REU Student Research Papers – Summer 2025
University of Notre Dame – Department of Physics & Astronomy

REU Student Research Papers – Summer 2025
University of Notre Dame – Department of Physics & Astronomy

Accelerating Chemical Reaction Network Characterization: Machine Learning as a Quantum Analysis Replacement

MYLES CARLEY

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Brett Savoie, Dr. Ericka Miller

Abstract

Chemical Reaction Networks (CRNs) use edge-connected nodes to model the mutually interacting reaction pathways propagating within a chemical system. In this study, we test the performance of an Edge-featured Graph Attention Network (EGAT) at predicting the activation energies for chemical transformations in a β-D-glucose centered CRN. Our goal is to accurately reproduce the semi-empirical data produced by expensive quantum-chemistry techniques like density functional theory (DFT) for integrating to Yet Another Reaction Program (YARP), a CRN characterizer. The EGAT model performed within the bounds of irreducible error (4 kcal/mol) between testing sets. Root mean square error for glucose was 15.54 kcal/mol versus 11.75 kcal/mol for RGD1. 24.64% of reactants in the glucose dataset were accurately ranked in priority from least activation energy to most. The spearman's and kendall's correlation between the EGAT values and their accepted semi-empirical rankings are 0.7156 and 0.6379 respectively. We recommend finetuning the EGAT model to the β-D-glucose CRN to measure the degree of performance improvement. Other CRNs should also be explored for comparison, as RGD1 data is not a CRN, making it a poor benchmark for ranking accuracy.

1 Introduction

Modeling reaction propagation in chemical synthesis is becoming more important as demand for new drugs, biofuels, and materials increases. Since novel synthesis is often multi-species and multi-step in nature, it usually features many competing reactions. Each reaction generates a new product that introduces new permutations of competing reaction pathways. These new pathways tend to intersect and overlap, complicating and expanding the system into a chemical reaction network (CRN). This makes modeling reaction propagation both important and challenging for modern synthesis, as any computational model must generate, filter, and navigate this CRN [1]. The foundations for a CRN are the chemical species present and the activation energy between them. Activation energy is the minimum energy it takes for a reactant to transition into a specific product and the model must generate this in situ. While this has already been achieved, the industry standard has overwhelmingly employed quantum-chemistry techniques from the highest levels of theory to generate activation energies. These theory-based calculations are computationally expensive, often making large CRNs tedious for analysis [1].

This research aims to repurpose a machine learning model from Graph to Activation Energy

Models Easily Reach Irreducible Errors but Show Limited Transferability to replace the quantum-chemistry calculations for a β-D-glucose centered CRN. The goal is to rapidly generate activation energies with minimal increase in error. Program architecture and workflow were specifically designed for integration into the Savoie research group's Yet Another Reaction Program (YARP).

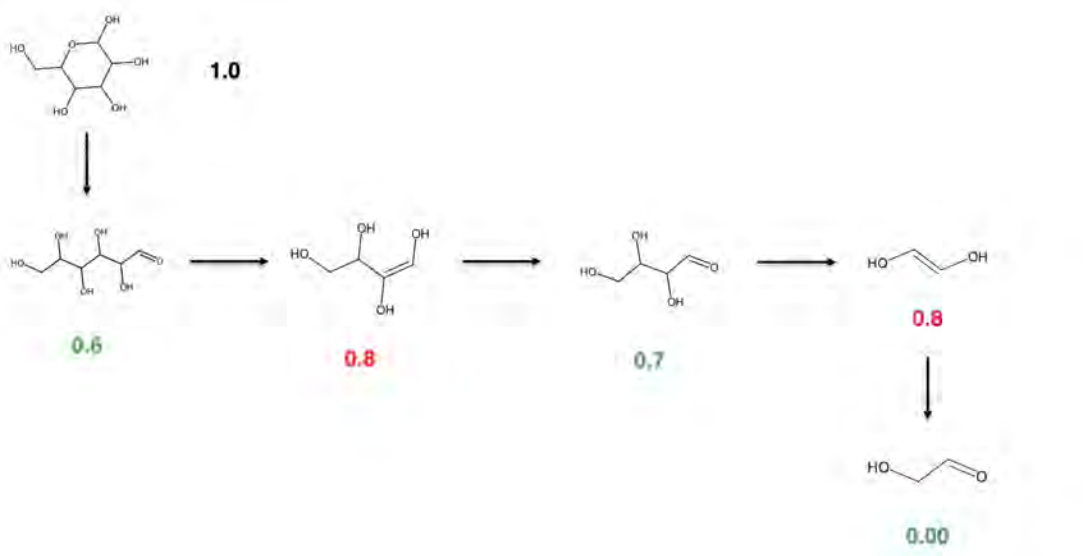

Figure 1: This image is a 1-dimensional example of how reaction products can function as reactants for other reactions, creating a sequential path towards a final product with the most favorable score (in this case the lower the score, the more favorable).

YARP is a CRN characterizer with three stages of automatic CRN characterization: product enumeration, network construction, and network analysis. The enumeration process functions under the break-n-form-n (bnfn) philosophy. If n is two, YARP performs a b2f2 analysis, which explores every combination of two bonds in the reactant. It then enumerates every molecule where those two bonds were broken, two new bonds were formed, and the Lewis octet rule has been satisfied. b2f2 and b3f3 enumeration are the only physically significant bnfn's which narrows down product possibilities substantially. However, with CRNs based on large organics, like β-D-glucose in this case, unreasonable quantities of potential products are formed that do not geometrically make sense even if they fulfill their octets [2]. These are unstable and would physically not exist under reaction conditions. YARP filters out these unreasonable products from the geometrically and energetically feasible ones [3]. YARP then iterates through the set of adjacent products, the

product neighborhood, and repeats the product enumeration process until the CRN has reached a given depth.

Once a CRN depth has been satisfied, reaction network construction iterates through the already generated nodes, and assigns activation energies to each of their edges. Two connected nodes represent a reactant and product pair; the edge is the forward activation energy for the transformation to occur, as depicted in Figure 3. YARP currently uses a transition state search to generate the activation energies, which requires computationally expensive density functional theory (DFT) originating from quantum-chemistry theory [2]. This is where EGAT, an Edge-featured Graph Attention Network, can substitute DFT driven transition state searches to dramatically improve program speed.

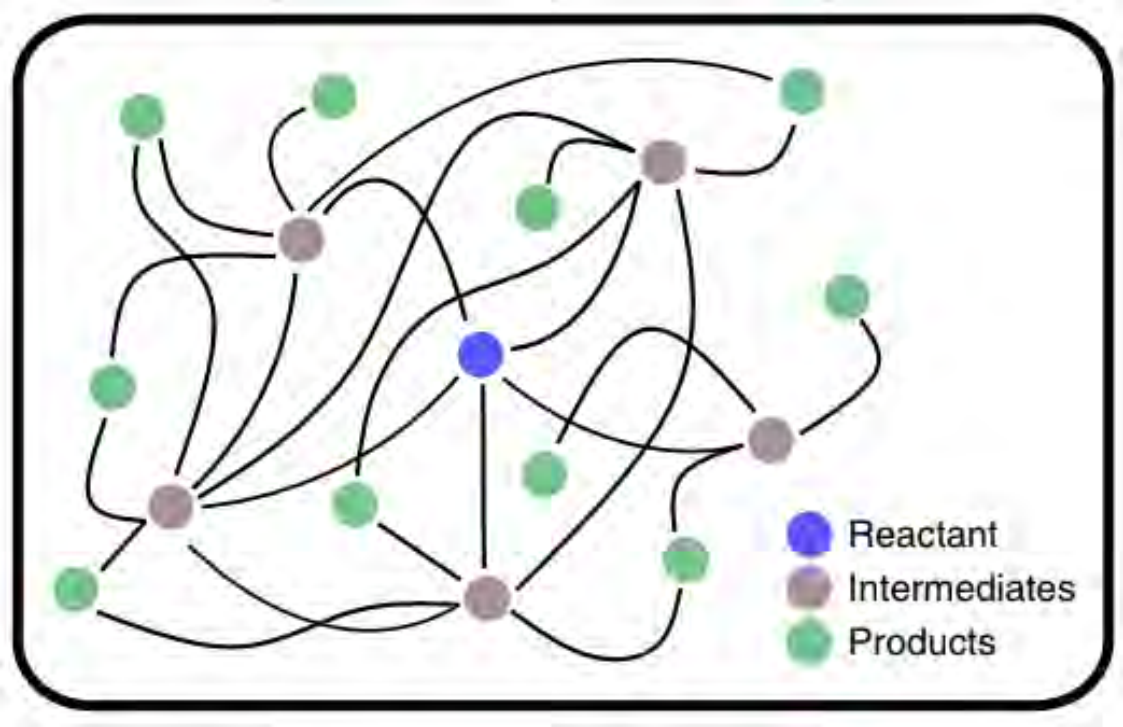

Figure 2: Depiction of how CRN intermediates, reactants, and products can intersect.

EGAT was trained to predict reaction properties, including activation energy, on the RGD1 dataset [5]. This dataset is one of the largest and most chemically diverse public datasets containing 176,992 chemically unique reactions and their characteristics. Since YARP is intended to be used on a broad range of chemical systems, RGD1 was selected as the best dataset for training due to

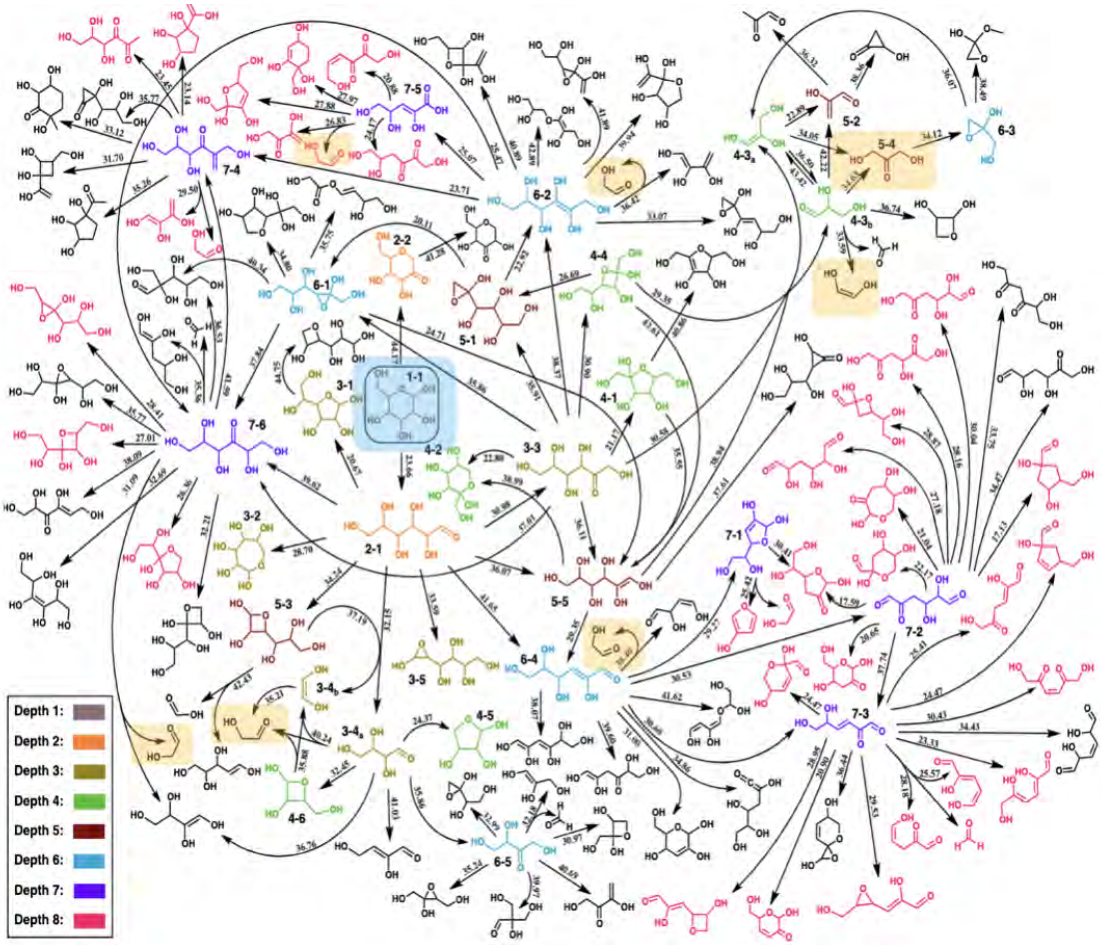

Figure 3: This image illustrates a selected sample from the β-D-glucose centered CRN. β-D-glucose is highlighted in blue left of center. The arrows and their attached weights represent chemical reactions and their activation energies. The color coding identifies the earliest depth at which a molecule occurred. Adapted from [4].

its inherent breadth [6]. 4 kcal/mol was preliminarily selected as an acceptable level of variance between datasets for the EGAT model because it is the accepted margin of irreducible error (error that is statistically impossible for a model to overcome through continued training) in machine learning based chemical property prediction. [5]. Additionally, the ranking of neighboring product nodes is more important than the accuracy of their energy values. During network analysis, the final stage of CRN characterization, YARP will search for the most energetically favored pathways between a given pair of reactants and products. YARP primarily needs to know which molecules to prioritize and explore first, not the explicit value associated with it. If the ranking of activation energies is correct even if the values show significant error, YARP's pathfinding algorithms can still

function. This makes precision more important than accuracy for EGAT-YARP integration [2].

This study investigates an already generated CRN propagating from β-D-glucose pyrolysis. β-D-glucose is of academic interest due to consistent failures to predict the high experimental favorability of hydroxymethylfurfural synthesis. This form of glucose is also easily accessible from cellulose and a ubiquitous biomass feedstock, making its pyrolysis academically relevant [4].

2 Methods

The β-D-glucose dataset generated in Deep Reaction Network Generation of Glucose Pyrylosis was the test set for activation energy prediction using EGAT. It features data on glucose pyrolysis which is already being explored in academia. The data contained 855 reactant and product pairs, each representing a single step in a reaction pathway. The same paper produced these data through DFT calculated transition state searches of the chemicals' potential energy surfaces [4]. This semi-empirical data was the benchmark of comparison for the EGAT model.

The EGAT model was pulled from Graph to Activation Energy Models Easily Reach Irreducible Errors but Show Limited Transferability. Out of prediction models produced in this study, the one reported with the best RGD1 performance was selected for β-D-glucose CRN benchmarking [5]. After data was generated, mean squared error was used as the primary assessment of accuracy. Spearman's rho and Kendall's tau correlations were used as non-parametric tests to assess the accuracy of the final rankings. Auxiliary python scripts enumerated the product neighborhoods (all the potential products) for a given reactant in order of increasing semi-empirical activation energy. This served as the "accepted" ranking from which spearman and kendall correlations compared the EGAT values. If they tend to monotonically increase with the semi-empirical activation energies, rho and tau approach a value of one, reflecting a more accurate ranking. Additionally, the percentage of reactants whose neighboring products were ranked perfectly with EGAT data was calculated. These steps were repeated on a comparably sized sample (177 reactions) from the RGD1 dataset for comparison.

3 Results

EGAT's performance on the glucose dataset is inferior to its performance on the test set from RGD1. For the product neighborhoods, EGAT had a 24.64% and 99.43% perfect rank accuracy for the glucose CRN's and RGD1's datasets respectively. Those are the percentage of reactants whose product neighborhoods were ranked with no errors. Spearman's rho and Kendall's Tau for the glucose set were 0.7156 and 0.6379 for the glucose set. Meanwhile, the RGD1 test set scored perfectly.

The RGD1's high performance, as illustrated in Figure 4, is to be expected as it is not a CRN with reaction pathway intersection, nor does it have the glucose dataset's depth. Many RGD1 nodes only have one neighboring product, favoring higher scores for rho and tau [6]. Moreover, because the glucose set's reactants on average had three to four neighboring products, 24.64% success rate is promising in comparison to the $\frac{1}{6}$ and $\frac{1}{24}$ random chance of randomly predicting the rankings [4]. While RDG1's sample set is low-depth and not a great benchmark for rank accuracy, it is the most apt benchmark for raw activation energy prediction because the data was sampled from EGAT's original training data. Root mean square error for the glucose set was 15.54 kcal/mol versus the RGD1 test set's 11.75 kcal/mol. Since the glucose performance is within the 4 kcal/mol range of the RGD1 dataset, this is an acceptable margin of error [5].

4 Conclusion

The EGAT model performed within the bounds of irreducible error (4 kcal/mol) between the β-D-glucose data and training data sample. Root mean square error for glucose versus RGD1 was 15.54 kcal/mol and 11.75 kcal/mol respectively. For the glucose dataset, only 24.64% percent of the rankings were completely accurate, though rankings still trended towards accuracy with spearman's and kendall's correlation proportions of 0.7156 and 0.6379. Since the β-D-glucose dataset on average had product neighborhoods between three and four molecules, these rank correlations are acceptable compared to the $\frac{1}{6}$ to $\frac{1}{24}$ chance of randomly predicting rank accurately. This discrepancy

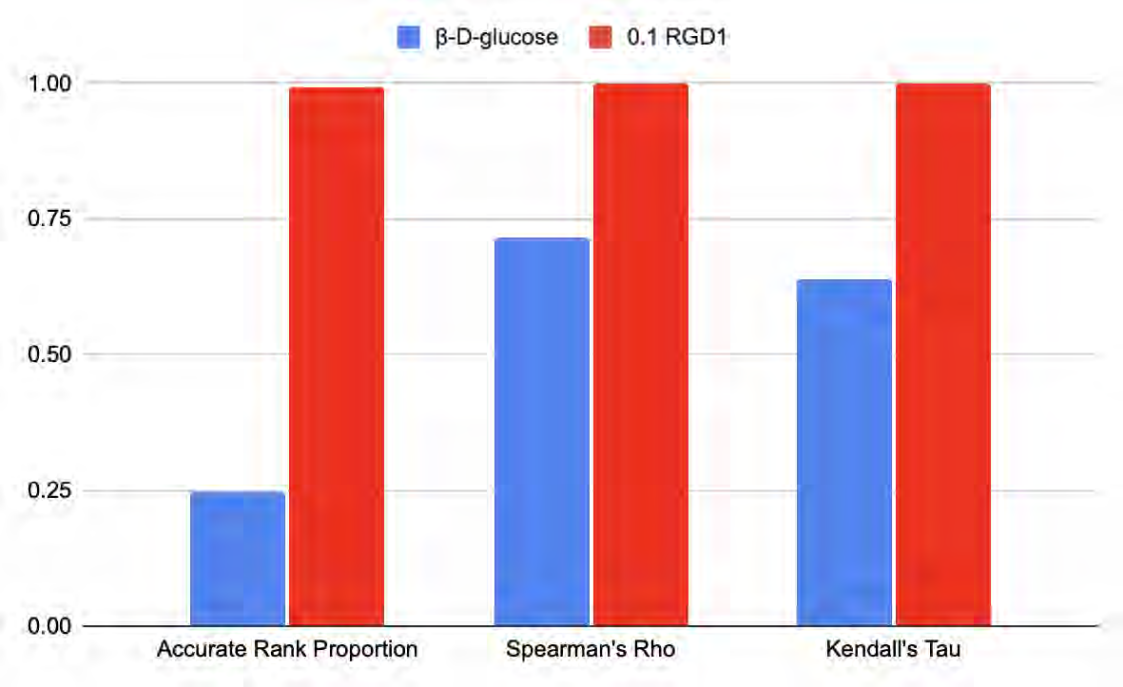

Figure 4: Column comparison of EGAT's ranking performance on the RGD1 sample data (red) and the β-D-glucose data (blue). The accurate rank proportion is the proportion of reactants whose product neighborhoods were correctly ranked. Spearman's and Kendall's correlations are also depicted on the same axis.

between data sets likely stems from the lack of transferability of the EGAT model.

Transferability is the ease with which a model can be used on data that is similar to but distinguished from its training data. This study should be repeated with a version of EGAT that is finetuned to the β-D-glucose CRN. Finetuning is the micro-training of an already completed machine learning model to adjust it towards a more specific data set. This could dramatically improve EGAT's performance on CRNs especially since most CRNs have a greater degree of chemical similarity and specificity, lending themselves to finetuning. While finetuning is computationally cheaper than starting training from scratch, it is not trivial. Finetuning a new EGAT variant for every incoming dataset would not align with the goal of trivial activation energy generation as it would still take large computing resources. However, the process of finetuning a new EGAT model for each new dataset could still be less expensive than using a rigorous DFT-driven transition state search [5].

Due to its lack of depth, the RGD1 data is not a good comparison for ranking accuracy so another

Figure 5: Column comparison of EGAT's activation energy prediction performance on the RGD1 sample data (red) and the β-D-glucose data (blue). Data is root mean squared error. The displacement bar is 4 kcal/mol.

CRN should be used for comparison [6]. A pyrolysis CRN of a similarly sized organic molecule would be promising, though published data sets could be limited in number.

YARP at its most robust uses an assortment of data alongside activation energies to model CRN propagation, so the error observed in this study may be insignificant [7]. Since EGAT appears semi-viable, it should be integrated into the YARP workflow to investigate the CRN characterizer's performance with EGAT activation energies.

5 Acknowledgments

I would like to thank Prof. Brett Savoie and the Savoie Research Group for incorporating me into a delicate but rewarding part of their workflow. Special thanks to Dr. Ericka Miller and Bryan Piguave for mentoring me in YARP, EGAT, and professional research practices.

References

[1] M. Wen, E. W. C. Spotte-Smith, S. M. Blau, and P. L. Freddolino, Nat Comput Sci **3**, 12 (2023).

[2] Q. Zhao and B. M. Savoie, Nat Comput Sci **1**, 479 (2021).

[3] Q. Zhao, H.-H. Hsu, and B. M. Savoie, J. Chem. Theory Comput. **18**, 3006 (2022).

[4] Q. Zhao and B. M. Savoie, Proc. Natl. Acad. Sci. U.S.A. (2023).

[5] S. M. Vaddadi, Q. Zhao, and B. M. Savoie, J. Phys. Chem. A **128**, 2543 (2024).

[6] Q. Zhao, S. M. Vaddadi, M. Woulfe, and B. M. Savoie, Sci Data **10**, 145 (2023).

[7] M. Woulfe and B. M. Savoie, J. Chem. Theory Comput. **21**, 1276 (2025).

Fabrication and Characterization of a Superconducting Resonator for ESR Measurements

LAUREN CATALANO

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Dafei Jin

Abstract

The development of scalable quantum computing platforms relies on the ability to achieve long coherence times in qubit systems. Electron spin resonance (ESR) offers a powerful method for probing spin dynamics and coherence, particularly when implemented with high-quality microwave resonators. In this project, we fabricated and characterized a superconducting microwave resonator designed for ESR measurements of electron spins on solid neon—a promising qubit platform developed by Prof. Dafei Jin's group. The resonator was patterned via optical lithography and reactive ion etching on a niobium-coated substrate. Characterization using a dilution refrigerator and vector network analyzer revealed a resonance frequency near 6.14 GHz and a distinct ESR signal at 216.8 mT, with a peak internal loss rate of approximately 203 kHz. These results support the potential of spin-based qubits on solid neon and contribute to efforts toward scalable, high-coherence quantum computing.

1 Introduction

The development of practical quantum computers is critically dependent on the ability to preserve quantum information over time. One of the key challenges in this pursuit is maximizing the coherence time, the time during which a quantum bit, a qubit, can maintain its quantum state to a high degree before it decays significantly. During this time the qubit maintains its ability to exist in a superposition of states and to be entangled with other systems. Among various physical platforms for qubits, spin-based systems, such as electron spins in solid-state environments, are particularly promising due to their long potential coherence times.

A powerful technique for probing and extending spin coherence involves utilizing electron spin resonance (ESR), which enables precise control and measurement of spin dynamics. Central to this method is the use of microwave resonators, which enhance the interaction between spins and electromagnetic fields, allowing for sensitive ESR detection and manipulation.

This project focuses on the fabrication and characterization of a custom microwave resonator tailored for ESR experiments. The goal is to optimize the resonator's performance to facilitate high-coherence spin measurements, contributing to the broader effort of building scalable and reliable quantum computing architectures.

1.1 Zeeman Splitting and Spin Resonance

Electrons with opposite spins occupying the same energy level will experience a split in energies with the application of an external magnetic field. This is known as the Zeeman effect, and the energy difference is given by the following equation, where g_e is the electron's "g-factor," μ_B is the Bohr magneton (a constant), and B_0 is the magnetic field.

$$\Delta E = g_e \mu_B B_0 \tag{1}$$

In continuous-wave measurements of microwave reflection, microwave photons with a given frequency f are sent in to a resonator and the magnetic field is varied. Once the magnetic field reaches the condition where $hf = g_e \mu_B B_0$ (the energy of a photon equals the energy difference in Zeeman splitting), an electron will absorb the energy from that photon. Since the energy is absorbed by the electron, it is lost from the resonator. This is indicated by κ_i, the internal loss rate. The magnetic field at which the internal loss rate peaks is the value of B_0 that satisfies the ESR condition. The peak of κ_i indicates maximum absorption of energy by the electrons.

2 Background

Prof. Dafei Jin's group at Notre Dame recently made an experimental breakthrough in developing a new kind of quantum bits (qubits) – the fundamental building blocks of quantum computers. They managed to trap single electrons on an ultraclean solid neon surface and use microwave photons in a superconducting quantum circuit to control and read out the qubit states [1]. Their electron-on-solid-neon (eNe) qubits have shown remarkable performance that is orders of magnitude better than many existing qubits and rivals the best ones made by Google and IBM [2]. Their eNe qubit platform has shown great promise in the development of a scalable architecture for future quantum computers at a much lower cost than any other platform.

Measuring the electron spin resonance and obtaining the spin coherence time on a solid neon surface is the most important step toward building electron spin qubits in this system. If a spin

coherence time on the order of 0.1s can be verified, as that has been theoretically predicted, it will represent a 1000-fold performance improvement of the electron charge qubits in this system that has been demonstrated by Prof. Jin's group [3]. Electron spin qubits with such a long coherence time will have orders of smaller operation errors than any existing qubit platforms have.

3 Methods

A series of nanofabrication techniques can be employed to actually make the resonator. First, a wafer of sapphire or silicon coated in Niobium (Nb), a superconductor, is cleaned and evenly coated with a layer of photoresist. Optical lithography can be used to transfer the desired pattern onto the sample. With the maskless aligner lithography machine in the Notre Dame Nanofabrication Facility, an array of micromirrors is positioned such that the sample will be exposed to light only where your pattern specifies. Developing the sample removes the photoresist in the exposed areas, revealing the Nb underneath that is to be etched.

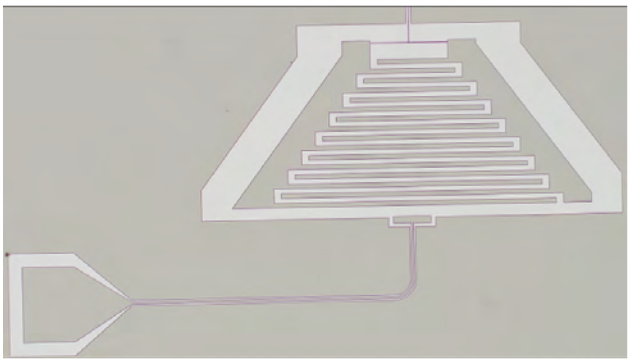

Figure 1: A microscopic view of part of the resonator, an interdigitated capacitor and inductive wire, before etching.

Reactive Ion Etching (RIE) is the preferred method of removing Nb from the substrate. The sample is placed on a platter in a vacuum chamber with a plasma that is created by applying an RF signal. Reactive radicals bounce around the chamber and are deposited on the wafer platter, leading to a negative charge on the platter. This creates a voltage difference and attracts the ions to the wafer platter, where they collide with the sample to be etched. The ions react with the material (Nb), removing it from the sample. RIE is more anisotropic than other wet etching techniques;

it etches mostly down instead of in all directions, due to the ions traveling directly down before colliding with the Nb. This makes it useful for high precision etching.

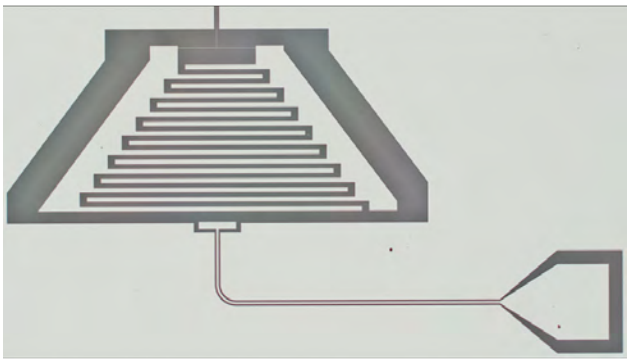

Figure 2: A microscopic view of part of the resonator, an interdigitated capacitor and inductive wire, after etching.

4 Results

Using the dilution refrigerator in Prof. Jin's lab, we were able to characterize the fabricated resonator and measure its properties. A vector network analyzer (VNA) was used to conduct the microwave transmission/reflection measurements.

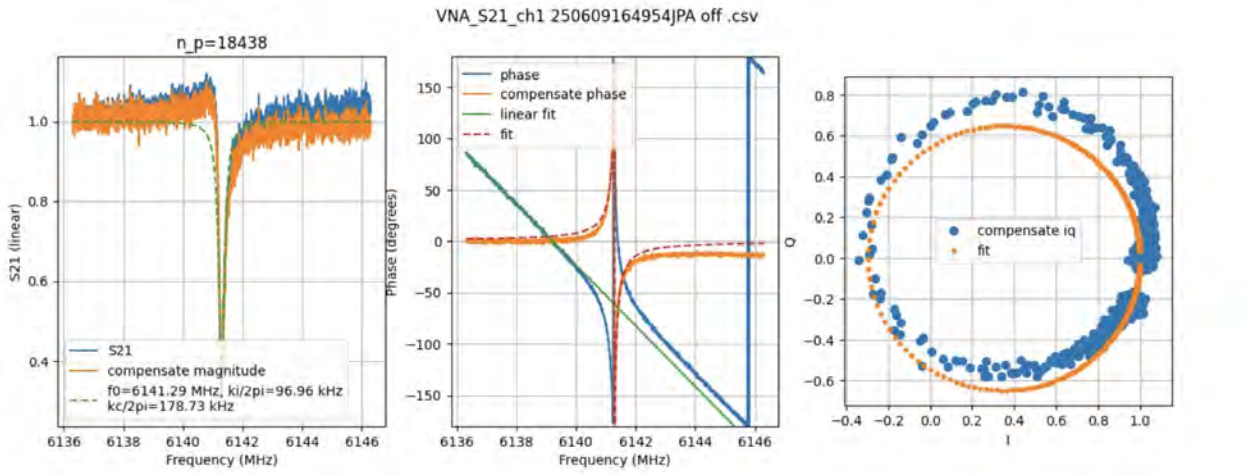

Figure 3: Input-output characterization of resonator at zero magnetic field

Figure 3 graphically represents the resonance frequency of the resonator. S_{21} measures how much of the input signal at port 1 makes it through to port 2. This shows a clear resonance fre-

quency of 6141.29 MHz, indicating energy is being absorbed or reflected at that frequency — the eigenfrequency of the resonator.

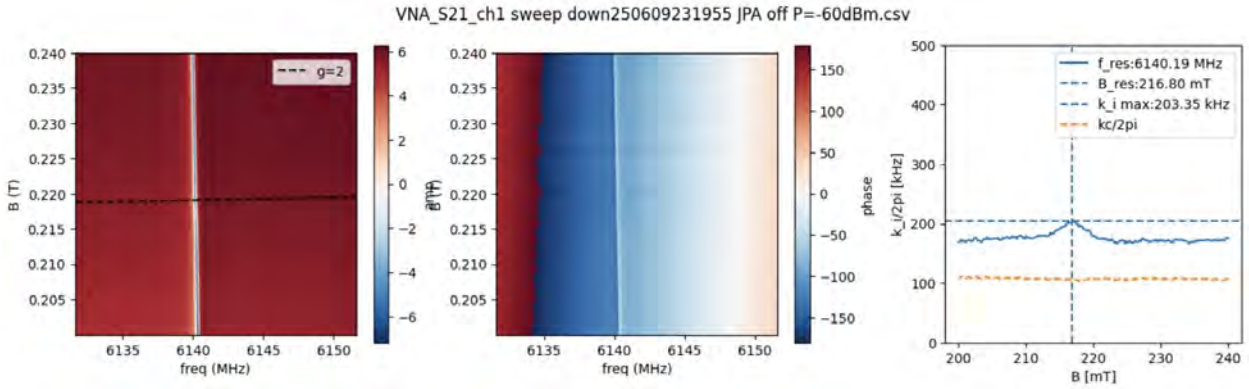

Figure 4: Continuous wave measurements of the microwave reflection as a function of the magnetic field

Figure 4 represents the results of the reflection measurements while sweeping the megnetic field. The leftmost plot shows the magnitude of the transmitted microwave signal (S_{21}) as a function of both frequency and applied magnetic field. A bright vertical line indicates the resonator's fixed resonance frequency. A horizontal dashed line near 0.216 T marks the magnetic field where the electron spin resonance (ESR) condition is expected for $g_e = 2$, according to the resonance condition $hf = g_e \mu_B B_0$. The middle plot is similar, but shows the phase of the transmitted signal instead of its amplitude.

The plot on the right is the ESR spectra. This graph plots the internal loss rate κ_i, which corresponds to how much energy is absorbed from the resonator, as a function of magnetic field. A distinct peak near $B = 216.8\,\text{mT}$ indicates resonant absorption of microwave energy by electron spins, which is the signature of ESR. The corresponding resonance frequency is $f_{\text{res}} = 6140.19\,\text{MHz}$, which matches the features observed in the top plots. The vertical dashed line at $B_{\text{res}} = 216.80\,\text{mT}$ marks the ESR condition. The internal loss rate κ_i peaks at 203.35 kHz, but subtracting the baseline value of approximately 180 kHz reveals that the ESR-induced absorption contributes an additional loss of $\Delta\kappa_i/2\pi \approx 23.35\,\text{kHz}$.

A flat dashed orange line shows the coupling loss rate $\kappa_c/2\pi$, which stays constant since it

doesn't depend on magnetic field.

5 Conclusion

The successful fabrication and measurement of a superconducting microwave resonator optimized for ESR experiments mark a critical step in supporting the development of spin qubits on solid neon. The observed ESR signature at 216.8 mT and resonance frequency of 6.14 GHz confirms coupling between the resonator and spin ensemble, enabling precise spin readout. The peak internal loss rate provides evidence for resonant energy absorption, validating the design and functionality of the device.

These results lay the groundwork for future experiments aimed at measuring spin coherence time and optimizing spin control in the eNe platform. However, more work remains to be done to refine resonator quality factors and reduce sources of decoherence. Continued development and characterization will be essential to fully realize the potential of this platform for scalable quantum computing.

6 Acknowledgments

I would like to thank the University of Notre Dame College of Science for supporting this research through the 2025 College of Science Summer Undergraduate Research Fellowships (SURF) program. I am especially grateful to Professor Dafei Jin for his mentorship, and to Dr. Shan Zou for her consistent guidance throughout this project. I also acknowledge the use of the Notre Dame Nanofabrication Facility, where all fabrication was performed. This work would not have been possible without access to its equipment.

References

[1] Zhou et al., "Single electrons on solid neon as a solid-state qubit platform," **Nature 605**, 46–50 (2022)

[2] Zhou et al., "Electron charge qubit with 0.1 millisecond coherence time," **Nature Physics 20**, 116–122 (2024)

[3] Qianfan Chen et al., 2022 *Quantum Sci. Technol.* **7** 045016

[4] A. Salari, *Microwave Techniques in Superconducting Quantum Computers* (Artech House, Boston, 2024)

Single Crystal Synthesis and Characterization of $HoMn_6Sn_{6-x}Ga_x$

BREANNA DRAUDEN

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S):
Prof. Nirmal Ghimire
Dr. Sekh Jamaluddin
Mohamed El Gazzah

Abstract

Materials in which magnetic atoms are arranged in a kagome lattice form an intriguing class of compounds, as they can host a variety of exotic electronic and magnetic properties. Among these, the $R\mathrm{Mn_6Sn_6}$ family (where R is a rare-earth element) has been extensively studied. These compounds have drawn significant attention due to their unique crystal structure, which can be tuned to stabilize complex magnetic textures such as skyrmions. Recently, the Ghimire group engineered skyrmion bubbles in $\mathrm{TmMn_6Sn_6}$ by partially substituting Sn with Ga. In this system, skyrmion bubbles were observed only in very thin samples (with thicknesses below 100 nm), and an external magnetic field was required to stabilize the skyrmion lattice. The formation of skyrmions was attributed to spin reorientation induced by Ga doping. In the present study, we extend this investigation to a closely related compound, $\mathrm{HoMn_6Sn_6}$, to explore whether a similar spin-reorientation effect can be achieved. To this end, we synthesized single crystals of $\mathrm{HoMn_6Sn_{6-x}Ga_x}$ and investigated their structural and magnetic properties. Rietveld refinement of powder X-ray diffraction data confirmed that the crystals retain the hexagonal structure upon Ga substitution. Magnetic susceptibility measurements revealed that Ga doping in $\mathrm{HoMn_6Sn_6}$ leads to an increase in the magnetic transition temperature, suggesting that chemical tuning could again play a key role in manipulating the magnetic ground state.

1 Introduction

Kagome magnets have emerged as fertile platforms for exploring exotic quantum states in condensed matter physics. The unique geometry of the kagome lattice induces exotic quantum interactions, including geometrical frustration. These led to a variety of intriguing phenomena, including the Dirac point, flat bands, charge density wave, and non-trivial topological spin textures such as magnetic skyrmions [1–11]. The kagome magnet is characterized by a two-dimensional lattice composed of corner-sharing triangles arranged in a hexagonal pattern, as shown in Fig. 1. Recently, various classes of kagome materials have been extensively studied; among them, the $R\mathrm{Mn_6Sn_6}$ family has attracted significant attention due to its potential to host a range of intriguing physical phenomena. In this class of materials, the Mn atoms form the kagome layer, while R can be substituted with rare earth elements, allowing for chemical and physical tunability [5; 7; 12]. Magnetic and electronic properties are strongly influenced by the selection of the rare earth element. Furthermore, the physical properties of these materials can be tuned by substituting Sn with other p-block elements [2; 13; 14]. Recently, it has been reported that Ga doping at the Sn site in

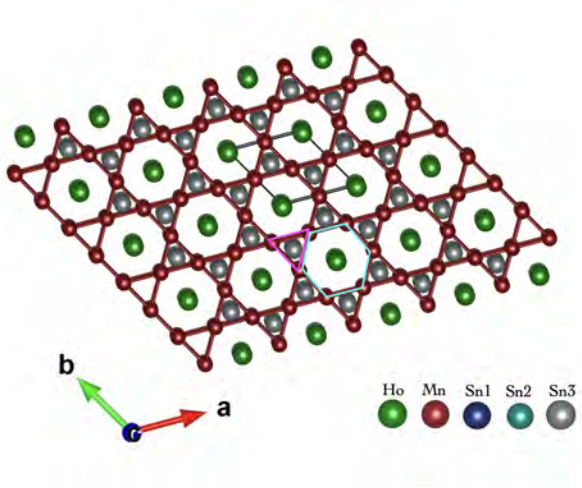

Figure 1: Illustrates the kagome lattice of $HoMn_6Sn_6$. The hexagon is highlighted in cyan and one of the corner-sharing triangles is highlighted in magenta.

