

Department of Physics and Astronomy
Heidelberg University

Bachelor Thesis in Physics
submitted by

Frederic Wolfgang Heinrich Böhm

born in Menden (Germany)

April 2025

Unbinned analysis of ^{163}Ho -spectrum endpoint region

This Bachelor Thesis has been carried out by Frederic Wolfgang Heinrich Böhm at the
Kirchhoff Institute for Physics in Heidelberg
under the supervision of
JProf. Dr. Loredana Gastaldo

Abstract

The aim of the Electron Capture in ^{163}Ho (ECHO) collaboration is to determine the effective electron neutrino mass by analyzing the endpoint region of the ^{163}Ho electron capture spectrum. The spectrum is measured using metallic magnetic calorimeters (MMC) enclosing ^{163}Ho and subsequently the data is reduced to avoid the presence of artifacts before further analysis can take place. Previously, a histogram-based approach has already proven to be a suitable choice for the analysis of the spectrum and, in particular, of the endpoint region. To further improve the sensitivity of the fitting algorithms to quantify the effect of tiny neutrino masses, we are testing methods of unbinned analysis like a Kernel Density Estimation (KDE) to mitigate potential artifacts of binning the continuous event energies of the low-intensity endpoint region close to the Q-value of the ^{163}Ho decay. We present the implementation of these algorithms in the analysis of the ^{163}Ho spectrum acquired within the ECHO-1k experiment and compare the results with those obtained with binned spectra.

Analyse der Endpunktregion des ^{163}Ho -Spektrums ohne Verwendung von Histogrammen

Das Ziel der Electron Capture in ^{163}Ho (ECHO) Kollaboration ist es die effektive Elektronenneutrinomasse über die Analyse der Endpunktregion des ^{163}Ho Elektroneneinfangspektrums zu bestimmen. Das Spektrum wird gemessen mithilfe von Metallischen Magnetischen Kalorimetern (MMC) in denen ^{163}Ho eingeschlossen ist, und anschließend werden die Daten reduziert um das Auftreten von Artefakten zu verhindern. Zuvor hat sich ein auf Histogrammen basierender Ansatz bereits als sinnvolle Methode bewiesen das Spektrum und insbesondere die Endpunktregion zu analysieren. Um die Empfindlichkeit der Kurvenanpassungsalgorithmen zur Quantifizierung des Effekts winziger Neutrinomassen weiter zu verbessern testen wir Analysemethoden die auf Verwendung von Histogrammen verzichten wie beispielsweise eines Kerndichteschätzers (KDE) um Artefakte des Einteilens kontinuierlicher Energien der Endpunktregion mit geringen Intensitäten nahe dem Q-Wert des ^{163}Ho Zerfalls in Klassen zu mildern. Wir präsentieren die Implementierung der Algorithmen in der Analyse des ^{163}Ho Spektrums welches im ECHO-1k Experiment gewonnen wurde und vergleichen die Ergebnisse mit denen eines Spektrums in Histogrammdarstellung.

Contents

1	Introduction	1
2	^{163}Ho endpoint region fit	2
2.1	Phase space factor	3
2.2	Gaussian Detector Response	3
2.3	Background	4
2.4	Atomic physics function	5
3	Kernel Density Estimation	6
4	Method of testing the KDE	9
4.1	Data reduction	9
4.2	Preparing the spectral shape	10
4.3	Fraction of pile-up events in the background	13
4.4	Detector response	14
4.5	Endpoint region fit	14
5	Results	16
5.1	Detector response	16
5.2	Endpoint region fit	16
5.2.1	Residuals	17
5.2.2	Parameters	18
6	Conclusions	20
7	Appendix	22

1 Introduction

As the Standard Model of particle physics describes, most of the fundamental particles we know in our universe have mass. One exception is the photon, which is massless, and the neutrinos were long believed to fall into this category as well, until evidence for the existence of neutrino oscillation was given and thus neutrinos carrying a mass was proven [1]. Determining the neutrino mass allows us to gain a better fundamental understanding of particle physics, including giving insight into the history of the universe as well as how it will evolve in the future.

As for how the neutrino mass is to be determined, there are multiple approaches. One of them would be the search for a neutrinoless double-beta decay. By finding the corresponding decay rate, it will lead us to direct conclusions concerning the neutrino mass, as the two are closely linked [2]. Another possibility is by examining the cosmic microwave background or the mass distribution in the universe [3].

Determining the neutrino mass through electron capture processes on the other hand has the benefit of being model-independent. If the energy available in a decay involving the release of a neutrino is known, then direct measurements of the energy after the neutrino has been released leads to information on the mass of the neutrino thanks to energy and momentum conservation. As we will see, determining this energy difference requires us to look at the endpoint region of decay spectra, where only very few events can be found. This poses a challenge for the analysis due to the low amount of data available. It is the essence of this thesis to explore a method called Kernel Density Estimation (KDE) to help improve analysis methods, especially for the endpoint region, by smoothing out the limited data which is still available.

