

2026 NSF/REU
RESEARCH EXPERIENCE FOR UNDERGRADUATES
IN
PHYSICS AND ASTRONOMY



STUDENT RESEARCH PAPERS
VOLUME 37



**UNIVERSITY OF
NOTRE DAME**
College of Science

REU DIRECTOR
UMESH GARG, PH.D.

REU Student Research Papers – Summer 2026
University of Notre Dame – Department of Physics & Astronomy

Student	Research Title	Page No.
James Burke University of Notre Dame	Simulating Ion Capture and Release for the St. Benedict Paul Trap	1
Mary Kate Carroll University of Notre Dame	Synthesis and Characterization of 1:1:3 Compounds	11
Yunya Chen Xi'an Jiaotong University	Angular Momentum Structure of Even-Mass Oxygen Isotopes from <i>Ab Initio</i> Calculations	21
Akira Clark Indiana University South Bend	Commissioning and Characterization of Micro Channel Plate Detectors for Use in St. Benedict and BEARTrap	31
Sabina Day University of Arkansas - Fayetteville	Construction and Beam Simulation Analysis for DUNE	41
Maggie Dwyer Georgetown University	Single Crystal Growth and Magnetotransport Studies of Ce ₂ PdGe ₆	51
Vanshika Gupta Amherst College	Search for Nonresonant HH Production in the <i>bbWW</i> Single-Lepton Resolved Channel, and Extended ggHH Signal Sample Production	61
James Harris University of Notre Dame	Investigating the Electronic Properties of Cr _{0.5} Ga _{0.5} Te ₂	71
Tinghao Hu Xi'an Jiaotong University	Halogen-Based Selective Etching of MAX Phases: Routes to Halogen-Terminated MXenes	81
Coren Hughes University College Cork	Shielding Optimisation for the Neutron Irradiation Station	91
Elizabeth Jacoby University of Dallas	Recharacterization of the High Pressure Point-like Jet Gas Target (HIPPO)	101
Savannah Jones Randolph Macon College	The Chemical Abundances of The Hercules Ultrafaint Dwarf Galaxy	111
Isabel Lockhart DePaul University	Characterizing the Exoplanet Transit Capabilities of the Sarah L. Krizmanich Telescope	121
Aldir Moreira Embry-Riddle Aeronautical University	Observations and Phase Space Constraints of Leptonic Decays Using CMS Scouting Data	131
Ellie Stonecipher College of Saint Benedict	The ¹³ C(α , n) ¹⁶ O Reaction Cross Section as a Function of Beam Energy and Angle	141
Aaron Sullivan Eastern Washington University	Ironing Out the Details of a Novel DeFect in FeSe	151
Ruthie Tan Brown University	Mapping Binary Star Orbits Using Direct Imaging Data from Keck Observatory and ESO	161
Gabrielle L. Williams Jackson State University	Efficiency Computation of HPGe for ¹⁵² Eu samples	171
Felicity Wunsch University of Notre Dame	Ultra-Precise Differential Abundances in Globular Clusters	181
Noah Wyatt University of Tennessee - Chattanooga	Direct Studies of 610 and 559 keV Resonances in ³⁵ Cl(p, α) ³² S	191
Hanqing Zhao Xi'an Jiaotong University	Simulation-Based Inference of Population Size and Epistasis from Evolutionary Hysteresis	197

Simulating Ion Capture and Release for the St. Benedict Paul Trap

JAMES BURKE

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR: Prof. Maxime Brodeur

Abstract

The Superaligned Transition Beta-Neutrino Decay Ion Coincidence Trap (St. Benedict) at the Nuclear Science Laboratory of the University of Notre Dame is used to test the Standard Model via measuring the beta decay of certain nuclei. In order to best study this radioactive decay, ions must be prepared and injected in an optimal, low-energy state in the measurement Paul trap, which can confine ions in free space for extended durations of time. The purpose of this research was to simulate, collect and analyze data on how charged particles are expected to behave in the Paul trap. Depending on the voltages applied to different parts of the trap, the ions are captured and confined more or less optimally. These Paul trap simulations provide insights on how these values should be adjusted to maximize the capture, holding, and extraction efficiencies.

1 Introduction

The Standard Model is as close as physics has come to a complete picture of fundamental particles and interactions. However, it is currently incomplete and its shortcomings warrant experimental verification to ensure its accuracy. One critical aspect of the Standard Model is the unitarity of the Cabibbo-Kobayashi-Maskawa (CKM) quark mixing matrix, where non-unitarity would signal the presence of physics beyond the Standard Model [1]. The Superaligned Transition Beta-Neutrino Decay Ion Coincidence Trap (St. Benedict) seeks to scrutinize the validity of the CKM quark mixing matrix, particularly through the precise measurement of the kinematics of radioactive ions as they undergo beta decay, which results in the emission of a beta particle and a neutrino. To obtain the ions in a state conducive to study, a radioactive ion beam must be generated, prepared and stopped.

When ions enter St. Benedict, they are first decelerated and thermalized via collisions with helium atoms in the gas catcher; next, the extracted ions are transported through a region with a large pressure differential using first an RF carpet and then an RFQ ion guide. Then, the ions enter the RFQ cooler-buncher where they cool via collision with helium atoms followed by

their accumulation and release as an ion bunch. The final component of St. Benedict is a Paul trap, which manipulates the ions into a low-energy state (fractions of an electron volt) and traps groups of ions. The positron and daughter nuclei produced from the beta decay are then recorded using detectors surrounding the trap [2]. The Paul trap has many sub-components and is highly configurable; however, this means that significant work must be done to simulate ion behavior in the Paul trap to ensure optimal data collection. Through the usage of SIMION, an ion simulator, ideal parameters were found to be used as a jump-off point once the Paul trap is ready for commissioning [3].

We chose ions with a mass of 39 amu and charge of +1 because an off-line ion source of $^{39}\text{K}^+$ ions will be used for the commissioning. However, the simulations are easily applicable for other masses and are more dependent upon other initial conditions; for instance, the standard deviation for the ions' initial distribution along the three principal axes. These initial conditions are largely derived from results obtained from simulations of upstream elements of the St. Benedict beamline.

2 Methods

Examining the Paul trap through the usage of SIMION allows for the easy study of various parameters which impact the nature of the beam. Of these, the four most important are the entrance tube, entrance rod, exit tube, and exit rod potentials. The purpose of the entrance and exit tubes are to decelerate and accelerate the ions into and out of the Paul trap, respectively. Figure 1 depicts the way these components are laid out. While the entrance/exit tubes and center rod are kept at constant voltages, the entrance and exit rods are capable of switching “up” and

“down” between different set voltages. For instance, the entrance rod might begin at -57V, switch up to 10V, and then switch back down to -57V as needed; the incoming positive ions would be drawn to the entrance by the negative voltage but then trapped in a potential well once inside as the voltage switches up. Once the ions are released, the voltage would drop again to receive more incoming ions. Figure 2 illustrates this process. After the ions are released, they collide with a microchannel plate (MCP) further downstream. This is where final measurements are taken.




Figure 1: The interior of the Paul trap, as represented in SIMION. Ions enter from the bottom-left side of the figure, are trapped in the center of the Paul trap, and subsequently released to the top right.




Figure 2: A graphical representation of ion trapping principles. The ion enters the well in the pre-trapping stage. In the trapping stage, the entrance rod voltage is switched up, trapping it in an oscillating state. Once the exit rod voltage is switched down, the ion is free to fly out the exit of the trap.

On top of these constant voltage levels, the entrance, exit, and center rods all possess an additional RF voltage of amplitude 450V. The switch-up values for the entrance and exit rods were set to 10V, and the center rods were kept at -45V constantly to minimize effects on trajectories of daughter ions after decay (as informed by other simulations). This left four main values to find: the switch-down values for the entrance and exit rods, as well as the constant voltage levels to keep the entrance and exit tubes at.

	Capture Stage Voltage (V)	Trapping Stage Voltage (V)	Release Stage Voltage (V)
Entrance Tube	?	Same as capture stage	Same as capture stage
Entrance Rods	?	10	10
Center Rods	-45	-45	-45
Exit Rods	10	10	?
Exit Tube	?	Same as capture stage	Same as capture stage

Because the effects of any electrode fall rapidly with distance, the distance between the entrance components and the exit tube is such that the effect of the exit tube on the behavior of ions in the first half of the trap is negligible, and likewise for the entrance tube and the second half of the trap. This is quite fortuitous, because it means only two independent variables must be manipulated at a time. Furthermore, while SIMION is a reliable tool, it struggles with accurately simulating the ion cooling which occurs via collisions with background gas while the ion is trapped that dampens its oscillatory path over time. For these reasons the simulation of the entire system (from ion entrance into the trap to collision with the MCP) can and should be avoided. Instead we split it into the trapping and release sections, which manipulate the entrance rods/tube and exit rods/tube respectively.

There are a few additional notes to be made regarding initial conditions in the trap. Because ions may be at any point in their oscillation when they are released from the trap, they were released after a random period of time between 0 and 8 microseconds (roughly 0-2 periods of oscillation). Typically, 50 ions were run per data point so that any individual anomalous behavior due to RF or simulated gas collisions would not affect the data too drastically. The trap itself was assumed to be at 10^{-5} Torr and entering ions were assumed to be moving at 720 eV, as informed by prior simulations. Trapped ions were assumed to begin at 0.5 eV at the center of the trap. Additional simulations (tested values: 0.1 eV and 1 eV) showed that for the small energies expected of trapped ions, this initial kinetic energy did not significantly impact the behavior of the ions upon release. Lastly, ion initial positions were governed by Gaussian distributions along the trap's x, y, and z axes, with a standard deviation of 0.4 mm each unless specified to be otherwise. This value is reasonable considering the results of the capture simulations, erring on the high end for good measure.

3 Results

3.1 Capture Simulations

Initially, a multidimensional contour plot of ion capture was produced (see figure 3) to determine the exact voltage combination which yields the lowest kinetic energy for mid-trap ions without sacrificing beam efficiency. A primary region of interest was identified, roughly around (Entrance Tube Voltage, Entrance Rod Voltage) = (-750V, -100V), of note because beam efficiency is high, ion kinetic energy is low, and the entrance rods do not need to make a large jump. After this, several two-dimensional scans were performed for various entrance tube voltages, scanning the relationship between the entrance rod voltage and the ions' kinetic energy

once they had reached the center of the trap (see figure 4). It was concluded that the best combination of voltages was $-775(5)\text{V}$ and $-59(1)\text{V}$ for the entrance tube and rods respectively.




Figure 3: General results for how entrance tube voltage and entrance rod voltage affect beam efficiency and ion kinetic energy at the center of the trap. The blue dot pinpoints the primary region of interest.




Figure 4: an optimal voltage combination (left) and a suboptimal voltage combination (right). Any ion that is represented as an x rather than a dot is one that failed to be trapped successfully.

3.2 Exit Simulations

A similar process was performed for exit simulations. Figure 5 exhibits a large range of acceptable combinations, so two points were chosen for comparison: the “safe” settings deep in the range of acceptable combinations, where (Exit Tube Voltage, Exit Rod Voltage) = $(-2000\text{V},$

-100V), and the “risky” settings in the upper right corner of that range (-1000V, -60V). The theoretical benefits of the safe settings are that the higher voltages help guide ions that are further from the center of the trap in cases of large ion cloud size. The risky settings are near to where the beam efficiency suddenly drops off, hence the “riskiness”. Very high voltages on the exit tube and rods produce an effect where the ions are not centered as nicely and therefore are more spread out when they hit the MCP. Therefore, the lower voltages of the risky settings mean that ions are more centered upon collision with the MCP, which is highly desirable.

To test these benefits and drawbacks, simulations were run where the standard deviation of the initial ion positions along the axes orthogonal to the axis of the beam within the trap was performed; that is, the settings were tested against various spreads and sizes of trapped ion clouds. Figures 6 and 7 demonstrate the beam efficiency and output beam radius standard deviation respectively for both the safe and risky settings. While beam efficiency remains higher for longer under the safe settings, the beam radius standard deviation—a measure of the spread of ion collision positions on the MCP—is lower for longer under the risky settings.




Figure 5: Contour plots for beam radius and beam efficiency depending on voltage values for the exit tube and exit rods.




Figure 6: a side-by-side comparison of the beam efficiency for the safe and risky settings, depending on the initial standard deviation along the x axis (orthogonal to the beam axis).




Figure 7: a side-by-side comparison of the beam radius standard deviation for the safe and risky settings, depending on the initial standard deviation along the x axis (orthogonal to the beam axis).

Conclusion

Because the beam efficiency is roughly the same for low standard deviations and beam efficiency is significantly better under the risky settings, the risky settings were chosen as the final values for the exit tube and exit rods. With these results, the table of voltage values may be updated to the following:

	Capture Stage Voltage (V)	Trapping Stage Voltage (V)	Release Stage Voltage (V)
Entrance Tube	-775	-775	-775
Entrance Rods	-57	10	10
Center Rods	-45	-45	-45
Exit Rods	10	10	-60
Exit Tube	-1000	-1000	-1000

Of course, circumstances may occur where position distributions are broad enough within the trap that increasing the voltages on the exit rods and tube may be preferable. Learning whether ions ultimately end up in a cooled, densely-packed region during the capture process would rely upon a detailed simulation of ion cooling. SIMION is a program of moderate granularity, and while it excels in gathering large amounts of data across varying combinations of voltages, it has not been sufficiently accurate when collecting cooling data. Therefore, next steps include the use of real-world data to analyze the efficiency of the ion beam within the Paul trap to verify these results and inform a more final decision regarding the voltage values.

References

- [1] Leendert Hayen, Annual Review of Nuclear and Particle Science **74**, 497-528 (2024).
- [2] M. Brodeur, T. Ahn, D. Bardayan, D. Burdette, J. Clark, A. Gallant, J. Kolata, B. Liu, P. O'Malley, W. Porter, R. Ringle, F. Rivero, G. Savard, A. Valverde, and R. Zite, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms **541**, 79-81 (2023).
- [3] David A Dahl, simion for the personal computer in reflection, International Journal of Mass Spectrometry **200**, 3-25 (2000).

Synthesis and Characterization of 1:1:3 Compounds

MARY KATE CARROLL

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Nirmal Ghimire

Abstract

The RTX_3 family of intermetallic compounds hosts a rich variety of magnetic phenomena, including spin reorientation transitions that can be leveraged to design and manipulate skyrmions. This report summarizes work performed in the Ghimire Lab Group at the University of Notre Dame during the Summer 2026 Physics REU on the synthesis and structural characterization of 113 compounds, with a focus on NdCrSb_3 . Building from the well studied, magnetically “inert” parent compound LaCrSb_3 , single crystals of NdCrSb_3 were grown via a high temperature flux growth technique and structurally characterized by powder X-ray diffraction (PXRD) and Rietveld refinement in FullProf. An initial synthesis attempt produced a diffraction pattern with peak positions matching the expected NdCrSb_3 structure but with substantially mismatched intensities, indicating an incorrect elemental ratio. Adjusting the stoichiometry of the starting elements in a second growth attempt significantly reduced the refinement chi-squared value, though some structural deviation remains. This iterative synthesis refinement cycle establishes a pathway toward phase-pure NdCrSb_3 crystals suitable for future magnetometry and neutron scattering studies of its spin reorientation behavior.

1 Introduction

Magnetic materials that host tunable, temperature or field driven changes in their spin structure are of significant interest both for fundamental condensed matter physics and for emerging spintronic technologies. One such class of tunable behavior is spin reorientation, in which the energetically preferred (“easy”) magnetic axis of a material rotates as a function of temperature or applied field, often due to a competition between distinct magnetic sublattices. Spin reorientation transitions are of particular interest because they can act as a mechanism for stabilizing skyrmions, nanoscale topological spin textures with promising applications in spintronic devices.

The RTX_3 family of intermetallic compounds offers an especially versatile platform for studying these effects. By systematically substituting the rare-earth (R), transition-metal (T), and metalloid/post-transition-metal (X) sites, the underlying electronic and magnetic behavior of the resulting material can be deliberately tuned. This report describes the synthesis and structural characterization of members of the RTX_3 family carried out in the Ghimire Lab Group, motivated by the goal of identifying and eventually manipulating spin reorientation and skyrmion physics in a new 113 compound, NdCrSb_3 .

2 Background

Compounds in the RTX_3 family follow the general formula RTX_3 , where R denotes a rare-earth element (primarily from the lanthanide series), T denotes a transition metal such as chromium or manganese, and X denotes a metalloid or post-transition metal such as antimony or tin. Because the lanthanide series spans a range of ionic radii from roughly 1.95 Å for lanthanum down to about 1.6 Å for lutetium, substituting different rare-earth ions onto the R site provides a convenient chemical handle for tuning lattice parameters and, in turn, the magnetic exchange interactions that govern the material's behavior.

Members of the RTX_3 family crystallize in an orthorhombic structure described by the Pbcm space group. Figure 1 shows the characteristic unit cell. Table 1 lists the fractional atomic coordinates and Wyckoff positions for LaCrSb_3 , one of the most extensively studied compounds in this family, as reported in the Materials Project database.




Figure 1: RTX_3 unit cell simulated in Vesta shown from a, b, and c axis.

Label	Element	a	b	c
4c	Cr	3/4	1/2	0.090435
4c	Sb	3/4	1/2	0.498548
4d	Sb	0.504103	1/4	0.689482
4d	Sb	0.606263	1/4	0.936734
4d	Sb	0.999795	1/4	0.783089

2.1 Motivation and Spin Reorientation

The central motivation for studying spin reorientation in the RTX_3 family is its connection to skyrmions. Because their topological character makes them highly stable against small perturbations, and because they can be manipulated using lower current density than those required to manipulate conventional transistors, skyrmions are considered strong candidates for the bit elements of future dense, energy efficient spintronic memory.

Spin reorientation is a temperature or field driven rotation of a material's easy magnetic axis and it provides one concrete route to deliberately design and stabilize skyrmion textures: as the preferred spin direction rotates, the resulting competition between magnetic anisotropy terms can nucleate and stabilize localized swirling spin configurations. Identifying RTX_3 compounds that exhibit spin reorientation transitions is therefore a natural first step toward engineering skyrmion hosting materials within this family

2.2 Case Study: TbMn_6Sn_6

To motivate this project, published results on the related kagome ferrimagnet TbMn_6Sn_6 provide an instructive example of both the signature and the consequences of spin reorientation. Field-cooled magnetic susceptibility data reported by Jones and collaborators show the net magnetization measured parallel to the crystallographic c-axis (summing the magnetic moments oriented along c) and perpendicular to c (summing moments lying in the ab-plane) as a function of temperature.

At a spin reorientation temperature of $T_{sr} = 309$ K, the two susceptibility curves abruptly and simultaneously invert: the component parallel to c drops sharply while the perpendicular component rises, indicating that the great majority of magnetic moments rotate collectively into the ab-plane. Because a single susceptibility curve cannot distinguish an overall change in the number of ordered moments from a change in their direction, measuring both orientations simultaneously is essential for unambiguously identifying a spin reorientation transition.

A separate 2023 study by Li and collaborators used Lorentz transmission electron microscopy to directly image the real-space magnetic texture of TbMn_6Sn_6 near this transition. Near room

temperature, an ordered array of magnetic stripe domains appear. Under a weak applied magnetic field the stripe pattern remains intact, but as the field strength increases the stripes fracture and break apart into a dense array of magnetic bubbles known as biskyrmions. Taken together, these two studies demonstrate how spin reorientation can be exploited, via an applied field, to generate and manipulate skyrmion-like magnetic textures.

2.3 Target System: LaCrSb_3 and NdCrSb_3

Within the RTX_3 family, this project focuses on LaCrSb_3 and NdCrSb_3 . LaCrSb_3 was among the first 113 compounds grown and remains the most extensively studied member of the family, making it a useful reference point. Because lanthanum has an entirely empty 4f electron shell, La^{3+} carries no localized magnetic moment of its own; the ferromagnetism observed in LaCrSb_3 is therefore driven entirely by the itinerant 3d electrons on the chromium sublattice. In the classic measurements of Leonard and collaborators on polycrystalline samples, a sharp drop in magnetization consistent with the onset of ferromagnetic order was observed near 144 K under a weak 0.01 T field. A later study by Kumar and collaborators, performed on single crystals rather than polycrystals, placed the same ordering transition closer to 125 K. This 125 K is likely closer to the value of an ideal, defect-free lattice.

This clean, single sublattice baseline sets up the central question addressed in this project: what happens to the magnetic behavior of this framework when the magnetically passive lanthanum site is replaced with a rare earth ion carrying a strong, localized 4f moment? NdCrSb_3 was chosen to answer this question. Introducing neodymium adds a second, localized magnetic sublattice that couples to the itinerant chromium spins. According to prior reports, on cooling from high temperature the chromium moments in NdCrSb_3 first order along the b-axis; between roughly 9 K and 20 K, competition between the Cr and Nd sublattices drives a spin reorientation, and below 9 K the neodymium moments lock the easy axis along the a-axis.

3 Methods

Single crystals of the target 113 compounds were grown using a high temperature self-flux growth technique. Raw elemental starting materials were loaded into crucibles in ratios designed to promote crystallization of the desired RTX_3 phase. The loaded crucibles were sealed in quartz tubes and subjected to a multi-stage heating profile: an initial 12-hour ramp brought the furnace to a peak temperature of 1273 K, which was then held for 24 hours to ensure the elemental charge fully melted into a homogeneous liquid. The furnace was then cooled slowly and continuously to 923 K over 350 hours, giving single crystals of the RTX_3 phase ample time to precipitate of the molten flux before the remaining flux was removed via centrifuge.

4 Results

Crystals recovered from each growth were characterized using powder X-ray diffraction (PXRD) followed by Rietveld refinement of the resulting pattern in FullProf. For this measurement, selected crystals were ground into a fine powder and spread as a thin, uniform film on a sample holder. The machine directs an X-ray beam at the powder while the source and detector rotate together, systematically stepping through the incident angle. Because every crystalline material produces a diffraction pattern that reflects its unique atomic arrangement, this pattern serves as a structural fingerprint: the angular positions of the diffraction peaks are set by the size and shape of the unit cell, while the relative intensities of the peaks encode which atoms occupy which positions within that cell.

The first synthesis attempt at NdCrSb_3 produced a diffraction pattern whose peak positions aligned well with the expected NdCrSb_3 structure, but whose peak intensities deviated substantially from those predicted for the ideal, stoichiometric compound. This intensity mismatch indicated that, while the correct unit cell had likely formed, the crystal did not have the intended NdCrSb_3 stoichiometry — most likely due to an imbalance in the ratio of elements loaded into the initial flux growth. In response, the elemental ratios used in the flux growth recipe were revised based on

further literature review, and the growth was repeated.




Figure 2: Rietveld refinement of the powder XRD pattern from the first NdCrSb synthesis attempt. Peak positions match the expected structure, but intensity mismatches indicate an incorrect elemental stoichiometry.

The second synthesis attempt, using the revised elemental ratios, produced a diffraction pattern with a substantially improved fit to the expected NdCrSb₃ structure. Rietveld refinement of this second dataset yielded a significantly lower chi-squared value than the first attempt, indicating that the adjustment to the growth stoichiometry meaningfully improved phase purity. Some structural deviation between the observed and calculated patterns remains, suggesting that further refinement of the growth recipe is still needed to obtain a fully phase-pure NdCrSb₃ sample.

5 Conclusion

This report presented ongoing work in the Ghimire Lab Group to synthesize and structurally characterize NdCrSb₃, a member of the RTX₃ family of intermetallic compounds, motivated by the




Figure 3: Rietveld refinement of the powder XRD pattern from the second NdCrSb synthesis attempt, following adjustment of the elemental ratios in the flux growth recipe. The reduced chi-squared value indicates improved agreement with the expected structure.

search for spin-reorientation-driven skyrmion physics analogous to that reported in TbMn_6Sn_6 and related kagome magnets. Using the well-characterized, magnetically simple compound LaCrSb_3 as a reference point, single crystals of NdCrSb_3 were grown by flux growth and characterized by powder X-ray diffraction and Rietveld refinement. An initial synthesis produced crystals with an incorrect elemental stoichiometry; adjusting the growth recipe substantially improved the structural refinement. These results lay the groundwork for future magnetic and neutron scattering studies aimed at directly probing the spin reorientation transition in NdCrSb_3 .

6 References

- [1] Jones, D. Connor, et al. “Origin of Spin Reorientation and Intrinsic Anomalous Hall Effect in the Kagome Ferrimagnet.” *Physical Review B*, vol. 110, no. 11, American Physical Society, Sept. 2024, doi:10.1103/physrevb.110.115134.
- [2] Li, Y., et al. Lorentz transmission electron microscopy study of magnetic stripe domains and field-induced biskyrmions in TbMn_6Sn_6 . *Materials*, vol. 16, no. 4, 2023, doi:10.3390/ma16041736.
- [3] Leonard, Michael, Shibaji Saha, and Naushad Ali. “Magnetic Properties of RTSb_3 .” *Journal of Applied Physics*, vol. 85, no. 8, 1999, pp. 4759–4761, doi:10.1063/1.370472.
- [4] Kumar, N., Lamba, N., Gayles, J., Le, C., Vir, P., Guin, S. N., Sun, Y., Felser, C., and Shekhar, C. “Giant Anomalous Hall Conductivity in the Itinerant Ferromagnet LaCrSb_3 and the Effect of f-Electrons.” *Advanced Quantum Technologies*, vol. 4, 2021, art. 2100023, doi:10.1002/qute.202100023.
- [5] Materials Project. Structural and Wyckoff position data for LaCrSb_3 (Pbcm space group). next-gen.materialsproject.org.

Angular Momentum Structure of Even-Mass Oxygen Isotopes from *Ab Initio* Calculations

YUNYA CHEN

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Shwetha L. Vittal and Prof. Mark A. Caprio

Abstract

The oxygen isotopic chain provides a laboratory for studying shell evolution: adding neutrons to the sd shell reveals new magic numbers at $N = 14$ and $N = 16$. We use the no-core configuration interaction (NCCI) approach with the Daejeon16 interaction to examine how these shell closures manifest in $^{18,20,22,24}\text{O}$. We implement a Lanczos decomposition that simultaneously resolves the probability distribution $P(N_k, \hat{J}_k^2)$ over valence neutron orbitals by traditional shell-model configuration quantum numbers, within the $N_{\text{max}} = 0$ model space. The results establish a baseline for the *ab initio* wave functions: the $0d_{5/2}$ orbital fills first, followed by $1s_{1/2}$, while $0d_{3/2}$ remains empty. In all dominant configurations, neutrons in $0d_{5/2}$ couple to $J = 0$, reflecting the dominance of pairing. The ^{24}O ground state is remarkably pure (96.0%), consistent with its interpretation as doubly magic.

1 Introduction

The oxygen isotopes are a proving ground for nuclear shell evolution. With fixed proton number $Z = 8$, each added neutron occupies an orbital in the sd shell. Experiment and theory have shown that neutron-rich oxygen isotopes exhibit new subshell closures at $N = 14$ (^{22}O) and $N = 16$ (^{24}O), making these nuclei doubly magic [1–4].

Ab initio methods, building the nuclear Hamiltonian directly from nucleon-nucleon interactions, offer a first-principles approach to understanding these closures. We use the no-core configuration interaction (NCCI) approach with the Daejeon16 interaction [5] to investigate the ground-state wave functions of $^{18,20,22,24}\text{O}$.

Our central question is structural: to what extent do the *ab initio* wave functions reflect the simple shell-model picture of sequential orbital filling? To answer this, we have developed and implemented a Lanczos decomposition that resolves the joint probability distribution $P(N_k, \hat{J}_k^2)$ for each valence neutron orbital—i.e., a simultaneous decomposition by the traditional shell-model quantum numbers N_k and $\hat{J}_k^2$. This capability builds on earlier work by O’Donnell [6] (who decomposed by N_k and seniority ν_k) and allows direct comparison with the standard shell model. For the sd -shell orbitals considered here ($j \leq 5/2$), N_k and $\hat{J}_k^2$ provide a complete and non-redundant classification.

Figure 1 shows the schematic shell-model filling of the oxygen isotopes. We focus on the ground

states of $^{18,20,22,24}\text{O}$ within the $N_{\text{max}} = 0$ space. Our results establish a baseline for the *ab initio* wave functions, revealing sequential filling, dominance of $J = 0$ pairing, and a remarkably pure ground state for ^{24}O . Section 2 presents the NCCI formalism and our Lanczos decomposition implementation. Section 3 gives results for the ground-state decompositions, including dominant configurations, average occupations, and $J = 0$ pairing probabilities. Section 4 concludes with a summary and discussion of future extensions.




Figure 1: Schematic shell-model picture of the oxygen isotopes: ^{16}O with filled sd -shell core (closed major shell at $N = 8$); $^{18,20,22}\text{O}$ with partial filling of the sd shell; and ^{24}O with filled $0d_{5/2}$ and $1s_{1/2}$ subshells (closed subshell at $N = 16$). The $0d_{3/2}$ orbital remains empty.

2 Method

2.1 NCCI and the $N_{\text{max}} = 0$ Space

We work within the NCCI framework, expanding the many-body wave function in a basis of anti-symmetrized harmonic oscillator states. The basis is truncated by N_{max} , the maximum number of oscillator quanta above the minimal Pauli-allowed configuration. In this work, we restrict our analysis to $N_{\text{max}} = 0$ —the minimal shell-model configuration space. This is the lowest-order truncation, and all results presented here are obtained from wave functions in this space.

Future work will extend the decomposition to larger N_{max} spaces, jointly with N_{ex} (the number of excitation quanta) as in the approach of O’Donnell [6].

2.2 The Decomposition Operator

To extract the angular momentum structure, we define the operator

$$C := \sum_k (a_k \hat{N}_k + b_k \hat{J}_k^2), \quad (1)$$

where k labels the valence neutron orbitals $1s_{1/2}$, $0d_{3/2}$, and $0d_{5/2}$. Here $\hat{N}_k$ is the occupation number operator, and $\hat{J}_k^2$ is the total angular momentum squared of neutrons in orbital k . For N_k identical fermions in a single- j orbital, the eigenvalues of $\hat{J}_k^2$ are $J_k(J_k+1)$. The allowed values of J_k are not solely given by the usual triangle inequality; they are further restricted by the antisymmetry of the many-fermion state. For example, for two identical nucleons in the same orbital, only even J_k values are allowed [7]. This constraint is essential for correctly identifying the configurations.

The coefficients a_k and b_k are chosen to obtain a simultaneous decomposition with respect to the quantum numbers (N_k, J_k) , ensuring that each distinct eigenvalue of C corresponds to a unique set of these quantum numbers. The Lanczos decomposition of a target state $|\Psi\rangle$ with respect to C yields the probability distribution $P(\{N_k, J_k\})$ over these configurations [8; 9]. The idea of using a linear combination of Casimir operators to achieve simultaneous decomposition was introduced in Ref. [10]. In future applications to $N_{\max} > 0$ spaces, the operator can be extended to include the excitation quanta N_{ex} as well.

We note that for the orbitals in the sd shell ($j \leq 5/2$), the seniority quantum number v_k provides no additional information beyond N_k and J_k , because the mapping from (N_k, v_k) to (N_k, J_k) is one-to-one (for a given N_k , v_k uniquely determines J_k or vice versa). For higher- j orbitals ($j \geq 7/2$), a joint decomposition with v_k would be more informative (see, e.g., Suhonen Fig. 12.3 [7]), but this is beyond the scope of the present work.

2.3 Computational Details

We obtain NCCI wave functions for the 0^+ ground states of $^{18,20,22,24}\text{O}$ using the MFDn code [11]. The many-body Hamiltonian matrix is generated using the Daejeon16 two-nucleon interaction [5],

including the Coulomb term. The basis is the harmonic oscillator basis with oscillator parameter $\hbar\omega = 15$ MeV. All calculations are performed in the $N_{\text{max}} = 0$ space. The coefficients in C are optimized via a simulated annealing algorithm to maximize spectral separation.

3 Results and Discussion

We now examine the ground-state decompositions of the four oxygen isotopes. Figure 1 shows the dominant configurations for each isotope. Each bar represents a distinct shell-model configuration, labeled on the horizontal axis by its physical content (e.g., $(0d_{5/2})_{J=0}^2$ for two neutrons in the $0d_{5/2}$ orbital coupled to $J = 0$). Only configurations with probability exceeding 0.01% are shown.