$TmMn_6Sn_6$ stabilizes magnetic skyrmions by modifying the magnetic anisotropy [6; 14]. In this system, the formation of a magnetic skyrmion lattice is limited to thin samples and requires the application of an external magnetic field. This observation prompts an important question: can skyrmions be stabilized in bulk samples under zero-field conditions? Realizing stable zero-field skyrmions in bulk would be a major step forward for next-generation nonvolatile, high-density spintronic memory technologies [6; 9; 10].

In this study, we analyze a related compound, $HoMn_6Sn_6$, with Ga substitution at the Sn site, focusing on its structural and magnetic properties [1; 8]. We grew single crystals of several Ga concentrations and confirmed their structural properties. The magnetization measurements reveal a significant change with Ga doping.

2 Experimental Methods

2.1 Single Crystal Synthesis

Single crystals of the parent compound $HoMn_6Sn_6$ and Ga-doped samples were grown using the self-flux method with Sn as flux. The raw elements Ho pieces, Mn pieces, and Sn shots were loaded onto a crucible in a 1:6:20 molar ratio and sealed in a fused silica ampule under vacuum. Figures 2

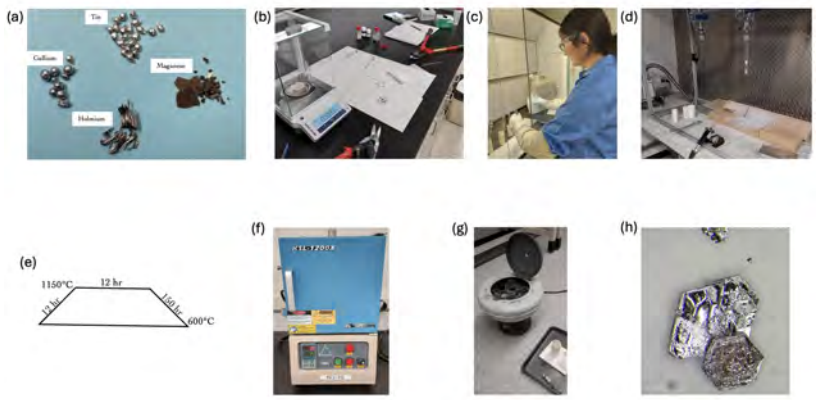

Figure 2: (a) Raw elements used for crystal synthesis. (b) The weighing process for each of the elements before adding them to the crucibles. (c) Process of adding quartz wool to the quartz tubes holding the crucible, serving the purpose of filtering out the flux during centrifuging. (d) Vacuum sealing process of the quartz tube. (e) The phase diagram is used for the temperature profile for the crystal growth. (f) High-temperature box furnace used for crystal growth. (g) Centrifuge to separate the extra flux. (h). Single crystal after centrifuging.

(a) and 2 (b) show the elements used and the weighing process of the grams needed for each element for self-flux growth. Once the elements were weighed and added to the crucible, the crucible was placed in a quartz tube that also contained quartz wool that would serve as a filter for the flux before sealing it (see Figs. 2 (c) and (d)).

Using a phase diagram (Fig. 2 (e)), we determined the times and temperatures needed for the compounds to grow in the box furnace (see Fig. 2 (f)). The tubes were heated to 1150°C in 12 h and remained at this temperature for 12 h. Subsequently, the furnace was cooled at a rate of 3.67°C/hour to 600°C. At this point, the tubes were centrifuged to separate the molten flux from the crystals. The resulting single crystals were then collected (see Fig. 2 (g) and (h)).

3 Structural Characterization

3.1 Powder XRD

The crystal structure of the grown single crystals was characterized using powder X-ray diffraction (XRD). Powder XRD works according to the principle of Bragg's law, which is expressed as

$$2d\,sin\theta = n\lambda$$

where d is the interplanar spacing between two adjacent planes, θ is the angle of incidence and λ is the wavelength of the X-ray. When a beam of monochromatic X-rays incident on the powder sample, they are reflected by the lattice planes, as shown in Fig. 2 (a). The reflected beams produce constructive interference once they satisfy Bragg's law. The diffraction patterns consist of multiple peaks at different positions of θ, and the value of d corresponding to the lattice plane can be determined from the known value of θ and λ.

A mortar and pestle were used to crush a polished single-crystal sample into a polycrystalline powder. This powder was placed inside the Rigaku MiniFlex diffractometer ($\lambda = 1.54$Å) to collect the XRD data. The collected XRD data was refined using the Rietveld refinement implemented in the *FullProf* software suite. The experimental data were compared with the simulation of a known structure acquired from a data database such as *Materials Project*. The aim is to have a minimum difference between the measured and calculated data through Rietveld refinement. This helps confirm the accurate structure as well as the atomic positions.

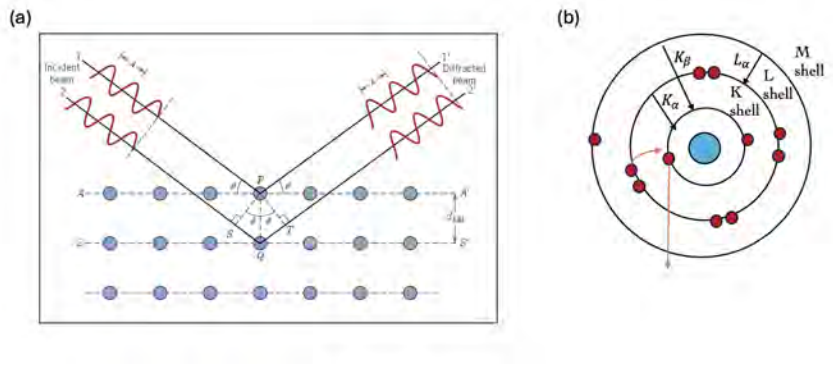

Figure 3: (a) Schematic of Bragg's law and how the incident beam diffracts from the crystal lattice. (b)Illustrates the EDS process, where an incident electron beam ejects an inner-shell electron, creating a vacant state. The higher energy cell atom jumps to the vacant lower energy state, emitting X-rays.

3.2 Energy Dispersive X-ray Spectroscopy

Energy Dispersive X-ray Spectroscopy (EDS) is used for the quantification of the composition of the elements present in a crystal. A high-energy electron beam was bombarded on the polished

face of the sample. Once the electrons interact with the sample surface, X-rays are produced. The X-rays are produced when the primary electrons knock out the inner shell electron, leaving a vacant position that is filled by another electron from the outer shell as depicted in Fig. 2 (b). This resulted in the emission of extra energy in the form of X-rays. These emitted X-rays give information about the composition of the elements present in the sample.

3.3 Magnetic Properties Measurement

Magnetic properties were measured using the Magnetic Properties Measurement System (MPMS). We have performed temperature-dependent and field-dependent magnetization measurements at several temperatures. The susceptibility is measured at varying temperatures but at a constant applied magnetic field of 0.1 T [15]. Magnetization is measured at varying applied fields (+7 T to -7 T) at a constant temperature [15].

4 Results

4.1 Structural Characterization

Powder XRD was used on all samples that were synthesized, including the parent compound and the doped compounds that had varying doping with x = 1, 2, 3, and 4 for $HoMn_6Sn_{6-x}Ga_x$. After completing the Rietveld refinement for all of them using *FullProf*, we were able to confirm that the crystals fit best with the space group P6/mmm. Figs. 4 (a) and (b) show the Rietveld refinement for $HoMn_6Sn_6$ and $HoMn_6Sn_2Ga_4$, respectively. From the refinement it is confirmed that Ga doping was incorporated into the third Sn site, as shown in Figs. 4 (c) and (d) (the pink portion of the gray Sn atom). In both structures, the Mn atoms form the kagome layer, as shown by the red atoms.

EDS revealed that the elemental ratio in the parent compound is approximately $HoMn_{6.59}Sn_{6.63}$, as shown in Figure 5 (a). Figs. 5 (b) and (c) show Ga peaks, confirming successful Ga doping. Figure 5 (b) presents the EDS results for $HoMn_6Sn_2Ga_4$, with an actual composition of $HoMn_{6.85}Sn_{5.08}Ga_{1.90}$, indicating that Ga has nearly fully substituted for Sn at the Wyckoff 2c site.

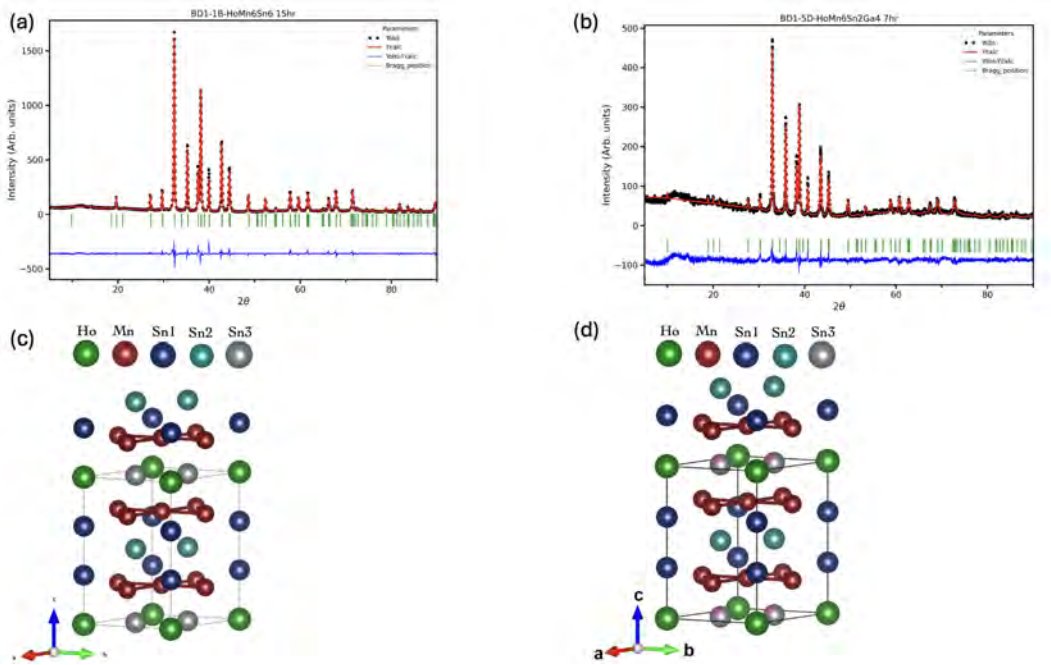

Figure 4: Structural characterization of $HoMn_6Sn_{6-x}Ga_x$. (a) Rietveld Refinement of powder XRD for parent compound. (b) Rietveld Refinement of powder XRD for doped compound $HoMn_6Sn_2Ga_4$. For (a) and (b) the black line represents the measured intensities from Powder XRD and the red line represents the actual intensities obtained from the simulations. The blue line depicts the difference between the measured intensities and the actual intensities obtained from the simulations. The green lines represent the expected Bragg peaks. (c) Schematic of crystal structure of parent compound $HoMn_6Sn_6$. (d) Schematic of crystal structure of doped compound $HoMn_6Sn_2Ga_4$.

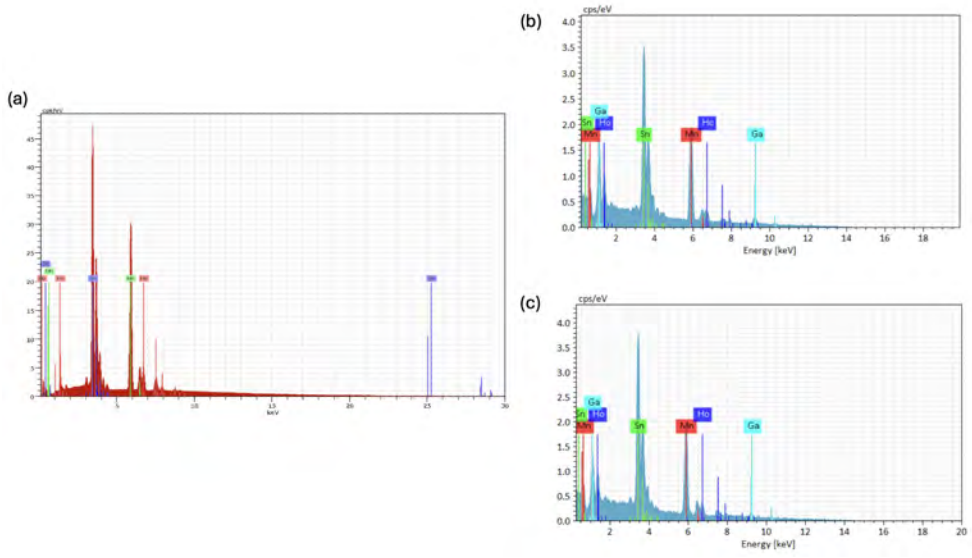

Figure 5: Graphs depicting the elemental composition of $HoMn_6Sn_6$. The peaks represent the K radiation detected after firing an electron beam on the sample. Each element has its own K_α and K_β allowing for EDS to recognize what elements are in the sample. (a) EDS spectra for the parent compound with no doping. (b) EDS spectra for $HoMn_6Sn_2Ga_4$. (c) EDS spectra for $HoMn_6Sn_3Ga_3$.

Figure 5 (c) shows the EDS data for $HoMn_6Sn_3Ga_3$, where the elemental ratio is $HoMn_{6.80}Sn_{5.39}Ga_{0.95}$, suggesting that Ga has replaced approximately half of the Sn atoms at the same site.

4.2 Magnetic Characterization

Figs. 6 (a) and (b) show the magnetization of the parent compound when the magnetic field is applied in-plane and out-of-plane. Figure 6 (e) shows the susceptibility for the $HoMn_6Sn_5Ga$. In the doped compound, the magnetic transition temperature increases by 20 degrees compared to the parent compound. The magnetic transition temperature of the parent compound is around 360 K as seen in Figs. 6 (a) and (b), while for the doped compound, the magnetic transition temperature is around 380 K in Figure 6 (e). The in-plane magnetic moment (red curve, Fig. 6 (e)) remains larger than the out-of-plane component over the entire temperature range. At low temperature, the in-plane magnetic moment decreases, and the out-of-plane magnetic moment increases. This indicates a gradual spin reorientation from the in-plane toward the out-of-plane direction as the temperature decreases. Although full spin reorientation is not achieved as observed in the parent compound [1; 8].

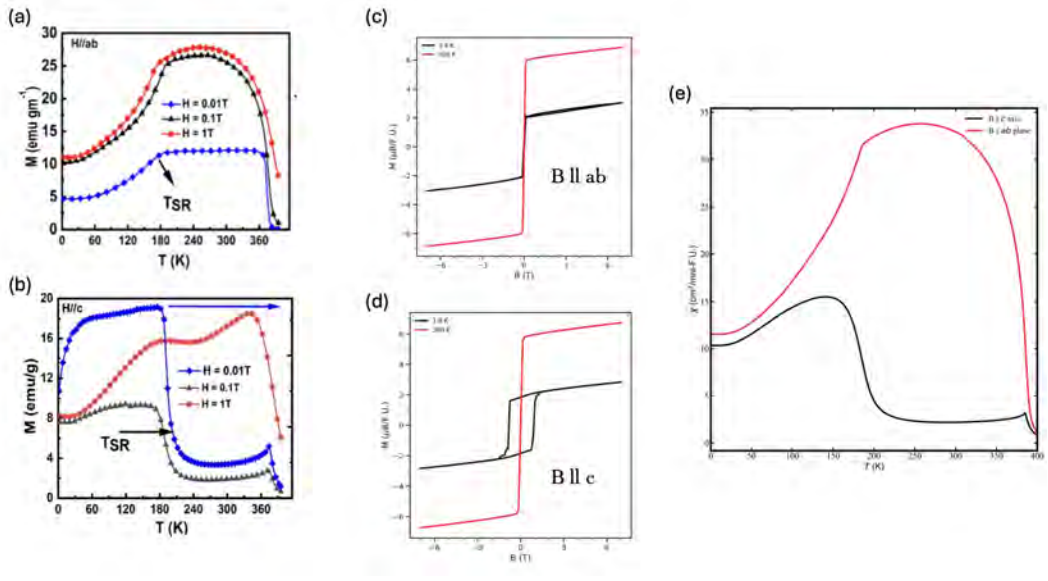

Figure 6: Magnetic properties of $HoMn_6Sn_{6-x}Ga_x$ with and without doping. (a) and (b) Temperature-dependent magnetic moment of $HoMn_6Sn_6$ with field applied along ab and c-axis, receptively [1]. (c)-(d) Magnetization vs. field at selected temperature with field applied along ab and c-axis, receptively. (e) Magnetic susceptibility as a function of temperature for field applied along ab (red curve) and c-axis (black curve).

Figs. 6 (c) and (d) show magnetization vs. field at selected temperature with field applied along the ab and c-axis, respectively. The low-temperature magnetization with field applied long the c-axis shows hysteresis, implying highly anisotropic magnetic properties.

5 Conclusion

We successfully grew single crystals of $HoMn_6Sn_{6-x}Ga_x$ with varying Ga doping levels and investigated the resulting changes in structural and magnetic properties. Powder X-ray diffraction confirmed that all samples crystallize in the hexagonal structure with space group P6/mmm (No. 191), with Sn substituted for Ga at the Wyckoff 2c site. Magnetic measurements revealed that Ga doping leads to an increase in the magnetic transition temperature, suggesting enhanced thermal stability of the magnetic phase. This improvement is promising for future technological applications, where thermal management remains a major challenge. Further studies will be performed to determine the changes in the magnetic and electronic properties for higher Ga-doped samples.

6 Acknowledgements

I would like to thank my advisor Dr. Nirmal Ghimire for his guidance and for giving me the opportunity to work in his lab this summer. I thank postdoc Dr. Sekh Jamaluddin and graduate students Resham Regmi, Abhijeet Nayak, and Mohamed El Gazzah for all their help over the course of this program. I also thank Dr. Umesh Garg and Kristen Amsler for watching over us and ensuring we had a great experience on the program.

References

[1] F. Kabir, R. Filippone, G. Dhakal, Y. Lee, N. Poudel, J. Casey, A. P. Sakhya, S. Regmi, R. Smith, P. Manfrinetti, L. Ke, K. Gofryk, M. Neupane, and A. K. Pathak, Phys. Rev. Mater. **6**, 064404 (2022).

[2] H. Bhandari, R. L. Dally, P. E. Siegfried, R. B. Regmi, K. C. Rule, S. Chi, J. W. Lynn, I. Mazin, and N. J. Ghimire, npj Quantum Materials **9**, 6 (2024).

[3] M. E. Gazzah, F. S. Yasin, S. Jamaluddin, H. Bhandari, R. B. Regmi, X. Yu, and N. J. Ghimire, arXiv preprint arXiv:2504.19045 (2025).

[4] D. C. Jones, S. Das, H. Bhandari, X. Liu, P. Siegfried, M. P. Ghimire, S. S. Tsirkin, I. I. Mazin, and N. J. Ghimire, Phys. Rev. B **110**, 115134 (2024).

[5] S. Riberolles, T. J. Slade, R. Dally, P. Sarte, B. Li, T. Han, H. Lane, C. Stock, H. Bhandari, N. Ghimire, *et al.*, Nature Communications **14**, 2658 (2023).

[6] M. E. Gazzah, P.-H. Chang, Y. Lee, H. Bhandari, R. Regmi, X. Zhou, J. F. Mitchell, L. Ke, I. I. Mazin, and N. J. Ghimire, Doping-induced spin reorientation in kagome magnet tmmn6sn6 (2025), arXiv:2505.02936 [cond-mat.mtrl-sci] .

[7] Y. Lee, R. Skomski, X. Wang, P. P. Orth, Y. Ren, B. Kang, A. K. Pathak, A. Kutepov, B. N. Harmon, R. J. McQueeney, I. I. Mazin, and L. Ke, Phys. Rev. B **108**, 045132 (2023).

[8] H. Zeng, G. Yu, X. Luo, C. Chen, C. Fang, S. Ma, Z. Mo, J. Shen, M. Yuan, and Z. Zhong, Journal of Alloys and Compounds **899**, 163356 (2022).

[9] Y. Zhou, National Science Review **6**, 210 (2019).

[10] A. P. Petrović, C. Psaroudaki, P. Fischer, M. Garst, and C. Panagopoulos, Reviews of Modern Physics **97**, 10.1103/revmodphys.97.031001 (2025).

[11] P. E. Siegfried, H. Bhandari, D. C. Jones, M. P. Ghimire, R. L. Dally, L. Poudel, M. Bleuel, J. W. Lynn, I. I. Mazin, and N. J. Ghimire, Communications Physics **5**, 58 (2022).

[12] T. Asaba, S. M. Thomas, M. Curtis, J. D. Thompson, E. D. Bauer, and F. Ronning, Phys. Rev. B **101**, 174415 (2020).

[13] H. Bhandari, Z. Ning, P.-H. Chang, P. E. Siegfried, R. B. Regmi, M. E. Gazzah, A. V. Davydov, A. G. Oliver, L. Ke, I. I. Mazin, and N. J. Ghimire, Three-dimensional nature of anomalous hall conductivity in ymn6sn6-xgax, x 0.55 (2024), arXiv:2411.12134 [cond-mat.mtrl-sci] .

[14] F. Canepa, M. Napoletano, C. Lefèvre, and G. Venturini, Journal of Magnetism and Magnetic Materials **285**, 254 (2005).

[15] S. Mugiraneza and A. M. Hallas, Communications Physics **5**, 95 (2022).

Evaluating Reconstructions of Simulated Events in the Deep Underground Neutrino Experiment's Liquid Argon Time Projection Chamber

John Duffy

2025 NSF/REU Program

Physics Department, University of Notre Dame

Advisor: Professor John LoSecco

Abstract

The neutrino is a particle of substantial interest and a potential pathway to new physics, which has spurred the development of the Deep Underground Neutrino Experiment. In advance of the installation and operation of DUNE's LArTPC detector, we must develop methods to reconstruct the events that produce data in the detector. To that end, various reconstruction fitters have been developed and are tested on data produced from simulated events against those events. This report analyzes the efficacy of these fitters in accurately reconstructing the simulated events by using the ROOT library to analyze several thousand simulated and reconstructed events and applying various comparative analyses. This report finds that the "Pandora Track" fitter is the most accurate in reconstructing the direction of motion of interacting particles relative to the detector. Additional methods of extracting comparative data from both the simulated and reconstructed events remains necessary in order to fully evaluate which of the reconstruction fitters is best equipped to accurately reconstruct particle events from real data.

Introduction

The existence of the neutrino was first hypothesized in 1930 by Wolfgang Pauli to resolve an energy conservation problem in beta decay, and subsequently given the name "neutrino" by Enrico Fermi in 1934 [1]. The neutrino is a neutrally charged lepton that interacts only with the weak nuclear force and gravity, making detection of the neutrino difficult due to its scarce interactions with matter. The neutrino has three flavors, those being the electron, muon, and tau neutrinos. These three flavor eigenstates are each a linear combination of the three neutrino mass eigenstates. The neutrino additionally has the ability to oscillate between flavors when traversing large distances, a phenomenon whose 1998 discovery by the Super Kamiokande collaboration

merited the 2015 Nobel Prize in Physics. Following 2015, work began on the Long-Baseline Neutrino Experiment, which serves as the engine for the Deep Underground Neutrino Experiment (DUNE). DUNE is a large collaboration aiming to study neutrinos with the goal of better understanding the physics of the neutrino, searching for new physics such as an explanation for matter-antimatter asymmetry, a theory for the Unification of Forces, and other pursuits [2]. The detector that DUNE will use is a Liquid Argon Time Projection Chamber (LArTPC), in which neutrino interactions with the detector medium will cause particle showers and other induced effects [3]. However, the detector will not be installed until 2027 and will not begin collecting data until 2029. Additionally, the neutrinos are not directly detected but have their properties inferred from the effects they induce in the detector medium. Therefore, it is necessary to develop a reconstruction method that accurately reconstructs the events that produced the data in the detector, so that these events can be analyzed in the search for new physics. To test these reconstruction methods, we must use simulated events to produce data, and then have the reconstruction tools reconstruct events from the simulated data. A reconstruction tool, or fitter, that reconstructs events with reasonable accuracy in comparison to the simulated events should have the potential to accurately reconstruct real events from the actual data that will be collected in 2029 onwards. Thus, the goal of this report is to ascertain which reconstruction codes, or fitters, most accurately reconstruct the simulated events in DUNE's detector based on simulated data.

Methods

When the detector becomes operational, it will initially only detect atmospheric neutrinos. To simulate this, a Monte Carlo simulation produces thousands of incident neutrinos, whose

interactions with the detector are modeled by Geant4. After the induced effects of the simulated neutrinos are rendered, the reconstruction codes will attempt to reconstruct the particle tracks and then the incident neutrinos themselves. This report compares four reconstruction codes named “Pandora Track”, “PM Track TC”, “PM Track”, and “PM Trajfit”. The data produced by the simulations and reconstructions are stored in files that utilize ROOT, a C++ computing library. ROOT was developed at CERN in 1994 and replaced FORTRAN in use starting in 2003. However, due to ROOT’s more narrow user base than more popular programming languages such as Python or C++ more generally, there are fewer available resources for learners to refer to when debugging their scripts. As a result of these challenges with operating the ROOT library, the comparative analyses of the reconstruction fitters was reduced from their initial scope. The fitters were compared primarily through two methods. First, the momenta of the electrons and muons that were reconstructed by each reconstruction fitter were placed into histograms to discern what energy regions were overrepresented or underrepresented by each fitter in comparison to the others. Second, the zenith angles of each reconstructed electron and muon relative to the detector was compared to the “true” electrons and muons via recording the cosine of the angle between their trajectories relative to the detector, henceforth the "comparative cosine”, and then placing those values in a histogram.

While the neutrino has 3 mass eigenstates and 3 flavor eigenstates, at the relevant energies tau neutrinos interact so little with the detector medium that we elected to approximate the physics by using a 2-component model in which all incoming neutrinos are either electron and muon neutrinos, with the probability of a neutrino of one flavor not oscillating to the other, also known as the survival probability, being governed by the equation

$$P_{survival} = 1 - sin^2(2\theta)sin^2(1.27\frac{\Delta m^2 L}{E}\frac{[eV^2][km]}{[GeV]}), \quad (1)$$

where θ is the neutrino mixing angle, Δm^2 is the square of the difference in mass between the mass eigenstates, L is the path length of the neutrino, and E is the energy of the neutrino. Atmospheric neutrinos scarcely interact with matter, so the path length of a neutrino can be inferred by its zenith angle relative to the detector and determining the distance from the detector to the atmosphere 15 km above the surface of the Earth, where atmospheric neutrinos form, along that direction [4]. The programs were further simplified by not introducing uncertainties in the values of physical constants or any sort of measurement error.

Results and Discussion

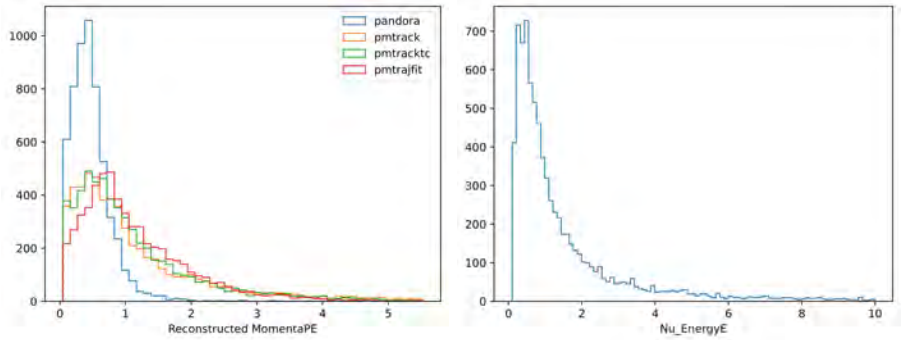

Figure 1 a-b: (a) Momenta of Reconstructed Electrons in GeV/c. (b) Energy of simulated incoming electron neutrinos in GeV.

The Pandora Track electrons are more centralized at lower energies than electrons in the other fitters, with a peak of about 1050 electrons with momenta about 0.4 GeV/c and electron momenta tapering out at about 2 GeV/c. The other fitters share the general structure of a plateau of about 400 to 500 electrons per energy region from about 0.5 to 1 GeV/c, with electron momenta tapering out around 5 GeV/c. PM Trajfit has a slight bias towards reconstructing electrons with larger momenta in comparison to PM Track and PM Track TC, which themselves

are quite similar. This means that, for the same data, Pandora Track is far more likely to reconstruct electrons with momenta below 1GeV/c, while the other fitters will reconstruct electrons with wider momenta distributions. Pandora Track also conforms more than the other fitters to the energy distribution of the incoming electron neutrinos considering the lower energy of its peak reconstruction momenta region.

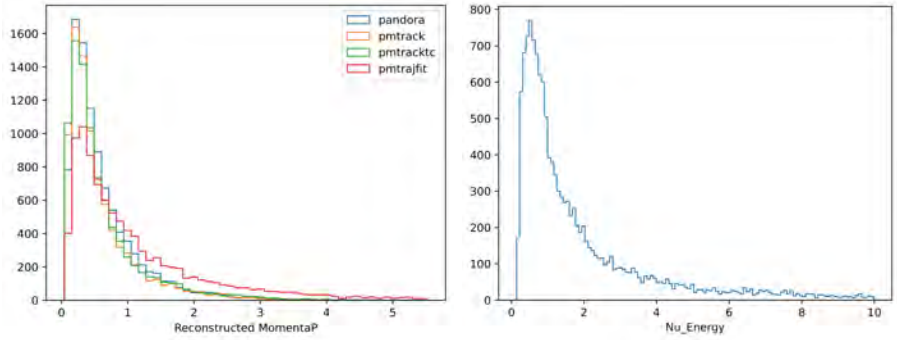

Figure 2 a-b: (a) Momenta of Reconstructed Muons in GeV/c. (b) Energy of simulated incoming muon neutrinos in GeV.

The Pandora Track, PM Track, and PM Track TC codes follow the same general structure of the majority of muons having momenta of less than 2 GeV, a peak reconstruction rate of about 1600 muons at about 0.3 GeV, and muon momenta generally tapering out at about 3 GeV. The PM Trajfit is a noticeable deviation, with a peak reconstruction rate of 1000 muons at about 0.4 GeV and muon momenta tapering out at about 5 GeV. This means that PM Trajfit, for the same data, will reconstruct muons with a wider momenta distribution than the other fitters. The three other fitters conform more closely to the shape of the energy distribution of the incoming muon neutrinos.

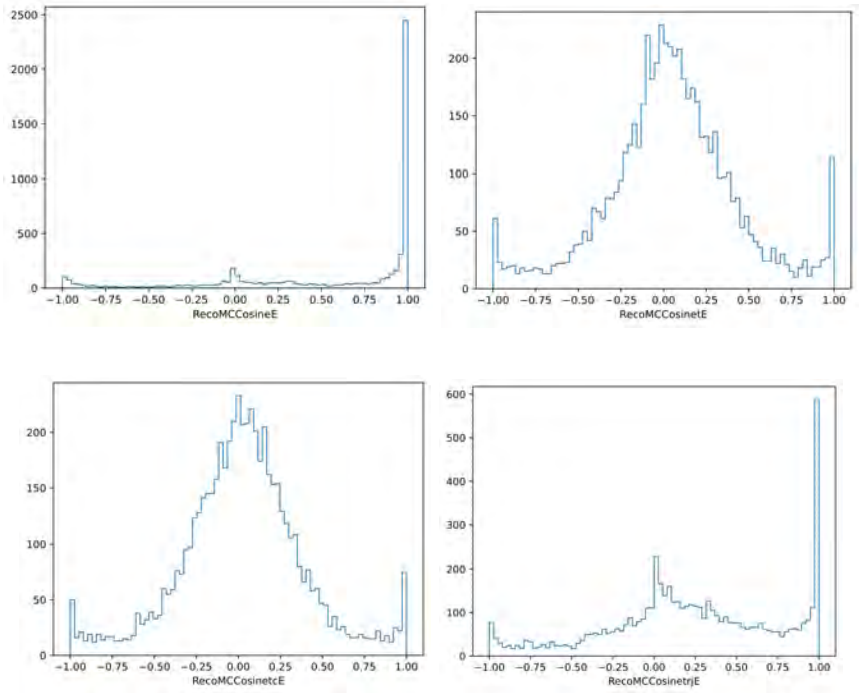

Figure 3 a-d: Comparative cosines of reconstructed electrons and Monte Carlo simulation electrons. (a) Top left, Pandora Track. (b) Top right, PM Track. (c) Bottom left, PM Track TC. (d) Bottom right, PM Trajfit.

Regarding the comparative cosine distribution, all fitters have 3 peaks in common across their plots. These are a peak at 1, which indicates that the simulated particle and reconstructed particle have the same trajectory, at -1, which indicates that the trajectory was reconstructed correctly except that the reconstructed particle moves in precisely the opposite direction as the simulated particle, and at 0, indicating that there are multiple reconstructed trajectories in opposite directions assigned to one simulation trajectory. About 2400 electrons reconstructed by Pandora have comparative cosines near 1, with much smaller peaks of 100 and 200 electrons about comparative cosines of -1 and 0, with relatively few electrons having other comparative cosine values. PM Track and PM Track TC are similar in having a peak of about 210 to 220 electrons about a comparative cosine of zero that gradually descends on both sides of the central peak, indicating that both fitters are widely erroneous in reconstructing electron trajectory, before

abruptly reaching smaller peaks at the cosines of -1 and 1. PM Trajfit's largest peak for comparative cosine values is about 600 electrons near a comparative cosine of 1, but it also has a peak of about 250 electrons near a comparative cosine of 0 that, similar to PM Track and PM Track TC, has a gradually descending slope that implies a widely erroneous reconstruction of electron trajectories. Pandora Track is evidently the fitter best suited to reconstruct electron trajectories on account of having the most reconstructed electrons with a comparative cosine near 1 and the fewest egregiously erroneous comparative cosines.

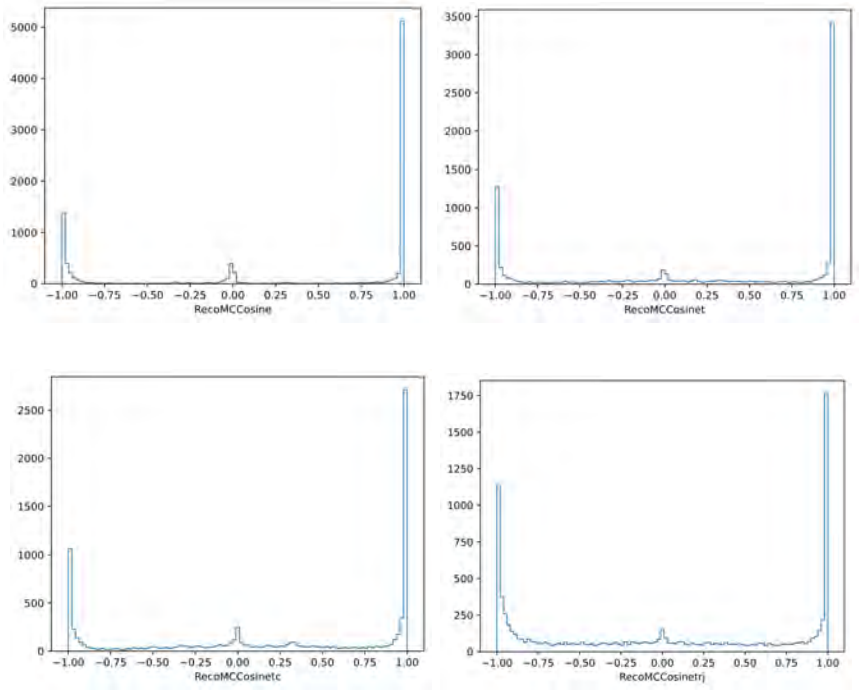

Figure 4 a-d: Comparative cosines of reconstructed muons and Monte Carlo simulation muons. (a) Top left, Pandora Track. (b) Top right, PM Track. (c) Bottom left, PM Track TC. (d) Bottom right, PM Trajfit.

Pandora Track has a peak of about 5000 muons near a comparative cosine of 1, followed by about 1400 muons near -1 and about 400 muons near a comparative cosine of 0, with relatively very few muons having other comparative cosine values. PM Track has a peak of about 3400 muons near a comparative cosine of 1, followed by about 1300 muons near -1 and about 200

muons near a comparative cosine of 0, with more muons compared to Pandora Track but few relative to the peaks of PM Track's plot having other comparative cosine values. PM Track TC has a peak of about 2700 muons near a comparative cosine of 1, followed by about 1100 muons near -1 and about 300 muons near a comparative cosine of 0, with notably more muons having comparative cosines not within the immediate vicinity of -1,0, or 1. PM Trajfit has a peak of about 1750 muons near a comparative cosine of 1, followed by about 1150 muons near -1 and about 180 muons near a comparative cosine of 0, and notably has at least 50 muons in every comparative cosine region, indicating a wide distribution of erroneous muon trajectories. Pandora Track is evidently the fitter best suited to reconstruct muon trajectories on account of having the most reconstructed electrons with a comparative cosine near 1. Notably, the comparative cosine plots do not account for a potential relation between a particle's energy and the resolution of its trajectory from its induced effects. A more energetic particle may induce less ambiguous effects, making it less difficult for fitters to accurately reconstruct its trajectory.

Conclusion

From the comparative analyses, Pandora is the best reconstruction fitter at correctly inferring the direction of motion of interacting particles. However, more methods of comparative analysis, in both breadth and depth, is necessary to fully compare the fitters. Furthermore, methods by which to extract comparative data from the simulated events for these new contexts must be developed in order to compare the fitters so that their accuracy can be thoroughly tested. Additionally, a potential relation between particle energy and trajectory reconstruction resolution should be investigated. Finally, measurement error and uncertainties in physical constants should be introduced and propagated. Even with the limited developments in this report, we are still

progressing towards reasonable accuracy in the reconstructed events from DUNE's data in 2029 and beyond, from which we will learn more about the neutrino and may potentially discover new physics.

