

NOvA *as a*  
SUPERNOVA OBSERVATORY

Justin A. Vasei

Submitted to the faculty of the University Graduate School  
in partial fulfillment of the requirements  
for the degree  
Doctor of Philosophy  
in the Department of Physics,  
Indiana University  
2021

Accepted by the Graduate Faculty, Indiana University, in partial fulfillment of the requirements for  
the degree of Doctor of Philosophy.

Doctoral Committee

Mark Messier, Ph.D.

Rex Tayloe, Ph.D.

Stuart Mufson, Ph.D.

Charles Horowitz, Ph.D.

15 December 2021

Copyright © 2021  
Justin A. Vasel







*To Duffy "Moocow Craken," whose soul outshone galaxies.  
2002–2015*



# ACKNOWLEDGMENTS

This thesis has been eight years in the making if counting from the moment I joined the program at Indiana University—nearly a quarter of my life so far. As far back as I can remember, my academic performance on paper has been at best mediocre and at worst insufficient. But in the twenty years since I decided I wanted to be a physicist during my freshman year of high school, persistence has carried me through from one stage of my career to the next. Personally, this thesis represents an enormous milestone that at times I thought might never come. I am proud of what I’ve accomplished so far in my career as a physicist and particularly during my time at Indiana University, but I did not achieve this alone. So much credit goes to the many individuals and groups who have supported me throughout my journey. I am certain that I would not be where I am today without the encouragement, friendship, love, and support of countless others. There are many people who have contributed to my success over the years, and I will take some time in this section to acknowledge them, but the number is so large that I will inevitably leave some people out. If you, dear reader, are one of those people, (1) thank you for reading this thesis, and (2) reach out to me and I’ll buy you a drink of your choice to show my appreciation.

First, I would like to thank my thesis committee. Rex Tayloe was the first IU faculty member that I spoke with, and as my academic advisor and director of graduate studies for the department, he has submitted many a form on my behalf and is always quick to respond whenever I send him a frantic email about <insert arbitrary academic requirement here>. Chuck Horowitz somehow managed to shepherd me through the crucible that is J. D. Jackson’s *Classical Electrodynamics* and I lived to tell the tale. I still don’t know how he did that. Stuart Mufson provided useful input on my updates during our research group meetings at a time when I barely knew what I was doing and could use all the input I could get. I am grateful to have learned from them. I also wish to thank Joh Urheim who provided useful feedback on my work during our regular research group meetings.

A heartfelt thank you to my advisor, Mark Messier. Of everything I’ve learned over the past eight years, a significant amount came directly from him. Mark taught me how to think like a physicist, and specifically like a particle physicist. Since I began working with him, I have always admired his “at-the-fingertips” knowledge. This encouraged me to think beyond the nuances of high-energy physics—which can easily become overwhelming—and to focus more on the fundamentals when it came to asking questions and developing a research strategy. Mark taught me how to present information in a plot: minimize the amount of ink used and maximize the information conveyed. I tried to implement that lesson in this thesis; I hope I was successful. You, dear reader, can be the judge of that. Either way, it’s a lesson—among many others—that I will carry with me through the rest of my career. In addition to being a good teacher, Mark is a great advisor. He has always been patient with me, even when progress

was slow, and it is clear that he cares deeply for the well being of everyone in our research group.

Before I came to Indiana, I spent two delightful but frigid years at the University of Minnesota Duluth in pursuit of a master's degree. My advisor, Alec Habig, taught me how to navigate a linux command line for the first time, and has continued to mentor me ever since. Even in my most unconfident moments and on days when I had little or nothing to show during updates about my work, Alec would offer encouragement and support. He never lost hope in me, so I never lost hope in myself. I learned to take a step back and see the bigger picture. If it were possible to enumerate everything I have learned from him, I would. But alas... All I can say is that I wouldn't be half the scientist I am today without his mentorship, and I likely wouldn't be writing this thesis on a ThinkPad X220 either.

I am grateful to have been a member of multiple collaborations in the neutrino community. Through the SNEWS and HALO collaborations I have had the privilege of working with many talented individuals. In particular, I wish to thank Clarence Virtue and Kate Scholberg for their support and mentorship.

The NOvA collaboration will always hold a special place in my heart. I am grateful to call many of my NOvA colleagues friends. Thank you to the NOvA Exotics working group for providing a wealth of useful discussions and suggestions. I wish to thank Matt Strait whose eye for detail and keen intuition of the underlying physics made this thesis much stronger than it otherwise would have been. This thesis would also look much different if it weren't for the incredible efforts of Andrey Sheshukov, the other doctoral student working on supernova research at NOvA who designed the simulation and triggering infrastructure on which my analysis depends. It was a pleasure working on Team Supernova with him.

Thank you to Andrew Norman for teaching me about the timing and DAQ systems at NOvA. I asked Andrew many questions and he always provided enthusiastic and insightful answers. I learned a lot about computing and programming through his example. I wish to thank Peter Shanahan for his support and friendship over the years, as well as other members of NOvA's leadership with whom I've worked closely, Tricia Vahle and Jeff Hartnell. It was a pleasure.

I owe many thanks to Louise Suter, who was patient with me in my early days on NOvA while I learned the ropes in the Data Quality group. We have worked closely for many years in a variety of contexts, from detector operations and maintenance to data quality to advocacy work. Louise has created many opportunities for me to develop my skills both technically and professionally, and I am grateful to her for that and for her friendship.

Joseph Zennamo never fails to make me laugh. Like, even for a second. Seriously, let me catch my breath. I know he is very happy for me earning this degree, but I suspect he is equally sad for me knowing that I will never perform another DSO scan.

Thank you to Young NOvA, an organization that is near and dear to my heart. I am humbled to have had the opportunity to lead it for a year, and it has been wonderful to watch how it has grown and evolved since I first joined in 2014. I have made dozens of friends who made working in this field an absolute blast. I am grateful for everyone in Young NOvA, past, present, and future, and will miss attending Young NOvA Nights.

I wish to thank Fermilab, which is a gem of an institution. I had the pleasure of spending half a year living in the Fermilab Village and it was an incredible experience that I will never forget. The beautiful prairie lands vistas, endearing bison heard, and unique sculptures and architecture never fail to take my breath away. I still get chills whenever I walk into the Wilson Hall atrium. But the best part is the

community. Thank you to Tania Claudette for creating a fun and inclusive environment at the Users' Center.

I have been able to get to know many members of the high-energy physics community through social media, particularly Twitter. Many of them I have not yet met in person, but they bring joy and positive energy into my daily life regardless. Shout out to the physics Twitter fam!

Thank you to Jost Migenda for the neutrino fidget spinner, and thank you to Linda Cremonesi for bringing it to me from across the ocean.

Another community that I am thankful to have been a part of is that of my Ph.D. cohort. We spent many long days and nights together in the office stressing over homework problems and studying for the quals. Thank you to Jon Zarling and Bill Solyst for their friendship. Thank you to Sam Cohen and Matt Caplan, who have both supported and encouraged me from the first to the last days of this process.

A little more than a year ago, I began a new career in the national security space. I do not recommend starting a full time job in a new career field while simultaneously writing a doctoral thesis, but thanks to the support of my new friends and colleagues, I somehow made it work. Thank you to those I worked with in the Technology and National Security Fellowship: Anna Melker, Jenn Robinson, Chris Painter, Jorge Narvaez, Naren Chaudhry, Joe Schuman, and Nourhan Ibrahim. Thank you to my colleagues at the Department of the Air Force who have also supported me since I joined the operational energy team: Bert Guerrero, Dan Pike, Miguel Valle, Jordan Eccles, James Olden, Andy Umanzor, and Shelly Black.

Thank you to my beloved cat, Ash, who kept me sane during the pandemic lockdown and helped with the writing of this thesis by laying on my keyboard at the most inopportune `zzxhkzccccccz`.

I can't spend too much time writing from home before I start to go insane. The staff at Switchyard Brewing Company in Bloomington knows this—I have spent many a post-vaccination evening with them alternating between reading a book and furiously typing this thesis on my tiny laptop while sipping the finest craft beers that Bloomington has to offer. I consider many on the Switchyard crew to be friends, and I am grateful for their support and also for the business which contributes so much to the culture that makes Bloomington such a wonderful town to live in.

I feel fortunate to have been surrounded by a wonderful group of friends in the neutrino research group at Indiana University ("Team IU"). The group has evolved quite a bit since I originally joined. I got to know Michael Baird and Evan Niner while they were still graduate students here. They are inspirations and good friends. I owe Evan in particular a great debt of gratitude for passing the torch and transferring the knowledge of maintaining the NOvA timing system.

I am grateful to Ryan Murphy as a colleague and a friend. Even when we were tired and troubleshooting data quality issues, we always had a good time. Bruce Howard, Erin Ewart, Erica Smith, Ashley Back, and Maria Manrique made our shared office space—or any space, really—an absolutely delightful place to be. I could fill this space with many cryptic references to inside jokes, but in the interest of not annoying everybody else, I will refrain. I am thankful to them all for their friendship over the years. A special shout out to Maria for feeding and spoiling my cat while I am out of town.

There is no doubt that the work presented in this thesis would have taken even longer if it weren't for Gavin Davies. Gavin is one of those people who puts in a lot of effort behind the scenes to make sure things continue to run smoothly for everyone else. That can be thankless work, but not today.

Gavin has helped me on countless occasions to resolve computing and programming issues and was always quick to offer sound advice on everything from my code to my career, and I am thankful for that. He is also a great friend who brings joy and humor to any situation.

Micah Groh is another Team IU alumn whom I feel fortunate to call my friend. As a machine learning expert, I sought his advice on several occasions which ultimately influenced the analysis strategies I pursued for this thesis. Micah also makes a great travel partner. Our adventures have taken us from the cliffs of western Ireland to an historic World Cup football match in France to a perilous trek through Hurricane Pass in Grand Teton National Park. Adventures like these are perhaps in some way analogous to the treacherous path one takes to obtain a Ph.D. In either case, I am glad to have shared parts of that path with Micah.

I cannot imagine taking on this Ph.D. journey without Teresa Lackey. As operations, DAQ, and data quality experts at NOvA, we have teamed up to solve many a problem, and we make a great team. Beyond academics, Teresa is an incredible friend. She has been a vital part of my support system for the majority of my time in this program, and I can only hope that I've returned the favor. Teresa has inspired me in many ways, and I am forever grateful to her. Also, her cat Mindy is an angel.

