distribution of collision sites in dimension l in that region:

$$\sigma_l = \sqrt{\frac{1}{N_c} \sum_{i=1}^{N_c} (X_{l,i})^2 - \left(\frac{1}{N_c} \sum_{i=1}^{N_c} X_{l,i}\right)^2}. \tag{8}$$

Bandwidths obtained using this equation will here-after be termed location-based bandwidths after their reliance on collision locations.

These location-based bandwidths have several flaws that are revealed with a simple thought experiment. Take a slab consisting of an infinitely massive absorber (inifite absorption cross

section) with a beam of mono-energetic neutrons incident on each side of the slab. The location-based approach of obtaining an optimal bandwidth would have the mean collision location in the center of the slab, thus creating a large standard deviation of collision distributions depending on the thickness of the slab, in turn creating a large bandwidth. This leads to an estimator that will spread the distribution of flux and reaction rate into the interior of the slab while in reality the neutrons will be absorbed in the skin of the slab. Additionally, now imagine that the beam is no longer mono-energetic and the slab has a cross section that varies with energy. Neutrons will now travel a variety of distances before colliding in the slab of absorber, thus creating a unique distribution for each unique cross section. Since the optimal bandwidth is dependent upon the distribution being estimated, as seen in Eq. (5), the bandwidths should vary to accomodate this energy dependence of the distribution. Therefore, a new method of calculating the bandwidth is required.

The shortfalls of the location-based bandwidth lead to the development of an energy-dependent physics-based bandwidth that is based on the distribution of mean free paths traveled in each direction in a region rather than the distribution of collision sites in that region. The physics-based bandwidth is defined as

$$h_l\left(\Sigma_t(\mathbf{X},E)\right) = \frac{1}{\Sigma_t(\mathbf{X},E)}\left(\frac{4}{(2+d)N}\right)^{1/(4+d)}\sigma_l, \tag{9}$$

with σ_l now defined as

$$\sigma_l = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\Delta X_l \Sigma_t(\mathbf{X_i},E)\right)^2 - \left(\frac{1}{N}\sum_{i=1}^{N}\Delta X_l \Sigma_t(\mathbf{X_i},E)\right)^2}. \tag{10}$$

where ΔX_l is the distance the particle traveled in the l dimension in a particular region between collisions and N is the number of track lengths created in each region. Thus, the physics-based bandwidth creates kernels whose support regions reduce in size for tracks or collisions created with large cross sections and increase in size for small cross sections. Furthermore, the bandwidth is no longer dependent on where collisions occur but rather on the distribution of mean free paths traveled by a neutron in a region, a more physically based quantity.

KDE methods suffer from an increased bias at nodes whose kernels overlap with the boundary region. This bias comes from an increase in the error from $O(h^2)$ to $O(h)$ due to a loss of symmetry of the kernel since a portion of the kernel exists outside the bounds of the problem. There are numerous methods that have been presented for handling this loss of accuracy at boundaries. The boundary kernel method used by Banerjee for KDE tallies in neutron transport simulations [Banerjee, K.] was shown to obtain accurate answers at non-curvilinear problem boundaries when using a multivariate kernel consisting of the products of 1D kernels. The boundary kernel method involves using a boundary kernel consisting of a linear combination of the kernels $K(u)$ and $uK(u)$, which has a bias of $O(h^2)$. The boundary kernel is defined as

$$K_b(u) = \frac{a_2(p_1,p_2) - a_1(p_1,p_2)u}{a_0(p_1,p_2)a_2(p_1,p_2) - a_1^2(p_1,p_2)}K(u) \tag{11}$$

where $a_l(p_1,p_2) = \int_{p_1}^{p_2} u^l K(u)du$, $p_1 = \frac{x-xmax}{h}$ and $p_2 = \frac{x-xmin}{h}$. Thus, this kernel effectively handles overlap with both the minimum and maximum boundaries. At interior points, this boundary kernel reduces to the original kernel $K(u)$. Another benefit of this correction method

is that this boundary method can handle any type of boundary condition: vacuum, reflected, etc. While this boundary kernel will result in negative contributions to the score, no practical problems have resulted from this in practice.

IMPLEMENTATION

Fully-functioning collision and track-length KDEs have been implemented in OpenMC for eigenvalue problems. The KDE uses OpenMC's native tally structures for accumulating scores, inserting additional filters, and writing output. The KDE implementation is capable of obtaining results along user-defined lines, structured meshes, or at points defined by the user in a separate file. This allows the user to generate an unstructured mesh in their favorite CAE software that can then be loaded into OpenMC and used to obtain flux or reaction rates at the unstructured mesh node locations.

A boundary correction method is employed at nodes that have kernels that overlap with external boundaries since the smoothing properties of the KDE depress scores at those nodes. The OpenMC KDE implementation handles nodes that have kernels that overlap with external boundaries via the boundary kernel method stated in Eq. (11). The 1D results presented later also have the boundary correction applied at material interfaces. This is not ideal as it allows slight discontinuities in the flux at material boundaries. Even so, the errors produced this way are not significant. One limitation of this method when used with the multivariate kernel described in Eq. (1) is that it is currently limited to non-curvilinear boundaries. While this is presently only boundary correction method implemented in OpenMC, the modern programming practices employed in OpenMC make adding new methods a straightforward procedure.