References

1. "The History of Neutrinos". *Fermi National Accelerator Laboratory*. Retrieved July 31, 2025 from https://neutrinos.fnal.gov/history/.

2. "An International Experiment for Neutrino Science". *Deep Underground Neutrino Experiment*. Retrieved July 31, 2025 from https://www.dunescience.org/.

3. "DUNE at LNBF: Detectors and computing". *Fermi National Accelerator Laboratory*. Retrieved July 31, 2025 from https://lbnf-dune.fnal.gov/how-it-works/detectors-and-computing/.

4. "Atmospheric neutrinos". *Fermi National Accelerator Laboratory*. Retrieved July 31, 2025 from https://neutrinos.fnal.gov/sources/atmospheric-neutrinos/.

Acknowledgements

I wish to thank Professor LoSecco for mentoring me throughout this project. I also wish to thank the National Science Foundation for funding and the University of Notre Dame for organizing this research project.

Improvement and Characterization of HIgh Pressure POint-like Gas Jet Target (HIPPO)

JULIA DUNN

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Manoel Couder

Abstract

Accurate characterization of a target in experimental nuclear astrophysics is imperative because analysis depends heavily on the number of atoms in the target. Traditional setups studying alpha capture reactions, for example, use a solid, well-understood target and a light ion beam such as ^{4}He. However, background radiation often overshadows gamma rays that are the signature of the reaction. Inverting the setup, such that a heavy ion beam bombards a light helium gas target, allows for the detection of heavy recoil, instead of detecting the gamma rays alone. The St. George separator uses a HIgh Pressure POint-like gas jet target (HIPPO) to maintain a stable target of helium, so that inverse kinematics measurements, with a heavy beam, can be performed. The thickness of the gas target varies with the gas injection pressure and residual gas may contaminate the gas makeup. By developing a remote interface for a mass flow controller and a liquid nitrogen auto-fill system for a sorption pump, we were able to improve the injection pressure stability and the purity of the gas makeup. With the cleaned and stabilized helium target, we will next use elastic scattering measurements in inverse kinematics to characterize the thickness of the target. Thus, we will be able to make more precise measurements and reliable conclusions in future experiments with St. George.

1 Introduction

Background noise is a significant challenge in experimental nuclear astrophysics. Most experiments measuring the rates of alpha capture reactions, for example, study the gamma rays emitted, which are often difficult to accurately detect because of the minuscule rates of reaction compared to the abundance of environmental radiation.

Standard experimental setups consist of a beam of a “light” element, such as ^{4}He, bombarding a "heavy" element, such as ^{2}0Ne. A solution to the background issues is to, instead, shoot the "heavy" element at a stationary target of the "light" element such that the recoil particles travel forward and to be detected themselves. This is the goal of the St. George beamline of the 5U nuclear accelerator in the University of Notre Dame Nuclear Science Laboratory.

This unique setup is not without limitations. HIPPO (HIgh Pressure POint-like gas jet target), a subsystem within St. George, recirculates helium to maintain a target for beam interaction. Parts of this system are inaccessible or difficult to repair. Consequently, despite fundamental analysis depending on a precise understanding of the number of atoms in the target, uncertainty regard-

ing the properties of the jet remains. A sorption pump, kept at very low temperatures, has been demonstrated to remove contaminants from the helium gas. Thus, a system was constructed to automatically fill a dewar surrounding the sorption pump with liquid nitrogen, such that it remains cold over long periods of time. Additionally, by developing a system to remotely adjust a newly installed and untested mass flow controller, the gas target injection pressure was stabilized.

Next, with a pure and stable target, we will be able to use the improved HIPPO setup to complete an experiment studying elastic scattering at non-resonant energies to specifically describe the target's properties. This will allow for more precise measurements and reliable conclusions in future valuable experiments utilizing St. George.

2 Background

2.1 Nuclear Astrophysics

Stars fuse light elements into heavier elements through nuclear fusion reactions. This is the origin of most of the elements we have on earth, notably the oxygen we breathe and the carbon we eat. The goal of experimental nuclear astrophysics is to study the nuclear reactions that play important roles in fueling stars. Alpha capture reactions play a key role during the helium-burning phase of star evolution [1]. An approach to studying these reactions is inverse kinematics, such that a beam of a "heavy" element bombards a helium target. With this method, we are able to detect not only the gamma rays produced but also the recoil particles themselves. With more massive particles as the projectile, conservation of momentum ensures that no recoil particles "bounce" backwards, and thus all of the recoil move forwards towards detectors. This detection is helpful in calculating the cross section from the reaction yield.

2.2 St. George Beamline

The St. George Recoil Separator is designed to study low-energy helium-capture reactions in inverse kinematics. For this experimental setup, an ion beam from the 5U Santa Ana accelerator

bombards a helium gas target called HIPPO. Following the collision between the beam and the target, the products of the reaction, along with unreacted beam and target recoil, continue forwards through the beamline. Large quadrupole and dipole magnets direct a specific charge state of the products of reaction to the end of the beamline by selecting for a specific mass-to-charge ratio. Anything outside this ratio, such as beam particles or products with different charge states, is not directed towards the detectors at the end of St. George [1].

2.3 HIPPO

HIPPO maintains a helium target in the St. George beamline using a nozzle and a series of pumps and collimators for differential pumping. The nozzle ensures a target with a thickness of the order of millimeters, and the pumps ensure a vacuum in the surrounding beamline and chambers [2]. Before entering the beamline, the helium runs through a zeolite sorption pump, which has been demonstrated to improve the purity of the overall gas makeup. The percentage concentrations of both helium and nitrogen (a good proxy for the presence of air in the system) over time were tested with and without the helium running through the sorption pump. As shown in Figures 1 and 2, using the sorption pump minimized the rate of contamination.[1] The helium is then injected through a Brooks Instruments model 5851E mass flow controller which stabilizes the gas flow and, at constant temperature, the injection pressure.

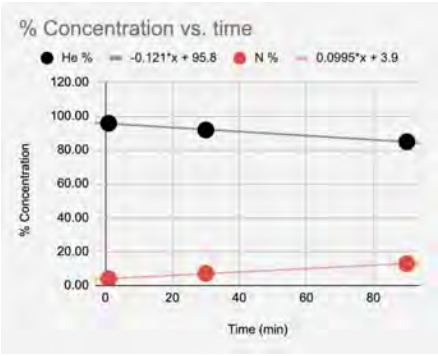

Figure 1: Without Sorption Pump

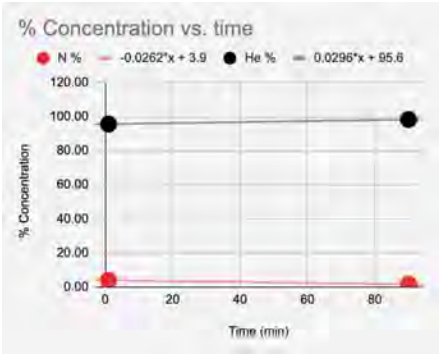

Figure 2: With Sorption Pump

[1] Data collection and visualizations done by ND NSL Graduate Student Adam Sanchez

3 Methods

Three main adjustments to HIPPO were made to improve the reliability of the target. Firstly, with the original setup, it was challenging to access central parts of HIPPO for maintenance and temporary experimental setup. Thus, the copper tubing and the helium injection panel were rearranged to improve safety and ease of access. A power source for the mass flow controller was modified so that it could be adjusted via the control computer, rather than manually in the target room. Finally, an automatic liquid nitrogen (LN2) filling system was developed and installed so that the sorption pump, which purifies the composition of the helium gas, is consistently operational.

3.1 Mass Flow Controller Power Source

A model 5851E Brooks Instrument Mass Flow Controller with a D-Connector was newly installed but untested in HIPPO. An existing power supply box, with both manual dial control and an isolated BNC input for remote control, was modified to be compatible with the device. The power supply for manual control was adjusted with a resistor so that the full range of the dial supplied 0-2V. A cable was created that supplied the device with the necessary inputs to the D-pin connector, as indicated in Table 1. An isolated BNC cable connects the remote BNC input to the analog output of an existing compactRIO along the St. George beamline. A labview VI was developed such that a selected value, between 0V and 5V, could be sent to the analog output as the command/drive, in place of the signal from the manual dial.

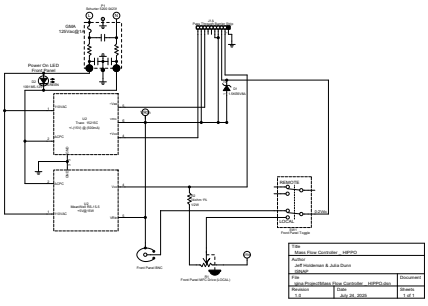

Figure 3: Mass Flow Controller Electrical Schematic

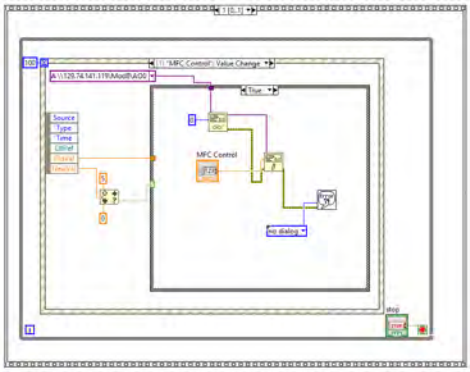

Figure 4: Mass Flow Controller LabVIEW VI Schematic

Color of Wire	Barrier Strip	Device (D-pin)	Function
White/Gray	1	14	Earth Ground
Orange	2	11	+5V Reference Output
White	3	8	Command/Drive
Dark Green	4 (jumped with 5)	1	+5V Ref Common / ±15V Common
Black	5 (jumped with 4)	9	+5V Ref Common / ±15V Common
N/A	6	N/A	N/A
Dark Blue	7	2	0–5V Signal Common
Light Green	8	6	–15V
Purple	9	10	0–5V Signal Output
Red	10	5	+15V

Table 1: Wired Connections for Cable Adapter

3.1.1 Long Term Testing

Since the mass flow controller will be used to constrain the gas for long periods of time, testing was necessary to ensure consistent stability. To do so, HIPPO was put in jet operation mode. Helium was injected into the system at 100 mbar higher than the desired injection pressure. Then, the voltage to the mass flow controller was slowly decreased until the gas was sufficiently constrained.

Every ten minutes the pressure before the mass flow controller, the injection pressure, and a rough estimate of the injection pressure were noted. Additionally, the ion gauge pressures, pressures in the chambers, and the pressure before the XDS35 pump were recorded. [2] When constraining the mass flow, there is a pressure difference between the pressure before the mass flow controller and the injection pressure (i.e. a gap), as gas builds up before the opening.

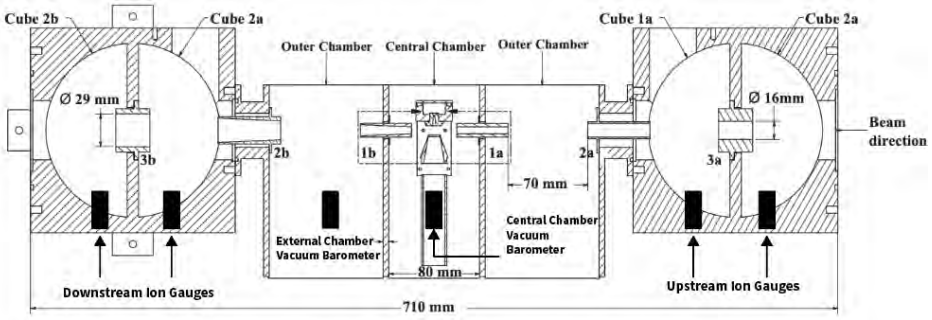

Figure 5: Locations of Ion Gauges and Barometers

[2]Figure of ion gauges and barometers is a modified figure from [2] which demonstrates the half section of the main gas target chambers

First, the target was in jet operation mode for four hours, with a desired injection pressure of 1000 mbar. During this run, the constraint of the mass flow controller was adjusted every hour to determine how different amounts of gas build up (i.e. different sized gaps) impact the system. During the second run, the target was in jet operation mode for six hours at the desired injection pressure of 1000 mbar and a gap of roughly 450 mbar. Finally, the target was in jet operation mode for four hours at a lower desired injection pressure of 150 mbar and a gap of roughly 200 mbar.

3.2 Liquid Nitrogen Auto-fill System

To avoid hand filling a LN2 dewar (about every 20 minutes) throughout a week-long experiment, an automatic system was developed using two temperature-sensitive LEDs. Each LED has a voltage difference of 2V at room temperature, but when submerged in LN2 they have a voltage difference of 5V. Using this signal, a circuit was created such that when the level is below the bottom LED, a valve connected to a large LN2 dewar is opened, and when the level is above the top LED that valve is closed.

Leading from the control box to the LN valve is a cable with a 20V signal, return, and ground. The valve itself is connected to a large LN2 dewar via a cryogenic hose featuring a safety relief valve. Following the valve, copper tubing leads into the small open dewar surrounding the sorption pump. Inside of the small dewar is a strip of repurposed circuit board supporting each of the LEDs, encased in thick aluminum to protect from splashing. These LEDs are then connected to the main control box with kapton insulated wires designed to survive liquid nitrogen temperatures.

The control box itself contains a power switch, power indication LED, a switch between automatic and manual fill, and a valve indication LED. The electronic schematic of the circuit box is shown in Figure 6. To simplify the physical control box, the logic was developed in a labview VI, shown in Figure 7. The program reads the voltages across each LED through a compactRIO. When there is a voltage across the bottom LED less than 4.5V, it is no longer submerged in liquid nitrogen, and 20V is supplied, opening the valve. When the voltage across the top LED is greater than 4.5V, it is submerged in liquid nitrogen, and the valve then closes.

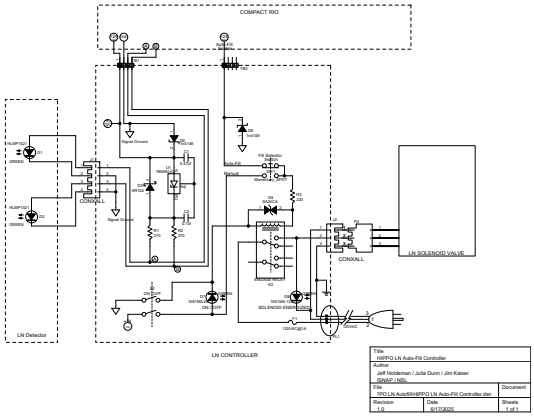

Figure 6: Autofill Electrical Schematic

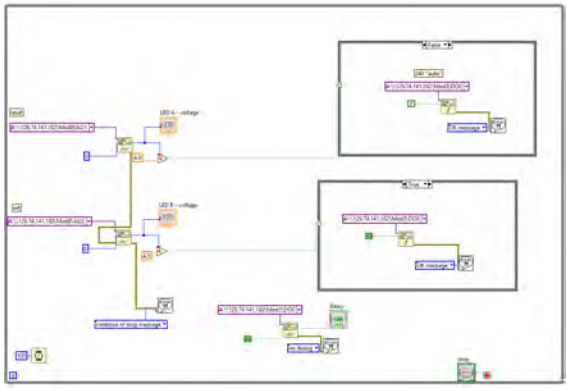

Figure 7: Autofill LabVIEW VI Schematic

4 Results

The rearranged setup of the helium injection panel within HIPPO has streamlined the overall design so that its components are easily accessible. In addition, it was shown that no substantial leaks were created when installing the equipment, since a decent vacuum was still achieved and maintained in the surrounding areas.

The mass flow controller can be adjusted remotely from the St. George labview control room. The injection pressure stability improved to roughly 2 mbar, compared to the previous variability of 7 mbar. When the injection pressure increases or decreases over time, it can be constrained back to the desired pressure without injecting or releasing gas.

Long-term tests demonstrated that the pressure buildup from constraining the gas (i.e. the gap) does not negatively impact the vacuum of areas surrounding the target. In fact, it was shown that a gap of 400 mbar minimizes the variation of the injection pressure. The surrounding vacuum in the ion gauges (two upstream and two downstream from the target), the chambers, and before the XDS35 pump all improved over time, regardless of the gap in pressure.

Finally, the liquid nitrogen automatic refill system for the HIPPO sorption pump was successful. It is able to initially fill the small dewar with the manual control switch. Then, when set to “AUTO,” the system automatically opens and closes the valve according to the logic.

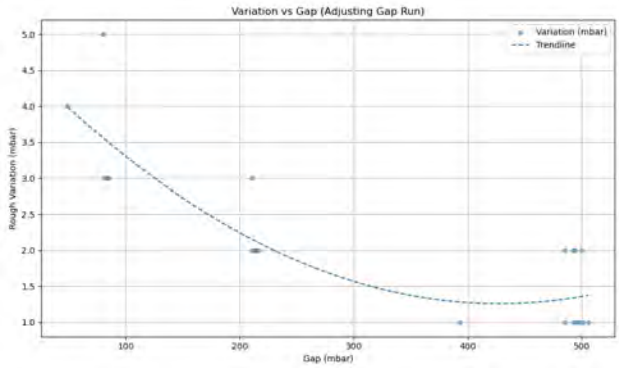

Figure 8: Variation vs. Gap (Adjusting Gap Run)

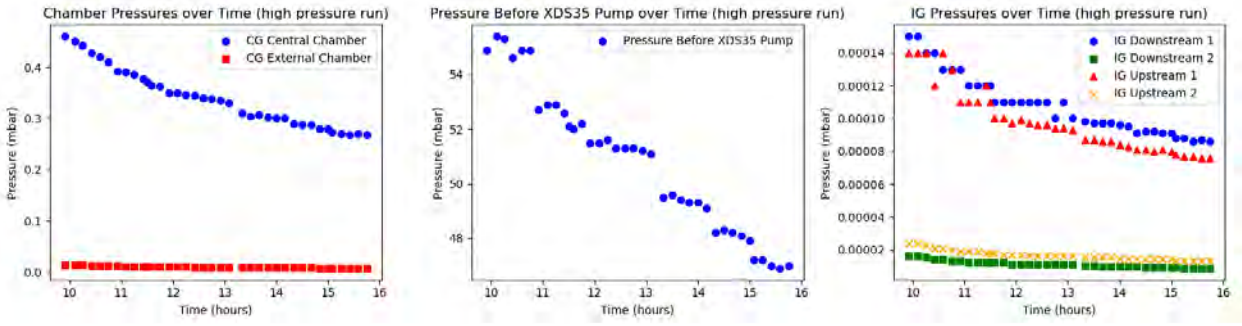

Figure 9: Surrounding Pressures over Time (High Pressure Run)

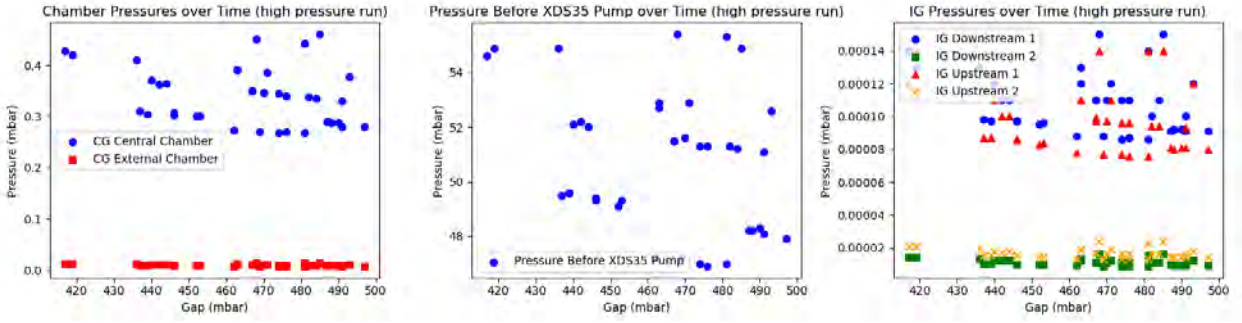

Figure 10: Surrounding Pressures vs Gap (High Pressure Run)

5 Conclusion

In conclusion, the mass flow controller power source and the automatic LN2 refill system improve the pressure and gas composition of our helium jet target. The automatic refill system keeps LN2 surrounding the sorption pump, so helium is consistently purified. The mass flow controller can be remotely adjusted to improve injection pressure, keeping the target consistent during experiments.

Moving forward, we will be able to use inverse kinematics to characterize the thickness of the improved target. By shooting a beam of non-resonant energy at the target, we will be able to understand and account for any variations from the properties we expect. This more explicit characterization provides valuable insight into our experimental setup such that more precise reaction rate calculations can be made for astrophysical reactions of more obvious significance.

6 Acknowledgments

I would like to thank my advisor Dr. Manoel Couder for his support and guidance throughout the last year and a half. I also thank Dr. Edward Stech, Jeffery Holdeman, Jerry Lingle, and the St. George group, especially Adam Sanchez, for their help on this project and beyond. Thank you Dr. Umesh Garg and Kristen Amsler for organizing the ND REU program. I also thank the ND COS for the Quazi and Shaheen Islam Summer Undergraduate Research Fellowship, supported by the Endowment for Excellence, the Glynn family, and the NSF for funding this research.

References

[1] M. Couder, G. Berg, J. Görres, P. LeBlanc, L. Lamm, E. Stech, M. Wiescher, and J. Hinnefeld, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **587**, 35 (2008).

[2] A. Kontos, D. Schürmann, C. Akers, M. Couder, J. Görres, D. Robertson, E. Stech, R. Talwar, and M. Wiescher, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **664**, 272 (2012).

CVD Growth and Characterization of a $MoTe_2$ and $PdTe_2$ Superconducting Bilayer

RUBEN V. FRANCHINI

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Badih Assaf

Abstract

The critical temperature and field of a superconductor depends on the material it is made from and the thickness of that material. A superconducting bilayer can exhibit unique superconducting properties due to the interface between the layers. The effects that this interface has on the critical field and temperature changes as layer thickness is varied. The bilayer of $PdTe_2$ and $MoTe_2$ has not been previously experimentally tested to determine critical temperature and field as a function of the thickness of the $PdTe_2$ layer. Upon varying the thickness of the palladium film, the critical temperature increases while the out of plane critical field decreases. The in plane critical field appears to be larger than that of $PdTe_2$ without the $MoTe_2$ interface.

1 Introduction

Superconductivity is a physical phenomenon that has applications in the transportation, communication, and medical fields. It is a property that allows materials to conduct electricity with zero resistance when cooled below a critical temperature (T_c). This temperature differs for every material and is affected by the thickness of the material. When in a superconducting state, materials also exhibit a behavior known as the Meissner Effect, where the material will expel magnetic fields from its interior up to the strength of the critical field (H_c). Superconductivity can be broken by heating the material past T_c or by applying a magnetic field stronger than H_c.

A superconducting bilayer can be formed by creating a heterostructure that is composed of two superconducting layers. Superconducting bilayers are of interest due to unique properties that arise from the superconducting interface[1]. We are interested in a bilayer composed of $PdTe_2$ grown on $MoTe_2$ through the process of chemical vapor deposition (CVD). The T_c and H_C of this bilayer as a function of thickness has not been previously experimentally determined. By varying the thickness of the $PdTe_2$ layer, the T_c and H_C of the material can be tuned to fit the need for specific applications. This material will be characterized through the use of Raman spectroscopy, X-ray diffraction (XRD), low-temperature resistance measurements, and magneto-resistance measurements.

2 Background

The phenomenon of superconductivity is caused by the formation of cooper pairs, electrons of opposite spin and momentum that have been bound together. These electron pairs are able to pass through a material without scattering, therefore, with zero resistance. The wave functions of these cooper pairs can penetrate out of the superconducting material leading to superconducting proximity effects. These proximity effects are of particular interest in van der Waals heterostructures, where other materials are in very close proximity with the superconductor[1].

The typical approach to constructing van der Waals heterostructures is to individually grow each layer of the material and mechanically stack them afterwards. This has the benefit of having more freedom as to what materials can be used and their ordering, but it comes with several downsides. The mechanical stacking process often results in imperfect van der Waals interfaces alongside poor environmental sensitivity and structural stability[2]. Instead of relying on mechanical stacking, a high to low temperature CVD growth method will be used to grow the superconducting bilayers. This method involves growing materials directly on top of one another without disturbing the lower layers. This method results in high crystal quality and good environmental stability that can be scaled up to wafer scale, with the only downside being that the ordering of materials is limited, since those with a higher growth temperatures will have to be on the bottom[2]. For this reason, $PdTe_2$ will have to be the top layer that is varied in thickness.

2.1 Growth and Characterization

CVD is a process that is performed in a tube furnace where the gas form of a certain element is introduced, in our case tellurium, and allowed to interact with a heated metal film in order to form the desired material. This process is typically performed in either a vacuum or in a chamber with an inert gas flowing through to carry the element to the metal film. Any unwanted gasses left inside of the chamber during growth can result in the creation of other materials that are not desired or even cause the material to be evaporated at lower temperatures than expected. Since there are several ways that the CVD growth process can go wrong, Raman spectroscopy and XRD will be used to

confirm the quality of material post-growth.

The first characterization technique that will be used is Raman spectroscopy which is based on the concept of inelastic light scattering. When laser light hits a molecule, most of it is scattered elastically through a process called Rayleigh scattering, in which the emitted light has the same energy as the absorbed light. However, a portion of the light is scattered inelastically, with energy different from that of the incident light. This occurs because the incident photon excites the molecule into a virtual energy state. The molecule then returns to a vibrational energy state that is different from its original one, releasing a photon of a different wavelength. Every molecule has its own unique vibrational states so by examining the wavelengths of the inelastically scattered photons the molecule can be identified. This makes Raman spectroscopy a great tool for identifying the molecules that are present in a material.

The other characterization technique is XRD, which is used to determine the atomic structure of a crystalline material. This is done by directing a beam of X-rays at a sample where the atoms in the crystal will scatter the rays. At certain incidence angles, the intensity of the scattered rays will be maximized because they constructively interfere with each other. The angles at which this occurs are determined by the spacing between the atoms of the crystal and the energy of the X-rays. By using a beam with known energy and varying it across several angles, the spacing between the atoms can be determined, and this information can be used to determine crystal structure, quality, and orientation.

3 Methods

The superconducting bilayer is grown on sapphire because of its smooth surface, high melting point, and insulating properties, which ensure it will not interfere with the superconductivity of the bilayer. The first part of the growth process is to bake the sapphire in a furnace at $1200°C$ for 1 hour to ensure that any oxygen trapped within the sapphire is removed prior to CVD. If this part of the process is skipped, the trapped oxygen will instead be released during CVD and will result in some of the material being evaporated.

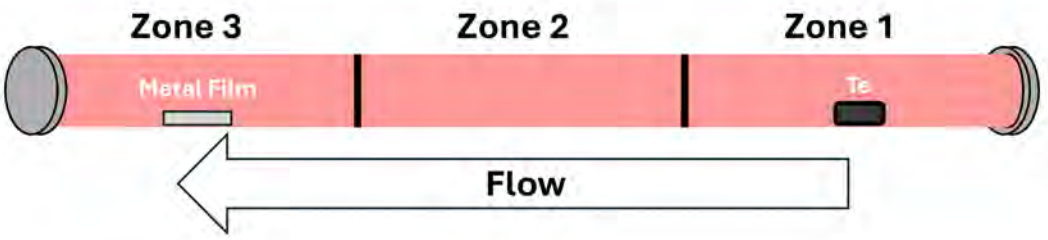

Figure 1: Schematic showing the tellurization process in the three zone CVD furnace.

After baking the sapphire, it is taken to a class 1000 cleanroom to deposit a 12nm thin film of molybdenum onto it through DC magnetron sputtering. The thin films were then cleaned in an ultra sonic bath using acetone and isopropyl alcohol before being dried off using compressed nitrogen gas. The cleaned thin films of molybdenum were then transported out of the cleanroom to the CVD furnace in order to grow $MoTe_2$. The CVD furnace is a horizontal three-zone tube furnace, 6 cm in diameter and 120 cm in length. In order to grow $MoTe_2$, H_2 gas has to be present in order to form H_2Te. Tellurium gas on its own is not very reactive with molybdenum but with the addition of hydrogen gas into the CVD furnace, H_2Te can be formed[3]. This H_2Te gas is much more reactive with the molybdenum film and allows $MoTe_2$ to grow with much better crystal quality. For the growths that were performed, a gas mixture of 10% H_2 and 90% N was pumped into the system at a rate of 0.32 L/hr. Before heating begins and a vacuum is pulled on the system, tellurium powder is placed in zone 1 and the molybdenum film is placed in zone 3 as far away from the tellurium as possible. The system is then sealed and a vacuum is pulled on it. This vacuum will reach around 30 μbar before the tube is purged with a high flow rate of argon gas. The purge and vacuum process is repeated 3 times, at which point the vacuum should reach around 4 μbar. The flow of N/H_2 gas can begin at this point and the heating process can begin.

Zone 1 is heated to 500°C at a rate of 16°C per minute while zone 3 is heated to 470°C at a rate of 12°C per minute. When growing $MoTe_2$, zone 2 functions as an intermediary between zone 1 and 3. It is also heated to 500°C but matching the rate of zone 1, 12°C per minute. These temperatures are then held for 30 minutes so that the growth can occur. The CVD furnace has

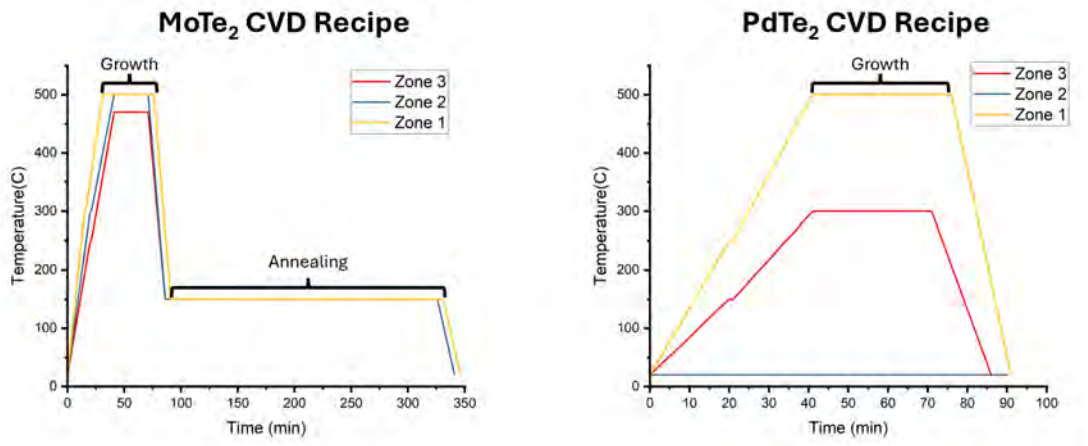

Figure 2: Graphs that show the recipes used in the CVD furnace for growing $MoTe_2$ and $PdTe_2$.

no ability to control cooling rate so the sample was passively cooled. The post-growth annealing process was performed in order to remove any large particles of tellurium that were left on the surface of the sample. It was held for 4 hours to account for the time it would take for the sample to cool to the annealing temperature.

Once the $MoTe_2$ samples were grown, it was tested for conductivity and Raman spectroscopy as well as X-ray diffraction were performed on it. The sample was then taken back to the cleanroom where an evaporator was used to deposit various thicknesses of palladium on the samples at a rate of $0.1 \mathring{A}/sec$: 3 nm, 5 nm, 7 nm, and 10 nm. The samples with palladium films were then returned to the CVD furnace where the purging process was repeated and the recipe for growing $PdTe_2$ was used. The biggest differences between this recipe and the previous one that was used for $MoTe_2$ is the exclusion of zone 2 and the lack of airflow. Not heating zone 2 allows it to function as a filter that prevents large particles of tellurium from reaching zone 1 and reacting with the sample. This eliminates the need for post-growth annealing since all of the large particles would have already been captured in zone 2. $PdTe_2$ is grown in vacuum as the pure tellurium powder is reactive enough with the palladium film that there is no need to introduce hydrogen and the temperature difference between zones 3 and 1 is sufficient to cause the tellurium gas to flow so there is no need for an inert gas.

Once the superconducting bilayer is completely grown, Raman spectroscopy and X-ray diffrac-

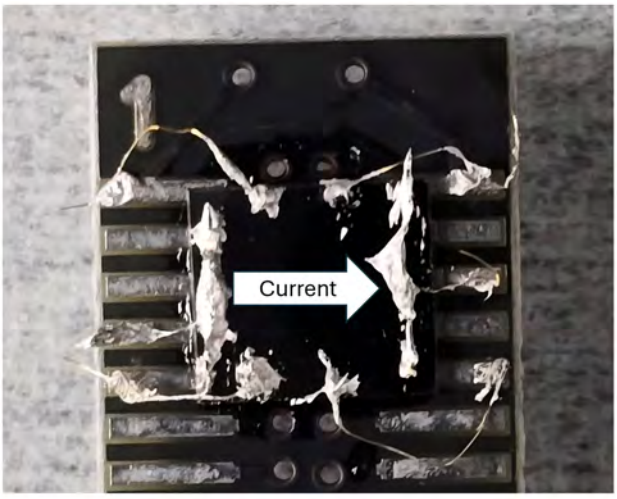

Figure 3: A superconducting bilayer sample soldered for low temperature resistance measurements.

tion are done again before the sample is prepared for low-temperature resistance measurements. The sample is connected to pins using gold wire in order to run current across it and measure voltage. The solder used is indium because it has better conductivity than other traditionally used solders. The sample is then cooled down to 1.4 K, the lowest temperature that can be reliably achieved on the set up, while measuring resistance to determine T_c. The sample is then held at temperature and a varying magnetic field is applied perpendicular to the surface of the sample to determine the value of the out of plane H_c.

4 Results

Four different thicknesses of palladium were deposited and then used to grow $PdTe_2$. The tellurization process increases the thickness of a layer by a factor of 5, for example, 10 nm of palladium will result in 50 nm of $PdTe_2$. Raman spectroscopy and XRD showed that as the thickness of the $PdTe_2$ layer was increased, the intensity of the peaks correlated to $PdTe_2$ increased while the intesnsity of the $MoTe_2$ peaks decreased. Generally, the peaks of the $MoTe_2$ were weaker than those of the $PdTe_2$ due to the lack of a pure hydrogen environment making $MoTe_2$ growth more difficult. In Raman, this issue is amplified as the beam is not able to penetrate deep enough as the $PdTe_2$ layer gets thicker to excite the atoms of the $MoTe_2$ drowning out the peaks even more. Despite this

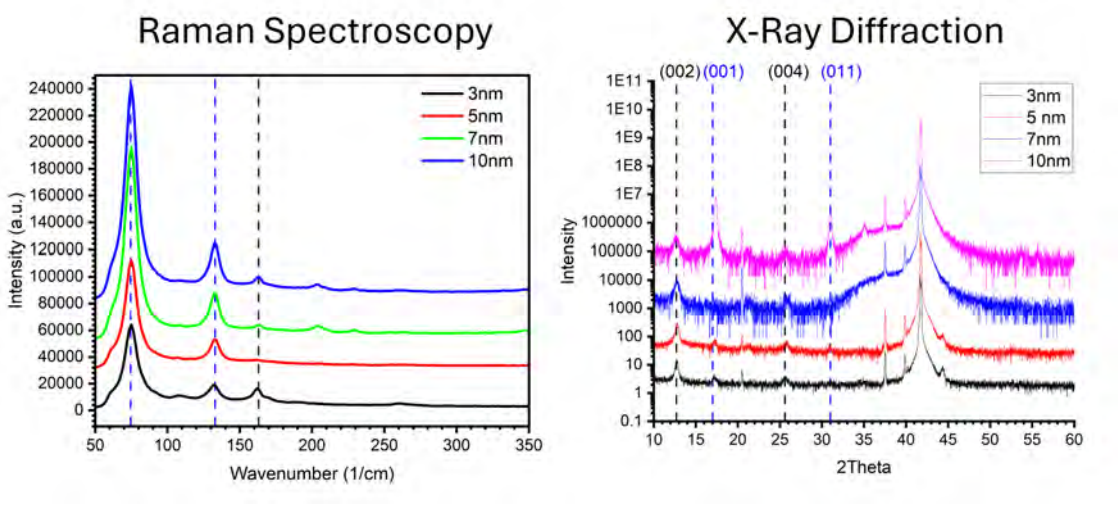

Figure 4: Raman spectroscopy and X-ray diffraction measurements for samples with varying thicknesses of palladium. Blue dashed lines correspond to peaks resultant from $PdTe_2$ and black dashed lines correspont to $MoTe_2$ peaks.

though, the peaks are still present confirming that the bilayer that was grown is $PdTe_2$ on $MoTe_2$.

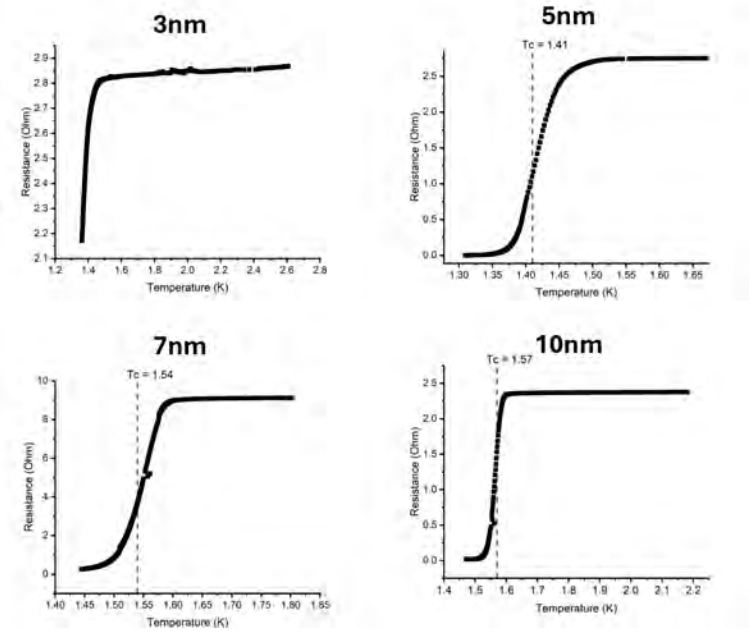

Figure 5: Temperature dependent resistance measurements at different thicknesses to determine critical temperature.

The set up that was used to measure resistance at low temperatures can only comfortably go to 1.4 K so it was not possible to see superconductivity in the 3nm sample, but the sample was

beginning to showcase superconducting behavior. All other thicknesses did reach a superconductive state and showed a trend where T_c would increase as the thickness of the $PdTe_2$ layer would increase.