Thank you to my dear friend Brigitte Botton, whom I met long before I decided to go to graduate school. We pursued different careers and rarely find ourselves in the same city, but when we do, it's like we never spent any time apart. Much has changed between when we used to hang out in the dorms at Winona State University and now, but the strength of our friendship has not. I'm thankful to her for her support throughout the years.

Conducting research and writing the thesis are two very different things, and Pruthvi Mehta saw me through most of the latter. Always quick with encouragement, Pruthvi kept me grounded while I attempted to put the finishing touches on a thesis during the height of an historic global health crisis. I don't know how everyone else who graduated in the past two years did it, but I know that Pruthvi's love and support was a major part of how I did.

History was made in August of 2011 when I met Fernanda Psihas in Duluth, Minnesota. Since that time, we have been inseparable. Over the past ten years, so much has changed, but Fernanda has been there for me every step of the way. We are best friends and we share everything—including thesis advisors, apparently. I have always admired her scientific acumen, and sought to emulate her example. I could say that I would not be where I am now without her, but that would be an understatement. She is the reason I ended up at Indiana University, and she is reason I chose to join the NOvA experiment. We earned our master's degrees together, and although it took me longer to get my doctorate degree, she always considered it fact that I wouldn't be far behind, even when it felt more like fantasy to me. Regardless of our phase difference, there's no one I'd rather be propagating through this life with than Fernanda Psihas, and I am eternally grateful for her everlasting support.

A little over twenty years ago, someone decided to make friends with a nerd obsessed with *The X-Files*. Today, Matthew Reynolds is friends with an X-Files fan obsessed with being a professional nerd. Matt has been a brother to me for most of my life, and he has been an invaluable pillar of support while I pursued my Ph.D. I always felt like I kept pushing back my timeline. "No, I'm not going to graduate this year unfortunately, but hopefully next year," I would say. He would always respond the same way: "Don't worry about it. You know what you're doing and you're working super hard; You are going to

do this.” At times, it was hard to believe in myself, but after Matt’s pep talks, it was hard not to. The story of our friendship is one of immense personal growth for each of us; at long last this chapter is closed. What will the next bring?

I wish to express my lifelong appreciation for my aunt and uncle, Deb and Dan Freund. I was barely three years old when both my parents received a letter informing them that the U.S. Air Force was calling them back to duty during the Gulf War. My aunt and uncle took me in for several months while my parents served. Since then they have continued to support me in my personal and professional life, and for that I will always be grateful.

I am the oldest sibling of three by eight years. When my brother and sister, Alex and Paige Vasel, were young, I used to teach them everything I could about science; why the sky is blue, how weather systems worked, what causes a helium balloon to rise. You name it. My hope at the time was that they would be so amazed by the science like I was that they might want to pursue a career in science too. Well, they both grew up to be their own non-scientist individuals, and I couldn’t be more proud of them. And it turns out they couldn’t be more proud of me either. I am deeply grateful that my siblings and I can all share in each other’s successes and cheer each other on.

My mother, Jennifer Olson, has always been a source of inspiration for me. As a child, she taught me the importance of empathy, and as an adult, she taught me the value of being a team player and doing the right thing. She has always been able to strike this balance of caring for the people around her while also striving for her own goals. Once she completed her graduate program to become a nurse practitioner, I thought “well, if she decides to become a doctor she’ll probably get there before I do!” She probably would have too; she’s a hard worker. I learned that work ethic from her and applied it to this Ph.D. effort. I am proud of my mom, and I know she is proud of me. I could not have done this without her love and support.

If somebody asked me what made me want to become a scientist in three words or less, I would respond “My dad, definitely.” Steve Vasel instilled in me an appreciation for the wonders of the Universe at a young age. Growing up, I remember paging through books in our home with beautiful photos from NASA spacecraft and telescopes. I was mesmerized by the wispy clouds of Jupiter, enchanted by the granular texture of the sun. When I decided I wanted to be a physicist after reading “A Brief History of Time” by Stephen Hawking in high school, he was all for it. He has been my most vocal cheerleader ever since, and he never misses an opportunity to tell me how proud he is of me. I know that my earning this degree is in large part thanks to his unwavering support.

I have addressed here only a handful of individuals that have been a part of my journey, but there are many others. While I am deeply proud of accomplishing what has been a goal of mine for so many years, I recognize that those who supported me set the stage for my success. To everyone who has cheered me on along the way: thank you.





Justin A. Vassel

NOvA *as a*  
SUPERNOVA OBSERVATORY

Extra-solar system neutrino astronomy was born in February 1987 when a supernova in a nearby satellite galaxy deposited a couple dozen neutrino events across three solar neutrino and nucleon decay experiments. This marked the first-ever observation of supernova neutrinos, as well as the first supernova in our galactic neighborhood visible to the naked eye in nearly 400 years. There is still much we do not know about the dynamics of these powerful explosions, but one thing is certain: neutrinos play a central role in driving a supernova. They also provide a unique opportunity for probing the interior conditions of a collapsing stellar core. Since 1987, the number of neutrino detectors around the world has grown. When the next supernova occurs, detectors must be ready to seize the moment and record data from this rare event that only happens a couple times per century. The NOvA experiment has been prepared for this for some time, but it has been unclear how well NOvA will be able to separate supernova-like events from the large cosmogenic backgrounds and extract meaningful physical insight. In this thesis, I show that the cosmic-induced backgrounds in the NOvA detectors—especially the far detector—can be reduced to an acceptable level for core-collapse supernovae within  $\sim 15$  kpc and that a determination of the neutrino mass ordering can be made for nearby supernovae under certain conditions. This thesis paves the way for future sensitivity studies and eventual analyses of real supernova data, and shows that NOvA’s liquid scintillator neutrino detectors are a capable supernova observatory.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>I</b>	<b>NEUTRINOS &amp; SUPERNOVAE</b>	<b>5</b>
<b>2</b>	<b>The Neutrino</b>	<b>7</b>
2.1	A brief history . . . . .	7
2.2	Neutrinos Physics . . . . .	13
<b>3</b>	<b>Neutrinos from Core-Collapse Supernovae</b>	<b>17</b>
3.1	Physics of a collapsing core . . . . .	19
3.2	Neutrino emission . . . . .	24
3.3	Simulations . . . . .	25
3.4	Looking ahead . . . . .	27
<b>4</b>	<b>Supernova Neutrinos at Earth</b>	<b>29</b>
4.1	Supernova 1987A . . . . .	30
4.2	Supernova signal characteristics . . . . .	30
4.3	Flavor transitions . . . . .	32
4.4	Supernova neutrino detection technologies . . . . .	35
4.5	Physics opportunities . . . . .	37
4.6	Supernova Early Warning System and global readiness . . . . .	38
<b>II</b>	<b>NO<math>\nu</math>A AS A SUPERNOVA OBSERVATORY</b>	<b>43</b>
<b>5</b>	<b>The NO<math>\nu</math>A Experiment</b>	<b>45</b>
5.1	Detector technology and design . . . . .	46
5.2	Data acquisition . . . . .	53
5.3	Timing system . . . . .	54
5.4	Triggering system . . . . .	60
<b>6</b>	<b>Simulation and Supernova Triggering</b>	<b>61</b>
6.1	Supernova neutrino interactions in NO $\nu$ A . . . . .	61

6.2	Flux models . . . . .	64
6.3	Simulating a supernova in the NOvA detectors . . . . .	64
6.4	Data-driven supernova trigger . . . . .	66
6.5	Long readout monitoring and data quality . . . . .	70
<b>7</b>	<b>Background Rejection and Hit Clustering</b>	<b>79</b>
7.1	Calibration and initial reconstruction . . . . .	79
7.2	Background rejection . . . . .	82
7.3	Clustering supernova neutrino interactions . . . . .	94
<b>8</b>	<b>Cluster Selection and Reweighting</b>	<b>101</b>
8.1	Cluster reweighting . . . . .	101
8.2	Cluster selection . . . . .	102
<b>9</b>	<b>Analyses and Results</b>	<b>121</b>
9.1	Modeling and fitting the time profile . . . . .	122
9.2	Sensitivity to supernovae . . . . .	125
9.3	Mass ordering . . . . .	128
9.4	Discussion and summary . . . . .	130
<b>10</b>	<b>Summary</b>	<b>133</b>

## CHAPTER 1

# Introduction

It happened in the year 185, and again in 393. Then it happened in 1006 and 1054 and 1604 and many times in between. In each case, a bright object appeared in the sky where none had existed before, confounding all who observed it. These “guest stars,” as some called them, did not move across the night sky as comets were known to do, and they twinkled like stars. Over a period of weeks or months, they would fade and eventually disappear, as if they had never been there at all.

We now know these as supernovae, a term coined in 1934 to describe the dramatic death of large stars. The cores of massive stars collapse under their own weight once they run out of nuclear fuel to fuse, and when they cannot collapse any further, the in-falling stellar material rebounds and blows the star apart, leaving behind either a neutron star or black hole and an expanding nascent nebula. Supernovae are extremely luminous events, often outshining their home galaxies. But despite their prolific photon emission, core-collapse supernova are primarily neutrino emitters.

The neutrino is a weakly-interacting chargeless lepton that exists in three flavors like its charged counterparts, the electron, muon, and tau lepton. For decades, neutrinos were thought to be massless like photons, but the observation of oscillations—a process by which a neutrino produced as one flavor interacts as another flavor after having traveled some distance—in the late 1990s confirmed that neutrinos must have a non-zero albeit small mass [1,2].

The proposition of core-collapse supernovae as neutrino emitters was first made in the early 1940s [3], but it wasn’t until the mid 1960s that they were understood to be the primary mechanism by which energy from the core is transferred to the stellar mantle during the explosion [4]. This was the first indication that neutrinos are not simply a by-product of core collapse, but play a pivotal role in the dynamics and evolution of the explosion.

The exact nature of their role is not fully understood today, and this is largely due to a lack of nearby core-collapse events in recent history against which theoretical models can be compared and constrained. This type of supernova is believed to be rare on a human timescale in or near a galaxy like ours, occurring only 2 or 3 times per century on average. The last such event occurred in February of 1987.

