PRELIMINARY TESTS OF THE PEAKDOCTOR GAMMA RAY SPECTRAL FITTING ROUTINE

S.C. Myers, S.E. Garner
Los Alamos National Laboratory
P.O. Box 1663, Los Alamos, NM 87545

N.A. Sandoval
Texas Tech University
2500 Broadway, Lubbock, TX 79406

ABSTRACT

PeakDoctor is a new, highly sophisticated, gamma ray spectrum fitting routine that has been developed by the Advanced Nuclear Technology group at the Los Alamos National Laboratory. The routine is focused on precisely evaluating spectral data from both low resolution and high resolution detectors in order to maximize the accuracy of two key components of spectral analysis: 1) the net counts in all photopeak regions, and 2) nuclide identification. In order to accomplish these goals the software is highly interactive, thus providing the skilled analyst an opportunity to guide the routine through a number of different steps.

Questions arise about the consistency and repeatability of photopeak fitting results since they are strongly influenced by the analyst's choices. To answer these concerns we had eight different analysts, all with varying degrees of expertise and familiarity with the software, independently produce spectrum reports for 14 individual high-resolution spectra. Their results were compared to those produced by the creator of the routine, our designated 'expert' fit. Seven of the spectra were relatively simple to fit, having a limited number of photopeaks and few, if any, multiplet regions. The other seven spectra had very complex multiplet regions, changing continuums, and dynamic step peaks which made them very challenging to analyze. One highlight of these tests was that the energy calibration feature produces extremely accurate results for photopeak energies (generally within ± 0.2 keV of the known). The energy calibration uses multiple data points that are fit to a cubic polynomial across the entire energy span of the spectrum. This allows the assignment of fuzzy-logic scores for each potential nuclide match, which in turn enhances the strength of the nuclide identification process. The results for determining the net photopeak counts were very consistent for the each of the seven 'easy' spectra, regardless of user expertise. The photopeak results were also mostly consistent on the 'difficult' spectra with one exception; there were substantial discrepancies in the attempts to fit very complex x-ray regions from 90 to 120 keV in high-activity plutonium and/or uranium spectra. These discrepancies (which were frequently on the order of 30-50%) existed regardless of the level of user experience. Other than the complex x-ray regions, the average bias in the photopeak counts was generally within 1-3% of the expert evaluation regardless of user experience. We used these results to identify specific areas of future training for potential users of the routine. Finally, we also determined that simple spectra could be reliably evaluated in *PeakDoctor* with a non-interactive process that is faster and simpler to use.

INTRODUCTION

The ability to accurately evaluate raw spectral data is an important first step in producing high-quality radioassay results from gamma ray spectroscopy data. A number of different spectral events can contribute to the difficulty of evaluating the raw data, including overlapping peaks that form complex multiplet regions, large step peaks, backscatter regions, Compton edges, and inelastic scattering from neutron capture events. *PeakDoctor* is the name given to a new spectrum fitting routine that attempts to solve many of the inherent complexities in gamma ray spectra with a user directed platform. PeakDoctor has been under development by members of the Advanced Nuclear Technologies Group (N-2) at the Los Alamos National Laboratory (LANL) for the last two years and is being extensively used and tested by the Radioactive Waste Characterization Team of the LANL Solid Waste Operations Group (FWO-SWO).

PeakDoctor has four fundamental steps which allow user interaction: 1) energy calibration of the raw data, 2) continuum fitting below photopeaks, 3) photopeak fitting, and 4) nuclide identification (Note: Although it promises to be a very powerful tool, the nuclide identification process is still under development and was not evaluated in this study). That PeakDoctor is an interactive program which allows the analyst to guide it through the spectrum fitting process is a very desirable feature. However, the interactive feature also raises questions about whether the results are reproducible, and whether different analysts with different levels of spectrum fitting expertise can produce similar results. Is it likely that a bad result could be obtained by a poorly trained or inexperienced user? Another question is whether the total measurement uncertainty should also consider a separate peak fitting bias among the systematic biases of the analysis?