In chapter 2, the general concept of the fit that is performed in the endpoint region of the ^{163}Ho decay spectrum is portrayed. The different parts making up the fitting function are introduced and explained.

In chapter 3, Kernel Density Estimation as a mathematical model for estimating underlying probability density functions from a data set is introduced. The benefits and challenges with the use of KDEs are portrayed.

In chapter 4, the method that is used to test the viability and quality of the KDE approach is presented, as well as a brief overview of the data reduction which is performed on the raw data before a KDE is applied.

In chapter 5, the findings concerning the tests mentioned in the foregoing chapter are presented.

2 ^{163}Ho endpoint region fit

With the help of the electron capture spectrum of ^{163}Ho measured in the ECHo experiment, it is possible to determine the effective electron neutrino mass. It is known from nuclear physics that a specific maximum amount of energy is available in the decay of ^{163}Ho , known as the Q-value $Q = (2863.2 \pm 0.6) \text{ eV}$ [4]. The analysis of the electron capture spectrum will yield the part of this energy deposited in the atomic excitation of the daughter ^{163}Dy atoms. This means the difference between the Q-value of the ^{163}Ho decay and the endpoint determined in this experiment results in information on the maximum amount of energy available to the electron neutrino. Therefore this analysis gives an upper bound on the effective electron neutrino mass. The majority of the electron capture spectrum shows very high intensities of detected events, in turn making the spectrum in these areas very precise. Unfortunately, these areas of the spectrum are not highly dependent on the effective electron neutrino mass and because of this cannot be used to gain information on it. Instead, the area that shows the highest effect in this regard is the endpoint region, close to the maximum amount of energy available in this decay. The comparison on the exponential model between an effective neutrino mass of 0 eV and 5 eV, as is modeled in 1, shows that the spectral shape experiences a visible difference near the endpoint.




Figure 1: Endpoint region close-up of exponential model for a) $m_\nu = 0 \text{ eV}$ and b) $m_\nu = 5 \text{ eV}$

Analysis in this area proves challenging because of the very low amount of events, multiple orders of magnitude lower than the majority of the spectrum. This is why the development of analysis methods specifically tailored to this problem is vital. A previous attempt at this utilized a histogram to quantify the shape of the spectrum from the list of energies detected in the events of the experiment. In the endpoint region, a fit was made on the spectral shape which could ultimately determine the

Q-value of the decay, leading to information on the effective electron neutrino mass as described before. I will now introduce the general structure of the fitting function $\frac{dN}{dE}$ and break down its several parts:

$$\frac{dN}{dE} = C \cdot [A(E) \cdot PS(E)] * g(E) + BC(E) \quad (2.1)$$

Equation 2.1 describes the shape of the energy spectrum and is comprised of the phase space factor $PS(E)$, the Atomic Physics function $AP(E)$, a Gaussian function $g(E)$ representing the detector response with a finite energy resolution and a constant background $BC(E)$. The constant C scales the spectrum to take into account the total amount of events.

2.1 Phase space factor

The phase space factor PS (2.2) models the phase space becoming narrower when approaching the endpoint from lower energies due to less energy being available for the decay itself, which leads to events becoming increasingly unlikely. Events at $E_{EC} = Q_{EC}$, when the detected electron capture energy is equal to the Q-value, the phase space factor reaches 0. At energies beyond the Q-value, the phase space factor remains 0 as events such as this are energetically impossible.

$$PS(Q_{EC} - E_{EC}) = H(Q_{EC} - E_{EC} - m_\nu) \cdot (Q_{EC} - E_{EC}) \cdot \sqrt{(Q_{EC} - E_{EC})^2 - m_\nu^2} \quad (2.2)$$

As we will see, the undesired background in the spectrum is not directly confined by the Q-value, the phase space factor only applies to the atomic physics and is therefore multiplied with the atomic physics function before the background is added on.

2.2 Gaussian Detector Response

The detector used to collect the experimental data for this analysis, like all detectors, operates at a finite resolution. This means that even for perfectly controlled events, there is always a certain statistical error in the measured data. This is modeled by convolving the product of the phase space factor and the atomic physics function with a Gaussian kernel. The width of the kernel, or rather the resolution of the detector, is determined as part of this analysis in section.

2.3 Background

Shielding of a detector to prevent a background disturbing the experimental data is vital, however, it is not possible to perfectly remove all background sources. For this experiment, two background sources need to be considered in particular. Firstly, cosmic radiation such as muons and natural radioactivity will be detected. In the energy ranges relevant to this experiment, this background is assumed to be constant across all energies, as no structures can be identified to attribute events to any possible background source.