In order to accurately capture flux gradients at material interfaces, locally adaptive bandwidths are required. The region-based bandwidth approach described earlier has been implemented in OpenMC, enabling users to assign collections of cells to a single region that will have a set of unique bandwidths. Scores are accumulated at each node in a region based upon the bandwidth in that region. That is, for the location-based bandwidths, a collision/track-length occuring in one region can contribute to the score at a node in a neighboring region if the collision occurred within one bandwidth of a node in that neighboring region. For the physics-based bandwidth, the score in one region is dependent on the number of mean free paths between the collision/track-length and the node. As such, ray-tracing would be required between each collision/track-length and each node within the support region of the kernel if the node and kernel existed in different regions. Since the KDEs discussed in this summary are conducted in eigenvalue problems, the standard deviations used to compute the optimal bandwidths are re-calculated at the start of each active batch using statistics from the previous batch via Eq. (8) or Eq. (12). For the test problems considered in this summary, this method produced bandwidths that fluctuate less than 1% from batch to batch. The capability to do 1D and 2D KDE tallies has also been added to OpenMC. That is, the results can be limited to the upper and lower limits of a dimension, similar to binning results in a particular dimension.

Since kernels with finite support regions are commonly used for KDE tallies, a nearest-neighbor list was created to reduce computing times. The nearest neighbor list breaks up the problem domain into a structured mesh with element widths equal to the size of the largest bandwidth in that direction among all the regions for the location-based bandwidth and the largest average bandwidth in that direction among all regions for the physics-based bandwidth,

where the average physics-based bandwidth for a region is defined by

$$\bar{h_l} = \left(\frac{1}{N} \sum_{i=1}^{N} \Sigma_t(\mathbf{X_i}, E) \right)^{-1} \left(\frac{4}{5N} \right)^{1/7} \sigma_l. \tag{12}$$

Nodes are then sorted into this structured mesh based upon their location. When a collision occurs, only the nodes in the neighborhood elements immediately in or adjacent to the element that the collision occurred in are tallied. Since the physics based bandwidth changes for each particle collision/track, a simple routine was added that allows for the number of neighborhood boxes to be searched in each dimension to adapt to the current bandwidth. Thus, the nearest neighbor list reduces the computational burden by only attempting to compute scores for nodes near the support regions of the kernels instead of over the entire problem domain. Since the speed increase obtained from the nearest-neighbor list is problem dependent, a simple figure of merit comparison will be presented later with the results.

Test Problems

The physics-based bandwidth methods should produce results similar to the location-based bandwidth methods for a one-group problem. Thus, Banerjee's simple 1D, one-group representation of a fuel lattice containing a strong absorber was modeled, a depiction of which can be seen in Figure 1 [Banerjee, K.]. The model consists of alternating slabs of fuel and moderator 1 cm thick with a central slab of absorber material and slabs of moderator 0.5 cm thick at the two boundaries. A similar test problem was created to show the continuous-energy capabilites of the KDE in a heterogeneous environment. The problem uses the same geometry as the one-group problem, but the materials consist of light water, 3.1% enriched UO_2, and a central slab of B_4C with cross sections obtained from the ENDF/B-VII.0 library at 293.6 K. Both problems are modeled in 3D with reflecting boundary conditions on all external boundaries, however they are effectively 1D problems as the materials only vary in the x-direction.

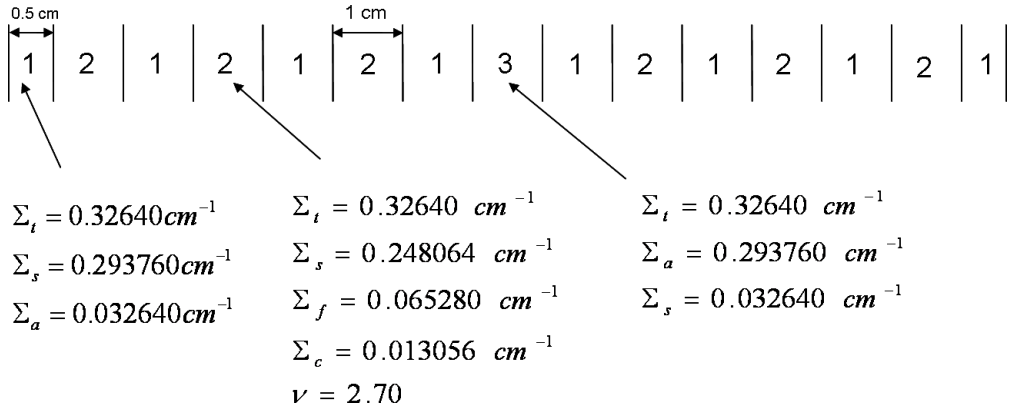

Figure 1: Depiction of 1D, one-group test problem.

All results were obtained using 20,000 histories per batch with 20 inactive batches and 300 total batches. The histogram tallies used a 1D bin structure with a total of 420 uniform bins over the 14 cm distance for the one-group problem and 1680 uniform bins for the continuous-energy problem, with bin-edges coinciding with material interfaces. The KDE for each problem was constructed such that point-wise results were obtained at the centers of each histogram bin.

The KDEs treat each cell as a unique region, thus making 15 regions with unique bandwidths in both test cases. Additionally, the Epanechnikov kernel, defined as

$$k(u) = \frac{3}{4\sqrt{5}}\left(1 - \frac{u^2}{5}\right), |u| \leq \sqrt{5} \tag{13}$$

was used for all KDE tallies.