Prior to opening a spectrum, the user selects the detector that was used to collect the data. The detector selection provides PeakDoctor with the following response information that will help with the upcoming analysis: 1) preliminary energy calibration (keV/channel), 2) expected photopeak resolution (FWHM vs. energy), and 3) intrinsic efficiency vs. energy. After choosing the detector, the raw channel by channel data is loaded into PeakDoctor and the preliminary energy calibration is called upon to initiate a much more precise energy calibration. This involves fitting a cubic polynomial to a set of valid, pre-selected calibration peaks that are stored from a list of likely spectral candidates. The list of energy calibration peaks is fully editable and can include any photopeak energies that are available in the full photopeak library. In general the list of energy calibration peaks includes common background related gammas, likely fluorescent x-ray events from shielding materials, and the most common source/item photopeaks encountered during operational measurements. In the end, the process normally uses from 5 to 15 photopeaks throughout the range of data and results in a very accurate estimate of the true energy fingerprint of the spectrum.

The next step involves identifying the two fundamentally different types of channels in the spectrum: 1) those that contain photopeak counts and 2) those that contain Compton continuum counts. The Reveal Peaks step attempts to correctly identify each peak channel with an automated function that calculates a "running" second derivative to ordered data sets. Although this function is usually successful, the user is fully capable of changing these automated choices in order to enhance the precision of the upcoming continuum fit. This interactive process is often a source of minor discrepancies in the end results between different users. Once the peak

and continuum channels are separately identified, the analyst chooses from an editable list of spectrum types to determine what spectral features will need to be dealt with. Fitting the continuum involves fitting a series of many cubic polynomials that meet at pre-defined knot locations. The knot locations are initially determined by the step peaks, Compton edges, backscatter peaks, and manual knot locations that are listed in the spectrum type dataset. Although the pre-selected list of spectrum types usually has adequate information for continuum fitting, all of the choices are completely editable by the analyst. Once again, this frequently becomes a minor source of variation in the final output of the fitted photopeaks, especially in highly complex spectra.

The next step after achieving an acceptable continuum fit is to fit the photopeaks. The initial fit is performed automatically by the program, but then additional photopeaks can be "force fit" by the analyst if desired. The extra peaks that the analyst decides to fit are usually small peaks below a pre-selected threshold, or they are sometimes smaller peaks that are part of a larger multiplet region. The user is visually aided by a residuals scale that makes it easy to pick out the location of small peaks. This additional "force fit" step is also a source of variation in the final spectral report of different analysts.

Methods

Eight different individuals were asked to use PeakDoctor to fit 14 distinct spectra that were obtained with four different high-purity germanium (HPGe) detectors. Each of the individuals had at least some familiarity with the program so that they could independently fit each spectrum without requiring help from an expert analyst. The experience level with the software ranged from one individual who had been using the routine for over two years to a summer student who had only recently become familiar with its use. The analysts were instructed to fit each spectrum in the same manner as they would if they were to produce a customer report from the data. The results of each analyst were compared to our designated "expert" fit. The expert fit was performed by one of two individuals who are intimately familiar with the program and have been beta testing it since the inception. Although the expert fits are likely to produce the highest quality fits to the raw data, there is still a small qualitative aspect to those results that is impossible to evaluate. In other words, it is possible that fits from one or more analysts may occasionally be better than the fits from the expert analyst.

All of the spectra that were analyzed came from actual measurements of waste items with the exception of one quality control measurement of a small check source. Most of the spectra have between 40 and 100 photopeaks, although a couple of them have 25 or less. The spectra were qualitatively considered to be "difficult" or "easy" based upon the following criteria: 1) the total number of peaks, 2) the number and complexity of multiplet regions, 3) the size and number of step peaks, Compton edges, and backscatter regions. Of the 14 spectra, seven were considered to be "difficult" and seven were categorized as "easy". From a qualitative perspective, some of the difficult spectra were extremely difficult in terms of getting a good fit to the continuum, especially in the complex x-ray regions of plutonium and uranium spectra from 90 - 120 keV.

Several comparisons of the analyst's results were performed in order to determine the consistency of the fits and the accuracy of the results:

1. An evaluation of the accuracy of the energy calibration function.
2. The fitting bias of major photopeaks compared to the "expert" fit.
3. The fitting bias in the complex x-ray region compared to the "expert" fit.
4. Differences in the fitting bias between "difficult" and "easy" spectra.
5. Differences in the fitting bias between "experienced" and "novice" analysts.
6. A comparison of a more automated fitting approach, called the "One-Step" method, to the fully interactive methods of the eight analysts.