Figure 2: ^{163}Ho spectrum and unresolved pile-up spectrum

Secondly, another background source is the unresolved pile-up spectrum as shown in 2. While the spectral data are reduced prior to analysis to remove detected events unrelated to the experiment, or environmental noise sources, one effect that cannot be mitigated is when two events pile up in such a short time frame that they cannot be told apart. This leads to the detection of an apparent event with the combined energy of the decay events that piled up. As any combination of energies between the two events is possible, this pile-up spectrum is modeled by the autoconvolution of the electron capture spectrum. Due to the relatively low intensity and little number of peaks in the pile-up spectrum near the endpoint region, a constant background BC will be modeled and quantified, and is then added to the fitting function.

2.4 Atomic physics function

The atomic physics function is responsible for modeling the overall shape of the electron capture spectrum of the ^{163}Ho decay, such as the location of the different peaks that can be found. However, no analytical solution for this is known as of yet. As we have previously established, the greatest effect of the neutrino mass is to be seen in the endpoint region and therefore fitting of the spectrum will be limited to this range. This also allows us to choose simpler fitting functions. We will be testing the analysis methods using both the exponential function in 2.3 and the quadratic function in 2.4 to model the atomic physics.

$$f_{\text{exp}}(x) = A \cdot \exp(\tau \cdot x) \quad (2.3)$$

$$f_{\text{quad}}(x) = a \cdot x^2 + b \cdot x + c \quad (2.4)$$

3 Kernel Density Estimation

Displaying a spectrum in the form of a histogram is an effective method for reconstructing the underlying probability density function of the decay that is observed. While it works well in high-intensity regions of the spectrum and can be tweaked by altering the bin width used to group the measured energies, especially in low-intensity regions we are met with a dilemma. While decreasing the bin width increases the resolution and accuracy of the representation of the spectrum, it comes at the risk of generating empty bins, so bins in the histogram into which not a single energy has fallen. This effect is undesirable, as the empty bins appear statistically, and therefore do not represent a true dip in the probability density function. As we have established before, it is exactly this low-intensity endpoint region in which the effective neutrino mass shows the largest influence, so it is vital that the spectrum is as close to the real distribution of event energies as possible. A potential way to mitigate the adverse effect of histograms worth exploring is choosing a method of unbinned analysis, like the Kernel Density Estimation (or KDE).

While a histogram has a predefined bin width and with that a regular step size at which the raw energy data points are grouped together, which leads to a certain loss in accuracy, a KDE takes each data point into account at its exact position. This is done by placing a normalized kernel of a certain shape centered around each data point, and then adding up all kernels resulting in a function resembling the probability density function. Both the shape of the kernel as well as the corresponding width can be chosen freely in theory, however a common choice is a Gaussian kernel, due to the measured data points always underlying statistics. As we will see, the choice of the width of a kernel, also referred to as the bandwidth of the KDE, is both highly relevant to the overall quality of the results, but also a challenging task.

When the KDE is generated from a set of n data points, it is a continuous curve made up of n overlaid kernels. Before the KDE can be used for analysis in similar ways as a histogram, it needs to be discretized to a finite amount of points at which it is evaluated. The amount of these points, as well as their positions and the step width between adjacent points can be chosen freely, and can even change along the axis of the data points within the same KDE. In figure 3 we can see an example of a KDE generated from a set of data points that are Normal-distributed, as well as a histogram of the same data set. Both of these representations give a reasonable impression of the underlying distribution, however the histogram shows far more sudden jumps in counts, which are instead smoothed out a bit on the KDE. If one wanted to smooth out the histogram, there would be no other choice than to increase the bin width and thereby decreasing the overall amount of bins. This is a loss in accuracy and results in fewer points available for analysis, for example through curve fitting.




Figure 3: Data set underlying Normal distribution, displayed as histogram and KDE (bandwidth = 0.1 eV)




Figure 4: Data set underlying Normal distribution, displayed as histogram and KDE (bandwidth = 1.5 eV)

The selection of a bandwidth depends strongly on the respective data set that is to be modeled. For every distribution, if the resulting estimation shows very narrow peaks around some data points, then the selected bandwidth is too small. On the other hand, selecting too large of a bandwidth (see figure 4) results in the shape of the underlying distribution not being shown correctly. Individually for each application of a KDE a bandwidth needs to be chosen that appropriately models the distribution, depending on how important the exact location of a single data point is or how much statistical

imbalances are supposed to be combated.

It is apparent that a KDE has benefits over a histogram when used to represent a distribution, such as the configurable resolution which is entirely unrelated to the bandwidth, whereas in a histogram the resolution and bin width are intrinsically tied together. But a KDE also comes with its challenges, such as far longer processing times. The steps required to build a KDE increase linearly with both the amount of points in the data set, as well as the amount of points at which the KDE is being discretized. Generating a histogram is a much simpler task, as data points are merely being counted in regular steps. Another challenge of a KDE is the selection of a reasonable bandwidth, as it will strongly influence the result.