The pattern follows the expected shell ordering. For ^{18}O , the ground state is overwhelmingly $(0d_{5/2})_{J=0}^2$ with a probability of 96.9%. This is precisely the configuration predicted by the pairing interaction for two identical nucleons in a single orbital.

Adding two neutrons to reach ^{20}O , the dominant configuration becomes $(0d_{5/2})_{J=0}^4$ at 81.5%. However, a non-negligible admixture of $(1s_{1/2})_{J=0}^2(0d_{5/2})_{J=0}^2$ appears at 17.2%, indicating that the nucleons also occupy the $1s_{1/2}$ orbital.

For ^{22}O , the $0d_{5/2}$ orbital reaches its capacity of six neutrons, producing the configuration $(0d_{5/2})_{J=0}^6$ at 76.4%. An admixture of $(1s_{1/2})_{J=0}^2(0d_{5/2})_{J=0}^4$ appears at 22.4%. This significant mixing indicates that the energy difference between the configurations with two neutrons in $1s_{1/2}$ versus two additional neutrons in $0d_{5/2}$ is relatively small compared to the residual interaction strength that couples them. It may reflect the near-degeneracy of the $1s_{1/2}$ and $0d_{5/2}$ single-particle energies in this mass region, but the two-body interaction also plays a crucial role in determining the final energies.

The ^{24}O ground state is strikingly pure. The dominant configuration $(1s_{1/2})_{J=0}^2(0d_{5/2})_{J=0}^6$ accounts for 96.0% of the wave function. This exceptional purity is strong evidence for a subshell closure at $N = 16$, consistent with the doubly-magic character of ^{24}O .

Figure 2 shows the average occupation numbers $\langle \hat{N}_k \rangle$ obtained by summing $N_k \times P(\text{configuration})$ over the full decomposition. The pattern is consistent with the discrete configuration analysis: $0d_{5/2}$




Figure 2: Dominant configurations in the $N_{\max} = 0$ space for the ground states of ^{18}O through ^{24}O . The horizontal axis labels indicate the six-tuple $(\hat{N}_{1s}, \hat{J}_{1s}^2, \hat{N}_{0d_{3/2}}, \hat{J}_{0d_{3/2}}^2, \hat{N}_{0d_{5/2}}, \hat{J}_{0d_{5/2}}^2)$. The values shown for $\hat{J}^2$ are the eigenvalues $J(J+1)$.

occupation increases steadily from ^{18}O to ^{24}O ; $1s_{1/2}$ occupation is negligible until ^{24}O , where it rises notably as the $0d_{5/2}$ shell reaches capacity; and $0d_{3/2}$ remains effectively unoccupied across the entire chain, underscoring the persistence of the $N = 16$ shell gap at this truncation level.

A key feature of the decomposition is that it resolves not only the occupation numbers but also the coupled angular momentum J for each orbital. In all dominant configurations shown in Figure 1, the neutrons in the $0d_{5/2}$ orbital appear with $J = 0$. This is not a trivial result: for four or six particles, $J \neq 0$ couplings are permitted by angular momentum algebra. The fact that $J = 0$ dominates speaks to the strength of the pairing interaction [7].

To quantify this, we compute the total probability that the $0d_{5/2}$ orbital couples to $J = 0$,




Figure 3: Average occupation numbers $\langle \hat{N}_k \rangle$ for the valence neutron orbitals in the ground states of ^{18}O through ^{24}O . The $0d_{5/2}$ orbital dominates across all isotopes, with occupation increasing monotonically. The $1s_{1/2}$ orbital shows a modest rise only in ^{24}O , while the $0d_{3/2}$ orbital remains essentially empty throughout. These trends confirm the shell-model ordering.

summed over all configurations. The results are shown in Figure 3. For $^{20,22,24}\text{O}$, this probability remains near 100% (98.7%, 99.2%, and 98.5%, respectively), and is lower for ^{18}O , at 86.6%. The lower value in ^{18}O reflects the fact that with only two neutrons in the orbital, the $J = 0$ coupling is the dominant one, but other J values can also contribute; however, the overwhelming majority is still $J = 0$. The pairing structure is remarkably robust in the $N_{\text{max}} = 0$ truncation.

These results establish a quantitative baseline for the oxygen isotopes at the lowest level of the NCCI expansion. The appearance of $1s_{1/2}$ admixtures already at $N_{\text{max}} = 0$ in $^{20,22}\text{O}$ indicates that the Daejeon16 interaction generates correlations beyond the simple filling picture—even before $2\hbar\omega$ excitations are included. The purity of ^{24}O reinforces its status as doubly magic.

4 Conclusion

We have analyzed the ground-state wave functions of $^{18,20,22,24}\text{O}$ from NCCI calculations at $N_{\text{max}} = 0$ using the Daejeon16 interaction [5]. The Lanczos decomposition reveals that the *ab initio* wave




Figure 4: Probability that the $0d_{5/2}$ neutrons couple to $J = 0$, summed over all configurations, for the ground states of ^{18}O through ^{24}O .

functions follow the naive shell-model ordering— $0d_{5/2}$ fills to capacity, followed by $1s_{1/2}$, while $0d_{3/2}$ remains empty—and that the $0d_{5/2}$ orbital consistently couples to $J = 0$, reflecting the dominance of pairing. The exceptional purity (96.0%) of the ^{24}O ground state provides strong evidence for the $N = 16$ subshell closure (filled $1s_{1/2}$ and $0d_{5/2}$ subshells), consistent with its doubly-magic character.

These findings establish a quantitative baseline for the $N_{\text{max}} = 0$ truncation, which is essential for interpreting the convergence behavior of the NCCI expansion. A key limitation of the present work is its restriction to the lowest-order model space; higher- N_{max} calculations are required to fully capture the mixing with intruder configurations that lie outside the $0\hbar\omega$ space [12]. When we go to higher N_{max} spaces, we will need to simultaneously decompose with N_{ex} to distinguish $N_{\text{ex}} = 0$, $N_{\text{ex}} = 2$ (intruder), and higher configurations, thereby providing a quantitative test of the convergence of the NCCI expansion for these nuclei. These extensions will also be applied to excited states and odd-mass isotopes, offering a more complete microscopic perspective on shell evolution in the sd shell.

Acknowledgements

The author sincerely thanks Shwetha L. Vittal and Prof. Mark A. Caprio for their patient and meticulous guidance throughout all stages of this project. Special thanks are due to Prof. Umesh Garg and Ms. Kristen Amsler for organizing the University of Notre Dame Physics REU program and making this research opportunity possible. Computational resources were provided by the Notre Dame Center for Research Computing. The author also gratefully acknowledges Xi'an Jiaotong University for providing financial support for the travel to the United States.

References

- [1] G. R. Jansen, J. Engel, G. Hagen, P. Navrátil, and A. Signoracci, *Physical Review Letters* **113**, 142502 (2014).
- [2] X. Pereira-López, B. Fernández-Domínguez, F. Delaunay, *et al.*, *Physics Letters B* **811**, 135962 (2020).
- [3] M. Stanoiu *et al.*, *Physical Review C* **69**, 034312 (2004).
- [4] O. Sorlin and M.-G. Porquet, *Progress in Particle and Nuclear Physics* **61**, 602 (2008).
- [5] A. M. Shirokov, I. J. Shin, Y. Kim, M. Sosonkina, P. Maris, and J. P. Vary, *Physics Letters B* **761**, 87 (2016).
- [6] K. O'Donnell, 2024 NSF/REU Program, Department of Physics and Astronomy, University of Notre Dame, NSF REU Student Research Papers, University of Notre Dame (2024).
- [7] J. Suhonen, *From Nucleons to Nucleus: Concepts of Microscopic Nuclear Theory* (Springer, Berlin, 2007).
- [8] R. R. Whitehead, in *Theory and Applications of Moment Methods in Many-Fermion Systems*, edited by B. J. Dalton, S. M. Grimes, J. P. Vary, and S. A. Williams (Plenum, New York, 1980) pp. 235–255.

- [9] C. W. Johnson, Physical Review C **91**, 034313 (2015).
- [10] M. A. Caprio, Bulgarian Journal of Physics **49**, 57 (2022).
- [11] P. Maris, H. M. Aktulga, M. A. Caprio, U. V. Çatalyürek, E. G. Ng, D. Oryspayev, H. Potter, E. Saule, M. Sosonkina, and J. P. Vary, Journal of Physics: Conference Series **403**, 012019 (2012).
- [12] C. W. Johnson, M. A. Caprio, and S. L. Vittal, Challenges for first-principles nuclear structure: ^{11}Li and ^{29}F (2026), arXiv:2511.00313v3 [nucl-th] .

Commissioning and Characterization of Micro Channel Plate Detectors for Use in St. Benedict and BEARTrap

AKIRA CLARK

2026 REU Program

Department of Physics

University of Notre Dame

ADVISORS: Dr. Maxime Brodeur and Sam Porter

ABSTRACT

The commissioning and characterization of micro-channel plate (MCP) detectors for use in a variety of beta-decay experiments is integral for providing precise and accurate determinations of the decay properties of daughter ions. A set of position-sensitive, $80 \times 100 \text{ mm}^2$ MCPs is being commissioned for use in both the BEtA Recoil ion trap (BEARTrap) at Argonne National Laboratory for beta-delayed neutron emission spectroscopy and the Superaligned Transition Beta-Neutrino Decay Ion Coincidence Trap (St. Benedict) experiment at the University of Notre Dame. Characterization of these MCPs, including evaluating timing resolution, spatial responses, and systematic uncertainties affecting time-of-flight measurements, is ongoing, and the status and results of these tests will be presented. More specifically, applications for measurements at BEARTrap for beta-delayed neutron spectroscopy and at St. Benedict for beta-neutrino angular correlation parameter determination will be discussed.

1 INTRODUCTION

In nuclear physics, micro-channel plate (MCP) detectors are one of the most common experimental tools. They offer a wide range of uses such as time-of-flight measurements, beam diagnostics, and mass spectrometry [5]. MCPs are often chosen over other devices due to their fast timing, high spatial resolution, and compact size [5]. Because of their consistent use in nuclear physics to detect low-energy charge particles, MCPs need to be well characterized. In the Superaligned Transition Beta-Neutrino Decay Ion Coincidence Trap (St. Benedict) experiment, the MCPs are expected to be under high voltages up to approximately **-4 kV**, while under vacuum conditions near **10^{-8} Torr**. With these conditions in mind, the need to apply high-voltage bias to

the MCPs, monitor their current output, and reduce outgoing signal noise are the primary focus of characterization and commissioning.

2 BACKGROUND

2.1 MCPs

MCPs, in essence, consist of 1 or more thin plates, usually made of glass, with extremely narrow channels [5]. These micro-channels are usually around 5-20 microns in diameter and run through the plates at an angle between 8-13° [5]. When a particle travels into one of these micro-channels, it will collide with the walls, releasing electrons, which will also collide with the walls [5]. This creates a cascade of electrons as the particle and electrons bounce off the walls. If there is a large potential difference between these plates, the shower of electrons becomes larger, producing a signal orders of magnitude larger than the initial interaction alone would produce [5]. By stacking more plates, the signal can become stronger, though eventually the signal will become too noisy. The MCPs currently being used are considered “Z-stack MCPs.” This means the plates are stacked in a way that the channels create a “Z” pattern, leading to more collisions and a higher gain [5]. A ceramic anode behind the MCP stack collects the electron cloud and provides position information

2.2 Detector Biasing

It is important to bias the MCPs as part of the commissioning process. In this context, biasing refers to applying and holding a large potential difference across the MCP stack for an extended period. Over the course of days and weeks, while conditioning the MCP, increasing the potential difference will test the MCPs ability to be held at higher potential differences. During this time, it is crucial that the MCP current maintains a steady relationship with the increasing voltage. If there are peaks in current, it signals there is potential sparking, which may be due to defects or

buildup on the MCP. Dips in the current can stem from leaks or shorts in circuitry or connections.


2.3 Signal Output

Several quantities were monitored during characterization, including applied voltages, current, pick-off signal, corner signals, and noise. Focusing on the corner and pick-off signals, it is expected that the pick-off signal shows up slightly after the corner signal when viewing them with an oscilloscope. For these signals, timing precision and clean signals clarity are the primary goals. With clean signals and prompt pick-off signals, an MCP can be expected to make accurate time-of-flight measurements [5]. The pick-off signal is an indication that a decay event has occurred and is the time variable. The four corner signals relay position information and will become more useful later on. The sum corner signal is a combination of all four corner signals, it is used to compare timing from the pick-off signal.

2.4 Noise Reduction

Reducing electronic noise is often a challenge in detector experiments. In the case of MCPs, noise is often generated from radiofrequency (RF) signals coming from circuits, wires, and even occasionally the MCP itself. These frequencies interfere with the signal's accuracy. In most cases, RF noise is reduced using a combination of foil covering, common grounding, and source-isolation techniques.

3 SETUP



(Figure 1) A flow chart of the setup process for offline MCP testing

The current offline setup used to test MCPs consists primarily of a vacuum chamber housing the test MCP, an alpha radiation source, and an external high-voltage supply. From the high-voltage supply, voltage is received by a voltage-divider box, dividing the incoming voltage across the MCP plates. The front plate holds the greatest potential difference, while the middle and back plates hold less potential difference as you move towards the anode. After the alpha source creates a signal by interacting with the MCP, that signal is relayed to a preamp box to further boost the signal. This preamp is powered by an outside **12 V** supply. The preamp also connects to the pick-off box which will read the signal and provide an accurate relay of decay interactions. Through flange cap connectors, the corner signals join together in a separate preamp box. They then connect to an oscilloscope to read the timing along with the pick-off signal.

4 PROCEDURE

4.1 Voltage Biasing

Over the course of several weeks, the MCP was biased under increasing voltage beginning at - 2.7 kV. To facilitate steady buildup, the voltage was increased at a rate of -100 V / 30 mins after that. At times, the current was recorded to compare to the voltage and ensure a steady trend between the two. As voltage was continuously increased through bias testing, recordings of the

pick-off and corner signals were recorded as well. In the low -3 kV range, it was suggested to test the potential differences between the plates through the corner signals. This was also recorded at the time.


4.2 Noise Chasing

As mentioned previously, noise chasing is often a part of experimental procedures. To improve the clarity of the pick-off signal, the original, thin aluminum housing unit was replaced for a thicker, RF-blocking housing. A majority of cables connecting the corner signals to the oscilloscope were covered in foil or replaced with shorter cables to reduce RF interference. The voltage-divider box, preamp, and pick-off box were put into contact with each other to provide them with a common ground. Finally, more foil shielding was placed around the pick-off, preamp, and voltage-divider box to reduce noise coming from their circuits.


5 RESULTS

5.1 Voltage Bias

Over the course of 10 weeks, the MCP was biased multiple times. After a few weeks, the high-gain MCP was exchanged for a low-gain MCP. Biasing restarted on the low-gain MCP and continued after. After biasing consecutively to -3 kV, some voltage and current recordings were conducted, see figure (3). The voltage tests were conducted on the front, middle, and back plates, while the current tests were conducted through the whole circuit. Figures (1) and (2) show a couple of possible sparks between the middle and back plates, but in both cases, the voltage returned to a steady decline. The current also showed a direct linear correlation to voltage, which is expected.




(Figure 2) Voltage test at -2 kV. Middle (blue) and back (red) in -kV. Current (yellow) in -mA.




(Figure 3) Voltage test at -3 kV. Middle (blue) and back (red) in -kV. Current (yellow) in -mA.

5.2 Pick-off and Corner Signals

With some preliminary testing below -3 kV, it was clear that the MCP did not hold a large enough potential difference between the plates to read a pick-off signal. Though, it was high enough to read the sum signal of the corners. Continuing to increase the voltage did start to show the indication of a pick-off signal, at around -3.2 kV. The voltage was continuously increased until -3.7 kV, where the pick-off signal was more obvious, but also considerably noisier. That is not unexpected, as increasing the gain through voltage has a clearly defined affect on the outgoing signal. [5] The difference is clearer to see looking at figure (4) and figure (5). The smaller focus now comes down to the time distance between the corner sum and pick-off signal. It is optimal for these two signals to be occurring as close to each other as possible. In figure (4), it is easier to see that the pick-off signal is about 150ns after the sum signal. This is not inherently bad, but shorter is better



(Figure 4) Sum and pick-off signal at -3 kV. Ch1 (yellow) is the sum signal, Ch2 (blue) is the pick-off signal. The x-axis indicates time in blocks of 500 ns. The y-axis indicates amplitude in mV.



(Figure 5) Sum and pick-off signal at -3.7 kV. Ch1 (yellow) is the sum signal, Ch2 (blue) is the pick-off signal. The x-axis indicates time in blocks of 500 ns. The y-axis indicates amplitude in mV.

5.3 Noise

While noise is inherently always going to be an issue, there are many ways to combat it. Foiling was the first method. Any SHV cables, the voltage-divider box, pick-off box, and connecting preamp were all covered with aluminum foil. It did not show obvious effects but was kept just in case. The next major noise fix was to rehouse the pick-off box into a thicker, RF-blocking box. This seemed to make a noticeable difference and turned out to be the only major contribution to noise reduction. It is important to note that these tests were done on a low-gain MCP, as the high-gain MCP was no longer available for use. This could certainly cause an increase in noise, but one that is currently untestable.

6 DISCUSSION

With the results gathered, it is appropriate to continue with commissioning the MCP. Both pick-off and sum signals are readable, and not unreasonably distant. The next goal would be to take position and decay count data. These tests will be sure to root out any physical issues with the

MCP. If these tests confirm no significant issues, the MCP will be ready to be used in St. Benedict. If not, then the characterization will continue until those issues are solved.

REFERENCES:

- [1] W. S. Porter, A. T. Gallant, M. Brodeur, P. D. O'Malley, F. Rivero, and R. Zite, *Nuclear Instruments and Methods in Physics Research Section B* 569, 165893 (2025).
- [2] W. S. Porter, M. Brodeur, O. Bruce, D. P. Burdette, J. A. Clark, O. Eversole, A. T. Gallant, A. M. Houff, J. J. Kolata, E. Markel, P. D. O'Malley, F. Rivero, G. Savard, D. Schroeder, A. A. Valverde, R. Zite, D. W. Bardayan, S. R. Carmichael, C. Dembski, P. L. D. Magro, and W. W. von Seeger, *Physica Scripta* 100, 085301 (2025).
- [3] M. Brodeur, T. Ahn, D. W. Bardayan, D. P. Burdette, J. A. Clark, A. T. Gallant, J. J. Kolata, B. Liu, P. D. O'Malley, W. S. Porter, R. Ringle, F. Rivero, G. Savard, A. A. Valverde, and R. Zite, *Nuclear Instruments and Methods in Physics Research Section B* 541, 79–81 (2023).
- [4] W. S. Porter, D. W. Bardayan, M. Brodeur, D. P. Burdette, J. A. Clark, A. T. Gallant, A. M. Houff, J. J. Kolata, B. Liu, P. D. O'Malley, C. Quick, F. Rivero, G. Savard, A. A. Valverde, and R. Zite, *Atoms* 11, 129 (2023).
- [5] J. L. Wiza, "Microchannel plate detectors," *Nuclear Instruments and Methods* 162, 587–601 (1979).

Construction and Beam Simulation Analysis for DUNE

SABINA DAY

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Dr. Laura J. Fields

Abstract

The Deep Underground Neutrino Experiment (DUNE) will help determine three mysteries about our universe: why antimatter lost the battle to matter, why supernovae lead to the development of the basic materials required to make up living things, and possibly whether protons have an infinite or finite lifespan. In this research project, DUNE photon detection system modules were built along with a detailed catalog of the kinds of nuclear interactions that create neutrinos in accelerator-based neutrino beams as a means to help answer these questions. This catalog provides key input to hadron production experiments that will make measurements allowing full exploitation of the DUNE neutrino beam.

1 Introduction

1.1 Basics of DUNE

DUNE consists of two neutrino detectors at two different locations at opposite ends of the world's most intense neutrino beam: an underground particle detector at the Sanford Underground Research Facility in South Dakota at the Long Baseline Neutrino Facility (LBNF) and at a proton accelerator at Fermi National Accelerator Laboratory (FNAL) in Illinois. FNAL will record particle interactions near the source of this beam, while the larger detector is placed 1,300 km away from the source and 1.5 km underground at LBNF (*Deep Underground Neutrino Experiment, n.d.*).

The Photon Detection System (PDS) of the Deep Underground Neutrino Experiment (DUNE) consists of photon detectors integrated into cathode planes and cryostat membranes; therefore, kits were constructed for DUNE photon detection system modules, and QC'ing, cleaning, and kitting were conducted for around 700 kits. Eventually, these pieces will become part of the DUNE far detector in the Sanford Underground Research Facility in South Dakota. Sanford hosts an underground particle detector that collaborates with a proton accelerator at Fermilab in Illinois (*Botogoske, et al.*). Additionally, beam simulation analysis was created using

standard HEP data analysis software (linux, C++, Python, and ROOT) and plots were created from DUNE simulation output files.

1.2 Physics Goals

DUNE aims to answer several different questions, including the case of matter versus anti-matter and determining if neutrinos are the primary reason for why our universe is matter-dominant; thus, DUNE will analyze the behavioral differences between neutrinos and anti-neutrinos to understand more about the development of our universe. Furthermore, DUNE also intends to evaluate the massive streams of neutrinos that result from supernovae, as these exploding stars could allow observation for the formation of black holes and neutron stars.

2 Background

2.1 DUNE Photon Detection System (PDS)

The main scientific objective of the Photon Detection System (PDS) for DUNE is to enhance the research into fulfilling the goals of the experiment: investigating proton decay, supernovae neutrino bursts, and solar neutrinos. The Photon Detectors (PDs) are primarily located in the Far Detector modules at the FNAL location where neutrinos are produced. Here, the PDs observe neutrino interactions by essentially trapping 128 nm scintillation photons. Fundamentally, the total amount of detected photons, called the total Light Yield (LY), is one of the most important properties of the PDs.

The PDs are made up of several different components all integrated in anode planes (APA) in a cathode frame along membrane walls.

2.2 Neutrino Beam

The Main Injector, a powerful particle accelerator at Fermilab, extracts protons and simultaneously propels those particles to 120 GeV. The Main Injector then guides the protons towards DUNE’s near and far detectors in Illinois and South Dakota. At these locations, the protons are smashed into a cylindrical graphite target, colliding with nuclei to produce new particles such as (+) and/or (-) pions, as well as neutrons. From the electric charges of these produced particles, a series of magnetic horns controls their path of direction by selecting positive pions and kaons and focusing them into a narrow beam.

Throughout this beam, positive pions continue to spontaneously decay into other positive muons and neutrinos. These newly created pions and kaons eventually decay into more neutrinos, in which the particles can no longer be directed by the magnetic horns. Along this beam line, a block of concrete and steel is in place to slow and absorb the muons; however, neutrinos are virtually “ghost-like” in nature and can pass through any form of matter. Thus, a neutrino beam is created once the muons are absorbed by the concrete and steel block.

Afterwards, the newfound neutrino beam passes through Fermilab’s Particle Detector, proceeding to travel ~1,300 km (800 miles) underground to South Dakota’s liquid-argon detector. Conclusively, the neutrino beam is used to search for and analyze rare neutrino-matter interactions (*DUNE at LBNF, n.d*).

2.3 Hadron Production Measurements

Hadron production measurements for DUNE are being performed at the NA61/SHINE spectrometer at CERN’s Super Proton Synchrotron and FAIR at GSI in Germany (*Ilieva, S., 2020*). FAIR, a proposed facility, at GSI hopes to conduct particle interaction measurements of hadron beams at energies 0.5 – 2.5 GeV. NA61/SHINE, however, uses 120 GeV/c proton beams

(DUNE's primary proton energy) to collide with thin carbon targets. NA61/SHINE also proposes to measure hadron beam interactions with 2 – 20 GeV particle beams.

Currently, these two Hadron Production experimental facilities have voiced their need to obtain measurable data to quantify the Neutrino flux uncertainty of DUNE. To quantify this flux, these groups primarily need to understand what kind of particle interactions need to be measured. Most importantly, the nature of what interactions are forming, or decaying into, neutrinos at the 2nd oscillation maximum ($0 \rightarrow 1$ GeV) and at the 1st oscillation maximum ($0 \rightarrow 6$ GeV) of DUNE's beamline must be understood.

Conclusively, the results of this project's simulation analysis will determine whether these experimental facilities could potentially improve on the current knowledge of the DUNE Flux.

3 Methods

3.1 DUNE Photon Detection System Modules Construction

Throughout this project, PDS Modules were constructed with various PDS components, including cleaning and quality control methods of said components. To elaborate, photon detector parts were bagged and sorted as can be seen below in **Figure 2**.



Figure 1: [Left] Counting screws to be bagged into labeled detector parts bags; [Right] sorting and

organizing of bagged parts.

3.2 HEP Data Analysis Software, DUNE Simulation Output Files

3.2.1 DUNE Oscillations

The DUNE Neutrino Energy Spectrum consists of four types of neutrinos: ν_μ , ν_e , $\bar{\nu}_\mu$, and $\bar{\nu}_e$. As neutrinos travel over long distances, these particles change “flavor”, or type. In **Figure 3** below, the energy spectrum of each neutrino type can be seen [Left], as well as their oscillation maximums [Right].




Figure 2: [Left] Energy spectrum of ν_μ , ν_e , $\bar{\nu}_\mu$, and $\bar{\nu}_e$ neutrinos; [Right] oscillation maximums of ν_μ , ν_e , $\bar{\nu}_\mu$, and $\bar{\nu}_e$ neutrinos.

Throughout the 1300-km distance that runs along DUNE’s beamline, these neutrinos begin their travel during the neutrino production process that takes place after the proton accelerator at Fermilab in Illinois. Once reaching Sanford in South Dakota, these neutrinos begin their oscillations.

The region where neutrinos possess energies > 1 GeV has been relatively well studied, with neutrinos possessing energies < 1 GeV having been mostly ignored. One of the main goals of this project aimed to properly catalog the particle interactions that decayed into neutrinos in this ignored region ($E_\nu < 1$ GeV). In other words, the particle interactions of the 2nd and 1st ($E_\nu < 6$ GeV) oscillation maximums, the first two peaks of the *Neutrino Energy (GeV)* vs.

$P(\nu_\mu \rightarrow \nu_e)$ plot {Figure 3, [Right]}, were studied to determine what kind of particle interactions create neutrinos in these regions.

4 Results

4.1 Neutrino Beam Simulation Analysis

2nd Oscillation Maximum, Low Energy ($E_\nu < 1 \text{ GeV}$)




Figure 3: 2D plot [Left] of Pion ($\pi^\pm$) particle interactions that produce neutrinos with energies $E_\nu < 1 \text{ GeV}$ and its corresponding momentum vs. angle ($P_{anc}(\text{GeV}/c)$ vs. $\theta_{anc}(\text{Degrees})$) plot [Right].

In **Figure 4**, the 1D plot [Left] represents ancestor Pion-particles that decayed to produce neutrinos with $E_\nu < 1 \text{ GeV}$. As can be seen in **Figure 4**, these Pion-particle interactions peak around $3 \text{ GeV}/c$. With the 2D momentum vs. angle plot [Right], it can be deduced that high momentum parents decay at small angles to yield low E_ν . With these kinds of 3D plots, the NA61/SHINE and FAIR experimental groups can begin to plan measurements of the energies and angles of these produced particles from Pion interactions.

<i>Interaction Summary</i>		
Primary Particle	All Neutrinos	Neutrinos $E_\nu < 1 \text{ GeV}$

Proton Interactions	46%	46%
Pion Interactions	48%	35%
Kaon Interactions	2%	2%
Other (e.g. neutrons)	4%	17%

Table 1: Table of interactions for ancestor particles that produce all Neutrino energies and Neutrinos $E_\nu < 1 \text{ GeV}$.

Table 1 lists the percentages for each interaction type for a specific primary particle that produces an energy-dependent neutrino, which in turn helps understand what kind of interactions make neutrinos at DUNE. The amount of interactions in the neutrino ancestor chain makes up all interactions, or in other words, 100% of total interactions. From this table, it can be interpreted that 3 nuclear interactions produce a single neutrino on average. Additionally, most of these interactions include protons (46%) or pions (35%) as the primary particle, with a large amount of other particles (17%) making up other interactions.

1st Oscillation Maximum, High Energy ($E_\nu < 6 \text{ GeV}$)




Figure 4: 2D plot [Left] of Pion ($\pi^\pm$) particle interactions that produce neutrinos with energies $E_\nu < 6 \text{ GeV}$ and its corresponding momentum vs. angle ($P_{anc}(\text{GeV}/c)$ vs. $\theta_{anc}(\text{Degrees})$) plot [Right].

In **Figure 5**, the 2D plot [Left] represents an ancestor Pion-particle that decayed to produce a neutrino with $E_\nu < 6 \text{ GeV}$. From this plot, the Pion momentum distributions peak around 5 – 6 GeV/c, which could understandably be deduced, in comparison to **Figure 4**, that

particles move faster with higher neutrino energies. The 3D momentum vs. angle plot [Right] simultaneously shows that higher energy particles are a result of parent particles primarily decaying at wider angles (yellow).

<i>Interaction Summary</i>		
Primary Particle	All Neutrinos	Neutrinos $E_\nu < 6 \text{ GeV}$
Proton Interactions	46%	48%
Pion Interactions	48%	39%
Kaon Interactions	2%	1%
Other (e.g. neutrons)	4%	12%

Table 2: Table of interactions for ancestor particles that produce all Neutrino energies and Neutrinos $E_\nu < 6 \text{ GeV}$.

Similar to **Table 1**, **Table 2** also helps understand what kind of interactions make neutrinos at DUNE. When compared to **Table 1**, however, fewer interactions per neutrino are observed (2.8 Ints/Neutrino). Additionally, Pion interactions are shown to dominate at higher neutrino energies (39% vs. 35%). On the other hand, the other particle interactions now only make up 12% of all interactions in contrast to the 17% that was seen for neutrinos with energies $E_\nu < 1 \text{ GeV}$.

4.2 Discussion

In retrospect, many of these Pion-neutrinos have energies at the 2nd and 1st oscillation maximums. The juxtaposition between the two maximums show that Pion-neutrino interactions were seen to have peaked at 3 GeV/c when neutrinos had energies $E_\nu < 1 \text{ GeV}$, but peaked higher at $\sim 5 \text{ GeV/c}$ when neutrinos had energies $E_\nu < 6 \text{ GeV}$.

These interactions at these maximums show potential to be measured at FAIR, NA61, and/or SHINE as Pion beam data is very important for understanding the DUNE Neutrino Flux Uncertainty data below 1 GeV.

5 Conclusion

In conclusion, to measure neutrino oscillations, DUNE will need a precise estimation of the neutrino beam flux before oscillation. From this project, we have now studied the particle beams needed to precisely estimate the neutrino flux below 1 GeV at DUNE's 2nd oscillation maximum. Additionally, we have shown that proposed experimental facilities at GSI and CERN are well suited to provide the measurements that DUNE needs in these areas.

In a continuation of this project, it would be best to determine what those unknown particles making up 17%/12% of total interactions are, and more importantly, how they affect the measurements needed to quantify the Neutrino Flux Uncertainty.