For the sake of consistency, we defined a major photopeak as meeting the following criteria: 1) a photopeak which would be likely to be used to produce an assay result and had at least 300 net counts, or 2) other photopeaks in the spectrum (e.g., background peaks and secondary gammas) that were not part of the complex x-ray region and had at least 1,000 net counts. To compare experienced and novice users we selected the three most experienced users to compare to the three least experienced. Hence, two analysts of medium experience were not included in that comparison.

The "One-Step" method removes most of the interactive opportunities by automating the steps for revealing the peaks, fitting the continuum, and fitting the peaks. However, all the choices for the spectrum type, step peaks, Compton edges, backscatter peaks, and manual knot locations must be made prior to initiating the routine.

Results

The eight analysts were instructed to fit each spectrum in the same manner as they would if they were to produce a customer report on the data. There were clear differences between the analysts in how rigorously they used the interactive features of PeakDoctor to fit the raw data. Some were very careful about how each peak in each spectrum was fit, while others were only concerned about getting a quality fit to photopeaks that would actually be used to generate a report. However, all the analysts were able to consistently produce very accurate energy calibrations to the raw data. As depicted in Fig. 1, the energy calibration of major photopeaks is nearly always within ± 0.2 keV of the known energy. An evaluation of 108 different photopeaks that were used in the energy calibration of the 14 spectra indicated that the intragroup variation among the 8 analysts had an average relative standard deviation (RSD) of just 0.016%. This confirms that the energy calibration, though interactive, is still very accurate and very stable as well.

The fitting bias of major photopeaks compared to the expert fit produced some interesting results. The average absolute percent difference for each spectrum of each of the eight analysts is presented in Table I. The average over all spectra ranges from 1.38% to 2.82%. On occasion, some fits of major photopeaks deviated from the expert fit by more than 40%. These were nearly always in the more difficult spectra and they usually occurred in complex regions as well. In plutonium spectra the first two major photopeaks to the right of the complex x-ray region are the 125.3 keV and 129.3 keV peaks. The calculated net counts in those peaks are greatly affected by the slope and direction of the continuum as it comes out of the nearly continuum-free region from 94 to 117 keV. Therefore, there tended to be a somewhat higher bias in those peaks than in others. Fitting large 59.5 keV peaks also proved to be more challenging than expected. The 59.5 keV peak frequently has more than 10^6 counts and a very small statistical uncertainty. However,

because of the size it is a dramatic step peak with a significant backscatter peak around 45 to 50 keV. Although the fitting bias for other very large peaks was normally less than 2%, the bias for 59.5 keV peaks was often between 2 – 8%.