Supernova 1987A was humanity’s first opportunity to detect neutrinos from a nearby core collapse, and was therefore the first multi-messenger transient astronomical event. The supernova occurred outside of the Milky Way in the nearby Large Magellanic Cloud, but it was close enough to be detectable by three detectors dedicated to the search for proton decay. Collectively these detectors registered 25

neutrinos [5–7].

Since then, neutrino detectors have proliferated around the world. From the subterranean caverns under the Apennine mountains in Italy to crystal-clear ice below snowdrift dunes at the South Pole to the North Woods of Minnesota, a dozen neutrino observatories stand ready for the next galactic core-collapse supernova, and that number will grow in the coming decades. These detectors are connected to one another through the Supernova Early Warning System (SNEWS), a global coincidence network that allows individual detectors to crowd-source triggering decisions and to disseminate an alert when the supernova signal is finally detected. SNEWS will play a critical role in coordinating the global response in the moments after the supernova neutrino burst arrives at Earth, and every capable detector must be prepared to make their contribution [8].

The NOvA experiment has two such capable neutrino detectors. They both use liquid scintillator as an interaction medium and together the detectors are 14.3 kt in mass. The two-detector design is ideal for studying the phenomenon of neutrino oscillations by sampling Fermilab’s NuMI beam both near and far away from the source, but supernova neutrinos pose some challenges. For example, the accelerator neutrinos they were designed for are roughly one-hundred times more energetic than those produced in a supernova, which causes several problems. For one, the spatial resolution of the detector is too coarse for detailed reconstruction of these low-energy events. Additionally, some supernova neutrino energies fall below detection threshold, truncating the lower end of the energy spectrum. Another challenge is related to the placement of the detectors. The far detector, which contains the majority of NOvA’s detector mass, sits on the Earth’s surface with a modest overburden and is exposed to a high rate of cosmogenic background activity.

Repurposing NOvA as a supernova observatory is not trivial. It is truly a needle-and-haystack type of task: a smattering of low-energy depositions with little-to-no distinguishing topological characteristics awash in a sea of cosmic backgrounds, detector effects, and noise. But there is hope. What the NOvA detectors lack in low-energy spatial resolution they make up in fine timing; individual hits can be timestamped with nanosecond precision. Additionally, while most of NOvA’s mass is on the surface, not all of it is. The 300 t near detector is 100 m underground, providing a quieter environment that is largely free from the cosmogenic bombardment present at the far detector.

Despite the difficulties, the NOvA detectors are capable supernova observatories. In this thesis, I will demonstrate that non-supernova activity at the far detector can be reduced to an amount that makes identifying supernova neutrino events possible, and that by combining data from both detectors, NOvA will be able to make meaningful contributions to physics if the experiment is still running when the next galactic core-collapse supernova occurs.

This thesis is organized into two parts. The first part (chapters 2–4) centers on the physics and history of neutrinos and core-collapse supernovae. The second part (chapters 5–9) focuses in on the NOvA detectors. Chapter 5 describes the design of the detectors and the DAQ, timing, and triggering systems that support its operation. Chapter 6 delves into modeling supernova neutrino interactions in the NOvA detectors and how we used simulations to develop a data-driven supernova trigger. In that chapter, I also describe a system that I developed to continuously monitor the stability and integrity of data-driven and event-driven supernova triggers. Chapters 7–9 contain the bulk of original work that I did for this thesis, including the rejection of background activity at both detectors, clustering

supernova-like hits together, selecting clusters as supernova neutrino candidates using machine learning techniques, and demonstrating NOvA's ability to constrain progenitor distance and signal arrival time for several supernova model scenarios. Also in chapter 9, I demonstrate using two methods that NOvA is sensitive to the neutrino mass ordering under certain assumptions.





PART I

NEUTRINOS & SUPERNOVAE



## CHAPTER 2

# The Neutrino

Neutrino history began at the moment of the Big Bang, when they were produced in copious amounts within the maelstrom of a hot and dense universe. Approximately one second later, the universe had expanded and cooled enough to become transparent to neutrinos, allowing them to decouple from the surrounding medium and travel freely through space. The photons—which today constitute the cosmic microwave background (CMB)—would not have their freedom for another 380,000 years. Humanity’s awareness and understanding of the neutrino began roughly 13.8 billion years later, beginning in the early twentieth century and continues to the present day.

Neutrinos are an abundant particle in the universe, second only to photons, but their small masses, lack of electrical charge, and low interaction probabilities make them notoriously difficult to detect. With typical cross sections on the order of  $10^{-44} \text{ cm}^2 (E_\nu/1 \text{ MeV})$ , a neutrino can pass through a light-year of lead and still have a 50% chance of emerging from the other side. It’s no wonder these ghostly particles have evaded the searchlight of scientific discovery for so long.

Twenty-six years elapsed between the first postulation of the existence of the neutrino and the first detection. Another sixty-four have passed since then, and our efforts to understand the neutrino are still underway. Many questions remain unanswered about their nature and their relevance to other unsolved mysteries in physics and cosmology. In this chapter, I will lay out the theoretical foundations of neutrino physics as we know them today and explain some of the more impactful experimental developments that have lead us to this understanding.

### 2.1 A brief history

If humanity’s experience with the neutrino was written as a novel, the prologue would be set in Paris in 1896. Wilhelm Röntgen had just discovered x-rays the year before, and the French physicist Henri Becquerel was investigating emissions from phosphorescent uranium salts. These materials glowed after being exposed to sunlight, and Becquerel wanted to know if they produced x-rays. His experiment involved exposing these uranium crystals to the sun, wrapping them in black paper, and placing them next to a photographic plate, but one cloudy week in Paris thwarted his measurements. He placed the salts wrapped in their paper on the photographic plates in a desk drawer. Serendipitously, he decided to develop the photographic plates anyway, and they revealed that the uranium compound had continued to expose the plates; the uranium was producing radiation by itself.

Becquerel published his findings [9], and they were soon followed up by additional studies by Marie

and Pierre Curie, among others. Marie Curie was the first to coin the term “radioactivity” to describe the phenomenon, and she soon discovered that uranium was not the only substance to exhibit it; thorium was also radioactive. She and her husband began to systematically study uranium ores, which were curiously more radioactive than pure uranium alone, a sign that there were other radioactive elements present in the ores. They discovered two new elements, one which Curie named after her native Poland—polonium—and another which was much more radioactive than any substance known at that point—radium. Radium was difficult to produce, but it proved to be much more convenient for research than uranium and ultimately had many applications beyond the lab. Radium would also go on to have a disturbing legacy as humanity’s first experience with the devastating health effects of radiation on the human body.

Soon, the study of radioactivity was a busy area of research internationally. Three distinct types of radiation were discovered, characterized by their penetrating power. Alpha radiation was known to be a form of neutral radiation with the lowest ability to penetrate materials, which would turn out to be align helium isotope emitted from a nucleus. The second form, beta radiation, was a lightweight charged particle of some sort. It would bend in a magnetic field, and was eventually discovered to be an electron (or its antiparticle, the positron). The third and most penetrating form is that of gamma radiation, high-energy photons.

### **The beta decay problem**

It was an exciting time for physics research, but a conundrum soon emerged. Chadwick showed in 1914 using a magnetic spectrometer and an electron counter that in beta decay processes, the energy spectrum of emitted beta particles was continuous [10]. The decay was known to produce a proton and an electron, and such a two-body decay should have produced mono-energetic electrons. The continuous nature of the distribution suggested that some energy was missing. But where else could it have gone? One possible explanation was that one of the most bedrock pillars of physics, the conservation of energy, did not apply to this phenomenon. This result was concerning. In a letter to Ernest Rutherford in June 1914, Chadwick admitted that “there is probably a silly mistake somewhere.” [11]. He never managed to find one, and eventually Ellis and Wooster definitively confirmed the result (Fig. 2.1) in 1927 [12]. In 1930, Lise Meitner wrote to Ellis to say that she had also confirmed his result, “but I do not understand this result at all.”

For the first time since it had been established, the law of energy conservation was up for debate. Neils Bohr emerged as a prominent skeptic of energy conservation, while Wolfgang Pauli considered abandoning the “energy law” to be sacrilege. Pauli teased Bohr about this in their personal correspondence: “What is the current status of your new ideas?” Pauli asked in a postcard to Bohr in March 1929 [13]. “Do you intend to continue with the maltreatment of the poor energy law?” Later that summer, Pauli implored him to refrain from publishing a paper on the topic, advising him to “...let the note rest for a long time. And let the stars shine in peace!” [14].

### **Dear radioactive ladies and gentlemen**

Pauli soon came upon an alternative explanation, but it was a long shot and he knew it. He postulated the presence of a hitherto undiscovered particle, which, if emitted along with the beta particle, would carry some of the energy away with it and produce a continuous beta spectrum as seen by the experi-




Figure 2.1: The continuous spectrum of beta particles from the decay of  $^{210}\text{Bi}$  as measured by Ellis and Wooster in 1927 [12].

ments. This idea was perhaps even more absurd than abandoning the principle of energy conservation. At the time, the only known elementary particles were the electron and the proton<sup>1</sup>. Inventing a new particle to explain an experimental result was not commonplace like it is today. Nevertheless, the situation was dire and called for out-of-the-box thinking. Pauli introduced his idea in a now-infamous letter<sup>2</sup> [15] to Lise Meitner which was read aloud to colleagues at a workshop in Tübingen, Germany in December of 1930:

*Dear radioactive ladies and gentlemen,*

*I have hit upon a desperate remedy to save the 'exchange theorem' of statistics and the law of conservation of energy. Namely, the possibility that in the nuclei there could exist electrically neutral particles, which I will call neutrons, that have spin 1/2 and obey the exclusion principle and that further differ from light quanta in that they do not travel with the velocity of light. The mass of the neutrons should be of the same order of magnitude as the electron mass and in any event not larger than 0.01 proton mass. — The continuous beta spectrum would then make sense with the assumption that in beta decay, in addition to the electron, a neutron is emitted such that the sum of the energies of neutron and electron is constant.*

The “neutron” that Pauli refers to here is what would soon come to be known as the neutrino, after Chadwick discovered the neutron in 1932. There was some confusion as to whether Chadwick had de-

<sup>1</sup>The photon was known at the time as the “light quanta”, but it was not considered a particle. The proton would eventually be understood to not be elementary, but rather a collection of quarks and gluons.