RESULTS

One Group

Figure 2 shows comparisons between the flux and the fission, scattering, and absorption reaction rates obtained from histogram tally and the collision and track-length KDE tallies for the one-group problem using the location-based bandwidth. From Figure 2, it is clear that the location-based bandwidth KDE results agree with the histogram results. The maximum percent error between the KDE results and the histogram results is less than 1% for estimates of the flux and all reaction rates using the location-based bandwidth. Furthermore, the flux and reaction rate profiles obtained using the physics-based bandwidth for the one-group problem are almost identical to those obtained using the location-based bandwidth. Comparing the distributions obtained using the location-based bandwidth and physics-based bandwidth produces a maximum percent error of less than 0.2%, using the physics-based bandwidth results as reference values. This is not surprising since each material in the one-group problem has the same total macroscopic cross section, thus the distribution of distances traveled between collisions will not change between material and the local adjustment of the bandwidth by the total cross section in the material will have no affect. However, the point-wise results do not agree within statistics in areas of steep flux gradients for both methods of obtaining optimal bandwidths. This discrepancy is due to the volume-averaging nature of the histogram estimator, as will be demonstrated in the continuous-energy problem.

Continuous Energy

While the KDE perfromed well for the one-group problem, issues appear when applying the location-based bandwidth to continuous-energy problems. Figure 3 shows the comparisons between the flux, and the fission, scattering, and absorption reaction rate densities and for the continuous-energy problem using the location-based bandwidth. Figure 4 shows a comparison between the absorption reaction rate obtained using the histogram tally and location-based bandwidth KDE tallies as well as their corresponding errors for one side of the central absorber.

From Figure 3 it is apparent that the collision and track-length KDEs are capable of capturing flux profiles and reaction rate densities in continuous-energy simulations. While results obtained from the KDE tallies and the histogram tally agree at interior points of the slabs, they do not agree near material interfaces when the quantity of interest exhibits steep gradients, as seen in Figure 4. As a result, the maximum percent error, using the histogram result as the reference value, between the flux and reaction rate densities for all results exceeds 1%, with the maximum percent error of the absorption reaction rate exceeding 20%.

In order to study this discrepancy further, the density distributions were broken down into thermal and above-thermal results using the cadmium cutoff of 0.5 eV to distinguish between