4 Method of testing the KDE

In order to test the effect of Kernel Density Estimation on the analysis of the ^{163}Ho electron capture spectrum, the approach is to make direct comparisons with the same analysis performed on the spectrum displayed through a histogram. For this, a program set up by Daniel Behrend was used and built upon, to enable the use of a KDE while being able to simply switch between the operational modes for using either a histogram or a KDE for analysis. I will now lay out the steps which are taken in the program, leading up to the final analysis.

4.1 Data reduction

Before any form of endpoint region analysis can take place, the data taken in the ECHo experiment first needs to be reduced. This process follows the aim of reducing undesirable effects like various background sources or data points that do not accurately represent the spectrum. While these steps were not part of this thesis, I will briefly introduce some of the steps taken for data reduction.

In a first step, a Holdoff filter is applied to the raw data points. Due to the MMC's nature of calorimetrically measuring excitations, the heat deposited in the calorimeter slowly dissolves into a thermal reservoir. Because of this, the voltage measured by a SQUID also only falls slowly over time. If a second event hits the same channel after only a short amount of time, the peak of the new event lands on top of the tail of the previous event - called a pile-up event - and the voltages stack. This throws off the energy reading and so the Holdoff filter's task is to wait a certain amount of time after an event is triggered before accepting another events. A typical Holdoff cut is on the order of $\Delta T_{\text{Holdoff}} = 15 \text{ ms}$.

Next up, a Burst filter is used. Its purpose is to detect quickly repeating triggers of events, indicating noise in a single channel, and discard the corresponding traces. With ^{163}Ho decay and the current setup, on average about 1 event per second is to be expected. The Burst filter first arranges the detected events into a histogram based on their timestamps, and then checks whether each bin's event count N_{Ch} deviates by more than 4σ from the average. If a bin is found that triggers this condition, a test is performed where one of the neighboring bins has to exceed a deviation of 2σ from the average. If this condition is also met, the Burst filter considers these events as noise and removes the bins.

A Coincidence filter is also used. Instead of acting per channel like the aforementioned filters, the Coincidence filter acts globally to detect background radiation, which might have sent a shower of particles over the entire detector, hitting different channels with only a very short amount of time of difference. This is why this filter will check if any

events across the whole detector were triggered with a time difference $\Delta T < 8 \mu\text{s}$, and if so, these events are recognized as background radiation and will be discarded.

4.2 Preparing the spectral shape

Now that the acquired data has been stripped of as many faulty data points as possible, it needs to be brought into a usable data format for analysis. Right now the spectrum is only a list of energy values, each entry describing exactly one event that remained after data reduction. In order to display the spectrum graphically, a common choice is a histogram, binning data points into ranges on the order of 1 eV. Due to the simplicity of how the histogram delivers us a representation of the probability density function governing the ^{163}Ho decay, the spectrum is now already in place for use in subsequent steps of the analysis.

In the case of a Kernel Density Estimation, more care is required to produce a usable result in this step. A naive approach, even if it seems sensible at first, is to simply perform a KDE on the entire data set at once, using only a single bandwidth value for the entire range from 0 eV to 7000 eV. While the bandwidth can and has to be chosen freely and through trial and error, it quickly becomes apparent that this produces sub-par results. The reason for this is the general shape of the spectrum, being very unlike a normal distribution itself. If it were purely a normal distribution, a KDE using a Gaussian kernel could reasonably match this shape as seen in figure 3. However, the ^{163}Ho spectrum shows peaks of very high intensity, while the endpoint region naturally goes down to very low intensities, and it does so in a very rapid manner over multiple orders of magnitude. The result is the extended tails of the millions of overlaid Gaussians in the high-intensity regions bleeding into the lower-intensity regions, and in the process strongly obscuring the spectrum in these areas. An exaggerated example is shown in figure 5.

A first attempt at mitigating this effect was decreasing the global bandwidth to values in the milli electronvolts. While this does fix the problem and still properly preserves the spectrum in the high-intensity regions, we are then met with another issue. A close look at the endpoint region (see figure 6) reveals, tiny bandwidth values lead to very steep peaks around the few events that can be found here, with no notable width and larger gaps between the events. The original idea behind using a KDE was being able to smooth out the few scattered events of the endpoint region through the Gaussian kernels extending left and right from the data point, and thus this fix defeats the purpose of the KDE in the first place.

In the end, two possible solutions were found that deliver usable spectral shapes. The first one involved sticking with a constant bandwidth across the entire spectrum, however the spectrum was evaluated in steps of a certain width, usually around the order




Figure 5: KDE of spectrum applied with constant bandwidth of 0.1 eV

of 50 eV, and the tails of all the Gaussians inside each step were cut at the edges and could therefore no longer bleed into the next step. This prevents the high-intensity regions of extending their range over hundreds of electronvolts, while almost entirely preserving the shape of the spectrum inside each step respectively. This solution was ultimately discarded though, as stepping through the spectrum in these fixed steps and thereby cutting off the tails of the Gaussians means not only modifying the actual Gaussian shape of the kernels, but it even does so irregularly depending on where inside the respective step each data point is. Data points close to the left boundary lose a larger part on the left side, and analogously for data points close to the right boundary. It was also noticeable that in this approach the graph evaluated from the KDE sometimes did not match up on step boundaries and therefore introduced minor artifacts that further have no physical interpretation. The mathematical validity of a KDE modification such as this is questionable at best.