<sup>2</sup>An English translation of the letter is available at [https://www.symmetrymagazine.org/sites/default/files/legacy/pdfs/200703/logbook\\_letter\\_translation.pdf](https://www.symmetrymagazine.org/sites/default/files/legacy/pdfs/200703/logbook_letter_translation.pdf)

tected Pauli's "desperate remedy", but it would turn out that Chadwick's particle was much too massive.

Pauli goes on in his letter to describe some additional predictions about his new particle before addressing the elephant in the room:

*But so far I do not dare to publish anything about this idea, and trustfully turn first to you, dear radioactive people, with the question of how likely it is to find experimental evidence for such a neutron if it would have the same or perhaps a 10 times larger ability to get through [material] than a gamma-ray. I admit that my remedy may seem almost improbable because one probably would have seen those neutrons, if they exist, for a long time. But nothing ventured, nothing gained...*

He concludes by expressing his regret for his absence at the workshop in person; he had to attend a ball in Zürich on the same night. Pauli's reluctance to publish his idea demonstrated his hesitance to make such a proposal with no direct experimental evidence to support it, but his idea had now made it out of his mind and out into the world, and it attracted interest. Efforts were soon underway to understand and detect this elusive particle.

### From theory to discovery

The Italian physicist Enrico Fermi took Pauli's idea and ran with it. Incorporating Pauli's particle and renaming it the "neutrino", meaning "little neutral one", to distinguish it from Chadwick's neutron, Fermi developed and published a theory of beta decay in 1934 [16]. This was a monumental work, and would later be generalized to describe other weak interaction processes. His theory described a neutron converting into a proton, and in the process emitting an electron and a neutrino<sup>3</sup>:

$$n \rightarrow p + \nu + e^{-}. \quad (2.1)$$

Pauli's invention of the neutrino as a product of beta decay had begged the question: where in the nucleus does the neutrino hide? The assumption behind this question is that the neutrino exists in the nucleus and is simply ejected during the decay, but Fermi's theory offered the alternative that the neutrino may be created at the time of emission, much like a photon. The experimental data favored Fermi's theory, and in that same year, Hans Bethe and Rudolf Peierls expanded upon this idea. If a neutrino can be created, there must be a corresponding annihilation process. They imagined running the beta decay process in reverse and calculated the neutrino cross section to be  $\sigma < 10^{-44} \text{ cm}^2$  [17]. While they noted that this cross section would grow with increasing neutrino energy, the interaction probabilities would still be too low to ever produce any appreciable signal. "...there is no practically possible way," they concluded, "of observing the neutrino."

But that did not stop people from trying. M. E. Nahmias attempted to detect neutrinos in December of 1934. This first underground neutrino experiment took place in the Holborn station of the London Underground. Nahmias placed a radioactive source near two Geiger-Müller counters shielded on all sides by lead. The hope was that, despite neutrinos having no electrical charge, if they had a sufficiently large magnetic moment, they could still produce ionization as they passed through the detectors. The result was null, and the neutrino magnetic moment was estimated to be smaller than 1/7000 Bohr magnetons [18]. More recently, a reporter asked Underground commuters for their thoughts on the

---

<sup>3</sup>This process specifically emits an electron antineutrino  $\bar{\nu}_e$ , but it was not yet known at the time that neutrinos have flavors or antiparticles

Holborn station having been used for a famous physics experiment, to which one woman cheekily replied “I loathe the Holborn station... If I was a neutrino, I wouldn’t stop there either.” [19]

The following decade saw a precipitous drop in theoretical and experimental developments in all areas of modern physics, including neutrino research, due to the deteriorating political situation in western Europe. The time between world wars saw a rapid modernization and institutionalization of physics, which until then had mostly been pursued by individuals or small groups in private labs. Physics research was increasingly seen as an important instrument of national power. Germany was arguably the center of gravity of that research in the early 1930s, and the Nazi’s rise to power in 1933—and subsequently World War II—would have a profound effect on scientific output for years to come. Neutrino physics was no exception; many individuals were forced to upend their work and their lives to escape the Nazi regime. Peierls fled in 1933 to the United Kingdom. Bethe left Germany for the United States in 1935. Fermi left Rome in 1938, accepted his Nobel Prize in Stockholm, and promptly fled to the United States. Also in 1938, Meitner made a daring and illegal escape from Berlin to Stockholm to save her own life. Pauli temporarily moved to the United States in 1940, fearing a German invasion of Switzerland. These are only a few examples; there are many others.

By the end of World War II, there was still no experimental evidence for the existence of the neutrino, or that obtaining experimental evidence was even plausible. But the Italian-Soviet nuclear physicist—and protégé of Enrico Fermi—Bruno Pontecorvo had an idea. In 1946 [20], he suggested that neutrinos might be detected through the inverse beta process

$$\nu + Z \rightarrow e^{-} + (Z + 1). \quad (2.2)$$

If  $^{37}\text{Cl}$  were used as the target,  $^{37}\text{Ar}$  would be produced by the process and one could look for its decay products. Pontecorvo also mentioned the need for a powerful neutrino source and recommended the sun or a nuclear reactor. This idea was picked up by Frederick Reines and Clyde Cowan at the Los Alamos Scientific Laboratory in the United States. Their idea was an experiment consisting of a target of roughly 400 liters of water and calcium chloride. The interaction of an antineutrino with the target would produce a positron and a fast neutron. The positron would quickly annihilate with an electron creating a prompt flash of light, while the neutron would take some time to thermalize before being captured by a nucleus, which would then de-excite and produce a secondary flash of light. These flashes of light would be detected by an array of photomultiplier tubes. The experiment would be placed next to a nuclear reactor at the Hanford Site in Washington State.

Reines and Cowan proposed their experiment—called “Poltergeist”—in the winter of 1953. By spring, it was taking data, and in the summer, they published their results. They may have detected the neutrinos, but the backgrounds were too large to be conclusive [21]. Undeterred, they moved their experiment a couple years later to the Savannah River nuclear plant in South Carolina, which offered better shielding. In 1956, they did what Bethe and Peierls had once thought would be impossible: they had detected neutrinos [22]. While this result, twenty-six years in the making, was celebrated and would eventually win Reines a Nobel Prize in 1995<sup>4</sup>, it did not come as much of a surprise. At this point, the existence of the neutrino had been broadly accepted. Upon hearing the news, Pauli casually replied to Reines and Cowan by telegram: “Thanks for the message. Everything comes to him who

---

<sup>4</sup>Clyde Cowan died in 1974.

knows how to wait.”

The neutrino that Reines and Cowan observed in 1956 came from a nuclear reactor and were electron neutrinos. At the time, only one neutrino flavor was known, but there was the question about whether there were more types of neutrinos. Pion decay had been observed in the 1940s, which produces a neutrino and a muon. So are these two neutrinos the same? Melvin Schwartz, Leon Lederman, and Jack Steinberger set out to find out. They built an experiment at Brookhaven National Laboratory which used a 15 GeV beam of protons colliding with a Be target to produce pions. The detector was placed behind several layers of concrete, steel, and lead to absorb the muons. The detector was a spark chamber and what they found were muon tracks coming away from the neutrino interaction vertices. They thus concluded that the neutrinos they observed must be distinct from those observed by Reines and Cowan in 1956 [23]. Within six years (1962) of discovering the neutrino, a second flavor of neutrino had been discovered, for which they were awarded the 1998 Nobel Prize in Physics. This was also the first instance of studying an accelerator-based beam of neutrinos.

Two families of neutrinos and their associated charged leptons had been discovered and by the mid 1970s, a second family of quarks had been discovered with the observation of the  $J/\Psi$  particle at Brookhaven and SLAC independently. A picture was beginning to emerge of a connection between the quarks and the leptons.

In 1975, Martin Perl et. al. discovered the  $\tau$  lepton at SLAC and thus discovered a third family of charged leptons [24]. Only months later, the  $b$  quark was discovered, confirming the existence of a third quark family. This triggered the search for a third neutrino family: the  $\tau$  neutrino. Detecting  $\tau$  neutrinos can be tricky due to the large energies involved, so it took a while for the discovery to be made. But finally in 2001,  $\tau$  neutrinos were observed by the DONUT experiment at Fermilab [25].

In 1989, the LEP experiments at CERN used an electron-positron collider to produce and study  $Z$  bosons. The width of the  $Z$  mass is expected to increase with the number of lepton families, as it mediates weak interactions for all of them. LEP was thus able to deduce the total number of light active neutrino families  $N_\nu$  based on their measurement of the  $Z$  width. Combined, four experiments measured  $N_\nu = 3.10 \pm 0.10^5$  [26–28] Despite having discovered all flavors of neutrinos that were thought to exist, the mysteries surrounding this ubiquitous-yet-elusive particle only continued to grow.

## The solar neutrino problem

Detecting neutrinos was not the only difficult aspect to understanding their nature. In 1939, Hans Bethe and others had proposed that the sun produces its energy through nuclear fusion of primarily hydrogen into helium [29], for which he would win a Nobel Prize in 1967. This process would produce a copious amount of neutrinos. John Bahcall developed a solar model and predicted the expected flux of solar neutrinos at Earth in 1964 [30, 31].

Soon thereafter, Ray Davis set out how test Bahcall’s prediction by setting up a 390,000 liter tank of  $C_2Cl_4$  underground at the Homestake gold mine in South Dakota. The idea was simple: neutrinos from the sun interacting with the  $^{37}Cl$  to produce  $^{37}Ar$ . Measure the amount of  $^{37}Ar$  present in the tank at the end of the experiment and deduce from that the neutrino flux. According to Bahcall’s prediction, Davis expected to observe a few interactions each day [32]. He only saw about 0.2 each day [33], much less than predicted. This was still a major discovery, and it earned Ray Davis a Nobel Prize in 2002, but

---

<sup>5</sup>This value has been further constrained in the years since by additional LEP experiments, cosmology, and SN 1987A.



a new problem had been born: where were the missing neutrinos from the sun?

The eventual solution wouldn't come until the late 1990s, when the Super-K experiment measured neutrino oscillations for the first time [1]. The neutrinos, it turned out, could change their flavor after propagating for some distance. They may have all been born in the core of the sun as one flavor, but by the time they made it to Bahcall's detector in South Dakota, a significant fraction had transformed into another flavor. Final confirmation came from the SNO experiment near Sudbury, Canada, which had added  $^3\text{He}$  detectors to its water Cherenkov vessel [34] to gain sensitivity to neutral current interactions which are agnostic to neutrino flavor [2].