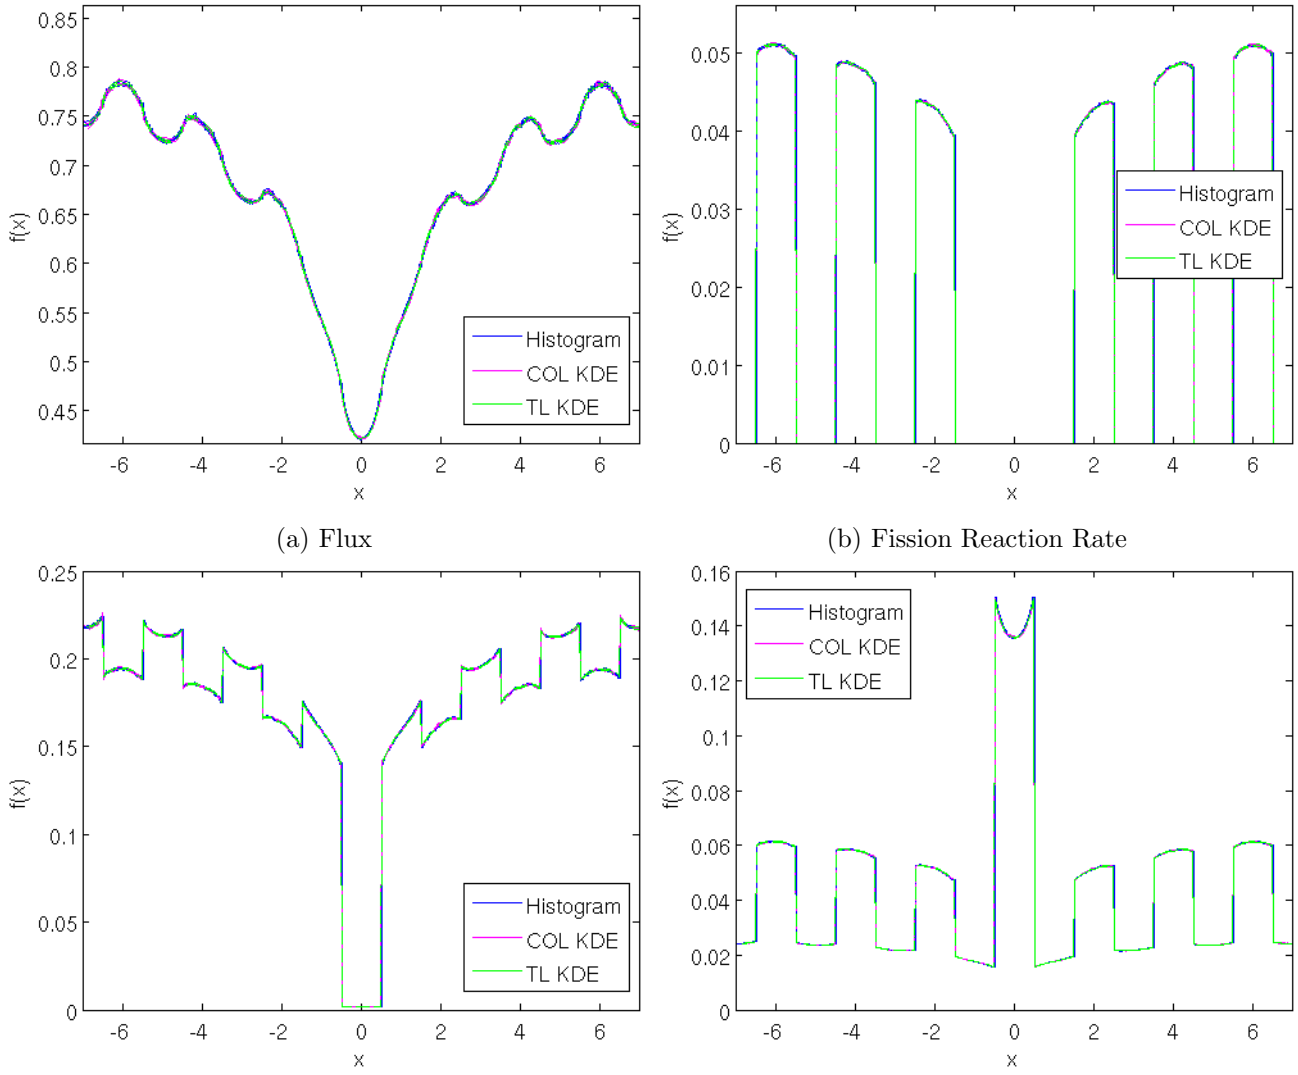

(a) Flux

(b) Fission Reaction Rate

(c) Scattering Reaction Rate

(d) Absorption Reaction Rate

Figure 2: Comparison of distributions obtained from histogram and KDE tallies for the one-group problem.

energy groups. Figure 5 shows the thermal (left) and above-thermal (right) absorption reaction rates using the location-based bandwidth. From Figure 5 it is clear that a constant bandwidth cannot capture the energy dependence of the problem. First, the absorption reaction rates in the central slab of absorber differ significantly between the thermal and above-thermal distributions. Since the optimal bandwidth is dependent on the shape of the distribution, different optimal bandwidths exist for each energy. Furthermore, the sharp peaks in the above-thermal absorption reaction rate distribution in the edge of the fuel require a bandwidth that's an order of magnitude smaller than the bandwidth obtained from Eq. 7. Such a small bandwidth would cause increased variance in the rest of the distribution, further showing the need for an adaptive bandwidth.

Several attempts to overcome these differences produced similar results. Simply breaking each slab into two regions to capture the distribution on each side of the slab seperately did not adequately resolve the issue. Furthermore, just using the track distance created by the