The second possible solution, and with that also the final one used in the rest of the analysis, is to perform a Variable Kernel Density Estimation, meaning the adapting of the bandwidth in specific energy ranges throughout the spectrum. By combining knowledge from previous guesses at a sensible bandwidth value, a very low bandwidth in the milli electronvolts is employed for the high-intensity region up to 2500 eV, a medium sized bandwidth is used going into the endpoint region up to 2800 eV and finally the bandwidths are further increased in the background past 3100 eV. The resulting spectrum closely matches the histogram in high-intensity ranges, and properly smooths out the endpoint region as well as the background (figure 8).




Figure 6: KDE of spectrum applied with constant bandwidth of 0.001 eV




Figure 8: KDE of spectrum applied with variable bandwidth, endpoint region

Seeing as the smoothed out KDE in the background still regularly drops in between individual events, tests should be made to further increase the bandwidth in these areas. Due to time constraints these could not yet be executed as part of this thesis and instead the rest of the analysis was run in this configuration.




Figure 7: KDE of spectrum applied with variable bandwidth

4.3 Fraction of pile-up events in the background

One of the background sources in the spectrum are pile-up events. As discussed before, if two events are detected too close together in time on the same channel, the second voltage peak stacks on top of the tail of the previous one, and the energy reading is inaccurate. The Holdoff filter is an effective method for reducing this occurrence. However, if the two events happen within a time frame shorter than the time resolution of the detector, then it is not possible to tell both events apart, a trace can be seen that appears to carry approximately an energy corresponding to the sum of the energy of the single events. The spectrum resulting from all these events is called the unresolved pile-up spectrum, and as it is not discarded during data reduction we need to take it into account in the analysis.

Mathematically, the shape of the unresolved pile-up spectrum is described with the autoconvolution of the ^{163}Ho decay spectrum, but the total intensity is yet to be determined. As the decay spectrum ends at the Q-value of 2863 eV, the highest possible pile-up energy is double of this, and beyond that point the background cannot be influenced through this effect. To be able to tell the unresolved pile-up spectrum apart from different background sources, we first determine the average count rate in the high-energy range from 5900 eV to 7000 eV to be approximately $1.36 \cdot 10^{-2} \frac{\text{events}}{\text{eV}}$. This background rate is assumed to be constant throughout the entire spectrum, so by counting the events in the medium-energy range from 2900 eV to 5900 eV and subtracting the expected amount of events based on the background rate, we can determine the amount of events in this range that originate from the unresolved pile-up spectrum.

We find 68 total events, approximately 27 of which we link to unresolved pile-up in this way. Because we know the shape, but not yet the amplitude of the pile-up spectrum, we use this count to estimate the fraction of pile-up events in the background to be $f_{\text{pu}} \approx 2.447 \cdot 10^{-6}$. Multiplying the normalized autoconvolution of the spectrum by this value as well as the total amount of events in the overall spectrum, we have properly scaled our unresolved pile-up spectrum and can use this, as well as the constant background determined in the high-energy region, to model the background of the experimental data.

4.4 Detector response

As the detector used in the experiment does not operate at an infinitely high resolution, the data points in the spectrum are statistically offset to their real values. We model this effect by convolving the product of the phase space factor and the atomic physics function with a normalized Gaussian of which the width corresponds to the detector response.

To determine the detector response, we look at the M2-line of the ^{163}Ho decay spectrum. This line can be described by a Voigt function (4.1), a convolution of a Lorentzian and a Gaussian function.

$$V(x; \sigma, \gamma) = G(x; \sigma) * L(x; \gamma) \quad (4.1)$$

The Lorentzian corresponds to the intrinsic lineshape, while the Gaussian is caused by the aforementioned detector response. We can fit a Voigt function to the M2-line of the spectrum to extract the respective parameters, the intrinsic line width γ from the Lorentzian and the detector resolution σ from the Gaussian.

4.5 Endpoint region fit

The final step of the analysis takes place in the endpoint region of the spectrum, where we hope to see the greatest benefits of the unbinned analysis due to the low amount of events. It is here that we apply the fitting function described in equation 2.1, in order to acquire the fitting parameters, which can be used to further model the endpoint region.

The fit that will be performed is confined to limited ranges of energies in the spectrum, since the model used is tailored specifically towards the endpoint region only, and does not describe the spectral shape in any other location. This limitation poses no issue, since we established that the effect of the electron neutrino mass is felt most significantly in the endpoint region. Therefore, for this analysis specifically, no model is required to be given for the rest of the spectrum. With the background of the spectrum

as well as the detector response already determined, the fit will now be applied on the parameters of the respective models used for the atomic physics function. The fit is performed for both the histogram-based representation of the spectrum as well as the KDE counterpart, using the same initial parameters to allow for direct comparisons to be made. The fit is also repeated on multiple varying energy ranges since the exact boundaries of where our chosen model is capable of well-describing the data are not known exactly. Also for the KDE, we repeat the process for varying resolutions, meaning different amounts of points at which the KDE is discretized - something that is not possible at all in binned analysis without altering the bin width itself.