References

- Botogoske, G., & the DUNE collaboration. (2025). DUNE Photon Detection System. *Journal of Instrumentation: An IOP and SISSA Journal*, 20(06), C06034.
<https://doi.org/10.1088/1748-0221/20/06/c06034>
- Deep underground neutrino experiment*. (n.d.). Dunesience.org. Retrieved July 30, 2026, from <https://www.dunesience.org/>
- DUNE Fluxes*. (n.d.). Crc.nd.edu. Retrieved July 30, 2026, from <https://glaucus.crc.nd.edu/DUNEFluxes/>
- Ilieva, S. (2020). Hadron production measurements at NA61/SHINE for precise determination of accelerator neutrino fluxes. In *arXiv [hep-ex]*. <https://doi.org/10.48550/arXiv.2011.00277>
- Session, R. P. (n.d.). *The Photon Detection System of the DUNE experiment*. Cern.Ch. Retrieved July 30, 2026, from https://indico.cern.ch/event/1283224/contributions/5590466/attachments/2724930/4735409/Hamza_CPAN_RENATA.pdf
- (N.d.). Harvard.edu. Retrieved July 30, 2026, from <https://ui.adsabs.harvard.edu/abs/2024APS..APREE2002T/abstract>

Single Crystal Growth and Magnetotransport Studies of Ce_2PdGe_6

MAGGIE DWYER

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Nirmal J. Ghimire

Abstract

Altermagnets are a new class of magnetic materials that share beneficial properties of ferromagnets and antiferromagnets. Only a handful of altermagnets have been verified experimentally, and thus the search for more remains an open and exciting challenge in the field of materials engineering. One signature of altermagnets is that they can give rise to an anomalous Hall effect, depending on symmetry conditions. This phenomenon has only reliably been observed in one altermagnet, MnTe. Ce_2PdGe_6 and Ce_2CuGe_6 are two materials, currently classified as collinear antiferromagnets, which exhibit AHE and appear to be strong candidate altermagnets. While previous studies have probed the magnetic and physical properties of Ce_2PdGe_6 , its magnetic structure remains unverified; thus, our understanding of this material is still evolving. The goal of this project is to grow high-quality single crystals of Ce_2PdGe_6 to investigate its magnetic and transport properties. In the future, neutron diffraction measurements will be performed for Ce_2PdGe_6 to determine the magnetic structure. Later, Ce_2CuGe_6 will also be studied. Angle Resolved Photo Emission Spectroscopy (ARPES) will be used to probe the electronic band structure of these materials, which can show altermagnetic characteristics like spin-split bands.

1 Introduction

Altermagnets, a class of magnets just recently discovered in 2019, are characterized by a specific combination of properties of both ferromagnets and antiferromagnets that arise due to broken symmetries in the underlying crystal structure (see section 2). Previous studies have indicated that the intermetallic compounds Ce_2CuGe_6 and Ce_2PdGe_6 hosts an interesting combination of magnetic and physical properties [1]. Based on these properties, and symmetry analysis of the known magnetic structure of Ce_2CuGe_6 , we realized that these two related compounds may be altermagnets. To date, both of these compounds have been classified as collinear antiferromagnets. The Néel temperature (T_N) of Ce_2PdGe_6 is 11.5 K and that of Ce_2CuGe_6 is 15 K. They crystallize in an orthorhombic structure and belong to space group Cmca [1]. These compounds are part of the $R_2T\text{Ge}_6$ (R : rare earth, T : transition metal) family. Magnetism in these compounds arises from the cerium 4f-electrons, which are responsible for creating local magnetic moments that collectively define this material's magnetic texture [2]. Both compounds have been shown to exhibit the anomalous Hall effect. The magnetic structure of Ce_2CuGe_6 was previously determined by neutron

powder diffraction [1]. Symmetry analysis this structure, provided in 1, suggests this compound is altermagnetic; According to this published structure, the Ce moments in the layers perpendicular to the crystallographic c -axis are ferromagnetic. When coupled along the c -axis, these layers alternate between ferromagnetic and antiferromagnetic coupling, resulting in broken time-reversal symmetry (see section 2). We expect Ce_2PdGe_6 may have a similar magnetic structure as Ce_2CuGe_6 , though at present its magnetic structure remains unverified.




Figure 1: The magnetic structure of Ce_2CuGe_6 , published in [1]

In order to rigorously investigate whether these materials may be altermagnets, their Hall resistivities will be measured to verify the reproducibility of the AHE. Other physical and magnetic properties, such as magnetic susceptibility and electrical resistivity, will also be verified. Later on, neutron diffraction studies will be conducted to verify the magnetic structure of Ce_2PdGe_6 . This will provide further insight into the intrinsic contributions of AHE in this material. Angle Resolved Photo Emission Spectroscopy (APRES) will be used to map out of electron band structure of these materials, which is necessary to confirm whether they exhibit spin-split bands, a hallmark feature of altermagnets. Combined, these studies will help increase our confidence in the classification of this material, and may indicate that this material needs to be re-classified. This paper will focus on Ce_2PdGe_6 , and the study of Ce_2CuGe_6 will be continued by other lab members in the coming months.




Figure 2: The crystal structure of Ce_2PdGe_6 , which is in the orthorhombic Cmca space group. This figure was adapted from [3]

2 Background

2.1 Altermagnetism

A primary property of altermagnets is the lifting of energy degeneracy between spin up and spin down bands in their electronic structure without the need for spin-orbit coupling. In magnetic materials, time-reversal symmetry (TRS) is the symmetry operation that reverses magnetic moments and spins. In ferromagnets (FMs), TRS is broken because of the net magnetization since all the moments point in the same direction. Therefore, no symmetry operation relates the two opposite spin directions, prohibiting degeneracy between spin-up and spin-down electronic bands and resulting in spin-splitting. In conventional antiferromagnets (AFMs), the two magnetic sublattices have opposite spins, but they are related by symmetries such as inversion combined with time reversal or time reversal combined with a translation. These combined symmetries restore the equivalence of the two spin channels and protect spin degeneracy. Therefore, conventional antiferromagnets do not exhibit spin splitting in the absence of spin-orbit coupling [4].

In altermagnets, opposite-spin sublattices are not related by inversion or translation combined with time reversal. Instead, they are connected by other crystal symmetries, such as proper rotations, screw rotations, mirror reflections, or glide operations. These symmetries transform momen-

tum space in a way that allows spin-up and spin-down bands to split in a momentum-dependent manner. Because the sign of the spin splitting changes between different momentum directions, these materials are called *alttermagnets* [4]. One of the consequences of such spin-splitting is the capability of producing the anomalous Hall effect (AHE) despite zero net magnetization [5].

2.2 Anomalous Hall Effect

The anomalous Hall effect (AHE) is a phenomenon, originally associated with ferromagnetic materials, in which conduction electrons are deflected to one side as they move through the material, creating a transverse voltage. Unlike the ordinary Hall effect, this deflection occurs in the absence of any external magnetic field, because it arises from the material's intrinsic magnetic structure [6]. Nowadays, it is understood that AHE can arise in antiferromagnetic materials as well. The appearance of AHE, alone, in a material with zero magnetization is not sufficient proof to conclude that the material is an *alttermagnet*; however, it does indicate broken TRS. Previous studies reporting an anomalous Hall conductivity of $100\Omega^{-1}cm^{-1}$ in Ce_2PdGe_6 [1].

3 Methods

3.1 Single Crystal Growth

Single crystals of Ce_2PdGe_6 were grown using the flux method with Bi as the flux. To prevent cerium from oxidizing, the elements were weighed in the glove box. High-purity cerium powder (thermo scientific, 99.9%), palladium granules (Alfa Aesar, 99.8%), germanium pieces (thermo scientific, 99.9999%), and bismuth pieces (thermo scientific, 99.9%) were loaded into a 2 mL aluminum oxide crucible in a molar ratio of 2:1:6:10. The total mass of all elements combined within the crucible was 5.1204 g. The crucible was then loaded into a silica-glass ampoule along with quartz wool and sealed under vacuum. The ampoule was then loaded into the MTI KSL-1200X muffle furnace and heated to 1150 °C over the course of 20 hours. The temperature was held 1150 °C for 12 hours, then cooled to 650 °C over 100 hours. Once cooling was complete, the ampoule

was removed from the furnace and centrifuged to separate the Bi flux from the single crystals. A razor was used to clean the crystals, which were then polished, to remove any excess Bi flux. The crystals obtained were flat and plate-like, with an average dimensions around $3 \times 3 \times 1 \text{ mm}^3$.




Figure 3: Picture of Ce_2PdGe_6 single crystal, grown using the flux-method. The sample shown here has dimensions $5 \times 5 \times 1 \text{ mm}^3$.

3.2 Powder X-ray Diffraction and Structural Characterization

Powder x-ray diffraction (PXRD) was performed, using the Rigaku MiniFlex instrument, to determine the phase of the crystals. A clean, polished single crystal was ground into a powder using a mortar and pestle. A fifteen-hour scan was performed to investigate the crystal phase, and the SmartLab Studio II search matching software was used to confirm the structure.

3.3 Magnetic Measurements

The sample was mounted onto a quartz holder using GE varnish. First, the Photonic Science x-ray Laue Crystal Orientation System was used to determine the orientation of the single crystal. Then, the 7-T Quantum Design Machine Property Measurement System (QD MPMS3) was used to measure magnetization and magnetic susceptibility. Magnetic susceptibility was measured over a temperature range of 1.8-300 K. Magnetization was measured under an applied magnetic field with a range of $\pm 7 \text{ T}$.

3.4 Thermal and Transport Measurements

The sample was polished into a Hall bar geometry, with a length:width ratio of 3:1. The four-probe method was used; Pt wires, 25 μm in diameter, were attached to the sample using silver epoxy. The epoxy was cured inside of the MTI KSL-1200X muffle furnace for 12 minutes at 110°C. The sample was then mounted to a rotator sample board using GE varnish, and the platinum wires were soldered to the contact pads. Electrical resistivity and Hall resistivity were measured using the 9-T Quantum Design Dynacool Physical Property Measurement System (QD PPMS). Data were collected over a temperature range of 1.8-200 K.

4 Results

4.1 Magnetic Susceptibility

The temperature dependence of magnetic susceptibility along [010] is shown in Figure 4. A small magnetic field of 100 Oe was applied along [010]. These data exhibit typical AFM behavior, with a kink at $T_N = 11.3$ K. Below T_N , the magnetic susceptibility approaches zero as the temperature approaches zero, suggesting that the material orders antiferromagnetically below 11.3 K. This value is in agreement with previously reported Néel temperature values for this compound [1]. The inverse susceptibility plot, shown in Figure 5, exhibits a linear trend for high temperatures, above 150 K. The Curie constant ($C = 1.42 \text{ cm}^3 \cdot \text{K} \cdot \text{mol}^{-1}$ per formula unit) and Weiss temperature ($T_0 = -38.2$ K) were computed after applying line of best fit to this high-temperature, paramagnetic region ($T > 150$ K). A negative Weiss temperature is expected for an antiferromagnet. The calculated effective moment for the Ce^{3+} ions is $2.38 \mu_B/\text{Ce}$. This value is in agreement with the theoretical value ($2.54 \mu_B/\text{Ce}$) and typical experimental value ($2.40 \mu_B/\text{Ce}$) [7, Table 1, p. 306].

4.2 Magnetization

Figure 6 shows a plot of magnetization with an applied magnetic field along [010] ranging from ± 7 T. The magnetization plots exhibit plateaus for $B < -2$ T and $B > +2$ T, indicating




Figure 4: Magnetic susceptibility over a temperature range of 1.8 K to 300 K. These measurements were taken parallel to [010], and with a constant applied magnetic field of 100 Oe. The plot labeled FC corresponds to magnetic susceptibility measured using field-cooled protocol; ZFC corresponds to magnetic susceptibility measured using zero-field-cooled protocol.




Figure 5: Inverse magnetic susceptibility over a temperature range of 1.8-300 K. These measurements were taken parallel to [010], and with a constant applied magnetic field of 100 Oe. These data were measured using field-cooled protocol.

that saturation occurs for fields larger than $|2 \text{ T}|$. For temperatures $< T_N$, there appears to be an additional plateau in the magnetization plot for $-2 \text{ T} < B < +2 \text{ T}$. This feature may indicate that the material has a non-trivial magnetic texture.

4.3 Electrical and Hall Resistivity

Under $H \parallel c$, a hysteresis is observed in field-dependent Hall resistivity measurements. This hysteresis is observed at 3 K (below T_N); at such a temperature, we expect extrinsic contributions to AHE, such as skew scattering and side-jumping, to be minimized. Extrinsic causes, therefore, likely do not account for this hysteresis. Instead, intrinsic causes that arise from the material's magnetic structure are suspected to be driving this observed AHE. The electrical resistivity clearly decreases below T_N , in agreement with the T_N from the magnetic susceptibility measurements. This follows the expectation that below T_N , once magnetic order has been established, scattering by the Ce atoms will be minimized.




Figure 6: Magnetization measured over a magnetic field range of ± 7 T. Data taken parallel to [010] at a set of six different representative temperatures, ranging from 1.8-20 K.




Figure 7: Electrical resistivity measured parallel to [110], over a temperature range of 1.8-300 K. A current was applied along the [110] axis.




Figure 8: Hall resistivity taken parallel to [110], over a magnetic field range of ± 2 T. Data were taken at 3 K

5 Conclusion

We have verified experimentally that Ce_2PdGe_6 is a strong candidate for altermagnetism. As discussed, this material exhibits antiferromagnetic ordering, with nontrivial magnetization plateaus, and also exhibits the AHE along [010], at temperatures below T_N . The fact that AHE arises within

this material suggests broken symmetries within its magnetic structure, which carries the implication that this material could be altermagnetic. However, in order to verify whether this material is an altermagnet, theoretical input is needed, such as density functional theory (DFT) calculations, and modeling of the electron band structure, to determine whether spin-splitting occurs.

In the future, the electron band structure of both Ce_2PdGe_6 and Ce_2CuGe_6 will be studied experimentally using APRES. Neutron diffraction studies, both under zero and non-zero magnetic field, will be conducted to determine the magnetic structure of Ce_2PdGe_6 . High-quality single crystals of Ce_2CuGe_6 have also been grown, and properties such as magnetization, magnetic susceptibility, electrical resistivity, and Hall resistivity will soon be measured to confirm reproducibility. Moreover, additional measurements for Ce_2PdGe_6 will be taken along [100] and [001]. Future studies might explore the origin of the additional plateau observed in the magnetization along [010].

References

- [1] H. Kotegawa, H. Tanaka, Y. Takeuchi, H. Tou, H. Sugawara, J. Hayashi, and K. Takeda, *Phys. Rev. Lett.* **133**, 106301 (2024).
- [2] H. Matsuda, R. Hibino, C. Tabata, K. Kaneko, N. Higa, T. Onimaru, H. Tanaka, H. Tou, H. Sugawara, J. Hayashi, K. Takeda, and H. Kotegawa, *Phys. Rev. B* **113**, 195116 (2026).
- [3] K. Momma and F. Izumi, *Journal of Applied Crystallography* **44**, 1272 (2011).
- [4] C. Song, H. Bai, Z. Zhou, L. Han, H. Reichlova, J. H. Dil, J. Liu, X. Chen, and F. Pan, *Nature Reviews Materials* **10**, 473 (2025).
- [5] L. Šmejkal, J. Sinova, and T. Jungwirth, *Phys. Rev. X* **12**, 040501 (2022).
- [6] N. Nagaosa, J. Sinova, S. Onoda, A. H. MacDonald, and N. P. Ong, *Rev. Mod. Phys.* **82**, 1539 (2010).
- [7] C. Kittel, *Introduction to Solid State Physics*, 8th ed. (John Wiley & Sons, Hoboken, NJ, 2005).

**Search for Nonresonant HH Production in the
 $b\bar{b}WW$ Single-Lepton Resolved Channel, and
Extended ggHH Signal Sample Production**

VANSHIKA GUPTA

2026 NSF/REU Program
Department of Physics
University of Notre Dame

ADVISOR(S): Prof. Abhisek Datta

Abstract. *Higgs boson pair (HH) production is one of the few processes at the LHC directly sensitive to the shape of the Higgs potential, and is the subject of an ongoing CMS search in the $b\bar{b}WW$ single-lepton final state. This report describes two intertwined contributions to that search carried out during the summer: a systematic study of how the output-node structure of the analysis’s neural networks affects its final sensitivity, and a private Monte Carlo production campaign that extends the ggHH signal sample the analysis is trained on. Twenty-four network configurations were trained and evaluated end-to-end; the four multi-class variants that produce final limits give median expected values in the range $26.1\text{--}29.5\times SM$, and the underlying discrimination, measured through the AUC-ROC of all 24 networks, shows that adding a single-Higgs node or splitting single-top into its own class each hurts sensitivity individually while cancelling partially when combined. The accompanying Monte Carlo campaign targets 2×10^6 generated ggHH events against the $\sim 10^5$ available centrally, and is running steadily on the Notre Dame Tier-3 cluster at the time of writing.*

I. INTRODUCTION

In the Standard Model, every fundamental particle acquires its mass through its interaction with the Higgs field. The 2012 discovery of the Higgs boson by the ATLAS and CMS collaborations [2, 3, 4] confirmed that this field exists, and a decade of measurements since has shown that the couplings of the Higgs boson to fermions and gauge bosons closely follow the Standard Model’s predictions. One central prediction remains untested: the Higgs field’s interaction with *itself*. The strength of this self-interaction sets the shape of the scalar potential—the “Mexican hat” whose form drives electroweak symmetry breaking—and it has never been measured directly. The self-coupling also has consequences beyond particle physics: it determines whether the electroweak vacuum is stable, and it enters models of the early universe’s phase transitions. The most direct experimental probe of the Higgs self-coupling is the production of *pairs* of Higgs bosons (HH). The self-coupling contributes to HH production at leading order, so measuring the HH cross section constrains the trilinear coupling modifier $\kappa_\lambda = \lambda_3/\lambda_3^{\text{SM}}$ and, through the vector-boson-fusion channel, the quartic coupling modifier κ_{2V} of the $HHVV$ vertex. The difficulty is that HH production is extraordinarily rare: its Standard Model cross section at $\sqrt{s} = 13.6$ TeV is roughly a thousand times smaller than that of single-Higgs production. No single measurement is yet sensitive enough to observe it, which is why the LHC experiments search for HH in many decay channels at once and statistically combine the results. The measurements described here are performed with the Compact Muon Solenoid (CMS), one of the two general-purpose detectors at the LHC. CMS surrounds the collision point with a silicon tracker, electromagnetic and hadronic calorimeters, and muon chambers, all inside a 3.8 T solenoid; a two-tier trigger reduces the 40 MHz collision rate to a recordable few kHz [1]. From these recorded events, a particle-flow algorithm reconstructs the individual leptons, jets, and missing transverse momentum that analyses like this one work with. This work contributes to the CMS search for HH in the $b\bar{b}WW$ decay channel [1], using the single-lepton final state in which one W boson decays to a charged lepton and a neutrino and the other decays to quarks. The contribution has two intertwined parts. The first is a systematic study of the analysis’s neural-network architecture: twenty-four network configurations were trained, evaluated, and—where applicable—carried through the full statistical fit, to determine how the choice of network output structure affects the final sensitivity. The second is a Monte Carlo production campaign: because the analysis trains its networks on simulated signal events, and because the centrally available ggHH sample contains only $\sim 10^5$ events, a private production extending that sample to 2×10^6

events was set up on the Notre Dame Tier-3 cluster and is running at the time of writing. The remainder of this report describes the physics background common to both efforts (Sec. II), the methods used in each (Sec. III), and the results obtained (Sec. IV).

II. BACKGROUND

A. Higgs pair production

At the LHC, HH production proceeds mainly through two mechanisms. Gluon–gluon fusion (ggF) dominates, and at leading order receives contributions from the two diagrams of Fig. 1: a “box” diagram in which two gluons couple to a top-quark loop that radiates two Higgs bosons, and a “triangle” diagram in which the top loop produces a single off-shell Higgs that splits into two through the self-coupling. The box amplitude scales with the square of the top Yukawa modifier κ_t^2 , while the triangle scales with $\kappa_t \kappa_\lambda$. Crucially, the two diagrams interfere *destructively*: the Standard Model HH cross section is small precisely because these two contributions largely cancel. This destructive interference is also what makes the process such a sensitive probe—any deviation of κ_λ from its Standard Model value upsets the cancellation and changes both the rate and the kinematics of HH production. The second mechanism is vector-boson fusion (VBF), shown in Fig. 2, in which




Figure 1. Leading-order Feynman diagrams for HH production via gluon–gluon fusion. Left: the box diagram, proportional to κ_t^2 . Right: the triangle diagram, in which an off-shell Higgs splits into two through the trilinear self-coupling, proportional to $\kappa_t \kappa_\lambda$. The two amplitudes interfere destructively in the Standard Model.

the incoming quarks radiate vector bosons that fuse into a Higgs pair. Although its cross section is roughly twenty times smaller than ggF, VBF is the only production mode with leading-order sensitivity to the quartic coupling κ_{2V} , through the contact diagram of Fig. 2 (right). The three VBF amplitudes also cancel almost exactly at the Standard Model point, so even modest deviations in κ_{2V} produce large enhancements of the VBF HH rate. The VBF topology is experimentally distinctive: the two scattered quarks appear as forward jets with a large rapidity separation and a large invariant mass. Among the possible HH decay modes, $HH \rightarrow b\bar{b}WW$ has the second-largest branching fraction, about 25%. In the single-lepton (SL) topology studied here, one W decays to $\ell\nu$ and the other to a quark pair, giving a final state of one lepton, missing transverse momentum from the neutrino, two b jets, and two light-quark jets. When the Higgs boson decaying to $b\bar{b}$ has low transverse momentum, the two b quarks are reconstructed as separate small-radius (AK4) jets—the




Figure 2. Leading-order Feynman diagrams for HH production via vector-boson fusion, with amplitudes proportional to $\kappa_V \kappa_\lambda$ (left), κ_V^2 (middle), and κ_{2V} (right). As in ggF, the diagrams interfere destructively at the Standard Model point.

resolved topology that this work targets. The main obstacle is the enormous $t\bar{t}$ background, whose decay products closely mimic this final state, followed by single-top (tW), W +jets, QCD multijet events with misidentified leptons, and single-Higgs production. The last of these is small in yield but kinematically Higgs-like, making it uniquely difficult to separate from the signal.

B. Why simulation matters

Everything in this analysis rests on simulated collisions. The neural networks that separate signal from background are trained entirely on Monte Carlo (MC) events; the templates fit to data to extract the upper limit are MC histograms; and the systematic uncertainties are propagated by re-evaluating those histograms under varied conditions. The simulation chain mirrors nature step by step: a matrix-element generator (POWHEG [5] at NLO for ggHH, MADGRAPH5_AMC@NLO [6] for VBF) draws the hard-scattering event from the theory’s probability distribution; PYTHIA [7] showers and hadronises the outgoing partons; and GEANT4 [8] simulates the passage of every particle through the CMS detector material. The output is a sample of events statistically indistinguishable, in principle, from real data—but with the true process identity attached to each event, which is what makes supervised training possible. This reliance on simulation is also the analysis’s Achilles heel: the precision of every template is limited by the number of simulated events behind it. The centrally produced ggHH sample for the Winter24 campaign contains only $\sim 10^5$ events, and once those events are divided among analysis categories and network score bins, the per-bin signal statistics become a leading source of uncertainty. This is the motivation for the production campaign described in Sec. III: enlarging the ggHH sample twentyfold so that the next iteration of the analysis is not statistics-limited on the signal side.

C. Deep neural networks as event classifiers

The tool this analysis uses to separate a rare signal from overwhelming backgrounds is a deep neural network (DNN). A DNN is a nested sequence of linear transformations and elementwise nonlinearities: an input vector x of event-level quantities (jet and lepton four-momenta, b -tagging scores, missing transverse momentum, angular correlations) is mapped through “hidden layers” of the form $\mathbf{h}_{k+1} = \text{ReLU}(W_k \mathbf{h}_k + \mathbf{b}_k)$, and a final *softmax* layer converts the last hidden representation into a probability for each output class. The weights W_k are adjusted by stochastic gradient descent to minimise the cross-entropy between the network’s predicted class probabilities and the true

(MC) labels. Trained this way, the network learns the correlations in the high-dimensional feature space that distinguish, say, a ggHH event from a $t\bar{t}$ event—correlations far beyond what any single kinematic cut could capture. Two properties of such classifiers matter for this work. First, the choice of *output classes* is a design decision: one can train a network to separate signal from an undifferentiated “background”, or give each major background process its own output node in the hope that the network learns each one’s signature more precisely. Whether finer output structure actually helps is not obvious *a priori*—each additional node divides the network’s capacity and the training sample—and answering that question quantitatively is the core of this work’s analysis contribution. Second, the quality of a classifier can be summarised independently of any single working point by the area under its receiver-operating-characteristic curve (AUC-ROC): the probability that a randomly chosen signal event receives a higher score than a randomly chosen background event. AUC-ROC is used throughout Sec. IV as the common yardstick across network configurations.

III. METHODS

A. Analysis regions and the fit pipeline

Following Ref. [1], events in the resolved SL region are grouped by AK4 jet multiplicity ($= 3$ or ≥ 4) and b -tag multiplicity ($= 1$ or ≥ 2). Within each group, a multi-class DNN assigns every event to one of five output nodes—ggHH, qqHH, $t\bar{t}$, tW , or W +jets—and the two HH nodes serve as signal regions while the three background nodes serve as control regions that constrain the background normalisations in the fit. A second-stage binary DNN sharpens the signal–background separation within each HH-enriched region. Combined with the four Run 3 data-taking eras (2022, 2022EE, 2023, 2023BPix), this yields 80 resolved analysis regions. The path from network scores to a physics result runs through five stages: DNN inference on every event; construction of per-era datacards binned in the network score; hierarchical combination of the four eras; combination with the $b\bar{b}\tau\tau$ and $b\bar{b}ZZ$ channels; and an asymptotic CL_s fit with the CMS COMBINE tool [10] that returns the 95% CL upper limit on the signal strength $\mu_{HH} = \sigma(HH)/\sigma_{HH}^{\text{SM}}$. As part of this work the era-combination stage was validated end-to-end, in the course of which a process-consistency bug—a background with zero yield, and hence no histogram, in one era but not the others—was identified, reported, and fixed upstream.

B. The multi-model scan

To test how the network’s output structure affects sensitivity, a driver was written that builds, trains, and evaluates 24 DNN configurations in a single run. All share the same backbone—a KERAS multilayer perceptron with three hidden layers of 256 nodes, ReLU activations, and a softmax (or, for binary heads, sigmoid) output—and the same input features as Ref. [1]. Two structural choices are varied independently: whether the single-top process receives its own output node (tWST) or is merged into the residual background (tW), and whether single-Higgs production receives its own node (the _H suffix) or not. Crossing these with the head type (binary_ggHH, binary_qqHH,

multi-class) and the jet-multiplicity region ($3j, 4j$) gives

$$3 \times 2 \times 2 \times 2 = 24 \text{ configurations.}$$

Each is trained with a 5-fold cross-application scheme, so every quoted metric carries a statistical uncertainty from the spread across the five passes. Only the four multi-class configurations feed the fit and yield an upper limit; the twenty binary heads are compared through AUC-ROC.

C. Extending the ggHH signal sample

The MC production campaign runs the same simulation chain as the central sample—the POWHEG ggHH gridpack, CMSSW_13_3_0 for the GEN and SIM steps, the CT18NNLO PDF error set [9], and PYTHIA shower weights—so that the private events are directly compatible with the analysis configuration. The target is 2×10^6 events, produced in tasks of 200 events each on the Notre Dame Tier-3 cluster (glados), orchestrated by the LOBSTER workflow-management framework [11] with workers submitted through a `work_queue_factory` on HTCondor. Each task launches `cmsRun` inside a CMSSW EL8 Apptainer container and stages its `SIM.root` output over XRootD to the Notre Dame storage element. Running LOBSTER’s Python-3 branch on the cluster’s current EL9 environment required a set of environment and stage-out fixes: sanitising the worker environment so that grid-middleware Python paths do not shadow the CMSSW software stack, pinning a binary-compatible `numpy` for the gridpack’s own tooling, making the container launch robust to the worker-node configuration, and repairing several corner cases in LOBSTER’s stage-out and reporting code. One property of the site shaped the design: worker tasks execute inside bind-mount-restricted containers from which most shared filesystems are invisible, even when writable from the login node. Stage-out therefore goes over the network via XRootD—using the double-slash absolute-path URL form the storage element requires—rather than through a mounted filesystem.

IV. RESULTS

A. The four multi-class configurations

The four multi-class configurations, fit with $3j$ and $4j$ combined, produce the median expected upper limits of Table 1. The nominal configuration—no extra nodes—is the most sensitive of the

Table 1. Combined ($3j + 4j$) median expected 95% CL upper limits on the HH signal strength μ_{HH} , for the four multi-class DNN configurations.

MODEL	EXTRA NODES	MEDIAN EXPECTED UL
multi_ggHH_qqHH_ttbar_tW_WJets	none	26.13
multi_ggHH_qqHH_ttbar_tWST_WJets	tW split	28.75
multi_ggHH_qqHH_ttbar_tW_WJets_H	H only	29.50
multi_ggHH_qqHH_ttbar_tWST_WJets_H	tW split + H	27.94

four, and adding either extra node degrades the limit. The ordering is not monotonic in the number of extra nodes: the H-only configuration is the worst of the four, yet adding the tW split *on top of* the H node recovers most of the loss. Four points are too few to interpret on their own, so the same

structural knobs are next tracked across all 24 configurations, using AUC-ROC as the common yardstick introduced in Sec. II.

B. Pass-to-pass and model-to-model variance

Because each configuration is trained five times, every metric carries a genuine uncertainty. The median standard deviation of AUC-ROC across the five passes of a given model is 0.0010, with a maximum of 0.0038 over all 24 models; the mean AUC-ROC across models spans 0.7640–0.9132, a range of ~ 0.15 . Model-to-model differences therefore exceed pass-to-pass noise by roughly two orders of magnitude: the rankings that follow reflect real structural differences, not training fluctuations. Figure 3 shows all 24 configurations ranked by mean AUC-ROC.




Figure 3. All 24 DNN configurations ranked by mean AUC-ROC. Error bars are the standard deviation across the five training passes. Binary heads (blue) sit above the multi-class configurations (navy) simply because two-class AUC is not comparable to ≥ 3 -class AUC; the ordering *within* each family is what the following subsections analyse.

C. Effect of the two extra nodes

For every matched pair of configurations differing only in the H node—same head type, signal target, tW -split choice, and jet multiplicity—the change in AUC-ROC from adding the node is computed; the analogous scan is performed for the tW split. The results appear in Figs. 4 and 5. Both knobs consistently degrade the discrimination: the H node shifts the mean AUC by up to ~ -0.04 and the tW split by up to ~ -0.02 , both large against the $\lesssim 0.004$ pass-to-pass noise of Sec. IV.B. The two effects are not independent. The H-node penalty is most severe when the tW




Figure 4. Change in mean AUC-ROC induced by adding the single-Higgs output node, computed as a matched pair for each combination of head type, signal target, tW -split choice, and jet multiplicity.

background is *not* already split off ($\Delta\text{AUC} \approx -0.04$) and much smaller when it is ($\Delta\text{AUC} \approx -0.01$); symmetrically, the tW -split penalty of ≈ -0.025 without the H node shrinks to ≈ 0 with it. Each extra node hurts most in isolation and is partially absorbed by the other when both are present. A physical reading is that single-Higgs and tW both share topology with $t\bar{t}$ -like backgrounds: when only one extra node exists, it must describe phase space it is not matched to, and borrows discriminative capacity from the signal nodes to do so; when both exist, each specialises. This interaction is precisely what produces the non-monotonic ordering of Table 1, with the “both extras” configuration sitting between the two “single extra” ones rather than below them. Consistently, the ordering of the four upper limits in Table 1 reproduces the ordering of the corresponding multi-class networks by mean AUC-ROC, supporting the use of the 24-model scan as a proxy for the structural effects that four limits alone are too few to establish.




Figure 5. Change in mean AUC-ROC induced by splitting the tW background into its own output node, computed as a matched pair for each combination of head type, signal target, H-node choice, and jet multiplicity.

D. Status of the extended signal sample

The MC production campaign that will lift the signal-statistics limitation identified in Sec. II is running steadily at the time of writing, with several hundred concurrent workers on the Notre Dame Tier-3. Order 10^5 events have been produced through the full GEN-SIM chain so far. Individual SIM.root files open cleanly in ROOT and contain the expected 200 events each with the full tree structure (MetaData, ParameterSets, Parentage, Events, LuminosityBlocks, Runs). Completion of the full 2×10^6 -event target is projected on a one-to-two-week timescale at the current sustained worker count, after which the SL resolved networks will be retrained on the enlarged sample—closing the loop between the two halves of this work by testing directly how much of the current upper limit is driven by signal-MC statistics.