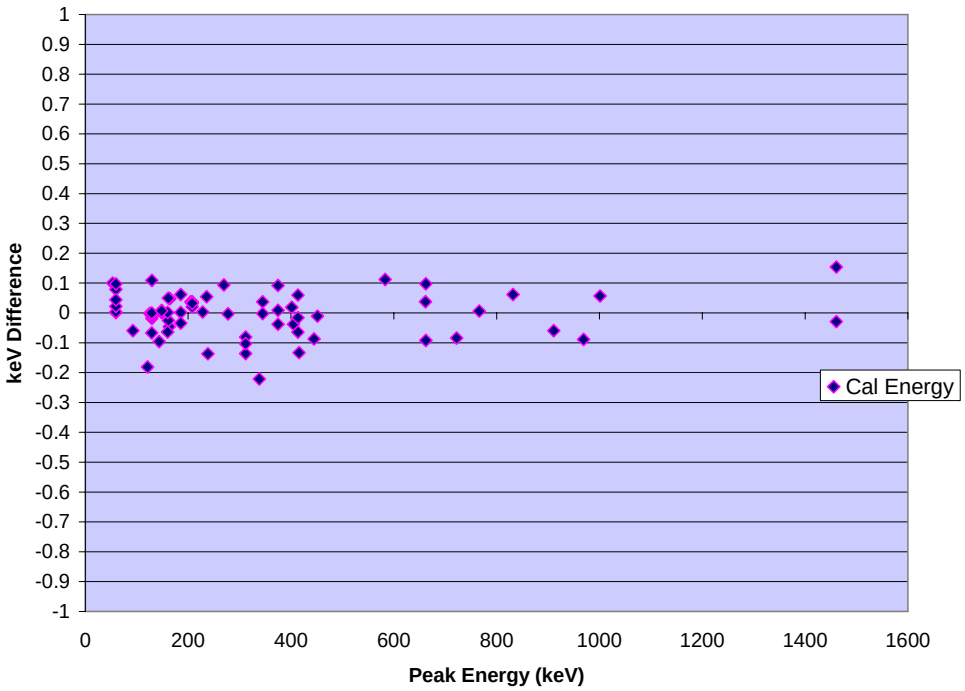

Fig. 1. Comparison of calculated energies to known energies

Table I. Average Fitting Bias for Major Photopeaks

Spectrum	1	2	3	4	5	6	7	8
1	1.48%	2.85%	6.22%	2.59%	3.27%	2.97%	1.26%	2.41%
2	2.55%	3.28%	3.41%	2.54%	4.59%	2.35%	2.04%	2.89%
3	2.37%	2.28%	2.70%	3.96%	5.81%	2.94%	2.74%	3.40%
4	1.67%	1.58%	1.33%	2.01%	3.15%	1.90%	2.10%	2.71%
5	3.03%	2.72%	5.68%	3.21%	5.74%	1.51%	3.70%	2.45%
6	0.25%	0.15%	0.53%	0.58%	0.72%	0.51%	0.12%	0.54%
7	0.69%	1.06%	1.04%	0.84%	1.25%	0.59%	0.94%	0.29%
8	2.53%	2.74%	2.95%	1.72%	7.96%	1.04%	6.24%	3.49%
9	0.00%	0.03%	0.00%	0.30%	0.30%	0.30%	0.03%	0.30%
10	1.80%	1.92%	2.49%	1.27%	2.83%	1.94%	1.23%	2.84%
11	0.99%	1.09%	1.80%	1.80%	1.67%	0.96%	1.35%	
12	0.02%	0.17%	0.20%	0.61%	0.72%	0.72%	1.82%	
13	0.68%	0.97%	1.42%	1.61%	1.00%	0.68%	1.29%	
14	1.20%	0.46%	0.56%	0.40%	0.53%	0.92%	0.64%	
Average	1.38%	1.52%	2.17%	1.67%	2.82%	1.38%	1.82%	2.13%

Note: Analyst 8 was unavailable to fit spectra 11 – 14.

Fig. 2 depicts the fitting bias of major photopeaks to the expert fit versus the %RSD of the photopeak area for all users. As expected, most peaks with a low statistical uncertainty also have a lower fitting bias. Deviations from this trend happened much more frequently with inexperienced analysts than with experienced ones. In fact, novice users tended overall to have a larger fitting bias than the more experienced users. The three most experienced users, all with at least two years of PeakDoctor experience, had an average fitting bias of 1.43% (ranging from 1.38% to 1.52%) while the three least experienced users had an average bias of 2.37% (ranging from 2.13% to 2.82%). Fig. 3 presents a comparison of fitting bias versus the %RSD of the photopeak for novice and experienced users.