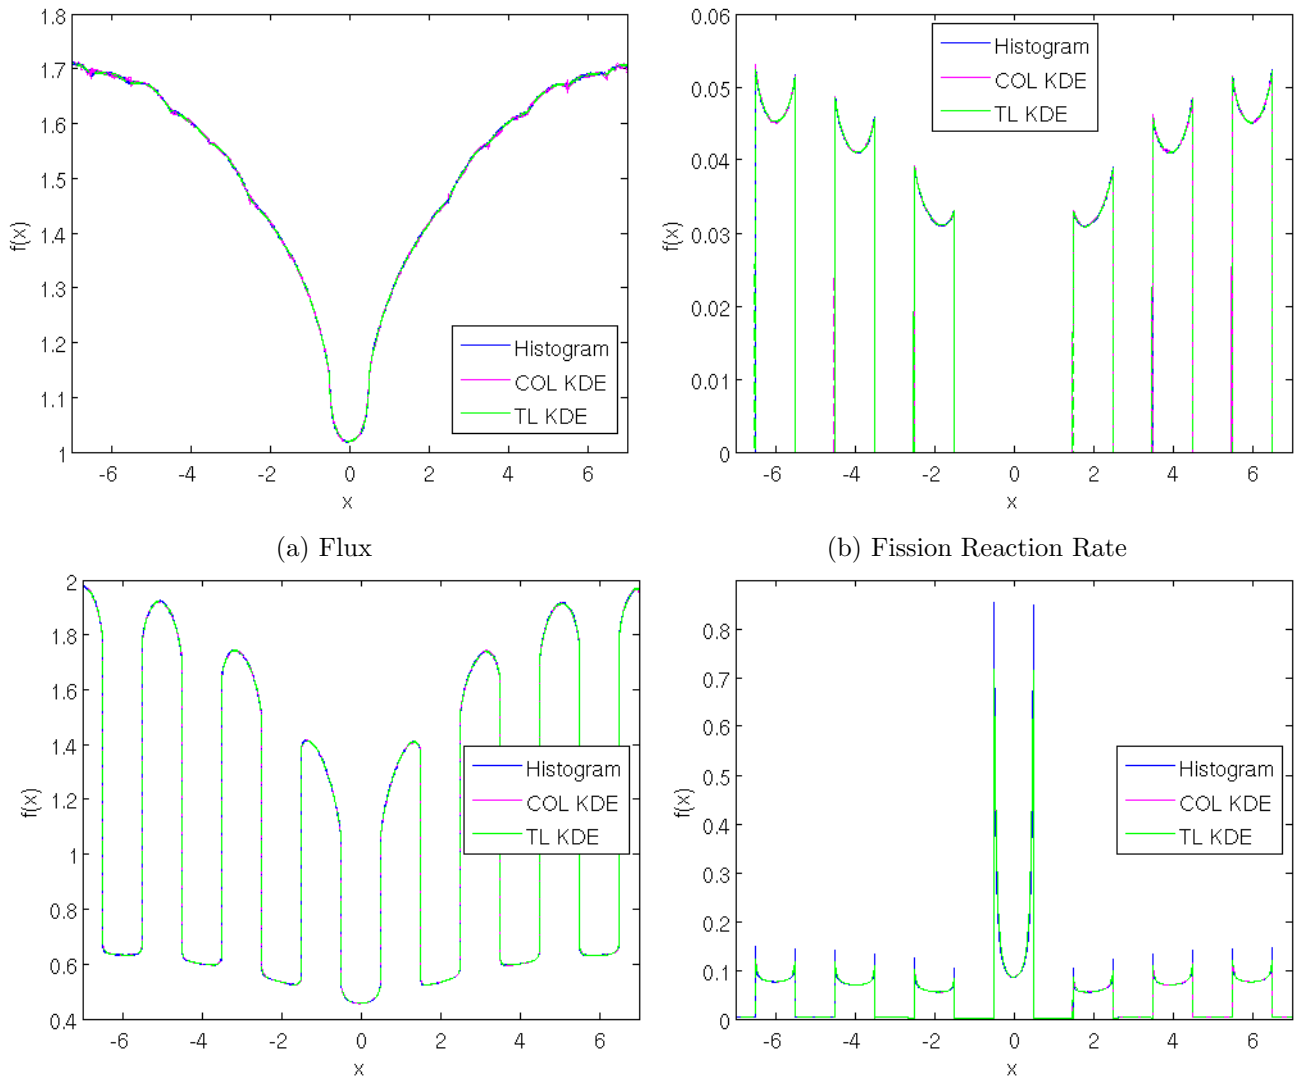

(a) Flux

(b) Fission Reaction Rate

(c) Scattering Reaction Rate

(d) Absorption Reaction Rate

Figure 3: Comparison of distributions obtained from histogram and the location-based bandwidth KDE tallies for the continuous-energy problem.

particle in each dimension for each collision produced results similar to those obtained using the location-based bandwidth. Since the differences in the distributions are a result of the energy dependence of the problem, it follows naturally that an energy-dependent bandwidth is required to resolve the various features in these distributions.

Physics-Based Bandwidth

A physical way to incorporate energy dependence into the KDE is to base the kernel on mean free paths traveled in a region instead of the location of collision points. This is exactly what the physics-based bandwidth accomplishes. Figure 6 shows the flux and the fission, scattering, and absorption reaction rate distributions using the physics-based bandwidth. With a simple visual comparison between Figures 6 and 3, it is clear that the physics based bandwidth more accurately captures local features in the reaction rate distribution resulting from the energy

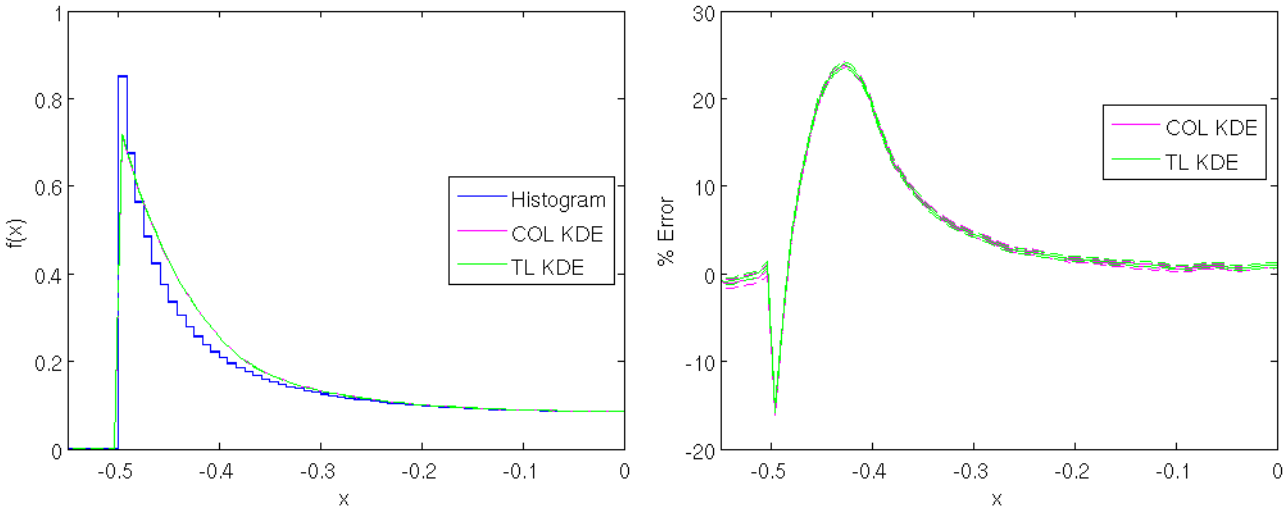

Figure 4: Absorption profiles and their corresponding error in the slab of absorber obtained from histogram and location-based bandwidth KDE tallies for the continuous-energy problem.