V. SUMMARY

Two intertwined contributions to the CMS $b\bar{b}WW$ SL resolved analysis [1] were made during the summer. A 24-configuration scan of the analysis’s DNN architecture showed that its rankings are statistically meaningful—pass-to-pass AUC noise of $\sim 10^{-3}$ against a model-to-model spread of $\sim 10^{-1}$ —and that adding a single-Higgs output node or splitting the tW background each degrades the discrimination individually while cancelling partially in combination. This interaction explains the non-monotonic ordering of the four multi-class upper limits ($26.1\text{--}29.5 \times \text{SM}$) and identifies the nominal, no-extra-nodes configuration as the most sensitive. In the course of the same work,

the era-combination stage of the datacard pipeline was validated and a process-consistency bug fixed upstream. In parallel, a private ggHH GEN-SIM campaign extending the central $\sim 10^5$ -event sample to 2×10^6 events was brought into steady production on the Notre Dame Tier-3 cluster, with completion projected within one to two weeks. The immediate next steps are to finish the production, retrain the SL resolved networks on the enlarged sample, and finalise the systematic-uncertainty treatment in preparation for the approval review.

REFERENCES

- [1] CMS Collaboration, “Search for Higgs boson pair production in the $b\bar{b}WW$ single-lepton channel in pp collisions at $\sqrt{s} = 13.6$ TeV,” *CMS-PAP-HIG-25-022* (2026).
- [2] ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC,” *Phys. Lett. B* **716**, 1 (2012).
- [3] CMS Collaboration, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC,” *Phys. Lett. B* **716**, 30 (2012).
- [4] CMS Collaboration, “Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV,” *JHEP* **06**, 081 (2013).
- [5] G. Heinrich *et al.*, “NLO predictions for Higgs boson pair production with full top-quark mass dependence matched to parton showers,” *JHEP* **08**, 088 (2017).
- [6] J. Alwall *et al.*, “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations,” *JHEP* **07**, 079 (2014).
- [7] C. Bierlich *et al.*, “A comprehensive guide to the physics and usage of Pythia 8.3,” *SciPost Phys. Codebases* **8**, 1 (2022).
- [8] GEANT4 Collaboration, “GEANT4—a simulation toolkit,” *Nucl. Instrum. Meth. A* **506**, 250 (2003).
- [9] T.-J. Hou *et al.*, “New CTEQ global analysis of quantum chromodynamics with high-precision data from the LHC,” *Phys. Rev. D* **103**, 014013 (2021).
- [10] CMS Collaboration, “The CMS statistical analysis and combination tool: COMBINE,” *Comput. Softw. Big Sci.* **8**, 19 (2024).
- [11] A. Woodard *et al.*, “Exploiting Volatile Opportunistic Computing Resources with Lobster,” *J. Phys. Conf. Ser.* **664**, 062018 (2015).

Investigating the Electronic Properties of $\text{Cr}_{0.5}\text{Ga}_{0.5}\text{Te}_2$

JAMES HARRIS

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR: Prof. Kateryna Foyevtsova

Abstract

As a particularly attractive member of a growing class of two-dimensional van der Waals ferromagnets, $\text{Cr}_{0.5}\text{Ga}_{0.5}\text{Te}_2$ is a semiconducting material which exhibits partial occupancies, rendering the true structure unknown. VASP simulations using Density Functional Theory are used to provide insight into how varying the Cr/Ga arrangements affect the electronic properties of the structures. Strong candidates for explaining the insulating behavior include Cr clusters which generate stable Te dimers. More work must be done to identify structures which match the stoichiometry constraints while replicating the experimentally observed insulating behavior.

1 Introduction

$\text{Cr}_{0.5}\text{Ga}_{0.5}\text{Te}_2$ is a newly synthesized two-dimensional (2D) van der Waals (vdW) ferromagnetic material. The two structures of which it is theoretically comprised, CrTe_2 and GaTe_2 both display metallic behavior (see Fig. 1), but $\text{Cr}_{0.5}\text{Ga}_{0.5}\text{Te}_2$ has been experimentally determined to be a semiconductor with a band gap of ≈ 0.3 eV [1]. The reason for this insulating behavior is unknown, but the structure of the material might provide some insight.




Figure 1: Density of states plots obtained by DFT VASP simulations for GaTe_2 and CrTe_2 , respectively, indicating that they are metallic.

1.1 Structure

Through X-ray diffraction techniques, it has been determined that the structure of the material consists of three atomic layers that are repeated throughout the crystal (see Fig. 2). The middle layer is a triangular lattice of Cr and Ga atoms where each lattice site is occupied by either a Cr or Ga atom. Over the entire structure (or equivalently, one unit cell), the Cr:Ga ratio is 1:1. This middle layer is sandwiched between two layers of Te atoms. Cr atoms are octahedrally coordinated by Te, while Ga atoms are tetrahedrally coordinated. As a result, the bottom Te layer is fully occupied and remains unchanged for different Cr/Ga arrangements, whereas the top Te layer varies significantly with the arrangement of Cr and Ga atoms. Specifically, changing the Cr/Ga arrangement in the middle layer




Figure 2: Three example layers of the material. Blue spheres represent Cr atoms, green spheres represent Ga, and yellow spheres represent Te. Half-filled spheres represent partial occupancies. Material image courtesy of Nayak, A. et al. [1].

alters the bonding pattern of the top Te layer. The goal is to identify Cr/Ga arrangements that produce a structure with the overall 1:1:4 stoichiometry while examining how the Te atoms bond, for different Te bonding configurations could affect how the electrons move through the system. Ultimately, the goal is to identify Cr/Ga arrangements that produce upper-layer Te bonding in a

way that replicated the insulating behavior we know the material to have while maintaining the 1:1:4 stoichiometry.

Although it is trivial to construct a unit cell with the correct overall stoichiometry from arbitrary Cr and Ga arrangements, we generally restrict our efforts to clusters of Cr surrounded by Ga, or vice versa. It has been proposed that these clustered arrangements "trap" electrons within them and reduce electrical conductivity.

1.2 Motivation

2D vdW ferromagnetic semiconductors belong to a class of materials which is quickly growing in number and in promise; its properties of polarized spin states alongside an electronic band gap lend itself well to applications in next-generation electronics such as spintronic devices. However, knowing that this material belongs to such a class is hardly useful without knowing why it displays these properties. Better understanding the behavior of the system will also help development of future, similar materials which might have even more attractive properties.

2 Background

$\text{Cr}_{0.5}\text{Ga}_{0.5}\text{Te}_2$ is a crystalline solid, so, at the microscopic level, it exhibits a periodic lattice structure; that is, the entire material can be modeled as a unit cell which is repeated in all directions. Thus, the position $\mathbf{R}$ of any atom in the structure is given by

$$\mathbf{R} = n_1 \mathbf{a}_1 + n_2 \mathbf{a}_2 + n_3 \mathbf{a}_3. \quad (1)$$

Here, n_i are integers scaling the primitive vectors $\mathbf{a}_i$ that describe the unit cell. It turns out, however, that working in what is known as the reciprocal lattice space is more computationally efficient. The primitive vectors can be related to the wave vectors, which are periodic to the lattice:

$$\mathbf{k} = \frac{m_1}{N_1} \mathbf{b}_1 + \frac{m_2}{N_2} \mathbf{b}_2 + \frac{m_3}{N_3} \mathbf{b}_3, \quad (2)$$

where m_i are integers and N_i are chosen such that $N = N_1 N_2 N_3$ (N is the total number of atomic sites in the material), and from which we can use a complete set of basis functions to build the wavefunction

$$\psi(\mathbf{r}) = \sum_{\mathbf{k}} c_{\mathbf{k}} e^{i(\mathbf{k} \cdot \mathbf{r})}, \quad (3)$$

a construction which is made valid by Bloch's Theorem.

Relying heavily on the topics just discussed, Density Function Theory (DFT) is foundational to areas of theoretical chemistry and physics in which it is necessary to solve the Schrödinger equation in order to gain information on lattice structures, bond lengths, density of states, band structures, and other electronic and magnetic properties of molecules or systems. The simplest form of the time-independent Schrödinger equation is given by $\hat{H}\psi = E_n\psi$, where the Hamiltonian operator $\hat{H}$ acts on the wavefunction ψ to yield the energy eigenstates E_n .

For a practical system, solving this equation becomes a many-body problem which is impossible to analytically solve; this can be avoided by considering the electron density function, which is given by:

$$\rho(\mathbf{r}) = 2 \sum_i |\psi_i(\mathbf{r})|, \quad (4)$$

where $\mathbf{r}$ describes the three-dimensional coordinate of the particle, typically an electron, and the factor of 2 accounts for spin. Using this electron density, we can write the Kohn-Sham equations:

$$\hat{H}_{KS}\psi_i = \left(\frac{-\hbar^2}{2m} \nabla^2 + \frac{e^2}{4\pi\epsilon_0} \int \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d^3\mathbf{r}' - \frac{e^2}{4\pi\epsilon_0} \sum_j \frac{Z_j}{|\mathbf{r} - \mathbf{R}_j|} + V_{xc}(\mathbf{r}) \right) \psi_i(\mathbf{r}) = \epsilon_i \psi_i(\mathbf{r}), \quad (5)$$

where the first term accounts for the kinetic energy, the second term describes the Coulombic repulsion between electrons, the third term gives electron-nucleus interaction, and the fourth term is given by

$$V_{xc}(\mathbf{r}) = \frac{\delta E_{xc}[\rho]}{\delta \rho(\mathbf{r})}, \quad (6)$$

denoting a functional derivative. The Kohn-Sham equations' usage of the electron density function means that they have to be solved in only three dimensions, so they are much easier to solve than

the many-body Schrödinger equation. The key to solving these equations, however, lies in that final term of the Hamiltonian, the exchange-correlation potential. It relies on $E_{xc}[\rho]$, which is a functional, or a function of a function, giving Density Functional Theory its name, and it is not known exactly. Approximations can be made, and it is these approximations which determine what kind of DFT calculation is being run.

Two theorems by Kohn and Hohenberg ensure that solving these equations gives a meaningful and physical result. The first states that the ground-state energy from Schrödinger's equation is a unique functional of the electron density, and the second that the electron density that minimizes the energy of the overall functional is the true electron density corresponding to the full solution of the Schrödinger equation. Knowing this, we can be sure that when the energy is minimized, we have found the electron density, and thus solved the Schrödinger equation, but how do we go about doing that? DFT calculations use an iterative, self-consistent method:

1. Guess a trial electron density.
2. Build V_{xc} based on the chosen exchange-correlation functional approximation and the trial electron density.
3. Solve the Kohn-Sham equations.
4. Construct the wavefunctions $\psi_i(\mathbf{r})$.
5. Rebuild the electron density based on the resulting wavefunctions.
6. Compare the new electron density with that of the last iteration, and repeat to a given error tolerance.

Again, by the Kohn-Hohenberg theorems, once the energy is minimized, we can be certain that we have successfully solved the Schrödinger equation and can use the information gained from that to analyze data such as free energy, magnetization, and density of states.

3 Methods

The search began with the smallest unit cells which can fit clusters in the triangular lattice. Only one such unit cell was found for clusters of four atoms (see Fig. 3(a)), and isolated clusters smaller

than four are ruled out. This four-site cluster produces a stoichiometry of 1:1:4.75. A stoichiometry of 1:1:4.6 is possible with a five-site cluster (see Fig. 3(b)), and with a six-site cluster (which is largely though not completely isolated), we are able to reach a stoichiometry of 1:1:4.5. Larger cells could achieve lower stoichiometry, but running calculations on these cells would increase computational resources, so we limit our focus to unit cells containing six-site clusters. Due to the




Figure 3: A four- (a), five- (b), and six-site (c) Ga cluster and their respective Te arrangements.

low stoichiometry of the six-site cluster, this structure was chosen as the candidate on which to run DFT calculations which were done using VASP [2]. First, "ENCUT" convergence tests were ran to identify the required accuracy at which to conduct these calculations moving forward; an ENCUT value of 320 was chosen, and KPOINTS were set at 8, 8, 14. Then, a static calculation was ran to determine some characteristics such as the free energy, the density of states, and stability of magnetic moments. Finally, a full ionic/unit-cell structural relaxation set to 50 ionic steps was ran to observe the forces acting on each ion. This same process was used to observe the properties and stability of other structures, including antiferromagnetic arrangements. The data obtained from these runs are out of the scope of this report and merely support the conclusions drawn from analysis of the six-site cluster candidate.

4 Results

Among the structures observed, the aforementioned six-site cluster (Structure "A") remains the most promising candidate. A plot of its density of states (see Fig. 4) demonstrates its insulating behavior, matching the experimental data. At a free energy of -167.25631000 eV (before structural relaxation), it is the lowest-energy structure that was tested. Its antiferromagnetic counterpart (with the center three Cr atoms spin down and the corner three Cr atoms spin up) had a free energy of -166.94437422 eV. This higher energy supports the true structure's ferromagnetic nature. The




Figure 4: The density of states plot for structure "A" after a full structural relaxation. The Fermi level is set at 0 eV, where an upward line is drawn to demonstrate the absence of energy states; ie, a band gap.

structural relaxation, however, was key to identifying the differences between structure "A" and the true structure (see Fig. 5). By studying how the Te atoms move and the bonds shift, we can notice arrangements that are unstable. Most noticeably, long chains of Te atoms which are bonded together move in such a way that the Cr atoms' octahedral coordination with its surrounding Te atoms, which

we expect to be very stable, are heavily disturbed. In general, these Te chains in any configuration violently separate, after which the energy is lower, indicating that these Te attachments are unstable and are not what the true structure contain. Te dimers, on the other hand, do not experience such a dramatic effect, merely lengthening when exposed to shifting Te chains, supporting that these are in fact stable and preferred by the system.




Figure 5: A full structural relaxation of structure "A". Te chains experience a net force and move, disrupting stable Cr octahedral coordination while Te dimers appear stable.

5 Conclusion

Despite the promise of structure "A", it features an incorrect stoichiometry (1:1:4.5 as opposed to the desired 1:1:4), so it is not the true structure of the material; as of yet, no studied structure features the correct stoichiometry. However, from studying these "wrong" structures, much has still been learned about the stabilities and preferences of the system. Most relevant to the stoichiometry is that a bigger unit cell produces a more accurate stoichiometry. In fact, a unit cell with a seven-site cluster has been identified which has a stoichiometry of 1:1:4.42857... (1:1:31/7). Furthermore, there is strong evidence that Te dimers are stable while Te chains or zig-zags are unstable. Though this narrows down the search significantly, what Cr/Ga arrangements cause Te dimers vs chains is still being investigated. Looking forward, creating Cr/Ga arrangements which lead to stable Te dimers while also lowering the stoichiometry will produce even better candidates for the true structure of $\text{Cr}_{0.5}\text{Ga}_{0.5}\text{Te}_2$, and hopefully, this structure can illuminate the mechanism by which this material exhibits its semiconducting properties.

6 Acknowledgments

I would like to thank Dr. Garg and Kristen Amsler for seamlessly running the show and providing a fantastic REU for everybody.

I also am grateful to the College of Science for their Summer Undergraduate Research Fellowship which financially supported my ability to participate in this REU, and I thank the Notre Dame Center for Research Computing for giving me the resources to complete this project.

Finally, I extend many thanks to my advisor, Professor Kateryna Foyevtsova for providing me the opportunity to conduct this research as well as for consistently guiding and supporting my efforts.

References

- [1] Nayak, A.; Jamaluddin, S.; Wu, F.; Yu, Z.; others Cr_{0.5}Ga_{0.5}Te₂: A near-room temperature semiconducting 2D van der Waals ferromagnet (2026)
- [2] G. Kresse and J. Hafner, Phys. Rev. B 47 , 558 (1993); *ibid.* 49 , 14 251 (1994).
- [3] G. Kresse and D. Joubert, Phys. Rev. 59 , 1758 (1999)
- [4] Density Functional Theory: A Practical Introduction, 2nd edition; Sholl, D.; Steckel, J. Wiley Publishing (2009)
- [5] Solid State Physics; N. Ashcroft, and N. Mermin. Holt-Saunders, (1976)
- [6] Williams, T., Kelley, C. Gnuplot 6.0: an interactive plotting program (2024)

Halogen-Based Selective Etching of MAX Phases: Routes to Halogen-Terminated MXenes

HU TINGHAO

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Laszlo Forro

Abstract

We systematically compare three halogen-based etching routes—direct gas-phase, *in-situ* gas generation, and molten salt—for producing halogen-terminated MXenes from four MAX phase precursors (Ti_3AlC_2 , Ti_2AlC , V_2AlC , Nb_2AlC) using five halogen sources. FeBr_3 and FeCl_3 give the best results on Ti_3AlC_2 , while I_2 etches only Ti_2AlC . Direct Br_2 gas yields poor products due to uncontrolled kinetics and over-etching to graphitic carbon. The gas-etched MXenes exhibit a larger (002) *d*-spacing than HF-etched MXene, consistent with the larger ionic radius of halogen terminations. The CuBr_2 system reveals a temperature-dependent Cu reduction pathway: at 300 °C, Cu^{2+} reduces to CuBr ; at 650 °C, to metallic Cu^0 , driven by Br_2 consumption by Al. Ti-based MAX phases are more susceptible than V- and Nb-based phases, with an optimal temperature window of 550–650 °C. These findings establish FeBr_3 and FeCl_3 as practical fluoride-free halogen sources and demonstrate that matching halogen identity to the MAX phase is critical for selective etching.

1 Introduction

Since the discovery of $\text{Ti}_3\text{C}_2\text{T}_x$ in 2011, MXenes—two-dimensional transition metal carbides and nitrides—have attracted intense interest for energy storage, electromagnetic shielding, and catalysis[1]. They are produced by selective etching of the A-element (typically Al) from MAX phase precursors ($\text{M}_{n+1}\text{AX}_n$). The standard approach—aqueous HF or LiF/HCl etching—has dominated for over a decade, but carries two drawbacks: hazardous fluoride waste, and a mixed $-\text{F}/-\text{O}/-\text{OH}$ surface termination that is difficult to control, limiting the precision of surface-sensitive property engineering.

These limitations have motivated fluoride-free, gas-phase etching routes. Three approaches have emerged: direct halogen gas etching, *in-situ* gas generation from heated metal halide salts (e.g., CuBr_2 , FeCl_3)[2], and molten salt etching where the halide salt serves as both halogen source and liquid reaction medium.

Several fundamental questions remain unresolved: the relative effectiveness of different halogens (Cl vs. Br vs. I) has not been systematically compared; the role of the MAX phase transition metal (Ti, V, Nb) and layer count ($n = 1, 2$) is unexplored; and the temperature window for selective etching is unknown for most halogen–MAX combinations.

Here, we systematically investigate four MAX phase precursors Ti_3AlC_2 , Ti_2AlC , V_2AlC , Nb_2AlC with five halogen sources: Br_2 , I_2 , CuBr_2 , FeBr_3 , and FeCl_3 . We use XRD, Raman spectroscopy, and XPS to track structural transformation and surface chemistry. We find that FeBr_3 and FeCl_3 give the best results on Ti_3AlC_2 ; I_2 etches Ti_2AlC but not other MAX phases, highlighting the importance of matching halogen to the MAX phase precursor.

2 Methods

MAX phase powders were used as received: Ti_3AlC_2 , Ti_2AlC , V_2AlC , and Nb_2AlC . Bromine (Br_2 , liquid), iodine (I_2 , solid), copper(II) bromide (CuBr_2 , powder), iron(III) chloride (FeCl_3 , powder), and iron(III) bromide (FeBr_3 , powder) were used as halogen sources.

All reactions were performed in sealed quartz tubes evacuated to 10^{-5} – 10^{-6} bar. Two tube sizes were used: 10 mm outer diameter (8 mm inner) and 6 mm outer diameter (4 mm inner). Reagents were either mixed in a single 10 mm tube or placed in a nested configuration with the halogen source in the smaller inner tube and the MAX phase in the outer tube (Fig. 1).



Figure 1: (a): MAX phase powder and halogen source loaded together in a single tube, used for direct Br_2 gas-phase and FeCl_3 molten salt etching. (b): MAX phase powder in the outer tube and halogen source in a semi-open inner tube, used for I_2 direct etching and *in-situ* gas-generation etching with CuBr_2 , FeCl_3 , and FeBr_3 .

Direct gas-phase etching. For Br_2 etching, liquid Br_2 was frozen in liquid nitrogen in a 10 mm quartz tube, a MAX phase pellet was added, and the tube was evacuated and sealed. For I_2 etching, solid I_2 was placed in a 6 mm semi-open inner tube with the MAX phase in the 10 mm outer tube, then evacuated and sealed. Tubes were heated at 450–850 °C for 6–24 h. The mass ratio of MAX to halogen was approximately 1:2 to 1:2.5.

***In-situ* gas-generation etching.** The metal halide salt (CuBr_2 , FeCl_3 , or FeBr_3) was placed in

a 6 mm inner tube with the MAX phase in the 10 mm outer tube, evacuated, and sealed[2]. Tubes were heated at 450–900 °C for 18 h, with a MAX-to-halide molar ratio of approximately 1:4 to 1:5. Control experiments with CuBr_2 + metal powders (Ti, Al, Zn) at 300–650 °C confirmed that both a reducing metal and sufficient temperature are required for Cu reduction to metallic Cu^0 .

Molten salt etching. MAX phase powder was mixed with FeCl_3 powder, loaded into a 10 mm quartz tube, evacuated, and sealed. Tubes were heated at 550–650 °C for 18 h. At these temperatures, FeCl_3 (m.p. 306 °C) forms a liquid reaction medium. After reaction, products were rinsed with deionized water to remove residual iron salts. The molar ratio of MAX to FeCl_3 was approximately 1:4 to 1:5.

Characterization. X-ray diffraction (XRD) patterns were collected on a Bruker D8 Advance diffractometer with Cu $K\alpha$ radiation ($2\theta = 5\text{--}80^\circ$, step 0.02° , 0.2 s/step). Etching was assessed semi-quantitatively by tracking the (002) peak shift from $2\theta \approx 10^\circ$ (MAX) to $6\text{--}9^\circ$ (MXene).

Raman spectroscopy (Witec Alpha 300R, -300 to 2500 cm^{-1}) identified halogen-specific M–X modes and reduction byproducts (metallic Cu, CuBr).

X-ray photoelectron spectroscopy (XPS, monochromatic Al $K\alpha$, 1486.6 eV) collected survey spectra and high-resolution scans of Ti 2p, Al 2p, C 1s, and halogen regions to probe surface terminations and chemical states.

3 Results

The (002) interlayer d -spacing of the gas-etched MXenes is larger than that typically reported for HF-etched MXene,[3] consistent with the larger ionic radius of halogen terminations ($-\text{Br}$, $-\text{Cl}$, $-\text{I}$) compared to $-\text{F}$, $-\text{O}$, and $-\text{OH}$. XRD patterns confirm a clear downshift of the (002) reflection from $2\theta \approx 9.5^\circ$ (MAX phase) to $5\text{--}7^\circ$ (MXene), indicating an expansion of the interlayer spacing upon Al removal (Fig. 2).

The three best-performing combinations were $\text{FeCl}_3 + \text{Ti}_3\text{AlC}_2$ (*in-situ*), $\text{FeBr}_3 + \text{Ti}_3\text{AlC}_2$ (*in-situ*), and $\text{I}_2 + \text{Ti}_2\text{AlC}$. In each case, XRD showed a pronounced (002) peak shift and near-complete disappearance of the MAX (104) reflection near $2\theta \approx 39^\circ$, and Raman spectra collected at room




Figure 2: XRD results.

temperature over $0\text{--}600\text{ cm}^{-1}$ displayed MXene-characteristic vibrational features (Fig. 3).

X-ray photoelectron spectroscopy (XPS) was performed to characterize the surface termination chemistry. Survey spectra revealed a substantial oxygen content on the surface of all etched products, although the origin of this oxygen remains unclear. Halogen-specific terminations ($-\text{Cl}$, $-\text{Br}$, $-\text{I}$) were detected in the corresponding samples, confirming that the gas-phase etching route successfully installs halogen surface groups (Fig. 4). The $\text{I } 3d$ signal was weak relative to Cl and Br , consistent with the poorer etching efficacy of I_2 and suggesting only partial Al removal and termination in the iodine-treated samples. Detailed elemental quantification is summarized in Fig. 5.

Control experiments with CuBr_2 revealed a temperature-dependent Cu reduction pathway. At $300\text{ }^\circ\text{C}$, CuBr_2 decomposed to release Br_2 with reduction of Cu^{2+} to Cu^{1+} , forming CuBr . At $650\text{ }^\circ\text{C}$, reduction proceeded fully to metallic Cu^0 .

Direct Br_2 gas etching gave poor results: the reaction initiated violently at the onset of heating, leading to inhomogeneous products (which are mostly graphite). I_2 etching was more difficult than Br - or Cl -based etching overall: at $550\text{ }^\circ\text{C}$, I_2 produced negligible etching across all MAX phases except Ti_2AlC . Ti-based MAX phases were generally more susceptible to halogen etching than V- and Nb-based MAX phases.




Figure 3: Raman spectrum of Ti_3AlC_2 etched with FeCl_3 at 650°C for 18 h. MXene-characteristic vibrational features appear in the low-frequency region ($200\text{--}600\text{ cm}^{-1}$), peaks around 1300 and 1600 cm^{-1} correspond to graphite. Other sharp peaks might be caused by cosmic rays.




Figure 4: XPS spectra of halogen-etched MXenes ($\text{Cl } 2p$, $\text{Br } 3d$, $\text{I } 3d$).

O1s	C1s	Al2p	Cl2p	Ti2p	Fe2p3	Br3d	I3d5	RSE	Corrected RSE	
0.733	0.314	0.256	0.954	2.077	1.964	1.149	6.302	RSE		
137.694	58.985	56.575	202.691	428.503	401.834	248.452	1347.178	Corrected RSE		
438.80	145.50		230.78	2,675.81			3,490.89			
3.19	2.47	0.00	1.14	6.24	0.00	0.00	13.04			
24.44%	18.92%	0.00%	8.73%	47.90%	0.00%	0.00%	0.00%	1.00		CL
366.90	177.39		1,492.08		364.45		2,400.82			
2.66	3.01	0.00	0.00	3.48	0.00	1.47	0.00	10.62		
25.09%	28.32%	0.00%	0.00%	32.79%	0.00%	13.81%	0.00%	1.00		BR
634.30	171.52	57.54		1,981.63			244.54	3,089.53		
4.61	2.91	1.02	0.00	4.62	0.00	0.00	0.18	13.34		
34.54%	21.80%	7.63%	0.00%	34.67%	0.00%	0.00%	1.36%	1.00		I

Figure 5: Elemental composition (at.%) of etched MXenes from XPS survey analysis.




Figure 6: Raman spectrum of Ti_3AlC_2 etched with CuBr_2 at 900°C for 18 h, showing MXene-characteristic features in the $200\text{--}600\text{ cm}^{-1}$ region from the *in-situ* gas-generation route.

4 Discussion

FeBr_3 and FeCl_3 both used in the *in-situ* gas-generation configuration were the most effective halogen sources for Ti_3AlC_2 . The CuBr_2 *in-situ* route also yielded positive results, though less extensive than the iron halides.

The performance of I_2 was strongly MAX-phase-dependent: it was ineffective for Ti_3AlC_2 , V_2AlC , and Nb_2AlC largely inert below 650°C yet effective etching was achieved with Ti_2AlC . This unexpected selectivity suggests that the $n=1$ structure of Ti_2AlC may present a lower kinetic barrier to iodine-based Al extraction than the $n=2$ Ti_3AlC_2 or the V- and Nb-based MAX phases.

The CuBr_2 system revealed a temperature-dependent Cu reduction pathway: at 300°C , Cu^{2+} reduces to Cu^{1+} (forming CuBr), while at 650°C reduction proceeds to metallic Cu^0 . The Br_2 released during CuBr_2 decomposition is consumed by Al in the MAX to form AlBr_3 ; this consumption of Br_2 drives further decomposition of CuBr to metallic Cu, with higher temperature favoring the forward reaction. The CuBr_2 route offers practical advantages over direct Br_2 gas: solid CuBr_2 is easier to handle than liquid Br_2 , the *in-situ* configuration provides self-metered halogen supply, and the metallic Cu byproduct serves as an internal marker for reaction progress. Direct Br_2 gas gave poor results because the reaction initiated too vigorously, leading to inhomogeneous products. Moreover, excess Br_2 can over-etch the Ti layers of the MAX phase, producing graphitic carbon

(Fig. 7).




Figure 7: Raman spectrum of Ti_3AlC_2 etched with direct and excess Br_2 gas at 650 °C for 18 h, showing characteristic graphitic carbon bands indicative of Ti over-etching.

The MAX phase identity strongly modulates etching susceptibility, and the effect is halogen-dependent. Ti_3AlC_2 responded best to FeCl_3 and FeBr_3 , while Ti_2AlC was the only phase etched effectively by I_2 . V_2AlC and Nb_2AlC proved more resistant across all halogen sources tested. This inverse selectivity between halogen types and MAX phases suggests that iodine follows a different rate-limiting step, potentially related to the larger ionic radius of I^- and its weaker Lewis acidity toward the M–Al bond.

5 Conclusions

We have systematically compared halogen-based etching routes: direct gas-phase, *in-situ* gas generation from metal halide salts, and molten salt etching across four MAX phase precursors (Ti_3AlC_2 , Ti_2AlC , V_2AlC , and Nb_2AlC). FeBr_3 and FeCl_3 , both deployed in the *in-situ* configuration, are the most effective halogen sources for Ti_3AlC_2 , outperforming direct Br_2 gas and CuBr_2 . Direct Br_2 gives poor results because the reaction initiates too vigorously and excess Br_2 over-etches the Ti layers to produce graphitic carbon. The *in-situ* configuration avoids this by providing a self-metered halogen supply. Ti_2AlC can be etched with I_2 , demonstrating that iodine-based etching is viable when matched to the appropriate MAX phase, despite being ineffective for Ti_3AlC_2 , V_2AlC ,

and Nb₂AlC. The CuBr₂ system reveals a temperature-dependent Cu reduction pathway: at 300 °C, Cu²⁺ reduces to CuBr, while at 650 °C reduction proceeds to metallic Cu, driven by consumption of released Br₂ by Al in the MAX phase. The gas-etched MXenes exhibit a larger (002) *d*-spacing than HF-etched MXene, attributable to the larger ionic radius of halogen surface terminations. Ti-based MAX phases are generally more susceptible to halogen etching than V- and Nb-based phases. These results establish halogen-based gas-phase etching as a viable fluoride-free strategy for MXene synthesis, in which the halogen identity serves as a tunable “chemical knob” for controlling surface termination chemistry.

References

- [1] M. Naguib, M. Kurtoglu, V. Presser, J. Lu, J. Niu, M. Heon, L. Hultman, Y. Gogotsi, and M. W. Barsoum, “Two-dimensional nanocrystals produced by exfoliation of Ti₃AlC₂,” *Adv. Mater.*, vol. 23, no. 37, pp. 4248–4253, 2011.
- [2] J. Zhu, S. Zhu, Z. Cui, Z. Li, S. Wu, W. Xu, T. Ba, Y. Liang, and H. Jiang, “Solvent-free one-step green synthesis of MXenes by ‘gas-phase selective etching’,” *Energy Storage Mater.*, vol. 70, 103503, 2024.
- [3] M. Alhabeb, K. Maleski, B. Anasori, P. Lelyukh, L. Clark, S. Sin, and Y. Gogotsi, “Guidelines for synthesis and processing of two-dimensional titanium carbide (Ti₃C₂T_x MXene),” *Chem. Mater.*, vol. 29, no. 18, pp. 7633–7644, 2017.

Shielding Optimisation for the Neutron Irradiation Station

COREN HUGHES

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

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

Abstract

The Neutron Irradiation Station (NIS) at the University of Notre Dame produces neutrons through the ${}^7\text{Li}(p, n)$ reaction for nuclear-astrophysics experiments. Neutron interactions can increase detector background and damage the germanium crystals. This project corrected and extended an existing Geant4 simulation of the NIS, producing a flexible model capable of testing different shielding materials and geometries. A shielding prototype was then designed and optimised using systematic simulation studies. Across detector angles from 30° to 120° , the shielding reduced the neutron count by an average of 99.93%. Preliminary simulations at angles above 140° showed similar neutron-reduction efficiency; however, the increased external gamma background requires further investigation. The completed simulation provides a reliable platform for continued shielding optimisation and experimental development.