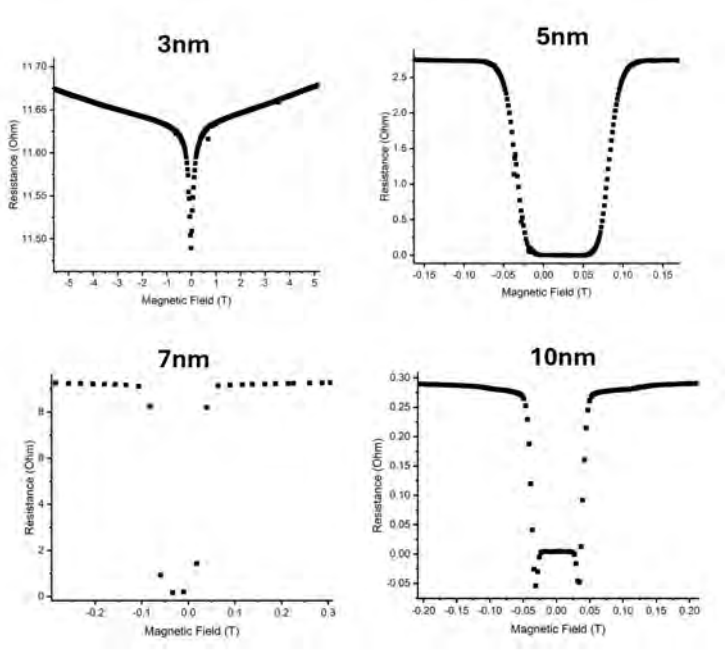

Figure 6: Critical field measurements done at $T/T_c = 0.93$ with the magnetic field perpendicular to the surface of the sample.

When H_c was measured, the temperature relative to T_c was kept constant and the strength of the applied magnetic field was increased until superconductivity was dissolved. Since the 3 nm sample was unable to be cooled below T_c, H_c was unable to be determined for it. From the other three samples, as the thickness of the $PdTe_2$ was increased, H_c would decrease.

Material	In Plane H_c (T)
40 nm $PdTe_2$	0.09
50 nm $PdTe_2$	0.05
50 nm $PdTe_2$ on $MoTe_2$	0.12
Bulk $PdTe_2$	0.04

A single in plane H_c measurement was done on the 10nm palladium sample which was shown to have a higher critical field than pure $PdTe_2$ of comparable thickness[4; 5].

5 Conclusion

Thickness (nm)	T_c (K)	H_c (T)
3	Unknown	Unknown
5	1.41	0.06
7	1.54	0.03
10	1.57	0.03

Using the high to low temperature CVD growth method, a bilayer of $PdTe_2$ on $MoTe_2$ was successfully grown. By increasing the thickness of the $PdTe_2$ layer, the T_c of the bilayer was increased and the out of plane H_c was decreased. This shows promise for being able to vary the thickness of the layers in a bilayer in order to tune the value of T_c and H_c of the bilayer for particular applications. The in plane H_c measurement shows promise for the bilayer being able to resist magnetic fields better than pure $PdTe_2$. In the future, more in plane critical field measurements will be made on the sample to examine the effect that the bilayer interface has on H_c.

References

[1] G. Brammertz, A. Golubov, P. Verhoeve, R. Hartog, A. Peacock, and H. Rogalla, Appl. Phys. Lett **80**, 10.1063 (2002).

[2] Z. Zhou, F. Hou, X. Huang, and et al., Nature **621**, 10.1038 (2023).

[3] M. Snure, M. Motala, T. Prusnick, E. Smith, D. Moore, C. Muratore, S. Vangala, and N. Glavin, J. Vac. Sci. Technol. A **40**, 042202 (2022).

[4] K. Yoshimura and et al., Paper in process of publishing (2025).

[5] H. Leng, C. Paulsen, Y. K. Huang, and A. de Visser, Phys. Rev. B **96**, 220506 (2017).

Optimizing CAD Model Conversion to GDML for Monte Carlo Simulations

ARAM GALSTYAN

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Ani Aprahamian, Prof. Wanpeng Tan

Abstract

Monte Carlo simulation-based codes such as Geant4 are widely used for particle transport and analysis. Their application extends beyond High Energy and Nuclear Physics to fields like Astronomy, Healthcare, and ground-based Radiation facilities. While the Engineering design of these environments is typically performed using CAD (Computer-Aided Design) tools, Geant4 does not support direct import of common CAD file formats. Instead, Geometry Description Markup Language (GDML) is required to pass CAD-based geometries in Geant4 successfully. However, reliable documentation and standardized workflows for GDML conversion remain limited. Several converters and commercial software were examined to test their effectiveness. Neutron background results from the NIS room were obtained by implementing all the beamline components, notably the closest lithium target holder, along with the massive dipole magnet and Wien filter. This is critical, as they are responsible for scattering neutrons produced at the lithium target location. Energy spectra were compared before and after the implementation of an additional iron box placed near the detector. The results will be presented.

1 Introduction

Monte Carlo simulations are a game-changer in the industry with broad applications in many environments. They can construct realistic simulations of any real-life frame. Monte Carlo simulations work based on a model of predicting the next step by extracting mean calculations from the previous results. One of the examples of such a program is Geant4, mainly used in Nuclear and High Energy Physics. Geant4 is an object-oriented software toolkit for the simulation of the passage of particles through matter. It tracks particles and models the processes through which they interact with the environment they traverse [1]. Such simulations are crucial for time efficiency, safety, as well as for validating results.

Geometry in Geant4 is defined in the DetectorConstruction header and source files. It first names all the solids that are going to be present in the simulation, and then goes into details about material, position, dimensions, etc. The main disadvantage of Geant4 is its limitation of geometry modeling. It supports a pre-defined list of solids, as well as union, subtraction, and intersection of those. However, this is often not enough to support describing complex solids that detectors may have. This is when it raises an extreme need to be able to make geometries in a Computer-Aided Design (CAD) software, which is the conventional environment for geometry architecture

and development.

The Geometry Description Markup Language (GDML) is a specialized XML-based language designed as an application-independent persistent format for describing detector geometries [2]. It serves to implement 'geometry trees' which correspond to the hierarchy of volumes, allowing identification of positions of individual solids, as well as to describe the materials. While Geant4 does not support the import of CAD files, conversion of them to GDML is possible with converter programs.

5U accelerator at NSL is a single-ended vertical pelletron, where the beam is sent through a 90-degree dipole magnet towards a room that includes a variety of different experimental stations. One such, the Neutron Irradiation Station (NIS), is a permanent beamline diverging by 45 degrees with a 5U main line after a bending magnet. It consists of several components, including the Faraday cup, turbo pump, a collimator, and the target. It uses the 7Li(p,n)7Be reaction for neutron production [3]. The neutron beam produced from the target could be directed to another material that would act as a study of interest for neutron-induced reactions, such as inelastic scattering or capture.

The detector is a high-purity germanium clover with four crystals, cooled by liquid nitrogen, and detects particles indirectly via gamma rays.

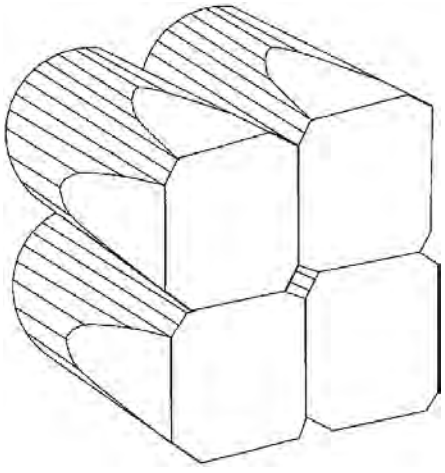

(a) Schematic view of a clover detector

(b) ORTEC - Germanium detector

Figure 1: HPGe detector

Neutrons are highly interactive with the environment, which makes the investigation sensitive

to any minor change in setup. For high-caliber results in beam neutron experiments, we want to make sure to exclude any neutron background present that may come from particles interacting with heavy equipment in the room. We apply the setup of the room and the beamline, and observe the changes before and after a heavy element placement near the detector. Geant4 simulations are ideal for studying and comparing different setups with varying beam sources and energies.

Figure 2: Neutron Irradiation Station (NIS)

2 Methods

2.1 Geometry design

To conduct a simulation on the neutron background in a room, we need to start with an appropriate geometry of the target room, beam line, and the detector. Particular importance was to include massive elements from the target line - the Wien velocity filter and the dipole magnet. In the actual room, heavy shelves are placed behind the detector, approximately 4 feet away. The detector, along with shielding, was placed another 4 feet away from the Lithium target. Our main interest is to compare energy spectra before and after putting a 1-foot iron box at the position of the shelves to mimic them.

Fusion was used for the construction and alignment of all of the components of the beam line. First, based on sketches and true-to-scale measurements, the copper-made target holder and block were aligned and adjusted to it's right dimensions of 3 inches in diameter, after which the rest of the

individual components were integrated. Radius centers, angles, and rotations, parallel and perpendicular planes, center of masses were used for comparison. The beamline had over 40 components in total, mostly consisting of high-vacuum tubes, 6-ways, and flanges. Some of the complex units as the gate valve, were sampled from official 3D sketches from the manufacturer.

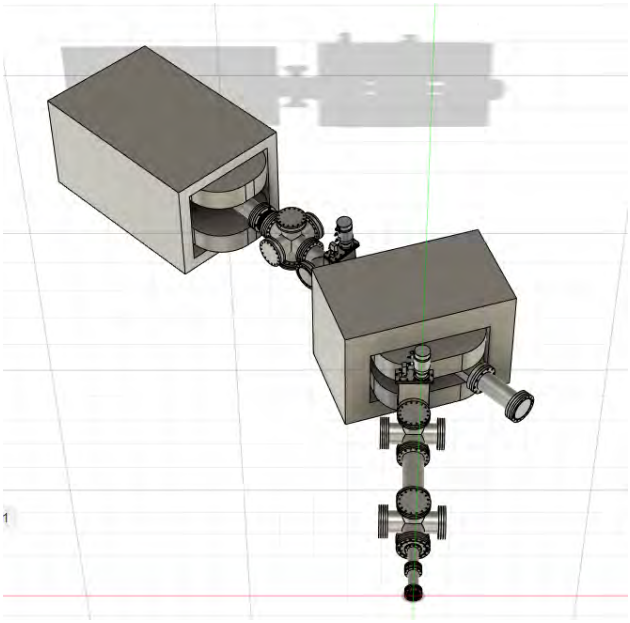

Figure 3: Design of the target line in Fusion

2.2 Conversion to GDML format

Once the framework is obtained, the next step is to convert it to a Geant4-readable file. For this, several converters were tested. Most of them accepted STL (only ASCII) or STEP format files. A separate Python code was developed for the STL binary to STL ASCII format change. Due to volume limitations, some of the constructions, consisting of the same material, were merged before relaying to the converter. After conversion, the geometry was restored with the right coordinates in Geant4.

2.3 Simulations

One of the Geant4 official examples is the G01, which is a code that allows to pass your GDML geometry and visualize it in a run window. This is to validate the accuracy of the GDML file before implementing it in an actual simulation. All of the components were authenticated through G01, as you can see in the picture shown below. For alignment of the target holder and the boxes with

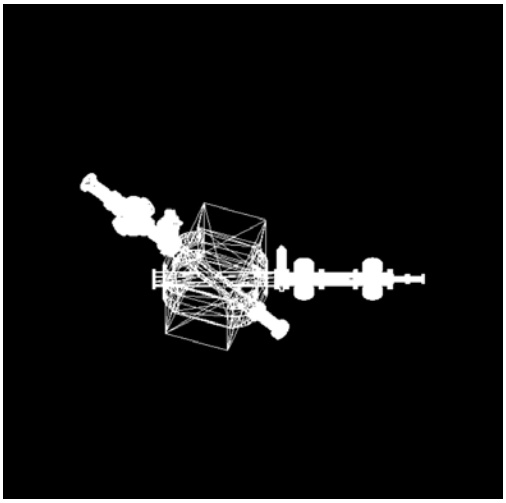

Figure 4: G01 run presenting the geometry

the rest of the beamline, as well as their relevant position in the room, some basic trigonometry was used. Translation vectors were calculated to shift the components with respect to each other. No electromagnetic effects were applied on the Wien velocity filter and dipole magnet, since our main interest was the mass effects on the background rather passage of the proton beam through the beamline.

The runs were concluded using a neutron beam source and mainly positioning the detector at 60° angles from the beamline. Then several angles and energies of interest were studied to investigate the background changes with specific parameters present.

3 Results

MRADSIM was the most efficient software, having user-friendly interfaces permitting the user to import STEP geometries with arbitrary size and complexity and convert them to GDML format [4].

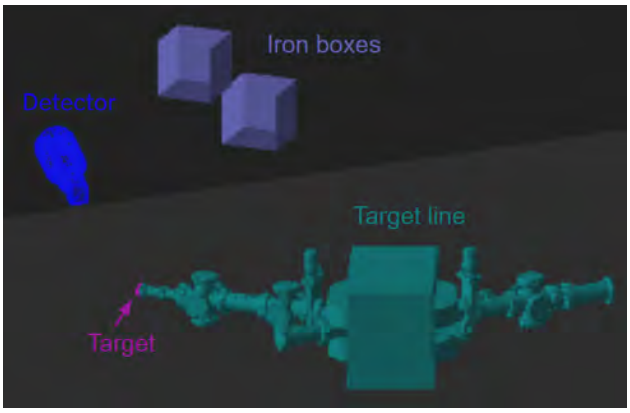

Figure 5: Setup components in Geant4

As well as it wasn't requiring any manual changes in the GDML code to avoid errors. FASTRAD, stl_gdml.py, pyg4ometry.py, all accepted smaller sizes of volumes as well as required material and position redifinitions after the conversion in the GDML file. No significant differences in runtime or memory usage were observed when using GDML files of the same geometry produced by different converters.

Another converter, ROOT2CAD, was tested, which was doing inverse conversion of GDML files, first making them a root file, then changing them to gITF, and lastly making it an obj file that can be easily imported to a CAD program. Although it does not keep the coordinate position of components, it allows us to actually get the dimensions of them, and allows us to make modifications to the initial geometry.

The simulation results were compared using the ROOT data analysis software. For each run, a separate histogram was plotted. The neutrons' entry number and energy distribution were surveyed. No distinguishable differences were registered.

After the full implementation of the setup, several runs were concluded to test the neutron background. The first run was with a source of three million neutrons, together with fixed 0.6 MeV energy, positioned slightly behind the target. The detector was placed at a 60 ° angle in relation to the beamline. In parallel, another run was introduced with the same framework, yet adding the iron box 4 feet from the detector. As for the results, there was little change in the energy spectra, and the box proved to have less impact than expected.

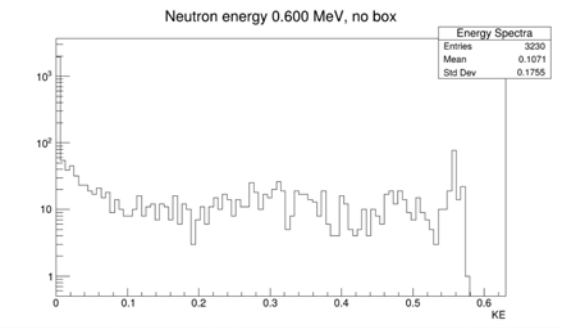

(a) Spectra without a box

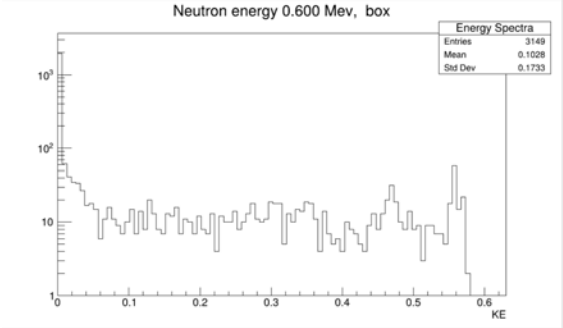

(b) Spectra with a box

Figure 6: Energy Spectra of neutrons at 60° degrees

Similar runs, with comparison before and after putting a box, were initiated for the detector placed respectively at 30° and 90° angles. The box was not moved with the detector. Only the latter showed a slight change in the number of neutrons entering the detector when a box is added. This suggests that examining at closer proximity reveals changes.

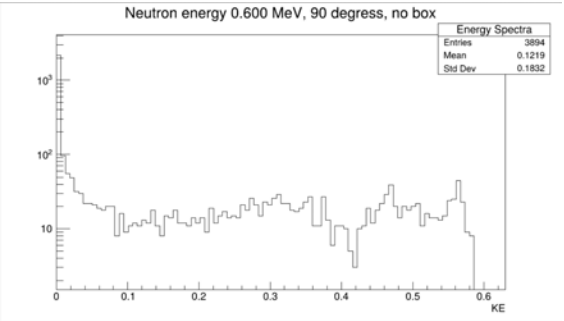

(a) Spectra without a box

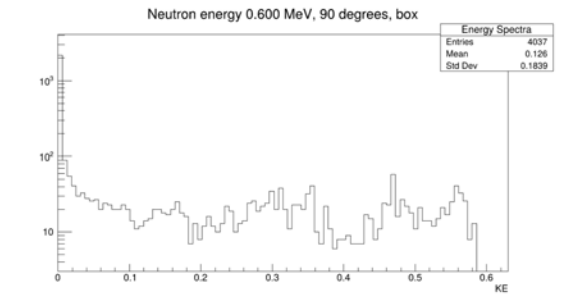

(b) Spectra with a box

Figure 7: Energy Spectra of neutrons at 90° degrees

Moreover, 2 runs with 0.4 and 1.5 MeV were resolved, demonstrating no evidence that the box caused any changes.

To extend the study, another box was located just one foot farther than the original, to see if additional equipment could create any disruption. However, even with 2 boxes near the detector, it again didn't yield to background changes.

Although the limited amount of data may not be sufficient to conclude on background changes, any significant effect would likely have already appeared. Given the small number of initial particles, minor variations can be attributed to statistical fluctuations. Therefore, the combined results

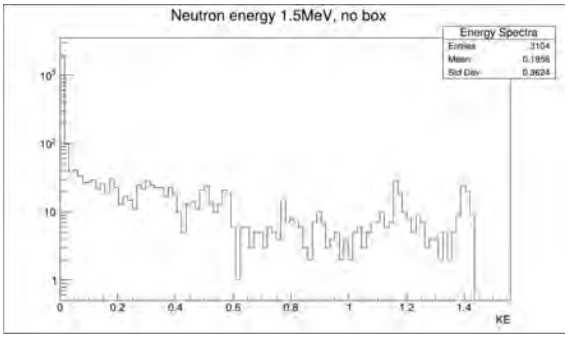

(a) Spectra without a box

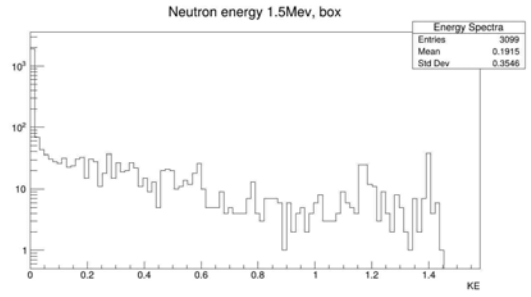

(b) Spectra with a box

Figure 8: Energy Spectra at 1.5 MeV

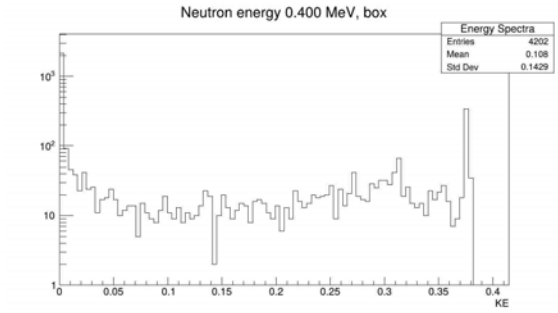

(a) Spectra with one box

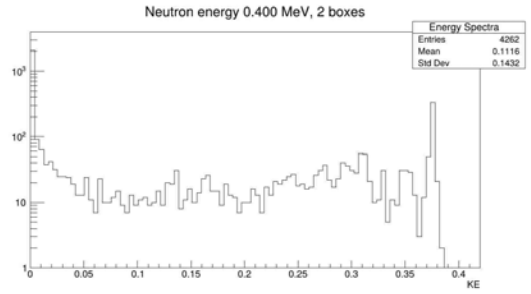

(b) Spectra with 2 boxes

Figure 9: Energy Spectra at 0.400 MeV

of all experiments indicate that no substantial background change occurs due to components near the detector of 1 foot in size.

4 Conclusion

GDML converters are highly useful for Monte Carlo simulations. They allow smooth import of geometry outside to Geant4. While some converters have volume or file size limitations, they still succeed in functioning in most of cases. Additionally, inverse converters help extract the geometries from the Geant4 GDML files.

The full setup was implemented, and the neutron background was obtained. The simulations showed that additional equipment in the room does not produce a significant effect on it. Changing detector angle, putting additional components, or changing energy does not produce any noticeable difference. Thus, during the experiments, they can be negligible.

For higher precision, further construction of the detector may be applied. As well as the adjustment of room size, with using ToF or LiDAR techniques. This will make the simulation vastly closer to the real experiment. This, with consideration of a broader range of parameters, would give the best results on the background study.

5 Acknowledgments

I want to thank my advisors, Prof. Ani Aprahamian and Prof. Wanpeng Tan, for their support and mentorship throughout the program. I also appreciate Kevin Lee for his assistance with any concerns. Special recognition goes to Ms. Kristen Amsler and Prof. Umesh Garg for making this program exceptional. Thank you to the Notre Dame Physics and Astronomy department faculty for actively championing the continuation and enrichment of the REU program. Many thanks to the fellow REU participants for making this summer unforgettable.

References

[1] T. Basaglia, Z. W. Bell, D. D'Agostino, P. V. Dressendorfer, S. Giani, M. G. Pia, and P. Saracco, Future Generation Computer Systems **159**, (2024).

[2] W. Pokorski, R. Chytracek, J. McCormick, and G. Santin, in *Proceedings of CHEP 2006 (Computing in High Energy and Nuclear Physics)* (Mumbai, India, 2006).

[3] M. Matney, D. Robertson, A. Simon, K. Manukyan, A. Aprahamian, D. Bardayan, R. J. deBoer, E. Stech, and M. C. F. Wiescher, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **1075**, 170406 (2025).

[4] A. B. Alpat, A. Coban, H. Kaya, and G. Bartolini, Computer Physics Communications **286**, 108688 (2023).

Revision of ^{27}Al(p,γ)^{28}Si level scheme and cascade branchings

ANNABEL GEIGER

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Anna Simon-Robertson

Abstract

Many nuclear capture reactions occurring in stellar environments are not well understood, and nuclear astrophysics research aims to characterize reaction cross sections to better understand stellar evolution and element synthesis. When these reactions are studied in the lab, efficiency calibrations must be performed to understand the detector's response to the signal of interest across the entire detected energy range. A well-known reaction producing γ rays up to the maximum energy of those produced by the reaction of experimental interest is used for efficiency calibrations. Many experiments use known resonances in the ^{27}Al(p,γ)^{28}Si reaction for this purpose. In the ^{27}Al(p,γ)^{28}Si reaction, level scheme and cascade branchings are incomplete for some resonances, in particular when high-energy γ rays are involved. This project created a new level scheme and cascade branchings for ^{27}Al(p,γ)^{28}Si based on available reaction data and compared simulated data from these new branchings with old branchings to determine which scheme more closely fits experimental data. Preliminary results from this work will be presented here. Future work will require adjusting the branching probabilities to further improve the simulated spectra.

1 Introduction

Stellar evolution and element synthesis is a large area of study within nuclear astrophysics, particularly heavy-element synthesis. One way in which the nuclear reactions responsible for heavy-element synthesis in stars are studied is by determining nuclear reaction cross sections - the probability of the reaction taking place - in the lab. This is commonly done by using an accelerator to generate a beam of particles, such as protons or alpha particles, which is sent to a target containing the reactant, resulting in nuclear reactions. The product is synthesized in an excited state and de-excites by emitting particles such as γ rays or neutrons. Detectors sense these particles and their energies, and data analysis is performed to determine nuclear reaction cross section, typically using γ-ray yields calculated in analysis.

1.1 γ-summing detector

The High EffiCiency TOtal absorption spectrometeR (HECTOR) is a 4π summing detector made of 16 NaI(Tl) crystals encased in aluminum. Each crystal has two photomultipliers (PMTs): one on the inside, closer to the target, and the other on the outside. Figure 1 shows a picture of the

outside of HECTOR. When a γ ray is generated, it scatters within one of the crystals, resulting in scintillation light emission, which is collected by the PMTs and converted to an electric signal for the acquisition system.[1]

A nucleus de-excites to ground state via a particular cascade by releasing multiple γ rays simultaneously. The energies of these coincident γ rays make up the particular cascade through which the nucleus de-excited, and their sum equals the energy level to which the nucleus was originally excited. When HECTOR detects coincident γ rays, their energies are added, resulting in a peak at the excited energy, also known as the sum peak.[1] Figure 2 shows an example sum peak for γ rays emitted from the radioactive ^{60}Co source and the corresponding level scheme, along with a sum of segments histogram, which plots the individual γ-ray energies detected instead of the sum energies in a cascade. The corresponding level scheme for ^{60}Co γ-decay is also shown.

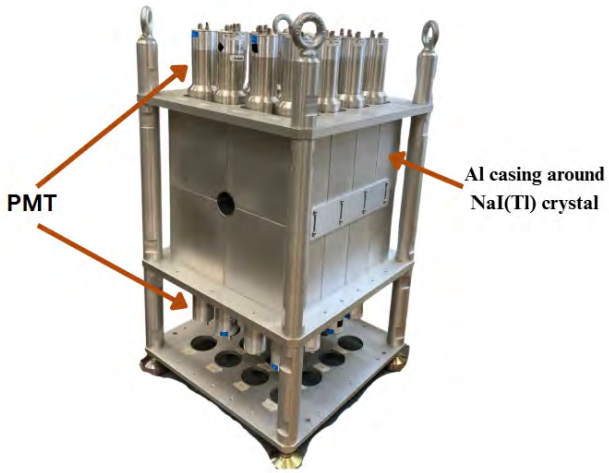

Figure 1: HECTOR as seen from the outside. Two PMTs protrude from each crystal segment. These segments are covered in aluminum shielding. At the intersection of the four crystal segments is the cylindrical hole through the detector to allow beam access to the target.

1.2 ^{27}Al(p,γ)^{28}Si for efficiency calibrations

The efficiency for a given energy is the fraction of γ rays detected at that energy, and when the cross section of a nuclear reaction is measured, it must be corrected by corresponding efficiencies at relevant energies. When performing an experiment with HECTOR or a similar instrument, an

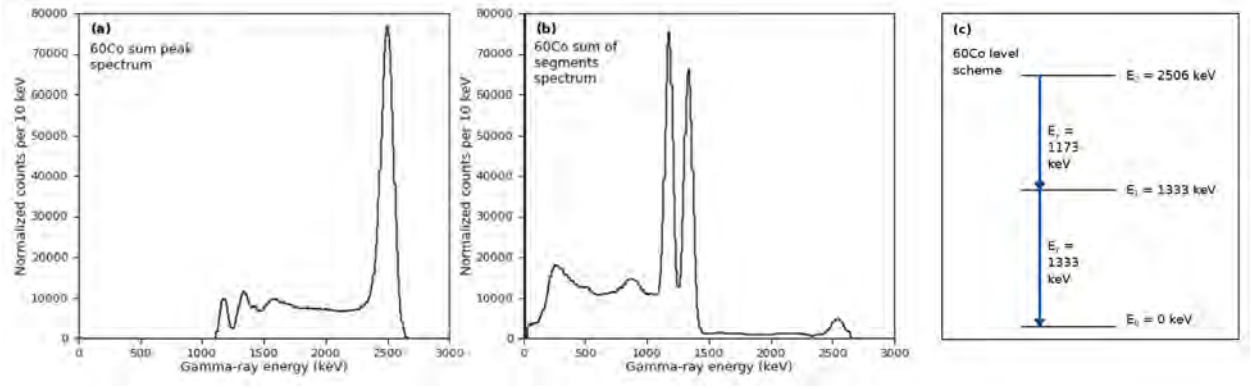

Figure 2: Example of sum peak (a) and sum of segments (b) histograms made from ^{60}Co γ-decay data. The far right peak in (a) is the sum peak, located at the energy of excitation in ^{60}Co. Two peaks can be seen in (b) at the γ-ray energies in the de-excitation cascade. The corresponding level scheme (c) shows the only possible decay to ground state, with γ rays emitted (arrows) and their energies. Levels are indicated by the lines.

efficiency calibration must be conducted so that collected data may be properly analyzed. The calibration should involve γ rays with energies up to the maximum energy the experiment will produce.

Previous experiments performed by the group required efficiency calibration for γ rays up to 12 MeV, and ^{27}Al(p,γ)^{28}Si was used for that purpose by taking data on certain reaction resonances. In nuclear physics, a resonance occurs when the projectile's energy (alpha particle, proton, etc.) matches the value of an energy state in the compound nucleus created after capture of the projectile. This causes a spike in the probability of capture. However, the energy level scheme and cascade branching probabilities for de-excitations to the ground state in ^{28}Si, the reaction product, turned out to not be well-known enough, especially for the high-energy γ rays involved in efficiency calibrations. This project aims to improve accuracy of the level scheme and cascade branching probabilities to improve efficiency calibrations using ^{27}Al(p,γ)^{28}Si which, if successful, would improve analysis of experimental data using these calibrations.

2 Methods

The following three ^{27}Al(p,g)^{28}Si resonances were studied, with energy value corresponding to lab beam energy: 2517 keV, 2711 keV, and 3674 keV. For each resonance, a simulation was created containing 200,000 events for comparison with experimental data.

2.1 Level scheme creation

^{28}Si de-excitation data from National Nuclear Data Center (NNDC), along with various publications [2–6], were used to create a file containing energy levels of ^{28}Si. For each of these levels, the file also contained possible energy levels to which the nucleus could de-excite via γ-ray emission and the corresponding probability of decay. This file, along with code from written by Sean Kelly [7], was used to create three input files, one for each resonance, with the possible cascades to ground state and their probabilities. These files are used with the Geant4 code that models the process of γ-ray interaction with the HECTOR array.

2.2 Simulations

Simulations were created using Geant4, which takes in input file of cascades from an excited state to a lower-energy state, their respective probabilities, along with the number of events Geant4 will run, and creates simulated events based on these probabilities and the construction of the detector. The simulation is compiled with a detector construction file containing materials in the detector, relevant physics for each material, and placement of these materials. When the simulation runs an event, a cascade from the input file is created, and each γ ray from the cascade is assigned a random direction. The simulation then tracks interactions between γ rays and detector materials. Once the energy has been deposited in a crystal, it is recorded in the segment corresponding to this crystal. These energy recordings are put into a root file as the output of the simulation.

After simulations were created, sum of segments and multiplicity histograms were created using simulation data. The sum of segments plots show individual γ rays detected by HECTOR for each event within the sum peak, and the multiplicity histograms show the number of γ rays in a cascade

detected by HECTOR. The sum of segments histograms are created by gating on the sum peak energy, eliminating any events for which coincident energies did not add to the sum peak energy. The remaining events are taken from individual γ-ray detection in each crystal segment. Event energies from individual segments, as opposed to the entire collection of events, are not summed since they are recorded as isolated events in each segment. These events are plotted together on one histogram, resulting in peaks at energies corresponding to individual γ rays involved in cascades.

Simulations and histograms were created using both the level scheme created for this project and a pre-existing scheme[8] to determine which was a more accurate representation of the experimental data.

A sum of segments histogram and multiplicity histogram were also created from experimental data taken in September 2024. For experimental data analysis, histograms were created for both on resonance runs - for which beam energy was equivalent to the resonance of interest - and off resonance runs - for which beam energy was far enough away from resonance energy that no resonance effects were observed. The off resonance histograms were subtracted from on resonance histograms, serving as background subtraction. The resulting histograms served as the reference to which the simulations should match. The counts in the simulation histograms were normalized to the counts in the experimental histograms.

3 Results

Table 1 shows the on and off resonance runs used for each resonance energy, along with the beam energy values recorded during the runs.

Table 1: Run numbers used for on resonance and off resonance data analysis, along with corresponding beam energies (keV) and literature resonance energies (keV).

On resonance runs	Resonance E (keV)	Beam E (keV)	Off resonance run	Beam E (keV)
628, 629	2517	2543.6	610	2538.6
594, 620, 621	2711	2736.6	597	2728.6
583, 584, 585	3674	3695	586	3693

Figure 3 shows the sum of segments histograms (a-c) and multiplicity histograms (d-f) generated from experimental data (black line), simulation data using the old cascade branchings (red dashes), and simulation data using the new cascade branchings (blue dots).

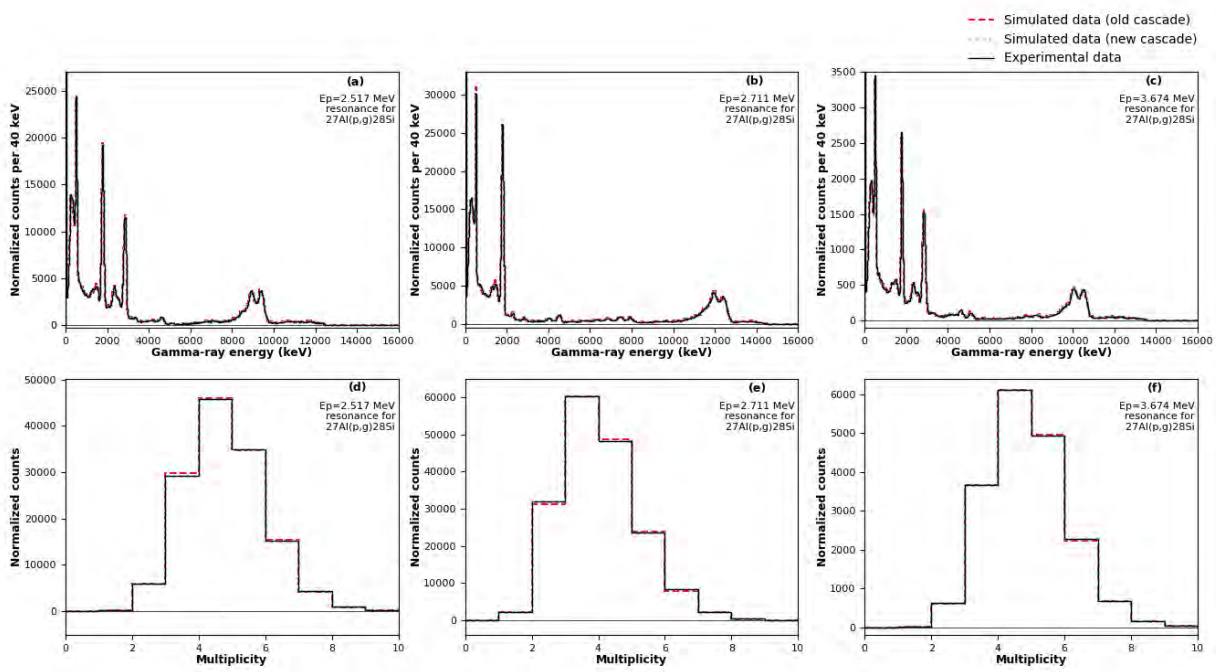

Figure 3: Sum of segments histograms are shown in a-c, and multiplicity histograms are shown in d-f. Experimental data is plotted as the black line, simulation data from the old cascade as the red dashed line, and simulation data from the new cascade as the blue dotted line.

Figure 4 shows residual plots for sum of segments and multiplicity histograms. These plots were made by subtracting a simulation histogram from the corresponding experimental histogram. The blue line represents the residual from subtracting new cascade simulation data, and the red dashed line represents the residual from subtracting old cascade simulation data.

Simulated data from the old cascade replicates experimental data better in some places, while the new cascade replicates the data better in others. Creating another cascade file combining the best parts of the two files used in this report may provide better replication of experimental data.

Both cascades over-predict low- and high-multiplicity events while under-predicting middle-value multiplicity events. These discrepancies point to directions for future work to improve the level scheme and cascade branchings for ^{27}Al(p,γ)^{28}Si. For example, branchings for the 2517 keV

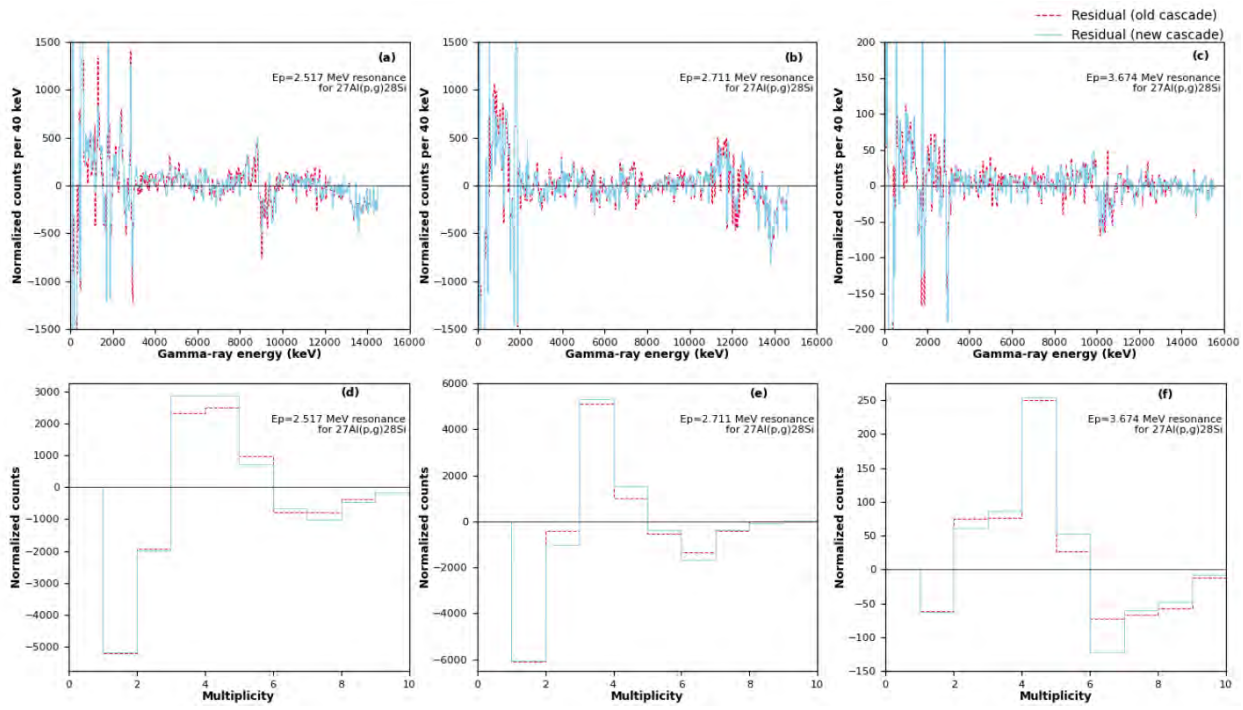

Figure 4: Residual sum of segments histograms are shown in a-c, and residual multiplicity histograms are shown in d-f. Simulation data was subtracted from experimental data to produce the plots, and the results of subtracting new cascade simulation data are shown as the blue line and of subtracting old cascade simulation data shown as the red dashed line.

resonance with multiplicities of 1, 2, 6, 7, 8, and 9 must have their probabilities lowered to correct the over-prediction of the current simulations. The residual plot for the 2517 keV resonance may point to exactly what energy γ rays are being over-predicted, indicating which specific branches must be adjusted. The same process is necessary for multiplicity 3, 4, and 5 branches, but these probabilities must be increased instead of decreased, and the residual sum of segments would be used to find under-predicted γ-ray energies.