The mystery was solved: neutrinos could oscillate into new flavors. This has an interesting implication. It means that neutrinos must experience the passage of time, and therefore travel slower than the speed of light, and therefore must have a mass. For a long time, neutrinos had been thought to be massless, but this discovery disproved that outright. As is typical with neutrino physics, this led to more questions than answers. The Standard Model predicts that the neutrinos should be massless, so how do they get their mass? If their mass is not zero, what is it? We still do not have definitive answers to these questions.

## 2.2 Neutrinos Physics

The Standard Model of particle physics is one of the most successful theories in terms of its predictive power. It is the product of decades of theoretical development and experimental study about the most fundamental constituents of matter and the forces that govern their interactions.

The Standard Model of today describes two classes of elementary particle (Fig. 2.2): the bosons—most of which mediate three fundamental forces—and the fermions—spin-1/2 particles which are the building blocks of matter. Of the bosons, the photon  $\gamma$  is the particle of light and mediates the electromagnetic force, the gluon  $g$  mediates the strong nuclear force, and the  $W^\pm$  and  $Z$  bosons are responsible for weak force interactions such as nuclear decays. The Higgs boson  $H$  is not associated with a fundamental force, but is responsible for giving particles their masses. The fermions are further split into two groups: quarks and leptons. Quarks constitute the bulk of ordinary matter by forming protons and neutrons. The leptons are the electron  $e$ , muon  $\mu$ , tau  $\tau$ , and their chargeless counterparts, the electron neutrino  $\nu_e$ , muon neutrino  $\nu_\mu$ , and tau neutrino  $\nu_\tau$ . Each of these fermions has an associated antiparticle.

### Mixing, mass, and oscillations

Neutrinos can be described by two categories of quantum mechanical eigenstates: flavor states  $|\nu_f\rangle$  which govern neutrino interactions, and mass states  $|\nu_i\rangle$  which govern neutrino propagation. The flavor states can be expressed as a linear combination of the mass states; that is,

$$|\nu_f\rangle = \sum_i U_{fi}^* |\nu_i\rangle. \quad (2.3)$$

This is known as “neutrino mixing”. Since there are three known active neutrino flavors,  $U$  is a unitary<sup>6</sup>  $3 \times 3$  matrix. The time-evolution of the mass eigenstates  $|\nu_i(0)\rangle$  can be deduced by solving the

---

<sup>6</sup> $UU^\dagger = U^\dagger U = \mathbb{1}$

	Fermions			Bosons
Quarks	<b>u</b> up	<b>c</b> charm	<b>t</b> top	<b>g</b> gluon
	<b>d</b> down	<b>s</b> strange	<b>b</b> bottom	<b>Z</b> Z boson
Leptons	<b><math>\nu_e</math></b> electron neutrino	<b><math>\nu_\mu</math></b> muon neutrino	<b><math>\nu_\tau</math></b> tau neutrino	<b>W</b> W boson
	<b>e</b> electron	<b><math>\mu</math></b> muon	<b><math>\tau</math></b> tau	<b><math>\gamma</math></b> photon
				<b>H</b> Higgs Boson

Figure 2.2: Elementary particles predicted by the Standard Model of particle physics.

Schrodinger equation. The solution for a plane wave is

$$|\nu_i(t)\rangle = e^{-iE_i t} |\nu_i(0)\rangle \quad (2.4)$$

where  $E_i$  is the energy. This describes how a neutrino propagates as a mass eigenstate, but when it is produced or when it interacts, it does so as a flavor eigenstate. It is therefore useful to express the above in flavor space as

$$|\nu_f(t)\rangle = \sum_{f'} \left( \sum_i U_{fi}^* e^{-iE_i t} U_{f'i} \right) |\nu_{f'}(0)\rangle. \quad (2.5)$$

This form allows one to compute the probability that a neutrino born as flavor  $f$  will later interact as flavor  $f'$ .

$$P(\nu_f \rightarrow \nu_{f'}) = \left| \langle \nu_{f'}(0) | \nu_f(t) \rangle \right|^2 = \left| \sum_i U_{fi}^* U_{f'i} e^{-i\frac{m_i^2}{2E} t} \right|^2. \quad (2.6)$$

Neutrinos are nearly always ultra-relativistic, so the substitutions  $E \approx p$  and  $t \approx L$  (in natural units where  $c = 1$ ) can be made where  $L$  is baseline, the distance traveled by the neutrino between production and interaction:

$$P(\nu_f \rightarrow \nu_{f'}) = \left| \sum_i U_{fi}^* U_{f'i} e^{-i\frac{m_i^2}{2E} L} \right|^2. \quad (2.7)$$

This can be expanded into more convenient forms such as the survival probability

$$P(\nu_f \rightarrow \nu_f) = 1 - 4 \sum_{i>j} |U_{fi}|^2 |U_{fj}|^2 \sin^2 \left( \frac{\Delta m_{ij}^2}{4E} L \right) \quad (2.8)$$




Figure 2.3: Normal (left) and inverted (right) neutrino mass orderings. The vertical axis represents mass. This cartoon is not to scale, and is meant only to illustrate the differences between the mass ordering scenarios: that there are either two light neutrinos and one heavy neutrino, or there are two heavy neutrinos and one light neutrino.

and the appearance probability

$$P(\nu_f \rightarrow \nu_{f'}) = \delta_{ff'} - 4 \sum_{i>j} \text{Re}[U_{fi}^* U_{fj} U_{f'i} U_{f'j}^*] \sin^2 \left( \frac{\Delta m_{ij}^2}{4E} L \right) \quad (2.9)$$

$$+ 2 \sum_{i>j} \text{Im}[U_{fi}^* U_{fj} U_{f'i} U_{f'j}^*] \sin^2 \left( \frac{\Delta m_{ij}^2}{2E} L \right). \quad (2.10)$$

where the quantities  $\Delta m_{ij}^2$  represent the mass-squared differences for the three neutrinos:  $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$ . Oscillation experiments measure these probabilities directly, which gives them access to the mass-squared differences. Solar neutrino experiments have measured  $\Delta m_{21}^2$  and its sign [35], but vacuum oscillations are agnostic to the sign of  $\Delta m_{32}^2$  to leading order. This leaves in ambiguity in the mass ordering of the neutrino mass states that the experimental neutrino community has yet to resolve. Either  $\nu_1 < \nu_2 < \nu_3$  (normal ordering) or  $\nu_3 < \nu_1 < \nu_2$  (inverted ordering) (Fig. 2.3).

For three active flavors, the matrix  $U$  can be parameterized in terms of the various mixing angles ( $\theta_{13}, \theta_{23}, \theta_{12}$ ) and a CP-violating phase ( $\delta_{\text{CP}}$ ):

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \quad (2.11)$$

$$= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (2.12)$$

with the notation  $c_{ij} = \cos \theta_{ij}$  and  $s_{ij} = \sin \theta_{ij}$  and  $\delta$  is the charge parity (CP) violating phase  $\delta_{\text{CP}}$ .

### Neutrino propagation in matter

The oscillation framework described above is specific to propagation in a vacuum. When neutrinos travel through matter, the presence of atomic electrons introduces an effective potential term  $\pm \sqrt{2}G_F N_e$  to the Hamiltonian, where  $G_F$  is the Fermi constant,  $N_e$  is the electron density, and

the sign depends on whether the mass ordering is normal (+) or inverted (−). This is known as the Mikheyev–Smirnov–Wolfenstein (MSW) effect [36,37], and it preferentially increases or decreases the effective mass of neutrino mass states with an appreciable electron flavor component and thus changes the effective mass splittings. Neutrino oscillations in matter therefore differ from those in a vacuum in a way that depends on the mass ordering. In situations where the matter density changes as the neutrino propagates, resonant flavor transformation can occur at certain critical electron densities.

## CHAPTER 3

# Neutrinos from Core-Collapse Supernovae

A star begins its life when gravitationally-bound gas and dust reaches a central density large enough to ignite nuclear fusion of hydrogen. A star's mass depends on how much gas is available in the immediate area and can range from less than that of the Sun to hundreds of solar masses. The mass determines much about how the star will evolve during the rest of its life. A larger stellar mass means hotter temperatures in the core, faster burning of its nuclear fuel, and a shorter lifetime. About 99% of the stars in the Milky Way have a mass of  $2 M_{\odot}$  or less [38], and will live on a timescale of billions of years while those with a very large mass ( $M/M_{\odot} \gtrsim 20$ ) may only live for several million. In any case, a majority of a star's lifetime will be spent fusing hydrogen as its primary fuel source.

Once a star exhausts its usable fuel stores, the end is near. For a low-mass star like the sun ( $M/M_{\odot} \lesssim 8$ ), a period of helium burning produces a higher outward pressure which inflates the outer envelope of the star and it enters the red giant phase. Much of the outer layers of the star will be expelled during this process, leaving a planetary nebula in its wake. After helium burning ceases, the remnant is a carbon-oxygen core called a white dwarf which will slowly cool and contract over time. This will be the fate of our sun.

High-mass stars meet their demise in a more spectacular manner: a core-collapse supernova. In this scenario, a star with a mass  $M/M_{\odot} \gtrsim 8$ –10 will burn through successively heavier elements (Tab. 3.1) from helium up to iron when nuclear binding energy is no longer released by fusion. The formation of iron in its nuclear furnace is the death knell of a large star. A rapid gravitational collapse of the core leads to a catastrophic explosion of the outer layers of the star, leaving behind a cooling neutron star as a remnant or a black hole if the star was massive enough.

Walter Baade and Fritz Zwicky coined the term “super-nova” in 1934. In two back-to-back proceedings, they postulated the existence of neutron stars, that supernovae would produce them, and that supernovae may be the origin of energetic cosmic rays [39, 40]. Nuclear physics and energy generation in stars were still nascent fields of study, so it's no surprise that their initial calculations of the energy released by supernovae were not accurate, and the neutrino did not appear in these first works. Nevertheless, Baade and Zwicky's postulation was an impressive insight for the time, and it kicked off a decades-long campaign to observe and understand supernovae which continues to this day.