The different configurations for the analysis are the following. One is the original analysis method, a histogram with a bin width of 2 eV, resulting in 3500 data points in total from 0 eV to 7000 eV. The new KDE method is also tested, but on different resolutions of a similar count of data points as the histogram (3501), double the amount of data points (7001), five times the amount of data points (17501) and also 12001 data points. The reason for the point count to always be 1 more is that the KDE is also evaluated at one final data point at the end of the 7000 eV total spectrum range, while the histogram is not. Since we are not fitting in that area this is of no concern though, especially as the spacing of the individual data points remains exactly the same as on the histogram in the case of the 3501 points, and the expected fractions of this for the higher resolutions. Each of these 5 different configurations is run both with a quadratic fit and also an exponential fit, and every single combination of this is run on 4 slightly different fitting ranges. The lower boundary is 2450 eV, 2500 eV, 2550 eV and 2600 eV respectively, while the upper boundary is always 2780 eV.

5 Results

5.1 Detector response

The first occasion for testing and comparing results between the binned and unbinned analysis methods respectively is the determination of the detector response. A Voigt-function fit was performed (see figure 9) in the energy range from 1832 eV to 1839 eV. Due to the relatively high intensity of events in this region thanks to the M2-line, a lower bin width of 0.5 eV was chosen for the histogram, resulting in 14 bins and therefore 14 data points falling into the fitting region. The KDE is evaluated at 100 evenly spaced points between 1830 eV and 1842 eV, resulting in 58 data points falling into the fitting region.

For the fitting parameter σ of the Gaussian portion of the Voigt function, we find the values

$$\begin{aligned}\sigma_{\text{Hist.}} &= (1.5 \pm 0.7) \text{ eV}, \\ \sigma_{\text{KDE}} &= (1.46 \pm 0.15) \text{ eV}.\end{aligned}\tag{5.1}$$

Analogously, for the fitting parameter γ of the Lorentzian portion of the Voigt function, we find the values

$$\begin{aligned}\gamma_{\text{Hist.}} &= (5.3 \pm 0.7) \text{ eV}, \\ \gamma_{\text{KDE}} &= (5.50 \pm 0.14) \text{ eV}.\end{aligned}\tag{5.2}$$

As can be seen, in the case of both parameters, the two different spectral methods of representing the spectrum produce values that are not significantly different within their error margins. However, the parameter errors are noticeably lower on the KDE fit.

For the intrinsic line width γ_{M2} , a literature value of 6.0 eV as well as an experimental value of 4.8 eV are given with an uncertainty of < 1.0 eV [5]. Our results for the fitting parameter σ correspond to the detector response, and were used in further parts of the analysis.

5.2 Endpoint region fit

For the endpoint region fit, we have different means for quantifying the difference in the results between the analysis methods. We will be looking at the χ^2_{red} of the fits, the relative residuals, the parameter values as well as the relative parameter errors.




Figure 9: Voigt function fit of ^{163}Ho M2-line

5.2.1 Residuals

In table 1 we see the reduced chi-squared of all fits that were performed, compiled from the fitting results found in the appendix from figure 11 to 20. It is notable that in all configurations, χ_{red}^2 decreases as the lower fit boundary increases, or put differently it decreases as the fitting range shrinks on the lower end. As we are looking at the reduced form of chi-squared, we are taking into account the degrees of freedom of the fit. Since the values are approaching a value closer to 1, this suggests an overall improvement of the fit. A possible explanation may be the fact that both the exponential and the quadratic function do not perfectly describe the spectrum in the endpoint region, but are getting better when approaching the Q-value.

As for a direct comparison between the histogram-based fit and the KDE-based fit, it is apparent that the KDE achieves a χ_{red}^2 closer to 1 in all cases, even in the case of 3501 total evaluation points of the KDE, where the amount of points used for the fit is the same as for the histogram. In fact, higher resolutions of the KDE did not further improve the reduced chi-squared.