4 Conclusion

Experimental data and simulation data of both old and new cascade branchings for ^{27}Al(p,γ)^{28}Si were compared using their sum of segments and multiplicity histograms. It was determined that the new branchings are a more accurate approximation of the true cascade branchings for ^{27}Al(p,γ)^{28}Si at certain γ-ray energies, but less accurate at others. This project will continue in the fall semester

to further improve accuracy of the cascade probabilities and level scheme by using experimental data to guide adjustments to branching probabilities. By continuing to improve the level scheme and cascade branchings for ^{27}Al (p,γ)^{28}Si, efficiency calibrations using this reaction will improve. The results of this work will be useful in the analysis of ^{89}Y(p,γ)^{90}Zr data completed with the Clovershare array at ND, as the detector efficiencies will be determined for energies up to 12 MeV as required by the experimental data.

References

[1] C. Reingold, O. Olivas-Gomez, A. Simon, J. Arroyo, M. Chamberlain, J. Wurzer, A. Spyrou, F. Naqvi, A. Dombos, A. Palmisano, T. Anderson, A. Clark, B. Frentz, M. Hall, S. Henderson, S. Moylan, D. Robertson, M. Skulski, E. Stech, S. Strauss, W. Tan, and B. Vande Kolk, The European Physics Journal A .

[2] J. Brenneisen, D. Grathwohl, M. Lickert, R. Ott, H. Röpke, J. Schmälzlin, P. Siedle, and B. Wildenthal, Zeitschrift für Physik A (1995).

[3] J. Brenneisen, D. Grathwohl, M. Lickert, R. Ott, H. Röpke, J. Schmälzlin, P. Siedle, and B. Wildenthal, Zeitschrift für Physik A (1995).

[4] J. Brenneisen, D. Grathwohl, M. Lickert, R. Ott, H. Röpke, J. Schmälzlin, P. Siedle, and B. Wildenthal, Zeitschrift für Physik A (1995).

[5] M. Meyer, I. Venter, and D. Reitmann, Nuclear Physics (1975).

[6] P. Endt, Nuclear Physics A (1998).

[7] S. Kelly (2020), unpublished manuscript.

[8] S. J. Quinn, *Capture Cross Sections for the Astrophysical P Process*, Ph.D. thesis, Michigan State University (2012).

Revisiting the gallium anomaly: using *ab initio* nuclear theory to establish the ^{71}Ga$(\nu_e, e^-)^{71}$Ge cross section

HATTIE HARROLD

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Ragnar Stroberg

Abstract

Experiments examining the ^{71}Ga(ν_e, e^-)^{71}Ge reaction have consistently recorded a deficit in the amount of ^{71}Ge produced. The discrepancy between the observed and expected ^{71}Ge production amounts, known as the "gallium anomaly," is of major interest as it has been interpreted as possible evidence for the existence of sterile neutrinos. The ^{71}Ga(ν_e, e^-)^{71}Ge cross section is fundamental for determining the expected ^{71}Ge production rate, and thus for refining the statistical significance of the discrepancy between observed and expected production. This work employs the valence-space in-medium similarity renormalization group (VS-IMSRG) technique to update values for the ^{71}Ga(ν_e, e^-)^{71}Ge cross section and its uncertainty through an *ab initio* approach.

1 Introduction

The "gallium anomaly" has bewildered scientists for nearly 30 years. First observed in the mid-1990s, the anomaly arises from the Soviet-American Gallium Experiment (SAGE) and GALLEX experiments, which were originally used in relation to the solar neutrino problem. In the experiments, neutrinos are released from a radioactive source into a target of liquid ^{71}Ga. When a neutrino interacts with a ^{71}Ga nucleus via the weak force, it converts one of ^{71}Ga's 40 neutrons into a proton to produce ^{71}Ge (Z=32, N=39), as shown in Fig. 1.

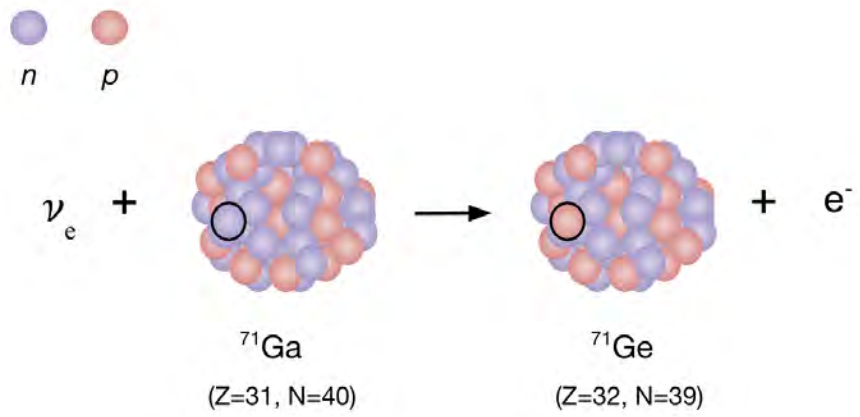

Figure 1: When a neutrino (ν) interacts with a neutron in ^{71}Ga, it converts the neutron into a proton to produce ^{71}Ge.

To physicists' surprise, however, the ^{71}Ga(ν_e, e^-)^{71}Ge reaction only produces approximately

80% of the amount of 71Ge expected, corresponding to a 2-3σ discrepancy [1]. Despite checks on potential experimental errors and updates to experimental setups, the ^{71}Ga(ν_e, e^-)^{71}Ge reaction continues to produce the same puzzling deficit of ^{71}Ge [2].

Though possible that the anomaly is attributed to some yet-undiscovered experimental error, many suspect that the anomaly– and other neutrino anomalies– may point to new physics, most notably, the existence of "sterile" neutrinos [1; 2]. Unlike the electron, muon, and tau neutrinos of the Standard Model (SM), these hypothesized sterile neutrinos would only interact gravitationally, and not via the weak force.

In order to better understand the extent of the anomaly and whether the sterile neutrino hypothesis is viable, it is important to refine the ^{71}Ga(ν_e, e^-)^{71}Ge cross section and its uncertainty, as this determines the expected production rate. In this work, we use an *ab initio* approach to estimate the cross section and its uncertainties "from first principles."

2 Background

Similar to how electrons inhabit different energy levels, or "shells," within an atom, so too protons and neutrons ("nucleons") within the nucleus follow a shell structure. A nucleus can therefore be in an "excited state" depending on the configuration of its nucleons, just as an atom can be excited depending on the arrangement of its valence electrons.

In the gallium-based experiments of interest, the neutrinos produced from radioactive sources of either ^{37}Ar or ^{51}Cr are at energies at which ^{71}Ga can transition not only to the ground state of ^{71}Ge but also to its first two excited states, as shown in Fig. 2. The ground state transition, ^{71}Ga(g.s.)(ν_e, e^-)^{71}Ge(g.s.), is proportional to its inverse reaction– electron capture (EC). Because the EC reaction is well understood experimentally, the cross section for the ground state transition is also precisely known [1]. In contrast, the excited state cross sections are poorly known and must be deduced from alternative means. This work aims to update these excited state contributions to the total cross section by modeling interactions between individual nucleons. These updates cannot solve the gallium anomaly, as a 2σ discrepancy would still exist even if only the ground state cross

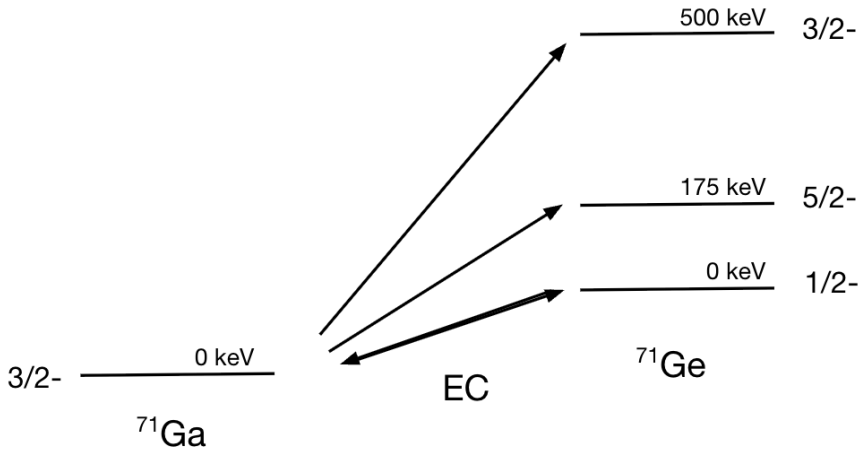

Figure 2: ^{71}Ga can transition to the ground state and the first two excited states of ^{71}Ge via ν capture. The electron capture (EC) reaction ^{71}Ge(g.s.)$(e^-, \nu_e)^{71}$Ga(g.s.) is the exact inverse of the ground state ν capture reaction, ^{71}Ga(g.s.)$(\nu_e, e^-)^{71}$Ge(g.s.).

section was considered [1]. Nevertheless, refining these excited state contributions also refines the statistical significance of the anomaly, potentially imposing stricter constraints on the hypothesized sterile neutrinos and providing insight into the possible source of the anomaly.

3 Methods

3.1 IMSRG Technique

The ^{71}Ga$(\nu_e, e^-)^{71}$Ge reaction is governed by the same Gamow-Teller (GT) transitions that describe β decay. Both the β decay rates and the neutrino capture cross sections are related (by some known kinematic factors) to the GT transition strength, B_{GT}. Previous works experimentally extracted B_{GT} values through the related forward-angle (p, n) and (^{3}He, t) reactions [1; 3]. However, mapping these values onto weak transitions can be problematic [1; 3]. Here, we instead derive the B_{GT} values directly from the transition strength matrix elements, M_{GT}. The matrix elements are defined as

$$M_{GT} = \langle \psi_f | \, | \hat{O}_{GT} | \, | \psi_i \rangle \,, \tag{1}$$

where $\hat{O}_{GT}$ is the GT operator (essentially, the operator responsible for converting a neutron to a proton) and the wavefunctions are derived from the time-independent Schrodinger equation,

$$\hat{H}\,|\psi\rangle = E\,|\psi\rangle\,. \tag{2}$$

We define the Hamiltonian, $\hat{H}$, by using the framework of chiral effective theory to model the "many-body" interactions between nucleons (see ref [4]).

For large nuclei like ^{71}Ga and ^{71}Ge, however, the Hamiltonian becomes too large and complex to fully diagonalize, making the many-body problem nearly impossible to solve. To make the problem more tractable, we use the in-medium similarity renormalization group (IMSRG) technique, specifically the valence-space IMSRG (VS-IMSRG), to approximate the Hamiltonian. The VS-IMSRG technique uses continuous unitary transformations to suppress the off-diagonal components of the Hamiltonian and drive it toward a block-diagonal form (see refs. [5] and [6] for more details). This allows us to approximately solve the many-body problem without actually constructing and diagonalizing the complete Hamiltonian.

The same transformations used on the Hamiltonian are then also applied to the GT operator, allowing us to derive the M_{GT} values using Eqn. 1. Finally, the B_{GT} values are related to the M_{GT} values by the relationship

$$B_{GT} = \frac{1}{2j_i + 1}|M_{GT}|^2. \tag{3}$$

3.2 IMSRG Inputs

Both the GT operator and the wavefunctions that result from solving the many-body problem with the approximated Hamiltonian are sensitive to different inputs used in the IMSRG process. We tested how the M_{GT} values are affected by changes to the following available inputs:

- Interaction: The interactions use different low-energy constants (LECs) and expand to different orders in the Effective Field Theory (EFT) expansion in order to encode unresolved physics while remaining consistent with experimental data [5].

Values: EM1820 [7], DNNLOgo [8], EM7.5 [9], N3LOlnl [10]

- e_{max}: We construct a single-particle basis from eigenstates of the 3D harmonic oscillator up to energy $e_{max} = 2n + l$, where n is the principle quantum number and l is the angular quantum number (corresponding to the s, p, d, f, etc. orbitals). Accordingly, e_{max} truncates the number of levels taken into account. Ideally, all levels would be considered by taking e_{max} to infinity; however, we cannot do this in our calculations. Instead, we increase e_{max} until the results converge (at $e_{max} = 11$).

  Values: 5, 7, 9, 11

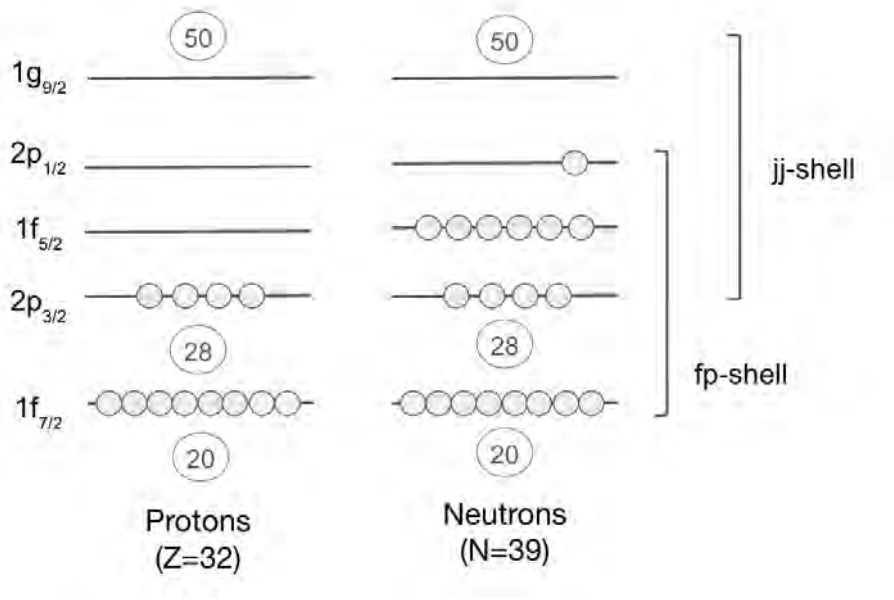

Figure 3: Different options for the valence space in IMSRG calculations. The circled values indicate "magic numbers," which correspond to filled nuclear shells. For ^{71}Ga (Z=31, N=40) and ^{71}Ge (Z=32, N=39), the valence space is the fp-shell. Alternatively, the "jj-shell" can be considered by disregarding the $1f_{7/2}$ level and instead including excitations to the $1g_{9/2}$ level.

- Model space: The valence space chosen for examining nucleon interactions. For ^{71}Ge (Z=32, N=39), the valence space for both the protons and neutrons is in the fp-shell, as shown in Fig. 3. Since there is a large gap between the completely filled $f_{7/2}$ level and the $p_{3/2}$ level, the $f_{7/2}$ nucleons likely won't have a large effect on the overall interactions. As a result, another option is to choose the "jj-shell," which excludes the $f_{7/2}$ level and instead includes the higher

$g_{9/2}$ level, as shown in Fig. 3.

Values: fp, jj

Finally, in addition to specifying the interaction, e_{max}, and model space inputs, we also tested the effects of including axial two-body currents ("2bc"), which arise in weak force interactions between two nucleons [11].

4 Results

The resulting M_{GT} values for each energy level are shown in Fig. 4, with the colors corresponding to the interaction and the symbol corresponding to the model space (circles for jj-shell; stars for fp-shell). Filled-in symbols indicate inclusion of two-body currents. Note that the two-body current results are slightly offset from the original results for clarity. As can be seen, including the two-body currents tended to bring the results for different interactions closer together. Increasing e_{max} also increased the convergence between interactions. This is expected since the interactions should all approach the same value eventually, but do so at different rates due to their use of different cutoff parameters.

As can be seen, the experimental M_{GT} for the ground state was within the spread of results for the theoretical interactions. However, the small discrepancy between the theoretical and experimental M_{GT} results became more pronounced as these values were extrapolated to the B_{GT} and cross section values. The M_{GT} values for the 5/2- transition provided little information, since they indicate that it is possible for this transition cross section to be effectively zero. $M_{GT}(3/2-)$, in contrast, appears to have a lower bound around 0.4. This corresponds to $B_{GT} = 0.04$, which is significantly greater than the $B_{GT}(3/2-)$ values obtained from the (p, n) and (^{3}He, t) reactions (0.011 ± 0.002 and 0.0176 ± 0.0014, respectively). As a result, these findings would have have a sufficient impact on the total cross section and the overall significance of the anomaly. However, because the ground state M_{GT} and B_{GT} values are also larger than expected, we cannot yet confidently conclude that $M_{GT} \geq 0.4$ and $B_{GT} \geq 0.04$.

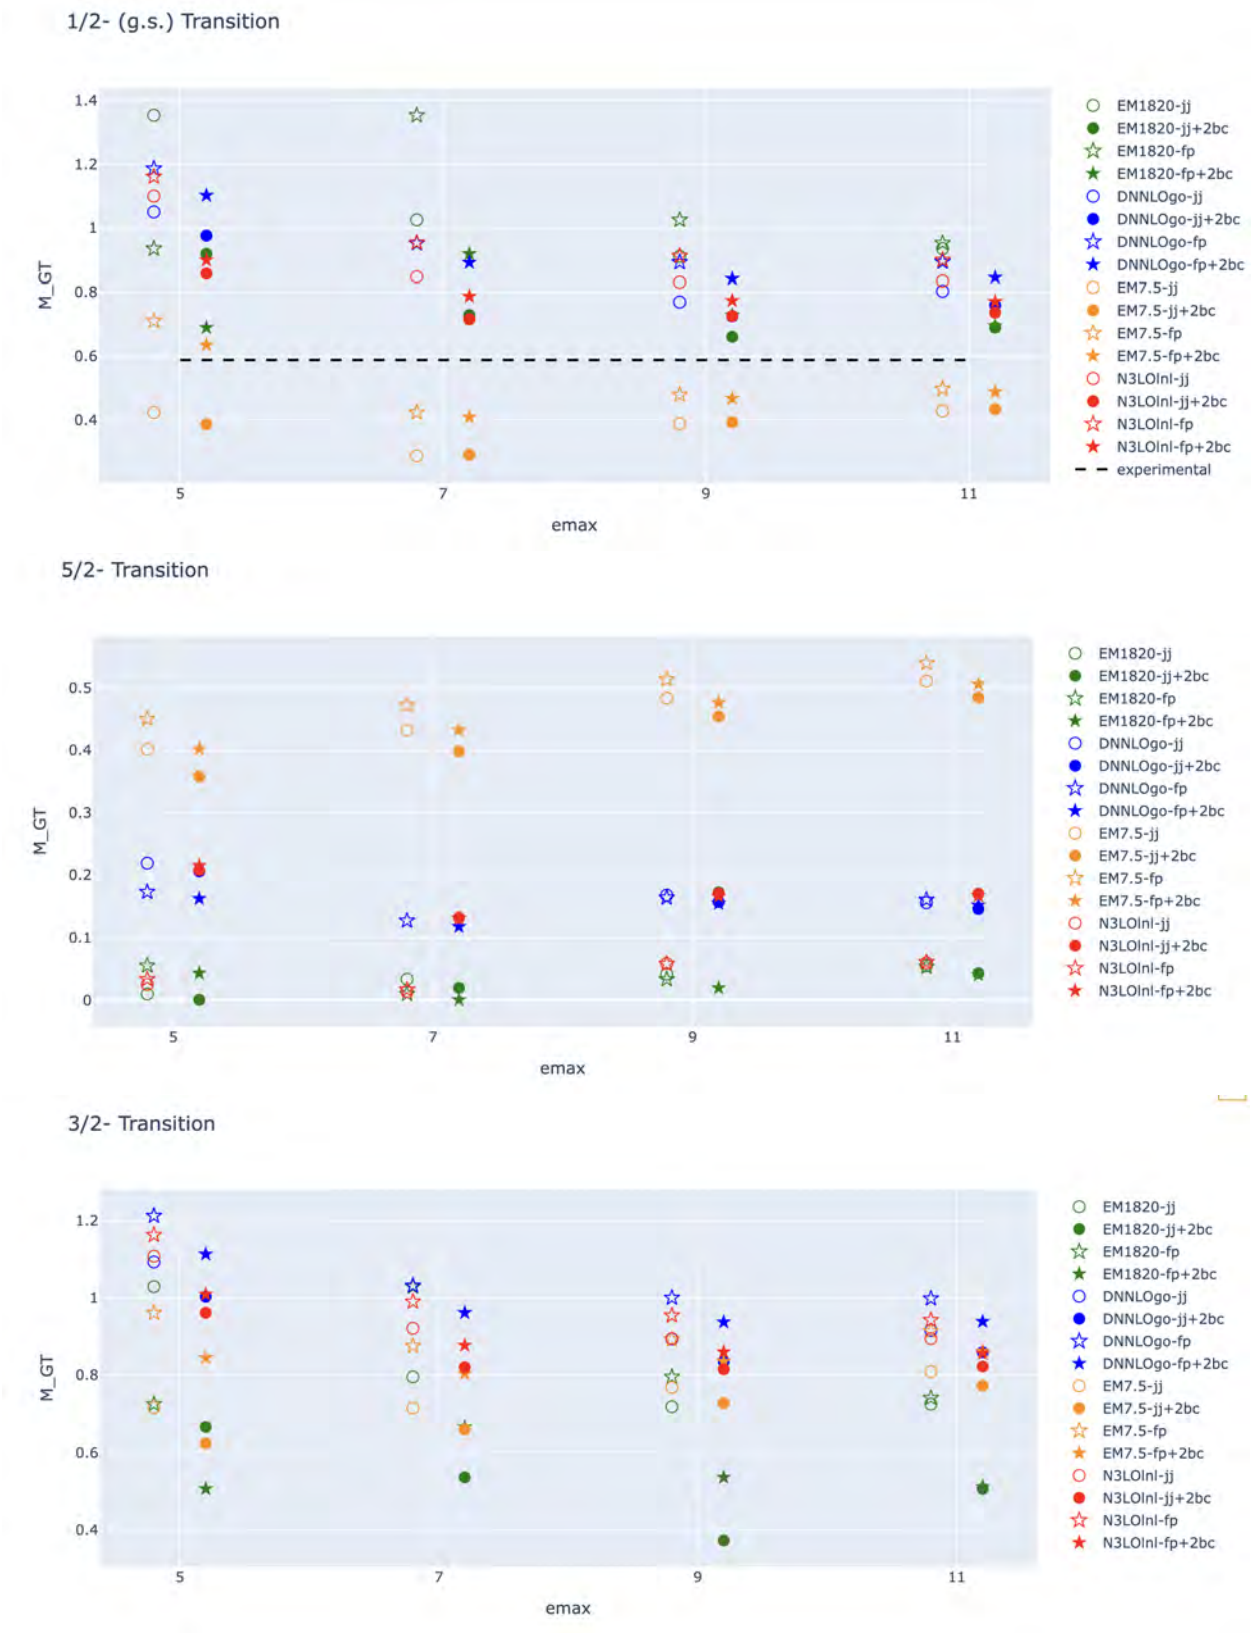

Figure 4: The transition matrix elements, M_{GT}, for each energy level. An expected experimental value (black dashed line) exists for the ground state M_{GT} value, but not for the excited states [1]

5 Conclusion

Though we did not make it to the point of calculating a reliable updated cross section, the work done here provides a foundation for making that calculation in the future. The current M_{GT} results for the 3/2- transition are suggestive of a lower bound around $M_{GT} = 0.4$. Future work can be done to confirm this by testing what modifications to the IMSRG calculations are required to obtain values below 0.04. If such modifications are not physically reasonable, this would allow us to conclude with more confidence that $M_{GT}(3/2-) \geq 0.4$.

Other work done here that can contribute to future research includes the evaluation of different IMSRG interactions. This was done by assessing their performance at producing observables for which experimental values are already known– namely, the magnetic moments and excitation energies. Upon examination, no best performers or outliers were identified.

We also explored the effects of the wavefunctions versus the GT operator on the final M_{GT} values and found that the wavefunction has a much stronger correlation, suggesting that the evolution of the Hamiltonian is of greater interest than that of the GT operator. Introductory work was started to explore how alterations to the wavefunctions impact the M_{GT} values. Continued work in this area can illuminate relationships between the wavefunctions and the M_{GT} values and provide insight into what changes should be made in the IMSRG calculations to produce wavefunctions that reduce the M_{GT} values.

Overall, this research provides a foundation for improving IMSRG calculations and establishing updated values for the $^{71}\mathrm{Ga}(\nu_e, e^-)^{71}\mathrm{Ge}$ excited cross sections in the future. Doing so will contribute to a better understanding of the gallium anomaly and its possible sources.

6 Acknowledgements

Thank you to my advisor, Professor Ragnar Stroberg, for developing and overseeing this research project. I am extremely grateful for his incredible patience, his willingness to answer questions, and his genuine investment in my research progress. I would also like to thank Dr. Umesh Garg and

Kristen Amsler for organizing this REU program and ensuring that we all had a safe, productive, and enjoyable summer. This research was funded by the National Science Foundation (Grants No. PHY-2050527 and No. PHY-2011890).

References

[1] S. R. Elliott, V. N. Gavrin, T. V. Ibragimova, and E. J. Rule, Physical Review C **108**, 035502 (2023).

[2] V. V. Barinov, B. T. Cleveland, S. N. Danshin, H. Ejiri, S. R. Elliott, and D. Frekers, Physical Review Letters **128**, 232501 (2022).

[3] N. Hata and W. C. Haxton, Physics Letters B **353**, 422 (1995).

[4] R. Machleidt and D. R. Entem, Physics Reports **503**, 1 (2011).

[5] H. Hergert, Frontiers in Physics **8** (2020).

[6] K. Tsukiyama, S. K. Bogner, and A. Schwenk, Physical Review Letters **106**, 222502 (2011).

[7] K. Hebeler, S. K. Bogner, R. J. Furnstahl, A. Nogga, and A. Schwenk, Physical Review C **83**, 031301 (2011).

[8] W. G. Jiang, A. Ekstrom, C. Forssen, G. Hagen, G. R. Jansen, and T. Papenbrock, Physical Review C **102**, 054301 (2020).

[9] P. Arthuis, K. Hebeler, and A. Schwenk, arXiv (2024).

[10] V. Soma, P. Navratil, F. Raimondi, C. Barbieri, and T. Duguet, Physical Review C **101**, 014318 (2020).

[11] A. Baroni, L. Girlanda, S. Pastore, R. Schiavilla, and M. Viviani, Physical Review C **93**, 015501 (2016).

A Graph Neural Network Approach for Momentum Reconstruction in the EMPHATIC Experiment

YUAN HE

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Fields

Abstract

Accurate momentum reconstruction is critical for precision measurements in high-energy physics experiments such as EMPHATIC. Traditional approaches like the Kalman filter are often limited by their assumptions of linear dynamics and Gaussian noise. In this study, we explore the use of Graph Neural Networks (GNNs) for downstream momentum prediction. Using simulated datasets at various beam energies, we evaluate the model's performance and generalization capability. We further propose architectural improvements and multi-domain training strategies to enhance the robustness of GNN-based momentum reconstruction.

1 Introduction

Track reconstruction in high-energy physics experiments is an important but challenging task. It includes recognition of a set of hits with a single particle and obtaining kinematic parameters of the particle through track fitting. In the EMPHATIC experiment, precise reconstruction of downstream hadron kinematics—such as momentum and scattering angle—is essential for accurate measurements of production and scattering cross sections, thereby providing reliable experimental inputs for neutrino physics research.

A commonly used method for track reconstruction in high-energy physics and nuclear physics is the Kalman filter [1], which is a recursive algorithm that estimates the state of a dynamic system from a series of noisy measurements. However, Kalman filters rely on two critical assumptions: linear dynamics and Gaussian measurement noise, which are often not satisfied in real detector environments [2]. For example, charged particle trajectories are inherently curved when passing through magnetic fields, which significantly deviate from the linear approximation required by the Kalman approach. Additionally, the Kalman filter's accuracy strongly depends on initial conditions; improper initialization can cause rapid divergence due to accumulated linearization errors. Such limitations make it challenging to achieve stable and precise track reconstruction using this method alone.

In contrast, Graph Neural Networks (GNNs) offer a powerful alternative for capturing non-linear and complex relationships inherent in detector data, especially when represented in a graph. GNNs are first introduced in [3] and have been widely applied in various domains such as social

networks, quantum chemistry, and physics. In the context of track reconstruction, GNNs enable connection of a sets of hits from a single particle by edge classification, reducing the computational cost significantly [4]. Besides, kinematic parameters of a track can also be predicted with graph regression model; by directly leveraging the relationships between detector hits without explicit linearization, this GNN-based approaches overcome many of the traditional limitations faced by Kalman filtering, enabling more robust and accurate estimation of particle kinematics.

In this report, we train and evaluate a GNN-based model to predict the downstream momentum of hadrons in the EMPHATIC experiment with 5 different beam energies, demonstrating the applicability of GNNs in momentum reconstruction.

2 Background

Our model is trained and tested using the dataset generated by simulation based on the EMPHATIC experiment geometry. EMPHATIC (Experiment to Measure the Production of Hadrons At a Test-beam In Chicagoland) is a compact, table-top-size experiment capable of high-rate data collection to measure hadron production and scattering cross sections. A top view of the experimental setup is shown in Figure 1. Details of the EMPHATIC experimental design and components are available in [5].

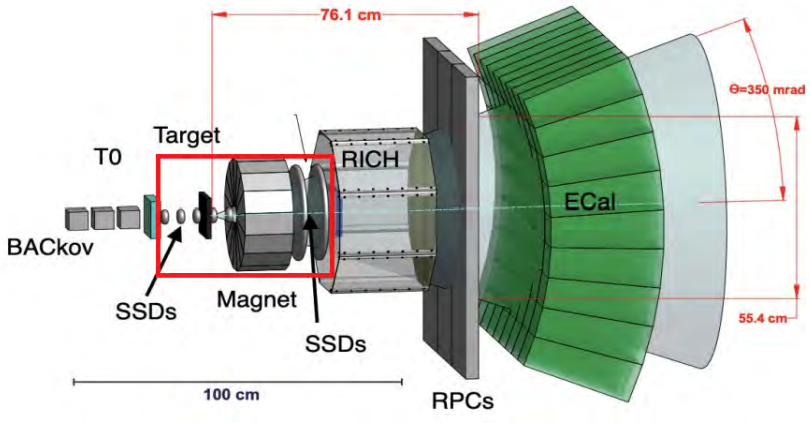

Figure 1: Top view of the EMPHATIC experimental setup.

The components relevant to our dataset collection is in the red box in Figure 1. After passing

through T0 and before hitting the target, a set of small-area silicon strip detectors (SSDs) will record the pre-target hit positions of hadrons for trajectory reconstruction. The target can be material such as carbon, iron, aluminum and water, After hitting the target, another sets of SSDs will be applied to record the hit positions. Then the charged particles traverse through a magnet, which bends the particles according to their momenta. A set of large-area SSDs is arranged after the magnet to record the hit positions. Data collected by the SSDs upstream and downstream of the target will be used to reconstruct the trajectories of particles for momentum calculation.

As an example, Figure 2a shows the configuration of SSDs in EMPHATIC phase 1c. The target is graphite. Each white box represents a SSD station, in which there are multiple planes(See Figure 2b). Each plane may consists of more than 1 SSD sensor and each sensor consists of many strips numbered from 0 to 640(See Figure 2c).

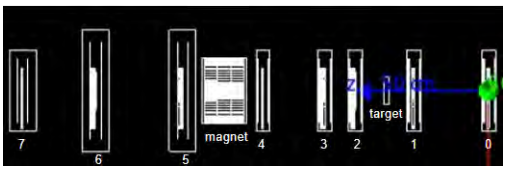

(a) Top view of SSD alignments in phase 1c

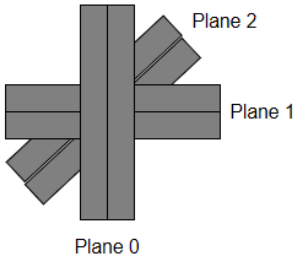

(b) Three planes in station 6

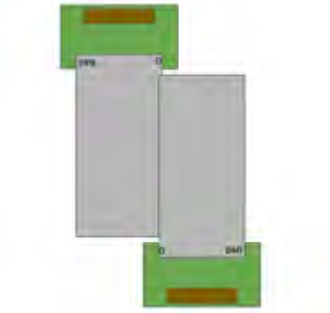

(c) Two sensors of plane 0

Figure 2: SSD geometry in Phase 1c and examples of internal structure.

3 Methodology

The goal of our model is to predict kinematic quantities—specifically, the downstream momentum vector of a particle—based on its associated detector hits and the features of each hit. This task falls under the category of graph-level learning, and more precisely, graph regression. In the following subsections, we first describe the dataset generation process and the selection of hit-level features in Section 3.1. Then, we present the architecture of our model in Section 3.2.

3.1 Dataset Generation

To obtain the ground-truth downstream momentum of particles—which serves as the target labels for training—we generate the dataset via a detailed Monte Carlo simulation based on the `Geant4` toolkit, using the geometry of the EMPHATIC experiment. We utilize the `EMPHATICSoft` framework [6], an open-source software developed for the EMPHATIC experiment, which integrates detector simulation, digitization, and data processing via the `art` framework. We can extract datasets from the final CAF file.

The simulation is based on the Phase 1c geometry of the EMPHATIC experiment(Illustrated in Section 2). We run five simulation corresponding to different inital beam energies which are chosen to match typical running conditions used in the recent EMPHATIC Phase 1c data-taking in 2023. The total momentum P follows a Gaussian distribution(See Table 1), while the transverse momentum components (p_x, p_y) are uniformly sampled and the angular spread $\frac{p_x}{p_z}$ as well as $\frac{p_x}{p_z}$ are constrained in $[-10^{-4}, 10^{-4}]$. The beam spot is uniformly distributed over a $4\,\text{cm} \times 4\,\text{cm}$ square in the transverse plane.

simulation	Pmean/(Gev/c)	Psigma/(Gev/c)
1	4	1
2	8	0.2
3	12	0.2
4	20	0.2
5	120	0.2

Table 1: Gaussian parameters of incidental proton beams corresponding to five simulations

For each beam energy, we simulate 20000 events, where each event represents the interaction of a single incident proton with the detector. Depending on the physics process, the proton may undergo inelastic scattering and produce daughter particles, or interact elastically (or quasi-elastically) without producing any secondaries. For our training and testing, we only select events where a single proton passes through the detector without generating any daughter particles. The momentum prediction of other charged particles would be the same if the model really learns the pattern from trajectories.

For each selected event, we extract hit-level features from the `truth.truthhits` branch, including the SSD station number, plane number, sensor number, strip number and SSD orientation, which contains the positional information of each hit. However, due to the phenomenon of charge sharing — where ionization at a single hit location may activate multiple neighboring silicon strips — the same station-plane-sensor triplet can be associated with multiple strip numbers. This introduces redundant nodes in the graph representation and increases noise in the dataset, potentially leading to overfitting in the early training stages and reducing training efficiency.

To alleviate the problem, we use cluster dataset from the `cluster.clust` branch, which aggregates nearby activations and uses the average strip number to represent a single effective hit. Additionally, we append three features to describe the SSD strip orientation based on the known alignment geometry of Phase 1c SSDs. The final input consists of 7 features per hit: (1) station, (2) plane, (3) sensor, (4) average strip number, and (5–7) SSD orientation.

The label for each event—the true downstream momentum of the proton—is extracted from the `truth.beam.Trajectory()` branch, specifically at the first recorded point after the proton exits the target.

3.2 Model Design

The model architecture, originally designed in [7] for upstream momentum prediction, has been adapted in this work to predict downstream momentum. The task aims to predict a three-dimensional momentum vector for each particle event, given a set of detector hits characterized by spatial features.

To enable efficient training, we adopt a mini-batch strategy, where the data is reshaped into tensors of shape [B, N, F], with B denoting the batch size(i.e., number of events per batch), N the number of hits per event and F the number of features per hit. The hit count per event is uniformly set to $N = 500$ using zero-padding, enabling consistent tensor dimensions across batches.

Figure 6 shows the architecture of the model and transformation of a batch tensor.

The model employs a hybrid architecture of MLPs and GNNs. Specifically, each input batch is

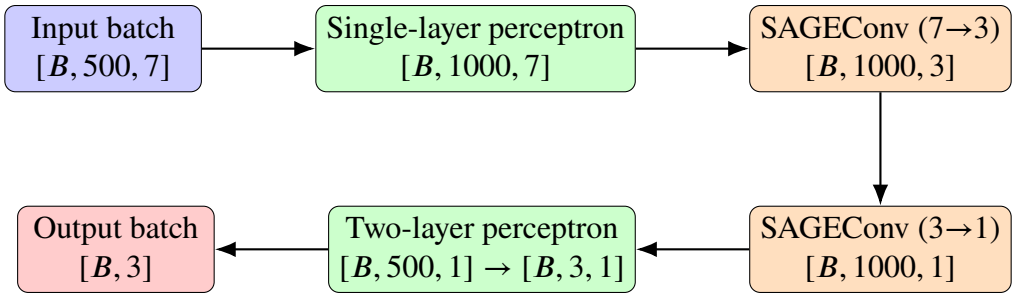

Figure 3: Architecture of the model and transformation of batch data.

first processed by a single-layer perceptron(i.e. linear transformation and activation), doubling the number of nodes from 500 to 1000 via trainable linear transformation to improve the expressiveness of the model.

Prior to applying the graph convolution layers, a unique graph structure is introduced: half of the nodes are fully interconnected to facilitate comprehensive global information exchange, while the other half remain isolated to preserve local information. Graph convolution are performed by GraphSAGE algorithm [8] , which is implemented by `SAGEconv` module in `Pytorch Geometric`. The constructed graph is processed through two successive `SAGEconv`layers, where the node features are updated by message passing and aggregation. The feature dimensionality is sequentially reduced from 7 to 3, and finally to 1.