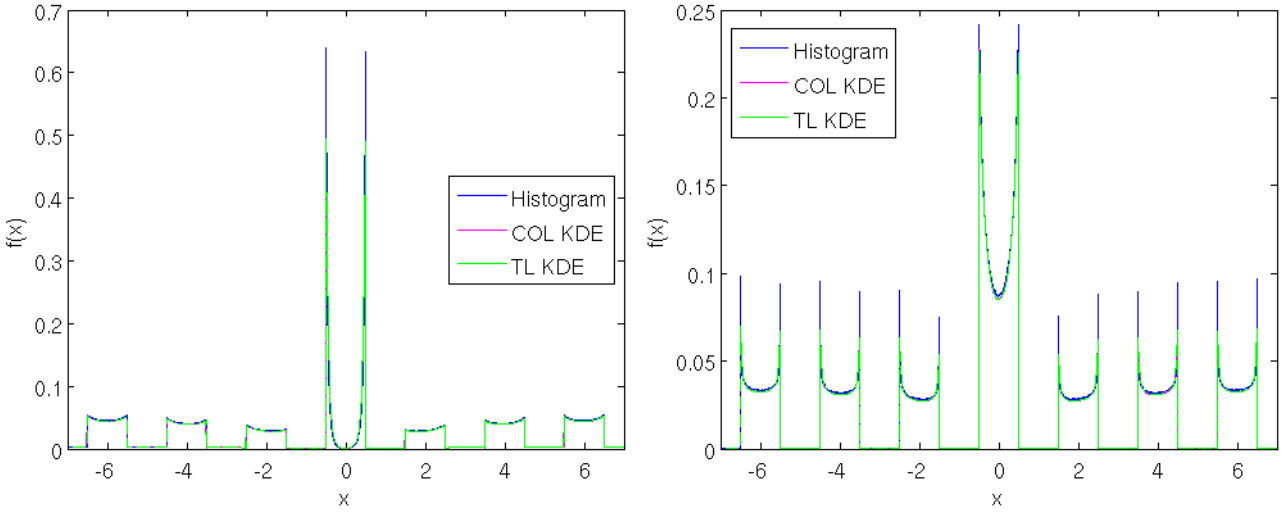

Figure 5: Comparison of absorption reaction rates between KDE and histogram tallies for thermal (left) and above-thermal (right) energies.

dependence of the solution. The KDE solutions and histogram solutions agree to within 1% for the flux, and all reaction rates at all locations, with the exception of the collision KDE estimate of reaction rates at nodes next to material interfaces and the track-length KDE estimate of the absorption reaction rates at nodes next to material interfaces. This is due to the volume-averaging nature of the histogram tally; densities that change slope significantly over the support of the histogram bin are poorly approximated by using the value in the histogram bin as the estimate of the density at the center of the bin. To demonstrate this, a simulation with 4x the resolution was conducted. Figures 7 and 8 detail the absorption reaction rate on the outer edge of a slab of fuel and the corresponding error for simulations conducted with a resolution of 1680 bins (left) and 6720 bins (right) over the whole problem domain. Again, KDE results were obtained at the center of each histogram bin for both resolutions.