The connection between neutrinos and supernovae was first made in 1941 by Gamow and Schoenberg [3] who recognized that the stellar interior would be hot enough to emit neutrinos, but another 25 years passed before Stirling Colgate and Richard White proposed in 1966 that the primary mechanism

Table 3.1: Shell burning stages for a  $15 M_{\odot}$  star. Neutrino loss refers to the luminosity of neutrino emission in units of solar luminosity. Table adapted from Ref. [42].

Stage	Time Scale	Fuel	Main Products	Temp. ( $10^9$ K)	Density ( $\text{g cm}^{-3}$ )	Neutrino Loss ( $L_{\odot}$ )
hydrogen	11 My	H	He	0.035	5.8	1,800
helium	2 My	He	C, O	0.18	1,390	1,900
carbon	2,000 y	C	Ne, Mg	0.81	$2.8 \times 10^5$	$3.7 \times 10^5$
neon	0.7 y	Ne	O, Mg	1.6	$1.2 \times 10^6$	$1.4 \times 10^8$
oxygen	2.6 y	O, Mg	Si, S	1.9	$8.8 \times 10^6$	$9.1 \times 10^8$
silicon	18 d	Si, S	Fe, Ni	3.3	$4.8 \times 10^7$	$1.3 \times 10^{11}$
core collapse	$\sim 1$ s	Fe	Neutron Star	$> 7.1$	$> 7.3 \times 10^9$	$> 3.6 \times 10^{15}$

through which energy is transferred from the collapsing core to the stellar mantle is the emission and subsequent absorption of neutrinos [4].

We know now that core-collapse supernovae are luminous neutrino emitters. The gravitational binding energy released by core collapse is approximately

$$E_b \sim E_g \approx \frac{3}{5} \frac{GM_{\text{ns}}^2}{R_{\text{ns}}} \approx 3.6 \times 10^{53} \left( \frac{M_{\text{ns}}}{1.5 M_{\odot}} \right)^2 \left( \frac{R_{\text{ns}}}{10 \text{ km}} \right)^{-1} \text{ erg.} \quad (3.1)$$

where  $R_{\text{ns}}$  and  $M_{\text{ns}}$  are the radius and mass of the remnant neutron star [41]. Of the  $\sim 10^{53}$  erg liberated by the collapse, only 1% takes the form of light (which can briefly outshine an entire galaxy) and kinetic energy of ejecta. The remaining 99% of that energy is carried away by  $10^{58}$  neutrinos which have energies of several MeV to several tens-of-MeV. The production of these neutrinos is the subject of the rest of this chapter, and the following chapter will focus on the properties and detection of the supernova neutrino signal at Earth.

## Supernova taxonomy

Supernovae are broadly classified into two types, type I and type II. These two types are distinguished by their observed spectral properties: Type II supernovae present hydrogen lines in their emission spectra while type I supernovae do not. Each type is then further subdivided based again on spectral properties, or in some cases the time-evolution of their light curves.

Type Ia supernovae are thought to originate from white dwarfs that accrete matter onto their surfaces from a nearby companion star which drives their masses over the Chandrasekhar limit ( $\sim 1.44 M_{\odot}$ ) and raises the temperature in the carbon core enough to ignite a runaway nuclear fusion process which blows the white dwarf apart. While the exact mechanism that drives these explosions is still uncertain (see Ref. [43] for a review), the emission properties of these explosions are known to be very uniform and are often used as standard candles for measuring the distances to galaxies. While these supernovae do release neutrinos, the vast majority of the emitted energy is in the kinetic energy of ejecta.

Types Ib, Ic, and II are core-collapse supernovae produced by massive progenitors. Type II supernovae are further divided into types II-P and II-L depending on whether their luminosity profiles plateau or decrease linearly over a period of time following the explosion. The collapse of an iron core is an effective engine for converting gravitational binding energy into neutrinos, so in contrast to the

type Ia supernovae, these core-collapse varieties release the majority of their energy in the form of neutrinos and present the greatest opportunity for supernovae neutrino detection at Earth.

### Core-collapse supernovae in the Milky Way

Two questions come to mind when planning to observe the next galactic core-collapse supernova: (1) when will it happen?, and (2) how far away from Earth will it be? Unfortunately, there is no way to answer these directly. We can, however, construct a probabilistic model that is informed by observations both within and outside of our galaxy.

One way to compute the supernova probability as a function of distance from Earth is to use the volume distribution of neutron stars in the galaxy as a proxy for the distribution of future core-collapse events. Using this method, Ref. [44] computes an average supernova distance of  $\langle d_{SN} \rangle = 10.7 \pm 4.9$  kpc. Another method is to simply count the number of known red supergiants, which are in their final stage of stellar evolution. Figure 3.1 shows probable future supernova distances as well as a map of the most likely nearby supernova progenitors.

Determining the average rate of core-collapse supernovae in our galaxy is not trivial. The last such supernova that was unquestionably visible to the naked eye from Earth was observed by Johannes Kepler in 1604, several years before the invention of the telescope. Modern telescopes have observed remnants of likely supernova explosions within our galaxy that can be used to infer the rate, but determining the age of the remnants can be difficult. While the remnants of past galactic supernovae can be observed by their emissions across many wavelengths, spotting one during the explosion is difficult when there is a large amount of dust to obscure our view. The emergence of neutrino telescopes during the past several decades will prove to be an invaluable tool for detecting supernovae even when the optical signal is diminished by dust.

The most recent attempt to compute the galactic core-collapse supernova rate combines a variety of independent calculations from recent literature along with a census of core-collapse supernovae in the Local Group. This approach yields a predicted rate of  $1.63 \pm 0.46$  century<sup>-1</sup> [46].

## 3.1 Physics of a collapsing core

The collapse of stellar core is a complex and dynamical interplay between all four fundamental forces. In this environment, general relativity, neutrino interactions, magnetohydrodynamics, and nuclear physics are significant contributors to the dynamics of the collapse. This section will explain our current understanding of how a stellar core collapses.

### Pre-supernova

Neutrinos are a byproduct of nuclear fusion and are therefore produced by all stars for the duration of their lives. But neutrinos begin to play a more prominent role as massive stars exhaust their nuclear fuel stores. By the time a star begins burning silicon, it is only days away from core collapse. During this phase, the temperature in the core continuously increases as does the production of thermal neutrinos. Neutrino emission during this phase is largely dominated by  $\nu_e$  and  $\bar{\nu}_e$  flavors. The average energy of emitted neutrinos also grows over time, beginning at the sub-MeV level before eventually peaking at 1 MeV–3 MeV [47].

