LA-UR-13-25055
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: 2013 American Nuclear Society Winter Meeting, 2013-11-10/2013-11-14 (Washington, District Of Columbia, United States)

Issued: 2013-07-09 (Rev.1)

Flux and Reaction Rate Kernel Density Estimators in OpenMC

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

[1]*Department of Nuclear Engineering and Radiological Sciences, Univeristy 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) [1]. 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. This summary introduces a method to obtain reaction rates using KDEs as well as the implementation of this method into OpenMC [2], a recently developed open-source Monte Carlo code primarily used for reactor-type problems. The results indicate that the KDE collision estimator is capable of accurately estimating the flux and reaction rate densities in one-group heterogeneous reactor-type problems, however inaccuracies exist in continuous-energy problems when the quantity of interest exhibits steep gradients at material interfaces.

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 summary is composed of the product of univariate kernels and is defined by

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

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, $h_l, ..., h_d$ 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 [3]. 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 [1]. 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.

Routines for obtaining the optimal bandwidth have been discussed by Banerjee in significant detail [1]. These optimal bandwidths for a 3-D problem are computed for every region by

$$h_i = \left(\frac{4}{5N_c}\right)^{1/7}\sigma_i, \qquad (2)$$

where σ_i is the standard deviation of the distribution of collision sites in dimension i occuring in a region and N_c is the number of collisions that occur in that region.

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:

$$\hat{f}(\mathbf{x}) = \frac{1}{Nh}\sum_{i=1}^{N}\sum_{c=1}^{c_i}\frac{w_{i,c}\Sigma_r(\mathbf{x},E)}{\Sigma_t(\mathbf{X}_{i,c},E)}k\left(\frac{\mathbf{x}-\mathbf{X}_{c,i}}{h}\right), \qquad (3)$$

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 or track-length, $\mathbf{x}$ is the location of the node 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 described in Eq. (1).

IMPLEMENTATION

A fully-functioning collision KDE has 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 [1]. The boundary kernel method employs a boundary kernel that is a linear multiple of the kernel being used at interior points whose shape is dependent upon the distance from the node to the external boundary. At nodes that lie farther than one bandwidth from an external boundary, this boundary kernel reduces to the kernel used at the interior points of the problem. Thus, the boundary kernel method is used at any node that falls within one bandwidth of an external boundary. This boundary method can handle any type of boundary condition: vacuum, reflected, etc. However, the use of this method with the multivariate kernel described in eq. (1) is limited to non-curvilinear boundaries. While this is currently the 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, a collision occuring in one region can contribute to the score at a node in a neighboring region if the collision occured within one bandwidth of a node in that neighboring region. Since the KDEs discussed in this summary are conducted in eigenvalue problems, the bandwidths are re-calculated at the start of each batch using collision location statistics from the previous batch via Eq. (2). For the test problems considered in this summary, this method produced bandwidths that fluctuate less than 1% from batch to batch.

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. 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 occured in are tallied. 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.

Test Problems

In order to verify that the region-based bandwith methods are implemented properly, Banerjee's simple 1-D, one-group representation of a fuel lattice containing a strong absorber was modeled [1]. The model consists of alternating slabs of fuel and water 1 cm thick with a central slab of absorber material and slabs of water 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.

The collision KDE results were obtained using 100,000 histories per batch with 20 inactive batches and 1000 total batches. The histogram tallies were obtained using 100,000 histories per batch with 20 inactive batches and 600 total batches. The histogram tallies used a 1-D bin structure with a total of 420 bins over the 14 cm distance for the one-group problem and 840 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 KDE nodes are located several bandwidths away from the y and z boundaries to eliminate any effect the boundary correction method may have on the answers at interior nodes. 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

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

was used for all KDE tallies.

RESULTS

Figures 1 - 3 show comparisons between the flux, fission reaction rate, and scattering reaction rate results obtained from histogram tally and the KDE tally for the one-group problem. The maximum percent error between the collision KDE results and the histogram results is less than 1% for estimates of the flux and all reaction rates. However,

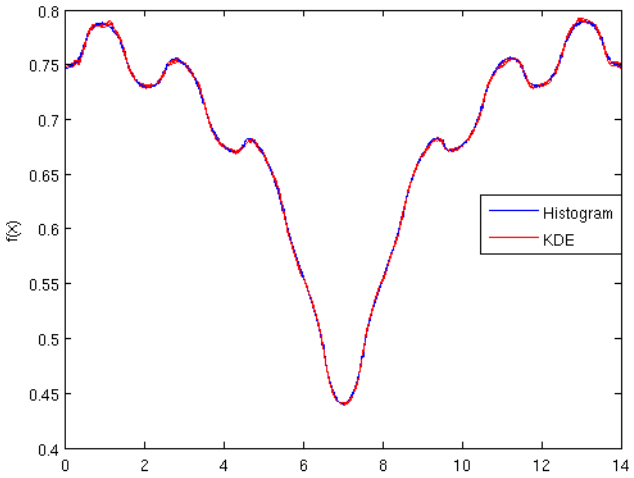

Fig. 1. Comparison of flux profiles obtained from histogram and KDE collision tallies for the one-group problem.