1 Introduction

The Neutron Irradiation Station (NIS) at the University of Notre Dame is part of the 5U particle accelerator beamline. It consists of a ${}^7\text{Li}$ target mounted on a water-cooled copper block. When protons from the beamline strike the lithium target, neutrons are produced through the ${}^7\text{Li}(p, n)$ reaction and used to bombard a secondary target [1].

The motivation for producing a neutron beam at the NIS is to study neutron-induced reactions and excited nuclear states under controlled experimental conditions. These measurements provide nuclear data relevant to nuclear structure and nuclear-astrophysics applications. Although reactor facilities can produce substantially higher neutron fluxes, the NIS offers the advantage of a tunable neutron energy. By varying the incident proton energy, the emitted neutron energy can be adjusted for a particular measurement.

Neutron-scattering reactions such as $X(n, n')X^*$ can selectively populate excited states in a target nucleus. The gamma rays emitted as these states decay are measured using a germanium detector, providing information on the nuclear level structure and allowing the lifetimes of short-lived excited states to be determined.

The main challenge is that neutrons are difficult to shield. Since they have no electric charge, they interact primarily through nuclear reactions rather than electromagnetic interactions. Neutrons can therefore scatter from and activate materials surrounding the experiment, producing secondary

radiation that contributes to the detector background.

This is particularly important at the NIS, where a germanium detector measures gamma rays emitted from the secondary target. Germanium detectors provide excellent energy resolution, but neutrons can interact with both the detector and nearby materials. Activation can produce additional gamma radiation, while neutrons entering the germanium crystals can cause displacement damage, degrading charge collection and detector resolution over time.

Shielding must therefore reduce neutron and gamma-ray backgrounds while protecting the germanium detector from damage. Simulating the design before construction allows different materials and geometries to be tested without the cost and risk of building each configuration experimentally.

This project centred on developing and using a detailed Geant4 simulation of the NIS. The existing model was assessed, corrected, and extended to represent the experimental geometry, particle transport, detector response, and relevant physical processes accurately. The improved simulation was then used to test and optimise shielding materials and geometries, with the objective of reducing the neutron flux reaching the detector without introducing additional gamma-ray background.

2 Methods

2.1 Developing the Geant4 Simulation

An existing Geant4 simulation of the NIS, developed by previous REU students, was used as the starting point for this project. Geant4 is a particle-transport toolkit developed at CERN [2]. The model contained several issues, including an inverted geometry, incorrect neutron scattering, missing radioactive-decay processes, and limited control over simulation parameters. These were corrected before the simulation was used to design the shielding.

The existing GDML geometry was inverted and imported into Geant4. The beamline, lithium target, secondary target, detector, and experimental room were then repositioned to match the physical setup. Their dimensions and separations were checked to maintain a one-to-one scale. A simplified representation of the reconstructed geometry is shown in Fig. 1.




Figure 1: Simplified geometry of the NIS. Protons strike the ^7Li target, producing neutrons that irradiate a secondary target. The emitted gamma rays are measured using a germanium detector positioned at 90° . The figure is not to scale.

Geant4 messenger commands were added so that the detector, target, source, and shielding could be controlled through macro files, allowing systematic scans without recompiling the code.

The simulation physics was also validated by comparing the built-in `Shielding_HP` and a custom physics list. Both produced the expected spectra after troubleshooting, but the custom list was faster and was therefore used.

2.2 Simulation Configuration

The neutron source was positioned at the lithium target and configured to represent neutrons produced through the $^7\text{Li}(p, n)$ reaction. Primary neutrons were emitted isotropically within a 10° forward cone towards the detector.

The simulation recorded histograms of the energies of neutrons and gamma rays entering the germanium detector, together with the total energy deposited in each of its crystals. Gamma rays produced inside and outside the detector were recorded separately, allowing the effect of secondary gamma-ray production within the shielding and detector materials to be investigated.

2.3 Designing the Shielding

With the Geant4 simulation corrected and fully functional, the shielding could be systematically designed and optimised. The initial configuration was based on previous shielding studies [3; 4]. Borated high-density polyethylene (HDPE) was placed around the neutron source to moderate fast neutrons through elastic scattering with hydrogen. Once slowed, the neutrons are more readily captured by ^{10}B , which has a large absorption cross section for low-energy neutrons.

A tungsten shadow shield was then positioned between the source and detector. Tungsten was selected because its high density and high atomic number allow a relatively compact shield to strongly attenuate radiation travelling directly towards the detector. The detector was surrounded by additional borated HDPE with an inner lead lining to reduce the remaining neutron and gamma-ray backgrounds. A simplified representation of this design is shown in Fig. 2.

Initial comparisons showed that source shielding produced the largest reduction in neutron background, whereas detector shielding alone had little effect and increased the gamma-ray background. This trend was consistent across the neutron energies tested. The dimensions and positions of the main shielding components were then varied systematically at a fixed neutron energy, beginning with the source shielding, to determine an approximately optimal geometry.




Figure 2: Simplified initial shielding design. Borated HDPE moderates neutrons near the lithium target and around the detector, while tungsten and lead reduce the radiation reaching the germanium detector. The detector is positioned at 90° to the secondary target. The figure is not to scale.

3 Results

3.1 Source Shielding Optimisation

The tungsten was placed as close to the ^7Li target as possible to minimise the required volume. However, increasing the borated HDPE radius moved the tungsten further from the target, requiring a larger block to cover the same line of sight to the detector.

The radius of the borated HDPE was varied from 3 to 7 inches, with the tungsten positioned directly outside it. Increasing the radius substantially reduced the neutron count, although the improvement became much smaller beyond approximately 4 inches. The external gamma-ray count remained relatively low throughout the scan. A radius of 4 inches was therefore selected as a compromise between shielding performance and the required tungsten volume.




Figure 3: Total neutron count as a function of the borated HDPE radius surrounding the lithium target over 5 million runs

Triangular, parallelogram, and trapezium tungsten cross sections were then compared. The triangular design produced the lowest neutron count, while the trapezium design produced the lowest gamma-ray count. As the difference in neutron performance was small and the trapezium design required considerably less tungsten, it was selected.

Increasing the wedge length reduced both neutron and gamma-ray backgrounds up to approximately 4.5 inches, after which the improvement became small. A similar trend was found for the wedge height above approximately 4 inches. The final wedge therefore had a length of 4.5 inches, a height of 4 inches, a width of 3 inches, and a 60° cutaway.




Figure 4: Top down view and dimensions of the final tungsten wedge.

3.2 Detector Shielding Optimisation

The size, length, and position of the detector shielding were varied. Increasing its length gave the largest reduction in both neutron and gamma-ray backgrounds, while changes in size and position produced smaller improvements.

Several borated HDPE and lead configurations were also tested. A front block gave the largest neutron reduction but slightly increased the gamma-ray background. A copper funnel was therefore added to shield the front of the detector, as copper can attenuate low-energy secondary radiation while producing less characteristic radiation than lead. Tests showed that copper outside the tungsten shadow scattered neutrons towards the detector, so the funnel was redesigned to remain within this shadow before flaring outward. An 8-inch-long funnel with a 2-inch flare was selected.

3.3 Final Shielding Performance

The final shielding was compared with the unshielded geometry at 60°, 90°, and 120°, using 180 million primary neutrons per configuration. As shown in Fig. 5, the unshielded neutron count decreased with increasing detector angle. With shielding, the count was lowest at 90°, falling from 8.79×10^6 to 2.5×10^3 , a reduction of 99.97%. It increased slightly at 120°, but remained several orders of magnitude below the unshielded case.




Figure 5: Total neutron counts with and without shielding at each detector angle.

The deposited-energy spectrum at 120° is shown in Fig. 6. At this angle, the shielding reduced the total deposited energy by 99.87% and the number of energy-deposition events by 99.93%.




Figure 6: Energy deposited in the detector with and without shielding at 120°.

4 Discussion

4.1 Interpretation of the Final Shielding Design

The largest reduction came from the high- Z tungsten shadow shield. Since tungsten is expensive, the final design required a balance between its size and the borated HDPE surrounding the source. Detector-side shielding had a smaller effect and could produce secondary radiation or scatter neutrons towards the germanium crystals. The copper funnel was therefore kept within the tungsten shadow before flaring towards the detector.

The shielding was most effective at 90° , where it had been optimised. Its weaker performance at 60° and 120° suggests that the tungsten shield and copper funnel may need repositioning for other angles.

4.2 Limitations and Uncertainty

The same initial random seed was used for each comparison, but particle histories still diverge and the results remain statistical. This is especially important for shielded cases with only 10^2 – 10^4 recorded events. The overall reductions are reliable, although small differences between similar designs may be affected by statistical uncertainty.

The optimisation also varied individual components while keeping others fixed, so coupled effects may have been missed. Further uncertainty arises from simplified geometry, materials, detector components, and neutron-source distributions. The final design should therefore be tested using higher statistics, a more detailed model, additional parameter combinations, and experimental measurements.

4.3 Future Work

Future work should compare additional shielding materials and combinations using the same neutron, gamma-ray, and deposited-energy metrics. Depleted uranium could also be investigated as an alternative shadow-bar material.

The shielding geometry should be adapted for detector angles above 140°. A preliminary compact design achieved comparable neutron reduction, but increased the external gamma-ray background, requiring further optimisation of the tungsten shadow bar.

5 Conclusion

An improved Geant4 simulation was used to develop a prototype shielding design for the NIS. The final geometry combined borated HDPE around the lithium target, a tungsten shadow shield, detector-side lead and borated HDPE, and a copper funnel.

At 90°, the design reduced the neutron count by 99.97%. At 120°, the deposited energy and number of energy-deposition events were reduced by 99.87% and 99.93%, respectively. Further work is required to adapt the geometry for a wider range of detector angles.

The project successfully implemented a fully functional and flexible Geant4 simulation of the NIS, capable of testing a wide range of shielding geometries and materials. This provides a reliable platform for future optimisation and development of the shielding design.

References

- [1] M. Matney, A. Simon, D. Robertson, K. Manukyan, A. Aprahamian, D. Bardayan, R. deBoer, E. Stech, W. Tan, and M. Wiescher, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **1075**, 170406 (2025).
- [2] S. Agostinelli *et al.*, Nuclear Instruments and Methods in Physics Research Section A **506**, 250 (2003).
- [3] J. C. Hopkins, J. T. Martin, and J. D. Seagrave, Nuclear Instruments and Methods **56**, 175 (1967).
- [4] J. I. W. Watterson, N. K. Mumba, and J. L. Makhabane, Nuclear Geophysics **9**, 333 (1995).

Recharacterization of the High Pressure POint-like Jet Gas Target (HIPPO)

ELIZABETH JACOBY

2026 NSF/REU Program
Department of Physics
University of Notre Dame


ADVISOR(S): Prof. Manoel Couder

Abstract

Understanding (α, γ) radiative capture reactions is important for the study of stellar helium burning and nucleosynthesis. A high pressure point-like helium gas target called HIPPO is used in conjunction with the St. George recoil separator at the Nuclear Science Laboratory of the University of Notre Dame to study (α, γ) reactions. After changes to the geometry of HIPPO, the point-like characteristic of the target needed to be investigated. If the helium target is larger than the original design, it will cause more energy to be lost by the beam and could complicate the recoil separator rejection. To determine the HIPPO target size, the $E_{\text{CM}} = 1606$ keV resonance of the $^{20}\text{Ne}(\alpha, \gamma)^{24}\text{Mg}$ reaction was measured. The goal of this experiment was to discover whether HIPPO deviated from its point-like character and produced an enlarged gas target. This paper reports that, while the exact size and jet properties require further investigation, HIPPO still produces a narrow target and not an extended gas cloud.

Introduction

Alpha-gamma (α, γ) radiative capture is a crucial step in nucleosynthesis and energy production during stellar helium burning. (α, γ) reactions occur when elements capture helium nuclei and release excess energy in the form of γ -rays. In nuclear physics, such reactions can be written as:



The Nuclear Science Laboratory of the University of Notre Dame possesses a High Pressure Point like helium gas target called HIPPO which is specially designed to study (α, γ) reactions at low energies. HIPPO sits on the St. George recoil mass separator beamline so that after the ion beam captures α -particles from HIPPO, these compound “recoil” nuclei can be separated from the ion beam—and any other contaminants—according to their masses. It is important that this gas jet is extremely narrow and “point-like” so that all recoil particles originate from the same place, which allows for optimal recoil separation.

HIPPO Design

HIPPO creates a supersonic gas jet which is injected downward into a chamber and then pumped out by a vacuum “catcher” located 8 mm below the injection nozzle (see Fig. 1.a). The helium gas is then purified and recirculated. Since its installation in 2012, changes have been

made to the geometry of HIPPO: the supports of the injection nozzle and catcher were extended, the shape of the support of the injection nozzle was changed from cylindrical symmetry to rectangular symmetry, and the catcher platform and nozzle container were attached together (see Fig. 1.b).




Figure 1: (a) Original design of HIPPO. (b) Updated design of HIPPO design.

A measurement taken in January 2026 reported nearly twice as much energy lost by the ion beam in the gas target than was predicted by *Kontos et al.*[1]. This behavior is consistent with an extended target, and it is feared that, due to the changes in geometry, HIPPO no longer produces a point-like gas jet but rather an extended gas target. An extended gas target would complicate measurements with HIPPO. A jet gas target produces recoil particles from nearly the exact same point in the gas. In an extended target, however, reactions can occur at any point in the gas “cloud.” This complicates the recoil separator rejection, since recoils no longer originate from the same point.

To test whether HIPPO produces a compact gas jet or an extended target, a measurement of the well characterized $^{20}\text{Ne}(\alpha, \gamma)^{24}\text{Mg}$ reaction was undertaken. The goal was to compare the γ -ray yields between measurements taken with and without shielding the face of the detector. If

the target was a compact jet, there should be a decrease in yield in the shielded measurement, but the shape of the yield curve (the γ -ray yield in function of beam energy) should remain unchanged; however, if the target was extended, the yield curve should noticeably decrease and change shape. This change occurs because reactions are not contained to the opening in the shielding but are taking place throughout the extended target and the γ -ray intensity is attenuated by the tungsten shielding.

Methods

This experiment was performed on the 5U accelerator and consisted of a ^{20}Ne ion beam of charge state 4+ and a helium gas target produced by HIPPO with injection pressure of 1000 millibar. A NaI detector was positioned outside of HIPPO to measure the γ -rays produced by the reaction. The detector was surrounded by lead shielding to reduce cosmic ray and beam induced background radiation (see Fig. 2). The 1606 keV (in the center of mass reference frame) resonance of the $^{20}\text{Ne}(\alpha, \gamma)^{24}\text{Mg}$ reaction was scanned. A resonance scan is the process of measuring a reaction at multiple beam energies near the resonance energy to map out certain characteristics of a resonance (such as width and strength).




Figure 2: Experimental setup.

To test whether HIPPO produced a 2 mm jet (as characterized in *Kontos et al.*[1]) or an extended target of a few centimeters, this measurement was performed twice: once with nothing shielding the face of the NaI detector and again with tungsten bricks covering its face with a 5 mm opening where HIPPO should be if it was a jet. If the target is a 2 mm jet, most of the γ -rays will reach the detector through the 5 mm gap in the shielding even with the tungsten reducing the solid angle of the detector. However, if the target is extended, $^{20}\text{Ne}(\alpha, \gamma)^{24}\text{Mg}$ reactions can occur at many locations and the subsequent γ -rays will be attenuated by the tungsten.



Figure 3: Gamma rays blocked by the tungsten shielding with point-like gas target (left) compared with extended gas target (right).

Magnesium created via $^{20}\text{Ne}(\alpha, \gamma)^{24}\text{Mg}$ reaction populates many excited nuclear states. However, it primarily de-excites by transitioning through the first excited state of ^{24}Mg and emitting a 1369 keV γ -ray. This makes the 1369 keV γ -ray the dominant signal detected by the NaI. The yield of the 1369 keV γ -ray was calculated and compared across the measurements with and without tungsten shielding. Yield is defined as the number of (α, γ) reactions per ^{20}Ne beam particle. The yield was calculated by summing all the γ -ray counts inside the 1369 keV peak and dividing the sum by the number of beam particles. This calculation was repeated for each measurement taken at a different beam energy and a yield curve was created.

$$Yield = \frac{N_{\gamma}}{N_{beam}} \quad (1)$$

Results




Figure 4: Gamma ray Spectra of reaction zoomed in on 1369 keV γ -ray. Both spectra from $E_{\text{beam}} = 1604.97$ keV (center of mass). (a) without shielding (b) with tungsten shielding.

Fig. 4 is an example of the spectra of the 1369 keV γ -ray used to calculate the yield. The decrease in counts in Fig. 4.b when compared with Fig. 4.a can be attributed to the reduction in the solid angle of the detector caused by the tungsten shielding which results in fewer γ -rays seen by the detector.

Analysis

To calculate the yield of the 1369 keV γ -ray, the peak was fitted with a gaussian plus a linear term and integrated. A linear background term was then subtracted from the measurements taken without the tungsten shielding. This integral was plugged into Eq. (1) and plotted against the beam energy of that measurement to generate a yield curve (see Fig 5.). This yield curve was then fit with the following equation:

$$Yield = A \left[\arctan\left(\frac{E - E_r}{\Gamma_1}\right) - \arctan\left(\frac{E - E_r - \Delta E}{\Gamma_2}\right) \right] \quad (2)$$

Eq. (2) is an analogy to the ideal yield curve expected for a narrow resonance in which ΔE is the energy that the beam loses in the target, Γ_1 and Γ_2 relate to width of the resonance, and A relates to the resonance strength [2]. However, because this measurement deals with the absorption from the tungsten, ΔE , Γ_1 , Γ_2 , and A are treated as free parameters in the fit.




Figure 5: Yield curve for the 1369 keV gamma as a function of beam energy. An uncertainty of 1 keV (Lab) in the energy.

To estimate the width of the gas target, the ratio between the two yield curves was calculated. This ratio provides the transmission of the 1369 keV γ -ray through the 5 mm gap in the tungsten as a function of beam energy. A Geant4 simulation, developed by A.T. Sanchez, provided the transmission of the 1369 keV γ -rays through a 5 mm gap based on the location that they originated from in an extended gas target. This simulation was used to translate the transmission from a function of energy to a function of position.




Figure 6: (a) Geant4 simulation of γ -ray transmission as a function of location of γ -ray emission. (b) HIPPO transmission data as a function of γ -ray emission position.

At Full Width at Half Maximum (FWHM), the target thickness was found to be approximately 1.6 cm. Unfortunately, using the transmission simulation does not provide a unique solution and will always yield a target thickness of 1.6 cm. This analysis method does not best reflect the data and was a deadend.

A new method to determine whether HIPPO produces a narrow or extended gas target was undertaken. This method involved comparing the widths of the shielded and unshielded yield curves. The width of the yield curve relates to the energy lost by the beam, ΔE from Eq. (2), which is a direct reflection of the target thickness. Specifically, the widths of the two yield curves at FWHM were compared. To make this comparison, the yield curves in Fig. 6 were scaled so that their FWHM aligned.




Figure 7: Yield curves scaled so that FWHM align.

Yield Curve	E_{low} (keV)	E_{high} (keV)	FWHM (keV)
Unshielded	1601.50 ± 0.17	1607.90 ± 0.17	6.4 ± 0.24
Shielded	1601.75 ± 0.17	1608.30 ± 0.17	6.55 ± 0.24

Table 1: Energy locations of FWHM of the unshielded and tungsten shielded yield curves.

Discussion

Fig. 7 and Table 1 demonstrate the widths of each yield curves at FWHM are nearly identical. These “matching” widths imply that, even with the tungsten shielding, a significant amount of the γ -rays are seen by the detector. As a result, it can be concluded that only a small fraction of the (α, γ) reactions are blocked by the tungsten. Therefore, the He gas target is narrow enough that most reactions occur within the 5 mm opening in the tungsten.

Conclusion

This experiment demonstrated that HIPPO still produces a narrow gas target on the order of a few millimeters; however, further analysis is required to establish the exact width of the target

and whether it remains a jet. Explaining the increased energy lost by the beam in the target (reported in January 2026) should be prioritized in future analysis, since the extended target hypothesis is no longer a viable explanation

Future work should also include correcting for the hysteresis of the 5U magnets, since the visible shift in energy between the two yield curves in Fig. 7 is due to the effects of hysteresis. Additionally, accounting for horizontal error would improve fits of the yield curves. If these efforts prove unsuccessful, a new measurement with increased statistics may be required. This will allow for a better mapping of the resonance—especially for the tungsten shielded measurement since the tungsten reduces the γ -ray statistics. Most importantly, a new method of analysis to unfold the true thickness of the target should be found. Knowing the exact size of the HIPPO target is important for the accuracy of St. George recoil rejection. Further characterization of HIPPO will improve measurements made with St. George and, consequently, impact nuclear astrophysics measurements employing gas targets and recoil separation.

References

- [1] 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).
- [2] C. Iliadas, “Nuclear Physics of Stars,” 2nd Edition, (Hoboken, New Jersey); Wiley-VCH, June, 2015.

The Chemical Abundances of The Hercules Ultrafaint Dwarf Galaxy

SAVANNAH JONES

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Evan Kirby and Lauren Henderson

Abstract

In this study, we present our findings on the chemical abundances of the Hercules ultrafaint dwarf galaxy. Using Keck-DEIMOS spectra with the 1200B diffraction grating, we obtained Fe, α , and neutron-capture abundances for 27 stars in Hercules. To reduce the data, we used the PyPeIt program, in contrast to the previously used reduction software, spec2d. We also employed chemfit, a program designed to fit model spectra to observed data. With these two programs, we were able to measure the α elements Mg, Ca, Si, and Ti. Most notably, we found the Mg abundances of 18 stars within the galaxy. Our abundance patterns of magnesium match previously published trends of ultrafaint dwarf galaxies. We were also able to measure the neutron-capture element Ba for 14 stars within the galaxy. Hercules stars have low abundances of Ba. This implies that the galaxy did not retain enrichment from many r -process events, which is common among ultrafaint dwarf galaxies.

1 Introduction

Galactic archaeology refers to the subfield in astrophysics that studies the history and evolution of galaxies by analyzing the properties of old stars within the galaxy. The chemical abundances of stars reveal when the galaxy formed stars and what nucleosynthetic events happened while stars were forming. Dwarf spheroidal galaxies and ultrafaint dwarf galaxies (UFDs) are well-suited for galactic archaeology because they contain relatively old populations of stars that are chemically close to the Big Bang. These galaxies in the Local Group contain old stellar populations, which have experienced some of the earliest supernovae in the universe. These events contributed to the stellar populations' chemical enrichment through nucleosynthesis. This chemical enrichment can originate from a few processes, such as different types of supernovae and star mergers. An important type of nucleosynthesis is α capture, which produces α elements, including Mg, Ca, and Si. Another type of nucleosynthesis is neutron capture, of which there are two types. The slow neutron-capture process (s -process) and the rapid neutron capture process (r -process) are responsible for forming elements heavier than iron. The s -process is where an element slowly undergoes neutron capture and occurs in low to intermediate mass stars. The r -process is where an element rapidly undergoes neutron capture from a high flux of neutrons. The r -process occurs in neutron star mergers, but it is theorized that the r -process can also result from core collapse supernovae. The elements that are

produced from these two processes are called neutron-capture elements and include Y, Ba, and Eu.

The Hercules ultrafaint dwarf galaxy is an ideal target for studying chemical abundances as it contains red giant stars and has been previously studied using Keck-DEIMOS data with the 1200G diffraction grating. The 1200G grating in particular is ideal for obtaining the velocity of stars to determine membership. However, it lacks the wavelength range needed to measure abundances of some elements, such as Mg, with a precision that is useful for chemical evolution models. The 1200B grating of the Keck-DEIMOS instrument allows access to the wavelengths of absorption lines that can be used to study the chemical abundances of stars within an ultrafaint dwarf galaxy. In this paper, we discuss the chemical composition of the Hercules ultrafaint dwarf galaxy studied using medium-resolution spectra from the Keck-DEIMOS spectrograph observed with the 1200B grating.

2 Background

Previous work on the Hercules UFD describes the chemical abundances of two stars within the galaxy [1; 2]. These works mainly studied the chemical abundances of Mg and Ba. Both studies were limited in the number of stars (two) observed, and the second, more in depth study was limited by the range of wavelengths captured during the observation (6100–6400 Å). In Koch et al. (2008) [1], the measurement of Ba was extremely low for both stars, and they were unable to find any Ba absorption lines above noise level for either star. They also found enhanced Mg and Na features in the two stars which led them to conclude that a Type II supernova is a progenitor of this galaxy. Further study was done on these two stars and compared to literature by Koch et al. (2013) [2], who concluded that Type Ia and Type II supernovae most likely contributed to the chemical abundances seen in the Hercules UFD; however, the study did not have sufficient data to assert this conclusion, specifically lacking in data around Ba absorption lines.

Previous research [3] was also performed using Keck-DEIMOS data with the 1200G diffraction grating. This study confirmed the membership of 19 stars within the Hercules UFD. This study observed alpha elements such as Mg, Si, and Ca. This 1200G grating however lacks efficiency over

valuable wavelengths such as the Mg triplet.

3 Methods

3.1 PypeIt

We analyzed spectra from 28 stars in Hercules taken at the Keck Observatory using the medium-resolution DEIMOS spectrograph with the 1200B grating. To reduce the raw data, we used a pipeline called PypeIt [4]. There are a few advantages to using PypeIt as opposed to the pipeline our research group formerly used, which was a modified version of the official DEIMOS reduction software, spec2d, written in IDL. Most importantly, PypeIt fits for wavelength across the blue–red CCD chip gap, whereas spec2d fits separate wavelength solutions. The simultaneous fitting provides an additional constraint on the fit that prevents the wavelength solution from wandering at the blue and red margins of each chip.

3.2 Chemfit

After data reduction, we interpolated and stacked the data before using our chemical analysis program, chemfit [5]. Chemfit is a program that fits spectra to models written by Roman Gerasimov, a post-doctoral member of the Kirby research group. This program works in two stages: gridfit and localfit. Gridfit matches the observed spectra data to models. This stage also requires initial guesses for some parameters before fitting the spectra to a model, such as photometric data and redshift. We set the initial redshift estimate to 45 km s^{-1} for each star. We used an empirical color–effective temperature relation to find the effective temperature of the stars within the spectra, [6]. The surface gravity was then computed by the Stefan–Boltzman law ($L = 4\pi R^2 \sigma_{\text{SB}} T_{\text{eff}}^4$) and the Newtonian equation for acceleration due to gravity ($g = GM/R^2$). Gridfit will also fit parameters such as effective temperature and surface gravity based on the fitted model. After the gridfit stage, localfit finds the abundances of individual elements. From this program, we were able to obtain the chemical abundances of stars within Hercules and estimate which stars are members of

Hercules.

3.3 Membership Selection

The radial velocity (v_r) of the Hercules ultrafaint dwarf galaxy is 45 km s^{-1} , so members of this galaxy have v_r of $45 \pm 15.3 \text{ km s}^{-1}$, where 15.3 km s^{-1} is three times the velocity dispersion of Hercules [7]. This criterion includes 99.7% of member stars. A second membership/quality selection was performed based on the qualities of stars fitted by chemfit, such as effective temperature, surface gravity, and metallicity. The effective temperature selection was performed first, eliminating any star with an effective surface temperature above 5500 K. Since the stars in the Hercules UFD are red giants, their effective temperatures should be relatively low. After this cut was made, there were 20 stars that remained potential members. These stars all had surface gravities that could classify them as red giants ($\log g < 3.6$).

4 Results

With the combination of PyPeIt and Chemfit, we were able to measure the chemical abundances of stars within the Hercules ultrafaint dwarf galaxy. Out of the 28 stars, only one star had a truly poor fit with chemfit, where no information on the chemical abundances could be used. We measured the Mg abundances of 18 likely member stars and the Ba abundances of 14 stars. Comparing our analysis done with the 1200B grating data to the previous analysis done with the 1200G data, we see in Figure 1 that we have more measurements of magnesium. The magnesium abundances also follow a downward sloping pattern as can be seen in Figure 2, which agrees with previously published work [8]. The downward sloping trend in the magnesium abundance plot is typical in dwarf spheroidal galaxies, where the decrease in $[\text{Mg}/\text{Fe}]$ with an increase in $[\text{Fe}/\text{H}]$ represents how iron is enriched faster than magnesium due to an increasing ratio of Type Ia to core collapse supernovae over time. The data points within Figure 2 highlighted by the red circle that do not follow the downward sloping trend are most likely spurious measurements and nonmembers which persisted through the membership selection.




Figure 1: Magnesium abundances measured with the 1200B grating (black points) and the 1200G grating (red points). The 1200G grating yielded some measurements of magnesium with poor accuracy, while the 1200B grating provides more accurate magnesium measurements. The x-axis describes iron metallicity of the stars, and the y-axis shows the ratio of magnesium to iron in the stars in comparison with the sun.

For the α elements, we measured Mg, Si, Ca, and Ti, as can be seen in Figure 3. While the trends of Mg, Si, and Ti agree with previous studies, the Ca abundances are elevated for this UFD compared to the 1200G measurements. Due to errant pixels and this elevated measurement of Ca, the overall abundance of the alpha elements measured in the 1200B data is systemically higher than the 1200G measurements, as can be seen in Figure 4.

For the neutron capture elements, we measured only Ba with accuracy, as can be seen in Figure 5. We had some measurements of Y, Eu, and La, but none had the precision needed to draw conclusions. The DEIMOS 1200B grating has a prominent Ba absorption line compared to the absorption lines of the other neutron capture elements. UFDs also lack neutron-capture elements as they do not have enough mass (i.e., gravitational potential) to hold on to the ejecta from the high energy




Figure 2: Magnesium abundances as a function of metallicity. The black data points are the measured magnesium abundances, and the red trend line is an approximation of the trend between magnesium and metallicity in UFDs. The points within the circle do not follow the trend.

events that produce the r -process.

5 Conclusion

The UFD Hercules is a typical metal poor ultrafaint dwarf galaxy with abundance patterns that agree with previous literature. As metallicity increases, we observe a decrease in the ratios of Mg, Ca, and Ba to Fe. In regards to the data, Keck-DEIMOS data using the 1200B grating captures the α elements' abundances but lacks the absorption lines to characterize the stars' neutron capture element abundances. However, the lack of neutron capture elements may be further explained by UFDs' natural lack of neutron-capture elements, as they do not typically experience the events that would produce such elements, or they lack the gravitational potential to retain those neutron-capture




Figure 3: The α elements we measured as a function of metallicity (Mg, Si, Ca, Ti).

elements.

Using a combination of the data reduction pipeline, PyeIt, and the model fitting software, Chemfit, we were able to measure the chemical abundances of an ultrafaint dwarf galaxy. Further research is needed to clarify the issue with elevated abundances of the α elements and in further membership selection.

References

- [1] A. Koch, A. McWilliam, E. K. Grebel, D. B. Zucker, and V. Belokurov, *ApJ* **688**, L13 (2008), arXiv:0810.0710 [astro-ph] .
- [2] A. Koch, S. Feltzing, D. Adén, and F. Matteucci, *A&A* **554**, A5 (2013), arXiv:1304.2764 [astro-ph.GA] .




Figure 4: α element abundances comparison between 1200B and 1200G gratings. The 1200B measurements of $[\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$ (black points) are systematically higher than the 1200G measurements (red points).