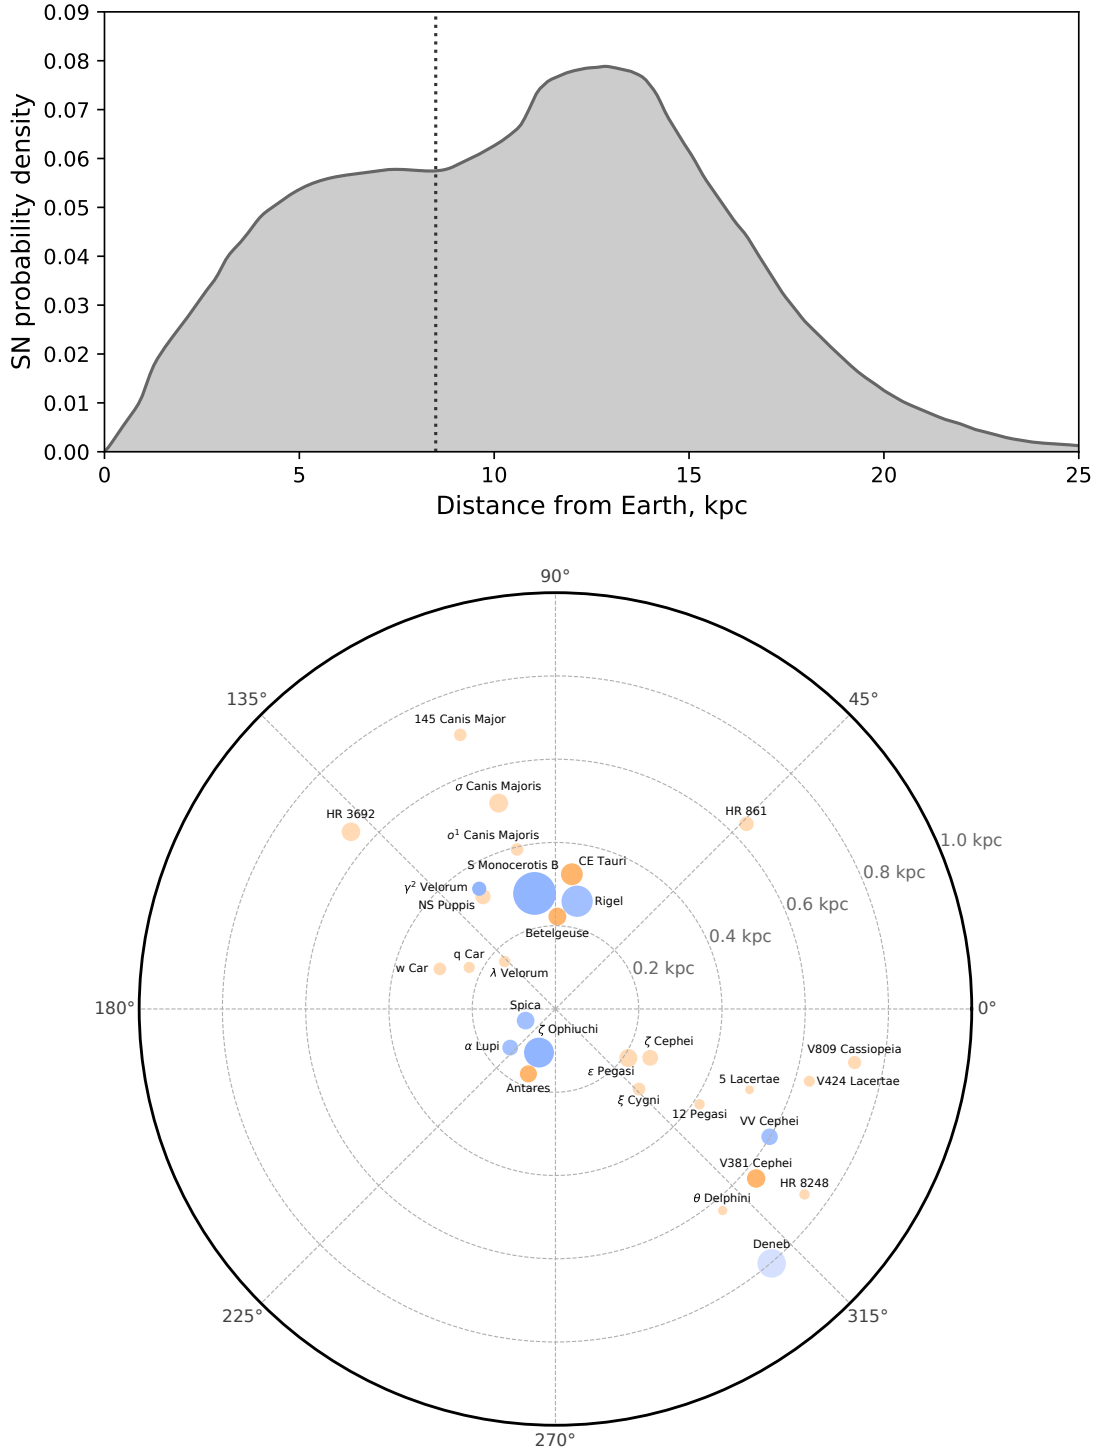


Figure 3.1: *Top*: Core-collapse probability distribution versus distance from Earth as calculated in Ref. [44]. The dotted line indicates the distance to galactic center (8.5 kpc). *Bottom*: Candidate red and blue core-collapse progenitors within 1 kpc of Earth (as reported in Ref. [45]). Area is proportional to the star's mass and the radial axis represents distance from Earth and angular axis is right ascension.



## Collapse

The silicon in the center of the core is replaced by its byproduct, iron, which is inert in this situation because nuclear binding energy can not be released by fusion. The internal source of outward pressure which has kept the star in a state of hydrodynamic equilibrium for millions of years disappears, and the star begins to collapse under its own enormous weight.

As the collapse proceeds, the mass of the iron core grows due to the presence of in-falling material and also because a thin shell of silicon surrounding the core continues to produce iron nuclei. Despite the rapidly dissipating outward pressure from nuclear fusion, the core is initially supported by electron degeneracy pressure and briefly resembles a white dwarf. However, the core soon reaches its effective Chandrasekhar mass as in-falling matter continues to accrete.

The rate of collapse is largely governed by the evolution of entropy inside the core and the electron fraction  $Y_e$  [48]. As the temperature rises, electrons begin to capture on heavy nuclei and free protons:

$$e^- + p \longrightarrow \nu_e + n, \quad (3.2)$$

$$e^- + (A, Z) \longrightarrow \nu_e + (A, Z - 1). \quad (3.3)$$

This “neutronization”, produces a large flux of  $\nu_e$  which initially stream out of the core unimpeded. The lowering of  $Y_e$  consequently reduces the electron degeneracy pressure and accelerates the collapse. As the collapse continues, temperatures rise to the point that photons are energetic enough to dissociate iron nuclei into alpha particles and neutrons, effectively undoing what the star took millions of years to achieve. This photo-dissociation process



is endothermic, and robs the core of energy that would otherwise be used to slow the collapse.

Roughly 100 ms after the start of collapse, densities in the core reach  $10^{12} \text{ g cm}^{-3}$  and neutrinos begin to diffuse rather than stream freely out of the core. The diffusion time exceeds the collapse time, so that the neutrinos are effectively trapped. The boundary between where neutrinos are trapped and where they are free is known as the “neutrinosphere” and is in many ways analogous to the photosphere of a star. Within the neutrinosphere, neutrinos are no longer able to carry energy away from the core, and the collapse in this region proceeds adiabatically until the inner core density reaches that of nuclear matter [49].

## Bounce and shock propagation

110 ms after the start of the collapse, the inner core achieves nuclear densities of  $\rho = 2.7 \times 10^{14} \text{ g cm}^{-3}$  and the equation of state stiffens abruptly [48]. The inner core is now supported by the internal pressure of a nucleon gas and is stable against further collapse. The momentum of collapsing material momentarily compresses the inner core beyond this equilibrium causing it to bounce back, sending a pressure wave outward which quickly develops into a shock front.

The temperature of the shock front increases as it absorbs kinetic energy from in-falling material from the outer regions of the core, allowing the photo-dissociation process described by Eq. 3.4 to take place in addition to further dissociation of the helium nuclei; as the shock propagates through the core, it encounters iron nuclei and leaves behind free neutrons and protons in its wake.

The capture of electrons on free protons produce copious numbers of  $\nu_e$ . Once the shock has passed the neutrinosphere boundary and its density has decreased sufficiently, the once-trapped neutrinos are free to escape. This event happens about 120 ms after the collapse begins, and produces a sudden burst of electron neutrinos: the “shock breakout” [50]. This results in a sudden decrease in electron-lepton number and allows for pair-production processes, which had been previously suppressed, to take place, leading to the thermal production of neutrinos of all flavors through processes like pair annihilation ( $e^- + e^+ \rightarrow \nu + \bar{\nu}$ ) and the photo-neutrino process ( $\gamma + e^- \rightarrow e^- + \nu + \bar{\nu}$ ). The presence of positrons also enables the production of  $\bar{\nu}_e$  via  $e^+$  captures on neutrons [51].

The dense inner core forms the proto-neutron star (PNS), which will continue to grow as more post-shock material rains down onto it. Neutronization and thermal processes will continue to produce neutrinos of all flavors which carry energy and lepton number out of the core. The PNS will either become the remnant neutron star or a black hole for progenitor masses  $M/M_\odot \gtrsim 25$ .

### Accretion and explosion

The shock front is the leading edge of the supernova explosion, pushing through the iron core before eventually blowing away the outer envelope of the stellar atmosphere. But there is a problem with this picture. The shock front loses energy as it traverses the dense core, primarily due to the ram pressure of in-falling matter as well as the endothermic photo-dissociation of iron group nuclei [49]. The energy of the shock front is not sufficient to overcome these barriers, and roughly 200 ms after the collapse began, the shock will stall well inside of the iron core and become a standing accretion shock [49].

Without some way to revive the shock, it is destined to fall back onto PNS and the explosion will have failed. At present, the most widely-supported mechanism for shock revival relies on neutrino heating [49]. When electron neutrinos and antineutrinos radiating out from the PNS encounter the dense stalled shock, a fraction of them deposit their energy behind the shock through the following reactions:

$$\nu_e + n \longrightarrow p + e^-, \quad (3.5)$$

$$\bar{\nu}_e + p \longrightarrow n + e^+. \quad (3.6)$$

The region where these interactions occur is known as the “gain region”.

Turbulent motions in region of the stalled shock may also contribute to its revival. One example is the standing accretion shock instability (SASI) [52], a global instability of the shock itself due to an unstable coupling of acoustic waves in the post-shock region and entropy and vorticity perturbations near the shock. These turbulent effects lead to a large-scale sloshing of the material between the shock and the PNS. This can help expand the radius of the stalled shock and increase the size of the gain region, allowing the neutrinos to deposit even more energy behind the shock [49]. If the combined pressure of turbulence and neutrino heating exceeds the ram pressure of in-falling matter, the shock will be revived and resume its march outward.

### Cooling

After the shock has been revived and mass accretion is largely complete, the PNS is left to contract and undergoes Kelvin-Helmholtz cooling via neutrino emission [51]. Any remaining proton and electron pairs will be converted into neutrons and electron neutrinos and thermal production of neutrinos of all flavors will continue over the course of tens of seconds [51].