A subsequent two-layers perceptron is applied, reducing the nodes number into 3. The resulting three nodes, each holding a single scalar feature, correspond directly to the three components of the predicted momentum vector.

4 Results and Analysis

4.1 Experimental Setup

We implement our model using `torch` 2.1.2, `pytorch-lightning` 2.1.3, and `pytorch_geometric` 2.4.0. The training hyperparameters are summarized in Table 2.

Training on the 4 GeV, 8 GeV, and 120 GeV datasets is performed on the University of Notre Dame's Center for Research Computing (CRC) using 4 GPUs. The experiments for the 12 GeV

Table 2: Training hyperparameters used in all experiments.

Hyperparameter	**Value**
Label Scaling	`MinMaxScaler` [0, 1]
Train/Validation/Test Split	0.85 / 0.10 / 0.05
Batch Size	1000
Gradient Accumulation	100 batches
Loss Function	`SmoothL1Loss`
Optimizer	Adam (lr = 10^{-3}, weight decay = 10^{-6})
Reload DataLoaders Every Epoch	`True`

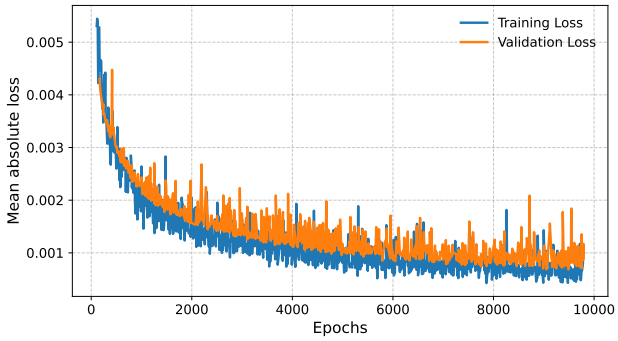

Figure 4: Training and validation loss curves

and 20 GeV datasets are conducted on Fermilab's Elastic Analysis Facility (EAF) with 1 GPU.

4.2 Model Performance

Due to the length limitations of this report, we only present the most representative among the five trained models — the one trained on the 4 GeV cluster dataset. It is trained for 4 days and 9700 epochs.

To evaluate the model's performance, we plot the training and validation losses to examine potential overfitting or underfitting, and to identify the checkpoint with the best validation performance. As shown in Figure 4, both the training and validation losses consistently decrease when training epochs increase, and there are no signs of overfitting (i.e., increasing validation loss) , suggesting that additional training epochs may further improve performance.

To further assess the model's generalization capability, we evaluate it on an independent dataset generated from a different simulation but with the same parameter configuration. Figures 5a and 5b

show the model's performance on data from the same simulation used in training, while Figures 5d and 5c present the results on the unseen dataset from a separate simulation.

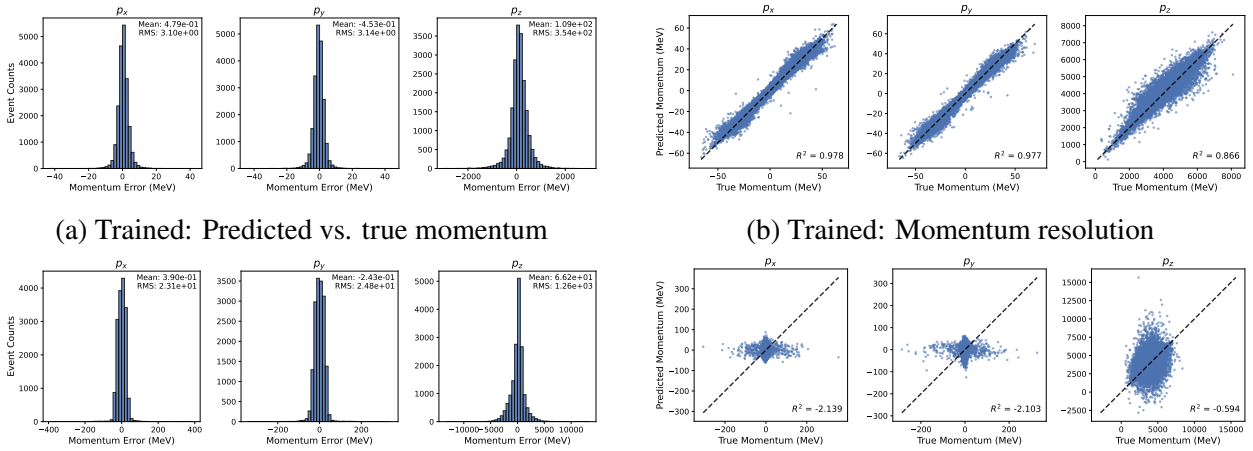

(a) Trained: Predicted vs. true momentum

(b) Trained: Momentum resolution

(c) Untrained: Predicted vs. true momentum

(d) Untrained: Momentum resolution

Figure 5: Performance comparison on trained (top) and untrained (bottom) datasets: (a, c) momentum resolution distributions; (b, d) predicted vs. true momenta.

As illustrated in Figure 5a and Figure 5b, the model achieves strong agreement between predicted and true momentum values on the dataset derived from the same simulation used in training. However, Figure 5d and Figure 5c reveal a notable performance drop when the model is tested on an independent dataset generated from a different simulation, albeit with the same parameter configuration. The increased root mean square error (RMS) and reduced coefficient of determination (R^2) across all momentum components suggest that the model's generalization remains limited.

This contrast highlights that the model, while capable of capturing dataset-specific statistical patterns—as evidenced by the steadily decreasing validation loss—fails to learn robust trajectory-level representations that are transferable across simulations. In other words, it does not fully extract underlying physical correlations invariant to simulation variations.

4.3 Further Steps

To improve generalization, we analyze the limitations of the current approach and propose modifications.

First, the current graph construction—fully connecting half the nodes while leaving the other half isolated—may weaken the model's ability to capture spatial ordering and local patterns. This structure likely limits the model's ability to learn consistent physical patterns across datasets. To address this, we propose a new model where the input hits are connected in a one-dimensional ordered chain, preserving trajectory structure(Figure 6).

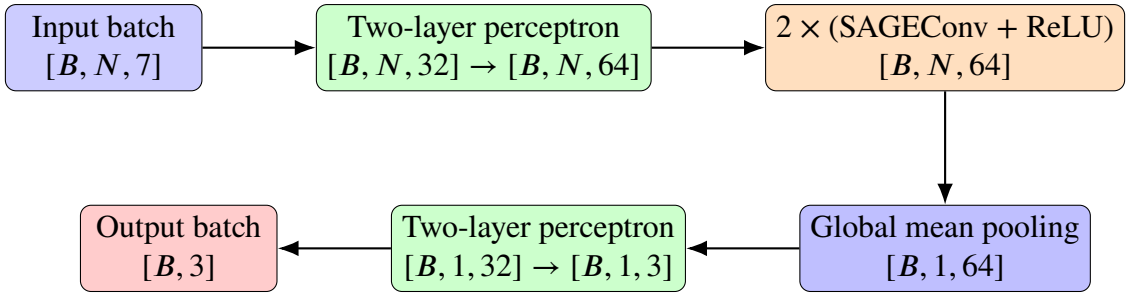

Figure 6: Proposed new model architecture with trajectory-preserving graph structure.

Second, incorporating training data from diverse simulations could prevent the model from overfitting to the statistical distribution of a single dataset. This multi-domain training strategy may improve the model's robustness and generalization.

Third, upstream hits are unnecessary for predicting downstream momentum, which is primarily determined by the bending trajectory after the target. Moreover, daughter particles lack upstream hits, so excluding them ensures consistency and broader applicability across events.

5 Conclusion

In this report, we have explored a GNN-based approach for downstream momentum reconstruction in the EMPHATIC experiment. The model shows strong predictive performance when trained and tested on data from the same simulation but struggles to generalize across different simulation instances. These results highlight both the promise and the current limitations of GNNs in this task. The model requires further refinement to capture robust, simulation-independent physics. To address this, we proposed architectural changes and multi-domain training strategies, paving the way for more generalizable reconstruction methods.

References

[1] R. E. Kalman, Journal of Basic Engineering **82**, 35 (1960).

[2] X. Ai, H. M. Gray, A. Salzburger, and N. Styles, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **1049**, 168041 (2023).

[3] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, IEEE Transactions on Neural Networks **20**, 61 (2009).

[4] M. D. Ju, Xiangyang *et al.*, The European Physical Journal C **81**, 10.1140/epjc/s10052-021-09675-8 (2021).

[5] T. Akaishi, L. Aliaga-Soplin, H. Asano, A. Aurisano, M. Barbi, L. Bellantoni, S. Bhadra, W. Chang, L. Fields, A. Fiorentini, *et al.*, arXiv preprint arXiv:1912.08841 (2019).

[6] E. Collaboration, Emphaticsoft: Simulation and reconstruction framework for the emphatic experiment, `https://github.com/EMPHATICSoft`.

[7] C. Woolford and E. Team, in *APS April Meeting Abstracts*, Vol. 2024 (2024) pp. M16–004.

[8] W. L. Hamilton, R. Ying, and J. Leskovec, Inductive representation learning on large graphs (2018), arXiv:1706.02216 [cs.SI] .

Direct Imaging of Binary Stars

MEGAN HEGARTY

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Jeffrey Chilcote

Abstract

The age of a star is important for creating evolutionary models for exoplanets around it. Rather than finding the ages of individual stars, it is possible to determine the age of a moving group using the binary stars within it. Observing the orbit of binary stars reveals their mass, which can then be used with their luminosity to determine age. Direct imaging of binary stars using Non-Redundant Masking can be used to study stars that are too close together to resolve. Placing a mask with 9 holes in the optical path causes the light from the stars to interfere, generating an interference pattern from which the contrast, separation, and position angle of the stars can be extracted. These values can later be used to determine the orbit of the binary stars.

1 Introduction

Determining the age of stars can reveal information about the planets orbiting them. Therefore, an inaccurate age for a star causes inaccurate evolutionary models for its planets. To resolve this error, moving groups of stars are opportunities to use binary stars to estimate the age of a large number of stars[1]. The age of a star is typically found through evolutionary models using its mass and luminosity. Although it is difficult to determine the mass of a single star, binary stars provide the opportunity to directly observe their orbit, which in turn reveals their mass [2]. The orbit is determined using a combination of radial velocity and direct imaging techniques.

Obtaining useful data from direct imaging involves both calibration and analysis. Calibration steps correct for systemic errors within the detector while analysis involves actually processing the data in the image[3]. HD 16760, HIP 14807, HD 30051, and 26 Gem are binary star systems observed using Non-Redundant Masking (NRM). Through this method, it is possible to determine the contrast, separation, and position angle of stars that are too close together to resolve using conventional photometry. These values reveal information about the position of the stars that can be used to help determine their orbit.

2 Background

HD 16760, HIP 14807, HD 30051, and 26 Gem are all binary stars which are part of different moving groups. A moving group is a collection of stars that move together, but are not gravitationally bound. Their similar movement implies that they were formed from the same material around the same time. Therefore, determining the age of specific stars within the moving group implies the age of the entire group[1]. Binary stars are particularly valuable for this type of analysis because it is possible to directly observe their orbit[2]. Using the resulting orbital period and separation, it is possible to make a dynamical mass determination. Mass in conjunction with luminosity can then be used to estimate age.

Accurately determining the orbit involves a combination of techniques, namely radial velocity and direct imaging. Radial velocity involves analysis of the Doppler shift of the stellar spectra. Direct imaging involves measuring the light coming directly from the targets[3]. It allows for direct measurements of the separation, position angle, and contrast, difference in brightness, of the stars in the binary system. HD 16760, HIP 14807, HD 30051, and 26 Gem in particular were observed using Non-Redundant Masking (NRM) because the stars are too close together to resolve. A mask with 9 holes was placed in the optical path to create an interference pattern with the light from the stars similar to one created from a double slit experiment. This pattern could be compared to the interference pattern of a well-understood reference star. An analysis of the differences in the interference patterns made it possible to model the target stars and determine their contrast, separation, and position angle[4]. This data, in combination with the information gained from radial velocity analyses, can inform a more accurate orbit.

3 Methods

3.1 Calibration

Before the raw data could be analyzed, it had to undergo a data reduction process utilizing two different types of calibration frames to account for the systemic error of the detector. Dark frames

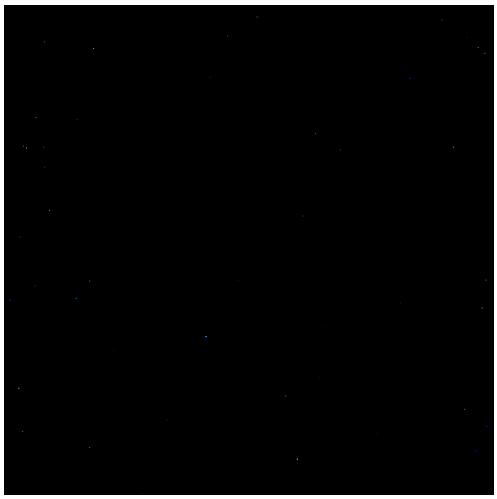

(a) Master dark for 10.0s integration time

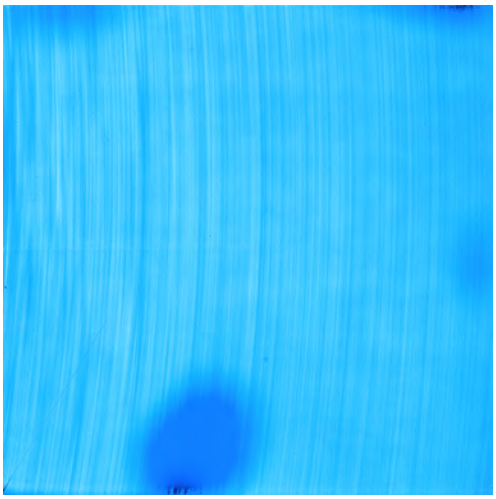

(b) Master flat with filter J

Figure 1: Examples of calibration frames.

are taken when there is no light shinning on the detector. They are a measure of the dark current, the systemic thermal effects of the detector. Because this effect is time dependent, a different set of dark frames is necessary for each integration time. Conversely, flat frames measure variability in pixel response. Ideally, the same amount of light would result in the same response from each pixel; however, not all pixels are identical and so do not have identical responses. These frames also account for anything that may be interfering in the optical path, such as dust. A separate master flat is created for each filter. Filters can be used to specify a range of wavelengths to detect.

A unique master dark was created for each integration time and image size to be used to calibrate each science image. After grouping the frames by configuration, they were median combined so for each pixel the median across all frames was retained. Median combination is preferable to mean combination because the median is less impacted by outliers such as those caused by cosmic rays. An example of a master dark is shown in Figure 1a.

Different master flats were generated based on the filter and presence of a coronograph. A coronograph blocks light coming from a star to reveal faint objects orbiting it. Although no coronograph was used for NRM frames, it was necessary to account for it when creating the master flats because the presence of a coronograph significantly impacts the flat. Flat frames are slightly more complicated to combine than dark frames. There are two types of flat frames taken, one with a lamp

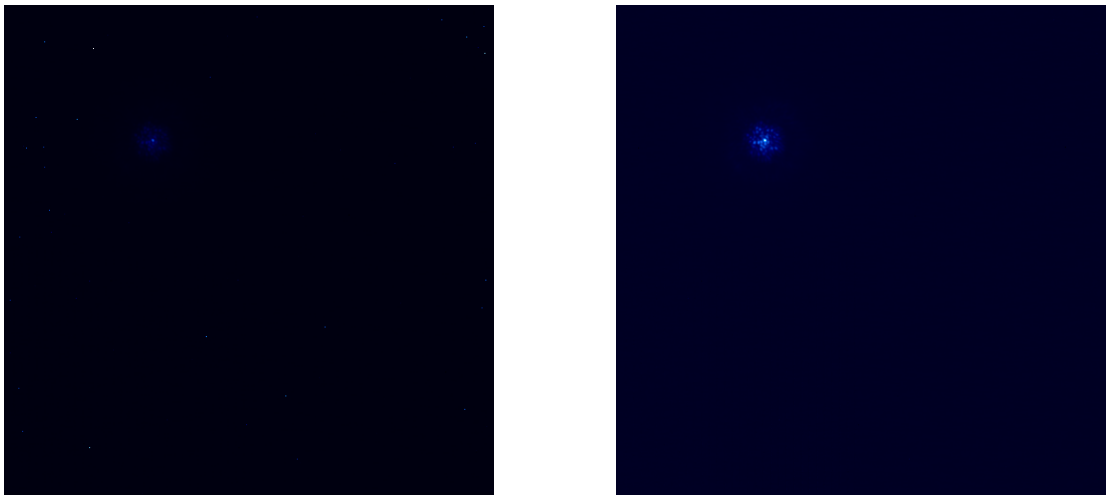

(a) Uncalibrated image (b) Calibrated image

Figure 2: HD 16760 calibrated with master flat J and 1.0s master dark.

on and the other with the lamp off. For each configuration, the group of lamp on and lamp off flats were median combined. The goal of the lamp is to illuminate the detector uniformly to reveal the relative sensitivity of each pixel. However, when the lamp is on, the output of the detector is not solely the result of the lamp, as it also includes the dark current. This is why it is necessary to also take flat frames with the lamp off, which will only include the dark current. Subtracting the on and off frames for each configuration resulted in a master flat impacted only by the lamp.

Master flats also need to correct bad pixels in the detector, namely saturated pixels that were exposed to too much light and no longer collect reliable information. This was done by identifying all pixels whose values are more than 5 standard deviations away from the mean. Those values were then replaced with a weighted average of the surrounding pixels. After correcting the bad pixels, the flat was normalized, meaning each pixel was divided by the median value, so it could adjust for pixel sensitivity. An example of a master flat is shown in Figure 1b

After the master darks and flats were created, the science images were calibrated. Calibrating an image required finding the correct master dark and flat for the specific configuration. The master flat needed to have the same filter and the master dark needed to have the integration time closest to that of the science image. The dark was then subtracted from the science image to remove the thermal effects of the detector and the image was divided by the flat to correct for non-uniform pixel

sensitivity. After the calibration was complete, the image should be fairly uniform, with the only bright spots being the target. NRM images should have a bright center surrounded by clearly defined fringes, both brighter areas of constructive interference and dark spots of destructive interference. An example of an NRM image before and after calibration is shown in Figure 2.

3.2 NRM Analysis

The first step in NRM analysis was quantifying the observed interference pattern. Specifically, a closure phase was calculated for each unique triangle of holes. The triangle index is the number assigned to each unique triangle to identify the specific closure phases. The closure phase is the phase difference between the light that reaches three particular holes in the mask. For the 9 hole mask used, there are 84 unique triangles and therefore 84 closure phases for each target. In order to find the closure phases, a unique reference star was used for each target, because the reference star needed to be a well understood star system with a similar composition. With this reference, it was possible to determine the closure phases of the target for each triangle index. A Monte Carlo Markov Chain simulation (MCMC) was then set up to analyze the closure phases. This type of simulation tests a wide variety of values within a specified range for each parameter by having a number of walkers independently model the interference pattern that would be created from certain choices of parameters. After testing a specific model against the measured interference pattern, each walker then moved independently to new parameter values. The processes continued until they converged on a probability distribution for each parameter.

Running the MCMC required a few inputs, namely priors and starting positions for the walkers. The priors contain known information in the form of probability distributions for possible values of each parameter. In this case, the priors were a flat probability within the allowed range for each parameter. Providing the starting positions for the walkers can improve the necessary runtime of the simulation. If the walkers start close to the actual values of the parameters, the simulation will converge faster. These starting values were selected from a χ^2 grid, which is a measure of how far the model is from the observed data for a certain choice of parameters. The walkers were initialized

to values surrounding the minimum of this grid (corresponding to the most probable parameters). Running the simulation with these inputs resulted in a probability distribution for each parameter.

4 Results

Target	Filter	Contrast	Separation (mas)	Position Angle (∘)
HD 16760	J	0.0364 ± 0.0030	30.20 ± 1.02	212.4 ± 1.6
	H	0.0454 ± 0.0025	28.13 ± 1.13	211.0 ± 1.1
	Kp	0.0769 ± 0.0157	26.24 ± 1.73	212.1 ± 0.7
HIP 14807	J	0.2448 ± 0.0466	53.97 ± 3.78	309.1 ± 11.4
	H	0.2859 ± 0.0168	55.51 ± 0.27	307.7 ± 3.3
	Kp	0.3157 ± 0.0013	56.56 ± 0.08	307.8 ± 0.1
HD 30051	J	0.2643 ± 0.0052	29.04 ± 0.23	105.6 ± 0.4
	H	0.3178 ± 0.0025	29.25 ± 0.13	106.05 ± 0.16
	Kp	0.3362 ± 0.0046	29.03 ± 0.12	105.6 ± 0.1
26 Gem	Kp	0.06764 ± 0.0086	33.71 ± 2.18	84.91 ± 1.05

Table 1: Contrast, separation, and position angle for HD 16760, HIP 14807, and HD 30051 in the J, H, and Kp bands and for 26 Gem in the Kp band.

The process described above was completed for four different binary stars: HD 16760, HIP 14807, HD 30051, and 26 Gem. Images taken with different filters were processed separately. Table 1 contains the exact values obtained from this analysis. The separation and position angle describe the relative positions of the stars which is important for determining their orbit. The contrast can help determine luminosity, which is also necessary to estimate the age of the star.

The results of the MCMC simulation are presented as a corner plot and as a comparison between the model and the data. Figures 3 and 4 show these results for single target. Figure 3 is a corner plot which shows the probability distribution of each pair of parameters. If the simulation runs correctly and converges, the probabilities should be fairly constrained and, in particular, the distribution of a single parameter, those along the diagonal, should resemble a gaussian. Where the corner plot shows if the simulation converges, Figure 4 reveals how accurate the model is to the data. It is a plot of the closure phase versus the triangle index for both the model and the data. This is important to consider because if the model does not align closely with the data, the simulation most likely did

not run correctly and the results are unreliable.

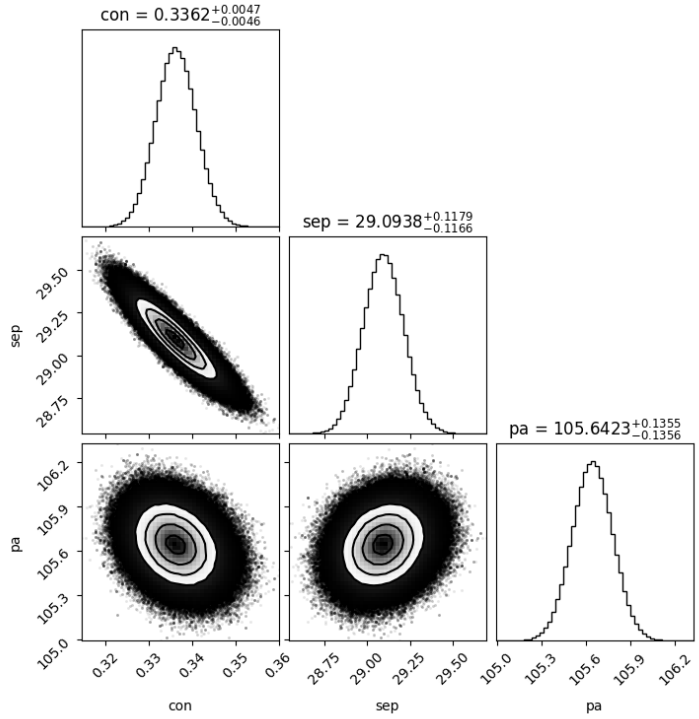

Figure 3: Probability distributions for the contrast, separation(mas), and position angle(∘) of HD 30051, Kp band.

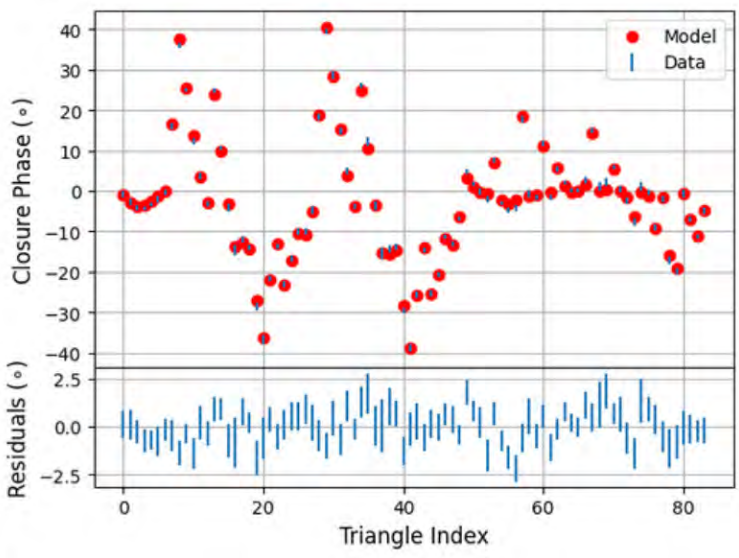

Figure 4: Compares the closure phases of HD 30051, Kp band for the model to the calculated closure phases from the data.

5 Conclusion and Next Steps

Using this analysis process it was possible to obtain the contrast, separation, and position angle of HD 16760, HIP 14807, HD 30051, and 26 Gem. These values are each a single observation that will be used with radial velocity measurements and other direct imaging observations to estimate the orbits of the systems and make dynamical mass measurements. The masses will then be used with the luminosities in evolutionary models to infer the ages of the stars. The ages of the individual star systems will be able to contribute to estimates of the ages of the moving groups they are part of.

References

[1] E. L. Nielsen, R. J. D. Rosa, J. Wang, J. Rameau, I. Song, J. R. Graham, B. Macintosh, M. Ammons, V. P. Bailey, T. S. Barman, J. Bulger, J. K. Chilcote, T. Cotten, R. Doyon, G. Duchêne, M. P. Fitzgerald, K. B. Follette, A. Z. Greenbaum, P. Hibon, L.-W. Hung, P. Ingraham, P. Kalas, Q. M. Konopacky, J. E. Larkin, J. Maire, F. Marchis, M. S. Marley, C. Marois, S. Metchev, M. A. Millar-Blanchaer, R. Oppenheimer, D. W. Palmer, J. Patience, M. D. Perrin, L. A. Poyneer, L. Pueyo, A. Rajan, F. T. Rantakyrö, D. Savransky, A. C. Schneider, A. Sivaramakrishnan, R. Soummer, S. Thomas, J. K. Wallace, K. Ward-Duong, S. J. Wiktorowicz, and S. G. Wolff, The Astronomical Journal **152** (2016).

[2] A. D. Thomas, E. L. Nielsen, R. J. De Rosa, A. E. Peck, B. Macintosh, J. Chilcote, P. Kalas, J. J. Wang, S. Blunt, A. Greenbaum, Q. M. Konopacky, M. J. Ireland, P. Tuthill, K. Ward-Duong, L. A. Hirsch, I. Czekala, F. Marchis, C. Marois, M. A. Millar-Blanchaer, W. Roberson, A. Smith, H. Gallamore, and J. Klusmeyer, The Astronomical Journal **166** (2023).

[3] K. B. Follette, Publications of the Astronomical Society of the Pacific **135** (2023).

[4] A. Sivaramakrishnan, D. Lafrenière, K. E. Saavik Ford, B. McKernan, A. Cheetham, A. Z. Greenbaum, P. G. Tuthill, J. P. Lloyd, M. J. Ireland, R. Doyon, M. Beaulieu, A. Martel,

A. Koekemoer, F. Martinache, and P. Teuben, Proceedings Volume 8442, Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave; (2012).

Determination of the $^{10}B(\alpha,\alpha')$ Cross Section Using Python

SHOSHANNA HERTZ

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Dr. Richard deBoer, Dr. Edward Stech

Abstract

Nuclear reactions between alpha particles and light nuclei have uses in a variety of fields, from the study of nuclear astrophysics and nuclear structure to non-destructive nuclear assay. However, reactions such as the $^{10}B(\alpha, \alpha')$ reaction have not been well-studied at low energies. Data from the $^{10}B(\alpha,n)$ experiment conducted at Notre Dame during June 2025 was analyzed to fit the 718.35 keV gamma peaks resulting from inelastic scattering. Using these yields, the differential cross section for $^{10}B(\alpha, \alpha')$ at 130°, 80°, and 36° was determined for energies between 2.71 and 4.24 MeV.

1 Introduction

Nuclei possess a shell structure similar to the electron shell structure in atoms. Reactions with high energy particles - such as those produced by particle accelerators - can excite the nucleus. The nucleus almost immediately de-excites by emitting one or several gamma rays as it returns to its ground state. By using equipment like high purity germanium (HPGe) detectors or a similar high resolution gamma ray detector to collect data and then analyzing the resulting gamma spectra, the reaction cross section - a measurement of the probability of a nuclear reaction occurring - can be determined. In this project, data from the $^{10}B(\alpha,n)$ experiment performed at Notre Dame was analyzed using a Python script in order to determine the $^{10}B(\alpha,\alpha')^{10}B$ differential cross section.

2 Methods

2.1 Experimental Setup

The Notre Dame 5U accelerator was used for the $^{10}B(\alpha,n)$ experiment. The 5U is a 5 MV vertical pelletron accelerator, often used for nuclear astrophysics experiments. For this experiment, three HPGe detectors were set up surrounding the 5U solid target line. Each of the three detectors were placed at different angles between 0° and 180°, with one detector at a forward angle, one at a backwards angle, and one central detector placed around 90°. The backwards detector was positioned at an angle of 130°, 43.2 cm away from the target. The forward detector was placed at an angle of 36° and 32.2 cm from the target. The central detector was at an angle of 80° and placed 33.0 cm

away from the target. Three detectors were used in order to capture as many of the gamma rays as possible. All three detectors were calibrated using Cesium-137 and Cobalt-60 gamma sources. The beam energy for the ^{10}B(α, n) experiment ranged from 1.44 MeV to 4.24 MeV. The 718.35 keV[1] gamma peak (corresponding to the first excited state in ^{10}B) was only able to be resolved in a small number of these runs, with beam energies ranging from 2.67 MeV to 4.24 MeV. Once the runs concluded, a 1.019 μCi ^{152}Eu source was used in order to determine the efficiency of the HPGe detectors. The ^{152}Eu source was left in the solid target line and data was collected overnight.

2.2 Peak Fitting

In order to find the yield from both the 718 keV gamma peak and the gamma peaks from the ^{152}Eu source, a Python script was written to handle three important processes. First, the script had to be able to extract the raw data from the detectors. Secondly, this data had to be displayed in a form that could be easily viewed and interpreted by the user. Lastly, this script had to be able to fit the peaks of interest in the gamma spectra.

The first two processes that the code handled were relatively simple. The script reads in the .mpa file from the detector and extracts the following information from each detector: run time, live time, calibration offset, calibration factor, and the number of counts in each channel. From there, the appropriate calibration offset and calibration factor were applied to the channels in order to obtain the energy values. The step function from matplotlib was then used to display the full gamma spectrum. The user could then easily zoom in on a specific region of interest.

In order to capture the shape, several parameters were required for each fit. These parameters were the upper and lower bound for the fit, the points that seemed to span 97.5% of the gaussian and the mean energy of the peak. These parameters were used to obtain the initial guesses for the slope and intercept of the background, the mean value, and the standard deviation, as well as the amplitude of the gaussian. An additional input was also added to account for skew, in an attempt to account for any Doppler shift. The curve fit function in Scipy was used to get an initial fit for the gamma peak. If the fit was unsatisfactory, the user then was given the opportunity to change the

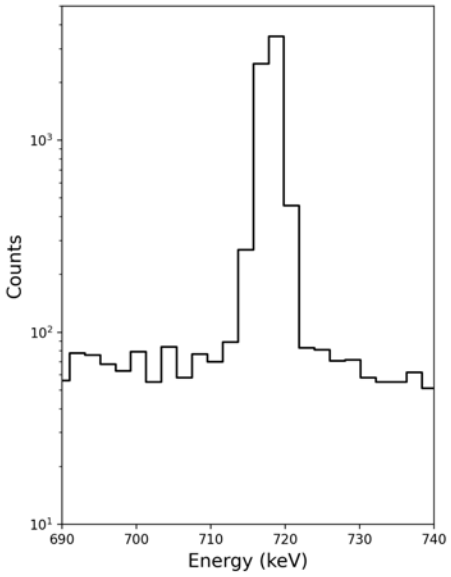

Figure 1: A section of the gamma spectrum from the forward detector, showing the 718 keV gamma peak.

initial parameters and rerun the curve fit function. An example of a fitted 718 keV gamma peak is seen in figure 2.

The gamma peak is fitted over a 6σ range, centered at the mean value for the peak. In that range, everything under the linear offset is taken to be part of the background, while everything above the linear offset is taken to be part of the gamma peak. The yield is determined by determining the total number of counts in the 6σ range (N_T), and subtracting the number of counts in the background (N_B) for that same range. The uncertainty in yield, or σ_Y for the gamma peak can be found through propagation of uncertainty.

$$\sigma_Y = \sqrt{N_T + N_B} \tag{1}$$

2.3 Efficiency Calculation

The gamma yield from the ^{152}Eu source was determined using the same Python script and fitting method described above. The gamma energies and branching ratios were taken from the Table of Radionuclides, Vol. 2 [2]. In order to get efficiency, the expected yield for each peak was deter-

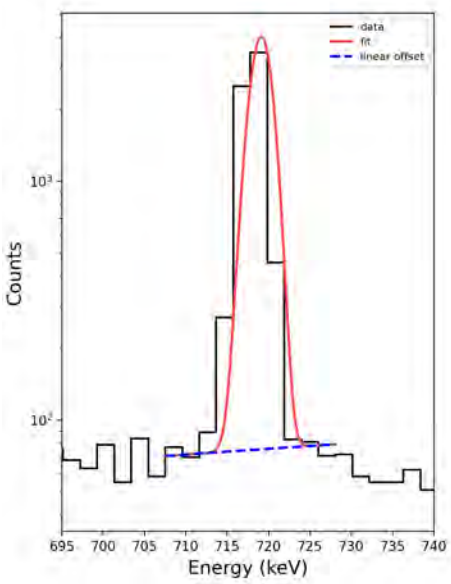

Figure 2: An image of the fitted 718 keV gamma peak for the forward detector.

mined using the current activity of the ^{152}Eu source, the live time for each run, and the branching ratio for each peak. The yield from the fit was then divided by the expected yield in order to determine the efficiency of the detector at different energies. The efficiency curve was then approximated using a variation of the Hyperion efficiency equation[3]. The original exponential term seen in the Hyperion efficiency equation was omitted as there was no turnoff point seen in the efficiency of the HPGe detectors.

$$\epsilon = a \times E^{-b} + c \tag{2}$$

The curve fit function from Scipy was used to obtain the best fit parameters for the efficiency curve, and based off of the residual, the uncertainty in the efficiency was estimated to be 15%.

2.4 Cross Section Measurement

A variation on the thin target yield equation was used to determine the cross section.

$$\sigma = \frac{A}{N_T N_P \epsilon} \tag{3}$$

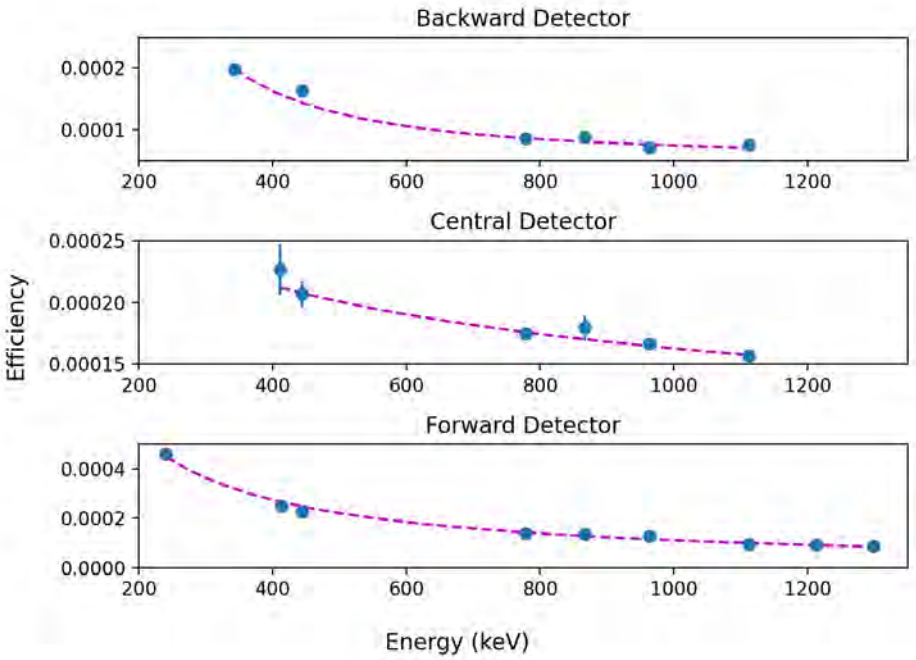

Figure 3: A graph of the efficiency curves for all three HPGe detectors.

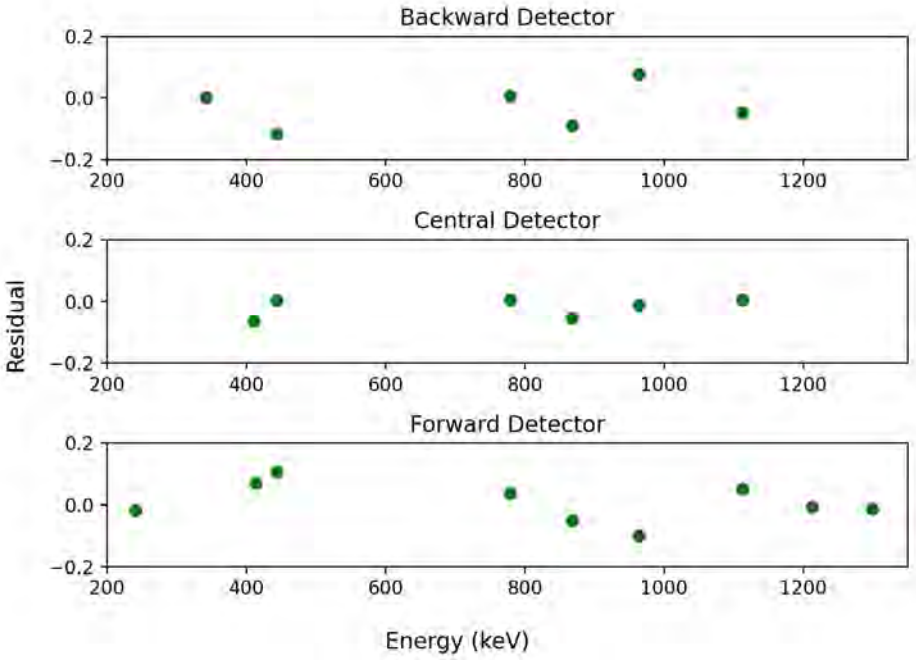

Figure 4: A graph of the residuals from the efficiency fit (in percent) for all three detectors.