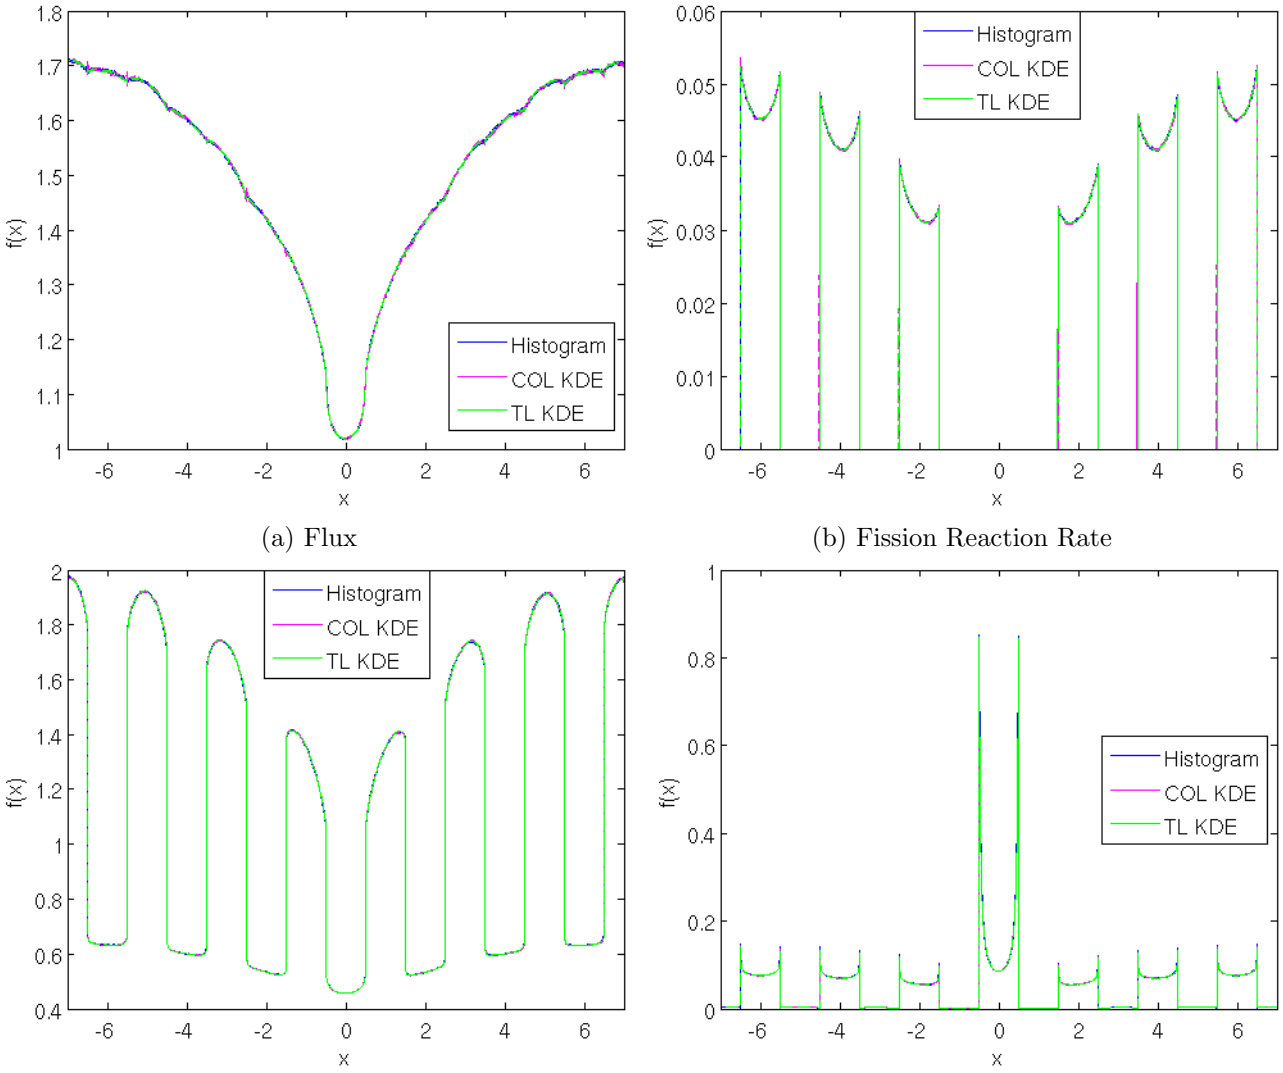

(a) Flux

(b) Fission Reaction Rate

(c) Scattering Reaction Rate

(d) Absorption Reaction Rate

Figure 6: Comparison of distributions obtained using physics-based bandwidth KDEs and histograms.

Figures 7 and 8 show that the finer the resolution on the histogram, the smaller the errors between the histogram tally and the KDE tallies. Thus, the systematic discrepancies between the KDE tallies and the histogram tallies are due to the resolution of the histogram, not a fault of the KDE. This showcases one of the benefits of the KDE tally: the accuracy of the KDE tally is not dependent on the density of KDE nodes. The KDE tally will obtain an accurate estimate of the flux or reaction rate at a point regardless of the density of KDE nodes in that region; changing the number of KDE nodes does not influence the results at any other nodes. Thus, if an accurate interpolating scheme is used and nodes are placed apppropriately, less nodes would be required to accurately capture density profiles as compared to the number of histogram bins required to accurately capture a density profile.

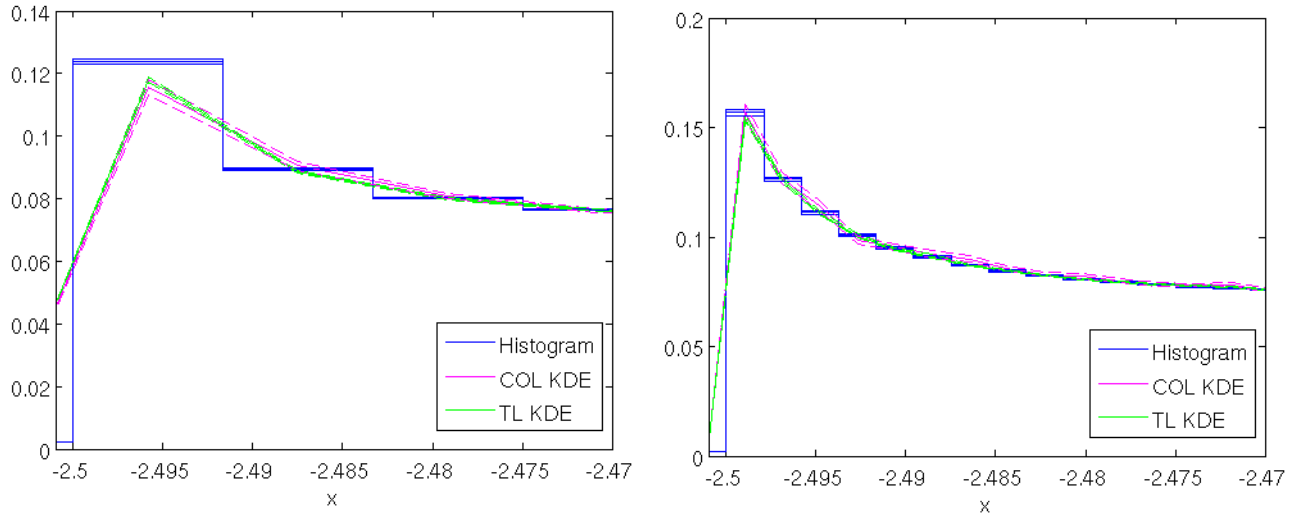

(a) 1680 bins and nodes

(b) 6720 bins and nodes

Figure 7: Absorption reaction rate in a slab of fuel resulting from histogram tallies and KDE tallies using 1680 bins (left) and 6720 bins (right).