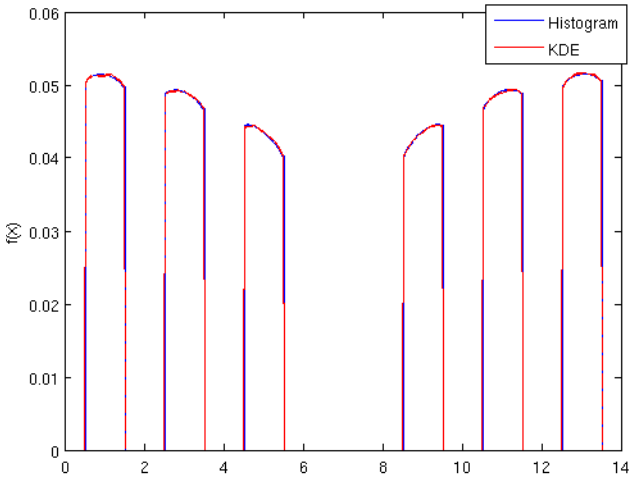

Fig. 2. Comparison of fission reaction rate profiles obtained from histogram and KDE collision tallies for the one-group problem.

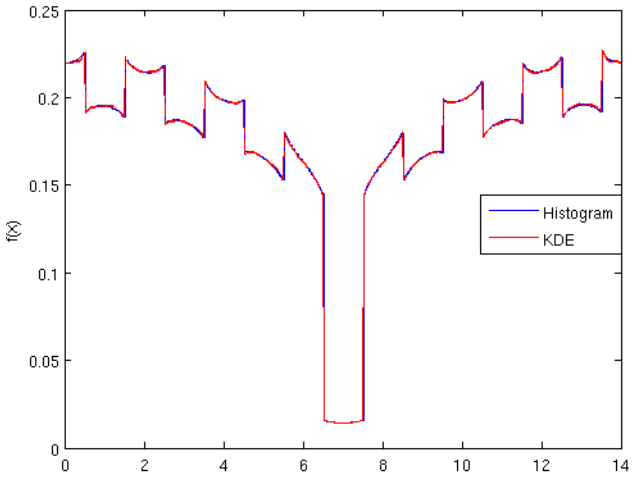

Fig. 3. Comparison of scattering reaction rate profiles obtained from histogram and KDE collision tallies for the one-group problem.

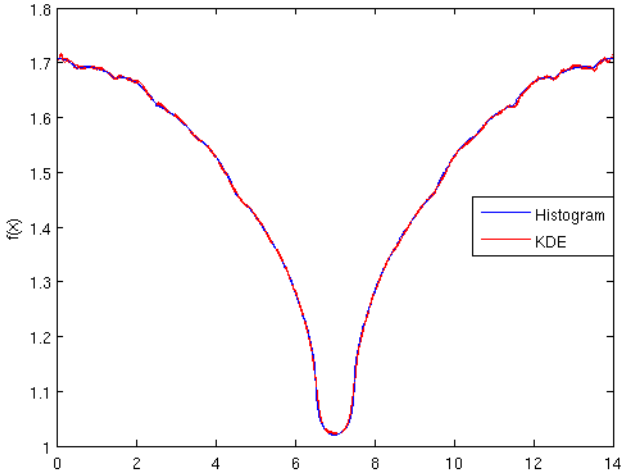

Fig. 4. Comparison of flux profiles obtained from histogram and KDE collision tallies for the continuous-energy problem.

the point-wise results do not agree within statistics in areas of steep flux gradients. The majority of this discrepancy is likely due to the volume-averaging nature of the histogram estimator. Thus, if a finer histogram bin structure were used then this discrepancy would likely decrease. Figures 4-6 show the comparisons between the flux and the fission and absorption reaction rate densities for the continuous-energy problem. Figure 7 shows the comparison between the fission reaction rates obtained using the histogram tally and KDE tally for a single fissionable slab. From figures 4-6 it is apparent that the collision KDE is capable of capturing flux profiles and reaction rate densities in continuous-energy simulations. However, the maximum percent error between the flux and reaction rate densities for all results exceeds 1%, and the results for the histogram tally and the mesh tally do not agree at material interfaces. Furthermore, while the point-wise estimates of the flux and reaction rate densities agree within statistics at the interior points of each slab, there is a difference between the two tallies at material interfaces that cannot be attributed to statistics when the quantity of interest exhibits steep gradients. This can be seen in Figure 7. The KDE tally overestimates reaction rates near material interfaces where flux and reaction rate gradients are steep.