The yield (A) was determined using the Python script outlined above. The number of target atoms (N_T) was determined from a scan of the 606 keV resonance in ^{11}B, and the number of incident particles (N_I) came from the charge integrator. The cross section was determined to have a total systematic uncertainty of 18.28%.

3 Results

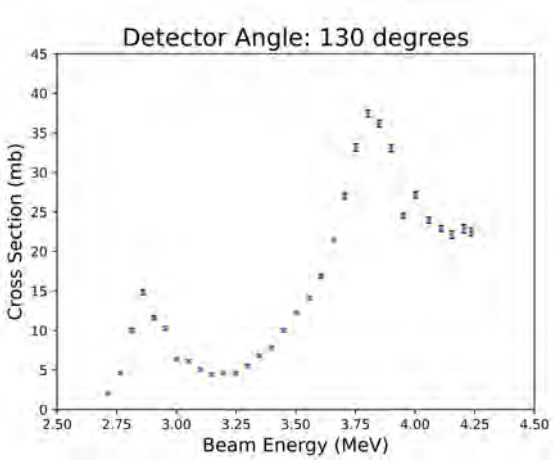

Figure 5: A graph of the excitation curve from 2.71 to 4.24 MeV for a HPGe detector at 130°.

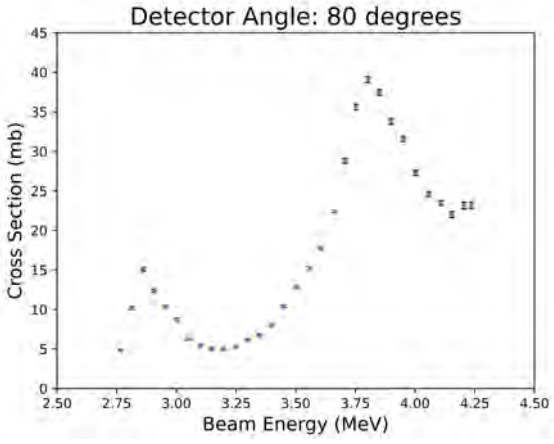

Figure 6: A graph of the excitation curve from 2.71 to 4.24 MeV for a HPGe detector at 80°.

The excitation curves for $^{10}B(\alpha,\alpha')$ at 130°, 80°, and 36° can be seen in figures 5, 6, and 7. The general trend for the cross section follows the trend seen in a 1969 paper measuring intensity

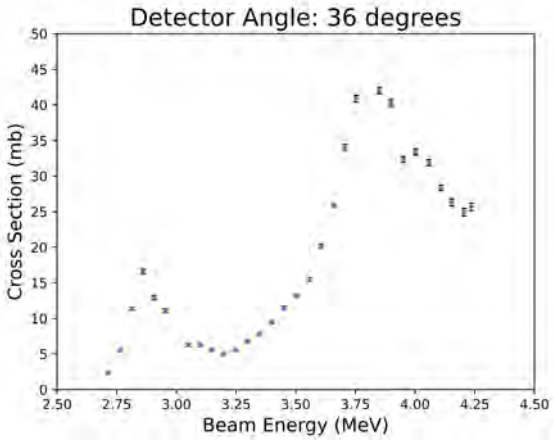

Figure 7: A graph of the excitation curve from 2.71 to 4.24 MeV for a HPGe detector at 36°.

of the 718 keV gamma rays between 1.0 and 3.5 MeV[4]. The same tail was seen around 2.90 MeV, but unlike in this experiment, they were able to resolve the 718 keV gamma peak at lower energies. This was not possible at energies lower than 2.71 MeV. At that point, the gamma yield dropped sharply, and it became impossible to fit the gamma peaks.The results shown in this paper also captures new features, such as the peak at 3.8 MeV, suggesting the existence of a resonance at that energy. Additionally, while there are some small differences in the cross section between angles, it all falls within uncertainty, suggesting that there is no angular dependence in the ^{10}B(α,α') cross section.

Much of the error in the cross section measurement came from uncertainty in detector efficiency. The uncertainty in detector efficiency at 718 keV (determined through error propagation) was unreasonably large: as much as five times the value of actual efficiency. Only 7 to 8 of the gamma peaks from the ^{152}Eu source could be fit for each detector, mainly due to the presence of room background. Many of the peaks from the source overlapped with room background. Since there was no background run done for this experiment, the yield from those peaks was not used when fitting for detector efficiency, as the contribution from room background may have artificially increased efficiency at those energies. This issue could have been addressed by taking a background run prior to the start of the experiment. Additionally, few peaks from the ^{152}Eu source could be fitted between 500 and 750 keV due to low yield. In the future, this gap in data could be addressed

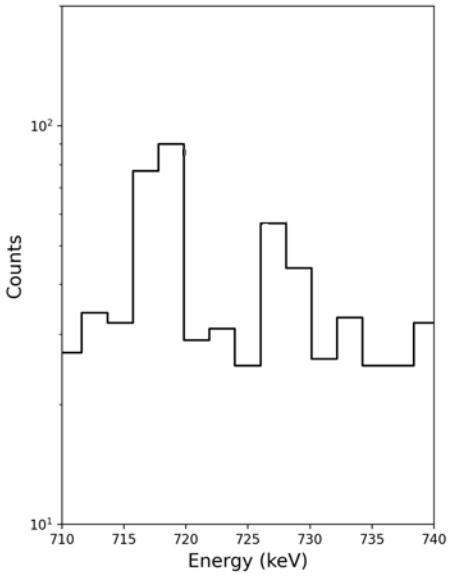

Figure 8: An image of the 718 and 727 keV gamma peaks seen in the forward detector for a beam energy of 2.62 MeV.

with a longer run with the ^{152}Eu source, or by using multiple gamma sources in order to cover the 500 to 750 keV energy range.

4 Conclusion

Reactions such as ^{10}B(α,α') and ^{10}B(α,n) provide useful information for studying nuclear astrophysics and nuclear structure, as well as having applications in non-destructive nuclear assay. The differential cross section measurements for ^{10}B(α,α') are valuable, but more investigation is needed, especially at energies below 2.7 MeV.

5 Acknowledgements

I would like to thank my advisors, Dr. James deBoer and Dr. Edward Stech, as well as Dr. Maxwell Sorenson for their help with this project. I would also like to thank Dr. Umesh Garg and Kristen Amsler for making this REU possible.

References

[1] D. Tilley, J. Kelley, J. Godwin, D. Millener, J. Purcell, C. Sheu, and H. Weller, Nuclear Physics A **745**, 155 (2004).

[2] M.-M. Bé, V. Chisté, C. Dulieu, E. Browne, V. Chechev, N. Kuzmenko, R. Helmer, A. Nichols, E. Schönfeld, and R. Dersch, *Table of Radionuclides*, Monographie BIPM-5, Vol. 2 (Bureau International des Poids et Mesures, Pavillon de Breteuil, F-92310 Sèvres, France, 2004).

[3] R. Hughes, J. Harke, R. Casperson, S. Ota, S. Fisher, J. Parker, C. Beausang, M. Dag, P. Humby, J. Koglin, E. McCleskey, A. McIntosh, A. Saastamoinen, A. Tamashiro, E. Wilson, and T. Wu, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **856**, 47 (2017).

[4] A. Gallmann, F. Hibou, and P. Fintz, Nuclear Physics A **123**, 27 (1969).

Visualizing the R-Process and Interpreting Nucleosynthesis Simulations

ROSE HEWALD

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Rebecca Surman

Abstract

The rapid neutron-capture process, or r-process, is the set of nuclear reactions responsible for the formation of roughly half of all elements heavier than iron, such as gold and uranium. However, the particular astrophysical site and nuclear conditions favorable for the r-process to occur remain unclear. To better understand this mystery, we use PRISM (Portable Routines for Integrated nucleoSynthesis Modeling) software which generates data that predicts the time-evolution and final abundances of isotopes under different conditions. From these simulations, we then create visualization tools and plots to gain a clearer understanding of the r-process data. These visualization tools include plots of the nuclear abundance patterns, abundance ratios of key nuclei in parameter space, strength of shell closures, fission flows as compared to decay rates, and r-process evolution animations. The purpose of these visualizations is to compare how variations in initial conditions shape the outcome of r-process nucleosynthesis, thus identifying physical trends and possible signatures that point to specific astrophysical origins.

1 Introduction

1.1 Astrophysical Nucleosynthesis Background

The rapid neutron-capture process, or r-process, is the set of nuclear reactions responsible for the formation of roughly half of all elements heavier than iron. There are several feasible sites for r-process nucleosynthesis, including compact binary mergers [1] [2] [3] [4] [5] and certain types of supernovae, excluding normal core-collapse supernovae [6] [7] [8] [9] [10]. Very neutron-rich nuclei that exist far from stability participate in the r-process, and the properties of these nuclei are much more unknown than those that exist near stability. Therefore, we must connect nuclear theory with astrophysical observation in order to constrain the conditions and environments where the universe's heavy elements are formed.

1.2 Project Overview

Using our understanding of the astrophysical conditions, such as entropy, neutron richness, and temperature, we predict how variations of these factors influence abundances of heavy nuclei formed during the r-process. This is done using data from nuclear reaction network simulations from PRISM software [11], which generates data that predicts the time-evolution and final abundances of

isotopes under different conditions. The final r-process abundance pattern is made up of three main peaks at mass numbers A = 80, 130, and 195, which corresponds to the 'magic' neutron numbers N = 50, 82, and 126. Magic numbers indicate the location of a neutron shell closure, which occurs when filled neutron shells lead to particularly stable nuclei, temporarily halting neutron captures, and accumulating abundance at these points. As part of the 2025 University of Notre Dame REU, I have worked with the PRISM output data in order to create a number of visualization tools, including plotting and animating the nuclear abundance patterns and peaks as a function of neutron and proton number, tracking the productions of key nuclei using ratios of their abundances, comparing how variations in initial conditions shape the outcome of r-process nucleosynthesis, and identifying physical trends and possible signatures that point to specific astrophysical origins. In this report, I will outline the process of completing PRISM simulations, the types of plots and analysis done with the data, and the directory of scripts that I have created for future analysis.

2 Methods

PRISM, which is a Fortran written modeling software, completes time-dependent analysis and tracking of nucleosynthesis reactions. Its primary purpose is to compute the abundances as a function of proton number, atomic mass number, and time, where the relative abundances are normalized as

$$\sum Y(Z, A, t) \times A = 1 \tag{1}$$

Where Y is the abundances, Z is proton number, A is mass number, and t is time.

The formulation of the system of equations that PRISM uses to produce the output data sets are dependent on the astrophysical inputs from the user. In its calculations, a number of assumptions based on physical principles must be made. This includes assuming that the system is closed, the nuclei and their abundances evolve over a series of timesteps, and the nuclei are uniformly distributed in both position and velocity.

The work of running PRISM simulations was completed by graduate student Pranav Nalamwar, and the input data that this process required were files identifying the initial state of the system, files explaining the thermodynamic evolution as a function of time, and files containing the nuclear physics. Once the calculations have been completed, the result is a number of output files, including three final composition files: one listing the relative abundances of the nuclei sorted by their proton and mass numbers, the next provides atomic number and the abundances of all nuclei associated with a particular Z value, and the last is similar but provides all nuclei associated with a particular atomic mass number A. PRISM also outputs time evolution information, mass-fractions, rates, and flow functions.

The PRISM calculations are run for different mass models, which are theoretical frameworks that fit the binding energy (mass) of a nucleus based on the physics principles or statistical fits to known data. These models help fill in the gaps where experimental data may not exist. It is important that we compare many different nuclear mass models because they uniquely influence the neutron separation energies (S_n), beta-decay rates, and fission probabilities, which, through comparison, reveals theoretical uncertainties, the model-dependence of results, and robust r-process signatures. In this report, we will focus on the mass models DZ33 (Duflo-Zuker) [12], FRDM (Finite Range Droplet Model) [13], HFB (Hartee-Fock-Bogoliubov) [14], and UNEDF1 (Universal Nuclear Energy Density Functional) [15].

The visualization tools and plots are used to help us better understand the vast amounts of output data that PRISM provides in each mass model. We use the various plots and animations to study specific conditions, properties, and calculations so that we can target patterns and characteristics of the r-process. The plotting is done using Python both to parse the data from the various PRISM files and to utilize the data in a way that makes it visually comprehensible and useful for our studies.

3 Results

3.1 Nuclei Abundance Ratios

Thus far in our work, we have produced a large number of plots focusing on many particular properties. One of our first investigations involved studying the abundance ratios of notable nuclei. A notable ratio of nuclei is the uranium-to-thorium (U/Th) ratio, which is significant because both are long-lived radioactive isotopes produced directly in the r-process, and their relative abundances can be used to estimate the time since nucleosynthesis occurred and late-time conditions. We also study the iridium-to-gold (Ir/Au) ratio because they are positioned in the third r-process peak, thus revealing how well a particular nucleosynthesis model reproduces the A 195 peak shape and relative abundances. Similarly, the tellurium-to-iodine (Te/I) ratio can be found in the second r-process peak region at A=130, so it can be used to study earlier r-process conditions for lighter nuclei.

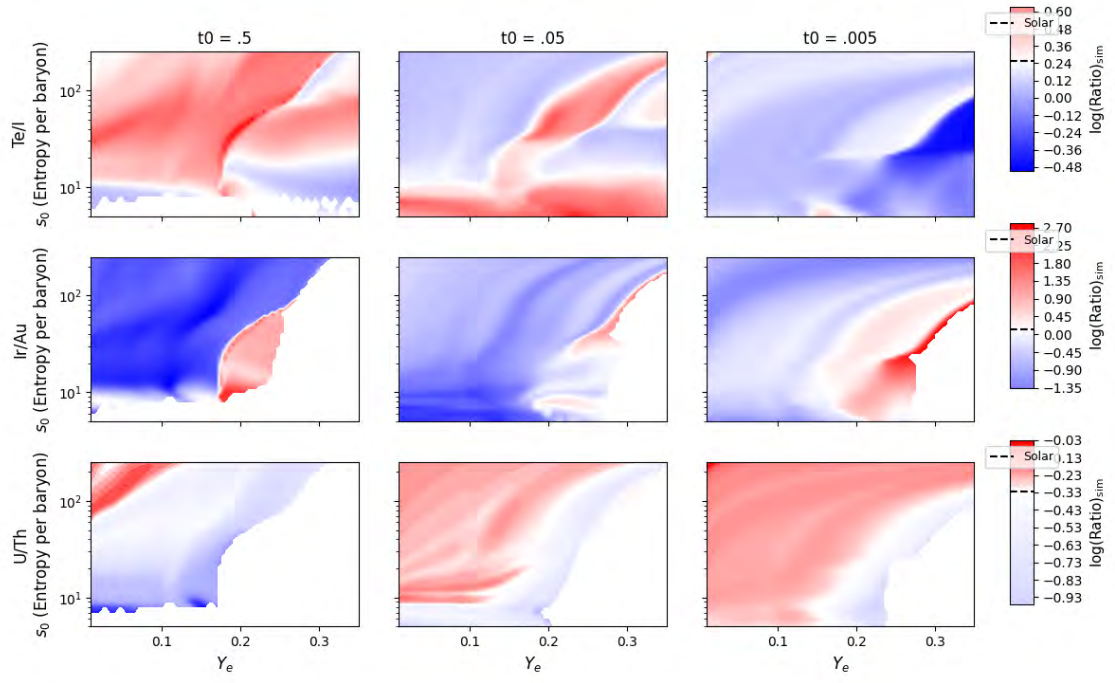

Figure 1: A comparison of nuclei ratios at different points in the r-process using FRDM. Te/I demonstrates properties of the second peak, while Ir/Au shows those of the third peak, and U/Th reveals information about the actinides. The black dashed line on the color bar represents the solar data value. The columns are divided by timescale, with the fastest timescale being 0.005 seconds on the far right and the slowest timescale being 0.5 seconds on the far left.

This data from FRDM highlights the general inability of the slowest timescale (0.5 s) to match the solar data [Figure 1]. Particularly in the second and third peaks, we see dramatic and widespread straying from the solar, and this pattern is repeated across the other mass models for this timescale. The slanted band structure that we see in most of these calculations is also a repeated pattern across the various mass models, and it reveals locations of fission cycling, which is when heavier nuclei undergo fission in the r-process, thus returning material to lighter elements and affecting the overall final abundances.

3.2 Neutron Separation Energy

Final r-process abundances are also heavily affected by the two-neutron separation energy difference (D2n), which is the difference required to remove two neutrons from a nucleus and that needed to remove two neutrons from a nucleus with initially two less neutrons. A sharp peak in this value can indicate a shell closure, or a build up in abundances, and it may predict a shifted peak or a change in peak height.

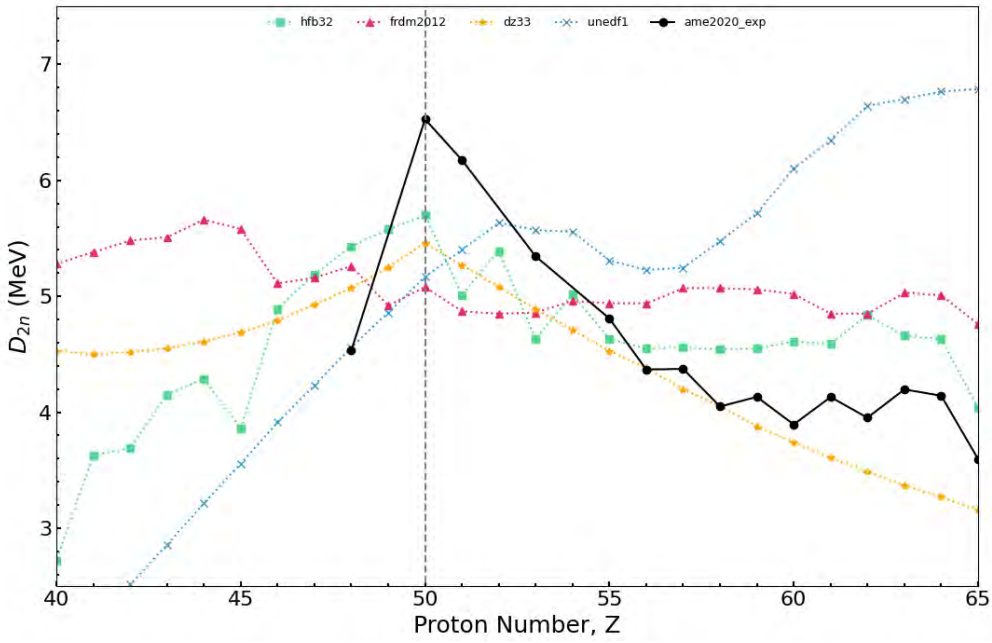

Figure 2: A plot showing the two-neutron separation energy as a function of proton number for N = 82. We see a peak at Z = 50, marked with a vertical grey dashed line, for the experimental AME2020 data, corresponding to a neutron shell closure at this point.

The shell closure at Z = 50 is clearly marked by the solid black AME2020 data line [Figure 2]. We expect the same sort of behavior to be replicated by the mass model data, but this sort of peaked trend is demonstrated by DZ33 at this point, and not by UNEDF1, FRDM, or HFB.

3.3 Actinides

The actinides are the radioactive elements with atomic numbers Z = 89 to 103 which are dominantly produced by the r-process and, therefore, can help identify sites where the r-process may have occurred depending on whether it is actinide-rich or actinide-poor. The production of actinides are very sensitive to nuclear physics inputs, and their final abundances can vary based on electron fraction, entropy, strength of shell closures, and fission cycling. When comparing the relative abundances of three different mass models for the 2nd peak, 3rd peak, and actinides, the difference in actinide production was most fascinating [Figure 3].

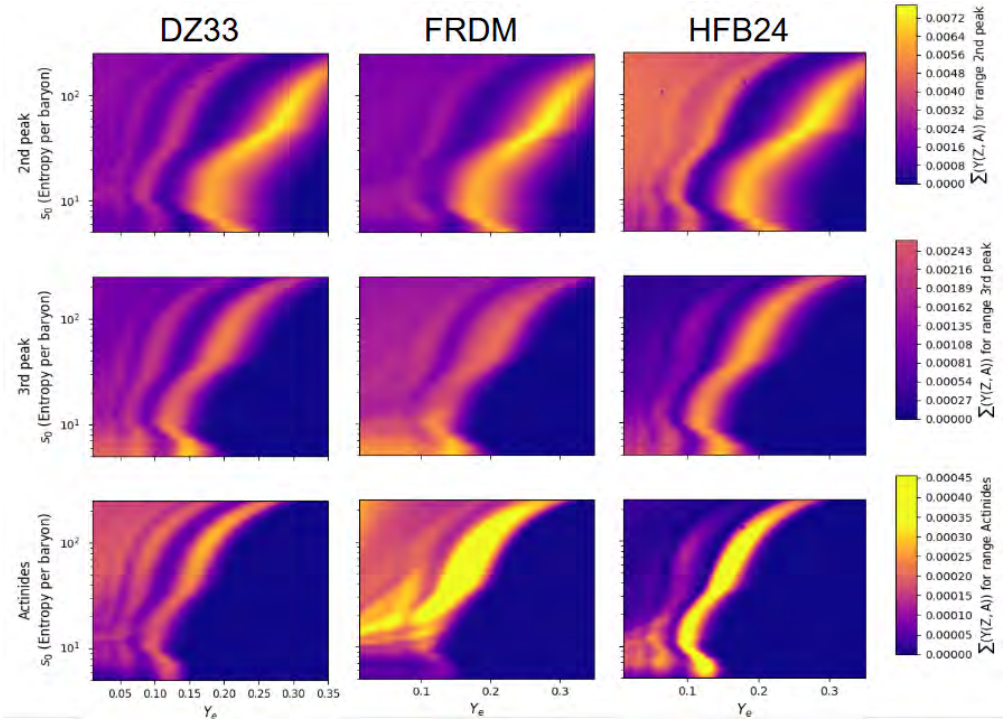

Figure 3: Contour plots showing the sum of the abundances at three different ranges: 2nd peak, 3rd peak, and actinides. This is done for the mass models DZ33, FRDM, and HFB24 all at tau=0.05 s. The yellow regions indicate the highest abundances, which are particularly high on the actinide region for FRDM and HFB24.

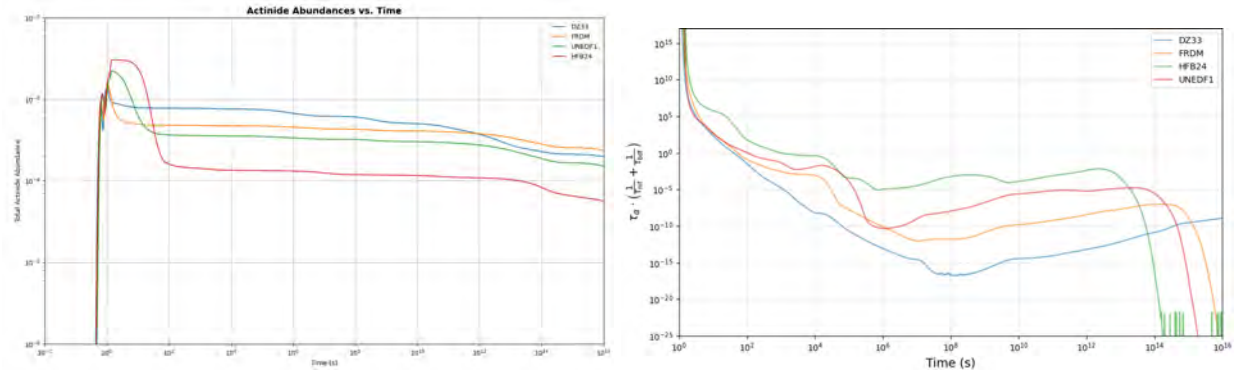

Figure 4: Plot showing the abundance of actinides over time for each of the four mass models DZ33, FRDM, UNEDF1, and HFB24 (left). As expected, DZ33 (blue) has the lowest initial actinide abundance and HFB24 (red) has the highest actinide abundance. As time goes on, however, we see HFB24 abundances drop significantly, resulting in FRDM having the greatest actinide abundance at late times. Additionally, HFB24 has a higher ratio of alpha decay versus fission rate (right).

This figure shows that the HFB24 mass model has a high initial abundance of actinides along with FRDM, and DZ33 has the lowest total abundance. When examining the results of the visualization closer at a particular point, we expect HFB24 to have a greater initial actinide abundance for an initial entropy of 51.0 and an electron fraction of 0.05. However, at later time conditions, we expect FRDM to surpass actinide production of HFB24 because the latter should demonstrate more fissioning while the former should demonstrate more alpha decay, thus leaving FRDM more material to populate the actinides.

The total actinide abundance from HFB24 does indeed begin higher than that of FRDM, then significantly decrease overtime due to experiencing more fission while FRDM undergoes more beta decay [Figure 4].

3.4 Script Directory

Throughout this summer REU program, I have created a large number of scripts that can produce dozens of unique plots and animations to visualize the r-process and its properties. These scripts are not only useful for the current work that we are completing, but can be used for other projects and investigations in the future. To make them more accessible for general group usage, I have compiled them into a directory that can be accessed at any point.

4 Conclusion

The sites, specific astrophysical conditions, and nuclear inputs that are favorable for the r-process to occur remain unconfirmed. In this study, we are using PRISM nuclear simulation data in order to create visualization tools that will solve the mystery behind the production of the heaviest elements, including gold and uranium, and elucidate r-process nucleosynthesis. We utilize key features of the r-process, including magic neutron numbers, neutron shell closures, and abundance pattern peaks, to hone in on its specific properties and patterns. From this project, I have not only produced the plots shown and described in this report, but I have also generalized the scripts associated with them and compiled them into a directory that others may access for their own use. These scripts will be beneficial for future analysis of other conditions, calculations, and metrics related to the search to understand the r-process.

References

[1] S. Rosswog, O. Korobkin, A. Arcones, F. K. Thielemann, and T. Piran, Monthly Notices of the Royal Astronomical Society **439**, 744 (2014).

[2] D. Grossman, O. Korobkin, S. Rosswof, and P. Tsvi, Monthly Notices of the Royal Astronomical Society **439**, 757 (2014).

[3] F. K. Theilemann, M. Eichler, I. V. Panov, and B. Wehmeyer, Annual Review of Nuclear and Particle Science **67**, 253 (2017).

[4] B. Wehmeyer, C. Fröhlich, and F. K. Côté, B.and Theilemann, Monthly Notices of the Royal Astronomical Society **487**, 1745 (2019).

[5] K. Farouqi, F. K. Thielemann, S. Rosswog, and K. L. Kratz, Astronomy Astrophysics **663** (2022).

[6] C. Winteler, R. Käppeli, A. Perego, A. Arcones, N. Vasset, N. Nishimura, M. Liebendörfer, and F. K. Thielemann, The Astrophysical Journal Letters **750**, L22 (2012).

[7] N. Nishimura, T. Takiwaki, and F.-K. Thielemann, The Astrophysical Journal **810**, 109 (2015).

[8] D. M. Siegel, J. Barnes, and B. D. Metzger, Nature **569**, 241 (2019).

[9] T. Fischer, M.-R. Wu, B. Wehmeyer, N.-U. F. Bastian, G. Martínez-Pinedo, and F.-K. Thielemann, The Astrophysical Journal **894**, 9 (2020).

[10] D. Yong, C. Kobayashi, G. S. Da Costa, M. S. Bessell, A. Chiti, A. Frebel, K. Lind, A. D. Mackey, T. Nordlander, M. Asplund, A. Casey, A. Marino, S. Murphy, and B. P. Schmidt, Nature **595**, 223 (2021).

[11] T. M. Sprouse, M. R. Mumpower, and R. Surman, Phys. Rev. C **104**, 015803 (2021).

[12] J. Duflo and A. Zuker, Phys. Rev. C **52**, R23 (1995).

[13] P. Moller, A. J. Sierk, T. Ichikawa, and H. Sagawa, Atomic Data and Nuclear Data Tables **109-110**, 1 (2016).

[14] S. Goriely, N. Chamel, and J. M. Pearson, Phys. Rev. C **88**, 061302 (2013).

[15] M. Kortelainen, J. McDonnell, W. Nazarewicz, P.-G. Reinhard, J. Sarich, N. Schunck, M. V. Stoitsov, and S. M. Wild, Phys. Rev. C **85**, 024304 (2012).

Refinement of ^{7}Li + α reaction data

ZACHARY HOCHSTETLER

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Richard deBoer, Prof. Wanpeng Tan, & Prof. Michael Wiescher

Abstract

The $^{7}Li + \alpha$ reaction results in many elements forming at various energy levels, and the analysis of the reaction data can reveal the nature of the elements that were formed. However, before this analysis can be done, further data refinement must be done in order to accurately obtain an accurate number of occurrences of each element. To do this refinement, I used a program that fits a Gaussian over a histogram to extract the counts of various reactions that occurred. In this paper, I will cover how the fitted Gaussian correctly extracts the correct data, and how I ensured the Gaussians were properly fitted. As a result, I extracted the reaction counts for grounded ^{7}Li, excited ^{7}Li, and grounded ^{19}F. The data extracted can then be used in conjunction with R-Matrix analysis to learn various things about each reaction, such as the cross section.

1 Introduction

During my time at the Nuclear Science Lab at Notre Dame, my main task was to look at histograms of $^{7}Li + \alpha$ particle reactions. Many runs at different energies of this same reaction were done to give as clear a picture as possible. Within each of these histograms, there were various peaks, or resonance energies, of the reaction that corresponded to various elements, but the elements on which I focused on were grounded ^{7}Li, excited ^{7}Li, and grounded ^{19}F. In these peaks, a program in ROOT automatically fit Gaussians over these peaks, and I had to go through and ensure the fits were adequate. If they were not, I changed the parameters to truly obtain an adequate fit of the resonance peak. These fits exist to obtain an accurate amount of counts of a certain reaction for a given energy in order to then use those counts for further analysis.

2 Background

Resonances in nuclear physics show the energy levels where interactions between particles are most likely to occur. These can lead to finding the cross-sections of reactions, or more simply, the likelihood of a reaction occurring at a certain energy. These cross-sections can then lead to even further analysis, such as the inner interactions between the particles within an atom and their interactions with outside factors such as other particles and reactions. In order to gain the best

understanding we can of the reaction, there are multiple runs that range from 2.525 to 5.000 MeV, and within each run, there are 22 detectors at various angles ranging from 20° to 150° to capture a wide range of reactions. All of the data for this experiment was captured by photodiode detectors. Normally, these detectors are used to detect light, but the silicon inside also allows them to detect charged particles such as the ones I worked with throughout my time.

You can see the full energy spectrum in figure 1. The various peaks on the graph correspond to the different resonance energies I discussed earlier. These resonances then correspond to certain reactions. Figure 1 also has the peaks for the specific reactions I worked with labeled for a random run.

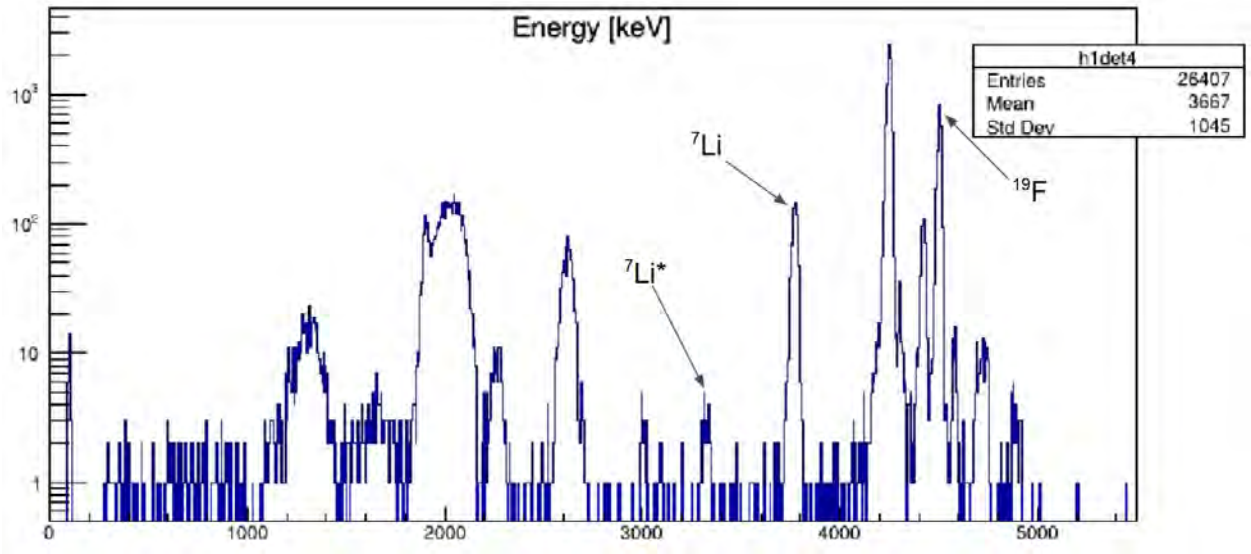

Figure 1: Full Energy spectrum with Labeled Peaks

Figures 2, 3, and 4 show a resonance peak for a run at an energy of 5.000 MeV and detector angle of 65°for excited ^{7}Li, ground ^{7}Li, and ground ^{19}F, respectively. You can see that the resonance energies between each reaction vary quite a bit. For excited ^{7}Li, the center of the peak is about 2.010 MeV, ground ^{7}Li is at about 2.415 MeV, and ground ^{19}F is about 3.895 MeV.

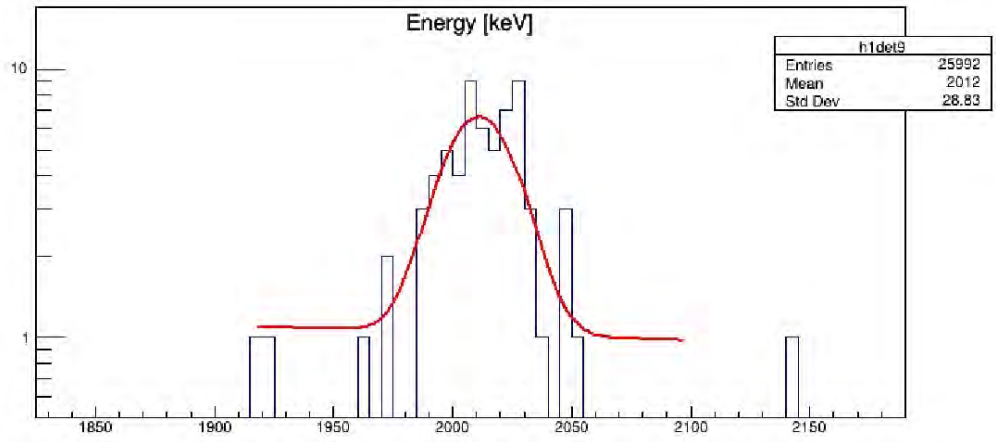

Figure 2: $^{7}Li \cdot$ Resonance Peak

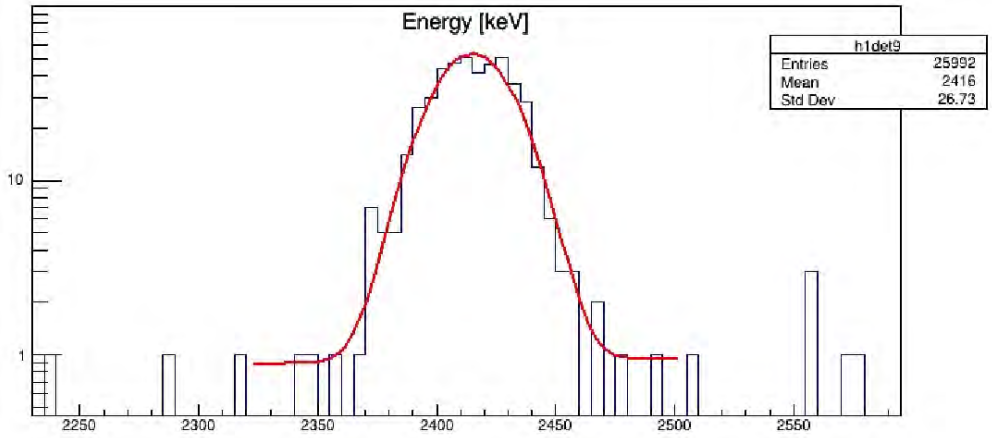

Figure 3: ^{7}Li Resonance Peak

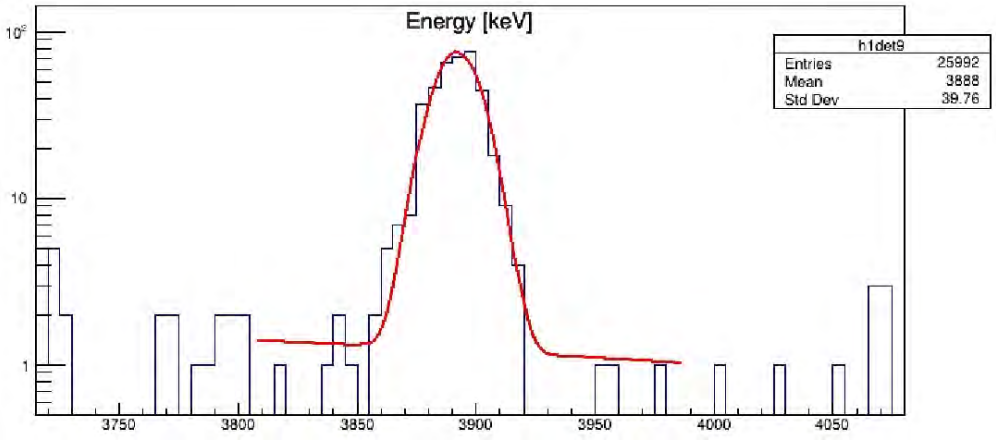

Figure 4: ^{19}F Resonance Peak

3 Methods

The vast majority of my work was done in ROOT, a program created by CERN based in C++ which was initially created for physicists to analyze large amounts of data. All of the data was already made into histograms showing the energy on the x-axis and the count of the number of reactions on the y-axis. The Gaussians were automatically fitted by a program in ROOT created by Dr. Wanpeng Tan. Sometimes, though, the program would not create a Gaussian that fits the resonance peaks as well as it should, in which case I manually entered parameters such as the end points, peak, and midpoint to adjust the Gaussian into a more proper fitting. Figure 5 is an example of what a proper fitting should look like.