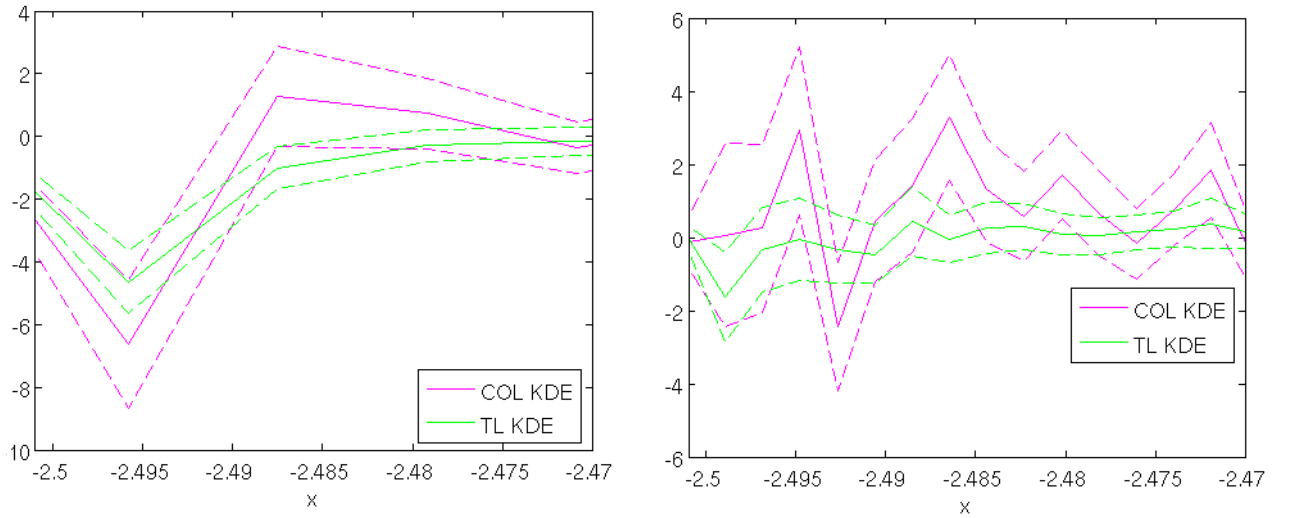

(a) 1680 bins and nodes

(b) 6720 bins and nodes

Figure 8: Percent error in the absorption reaction rate in a slab of fuel resulting from histogram tallies and KDE tallies using 1680 bins (left) and 6720 bins (right).

CONCLUSIONS & FUTURE WORK

A method for obtaining an optimal physics-based bandwidth for use in eigenvalue problems to obtain flux and reaction rate values at a collection of points has been introduced and implemented in OpenMC. The physics-based bandwidth KDE tallies were tested using collision and track-length KDEs in a one-group problem and a continuous-energy problem. The the results obtained from the histogram estimator and the KDEs agree well on both the one-group and continuous-energy problems with the use of the physics-based bandwidth. A simple comparison showed that as the number of histogram bins increased, the discrepancy between the KDE

tallies and the histogram tallies decreased.

While this method worked well in 1D, its implementation in higher dimension problems is more complicated due to the use of the boundary correction kernel at material interfaces. The current implementation of the boundary kernel method is not well suited to handle curvilinear geometry, so moving to simulations involving fuel pins would introduce greater complexity to the boundary kernel. The use of the boundary kernel at material interfaces could be circumvented by allowing collisions/track-lengths in one region to contribute to the flux at a point in another region. However, this would require ray tracing between collision points and nodes and would greatly increase the computational burden of the physics-based bandwidth KDE. Thus, further research is required to implement the physics-based bandwidth KDE in higher dimensional problems with curvilinear geometries without excessive increase in runtimes.

ACKNOWLEDGMENTS

Funding for this work was provided by the US DOE/NNSA Advanced Scientific Computing program. The authors would like to thank Dr. Kaushik Banerjee for providing insight into the implementation of the region-based bandwidth KDE methods.

References

[Banerjee, K.] K. BANERJEE, "Kernel Density Estimator Methods for Monte Carlo Radiation Transport", Ph.D. Thesis, University of Michigan (2010).

[Romano, P. K.] P.K. ROMANO, B. FORGET, "The OpenMC Monte Carlo Particle Transport Code," *Ann. Nucl. Energy*, **51** 274-281 (2013).

[L. Devroye, L. Györfi] L. DEVROYE, L. GYÖRFI, *Nonparametric Density Estimation*, p. 112, John Wiley & Sons, New York (1985).

[Dunn, K. L., Wilson, P.P.H.] K.L. DUNN, P.P.H. WILSON, "Kernel Density Estimators for Monte Carlo Tallies on Unstructured Meshes," *Transactions of the American Nuclear Society*, Vol. 107, San Deigo, CA, November 11-15, (2012).

[SILVERMAN, B.W] B.W. SILVERMAN, *Density Estimation for Statistics and Data Analysis*, Chapman and Hall, London (1986)

[E. PARZEN] E. PARZEN, "On estimation of a probability density function and mode," *Ann. Math. Statist.*, **33**, 1065-1076 (1986).