On the residuals of the histogram fits (in the appendix, figures 21 and 22), it is apparent for both the quadratic as well as the exponential fit and also across all fitting range configurations, that 1 residual very close to the endpoint has a value of around -4 , while all other residuals stay within the range of -1 to 1 . This outlier is the effect of a bin that got less events than the surrounding bins by statistical chance. It can be seen in figure 10 at 2770 eV. The KDE shows a small dip at this position, but the smoothing through the Gaussian kernels prevents the value from dropping nearly as far as the histogram. The same effect can be seen on multiple empty bins on energies

Table 1: Reduced chi-squared of fits

Spectrum	Fit	χ^2_{red} with lower boundary of			
		2450 eV	2500 eV	2550 eV	2600 eV
Hist, 3500 bins	Quad.	97.91	64.14	48.17	27.88
Hist, 3500 bins	Exp.	107.41	64.58	48.13	28.04
KDE, 3501 points	Quad.	60.17	32.38	9.71	6.15
KDE, 3501 points	Exp.	63.16	32.86	9.87	6.24
KDE, 7001 points	Quad.	62.81	36.72	9.71	6.14
KDE, 7001 points	Exp.	65.96	37.30	9.86	6.23
KDE, 12001 points	Quad.	62.93	36.95	9.76	6.17
KDE, 12001 points	Exp.	66.18	37.51	9.92	6.27
KDE, 17501 points	Quad.	63.09	37.09	9.79	6.18
KDE, 17501 points	Exp.	66.36	37.64	9.94	6.28

at or above 2800 eV. We can see on the relative residuals of all KDE configurations (figures 23 to 30) that the outlier at 2770 eV does not show up, all values stay within an interval from -0.7 to 0.3 .




Figure 10: Closeup view on the endpoint region with Histogram and KDE

5.2.2 Parameters

Finally, we will be taking a look at the fit parameters and their corresponding error values, which can be found in the appendix from figure 31 to 40. First of all one small note, the relative parameter error for the histogram fit of the quadratic function (see figure 31) found a parameter $a = 1.85 \cdot 10^{-13} \pm 0.08$ on the fitting range with

the highest lower boundary of 2600 eV, while the other ranges found a value for this parameter $a > 2 \cdot 10^{-2}$. This means the relative parameter error was multiple orders of magnitude above all other values and it was therefore hidden from the plot to benefit visibility.

As the parameters which are found for the different fitting ranges can sometimes be quite different, a more useful way of checking the results is by taking a look at the relative parameter errors. Due to less data points being used for fitting, we expect parameters determined on the shorter energy ranges to show larger relative errors, because less data results in lower significance of the results. For the parameters of the exponential fit, we see this happening in all cases for both parameters A and τ , both for the histogram-based fit as well as the KDE-based fits. On the other hand, while the quadratic fit indeed generally shows the largest relative parameter error for the smallest fitting range, for the histogram the second fitting range starting at 2500 eV is the smallest, and for the KDE the third fitting range starting at 2550 eV is either the smallest or at least smaller than the relative parameter error of the second range. We can see that in each of these cases the parameter in question has the highest absolute value out of the 4 fitting ranges, which of course leads to a decrease of the relative parameter error. As for why this happens in the first place, this is not entirely clear. It is possible that a local minimum is found in these cases which steers the parameter value away from what the other configurations are finding. Since this effect appears both in the histogram-based approach as well as the KDE-based one, one may assume that this is a general risk when using the quadratic function, and might hint at it not being an ideal fit function.

6 Conclusions

In summary, a concept has been worked out to implement the use of Kernel Density Estimation into the analysis of the ^{163}Ho decay spectrum endpoint region, with the aim of determining the effective electron neutrino mass. Different approaches were explored until a variable bandwidth KDE seemed most promising under the conditions of the spectrum. The detector response has been examined by performing fits on the M2-line of the spectrum. While not being the most rigorous analysis of the detector response, general improvements for the significance of the results over a histogram-based approach were apparent. Finally, fits of the endpoint region have been utilized to investigate the performance of a KDE in low-intensity regions. Through the inspection of fit residuals or the fit parameter errors, it was again possible to validate the benefits of a KDE over histograms in this scenario.

With these results, it is now clear that a Kernel Density Estimation is a viable alternative to a histogram for use as a means of representing the ^{163}Ho decay spectrum with the aim to analyze both high-intensity regions like the M2-line for the detector response, and also more importantly the low-intensity endpoint region which is crucial for determining the effective electron neutrino mass.

The performance of the KDE still has potential for improvements. Firstly, the evaluation of the KDE at a set of positions is a computationally intensive process. Loosely speaking, producing a histogram and a KDE on the same data set and resulting in an equal amount of points, the histogram might finish on the order of seconds, while the KDE might take an hour or longer. While this needs to be considered when choosing the representation method of the spectrum, fortunately it is not quite the end of the story. At the moment, each point of the KDE is evaluated in sequence, and for each point the Gaussian kernel of every single event is calculated before moving on to the next point. Since the order in which points are evaluated is irrelevant in the case of our method, it is possible to use multi-threaded computing to process as many points of the KDE in parallel as is possible on the hardware used for the analysis. Therefore, the generating of a KDE is highly scalable. Another task for optimizing the resulting KDE is to determine better methods of selecting a bandwidth near the endpoint region and in the background of the spectrum beyond the Q-value. Special care needs to be addressed to the fact that the endpoint region cannot be allowed to be smoothed out too far into the background, as this will lead to an incorrect fit and determination of the energy that is left to be attributed to the effective mass of the electron neutrino. Also, while it was possible to smooth out the background events to some extent, there are still parts in the background of the spectrum where the intensity fluctuates by multiple orders of magnitude due to larger gaps in energy where no events were detected at all. More testing with different bandwidths is necessary to ensure the highest possible

quality for the fits.