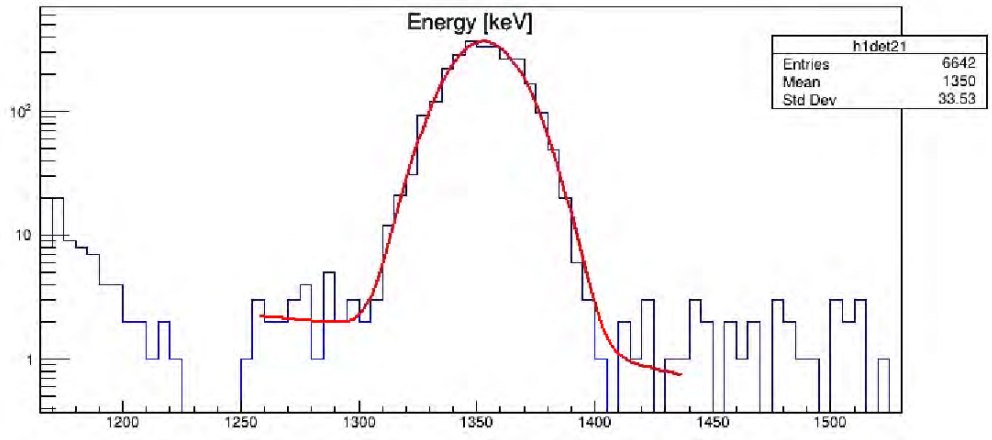

Figure 5: Proper Gaussian Fit (Ground ^{19}F)

However, sometimes there are multiple resonance peaks within a small energy range, making it nearly impossible to obtain a Gaussian that is reasonable and accurately gives a proper fit to a single resonance peak. Figure 6 shows an example of one of these instances. As you can see, multiple peaks are too close to distinguish between them, and as a result, the Gaussian created is unreliable. In this case, the data cannot be analyzed. Even adjusting the parameters as I previously discussed does not provide a Gaussian good enough to be able to take the data any further.

Now, the Gaussian curve is not actually as important as the linear fit that is created at the end-points of the Gaussian. Within the histograms, there is almost always background that is not a result of the reaction, and this needs to be accounted for. To accurately sum the counts of reactions of a

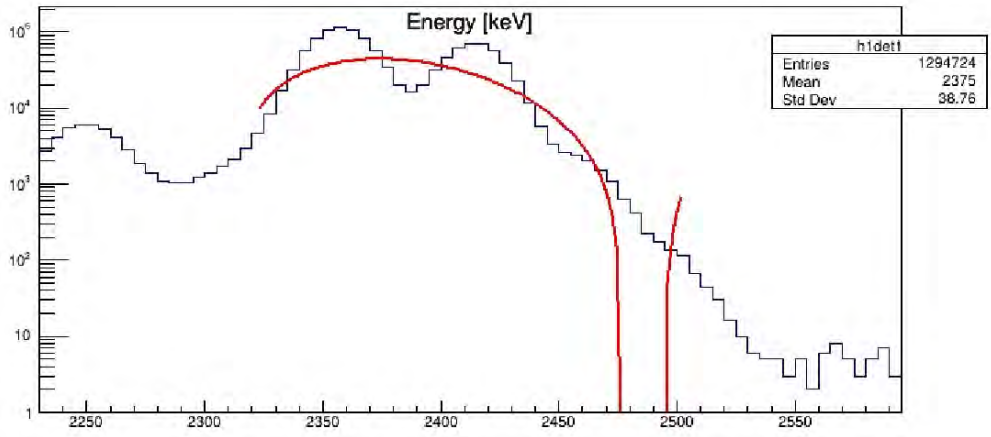

Figure 6: Resonance Peaks too Close to Analyze (Ground ^{19}F)

resonance peak, the linear fit is supposed to fit the background. From there, the sum of all counts under the Gaussian is summed, and then the counts under the linear fit are summed and subtracted from the Gaussian total counts. As a result, this effectively eliminates the background well enough where we do not have to be concerned with it in the rest of the data analysis.

4 Results

In total, I fit a total of 7,194 Gaussians over a total of 109 runs with different energies. Not all of the Gaussians were effective enough to make use of, so less than 7,194 fits will actually be used in further analysis.

5 Conclusion

Although I may not have immediate results because of the length of time the analysis of the data I refined will take, I have helped the lab's ability to speed up the analysis process. This will aid in their work, allowing more data to be obtained instead of refining the data they currently have. They can now analyze the $^7\text{Li} + \alpha$ reaction without worrying about data refinement. The counts of reactions of grounded ^{7}Li, excited ^{7}Li, and grounded ^{19}F are now available for further research and analysis. If I had slightly more time, I could participate in this, but unfortunately I do not. However,

further analysis could provide more information on the creation of heavier elements such as carbon or nitrogen in the early days of our universe, as explained here [1]. Lithium provides an important stepping stone to reach a full understanding of the creation of heavy elements, meaning most of the reactions I sifted through will be helpful in continuing to grow our knowledge. Furthermore, studying these reactions can help us understand other reactions as well. For example, a certain Beryllium decay resonance goes to $^{7}\mathrm{Li} + \alpha$, and studying the reaction discussed in the paper could help process that decay along with any others that may follow a similar pattern [2]. Overall, in conjunction with R-matrix analysis, the data extracted can be used for various purposes even outside the immediate $^{7}\mathrm{Li} + \alpha$ process.

References

[1] M. Wiescher, O. Clarkson, R. J. deBoer, and P. Denisenkov, Eur. Phys. J. A **57** (2021).

[2] Y. Ayyad, W. Mittig, T. Tang, B. Olaizola, G. Potel, N. Rijal, N. Watwood, H. Alvarez-Pol, D. Bazin, M. Caamaño, J. Chen, M. Cortesi, B. Fernández-Domínguez, S. Giraud, P. Gueye, S. Heinitz, R. Jain, B. P. Kay, E. A. Maugeri, B. Monteagudo, F. Ndayisabye, S. N. Paneru, J. Pereira, E. Rubino, C. Santamaria, D. Schumann, J. Surbrook, L. Wagner, J. C. Zamora, and V. Zelevinsky, Phys. Rev. Lett. **129** (2022).

Radial velocity as a method of exoplanet study

ISAAC HOWK

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Jeff Chilcote

Abstract

The study of exoplanets is vital in helping us to understand the universe. Studying these planets not only helps us to learn about their formation and evolution but also about the evolution of our own planet. To study exoplanets, we first need to find out how old they are. The best way to do so is to estimate the age of the planet's host star, in this case by using a technique known as radial velocity (RV). When applied to binary systems in moving groups, RV allows us to calculate the orbit, and thus the masses, of the system. Once we know the mass of a star, finding its age is much easier. There are two types of RV, absolute and relative, but for determining orbits relative is more convenient, and so that is what this paper will focus on. Unfortunately, no relative RVs have been produced yet due to a bug in the code being used. The final output of this experiment will still be relative, but those results have not been reached yet, and so in their place we will analyze the absolute RVs for the same observations. The hope is that by using the absolute values as a substitute, we can determine how helpful the final result will be in calculating orbits. This is difficult to determine without doing further calculations, but just looking at the plots of these RVs, many seem to be promising. Not only will the relative RV code need to be fixed, but several of the observed objects had far too few observations to draw any sorts of conclusions; gathering more data will greatly improve future results.

1 Introduction

Determining the age of a star is difficult, but hugely valuable for the study of exoplanets. Using the age of a planet's host star, we can extrapolate the age of that planet, which gives us valuable information about its properties. [1] While calculating the age of a star is difficult, our understanding of stellar evolution is advanced enough that we can reach an answer using only the star's mass and luminosity. Finding a star's luminosity is generally fairly simple, but mass is more difficult. Many of our methods for calculating a star's mass require specific conditions to be present, such as the star being part of a binary system. Radial velocity, or RV, a strategy that analyzes the Doppler shift of received spectra in order to calculate orbits, is one of such methods. [2] RV is only applicable for systems with more than one star, and has limitations when applied to systems that are not oriented optimally. Despite the issues with this method, it allows for the calculation of the age of a star, which, when applied to a star within a moving group, can give an estimate for the age of many more stars. [2]

2 Background

2.1 Radial Velocity and Orbits

Radial velocity (RV) is a tactic that can be used to calculate the orbits of binary star systems using the Doppler effect. [3] As a star orbits around its partner, it also moves relative to the Earth, meaning that the Doppler effect applies to the light we receive from the system. This results in a sinusoidal oscillation of the star's light spectra which can be analyzed to deduce information about the system. Although we can determine a lot using RV, other analysis methods are needed to provide a full picture. Motion perpendicular to our line of sight from the Earth does not cause a Doppler shift, and since very few systems are oriented perfectly parallel to this line of sight, we need another analysis method to tell us about the motion radial velocity can not see. This other type of analysis is direct imaging, which consists of directly observing the system and analyzing its motion with the pictures taken. [4] This method has the exact opposite problem as radial velocity, it can only see motion perpendicular to our line of sight, not parallel, and so by combining the two methods we can compensate for each of their weaknesses and effectively study the orbits of these systems.

Radial velocity is calculated using the following equation:

$$\upsilon = \Delta\lambda c / \lambda \tag{1}$$

where $\Delta\lambda$ is the shift in wavelength from the template to the current observation, λ is the wavelength of the template file, υ is the radial velocity, and c is the speed of light. There are two types of radial velocity: absolute RV, which measures the motion of a star relative to the earth, and relative RV, which measures the motion of a star relative to a previous observation of that star. Both approaches have their pros and cons, but this paper will focus on relative RV as it is the more convenient method for studying orbits.

2.2 Moving Groups

Moving groups are groups of stars that move together with similar velocities and are not bound together by gravity. Due to their matching velocities and lack of gravitational cohesion, moving groups can be assumed to have a uniform age, likely sharing a common origin. By focusing on binary star systems in moving groups, we can estimate the age of the entire group. [2] In this way we can determine the ages of many stars we otherwise would not be able to, making moving groups a powerful tool for astrophysicists.

2.3 exoplanets

One of the most appealing reasons for computing the age of stars is to better understand the exoplanets they host. The development of these planets is closely linked to their host stars, and so if we know the age of their star, we can estimate the age of the planet. With this quantity, we can improve our models of planetary evolution and better understand their formation, as well as eliminating objects that are not quite planets. [1] It can be hard to tell the difference between a planet and other similar objects, such as brown dwarfs. However, we do have a good understanding of how their evolution differs from that of planets. By combining the ages and luminosities of observed objects, we can separate planets from non-planets, isolating only what we want to focus on.

3 Methods

Objects were observed over the course of four years by the Las Cumbres Observatory. Before calculating any radial velocities, a template observation to serve as a baseline for the object's relative motion had to be selected. It functionally did not matter which was chosen, so for each object the observation with the highest signal-to-noise ratio (SNR), a value included in the header of each file, was chosen. After selecting the template, the RV for each other observation of that object was calculated, and the set of RVs was plotted against time. Before beginning to calculate the RV, an initial guess had to be made. To do so, an evenly spaced array of numbers was generated, centering

around the difference between the barycentric correction of the template file and that of the current file. Next, a chi2 function was run on each entry in the array using equation 1 to find the shifted template wavelength for the input RV. The function then interpolated the flux values of the template to extend to this shifted wavelength, and subtracted the flux of the current file at that point. Finally, it output the square of this difference. To finalize the initial guess, the minimum of these outputs was selected and its corresponding RV was returned, providing a starting guess for what the final result of the program should be.

Now that an initial guess was calculated, a chi2 minimizer was run on a set of 52 separate wavelength ranges with respect to the initial guess to determine the RV's for those wavelength ranges. Barycentric corrections was then applied to each of those values, followed by a five-standard-deviation sigma clipping preformed six times on the entire array of values. The error for this array was found by dividing its standard deviation by the square root of its length. Next, the sky correction was tackled. First an initial guess was made as with the RVs, but this time the guess was centered around 0 instead of around the barycentric difference. The initial guess was taken and again input into a chi2 minimizer, but the number of wavelength ranges this was preformed on was much smaller than before: only 16 entries. This array was then sigma clipped, and the median value was added to the median value of the RVs to calculate the final RV. The error for the sky corrections was calculated in the same way as for the RVs, by dividing its standard deviation by the square root of its length. Finally these two errors were combined by taking the square root of the sum of their squares, and used as the final error for the output RV.

4 Results

Unfortunately, due to a bug in the code used to calculate relative radial velocities, no real results have been found yet. In the absence of final results, the absolute radial velocities of these observations will be analyzed. This is not a perfect substitute, but the overall pattern should be similar. Looking at the absolute RV will indicate how viable relative RVs will be for calculating of orbit of each object, and also whether there is enough data to draw a conclusion about that object. There was a

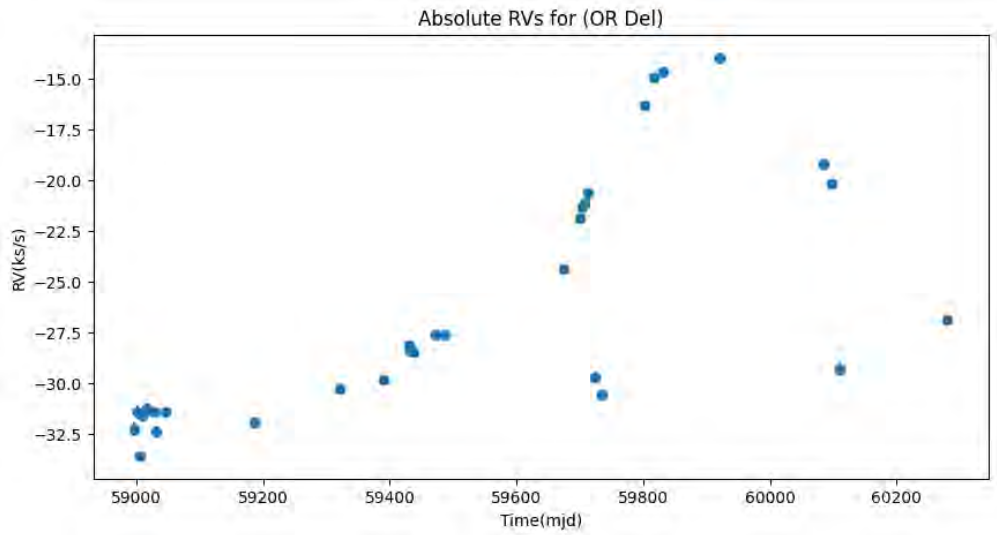

Figure 1: Absolute RVs vs time for the start OR Del

variety of results from these plots, ranging from objects with a clear RV oscillation to others that simply had too few data points to draw much of a conclusion. It will not be clear how useful this process has been until the relative RVs have been calculated and used to find the orbit of the system, but several of the absolute RV plots having recognizable oscillation patterns is a good sign.

One thing that raises questions on the usefulness of examining absolute RVs is the error given for these values. Both the absolute RVs and their errors were included with the observation files, but these errors are completely unrealistic. While absolute RVs generally ranged from -200 to 200 km/s, the error only ranges from about 0.1 to 0.8 km/s. Not only does such a small percent error draw concern, but there are very clearly outliers, as seen in fig. 1. Such small error bars, barely visible in the figure, indicate that something is wrong with the error calculations, but more importantly, they sow doubt about the RVs themselves. The error of these points does not actually matter for our purposes, as we will be using an entirely different set of data for further work, but the values being so clearly wrong raises concerns over whether the RVs are accurate.

5 Conclusion

Although no results have been produced yet, analyzing absolute RVs gives a good idea of how viable this strategy will be. Many systems showed visible patterns, and many of the ones that did not will still be useful for the next step of calculating the orbits of these systems. While the unrealistic errors of the absolute RVs raises concerns over whether this assessment is trustworthy, it is much more likely to be due to a problem with the error calculations than an indication of the absolute RVs being wrong. Future steps to be taken include fixing the relative RV code to generate real outputs, using those results to calculate orbits for these systems, and of course gathering more data. There is still much to be done, but preliminary signs indicate that the data used will be sufficient to gather significant information about the orbits of these binary systems.

References

[1] A. D. Thomas, E. L. Nielsen, R. J. De Rosa, A. E. Peck, B. Macintosh, J. Chilcote, P. Kalas, J. J. Wang, S. Blunt, A. Greenbaum, Q. M. Konopacky, M. J. Ireland, P. Tuthill, K. Ward-Duong, L. A. Hirsch, I. Czekala, F. Marchis, C. Marois, M. A. Millar-Blanchaer, W. Roberson, A. Smith, H. Gallamore, and J. Klusmeyer, The Astronomical Journal **166**, 246 (2023).

[2] E. L. Nielsen, R. J. De Rosa, J. Wang, J. Rameau, I. Song, J. R. Graham, B. Macintosh, M. Ammons, V. P. Bailey, T. S. Barman, J. Bulger, J. K. Chilcote, T. Cotten, R. Doyon, G. Duchêne, M. P. Fitzgerald, K. B. Follette, A. Z. Greenbaum, P. Hibon, L.-W. Hung, P. Ingraham, P. Kalas, Q. M. Konopacky, J. E. Larkin, J. Maire, F. Marchis, M. S. Marley, C. Marois, S. Metchev, M. A. Millar-Blanchaer, R. Oppenheimer, D. W. Palmer, J. Patience, M. D. Perrin, L. A. Poyneer, L. Pueyo, A. Rajan, F. T. Rantakyrö, D. Savransky, A. C. Schneider, A. Sivaramakrishnan, R. Soummer, S. Thomas, J. K. Wallace, K. Ward-Duong, S. J. Wiktorowicz, and S. G. Wolff, The Astronomical Journal **152**, 175 (2016).

[3] N. C. Hara and E. B. Ford, Annual Review of Statistics and Its Application **10**, 623 (2022).

[4] K. B. Follette, Publications of the Astronomical Society of the Pacific **135**, 093001 (2023).

Reconstruction of Merged Diphoton Signatures in Higgs Decays Using ML and ECAL Clustering in CMS

SEJLA KULDIJA

2025 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Marc Osherson

Abstract

In several extended Higgs sector models, light scalars (ϕ) decay into photon pairs ($\phi \to \gamma\gamma$). When ϕ is produced with high momentum, the resulting photons are very energetic and often merge in the CMS electromagnetic calorimeter (ECAL), complicating standard reconstruction. This study uses CMS Run II simulated samples to evaluate two reconstruction strategies for such merged diphoton signatures. A machine-learning–based photon object (MLPhoton), tested on Bulk Radion (Bkk $\to \gamma\phi$) events, recovers a significant fraction of signal events missed by standard diphoton reconstruction at low ϕ/Bkk mass ratios. A complementary analysis of $A \to \gamma\gamma$ decays Reconstructed clusters are matched to generator-level photons, which represent the original photons produced in the simulated decay, using angular separation (ΔR). Efficiency is studied as a function of photon boost (γ). Results show that single-cluster matches outperform two-photon methods at high boost, and spatial maps reveal inefficiencies near the ECAL center. These findings support the use of nonstandard reconstruction techniques to improve sensitivity in searches for new physics with collimated photon final states.

1 Introduction

The Compact Muon Solenoid (CMS)[1] is a major experiment at the Large Hadron Collider (LHC)[2], operated by CERN[3]. It studies high-energy proton collisions, which can produce a range of particles, including rare or short-lived ones such as the Higgs boson. The detector composed of concentric layers, each responsible for tracking charged particles, energy measurement, or muon identification.

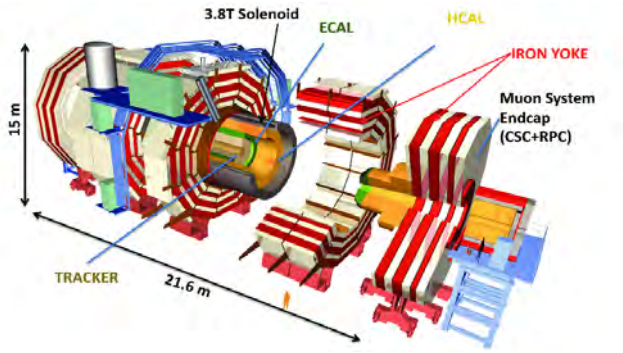

Figure 1: Cutaway view of the CMS detector.

The Electromagnetic Calorimeter (ECAL) [1] is the CMS component responsible for detecting

[1] https://cms.cern
[2] https://home.cern/science/accelerators/large-hadron-collider
[3] https://home.cern

electrons and photons. It is built from lead tungstate crystals which emit light when struck by high-energy electrons or photons. This light is converted into an electrical signal to measure the particle's energy, while the crystal's location provides spatial information. The ECAL includes a central barrel and two endcaps 2, with a preshower detector in front of the endcaps to help distinguish single photons from overlapping pairs.

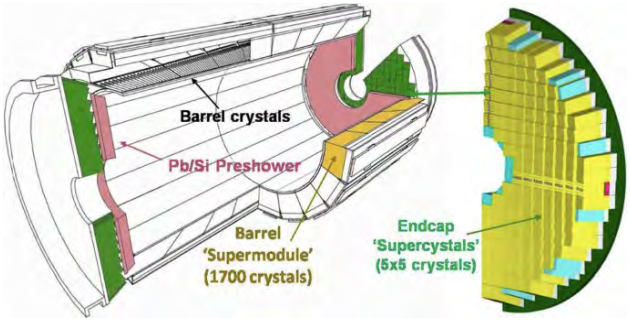

Figure 2: ECAL consists of a barrel and two endcaps.

When a high-energy photon strikes the ECAL, it initiates an electromagnetic shower that deposits energy across multiple crystals. Under standard conditions, isolated photons produce compact, symmetric showers, allowing CMS reconstruction algorithms to accurately measure their energy and direction. Figure 6(b) shows an example of such a resolved diphoton event, where both photons are individually reconstructed. However, many beyond-the-Standard-Model theories predict spin-0 particles (ϕ) that decay into two photons [2]. If the ϕ particle is produced with high momentum, for example from a heavier resonance, the two photons can be emitted in nearly the same direction (collimated), as illustrated in Figure 6(a). Their overlapping showers may then appear as a single energy deposit in the ECAL (Figure 6(c)).

This project aims to improve the reconstruction of such merged diphoton signatures using alternative techniques based on machine learning and ECAL energy clustering.

2 Background and Motivation

CMS uses a right-handed coordinate system with the origin at the collision point: the x axis points toward the LHC center, the y points upward, and the z follows the beam direction. The azimuthal

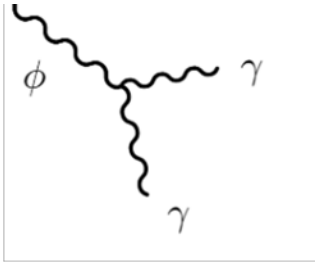

Figure 3: (a) Feynman diagram

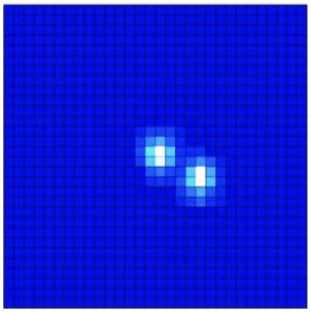

Figure 4: (b)

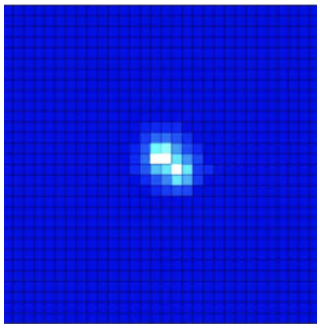

Figure 5: (c)

Figure 6: (a) ϕ decays to two photons. (b) Two well-separated ECAL clusters. (c) Overlapping clusters.

angle ϕ is measured in the transverse (x–y) plane from x axis , and the polar angle θ is defined from the z axis. Pseudorapidity is given by $\eta = -\ln(\tan(\theta/2))$ [1].

This study improves photon reconstruction in events where boosted resonances produce collimated photon pairs that merge in the ECAL. I evaluated two approaches using simulated proton-proton collisions at $\sqrt{s} = 13$ TeV from CMS Run II. The first uses a machine learning–based object, the **MLPhoton**, trained to identify merged photons using ECAL shower shapes and energy patterns [3]. The second uses raw **ECAL clusters**, which represent the local energy deposits recorded in the calorimeter crystals before any higher-level reconstruction is applied. To assess performance, both methods are matched to **generator-level photons**, which represent the "true" photons in the simulation. Matching is done using the angular distance $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. For comparison, standard CMS-reconstructed photons, **RECO photons** or **PAT photons**, are included as a baseline.

I used two signal models to analyze different merging regimes. The first is a bulk radion model, where a heavy resonance decays as $Bkk \to \phi\gamma$, followed by $\phi \to \gamma\gamma$, producing three photons. In this case, the ϕ often carries high momentum, causing its two daughter photons to merge. The second model is a scalar resonance, $A \to \gamma\gamma$, where the photons are typically less boosted and more easily resolved.

The analysis examines how reconstruction efficiency depends on the parent particle's boost, the ϕ-to-Bkk mass ratio, the photon separation ΔR, and the ECAL detector region (barrel vs. endcap).

3 Methods

I performed the analysis using simulated CMS Run II proton-proton collisions. Reconstructed photons were matched to generator-level photons using the angular distance metric $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$, with thresholds scaled to photon boost and supplemented by energy and mass consistency checks to reduce mismatches.

Three reconstructed object types were compared: standard CMS photons (PATPhotons), machine learning–based **MLPhotons** selected by diphoton score, and raw **ECAL clusters**, which reflect unprocessed crystal energy deposits. Generator-level photons served as the truth reference. Matching was done event-by-event, with reconstruction success defined by the number of matched photons, mass window agreement, and energy conservation relative to the parent.

I evaluated efficiency as a function of photon boost (γ), corresponding to the mass ratio of daughter to parent particles. ECAL performance was studied separately in the barrel and endcap regions to account for geometry. Two-dimensional heatmaps were also generated to show the spatial distribution of matched and unmatched clusters in ECAL coordinates.

All selection, matching, and efficiency calculations were implemented in Python using NumPy and `TLorentzVector` for kinematics. ROOT was used for histogramming and visualizing efficiency and failure modes.

4 Reconstructing Merged Photons in $Bkk \to \gamma\phi$ Events

To evaluate reconstruction performance in scenarios with overlapping electromagnetic showers, I analyzed a simulated process in which a heavy radion-like resonance (Bkk) with mass $m_{Bkk} =$ 900 GeV decays to a photon and a light scalar ϕ, followed by $\phi \to \gamma\gamma$. When ϕ is light, it receives a large boost from the Bkk decay, causing the decay photons to become highly collimated and their ECAL showers to overlap.

Reconstruction efficiency was studied as a function of the mass ratio m_ϕ/m_{Bkk}, which controls the boost and photon separation. Events were selected to contain exactly one generator-level ϕ

(PDG ID 9000025) decaying to two photons (PDG ID 22). The angular distance ΔR between the photons was computed to quantify their collimation.

Two reconstruction strategies were tested. The first uses standard CMS **Reco Photons**, matched to generator-level photons using a ΔR–based criterion. The invariant mass of the matched pair was required to fall within $\pm 20\%$ of the nominal ϕ mass to count as a successful reconstruction.

The second approach uses **MLPhotons**, machine learning–based objects trained to identify merged diphoton showers. MLPhotons with a diphoton score above 0.7 were selected, and a mass window of $\pm 2\sigma$ around the MLPhoton mass peak (derived from the sample) was applied. Events were considered matched if at least one MLPhoton satisfied these criteria.

Figure 7 shows the reconstruction efficiency for both methods as a function of the ϕ/Bkk mass ratio. At high mass ratios ($\gtrsim 0.02$), where the decay photons are well separated, standard diphoton reconstruction exceeds 80% efficiency. However, its performance drops sharply below 0.01 due to merging. In contrast, the MLPhoton method remains effective at low mass ratios, recovering up to 60% of signal events in regions where standard reconstruction fails. Its efficiency declines at higher mass ratios where merging no longer occurs.

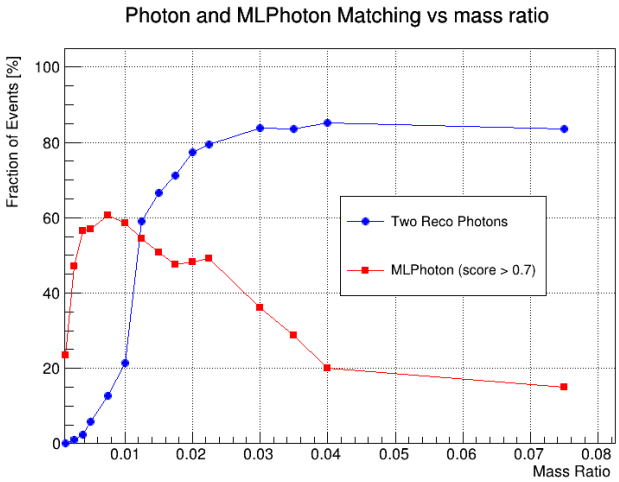

Figure 7: Reconstruction efficiency in $Bkk \to \gamma\phi$ events as a function of the ϕ/Bkk mass ratio.

5 Cluster-Based Photon Reconstruction in A → $\gamma\gamma$ Events

To study reconstruction in collimated diphoton topologies, I analyzed simulated scalar decays of the form $A \to \gamma\gamma$. When the parent scalar A is produced with high boost, the decay photons can become highly collimated, depositing overlapping energy in the ECAL.

I compared two reconstruction approaches: the standard CMS method using PATPhotons, and a clustering-based method using raw ECAL crystal energy deposits. Events with a single generator-level scalar A boson were selected. The decay photons were matched to reconstructed objects using angular distance ΔR.

In the PATPhoton method, both photons were required to match distinct PATPhoton objects, and their invariant mass had to lie within $\pm 20\%$ of the true A mass. For the cluster-based method, region-specific ΔR and p_T cuts were used to match clusters to gen-level photons. When only one cluster was matched, additional energy-sharing criteria were applied to identify merged photon signatures.

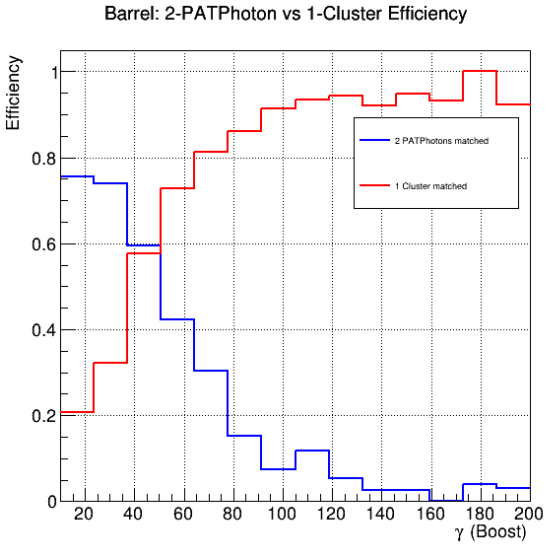

Figure 8: (a) ECAL Barrel

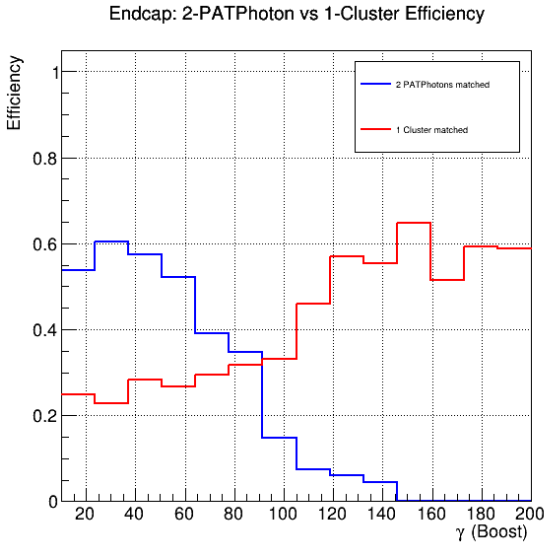

Figure 9: (b) ECAL Endcap

Figure 10: Reconstruction efficiency using two PATPhotons vs. one ECAL cluster as a function of parent boost γ, shown separately in the ECAL barrel (left) and endcap (right).

Figure 10 shows reconstruction efficiency in the barrel and endcap regions. The PATPhoton method performs well at low boost ($\gamma \lesssim 60$ in the barrel, $\lesssim 80$ in the endcap), where photons are well-separated. However, its efficiency drops sharply at high boost due to merging. In contrast, the cluster-based method maintains higher efficiency in the high-boost regime, where a single cluster captures both photons. A representative high-boost event is shown in Figure 11. The gen-level

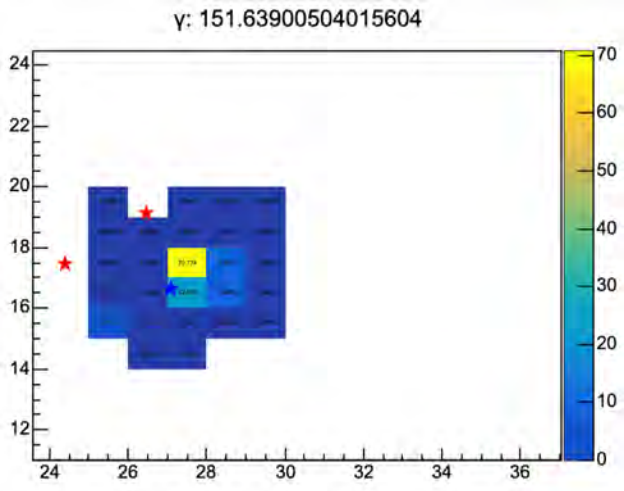

Figure 11: Crystal-level ECAL energy deposits with parent boost $\gamma = 151.6$ and total energy $E_H = 156.5$ GeV. The two gen-level photons have energies of 15.5 and 141.0 GeV. $\Delta R = 0.060$. Red stars show the photon directions; the blue star is the reconstructed cluster position.

photons, separated by only $\Delta R = 0.060$, deposit energy in adjacent crystals. Despite the asymmetry in photon energies (15.5 GeV and 141.0 GeV), the merged ECAL cluster (blue star) successfully reconstructs the combined energy and location.

To examine spatial behavior, I mapped the (x, y) crystal positions of matched and unmatched ECAL clusters in the endcap. Figure 14 shows normalized heatmaps for both cases. Matched clusters are broadly distributed, while unmatched clusters concentrate near the ECAL center, where a gap (the "hole") leads to reduced coverage.

To further confirm this behavior, Figure 15 shows the fraction of matched vs. unmatched clusters as a function of radial distance from the hole center. Matching efficiency increases with radius, exceeding 90% at large distances. Unmatched clusters dominate at small radii, confirming that local geometry significantly impacts reconstruction performance.

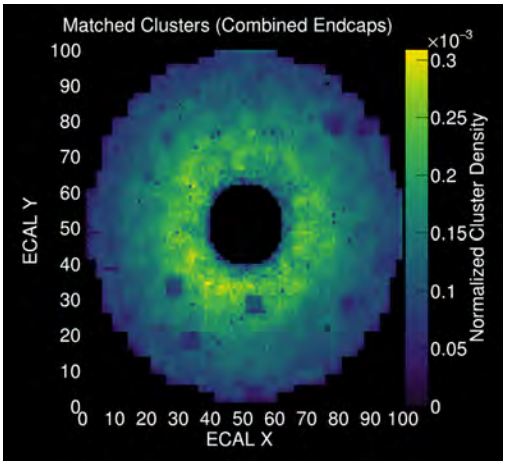

Figure 12: (a) Matched clusters

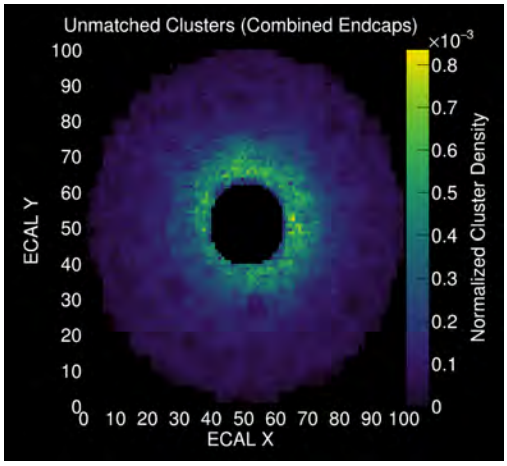

Figure 13: (b) Unmatched clusters

Figure 14: Heatmaps of matched (left) and unmatched (right) ECAL clusters in the endcap regions. The dark center represents the "hole" in the detector.

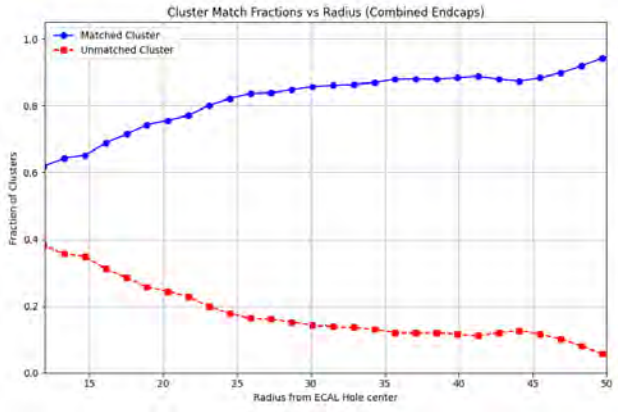

Figure 15: Fraction of matched and unmatched ECAL clusters in the endcap as a function of radial distance.

6 Conclusion

This study investigated photon reconstruction in scalar decays of the form $\phi \to \gamma\gamma$, focusing on high-boost scenarios where the two photons become collimated. Using simulated events, I compared standard CMS reconstruction based on PATPhotons with two alternative approaches: a machine learning–based diphoton tagger and an ECAL cluster-based method.

While PATPhotons perform well at low boost, their efficiency drops as photon separation decreases. In contrast, the ML- and cluster-based methods show improved performance in the high-

boost regime, recovering events that standard reconstruction misses. These gains are particularly notable in the ECAL endcaps, where geometric effects such as the detector hole reduce efficiency.

In general, the results demonstrate that alternative reconstruction strategies can enhance sensitivity to new physics involving merged photons and motivate further development of such techniques within CMS.

References

[1] CMS Collaboration. "Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC". In: *JINST* 16.05 (2021), P05014. DOI: 10.1088/1748-0221/16/05/P05014. URL: https://doi.org/10.1088/1748-0221/16/05/P05014.

[2] CMS Collaboration. "Search for new particles in an extended Higgs sector with four b quarks in the final state at $\sqrt{s} = 13$ TeV". In: *Phys. Lett. B* 835 (2022), p. 137566. DOI: 10.1016/j.physletb.2022.137566. URL: https://doi.org/10.1016/j.physletb.2022.137566.

[3] CMS Collaboration. "Search for new resonances decaying to pairs of merged diphotons in proton-proton collisions at $\sqrt{s} = 13$ TeV". In: *Phys. Rev. Lett.* 134 (2025), p. 041801. DOI: 10.1103/PhysRevLett.134.041801. URL: https://doi.org/10.1103/PhysRevLett.134.041801.