- [3] E. N. Kirby, J. G. Cohen, P. Guhathakurta, L. Cheng, J. S. Bullock, and A. Gallazzi, *ApJ* **779**, 102 (2013), arXiv:1310.0814 [astro-ph.GA] .
- [4] J. X. Prochaska, J. F. Hennawi, K. B. Westfall, R. J. Cooke, F. Wang, T. Hsyu, F. B. Davies, E. P. Farina, and D. Pelliccia, *Journal of Open Source Software* **5**, 2308 (2020).
- [5] L. Fisher, R. Gerasimov, E. N. Kirby, A. J. Burgasser, C.-C. Hsu, and L. M. Weiss, *Research Notes of the AAS* **8**, 227 (2024).
- [6] Z. Zhou, L. Casagrande, X. Zhang, J. Xiong, J. Li, Y. Guo, Z. Han, and X. Chen, *Monthly Notices of the Royal Astronomical Society* **549**, stag1009 (2026), <https://academic.oup.com/mnras/article-pdf/549/4/stag1009/68435592/stag1009.pdf> .




Figure 5: Neutron-capture element abundances as a function of metallicity (Ba, La, Y, and Eu).

[7] J. D. Simon and M. Geha, *ApJ* **670**, 313 (2007), arXiv:0706.0516 [astro-ph] .

[8] L. E. Henderson, E. N. Kirby, M. A. C. de los Reyes, R. Gerasimov, and V. Manwadkar, *The Astrophysical Journal* **983**, 117 (2025).

Characterizing the Exoplanet Transit Capabilities of the Sarah L. Krizmanich Telescope

ISABEL LOCKHART

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): [Dr. Lauren Weiss]

Abstract

Ground-based transit photometry provides an accessible method for detecting and characterizing exoplanets by measuring small decreases in stellar brightness that occur when a planet passes in front of its host star. However, the shallow depth of many exoplanet transits makes signal recovery dependent on photometric sensitivity and consideration of observational conditions to distinguish the planetary signal from sources of noise. We have assessed the extent to which the Sarah L. Krizmanich Telescope (0.8 m) at the University of Notre Dame can support this work. Because eclipsing binaries produce light curves governed by the same underlying physical principles as transiting exoplanets but typically exhibit much deeper eclipses, they provide an ideal intermediate step before attempting exoplanet transit observations. Observations were obtained over eight nights, with each night providing an opportunity to recover and characterize a known stellar eclipse or exoplanet transit event. Eclipse and transit depths were measured from the normalized light curves. We recovered binary star eclipses of progressively shallower depths, as well as two exoplanet transit events with published depths of 2.74% for TrES-3 [1] and 1.40% for HAT-P-23 b [2]. The median RMS of the photometric residuals was 0.015, and the SNR of our shallowest transit was 35. These results demonstrate that the Krizmanich Telescope is capable of recovering exoplanet transit signals when observing conditions are favorable and appropriate targets are selected.

1 Introduction

Exoplanets are planets that orbit stars outside of the Solar System. Because they are much smaller and fainter than their host stars, they are often indirectly detected through their effects on the host star. Over 6,000 exoplanets have been confirmed to date, with many being discovered through space-based transit surveys such as Kepler and TESS, according to the NASA Exoplanet Archive [3].

Ground-based observatories play an important role in confirming candidates, refining transit measurements, and monitoring known systems [4]. The Sarah L. Krizmanich Telescope (the 'Kriz') at the University of Notre Dame has the potential to contribute to ground-based exoplanet observations. The Kriz has previously contributed to exoplanet research through the identification of an exoplanet candidate using eclipse timing variations [5].

Detecting and characterizing transiting exoplanets requires the ability to measure small decreases in stellar brightness as a planet passes in front of its host star. Because eclipsing binaries

follow the same underlying geometry while typically producing deeper signals, they provide an intermediate step for evaluating the telescope’s ability to recover shallow transit events. By observing eclipsing binaries with progressively decreasing eclipse depths before transitioning to known transiting exoplanets, the transit detection capabilities of the Kriz can be evaluated.

2 Background

An eclipsing binary is a system of two stars orbiting a common center of mass. When the orbital plane is aligned with the observer’s line of sight, one star periodically passes in front of the other, producing a decrease in the observed brightness of the system. These changes produce a light curve, from which properties including the orbital period and eclipse depths can be determined.

Transiting exoplanets produce light curves through the same geometry, where a planet passes in front of its host star and blocks a portion of the stellar light. The transit depth is approximately equal to the ratio of the projected cross-sectional areas of the planet and star,

$$\delta \approx \left(\frac{R_p}{R_*} \right)^2, \quad (1)$$

where δ is the fractional transit depth and R_p and R_* are the radii of the planet and host star, respectively.

Despite the difference in signal strength, eclipsing binaries and transiting exoplanets are analyzed using similar photometric techniques. Differential photometry compares the brightness of a target star to nearby reference stars to reduce atmospheric and instrumental variations. Aperture photometry measures the flux of these stars, allowing relative changes in brightness to be used to construct light curves.

3 Methods

Potential eclipsing binary targets were identified using the Kepler Eclipsing Binary Catalog [1]. A Python pipeline was developed using the Astropy [6] and Astrolan [7] packages to predict future

eclipse events from the published orbital periods (Table 1) and corresponding reference eclipse times. The pipeline was then used to determine which eclipses would be observable from the Kriz based on the predicted eclipse midpoint, target altitude, and declination. Targets with declinations greater than the observatory latitude were excluded due to telescope cable-wrap limitations that can trigger safety errors during tracking. Targets were filtered using combinations of stellar K-band magnitude, eclipse depth, orbital period, and eclipse morphology (detached/contact eclipsing binary).

The target selection process was then adapted for known transiting exoplanets using the NASA Exoplanet Archive [8]. Candidate exoplanets were selected using predicted transit times together with stellar K-band magnitude, transit depth, target altitude, and declination to identify observable transit events.

Observations were conducted over eight nights. Exposure times, focus settings, and the CMOS detector cooling setpoint were adjusted to account for varying observing conditions. No optical filters were used during these observations. Images were calibrated using dark and flat-field frames collected each observing night. MaxIm DL Pro 6 was used to construct the master calibration frames using the median combine method before applying the calibrations to the science images.

Following calibration, aperture photometry was performed in MaxIm DL Pro 6. An aperture radius and sky annulus were selected for each target, after which the software automatically matched stars between images and generated light curves using reference stars. Each curve used 2–3 reference stars. The continuum magnitude was determined from the resulting light curves by selecting the brightest 10%-20% of observations and calculating the average magnitude of these points. This value was used as the continuum reference for converting the differential magnitudes into relative flux using

$$\frac{F_2}{F_1} = 10^{-0.4(m_1 - m_2)}, \quad (2)$$

where m_2 represents the continuum magnitude and m_1 represents the measured magnitude at each observation.

A polynomial fit was then applied to each flux light curve to estimate the eclipse depth. Deeper eclipses typically required higher-order polynomials to fit the full eclipse profile. The polynomial fits were not intended to represent a physical model of the eclipsing system or exoplanet transit.

For the two exoplanet observations, a slight downward trend in continuum flux was observed throughout the observation period. Because this trend was also present in the reference stars, it was not interpreted as a feature of the target stars and was not relevant to our scientific goals. Therefore, the data for the two exoplanet transits were detrended before analysis. The continuum data outside of transit was fit with a low-order polynomial, and this trend was removed from the entire light curve.

The noise level of each observation was calculated using the root-mean-square (RMS) of the residuals from the polynomial fit. The reduced chi-squared statistic was then calculated for both the polynomial fit and a constant baseline model:

$$\chi_{\text{red.,baseline}}^2 = \frac{1}{N} \sum_{i=1}^N \frac{(y_i - y_{\text{baseline}})^2}{\sigma^2}, \quad (3)$$

$$\chi_{\text{red.,fit}}^2 = \frac{1}{N} \sum_{i=1}^N \frac{(y_i - y_{\text{fit}})^2}{\sigma^2}, \quad (4)$$

where σ is the RMS noise of the light curve. The difference between these two values,

$$\Delta\chi_{\text{red.}}^2 = \chi_{\text{red.,baseline}}^2 - \chi_{\text{red.,fit}}^2, \quad (5)$$

was used to quantify the improvement of the eclipse fit compared with a flat baseline.

Finally, a transit signal-to-noise ratio (SNR) was calculated for each target. Although this is typically used for exoplanet transit observations, it was also applied to the eclipsing binary targets as a consistent comparison method. The SNR was calculated as

$$\text{SNR} = \frac{\delta}{\sigma} \sqrt{N}, \quad (6)$$

where δ is the measured eclipse depth, σ is the RMS noise, and N is the number of observations occurring during the eclipse or transit.

To further evaluate the recovered exoplanet transit signals, theoretical transit models were generated using the BATMAN transit modeling package [9]. Both BATMAN and polynomial models for transits are highlighted for comparison. BATMAN was used to model the expected transit shape based on published system parameters, including the transit midpoint, orbital period, semi-major axis, inclination, eccentricity, argument of periastron, and limb darkening coefficients. The resulting depths are compared with the literature values in Results.

4 Results

Eight primary observing targets were selected to investigate progressively shallower eclipse depths. These included six eclipsing binaries and two transiting exoplanets.

Target	Period (days)	Star Kmag	Weather	Setpoint (°C)	Vents?	Exp Time (s)	Focus	Type	Date UTC
KIC 7871200	0.24	12.8	Warm	-10	Off	5	76000	EB	22-Jun
KIC 4451148	0.74	11.2	Hot	-5	Off	1	78000	EB	30-Jun
KIC 2715417	0.24	14.07	Cloudy	-8	On	5	78000	EB	1-Jul
KIC 3836439	1.54	7.57	Warm	-10	Off	0.2	78000	EB	7-Jul
KIC 4547308	0.58	12.5	Hot	-5	Off	8	78000	EB	14-Jul
TrES-3	1.31 ^a	10.6 ^a	Hot	-5	Off	10	78000	Transit	15-Jul
HAT-P-23 b	1.21 ^b	10.8 ^b	Cool/Hazy	-10	Off	8	77500	Transit	20-Jul
KIC 4739791	0.90	14.7	Cool/Clear	-10	Off	20	76500	EB	23-Jul

Table 1: Summary of observed eclipsing binary and transiting exoplanet targets and corresponding observing conditions. EB magnitudes and periods obtained from the Kepler Eclipsing Binary Catalog [1]. Transit magnitudes and periods obtained from ^aSozzetti et al. (2009) [10] and ^bKokori et al. (2023) [2] as referenced in the NASA Exoplanet Archive [8].

The first observations focused on recovering eclipsing binary light curves with relatively large eclipse depths. KIC 7871200 was observed with an eclipse depth of approximately 41% [1], which provided an initial trial of the observing procedure. Additional eclipsing binaries were then observed with decreasing eclipse depths, including KIC 4451148 with a depth of approximately 13% [1], KIC 3836439 with a depth of approximately 6% [1], and KIC 4547308 with a depth of approximately 9% [1]. KIC 3836439 could not be recovered due to a lack of reference stars of similar magnitude to the target star. This observation highlighted the absolute importance of reference

Target	Lit. Depth (%)	Polynomial Depth (%)	Percent Diff. (%)	Polynomial Order	RMS Noise	Red. $\Delta\chi^2$	Transit SNR
KIC 7871200	41.05	41.93	2.12	10	0.0149	178.64	1113.88
KIC 4451148	13.40	9.36	35.50	5	0.0291	4.65	202.28
KIC 2715417	10.05	N/A	N/A	N/A	N/A	N/A	N/A
KIC 3836439	6.13	N/A	N/A	N/A	0.0471	N/A	N/A
KIC 4547308	8.83	5.33	49.44	3	0.0145	5.72	173.37
TrES-3	2.74 ^a	3.00	9.06	12	0.0150	0.54	49.29
HAT-P-23 b	1.40 ^b	1.88	29.27	8	0.0143	0.32	35.40
KIC 4739791	16.14	17.46	7.86	10	0.0293	8.69	141.11

Table 2: Comparison between literature and calculated eclipse/transit depths and analysis results. Literature eclipse depths for the eclipsing binary targets obtained from the Kepler Eclipsing Binary Catalog [1]. Literature transit depths obtained from ^aSozzetti et al. (2009) [10] and ^bKokori et al. (2023) [2] referenced in the NASA Exoplanet Archive [8]. Percent difference column displays difference between calculated and literature transit/eclipse depths.

stars in the data reduction process. The four successful recoveries with normalized light curves and the best fit model are shown in Fig. 1. Additionally, KIC 2715417 was observed during unfavorable weather conditions. Significant cloud coverage and running observatory vents prevented the recovery of a reliable eclipse signal.

Following successful recovery of eclipsing binary light curves, observations transitioned to known transiting exoplanets. TrES-3, with a published transit depth of approximately 2.74% [10], produced a measurable transit signal. The calculated depth based on the polynomial model is 3.00%, and the calculated depth based on the BATMAN model is 2.60%. The shallowest successfully recovered transit was HAT-P-23 b, with a published depth of 1.40%, a polynomial depth of 1.88%, and a depth calculated from the BATMAN model of 1.46%. Differences between the depths calculated from polynomial and BATMAN models may be attributed to the BATMAN models incorporating physical transit shape, whereas the polynomial fits do not and are approximations based on the data. The quality of each observation was evaluated using RMS noise, reduced $\Delta\chi^2$, and transit SNR, as shown in Table 2.

The observing conditions were found to have a significant impact on the achievable photometric precision. Several observations were affected by weather conditions including clouds, haze, and elevated temperatures. Warmer nights required higher camera setpoints to maintain detector operation. Reference star selection was another important factor in determining the quality of the recovered light curves. Targets with multiple suitable comparison stars produced more stable dif-




Figure 1: Recovered eclipsing binary light curves with polynomial models.

ferential photometry, while fields lacking reliable reference stars made recovery of transit signals more difficult.

Exposure times were generally adjusted based on target brightness. However, short exposures introduced additional noise, whereas long exposures increased the likelihood of image motion during an exposure. As the telescope passes near zenith, camera auto-tracking turns off due to a large error leading to image blur.

5 Conclusion

This project investigated whether the Sarah L. Krizmanich Telescope at the University of Notre Dame can reliably recover transiting exoplanet signals and characterize their light curves.




Figure 2: Recovered transiting exoplanet light curves. Row 1 shows polynomial and BATMAN model for TrES-3, row 2 shows polynomial and BATMAN model for HAT-P-23 b.

Four eclipsing binaries and two known transiting exoplanets were successfully recovered. These observations demonstrate that the Krizmanich Telescope is capable of recovering known shallow exoplanet transit signals under favorable observing conditions using differential photometry and standard data reduction techniques. The identification of several factors that strongly influence detection quality was also achieved. While eclipse depth is an important consideration, observing conditions, target brightness, exposure time, and the availability of suitable reference stars all contribute to the quality of the recovered light curve.

The results of this work provide a practical framework for future exoplanet observations with the Krizmanich Telescope. Future work will focus on observing additional transiting exoplanets

with progressively shallower transit depths to further define the detection limits.

References

- [1] B. Kirk, K. Conroy, A. Prša, *et al.*, The Astronomical Journal **151**, 68 (2016).
- [2] A. Kokori *et al.*, The Astrophysical Journal Supplement Series **265**, 4 (2023).
- [3] NASA Exoplanet Archive, Exoplanet and candidate statistics (2026), accessed: 2026-08-03.
- [4] K. Collins, S. N. Quinn, D. W. Latham, *et al.*, in *American Astronomical Society Meeting Abstracts #231*, Vol. 231 (2018) p. 439.08.
- [5] M. Leichty, P. Garnavich, C. Littlefield, A. Schwope, J. Kurpas, P. A. Mason, and K. Beuermann, The Astrophysical Journal **967**, 81 (2024).
- [6] Astropy Collaboration, A. M. Price-Whelan, *et al.*, The Astrophysical Journal **935**, 167 (2022).
- [7] B. M. Morris, E. Tollerud, B. Sipőcz, C. Deil, S. T. Douglas, J. Berlanga Medina, K. Vyhlířmeister, T. R. Smith, S. Littlefair, A. M. Price-Whelan, W. T. Gee, and E. Jeschke, **155**, 128 (2018), arXiv:1712.09631 [astro-ph.IM] .
- [8] J. L. Christiansen, D. L. McElroy, *et al.*, The Planetary Science Journal **6**, 186 (2025).
- [9] L. Kreidberg, Publications of the Astronomical Society of the Pacific **127**, 1161–1165 (2015).
- [10] A. Sozzetti, G. Torres, D. Charbonneau, J. N. Winn, S. G. Korzennik, M. J. Holman, D. W. Latham, J. B. Laird, J. Fernandez, F. T. O’Donovan, G. Mandushev, E. Dunham, M. E. Everett, G. A. Esquerdo, M. Rabus, J. A. Belmonte, H. J. Deeg, T. N. Brown, M. G. Hidas, and N. Baliber, The Astrophysical Journal **691**, 1145–1158 (2009).

Observations and Phase Space Constraints of Leptonic Decays Using CMS Scouting Data

Aldir Moreira

2026 NSF/REU Program
Department of Physics
University of Notre Dame

ADVISOR: Dr. Marc Osherson

Abstract

The study and testing of BSM models requires searching for theorized decays not predicted by the Standard Model. The Scouting Dataset created by the Compact Muon Solenoid experiment at CERN collects events filtered out by energy allowing the dataset to be used in the investigation of rarer, lower energy decays. This dataset is used to investigate the muon-producing η double Dalitz decay and the all-charm tetraquark family. Additionally, the phase spaces of decays resulting in pair permutations of the J/Ψ , $\Psi(2S)$, and $\Upsilon(1S)$ resonances as well as charge, momentum, and ‘mistaken-identity’ particle constrained spaces are investigated. η is reported at 0.5480(4) GeV with a width of 3.0 (22) MeV. $X(6600)$ and $X(7100)$ are observed without background removal at 6.46 (0.05) GeV and 7.02 (0.02) GeV with widths of 0.180 GeV and 0.241 GeV and significances of 3.25σ and 3.22σ respectively which are in agreement with ATLAS results. The investigated phase spaces are not found to strongly warrant further investigation, however, the all-charm producing spectrum includes various features that require further analysis.

1 Introduction

High energy physics in the 21st century has focused on Beyond Standard Model (BSM) theories which aim to explain the fundamental gaps that exist in the current understanding of particle physics. Studying and testing BSM models requires searching for theorized and anomalous decays through observations using particle colliders.

The Compact Muon Solenoid experiment is a muon detector consisting of a first trigger made up by electromagnetic and hadron calorimeters along with automated filters to reduce the event rate by 7 orders of magnitude which then allows the second trigger to create live reconstructions of the ongoing collisions. Part of the approximately 40,000,000 events lost is stored in a Scouting Dataset which has a lower energy threshold and captures more events at the cost of lower quality. This lower energy threshold is vital in studying less common decays that may be outside of the Standard Model (SM) would otherwise be lost in higher energy background and SM decays. Of interest among these decays are the all-charm tetraquark family and the η double Dalitz decay.

Through the use of the CMS Scouting Dataset we support the recent observations of these particles, provide comparable characteristic information that may be used to learn more about the internal structures of the all-charm tetraquark family, and set a baseline of accuracy for the

Scouting dataset. Additionally, we report on the null results and potential points of interest found in the investigation of muon decay phase spaces constrained by charge, momentum, count, and identity to determine whether such spaces justify further investigation.

2 Background

Particle resonances, among which the η , all-charm family, and more famously the Higgs boson are included, are short lived particles that may be studied through their lepton decay products. Such particles have been used to test the precision of the SM as well as create frameworks for BSM theories.²




Figure 1: A 2 muon decay spectrum rebuilt from Compact Muon Solenoid Experiment Scouting data used in this study. The expected dimuon resonances from 0-200 GeV appear clearly.

The η resonance is a pseudo-scalar particle composed of a linear up, down, and strange, quark-antiquark. As seen in Figure 1, the particle is readily visible in the 2 muon decay spectrum at $547.85 \text{ MeV}/c^2$. In the double Dalitz decay mode a pseudo scalar meson such as η decays to a pair of exchange photons which then decay to a lepton pair each. Of these possible leptons for the η decay, the electron end state has been observed and more recently the muon end state was observed (CMS Collaboration, 2023). This decay holds potential for BSM study as it requires the electromagnetic force to dominate when in this circumstance it is usually suppressed by the strong nuclear force.

The all-charm tetraquark family, $X(6600)$, $X(6900)$, and $X(7100)$, is a group of tenuous

exotic quark states where the latter 2 are radially excited states of X(6600). The particles are likely structured as charm diquark-antidiquark pairs and serve as useful tools to study hadronic matter structure in exotic states². The identification of this family in the Scouting data would support the notion that its higher volume is vital for precision and BSM decays.

3 Methods

The objective of these analyses hinges on the use of Lorentz vectors or four-vectors. For these calculations we use the CERN maintained python package ROOT. These properties are in terms of CMS coordinates, transverse momentum, pseudorapidity, and azimuthal angle. In addition to these tools, the grand majority of analysis was done using python and C++ on the Notre Dame Glados computing server.

For a given experiment the data was taken from the Scouting dataset and compiled into a ROOT tree for easier manipulation. All events below the specified muon per event number were then eliminated. Following this an event reconstruction size was chosen and a C++ function was used to construct every permutation of this size within a given event. The constructed events were then filtered by specified charge, for diresonance spectra where an initial particle must first decay to a specified particle pair, such as the $\Psi(2S)$ and $Y(1S)$, the events selected permutations are reconstructed into a second set of combinations of which those with masses matching the desired particle pair are passed. The particles are then checked to be coming from the same origin point or ‘vertexed’. Vertexing consists of collecting extracted position data of each muon’s origin and checking the likelihood that 2 extracted vertices are the same. For our likelihood factor we required that any same vertices must have χ^2 value divided the degrees of freedom equal to less than 5.

Having acquired a pair of muons in an event likely from the same origin, this vertexing process was repeated to ensure that all remaining events are composed of related particles. The events were then reconstructed with the ROOT MATH LorentzVector function allowing a singular permutation to be traced back to the particle that led to it. Lastly, these parent particles

were plotted as ROOT histograms with the invariant mass of each passed permutation as the x-axis and the event count at each mass as the y-axis.

4 Results

4.1 Phase Space Constraints

This set of results were based on constraints of particle count, charge, and identity. The first results observed are particle count where permutation sizes were allowed to vary between 2 and 6 muons such that the end product of a parent particle may be identified more easily from the background and other mixed end product amounts in nearby decays.




Figure 2: This figure displays the 3 muon spectrum from 0-200 GeV. The primary feature visible is the turn on and noise is relatively equal in all parts of the spectrum.

The 3 muon spectrum pictured above demonstrated no strong evidence of resonances. In this spectrum the still unobserved tau to 3 muon decay was searched for however no conclusive evidence was found to support such a decay. This spectrum was deemed to be a null result and further study is not likely necessary.




Figure 3: This figure displays the 4 muon spectrum from 0-200 GeV. This spectrum primarily consists of the turn on as well as a secondary broad feature near 60 GeV which is likely kinematic in origin. This spectrum is observed in finer regions in the following resonance result section.

The 4 muon spectrum in Figure 4 demonstrates considerably more interesting data than the other muon count spaces. The included all-charm tetraquarks and η decay are discussed in later sections, however, several other statistically significant peaks pointing to resonances are potentially present. The further study of this spectrum beyond the narrow mass windows observed was outside of the scope of this project but the high quality data and sheer amount available strongly suggest that there is more to be observed in the above spectrum.

Preliminary studies of decays leading to 5 and 6 muons were conducted leading to null results of exceedingly noisy turn ons similar in appearance to the 3 muon spectrum of Figure 2. We omit strong conclusions on these spaces as it is possible that not enough data was used to find peaks in these less common decays and such decays may not be obtained with our methods.

For each of the muon decay amounts specified above, every possible charge permutation was investigated under the possibility that a differently charged background may be hiding less obvious features. This process was observed to minorly change certain background features across all of the particle decay numbers, however, it did not reveal any statistically significant peaks and therefore, produced no usable results. Although it is possible that in the search of particles with known charges in a background that tends to a particular charge, such a method would be a viable tool of research, our work found that all spaces remained relatively similar. Thus, it was determined that all of the aforementioned null phase space results, 3, 5, and 6 muons, also failed to produce results across all charge specifications.

The final set of phase spaces observed was the mistaken identity phase spaces. This set of spaces is possibly different as the CMS experiment may misidentify pions, kaons, and other particles as muons leading to misplaced peaks and potentially hiding resonances³. The tested phase spaces consisted of all muon-pion and muon-kaon combinations of the 3 particle dataset as well as the 4 muon 2 pion combination of the 6 particle dataset. These combinations were all

tested across the available charge permutations. In the misidentified pion case no peak shift greater than 0.1 MeV was found and no peaks were revealed. Similarly, the greater mass of the kaon caused a peak shift on the order of 0.5 MeV however these shifts did not correspond to resonance masses of higher accuracy. This set of results was therefore deemed null and further study will not be conducted on misidentified particles in the 3 muon spectrum while further study will also not be conducted without motivation or a greater dataset on the 6 muon spectrum.

These constraints represent a total of 38 distinct phase spaces that were checked across the scope of this project. Beyond the well understood 2 muon spectrum, it was determined that all permutations of the 3 particle spectrum do not warrant further study while those of the 5 and 6 particle spectrum do not require further investigation unless a specific cause is given or a dataset that differs by magnitudes is obtained.

4.2 Diresonance Phase Spaces

This set of phase space results consists of resonances visible in the 2 muon spectrum, namely J/Ψ , $\Psi(2S)$, and $Y(1S)$, joined as pairs in the search for parent particles which first decay to a pair of quarkonium states. The individual muon pairs were also required to be electrically neutral to increase the likelihood that they formed from a dimuon resonance state.




Figure 4: These plots show the spectra of 4 muon decays which are first required to decay to a pair of J/Ψ resonances. The right graph is a cut from 6-12 GeV of the left graph which was created from 6-50 GeV. The plots are primarily background pair production noise with a broad feature from 8-9 GeV which may be kinematic in origin.

Figure 4 displays possible peaks at 26 GeV and more evidently at 9.8 GeV which were not investigated but may indicate points of interest. A double peak feature is also visible at the

peak of the turn on and may require further research. This diresonance phase space was deemed a null result and does not seem to strongly warrant further study.




Figure 5: This plot shows the $J/\Psi \Psi(2S)$ pair decay. Despite the large dataset exceeding a million events, very few points are visible in the spectrum.

In Figure 5, a possible peak appears at 10.5 GeV but there is a strong likelihood that it belongs to the noisy turn on. This result was null and further investigation is not suggested.




Figure 6: This figure displays the $\Psi(2S) \Psi(2S)$ pair decay. The points are extremely sparse with a minor turn on visibly peaking around 8 GeV.

Figure 6 displays a possible peak of note in a very strong feature at 7.42 GeV, however, this feature is just above the double $\Psi(2S)$ mass at 7.37 GeV and in conjunction with the lack of interference removal, this point may be dismissed. This phase space was thus concluded to lack any interesting data.

Lastly, the $J/\Psi Y(1S)$, $\Psi(2S) Y(1S)$, and $Y(1S) Y(1S)$ combinations were also investigated and produced nothing of note beyond noisy turn on spectra with noise ending before 80 GeV in every case. These 3 results were not considered to be worth further study, however, if a significantly larger dataset is made available they may be worth checking.

4.2 η and All-Charm Measurements

In the investigation of the 4 muon spectrum, the η double Dalitz decay was found at 548.0(4) MeV with a width of 3.0(22) MeV. These values are highly consistent with CMS Collaboration values of 547.9(18) MeV with a width of 2.89(83) (MeV)¹.




Figure 7: This figure displays the 4 muon spectrum from 0.5 GeV to 0.6 GeV with a RooFit 3 polynomial and gaussian fit around the suspected η decay which occurs before the turn on.

Our values are presented with a higher precision than CMS Collaboration values however this is due to the higher amount of systematic error that the latter took into account and any value narrowing that our characterization does should be understood in this context.

Similarly, the X(6600) and X(6900) particles were readily observed in the 6-9 GeV window. These had masses of 6.46(5) GeV and 7.02(2) GeV with widths of 180(3) MeV and 240(11) MeV respectively.




Figure 8: This figure shows the 4 muon spectrum from 6-9 GeV with X(6600) and X(6900) visible and compared to literature values with vertical lines. The green crosses represent experimental data while the black line represents a fit created using the ABCD estimation method. The bottom window plots significance on the y-axis.

These were found with significance values of 3.25σ and 3.33σ respectively. These values lay outside of the CMS Collaboration values of 6.64(23) GeV at 440(600) MeV width for X(6600) and 6.85(14) GeV with a width of 191(15) MeV for X(6900)². This incongruity may

be attributed to the background removal that was performed in the collaborative experiment and does not strongly suggest that the observed points are not part of the all-charm family.

Additionally, a rounded peak consistent with the shape and relative size of X(7100) is visible at 7.2 GeV but analysis is omitted for lack of strong evidence. Lastly, several peaks of interest, most notably at 8 GeV, are present and may be consistent with currently observed resonances; we strongly recommend that this Scouting data phase space be studied further.

5 Conclusion

Through the course of this experiment we have investigated 45 distinct phase spaces using CMS Scouting data and through resonance observations have supported the viability of this dataset for experiment work. Additionally, we have ruled out at least 41 phase spaces from further study and have indicated several areas which require more data or investigation due to anomalous outcomes. We have thus done the preliminary work for more in-depth investigation of the neutral 4 muon phase space and indicated specific points of interest with scientific potential for BSM research.

I would like to personally thank my advisor Dr. Marc Osherson for the guidance and knowledge he provided for this project. I would also like to thank Dr. Umesh Garg and Kristen Amsler for creating a productive and inviting environment during the 2026 Notre Dame Physics REU as well as Notre Dame for providing the opportunity to conduct research.

References

- [1] The CMS Collaboration. (2023). Observation of the rare decay of the η meson to four muons. *APS Physical Review Letters*, 131(091903).
<https://doi.org/https://doi.org/10.1103/PhysRevLett.131.091903>
- [2] The CMS Collaboration. (2025). Determination of the spin and parity of all-charm tetraquarks. *Nature*, 648, 58–63. <https://doi.org/10.1038/s41586-025-09711-7>
- [3] An, L., McNulty, R., & Mikhasenko, M. (2024). The determination of the spin and parity of a vector-vector system. *Journal of High Energy Physics*, 2024(7).
[https://doi.org/10.1007/jhep07\(2024\)092](https://doi.org/10.1007/jhep07(2024)092)

The $^{13}\text{C}(\alpha, n)^{16}\text{O}$ Reaction Cross Section as a Function of Beam Energy and Angle

ELLIE STONECIPHER

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Dr. James deBoer and Dr. Maxwell Sorenson

Abstract

There is a gap in consistent, relevant (α, n) reaction data, making it difficult to model, predict, and calculate neutron production in various environments. The $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction is a useful reference for other (α, n) experiments, on top of other applications like subtracting background from high sensitivity experiments. Over an energy range of 5.2-8 MeV, the $^{13}\text{C}(\alpha, n)$ reaction cross section was measured as a function of beam energy and angle. LaBr_3 detectors recorded γ -rays from ^{16}O de-excitations to the ground state. The data were used to find the cross section and angular correlation coefficients. Results are supported with comparisons to other experimental findings. Results will be available as a reference to other (α, n) reaction experiments, among other uses.

1 Introduction

The (α, n) reaction is a primary source of neutron production in various locations such as stars, nuclear energy plants, and underground laboratories. Neutron production can be modeled, predicted, and calculated from the cross section of a reaction, which is the probability of a reaction occurring. Accurate determination of cross sections requires new and consistent data. However, most of the data in use today can be dated back to at least the 1980s, and there are many inconsistencies among them [1]. The main reference reaction used in literature is the $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction, making the task of collecting precise $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction data a priority.

Furthermore, a thorough evaluation of the $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction cross section benefits specific fields of physics beyond being a reference. In neutrino and dark matter search experiments, α -particles released through decay in some equipment material can interact with naturally occurring ^{13}C deep below the surface of the Earth. Reactions like these add neutrons to backgrounds, depriving experiments of their sensitivity. In nuclear reactors, the inverse reaction occurs, where a neutron is captured and an α -particle is emitted. The emission of α -particles leads to a build up of helium which harms the reactor material through embrittlement. Solid understanding of the cross section allows the lifespan of nuclear reactor material to be predicted.

From the $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction cross section, accurate measurements of other (α, n) reactions can be made, neutron backgrounds in high sensitivity experiments can be reduced, and the accu-

mulation of helium in reactors can be predicted. For these reasons, the $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction was measured to collect new, consistent data which have been analyzed to determine the cross section as a function of beam energy and angle.