In the case of the histogram-based analysis, it was already clear that the best step to increase the accuracy of the results is to increase the total amount of events in the data set. This would allow for choosing smaller bin widths leading to more fitting points, while not risking the appearance of more empty bins. While the KDE does not run into the risk of producing empty bins in the same way a histogram does, it is still susceptible to regions of statistically lowered intensity. Further smoothing of the events can help, however higher accuracies are achieved with the lowest possible bandwidths that still result in a uniform curve while not yet being overly noisy. Therefore, an increase in the total event count detected in the experiment is still highly desirable for the significance of the results, regardless of the representation method of the spectrum.

7 Appendix




Figure 11: Histogram (3500 bins), quadratic fit




Figure 12: Histogram (3500 bins), exponential fit

Quadratic Fit / Autoconv + Const. Background




Figure 13: KDE (3501 points), quadratic fit

Exp Fit / Autoconv + Const. Background




Figure 14: KDE (3501 points), exponential fit

Quadratic Fit / Autoconv + Const. Background




Figure 15: KDE (7001 points), quadratic fit

Exp Fit / Autoconv + Const. Background




Figure 16: KDE (7001 points), exponential fit

Quadratic Fit / Autoconv + Const. Background




Figure 17: KDE (12001 points), quadratic fit

Exp Fit / Autoconv + Const. Background




Figure 18: KDE (12001 points), exponential fit

Quadratic Fit / Autoconv + Const. Background




Figure 19: KDE (17501 points), quadratic fit

Exp Fit / Autoconv + Const. Background




Figure 20: KDE (17501 points), exponential fit




Figure 21: Histogram (3500 bins), quadratic fit, relative residuals




Figure 22: Histogram (3500 bins), exponential fit, relative residuals




Figure 23: KDE (3501 points), quadratic fit, relative residuals




Figure 24: KDE (3501 points), exponential fit, relative residuals




Figure 25: KDE (7001 points), quadratic fit, relative residuals




Figure 26: KDE (7001 points), exponential fit, relative residuals




Figure 27: KDE (12001 points), quadratic fit, relative residuals




Figure 28: KDE (12001 points), exponential fit, relative residuals




Figure 29: KDE (17501 points), quadratic fit, relative residuals




Figure 30: KDE (17501 points), exponential fit, relative residuals




Figure 31: Histogram (3500 bins), quadratic fit, parameters




Figure 32: Histogram (3500 bins), exponential fit, parameters




Figure 33: KDE (3501 points), quadratic fit, parameters




Figure 34: KDE (3501 points), exponential fit, parameters




Figure 35: KDE (7001 points), quadratic fit, parameters




Figure 36: KDE (7001 points), exponential fit, parameters




Figure 37: KDE (12001 points), quadratic fit, parameters




Figure 38: KDE (12001 points), exponential fit, parameters




Figure 39: KDE (17501 points), quadratic fit, parameters




Figure 40: KDE (17501 points), exponential fit, parameters

References

- [1] Y. Fukuda¹, T. Hayakawa¹, E. Ichihara¹, K. Inoue¹, K. Ishihara¹, H. Ishino¹, Y. Itow¹, T. Kajita¹, J. Kameda¹ et al. (Super-Kamiokande Collaboration); Evidence for oscillation of atmospheric neutrinos; Phys.Rev.Lett.81:1562-1567; 1998.
- [2] J.D. Vergados (University of Ioannina, Ioannina, Greece), H. Ejiri (RCNP, Osaka University, Osaka, Japan), F Šimkovic (Laboratory of Theoretical Physics, JINR, Dubna, Moscow Russia); Neutrinoless double beta decay and neutrino mass; Int. J. Mod. Phys. E Vol. 25, 1630007; 2016.
- [3] J. Lesgourgues, S. Pastor; Massive neutrinos and cosmology; Physics Reports, Volume 429, Issue 6; 2006.
- [4] Ch. Schweiger et al., Nat. Phys. 20, 921–927; 2024
- [5] Ranitzsch, P. C. and Hassel, C. and Wegner, M. and Hengstler, D. and Kempf, S. and Fleischmann, A. and Enss, C. and Gastaldo, L. and Herlert, A. and Johnston, K.; Characterization of the ^{163}Ho Electron Capture Spectrum: A Step Towards the Electron Neutrino Mass Determination; Physical review letters 119, 122501; 2017.

Erklärung

Ich versichere, dass ich diese Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Heidelberg, den 01.04.2025,

A handwritten signature in black ink, appearing to read "J. Böhm". The signature is written in a cursive style with a horizontal line under the "h".