Additionally, some spurious oscillations can be seen in the collision KDE's estimate of the flux. Previous work

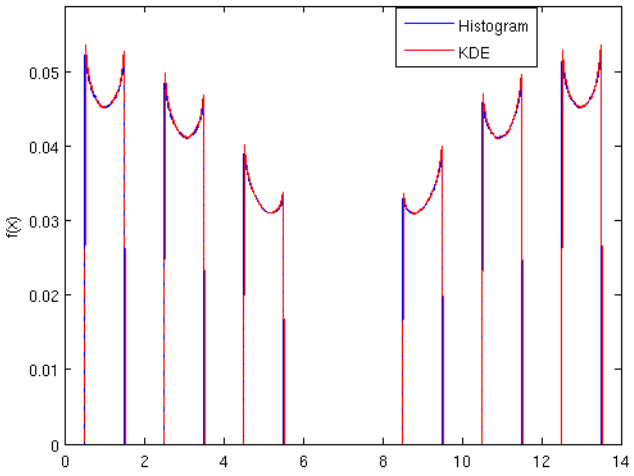

Fig. 5. Comparison of fission reaction rate profiles obtained from histogram and KDE collision tallies for the continuous-energy problem.

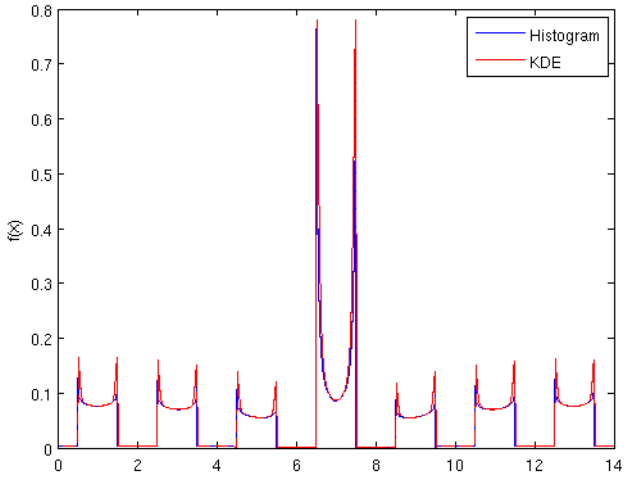

Fig. 6. Comparison of absorption reaction rate profiles obtained from histogram and KDE collision tallies for the continuous-energy problem.

has shown that these oscillations can be reduced by using a track-length KDE [1].

CONCLUSIONS & FUTURE WORK

A method for calculating reaction rates using region-based bandwidth collision KDEs has been introduced and tested using a one-group problem and a continuous-energy problem in OpenMC. While the comparison between the results obtained from the histogram estimator and collision KDE agrees well on the one-group problem, the collision KDE produces inaccurate reaction rates for the continuous-energy problem when the reaction rate exhibited steep gradients. This is most likely due to the lack of direct material information or energy-dependent information in the bandwidth calculation, as each material region in the one-group problem had the same total macroscopic cross section.

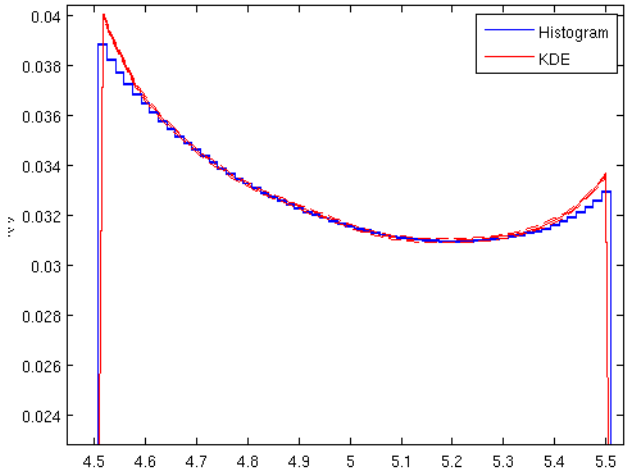

Fig. 7. Comparison of fission reaction rate in one slab of fissionable material for the continuous-energy problem.

Thus, either methods that exhibit locally adaptive bandwidths, methods that have energy-dependent bandwidths, or methods with region-based bandwidths that are more ideally optimized for energy-dependent solutions will be investigated to reduce inaccuracies at material interfaces with severe density gradients. Furthermore, a track-length KDE will be implemented into OpenMC to attempt to resolve some of the discrepancies between the histogram and KDE results.

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

1. K. BANERJEE, “Kernel Density Estimator Methods for Monte Carlo Radiation Transport”, Ph.D. Thesis, Univeristy of Michigan (2010).
2. P.K. ROMANO, B. FORGET, “The OpenMC Monte Carlo Particle Transport Code,” *Ann. Nucl. Energy*, **51** 274-281 (2013).
3. L. DEVROYE, L. GYÖRFI, *Nonparametric Density Estimation*, p. 112, John Wiley & Sons, New York (1985).