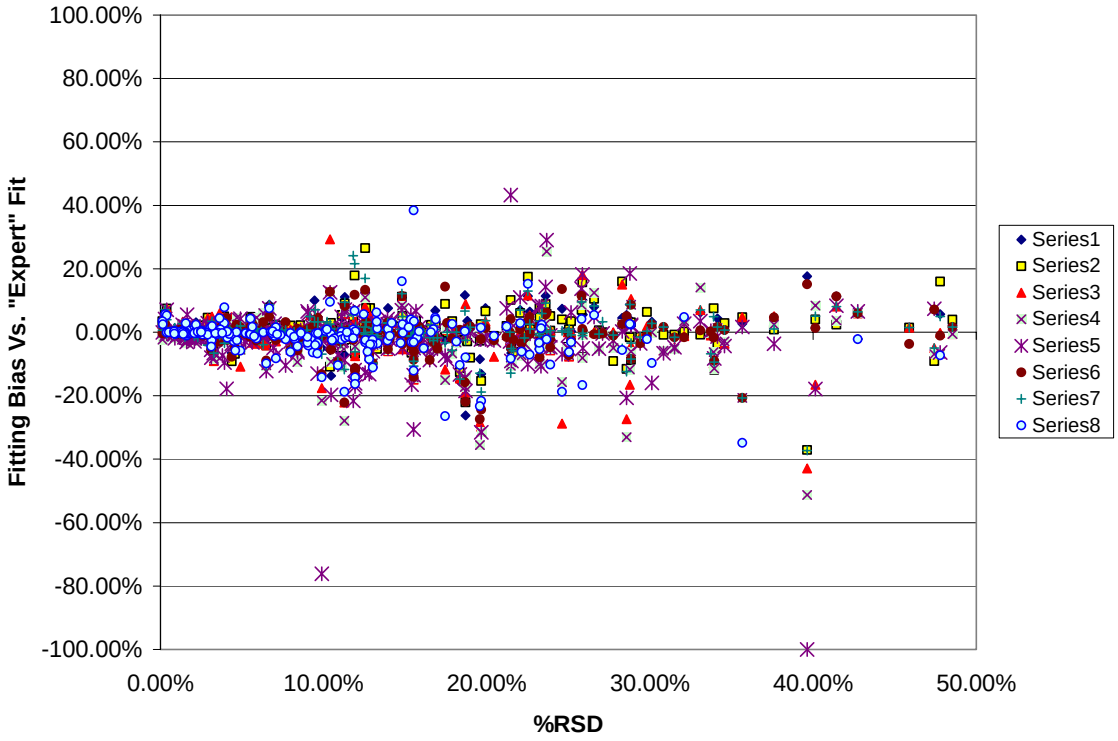

Fig. 2. Fitting bias of major photopeaks vs. %RSD of photopeak area

Another comparison we examined was the relative fitting bias in difficult vs. easy spectra. On average, the fitting bias of major photopeaks in difficult spectra was 2.73% while in easy spectra is was only 1.28%. The average fitting bias of difficult and easy spectra for each of the eight analysts is presented in Table II below. It is also very interesting to present the results of the non-interactive One-Step method compared to these actively fit results. For the easy spectra the One-Step method produced an average bias of 1.31% of which is very close to the mean value of the eight analysts. For the difficult spectra the One-Step method had an average fitting bias of 1.81% of which is better than any of the individual analysts. Clearly the One-Step method appears to produce adequate results regardless of the spectrum difficulty. However, this also assumes that all the important spectral features are properly identified (step peaks, Compton edges, backscatter peaks, etc.) before the process is initiated.