2 Methods

2.1 Experimental Setup

The Nuclear Science Lab (NSL) at the University of Notre Dame facilitated a $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction. The NSL's 5U accelerator produced a beam of α -particles. The energy of the α -particle beam E_α changed with each run, ranging from $E_\alpha = 5.2$ to 8 MeV. The beam deposited an average of 2 mA of beam current on a water cooled ^{13}C target. The target appears to have missing parts, exposing the backing of the target holder and possibly impacting the results of the experiment. The α -particles interact with the target, inducing $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reactions, thus causing the output of γ -rays. $^{14}\text{LaBr}_3$ detectors record these emitted γ -rays. These detectors, 2" x 2" in diameter, are placed 12" from the target in a circular arrangement to encompass multiple angles (see Figure 1), allowing the angular coefficients to be determined from analysis. The charge deposited on the target throughout a run is recorded using an ORTEC 439 Digital Current Integrator.




Figure 1: The schematic (left) vs the actual setup (right): detectors arranged around the target (highlighted in red), with the ORTEC 439 on the target

2.2 Energy Calibration

Every time a γ -ray hits a detector, that detector records a voltage pulse as one count. The maximum voltage of the pulse determines to which channel the count is assigned. Each channel has a different range of voltage pulses that it accepts. Counts and their corresponding channels are recorded in the file affiliated with that run. That file can be uploaded to a ROOT program, where the data is plotted as a channel vs counts histogram like Figure 2a). The histograms of these data files provide information on the cross section of the reaction; however, the data lack universality when it is sorted by channels. For this reason, the histograms are energy calibrated. Energy calibration is performed by scanning sources with easily identifiable peaks of well known energies. For this experiment, scans of ^{137}Cs , ^{27}Al , and ^{60}Co were held. These sources have a combined total of eight peaks serviceable for energy calibration. For each detector, channel numbers containing the centers of peaks were assigned the peak's energy value. These eight points were plotted as shown in Figure 3 and the trend line's equation was extracted. The trend line is essentially the energy calibration formula unique to a given detector, and provides the slope, offset, and quadratic term of the formula. With the incorporation of these values into a ROOT program that applies the energy calibration formula to the corresponding detectors data, histograms such as Figure 2b) are created. Now, the cross section can be calculated as a function of energy, making results widely applicable.




Figure 2: Two version of the same histogram: a) the raw histogram of channels vs counts and b) the histogram after energy calibration of energy vs counts




Figure 3: A detector's channel vs known energy graph with the extracted calibration formula listed

2.3 Cross Section Analysis

The cross section of any reaction is determined by solving

$$A = \epsilon N_p N_T \sigma \quad (1)$$

for the cross section σ , where A is the number of counts in the peak, ϵ is the detector efficiency, N_p is the number of beam particles hitting the target, and N_T is the number of particles in the target.

The number of counts in a peak A is known as the yield. Yields are found by fitting peaks with a ROOT program. The program searches for a peak within a given energy range, centered at a given value. The found peak then undergoes a fitting process where the background yield is calculated and subtracted from the net yield. The program loops through all of the runs, fitting the same peak. The program outputs a file for each detector, containing each run's energy, yield, yield error, and charge.

Each detector has a different efficiency at recognizing γ -rays and recording them. To find the efficiency of the detectors, a reaction that reliably outputs a prominent amount of high energy γ -

rays must be used. The charge deposited on the target during the run is used to determine the total amount of emitted γ -rays, which are compared to the amount of γ -rays a detector actually recorded. To find the latter, runs of the $^{19}\text{F}(\text{p},\alpha\gamma)$ reaction were executed. For each detector, the 6.1 MeV single escape peak's yield and yield error were found using the peak fitting ROOT program described earlier. In order to find the amount of γ -rays actually emitted, the charge deposited on the target throughout the run is recorded and saved in μC . Using the conversion factor from [2], the charge can be converted from μC to γ -rays, therefore being the total amount of γ -rays produced. The ratio of a given detector's γ -ray yield to the actual amount of γ -rays produced is the efficiency of that detector.

The number of beam particles hitting the target is found using the charge deposited on the target during a run, given by the current integrator. An α -particle beam consists of doubly ionized helium nuclei, which have a charge state of two. To differentiate the amount of particles hitting the target from the amount of charge deposited, the charge value is divided by two. The charge value is then converted from μC to C . For each particle, the charge of an electron $q = 1.6E-19 \text{ C}$ is recorded, so the charge value is divided by the electron charge to find the number of beam particles that hit the target.

To find the number of particles in the target, the target thickness must be known. The empirical target thickness was never measured, leaving the nominal thickness of $20.2 \mu\text{g}/\text{cm}^2$ to operate as the target thickness. The number of particles is a quantity measured in units of atoms/cm^2 , which can be obtained from target thickness through unit conversion.

2.4 Angular Correlation Calculation

The angular correlation coefficients are found using the angular correlation equation

$$W(\theta) = A_0 + A_2P_2(\cos\theta) + A_4P_4(\cos\theta) \quad (2)$$

where A_0 , A_2 , A_4 are angular coefficients and $P_2(\cos\theta)$ and $P_4(\cos\theta)$ are the 2nd and 4th order Legendre polynomials. The angular correlation equation gives information as to how the beam angle impacts the probability of detecting γ -rays. A_0 is independent of the beam angle. A_0 gives information on the intensity of the radiation, as it represents the radiation that is emitted in all directions uniformly. Higher order coefficients of the Legendre polynomials like A_2 , A_4 give further insight into the spin, shape, and decay of the nucleus.

The cross section data were fit to the angular correlation equation to extract the angular correlation coefficients. A fit is made for each energy, and the extracted A_0 points are plotted as a function of energy. The resulting plot displays the total cross section independent of beam angle.

3 Results

The cross section of the $^{13}\text{C}(\alpha,n)$ reaction was found as a function of energy for several angles. Plots for each angle were produced, as in Figure 4.




Figure 4: A graph of the $^{13}\text{C}(\alpha,n)$ reaction cross section (in units of millibarns) at 240°

Using data from another experiment held in 2020 at the NSL, a comparison of the 45° data could be made, as shown in Figure 5. There is apparent agreement between the shape of the datasets. However, the cross section found using data from this experiment must be multiplied by 1.6 to




Figure 5: Comparison of cross section data from two different experiments with detectors at 45° match the scale of the other experiment. This discrepancy could be due to the difference between the target's nominal thickness and actual thickness. The nominal thickness is a vague, estimated value given by the target's manufacturer. In the process of creating a target, variations exist causing fluctuations in the target's thickness. Therefore, a target's nominal thickness is typically used as a reference for comparing the thickness found through physical measurements.

For each energy level, the cross section data were fit to the angular correlation equation, producing a plot like Figure 6. Symmetric detectors were overlaid to show the consistency in γ -ray angular emission. Some points are evidently overlapping such as points at angles 30° and 120° (see Figure 1). Some points do not overlap, such as points at angles 15° , 45° , 90° , and 105° (see Figure 1). This could be due to detector misalignment with the target, as it appears in the actual setup (see Figure 1), the target is not exactly in the middle of the detector arrangement, but shifted slightly towards the upper half of the circle. This could also be due to the missing parts of the target.




Figure 6: Angular correlation fit at 6.43968 MeV across angular distribution with overlaying points representing detectors at symmetric angles




Figure 7: A graph of the the total cross section over all angles

Angular coefficients A_0 , A_2 , and A_4 are extracted from each fit. As a function of energy, the A_0 data were plotted with data from another NSL experiment for comparison. The comparison data were obtained using neutron detectors. Figure 7 shows the uniform features shared by both datasets. This agreement confirms the consistency between calculating cross sections with γ -ray data and neutron data. Likely due to the aforementioned reasons, the data had to be multiplied by a similar factor of 1.7 to match the scale of the comparison data.

4 Conclusion

Over the energy range of 5.2-8 MeV, the cross section of the $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction has been found as a function of beam energy and angle. In this work, we showed that our results, with a consistent scale factor, are analogous with two comparable NSL datasets. Our data will be available to the public when uploaded to EXFOR. Background calculations can be made using these data to increase sensitivity of neutrino and dark matter search experiments. This data is useful for future (α, n) experiments, as the $^{13}\text{C}(\alpha, n)$ reaction is a helpful reference for determining detector efficiency.

5 Acknowledgments

Thank you to Dr. deBoer and Dr. Sorenson for their constant support throughout this project. Thank you to Chloe Jones for her valuable writing advice for this report. Thank you to Dr. Garg, Kristen Amsler, and the NSF for this REU opportunity.

References

- [1] C. Romano, D. Brown, S. Croft, A. Favalli, L. Nakae, M. Pigni, S. Skutnik, M. S. Smith, W. Wieselquist, and M. Zerkle, *(α, n) Nuclear Data Scoping Study*, Tech. Rep. (Oak Ridge National Laboratory, 2020).
- [2] Croftl, S., *Nuclear Instruments and Methods in Physics Research Section A* (1991).

Ironing Out the Details of a Novel DeFect in FeSe

AARON SULLIVAN

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Prof. Kateryna Foyevtsova

Abstract

First principles calculations using density functional theory (DFT) to analyze ground state properties of monolayer FeSe were motivated by experimental images showing a novel defect in thin-film FeSe.[1] Three defect candidates (Se vacancy, N_2 adsorption, and O adatom), in a range of configurations were studied. DFT was then used to generate STM images for comparison to experiment, showing the O adatom to be the likely defect.

Introduction

FeSe is the simplest stoichiometric compound in the family of iron-based superconductors (FeSCS).[2] Because previous studies have shown that the presence of defects can greatly influence superconductivity in FeSe,[3] studying these defects can enhance our understanding of superconductivity in FeSCs. Accordingly, this project investigates the identity of a novel defect in 2D Fese.

This defect demonstrates current-induced hopping with directional anisotropy along the Fe–Fe bond axes, along with mutual annihilation upon collision. Using DFT to simulate STM images of the structure with a range of defects and configurations allows for the direct comparison of possible structures with experimental images. This approach is paired with total energy minimization, also using DFT, to exclude all candidates with appropriate STM images other than the structure with the lowest total energy.

Methods

All DFT calculations were done using the Vienna Ab Initio Simulation Package (VASP). Brillouin-zone integration was carried out using the Monkhorst-Pack (MP) method. The k-point sampling and ENCUT values are specified in the relevant subsections. To impose nematic order commensurate with the directional hopping anisotropy, we included spin-polarization for each structure by setting the initial magnetic moments to $\pm 5.0 \mu B$ and $0.0 \mu B$ for the Fe and Se atoms, respectively, with the spins of the Fe atoms alternating every line of Fe atoms.




Figure 1: Defect free FeSe cell showing antiferromagnetic (AFM) ordering.

Electron-ion interactions were described using the Projector Augmented Wave (PAW) method. For the electronic exchange-correlation, both the Perdew-Burke-Ernzerhof (PBE) and r^2 SCAN functionals within the generalized gradient approximation (GGA) and meta-GGA functional types were used.

Single unit cell

Convergence testing

Before starting any computationally intensive calculations, it is important to conduct convergence tests on a smaller, similar system. The general process for this is to vary one parameter at a time and perform a static calculation (fixed atomic positions) to determine the total energy of the unit cell. Once the change in energy between successive values of the test parameter is sufficiently small the calculation is considered converged with respect to that parameter.

The purpose of convergence testing is the determination of two important values in DFT calculations: k-point sampling and energy cutoff (ENCUT). K-point sampling determines the density of the grid in reciprocal space used to integrate over the first Brillouin zone. ENCUT is the kinetic

energy cutoff which determines the size of the planewave basis set used to represent the electron wavefunctions. The total wavefunction can be represented as a sum of planewaves, where accuracy and computational cost increases as ENCUT increases.[4]

We conducted convergence testing on two (orthorhombic: $\alpha = 90.0^\circ, \beta = 90.0^\circ, \gamma = 90.3016^\circ$, tetragonal: $\alpha = 90.0^\circ, \beta = 90.0^\circ, \gamma = 90.0^\circ$) defect free FeSe unit cells to determine both the number of k-points and ENCUT value needed to ensure sufficient accuracy and computational feasibility. For our calculations, k-point sampling was converged at 12x12x8 and the cutoff energy was converged at 320 eV, with $\Delta E < 10^{-3}$ for both.

STM images

In order to properly simulate experimental conditions, the unit cells used for convergence testing were altered to emulate the vacuum in an STM. The presence of the vacuum layer was modeled by increasing the size of the c lattice parameter from 5.521 Å to 26.000 Å. The density of the k-point mesh is inversely proportional to the length of the lattice parameters. Therefore, the initial 12x12x8 k-point mesh was altered to 12x12x1.

The initial STM images contained persistent visual artifacts because the system was not converged with respect to image quality. To address this, thorough comparisons between images with changes in k-mesh, applied current, van der Waals corrections, and magnetic ordering were made. We determined that the visual anomalies were caused by a combination of insufficient k-point sampling and the software (critic2)[5], [6] used to generate the images. Through another round of convergence testing, the k-point sampling for the STM images was determined to be converged at 24x24x1. The software issue was addressed by minutely shifting the coordinates of the atoms in the vacuum direction (c), which does not affect any of the calculated quantities of interest. Fortunately, the denser k-mesh only needs to be used for generating STM images.[7]

Supercells

For all supercells, k-point sampling was $2 \times 2 \times 1$ for calculating system energy and $6 \times 6 \times 1$ for generating STM images. All unit cells were 8×8 slabs. For each, we tested several high and low symmetry configurations to ensure the lowest system energy was the true ground state of the structure, rather than a local minima.

Results

Se Vacancy

Experimental STM images show the defect occupying a bottom Se site. Thus, an Se vacancy was one of our candidates. We tested three configurations using the PBE functional. After structure relaxation, the ground state (-398.589 eV) configuration is a symmetric configuration which does match the experimental STM images, but does not have the required hopping anisotropy.

N_2 Adsorption

Excluding the Se vacancy leaves only the possibility of the incorporation of external impurities as the source of the defect. Due to the abundance of N_2 in air, we tested multiple orientations of N_2 parallel to the FeSe surface, with the N_2 bonded to the Fe atoms surrounding the defect. We also tested an asymmetric configuration, with N_2 bonded to a single Fe atom. The N_2 calculations were done using r^2 SCAN functional. After relaxation, the ground state (-958.238 eV) was found to be an asymmetric configuration which does not match the experimental images.

O Adatom

Following the same logic that made N_2 a plausible choice, we also tested atomic oxygen (O). Atomic oxygen also offers an explanation for the annihilation upon collision behavior of the defect - when O atoms collide, they bond and desorb as O_2 . In addition to O adsorbed above a bottom layer Se atom, we also considered a range of interstitial O possibilities. All O adatom calculations

were done using r^2 SCAN functional. At this time, the ground state (-926.945 eV) is the high symmetry surface bonded O atom. There are several calculations still running, so this conclusion is not definitive.




Figure 2: Left column, descending: Se, N_2 , (current) O ground state configuration. Right column, descending: Se, N_2 , (current) O ground state STM.




Figure 3: Experimental STM image.

Discussion

Density functional theory candidate testing identified that surface bonded atomic oxygen is a prospective model for the observed defect, though its current ground state (-926.945 eV) remains non-definitive. This structure was evaluated alongside a selenium (Se) vacancy and molecular dinitrogen (N_2) adsorption, both of which were ultimately ruled out because their calculated ground states failed to match the full profile of the experimental data. Specifically, simulating these alternative candidates confirmed that while the Se vacancy ground state (-398.589 eV) successfully matches the experimental STM images, it lacks the required hopping anisotropy. Meanwhile, the N_2 adsorption ground state (-958.238 eV) settles into an asymmetric configuration that structurally mismatches the physical scans entirely. A clear limitation of the current dataset is that several interstitial oxygen calculations are still running, leaving the final ground-state symmetry unresolved. Consequently, future work should be determining if an alternative oxygen configuration resolves to the true ground state.

References

- [1] J. McKenzie, N. Sharma, S. Ran, and X. Liu, Electronic nematicity of FeSe revealed by anisotropic defect hopping, PowerPoint slides (2026), department of Physics, Notre Dame University.
- [2] J. Ge, Z. Liu, C. Liu, *et al.*, Nature Materials doi.org/10.1038/nmat4153 (2014).
- [3] D. Huang *et al.*, Nano Letters **16**, 10.1021/acs.nanolett.6b01163 (2016).
- [4] A. Staykov, *Ab Initio Simulations in Materials Science* (Independent, 2025).
- [5] A. Otero-de-la-Roza, M. Blanco, A. Martín Pendás, and V. Luaña, Computer Physics Communications **180**, 10.1016/j.cpc.2013.10.026 (2009).
- [6] A. Otero-de-la-Roza, E. Johnson, and V. Luaña, Computer Physics Communications **185**, 10.1016/j.cpc.2013.10.026 (2014).
- [7] VASP Developer Team, Part 3: Stm simulations (2026), developer tutorial.

Mapping Binary Star Orbits Using Direct Imaging Data from Keck Observatory and ESO

RUTHIE TAN

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISOR(S): Professor Jeffrey Chilcote, Joel Burke

Abstract

The Orbits of Moving Group Binaries (OMG Binaries) Project analyzes the orbits of binary star systems within nearby young moving groups in order to constrain the age of other young stars of that moving group. In this research project, direct imaging data obtained from the archives of Keck Observatory and the European Southern Observatory (ESO) were used to calibrate and analyze several star systems: HD 33999, NT Aquarii, HIP 12787, GJ 3305, and HIP 16853. Through calculating the separation and position angle (PA) for each binary system, future possibilities are opened up to fit the orbits of these systems.

1 Introduction

The Orbits of Moving Group Binaries (OMG Binaries) Project aims to determine the ages of young stars located in nearby moving groups by examining the binary members within those moving groups. Young moving groups are clusters of stars traveling together in space with the same velocity, and are believed to share a common origin and age, though as they grow older they drift apart and lose that association. By learning the age of one binary star system within a young moving group, the age of all the stars in the moving group can be inferred.

Age is a crucial piece of information about a star in a moving group that is necessary to determine the formation and mass of any planets within that group. Accuracy is vital in this situation: for example, an erroneous estimate of age for one star system in a moving group degrades the accuracy of the mass calculations for all the planets in that group, creating a systemic error.

In order to constrain the age of a moving group and the stars that reside within it, the project intends to study the orbits of binary stars within that group. The age constraint of a binary star system depends on orbital information that is acquired through direct imaging data and spectroscopic data [1]. One way in which the spectroscopy of a star can be examined is through the radial velocity method. Due to the gravitational influence of an orbiting star companion, the star will periodically move toward and away from Earth, which is detected in subtle blue- and red-shifts in its spectral lines. Using this method, the orbital period and size can be determined, among other features of the system. Direct imaging, on the other hand, maps the orbital path that the stars trace in the night sky and provides color information about both members of the binary pair. Using radial velocity data

alone, information cannot be extracted about the angle of inclination of the orbit or the position of the stars in the sky. Through direct imaging, points where the orbit must pass through are fixed. When direct imaging data are used in conjunction with spectroscopic data, an accurate orbit of a binary star system can be fit.

The focus of this research project is to home in on direct imaging data taken using the NIRC2 detector at Keck Observatory and NACO at the European Southern Observatory (ESO). Several binary star systems will be examined: HD 33999, NT Aquarii, HIP 12787, GJ 3305, and HIP 16853. The first four binary systems listed were all observed on December 28, 2025, at Keck Observatory, though they differ in which moving groups they belong to. GJ 3305 and HIP 12787 both belong to Beta Pictoris, a well-studied moving group that is estimated to be at an age of 26 ± 3 Myr [2]. HD 33999 belongs to the AB Doradus moving group, which has a much more controversial age estimate of 50 million to 150 million years. It is not yet certain which moving group NT Aquarii belongs to, whether it belongs to Beta Pictoris, Argus, or is a field star that is not a member of any moving group. The last binary system listed, HIP 16853, was observed at ESO on three separate nights in February and March of 2009. It belongs to the Tucana-Horologium Association, a moving group with an age of approximately 28 to 30 Myr [3]. More work remains to be done to constrain the orbits of these binary stars, especially with HD 33999 and NT Aquarii, which have highly imprecise age estimates.

Using the Python programming language, this project aims to characterize the orbits of these five binary star systems by calculating their separations and position angles (PA). Because the Keck Observatory data have not yet been analyzed through this lens, this project will open the door to orbital fitting and the age constraint of these binary systems.

2 Methods

2.1 Calibration

The entire analysis in this research project was done using Python due to its efficiency in ma-

nipulating the large arrays contained within the imaging data FITS files. Additionally, in the past decade, a portion of the astronomy community has invested significant time and resources into advancing its astronomy specific capabilities, leading Python to become the current favorable astronomy data reduction tool.

The first step of analyzing the direct image data was calibration. When light enters a telescope and strikes the detector, photoelectrons accumulate in potential wells, which form an array of pixels. The detector then counts the electrons in each pixel, with a higher electron count corresponding to a greater intensity of light [4].

However, telescopes and detectors are not perfect instruments: some pixels may be more sensitive than others, and problems such as dust and insects can often turn up in unexpected places. Similarly, the night sky is not a perfect backdrop. Cosmic ray hits can make unwanted appearances, and two stars may be close enough to where they interfere with each other's light and cannot be individually resolved. Therefore, two types of images were required to calibrate the raw images of the stars: dark frames and flat frames.




Figure 1: Binary Star System HIP 16853 Before and After Calibration Using Dark and Flat Frames

Dark frames are images taken with the cap over the aperture, so no light enters the telescope. When the telescope is in use, it heats up, emitting infrared light. This thermal energy adds a flow of additional electrons into potential wells, called dark current, which adds a certain background level, as well as creates noise in raw images. In order to correct for this, multiple dark frames of the same exposure time were taken and combined into one master dark frame, taking medians to exclude any potential cosmic ray hits or other outlying pixels. The dark rate was then calculated by dividing the master dark frame by its exposure time.

Flat frames, in contrast, are images taken of an evenly lit surface, such as the twilight sky or the inside of the dome. If there is dust in the detector or any variation in pixel sensitivity, the flat frame will show those imperfections. To use the flat frames, each flat was first dark corrected by subtracting a master dark frame with the same exposure time, then grouped together by wavelength filter bands to create master flat frames. Because the data were obtained from near-infrared detectors, the filters used were Ks, H, and J. As a last step, the master flat frame was normalized and resized to fit the raw frame.

The formula used to calibrate was

$$S = \frac{R - t\dot{D}}{F}, \quad (1)$$

where S represents the calibrated science frame, R represents the raw frame, t represents the exposure time, $\dot{D}$ represents the dark rate of the master dark frame, and F represents the master flat frame. It was important to ensure that the exposure times for the master dark frame and raw image matched, as well as the filter bands for the master flat frame and raw image.

As can be seen in Figure 1, calibration removed much of the background noise present in the raw frame on the left, leaving a cleaner image on the right.

2.2 Centroiding

Once the raw images were calibrated, centroiding was done to establish the true centers of the binary stars. A two-dimensional Gaussian fit was performed separately for each star because the




Figure 2: Demonstration of Separation and PA Using Binary Star System HIP 16853

stars in each binary system were separated enough to be resolved individually. The function used was

$$f(x, y) = Ae^{-\left(\frac{(x-x_0)^2}{2\sigma_x^2} + \frac{(y-y_0)^2}{2\sigma_y^2}\right)}, \quad (2)$$

where A is the amplitude, x_0 and y_0 are the coordinates of the center of the star, and σ_x and σ_y are the spread of the Gaussian fit. By feeding the code an initial estimate of these variables, it was able to output an optimization of these parameters via least squares fitting.

Then, using the peak of the fit as the true center of the stars, the separation and PA were calculated. As can be seen in Figure 2, the separation is the distance between the stars, and the PA is the angle measured from north to east of the primary star to the secondary star. Because astronomy

looks up from the ground towards the night sky, east is oriented so it points to the left, opposite its traditional orientation.




Figure 3: Comparison Between the Image of HIP 16853 With and Without the Defective Center

For one of the stars, HIP 16853, two of the images were defective such that the center of the brighter star was dark, so the star appeared to be a ring of light with a hole in the center, displayed in Figure 3. Likely, this was an issue of saturation. Looking at the scale, the image on the left is markedly brighter than the image on the right, and the left image had an exposure time six times longer than the right image. Therefore, image matching was required to accurately find the center of the star. By using the non-defective image as a model of the telescope's Point Spread Function (PSF) and matching it onto the defective image, the true center of the defective image was able to be found.

3 Results and Discussion

3.1 Data

Binary Star System	Date of Observation	r (mas)	σ_r (mas)	PA (deg)
GJ 3305	December 28, 2025	323.200	1.721	338.782
HD 33999	December 28, 2025	714.435	1.380	91.069
HIP 12787	December 28, 2025	144.493	0.401	0.472
HIP 16853	February 19, 2009	2704.582	0.101*	90.741
HIP 16853	February 21, 2009	2703.409	0.102*	90.869
HIP 16853	February 21, 2009	2703.122	0.102*	90.920
HIP 16853	March 14, 2009	2713.600	0.102*	90.747
NT Aquarii	December 28, 2025	87.028	1.260	184.290

Table 1: Table of the Separation, Separation Standard Deviation, and PA of the Examined Star Systems, Measured in Milliarcseconds and Degrees

*Separation Standard Error Propagation from Gaussian Fit

Using the techniques illustrated in the previous section, the separation and PA of each binary star system was found, as shown in Table 1.

For GJ 3305, HD 33999, HIP 12787, and NT Aquarii, there were over twenty usable images for each star system, all taken within the same night. To find the separation for each system, the x and y components of the separation were measured for each image, then averaged, before using the Pythagorean Theorem to calculate the r separation. However, there were only 4 images total of HIP 16853, taken over three different nights, so the separation and PA changes over time as the stars orbit each other. Therefore, to illustrate the movement of the system over time, the separation was calculated separately for each image, and the error was propagated from the standard error of the Gaussian fit.

3.2 Limitations

There was less precision in measuring the separation and PA of HIP 16853 for two important

reasons: first, there were fewer images being analyzed of the star, resulting in greater random error in the data; and second, the blurring of the light in the center of the primary star meant greater uncertainty in the true center of the star. Better quality and more imaging data would lead to more precise calculations regarding the orbit of the system. Nevertheless, the usage of orbital data from nearly 20 years ago provides a significant constraint if paired with present day data.

3.3 Future Work

Since analysis of direct imaging data has not yet been performed on these binary star systems, these new calculations will allow for future work to be done regarding the orbital fit of these star systems.

Additionally, because the age estimate of AB Doradus, the moving group of HD 33999, has an error of ± 50 Myr, fitting the orbit for HD 33999 and subsequently constraining its age will increase the precision of the age estimate of the entire moving group. Furthermore, a lack of information means NT Aquarii cannot be placed yet in a moving group. Therefore, using the data presented in this paper to constrain the age of NT Aquarii will allow more insight as to which moving group it might be a member of.

4 Conclusion

Through the calibration and centroiding of direct imaging data from the Keck Observatory and ESO, the separation and PA were found for five binary star systems, including one whose age estimate currently has a large error and one whose moving group membership is not substantiated.

Future work stemming from these results remains to be done to constrain the orbits of these systems, and thus constrain their ages, and HIP 16853 can be further examined through new imaging for a more accurate determination of its separation and PA.

Through the orbital constraints provided in this paper for these five star systems, a path has been opened to further understand the ages and planets of their associated moving groups.

References

- [1] J. Chilcote, T. Tobin, T. Currie, T. D. Brandt, T. D. Groff, M. Kuzuhara, O. Guyon, J. Lozi, N. Jovanovic, A. Sahoo, V. Deo, E. Akiyama, M. Janson, J. Knapp, J. Kwon, M. W. McElwain, J. Nishikawa, K. Wagner, K. Hełminiak, N. Skaf, and M. Tamura, *The Astronomical Journal* **162** (2021).
- [2] A. D. Thomas, E. L. Nielsen, R. J. D. 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] V. V. Makarov, *The Astrophysical Journal* **654** (2007).
- [4] F. R. Chromey, *To Measure the Sky: An Introduction to Observational Astronomy* (Cambridge University Press, 2017).

Efficiency Computation of HPGe for ^{152}Eu samples

Gabrielle L. Williams

University of Notre Dame | Nuclear Safeguard Workforce Development | Oak Ridge

Author Note

Abstract

HPGe's are detectors used for the analysis and capture of gamma ray emission data. These devices must be accurate as the data they extract can deeply disturb the value of the data obtained from running the detector. In this experiment, analyzing the efficiency of each detector over the energy of the given branching ratios allows for the confidence marking of the detector's accuracy in its results for ^{152}Eu isotope experimentation. Through this experiment, it is understood that issues with both efficiency calculation and error readings rely heavily on the strength of the beam and understanding the state of the germanium to ensure that all data collected is able to be understood and displayed as error in graphing revealed a potential issue with the germanium inside of the HPGe used for this analysis. Given this, it is imperative that the detectors be inspected and potentially ran through the same experimentation several times over to ensure the efficacy of the results given and to ensure data is well received by the software used.

Keywords: HPGe, efficiency, energy

Efficiency Computation of HPGe via Graphical Analysis

I. Introduction

The HPGe detector is the high purity germanium detector used in nuclear physics research and detects information from gamma ray emissions. The efficiency of this device is crucial to both the capture and analysis of critical data in nuclear physics research as its efficiency will drive the accuracy of the data obtained from experimental runs. HPGe is said to be resistant to signal degradation caused by changes in background radiation, shielding, interference and temperature changes as well as able to provide detailed information about the signals of interest capture through its software (Ortec). However, when analyzing particular isotopes, it is important to not simply

trust the efficiency of HPGe detectors. One should instead cross analyze with the known gamma emissions of the sample material to be certain that the results are in the range of expected energies to truly see the efficiency of the machine and the data it's collected.

In the specific case of ^{152}Eu , it is known to be susceptible to true coincidence summing (TCS) errors, which makes it a difficult piece of material to work with and requires corrective measures to ensure accuracy of its results (Necati Çelik et al.). Thus, it is imperative that human computation and analysis is performed post-HPGe to verify that the emissions found through experimentation align with the energy range and emissions found in an emission table for the given sample (^{152}Eu). Utilizing computer software reduces the amount of strain and time that performing this analysis would require (El).

II. Methods

In order to analyze and compare the efficiency of the HPGe at various emissions, an algorithm was created to calculate the efficiency of the detector using the following equation:

$$\text{efficiency} = \frac{\text{net rate}}{\text{current activity} \times \text{branching ratio}}$$

where net rate is the rate of the background is subtracted from the rate of the source, and branching ratio is the probability of the gamma emission. The rate of the background and the source was calculated with the following equation:

$$\text{rate} = \frac{\text{total counts}}{\text{live time}}$$

The algorithm was additionally used to produce histograms to visualize peaks of the counts in each run. The creation of this algorithm used a combination of previously created code as well as Gemini 3.5 Flash to break down the structure of the original codes to help utilize parts of its design in the new codes created.

Utilizing these explanations of the previous codes, a written algorithm was created, outlining the desired outcomes of both how to obtain the efficiency as well as the histograms of the runs, which are to be analyzed using the points of interest, where the photons per 100 disintegrations is >10 from a standard gamma emissions table.

	Energy (keV)	Photons per 100 disint.
Υ 1,0(Sm)	121.7817 (3)	28.41 (13)
Υ 1,0(Gd)	344.2785 (12)	26.59 (12)
Υ ,1(Gd)	778.9045 (24)	12.97 (6)
Υ 9,1(Sm)	964.079 (18)	14.50 (6)
Υ 9,0(Sm)	1,085.837 (10)	10.13 (6)
Υ 10,1(Sm)	1,112.076 (3)	13.41 (6)
Υ 13,1(Sm)	1,408.013 (3)	20.85 (8)

Table 1.

Three new codes were created with these desired outcomes and data in mind:

1. To extract data from sample and background runs to produce calibrated energy and summed energy histograms. This involved using the calibrated energy histograms created by the initially provided code, which displayed the number of counts over energy (keV),

ranging from 0 to 10,000 in order to find emission lines of interest where signal would be expected to be the strongest for the benefit of this experiment.