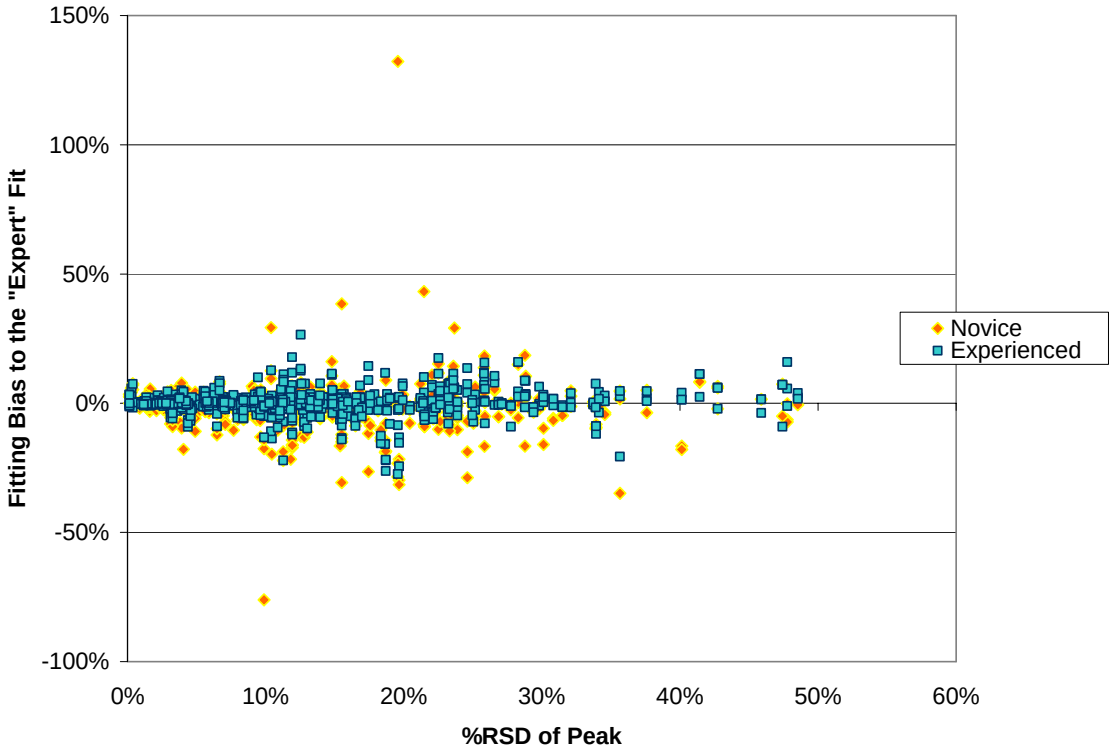

Fig. 3. Fitting bias of major photopeaks: Experienced users vs. novice users

Table II. Fitting Bias of Major Photopeaks in Difficult and Easy Spectra.

	Analyst 1	Analyst 2	Analyst 3	Analyst 4	Analyst 5	Analyst 6	Analyst 7	Analyst 8
Difficult	2.11%	2.29%	2.56%	2.73%	4.52%	2.32%	2.39%	2.95%
Easy	1.04%	1.07%	1.91%	1.61%	1.80%	0.85%	1.50%	0.49%

One area that proved to be quite problematic for the eight analysts was the complex x-ray region between 90 and 120 keV. The average fitting bias for the nine spectra that had x-ray regions was 15.03% with an average range from 11.91% to 19.86%. It was not uncommon for individual x-ray fits to deviate from the expert fit by more than 30% and sometimes by more than 100%. Fig. 4 presents the fitting bias in the x-ray region compared to the expert fit. A qualitative visual assessment of that chart makes it clear how dramatically these results varied from one analyst to the next. There are probably two reasons that explain the majority of these biases. First, although the x-ray region spans about 25 keV, there are only a few continuum channels between many peak channels. Therefore, the shape of the continuum below the peaks can be dramatically affected by the choices that can be made for knot locations between continuum splines. Another likely reason for some of the variability between analysts has to do with their focus, or lack of focus in some cases, on how carefully to fit the x-ray region. It is rare, if ever, that the peaks in that region are used to produce an assay result. As such, some analysts tend to ignore the quality of the continuum fit in that region. Others are simply concerned with the direction of the spline that exits the region out of the right hand side. Finally, the less experienced analysts have little training on how to deal with that complex region. Therefore, although it is a bit disconcerting

that there is such a large fitting bias in that region, it tends to have only a minor effect on the generation of quality assay reports.

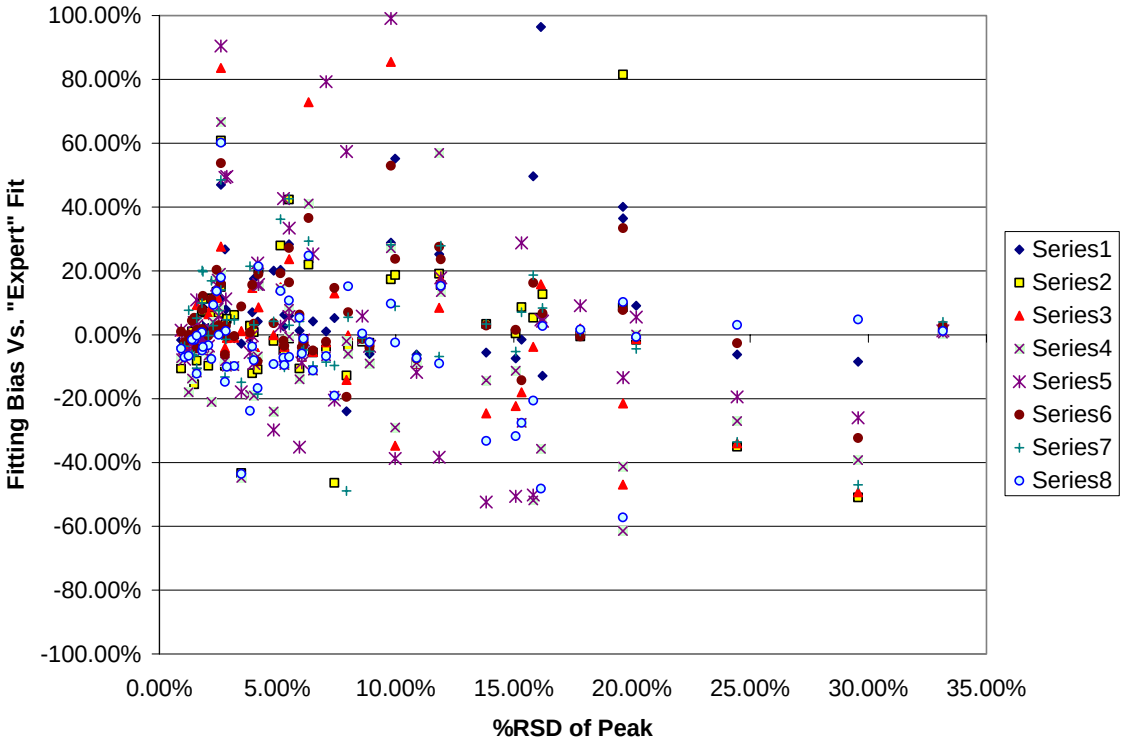

Fig. 4. Fitting bias of x-ray region peaks vs. %RSD of peak area

Discussion

In general, we found that different users of PeakDoctor were able to generate consistent results for the portions of the spectra that are critical for an accurate assay. One area that consistently performed well was the energy calibration feature. Although it can be interactive, our analysts found it to be very straightforward to use and our results strongly indicate that it is accurate and robust. The ability to consistently predict the true energy of a photopeak within ± 0.2 keV will greatly enhance the nuclide identification process of the routine.

Several processes that were evaluated produced results in the direction that we expected. There was a larger fitting bias of major photopeaks in difficult spectra compared to easy spectra. It appears that when there are more choices to be made about fitting step peaks, Compton edges, and backscatter peaks that there will ultimately be greater variability in the final result. However, it is important to point out that this variability does not just lie with the PeakDoctor routine. Most commercially available spectrum fitting routines can be manipulated by the user to produce different results, although most of them are not as innately interactive. A larger question concerning PeakDoctor was one that this study did not attempt to resolve: Can PeakDoctor produce better results that other currently available software? It would be wise to address this question in further studies.

Since experienced users had a lower fitting bias than novice users confirms that some minimum amount of training and experience is necessary in order to get excellent results. However, even our least experienced users were able to produce results that were relatively consistent with the expert results. Another expected trend was that the fitting bias varied more or less directly with the statistical error of the photopeak. Exceptions to this rule appear when the continuum below a large photopeak is difficult to fit consistently. Photopeaks to the immediate right of the complex x-ray region was one case were this trend was noticed. The 59.5 keV peak of Am-241 was another case as both the size of the step peak and the large backscatter peak to the left contributed to the challenge of fitting the continuum.

By far, the most inconsistent results occurred in the complex x-ray region of plutonium and uranium spectra between 90 to 120 keV. There tends to be only a handful of continuum channels in the entire region and small changes in the analyst's choices can lead to large changes in the magnitude of the final results. Fortunately, this undesirable inconsistency rarely leads to a poor result as the peaks in the x-ray region are not normally used to produce an assay value. Two exceptions are when the 92.6 keV peak of Th-234 is chosen for depleted uranium samples or when the 120.9 keV peak of U-234 is chosen for enriched uranium samples.

The most surprising result of this study was how well the automated One-Step process produced low bias results in all spectra, difficult and easy alike. It should take very little training to learn how to make use of this potentially powerful tool. Consistent results can be readily obtained as long as the spectrum type is properly identified along with other incidental spectral features. Another positive aspect of the One-Step method is that the total time to fit a spectrum is somewhat less than those that are actively fit. Furthermore, if a One-Step fit does not produce acceptable results then the analyst is free to re-fit the spectrum using the normal interactive steps.

All the users in this study appreciated the interactive features of PeakDoctor as it gives them excellent control of the fitting process. The visual quality of the routine was also found to be very appealing as well. Having the peak fits and continuum fits both clearly displayed along with the residual scale allows the analyst to make a very comprehensive qualitative assessment of the spectrum fit that is not present in other routines. Although interactive fitting usually takes longer than fully automated methods, analysts tend to have more confidence in the quality of the final result when they have greater control of the process. PeakDoctor appears to be a powerful tool when it is critical to get a highly accurate evaluation of your spectral raw data.