2. To utilize the energies obtained from the calibrated energy graphs in the form of user prompted input for the purpose of calculating the efficiency, the error, and the net calibration in the form of a multi-detector comparison graph for all the runs of each branching ratio (obtained from the gamma emission table).
3. To use the efficiency results to create scatterplots which display the relationship between the efficiency of the branching ratios over the energy of that efficiency.

A Root terminal was used to run the second code whereas a python terminal was used to obtain the scatterplots. This data was then exported from the ROOT terminal and the scatterplots from each run to GitHub as well as an internal document folder for saving.

III. Results

The efficiency of the detectors held a decay trend for all detectors, however branching ratios 28.41, 10.13 and 12.97 have runs omitted due to signal strength not being strong enough to deflect background signals.




Figure 1. Scatter plot of efficiency of each branching ratio over the gamma energy (keV) for detector 0




Figure 2. Scatter plot of efficiency of each branching ratio over the gamma energy (keV) for detector 1




Figure 3. Scatter plot of efficiency of each branching ratio over the gamma energy (keV) for detector 2




Figure 4. Scatter plot of efficiency of each branching ratio over the gamma energy (keV) for detector 3

IV. Discussion

The efficiency of each detector fits a mostly linear trend with efficiency decreasing with each detector run. The errors for one of the AFIT detectors (detector 1) and detector 3 experience large amounts of background interference, causing a lack of significant signal. This resulted in a lack of efficiency being achieved in the data results and those values were calculated to be zero further signaling this point. Despite a lack of data in the final graph, the trend still continues to align with the trends of previous graphs, though more dramatically. The errors of each point also trend similarly to the efficiency trends, with each branching ratio decreasing in a mostly linear trend with 28.41 having the higher +/- error and 10.13 having the lowest for most runs. This was mostly expected behavior given that the background of some of the runs of the branching ratios overlapped greatly with the detector causing disturbance in the signal's detected. The second highest peak in all of the graphs is consistently high, which points to their being a systematically off balanced issue with that peak. A possible explanation could be that the last Germanium had some issues, causing the fourth graph to only have three peaks.

V. Conclusion

The efficiency of the detectors follows mostly expected behaviors sans the last the detector which may have been caused by an issue with the Germanium. This is extremely relevant to the confidence and ability of the detector for future isotope research and beam-related experimentation in nuclear physics. Knowing the efficiency of the detectors is crucial to being able to conduct any kind of beam experiments as confidence in the efficiency directly related to

the confidence in the results produced by the detector. For this experiment specifically, readjusting and inspecting the detector and rerunning the experiment as a whole will help to reveal any potential issues the detector may have had as well as solidify the efficacy of the results derived from this computation. That way future replication of this experiment with a different isotope may lend to cleaner and more confident results.

Work Cited

El, Amri Lahcen. “Analysis of HPGe Gamma Spectrum Components and Instrumental Techniques for the Optimization of Spectrometric Measurements.” *Journal of Analytical Atomic Spectrometry*, vol. 41, May 2026, pp. 1602–1620, <https://doi.org/10.1039/d6ja00018e>, 10.1039/d6ja00018e.

Necati Çelik, et al. “A Combined Monte Carlo, Analytical, and Experimental Study of True Coincidence Summing Correction Factors for ^{152}Eu .” *Radiation Physics and Chemistry*, vol. 240, 2026, p. 113459, <https://www.sciencedirect.com/science/article/pii/S0969806X2500951X>, <https://doi.org/10.1016/j.radphyschem.2025.113459>.

Why High-Purity Germanium (HPGe) Radiation Detection Technology Is Superior to Other Detector Technologies for Isotope Identification Background Information on Radiation Detection and Identification.

Ultra-Precise Differential Abundances in Globular Clusters

FELICITY WUNSCH

2026 NSF/REU Program
Department of Physics
University of Notre Dame

ADVISOR(S): Dr. Evan Kirby

Abstract

Globular clusters have abundance patterns not found anywhere else in the universe, but there is no theory that explains the origin of these patterns that is consistent with the observations. This study investigated the astrophysical origins for the r-process elements (formed by rapid neutron capture), light elements, and iron-group elements in globular clusters. Previous differential abundance studies of clusters M15 and M92 found that an r-process event happened early in the cluster, but the light element patterns took millions of years to emerge. This project focused on corroborating these patterns in the more metal-rich clusters M2 and M13. High-resolution spectra of 12 stars in M2 gathered using the Keck Observatory were processed through a Python data reduction pipeline. Absorption line strengths were measured through spectral visualization tools, before incorporating differential analysis to derive chemical abundance variations. A strong Na-O anti-correlation, a hallmark of globular clusters, was observed in M2. No trend in r-process differential abundances from the first to second generation was observed in M2, indicating the pattern observed in M92 may not be universal. M2 also displays no or a weak Na-Fe correlation, which is different from the Na-Fe trends in M15 and M92, but the statistical strength of our finding is not conclusive. While this study has not yet identified the astrophysical sources of r-process and light elements in globular clusters, future analysis on this data set will use refined approaches to obtain more accurate abundances.

1 Introduction

Globular clusters exhibit unique abundance patterns whose origins remain largely unexplained by modern astrophysics [1]. This research focuses primarily on three categories of elements that make up almost all of the elemental abundances present in globular clusters: r-process elements, light elements, and iron-group elements.

Due to the old ages of globular clusters, most of the neutron capture happens rapidly. Thus, the rapid neutron capture process, or r-process, is responsible for forming all the elements with an atomic number (Z) greater than around 30, otherwise known as the heavy elements, in globular clusters [2]. The light elements have atomic numbers less than around 20 and are formed by high temperature hydrogen burning [3]. Since globular clusters were formed in the early universe, core collapse supernovae were responsible for much of the iron (Fe) and the iron-group elements found in globular clusters [1].

Several studies have been conducted that investigate the astrophysical sources of these groups of elements in globular clusters M15 and M92 [1]. Kirby et al. (2026) calculated the

differential abundances of r-process elements found in M15 and M92 by finding the difference between the elemental abundance of each star and the elemental abundance of a reference abundance for each absorption line in the spectrum. They found a relation between the light element abundances (Mg and Na) and the scatter in the differential abundances of the r-process elements. Stars that are first-generation (1G), or high in Mg and low in Na, had a wide scatter of r-process element abundances, while second-generation stars (2G) contained low Mg, high Na, and a constant amount of r-process elements in each star (Figure 1). Thus, Kirby et al. theorized that an r-process event such as a neutron star merger or magnetorotational supernova polluted the cluster with r-process elements around the time of the stars' first generation. The r-process abundances then mixed evenly throughout the cluster over time as the gas became homogenized due to dynamical mixing, causing the second generation to have no detectable r-process dispersion [1].




Figure 1: Average differential abundances of Eu versus Na. The gray squares represent stars in M15 and the red circles represent stars in M92. The gray left half of the graph indicates the first generation, while the white right half indicates the second generation. The statistics show the mean (μ) and standard deviation (σ). Adapted from [1].

Kirby et al. also discovered a relation between the light element abundances and the iron-group abundances. In M92, the light element abundances did not begin to change until Fe

was created, which puts a timescale on the formation of the light elements. This rules out several previous candidates for the type of star that was responsible for the creation of the light elements. For example, the light elements could not have originated from a very massive star because its lifetime is shorter than the core collapse supernovae that form iron, which must exist before the light elements [1].

This project investigates the effect of metallicity on the creation of these elemental abundances. The below formula can be used to calculate metallicity.

$$[Fe/H] = \log_{10}\left(\frac{N_{Fe}}{N_H}\right)_{\star} - \log_{10}\left(\frac{N_{Fe}}{N_H}\right)_{\odot} \quad (1)$$

M15 and M92 both have metallicities of around -2.4, while M13 has a metallicity of -1.53 and M2 has a metallicity of -1.65 [4]. This means that M2 is about 5.6 times more metal-rich than M15 and M92, while being about 45 times less metal-rich than the Sun (metallicity of 0 by definition). Thus, this project looks to see if the same elemental abundance trends that were observed in the metal-poor clusters M15 and M92 can be found in the more metal-rich clusters M2 and M13. Because it is more difficult to detect a dispersion in abundance when the baseline abundance is higher, the precision required for M2 and M13 is even higher than for M92.

2 Methods

High-resolution spectra of 12 stars in the M2 globular cluster gathered using the HIRES spectrograph at Keck Observatory were first run through a Python data reduction pipeline. This pipeline redshift-corrects and continuum-normalizes the spectrum and co-adds the echelle orders together. It also performs data cleaning steps such as removing bad pixels, sigma clipping outliers and high flux values, and trimming the order edges.

After this, the reduced data was read into hiresspec, a program written by Dr. Kirby that visualizes the spectrum and measures the strengths of the absorption lines, called equivalent widths. The absorption line fit was manually adjusted line by line where necessary. The equivalent width file produced by hiresspec and a model atmosphere file were then inputted into

a Fortran code called MOOG that solves the equation of radiative transfer through a one-dimensional atmosphere to find the elemental abundances [5].

Finally, differential analysis was performed. The differential abundances for each line in each star were calculated by subtracting the average abundance of that line across all 12 stars from the abundance in a certain star [1]. The average differential abundances for each element in each star were found by computing the mean differential abundance of all the lines of that element in the star. The standard deviation and error were also determined.

3 Results

A strong anti-correlation was observed between the differential abundances of the light elements sodium (Na) and oxygen (O) in M2, as shown in Figure 2. Since Na gets enhanced over time, it can be used as a proxy for time evolution [1].




Figure 2: Average differential abundances of O versus Na. The blue circles represent stars in M2.

The differential abundances of the r-process element europium (Eu) graphed versus the light element Na, which indicates the timescale, can be seen below in Figure 3. Unlike M92, the scatter in differential abundances does not show a significant overall trend from the first

generation (less Na than average amount) to the second generation (more Na than average amount) [1].




Figure 3: Average differential abundances of Eu versus Na. The blue circles represent stars in M2. The gray left half of the graph indicates the first generation, while the white right half indicates the second generation.

Also unlike M92, little to no correlation was observed between the differential abundances of the iron-group element iron (Fe) and the light element Na, as depicted in Figure 4 [1].




Figure 4: Average differential abundances of Fe versus Na. The blue circles represent stars in M2.

4 Discussion

The Na-O anti-correlation shown in Figure 2 is solid evidence of the Na enhancement and O depletion characteristic of all globular clusters [3]. These measurements are also a testament to the quality of the spectrum because O is a very difficult line to measure. No definite indications of an astrophysical source of the light elements in globular clusters can be deduced from the M2 data, but this work represents a step towards its identification.

As depicted in Figure 3, no significant pattern can be seen in the scatter in the differential abundances of the r-process element Eu over time, or as Na is enhanced. This result for M2 is different from the result in Kirby et al.'s study of M92, shown in Figure 1, which found that the scatter in differential abundances decreased from the first to the second generation [1]. This inconsistency could indicate that the pattern observed in M92 may not be a universal process, instead applying differently to all globular clusters or to more metal-rich clusters such as M2. Another explanation could be that more precision is needed to find the dispersion in M2, since its

higher metallicity and thus higher baseline abundances make variations more challenging to detect.

A weak to nonexistent correlation between the iron-group element Fe and the light element Na is displayed in Figure 4, opposed to the clear correlation between these two elements found in globular clusters M15 and M92 by Kirby et al (Figure 5) [1]. The Na-Fe trends in M2 are thus not consistent with those in M15 and M92. However, the statistical strength of our M2 findings is not conclusive.




Figure 5: Average differential abundances of Fe versus Na. The gray squares represent stars in M15 and the red circles represent stars in M92. The gray left half of the graph indicates the first generation, while the white right half indicates the second generation. The statistics show the mean (μ) and standard deviation (σ). Adapted from [1].

4.1 Research in M13

The same methods process described in the “Methods” section of this paper for cluster M2 was also completed for cluster M13, using high-resolution spectra of 5 stars gathered on the Automated Planet Finder (APF) at the Lick Observatory.

This was motivated by the fact that this study and other similar research require very precise, high-resolution abundances, with very subtle differences in abundance (less than 0.1 dex). However, observing these abundances can be costly financially, temporally, and

computationally. Due to APF's smaller size, it is cheaper and more available to book time on, enabling significant amounts of spectrographic survey data to be gathered and studied for a fraction of the cost and wait time compared to other larger observatories such as Keck.

Unfortunately, APF is limited in the stars it is able to access due to its smaller collecting area (4.5 m^2). Keck's larger collecting area (76 m^2), on the other hand, enables it to observe fainter stars lower on the red giant branch, making it easier to find many red giants with similar temperatures and luminosities. The stars Keck is able to detect are also warmer, causing the absorption lines to be weaker and less likely to overlap, making the lines more easily distinguishable. Finally, as shown in Figure 6, the quality of the data is rather low due to the high scatter and error. Due to the combination of all these limitations, including the number of stars able to be observed and the difference in temperature and luminosity between stars, the overall quality of APF data was deemed insufficient for this study, but could be sufficient for other studies requiring less precise abundances.




Figure 6: Average differential abundances of Eu versus Na. The blue circles represent stars in M13.

4.2 Conclusion

This study has not yet found clear, conclusive evidence as to decades-old questions of the astrophysical sources of the r-process and light elements in globular clusters, but the quality of the data observed and the trends and correlations look reasonable. Future research will use refined approaches to obtain more accurate abundances, including the hyperfine structure approach, which better distinguishes between close lines of the same element, and the synthesis approach, which does the same for lines of different elements.

References

- [1] E. N. Kirby, R. Gerasimov, A. Cai, B. Coco, P. Nalamwar, and L. Henderson, “Evidence of Supernova Between Formation of Stellar Populations in a Globular Cluster,” *arXiv* (2026). <https://arxiv.org/abs/2606.29435>.
- [2] J. J. Cowan, C. Sneden, J. E. Lawler, A. Aprahamian, M. Wiescher, K. Langanke, G. Martínez-Pinedo, and F.-K. Thielemann, "Origin of the Heaviest Elements: The Rapid Neutron-Capture Process," *Rev. Mod. Phys.* 93, 015002 (2021). <https://doi.org/10.1103/RevModPhys.93.015002>.
- [3] E. N. Kirby, A. P. Ji, and M. Kovalev, “r-process Abundance Patterns in the Globular Cluster M92,” *Astrophys. J.* 958, 45 (2023). <https://doi.org/10.3847/1538-4357/acf309>.
- [4] W. E. Harris, "A Catalog of Parameters for Globular Clusters in the Milky Way," *Astron. J.* 112, 1487–1488 (1996). <https://doi.org/10.1086/118116>.
- [5] C. Sneden, "The Nitrogen Abundance of the Very Metal-Poor Star HD 122563," *Astrophys. J.* 184, 839–849 (1973). <https://doi.org/10.1086/152374>.

Direct Studies of 610 and 559 keV Resonances in $^{35}\text{Cl}(\text{p},\alpha)^{32}\text{S}$

NOAH WYATT

2026 NSF/REU Program
Department of Physics and Astronomy
University of Notre Dame

ADVISORS: Prof. Dan Bardayan, Prof. Patrick O'Malley

Abstract

In classical novae, the production of intermediate-mass nuclei in the SCl cycle is constrained by the competition between $^{35}\text{Cl}(p,\alpha)^{32}\text{S}$ and $^{35}\text{Cl}(p,\gamma)^{36}\text{Ar}$. To understand this competition, two low-energy resonances in the $^{35}\text{Cl}(p,\alpha)^{32}\text{S}$ reaction at $E_{c.m.} = 610$ and 559 keV were directly studied. The resonance strength for the $E_{c.m.} = 559$ keV resonance was measured for the first time, the lowest energy resonance that has been directly studied. This experiment was carried out using the 5U accelerator at the University of Notre Dame. Our results and relevance to classical novae will be presented.

1 Introduction

White dwarf stars in binary systems can accrete hydrogen-rich material from their partner, leading to thermonuclear-powered classical novae. These stellar explosions reach peak temperatures of $T \leq 0.4$ GK. Observations of classical novae have shown increased concentrations of intermediate-mass nuclei in the Si-S region, which can only be explained by explosive hydrogen burning. This hydrogen burning is characterized by several key nuclear reaction cycles, which lead to the production of heavier elements [1]. Of particular interest to the production of intermediate-mass nuclei is the SCl cycle. This reaction is closed by the $^{35}\text{Cl}(p,\alpha)^{32}\text{S}$ reaction. Whether the reaction gets out of the SCl cycle is dependent on the competition between the (p,α) and (p,γ) reactions of ^{35}Cl . Sensitivity studies have shown that varying the strength of the $^{35}\text{Cl}(p,\alpha)^{32}\text{S}$ reaction rate has a significant effect on the production of nuclei in the S-Ar region [2]. This makes it clear that understanding the $^{35}\text{Cl}(p,\alpha)^{32}\text{S}$ reaction is important for modeling explosive hydrogen burning in classical novae.

Despite its importance, only one direct measurement of the (p,α) cross section for this reaction has been performed (see Table 1) [1]. Additionally, this study was only able to measure the cross section of the 610 keV resonance, and the yield was too low to determine the angular distribution leading to a greater uncertainty in the extracted resonance strength. Further direct measurements of the $^{35}\text{Cl}(p,\alpha)^{32}\text{S}$ reaction rate were needed to reduce this uncertainty and to measure the strengths of lower-lying resonances for the first time.

Table 1: Resonance energies and strengths for the $^{35}\text{Cl}(p,\gamma)^{36}\text{Ar}$ and $^{35}\text{Cl}(p,\alpha)^{32}\text{S}$ reactions in Gamow Window. Values are taken from ref. [3] except where noted.

E_r (MeV)	J^π	$\omega\gamma_{p,\gamma}$ (eV)	$\omega\gamma_{p,\alpha}$ (eV)
0.049	2^+	$(2.8 \pm 1.2) \times 10^{-24}$	$\leq 2.0 \times 10^{-25}$
0.165	1^-	$(3.4 \pm 1.6) \times 10^{-10}$	$\leq 4.0 \times 10^{-11}$
0.302	$(1, 2, 3^-)$	$(1.5 \pm 0.5) \times 10^{-5}$	$\leq 1.6 \times 10^{-6}$
0.380		$\leq 3.1 \times 10^{-5}$	$\leq 6.2 \times 10^{-6}$
0.404	2^+	$(6.6 \pm 3.3) \times 10^{-5}$	$\leq 1.0 \times 10^{-3}$
0.416		$\leq 1.3 \times 10^{-4}$	$\leq 7.8 \times 10^{-6}$
0.433	$(2^+, 3^-)$	$(1.4 \pm 0.3) \times 10^{-2}$	$\leq 3.2 \times 10^{-4}$
0.508	$(3^-, 5^-)$	$(1.2 \pm 0.3) \times 10^{-3}$	$\leq 4.5 \times 10^{-5}$
0.518	2	$(3.2 \pm 0.7) \times 10^{-2}$	
0.559	3^-	$(3.1 \pm 0.6) \times 10^{-2}$	$\leq 7.2 \times 10^{-3}$
0.610	1^-	$(7.6 \pm 1.8) \times 10^{-4}$	$(1.2 \pm 0.2) \times 10^{-2}$ ^a

^a Moazen *et al.* [1].

2 Methods

To study this reaction, a beam of protons was created by the 5U accelerator at the Nuclear Science Laboratory located on the campus of the University of Notre Dame. This accelerator is a single-ended vertical pelletron using four charging chains to reach a maximum voltage of 5 MV. The beam of protons had a current of ~ 40 nA, measured with a beam dump, and were impinged on a solid natural KCl target of thickness $96 \mu\text{g}/\text{cm}^2$ placed in an aluminum target ladder in a vacuum of 5×10^{-6} Torr. The α particles were detected with six silicon photo diode detectors at backwards lab angles relative to the beam axis of $\pm 120^\circ$, $\pm 140^\circ$, $\pm 160^\circ$. Each detector was located 115 mm from the target and had an aperture diameter of 5 mm. The setup is shown in Figure. 1. All six detectors were calibrated using a mixed α source of ^{241}Am and ^{148}Gd . The solid angle was measured to be 1.44×10^{-3} sr and the geometric efficiency was 2.29×10^{-4} . The detectors were biased at 50 V using the mesytec MVH-4 supply and the MSI-8 preamplifier and output pulses were measured using the ComTec FastDAQ. Data were gathered by the software MPANT. Expected α particle energies were acquired from CATKIN and energy loss inside the target was estimated by stopyt. The proton energy loss was not measured due to the resonance strength being very low at (0.99 ± 0.08) eV [1]. For α particle energy loss, the decays were assumed to take place in the very center of the target. This gave an approximate thickness for the 160° detectors to be $51.08 \mu\text{g}/\text{cm}^2$, the 140° detectors

to be $62.66 \mu\text{g}/\text{cm}^2$, and the 120° detectors to be $96 \mu\text{g}/\text{cm}^2$. The data analysis is being carried out using ROOT.




Figure 1: A schematic diagram of the experimental detection chamber setup is shown.

3 Results

Figure 2. shows the resonance scan over the 559 and 610 keV resonances. This scan shows that we were able to find the resonances and that the 610 keV resonance is stronger than the 559 keV resonance, which was expected.




Figure 2: Resonance Scan of 559 keV and 610 keV resonances. Error bars given as $\sqrt{N}$.

Figure 3. shows the expected α -energies at our 610 keV resonance. In the KCl target, we would expect to see α particles emitted from ^{35}Cl , ^{37}Cl , ^{39}K , and ^{41}K . Additionally, we would expect to see some $^{27}\text{Al}(p,\alpha)$ from the target ladder and $^{18}\text{O}(p,\alpha)$ from the target oxidizing between creation until it was used. Most of these reactions were off resonance and the energy does not align with the events of interest so they can mostly be ignored. Figure 3. shows an off resonance run. We see very low counts in our region of interest meaning, our background is seemingly low.




Figure 3: Spectrum resulting from $E_{lab} = 639 - 659$ keV at 120° . Regions of expected α particles are highlighted with ^{35}Cl peak in orange.




Figure 4: Off resonance spectrum resulting from $E_{lab} = 607$ keV at 120° .

4 Next Steps

Data analysis is ongoing, and will continue for the coming months. Off-resonance measurements will be studied in order to understand the background we may be seeing in our detectors. This will be subtracted giving us our proper number of counts to be used to find the cross section. With the cross section, we will use the program AZURE2, an R-Matrix software to give the angular distribution of the cross section. Once the cross section and resonance strength are calculated, we will determine the astrophysical implications of these new measurements in classical novae.

References

- [1] B. Moazen *et al.*, Eur. Phys. J. A **47**, 66 (2011).
- [2] C. Illiadis *et al.*, Astro. Jour. Supplement Series **105**, 142 (2002).
- [3] J. Ross *et al.*, Phys. Rev. C **52**, 1681 (1995).

Simulation-Based Inference of Population Size and Epistasis from Evolutionary Hysteresis

HANQING ZHAO

2026 REU Program

Department of Physics

University of Notre Dame

ADVISOR(S): Prof. Derviş Can Vural

Luke Piszkin

Abstract

Evolutionary hysteresis occurs when a population follows different trajectories as an environmental variable is cycled forward and backward. We tested whether a stochastic hysteresis loop can constrain population size and epistatic interaction strength. Loops were generated using a two-locus Wright–Fisher model across population size, epistasis, and known mutation rate, then represented by geometric features and observation-protocol descriptors. Conditional Gaussian and RealNVP models were trained to estimate the joint parameter posterior. On held-out synthetic test data, the Gaussian model achieved mean absolute errors of 0.158 in log-transformed population size and 0.0054 in epistatic strength, while RealNVP achieved a slightly better joint density fit. Posterior predictive checks reproduced the principal loop structure with approximately nominal coverage. Performance was weakest at the lowest mutation rate. These results demonstrate that evolutionary hysteresis contains recoverable information about hidden parameters within the simulated model.

1 Introduction

Evolutionary outcomes depend on both the current environment and the structure of the underlying fitness landscape. Epistatic interactions can create alternative favorable states separated by fitness barriers, causing a population’s response to depend on its previous history. When an environmental variable is cycled forward and backward, this history dependence can produce an evolutionary hysteresis loop. Piszkin and Vural characterized this behavior in a two-locus Wright–Fisher model and related evolutionary systems [1], suggesting that loop morphology may reveal otherwise hidden properties of the population and fitness landscape.

This possibility creates an inverse problem. Population size N controls the magnitude of genetic drift, while epistatic interaction strength J affects the barrier between alternative population states. Mutation rate u also changes the timing and probability of transitions. Loop area, branch separation, and transition structure may therefore carry information about N and J . However, identical parameters can generate different stochastic trajectories, while different parameter combinations can produce similar loops.

For such simulators, the likelihood of an observed trajectory may be unavailable or impractical to calculate. Simulation-based inference instead learns parameter–observation relationships from generated data [2–4], while neural posterior estimation produces distributions over plausible parameters and their uncertainty [4,5]. Here, we train Conditional Gaussian and RealNVP estimators on synthetic two-locus hysteresis loops and evaluate their ability to recover N and J using held-out parameter recovery, calibration, and predictive validation.

2 Model and Inference Framework

Figure 1 summarizes the pipeline from stochastic loop simulation and observation protocols to posterior estimation and held-out validation. Throughout this report, p14 denotes the frozen model-and-dataset configuration used for formal evaluation.




Figure 1. EH-v3 workflow from stochastic Wright–Fisher loops to joint posterior estimation and validation.

2.1 Evolutionary hysteresis model

We used the haploid two-locus, two-allele Wright–Fisher model of Piszkin and Vural [1]. The four haplotypes have relative fitnesses

$$\omega_{11} = 1 + 2s, \omega_{12} = \omega_{21} = 1 + s - J, \omega_{22} = 1,$$

where s is the environmentally controlled selective advantage and J determines the epistatic barrier between the pure haplotypes. Mutation permits allele transitions, while finite population size N controls genetic drift. For each (N, J, u) , s was swept from -0.5 to 0.5 and reversed, producing a stochastic hysteresis loop. Because repeated simulations yield different trajectories, the inverse model estimates a parameter posterior rather than a deterministic solution.

2.2 Loop representation and observation protocols

Each latent loop comprised two 181-point branches. Six observation protocols represented complete, nonuniform, asymmetric, randomly missing, continuously gapped, and edge-truncated

sampling (Figure 2). Eight geometric loop features, ten protocol descriptors, and the known $\log_{10}u$ formed a 19-dimensional input vector for inferring $\theta = (\log_{10}N, J)$




Figure 2. Six sampling protocols applied to the same latent hysteresis loop. Blue circles and orange open squares denote the down and up branches, respectively.

2.3 Conditional posterior estimation

We trained two estimators for $q_{\phi}(\theta | x)$: a correlated two-dimensional Gaussian and a conditional RealNVP normalizing flow [6]. Both were trained by minimizing the negative log-likelihood of the simulated parameters under the predicted posterior. Preprocessing, feature definitions, and model selection were frozen using validation data before the p14 test set was opened.

2.4 Evaluation

Evaluation combined parameter-recovery errors, physical-space negative log-likelihood, empirical

credible-region coverage, fixed- u benchmarks, and posterior predictive checks.

3 Results

3.1 Parameter Recovery and Posterior Model Comparison

Both models recovered information about N and J on the held-out p14 test set (Table 1). The Conditional Gaussian achieved lower point-estimation errors, with MAEs of 0.1577 in $\log_{10}N$ and 0.00540 in J . RealNVP achieved a slightly better physical-space negative log-likelihood but larger parameter errors.

Table 1. Performance on the held-out p14 test set

Model	Physical NLL	MAE in $\log_{10}N$	MAE in J	Normalized MAE
Conditional Gaussian	-3.9063	0.1577	0.00540	0.04618
Conditional RealNVP	-3.9434	0.1690	0.00564	0.04930

RealNVP was selected using frozen validation NLL before the test set was opened. The comparison shows that improved posterior density fit does not necessarily imply improved point estimates and that the Gaussian approximation captured most of the recoverable structure in this two-parameter problem.

3.2 Robustness, Calibration, and Predictive Validation

For the observation protocols illustrated in Figure 2, explicitly encoding the sampling process improved robustness. The protocol-aware model reduced the normalized error score from 0.1847 to 0.1122, corresponding to an improvement of approximately 39%.

Inference performance varied systematically across held-out fixed- u slices (Figure 3).




Figure 3. Performance across held-out fixed-mutation-rate slices. Both posterior estimators were least accurate at the lowest mutation rate. RealNVP generally achieved a slightly better physical-space negative log-likelihood, whereas the Conditional Gaussian model produced lower parameter-recovery errors.

Performance improved substantially above $u = 3 \times 10^{-4}$ and remained relatively stable across the three higher-rate slices, demonstrating that the lowest-rate regime was the principal source of cross- u difficulty.

Posterior predictive checks assessed whether inferred parameters regenerated the observed loop dynamics. In a representative continuous-gap case, both posterior medians reconstructed the

principal branch transitions despite incomplete sampling (Figure 4).




Figure 4. Posterior reconstruction from an incomplete hysteresis observation. Panels A and B show the complete latent loop and its continuously gapped observation; panels C and D show median loops generated from the Conditional Gaussian and RealNVP posteriors.

Across the test set, both models achieved approximately nominal 90% pointwise coverage. The Gaussian model produced lower predictive error, with predictive energy 8.8% above the true-parameter oracle compared with 15.6% for RealNVP.

4 Discussion

Evolutionary hysteresis retained information about both population size and epistatic interaction strength. Epistasis systematically alters branch separation, transition locations, and loop area, whereas population size affects stochastic variability through genetic drift. This distinction likely

explains why J was recovered more accurately than N : systematic geometric changes are easier to identify from one loop than stochastic fluctuations.

The two posterior estimators performed similarly. RealNVP achieved a slightly better density score, while the Gaussian model produced lower point-estimation and predictive errors. The limited advantage of RealNVP suggests that a correlated Gaussian captured much of the posterior structure in this two-parameter problem. Protocol descriptors improved robustness only for sampling patterns represented during training, and Figure 3 shows that mutation-conditioned inference remained least reliable at the lowest u .

The study is limited by its two-locus fitness model, known mutation rate, fixed summary representation, and evaluation on synthetic data generated by the same simulator used for training. The results therefore establish identifiability within this controlled model rather than direct empirical parameter recovery.

5 Conclusion

We developed a simulation-based inference framework for estimating population size and epistatic interaction strength from evolutionary hysteresis loops. Within the two-locus Wright–Fisher model, both conditional Gaussian and RealNVP posterior estimators recovered meaningful information about N and J from stochastic loop features. The Gaussian model produced more accurate point estimates and predictive results, while RealNVP achieved a slightly better joint density score. Explicitly representing the observation protocol improved robustness to simulated sparse and irregular sampling, although performance remained dependent on mutation rate and on the range of conditions represented during training.

These results demonstrate that evolutionary hysteresis can serve as a statistical probe of hidden population-genetic parameters within a controlled simulation model. They also show the importance of reporting posterior uncertainty rather than relying only on point estimates. Extending this framework to more complex fitness landscapes and experimentally matched observation models will be necessary before applying it to empirical evolutionary systems.

References

- [1] L. Piszkin and D. C. Vural, “Evolutionary Hysteresis: Cycling about in a Rugged Landscape,” arXiv:2511.14090 (2025).
- [2] S. Sheehan and Y. S. Song, “Deep Learning for Population Genetic Inference,” *PLoS Computational Biology* 12, e1004845 (2016).
- [3] K. Cranmer, J. Brehmer, and G. Louppe, “The Frontier of Simulation-Based Inference,” *Proceedings of the National Academy of Sciences* 117, 30055–30062 (2020).
- [4] G. Papamakarios and I. Murray, “Fast ϵ -Free Inference of Simulation Models with Bayesian Conditional Density Estimation,” *Advances in Neural Information Processing Systems* 29 (2016).
- [5] J. Min, Y. Ning, N. S. Pope, F. Baumdicker, and A. D. Kern, “Neural Posterior Estimation for Population Genetics,” *Genetics* 233, iyag107 (2026).
- [6] L. Dinh, J. Sohl-Dickstein, and S. Bengio, “Density Estimation Using Real NVP,” *International Conference on Learning Representations* (2017).