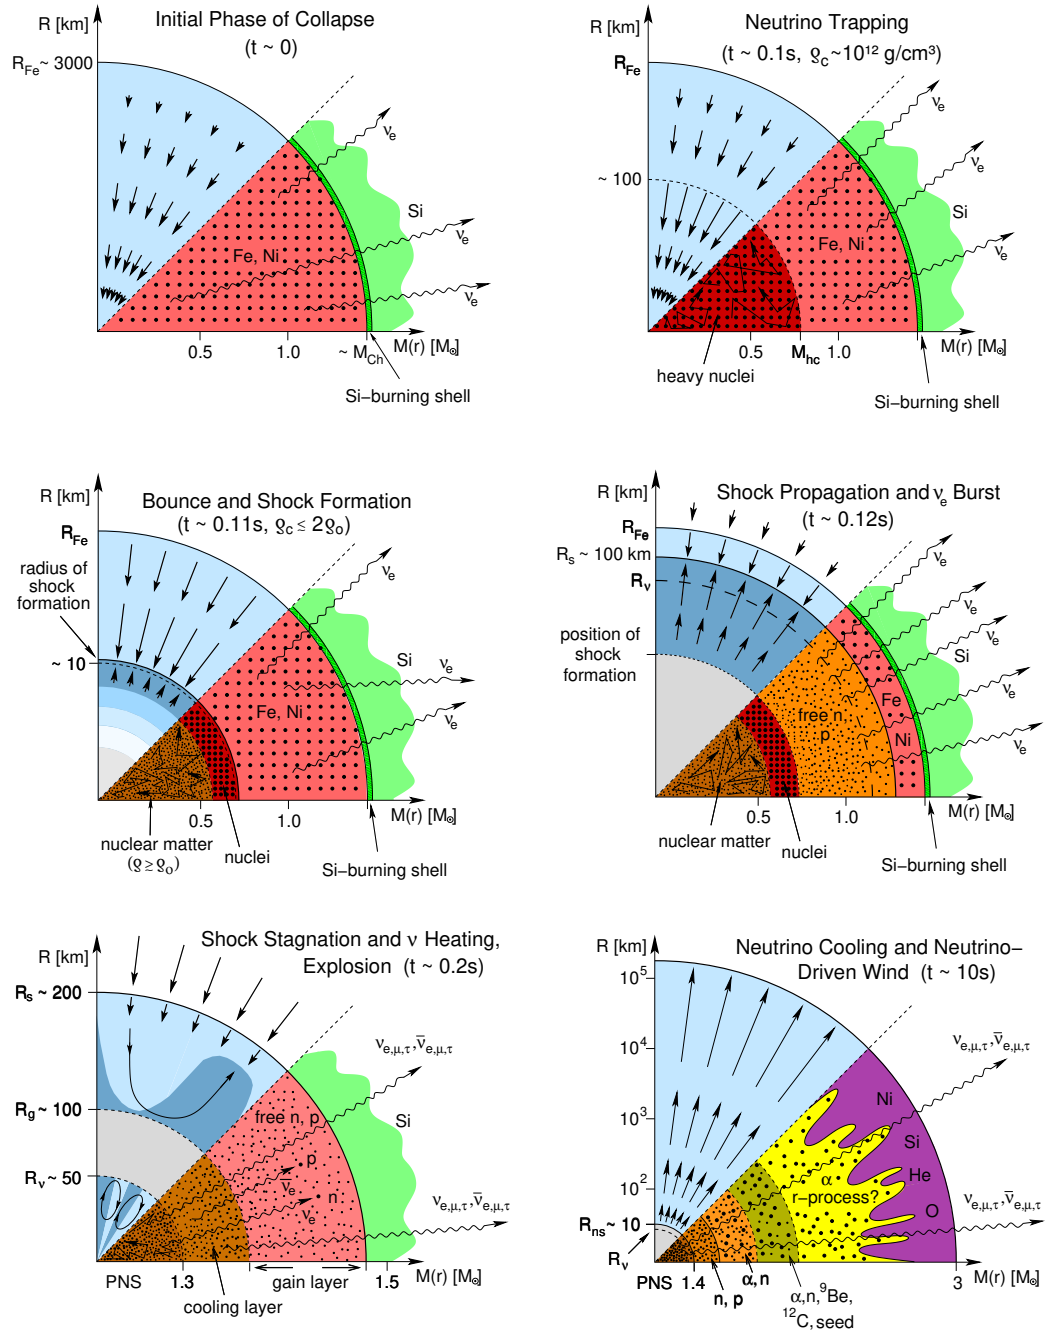


Figure 3.2: Stages of a core-collapse supernova. The center of the core is at the origin. Vertical axes represent the distance from the center of the core and horizontal axes represent the amount of mass enclosed by a sphere of that radius. The upper wedges depict the velocity of matter flows within the core, while the lower wedges show the composition.  $R_{\text{Fe}}$  is the radius of the iron core,  $R_\nu$  is the approximate radius of the neutrinosphere,  $R_s$  is the radial location of the shock,  $R_g$  is the gain radius, and  $R_{\text{ns}}$  is the radius of the proto-neutron star.  $M_{\text{Ch}}$  is the Chandrasekhar mass and  $M_{\text{hc}}$  is the mass enclosed by the homologously-collapsing core where  $v \propto r$ .  $\rho_0$  is the nuclear saturation density ( $2.7 \times 10^{14}$  g cm<sup>-3</sup>) and  $\rho_c$  is the density at the center of the core. Figure from Ref. [50].




Figure 3.3: Neutrino emission rates during each stage of a simulated  $27 M_{\odot}$  core-collapse supernova. The supernova model used to produce this plot is LS220-s27.0co from the Garching group [55].

### 3.2 Neutrino emission

Neutrino emissions from a core-collapse supernova occur in three main stages: burst, accretion, and cooling. Each successive stage spans a longer timescale than the one before and present distinctive signatures of the underlying physics involved. Figure 3.3 shows the instantaneous neutrino emission rate by flavor during each stage for a simulated  $27 M_{\odot}$  supernova.

1. *Burst: 5 ms before bounce – 20 ms after bounce*

A prompt burst of  $\nu_e$  from electron capture (neutronization) lasts 20 ms. This is a distinct feature that is largely model-independent [53].

2. *Accretion: 20 ms – 500 ms after bounce*

Neutrino emission is driven by matter accreting onto the stalled shock front. Emission rates are dependent on the radius of the shock, which can oscillate due to SASI effects. SASI can cause the neutrino flux to modulate at a frequency of several hundred Hz [54].  $\nu_e$  and  $\bar{\nu}_e$  luminosities are roughly equal, but higher than that of muon and tau flavors.

3. *Cooling: 500 ms – 60 s after bounce*

The nascent neutron star cools via the emission of thermal neutrino-antineutrino pairs. This stage lasts for a relatively long period of time, tens of seconds.

If the progenitor’s mass is sufficiently large, the proto-neutron star may accrete enough mass so that the repulsion by the strong nuclear force is no longer strong enough to prevent a complete collapse. In this scenario, the PNS will collapse into a black hole, and the neutrino emission will be abruptly cut off. This provides a distinct signature of black hole formation, provided that the collapse does not happen too far into the cooling phase when the flux has already fallen below our ability to detect it on Earth.

### 3.3 Simulations

Simulating the dynamics of core-collapse supernova is a challenge. The extreme environment in which the explosion evolves is subject to all four fundamental forces and spans many sub-fields of modern physics: general relativity is needed to properly handle gravitational forces due to the dense proto-neutron star; the strong and weak nuclear forces play critical roles in the neutronization and shock revival phases; the nuclear equation of state in the core can have a large impact on the dynamics of the collapse and bounce; rotational magnetohydrodynamics impacts the motion of hot plasma in the outer core and mantle; neutrino transport dictates how energy leaves the core and how it is deposited behind the accretion shock; turbulence at the stalled shock front is vital to its revival. Any one of these effects can be challenging to simulate on their own, never mind simulating all of them and their interplay at once. Add to that the scale of the simulation: from dynamics that play out on time scales of milliseconds to the long-term cooling phase which can last for tens of seconds.

Current supernova models [56, 57] include treatment of some physics processes while excluding certain others and approximations are often employed to reduce computational complexity. But how do these choices affect the accuracy of these simulations? One way to answer that is to run many simulations under different scenarios and looking at which features are common between most permutations and which are sensitive to these differences. The presence of the  $\nu_e$  neutronization peak, for example, is present in the vast majority of model scenarios. Whether the supernova successfully explodes or not is a different story. The limiting factor here is how fast new simulations can be produced to compare with one another, and that depends in large part on the dimensionality of the simulation. Today, in one year the community can produce many 1D simulations, a couple dozen 2D simulations, and a few 3D simulations. I won't go into detail about existing simulation codes here, but see Ref. [55] for a review. In a later chapter, I will describe the specific models that I used for this thesis in greater detail.

The problem of non-exploding simulations has haunted the supernova simulation community for years. The revival of the stalled shock front appears to be quite sensitive to the physics, especially the nuclear equation of state, the neutrino luminosity, and turbulence that occurs at the shock front itself. Models that do not explode on their own are still useful, but they require an artificial “kick” to keep the accretion shock from falling back onto the proto-neutron star. Although most simulated explosions fail in 1D, 2D, and 3D (The first successful 3D explosion was reported as recently as 2015 [58]), dimensionality still plays an important role. Figure 3.4 shows two comparisons of explosion profiles. One is a comparison between a 1D and 2D simulation and the other is between a 2D and 3D simulation. Each comparison includes the same physics and neutrino luminosity. In both cases, the lower-dimension simulation fails to explode and the other succeeds. While simulating a 3D supernova is not alone sufficient to guarantee a successful explosion, dimensionality is clearly important. After all, real core-collapse supernovae happen in three dimensions.

Despite these challenges, computational infrastructure continues to improve, our theoretical treatment of core collapse matures, and supernova simulations grow faster and more sophisticated; it is an active area of research, and there is reason to be optimistic about its future.

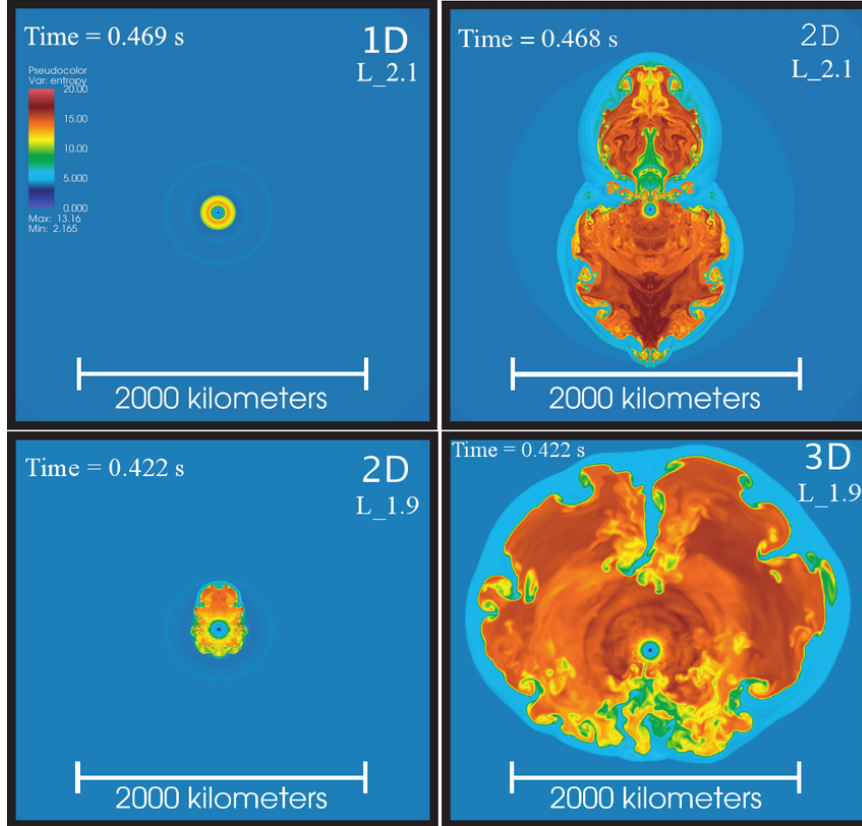


Figure 3.4: Comparisons of 1D to 2D and 2D to 3D supernova simulations under similar initial conditions. Color represents entropy. The top panels are driven by a  $\nu_e$  luminosity of  $2.1 \times 10^{52} \text{ erg s}^{-1}$ . The 1D simulation on the left has failed to explode while the 2D simulation has not. In the bottom panels, the  $\nu_e$  luminosity is reduced slightly to  $1.9 \times 10^{52} \text{ erg s}^{-1}$ . In this case, the 2D simulation on the left fails to explode while the 3D model on the right does not. Figure from Ref. [59].

### 3.4 Looking ahead

Our knowledge of supernovae and the pivotal role that neutrinos play in driving these explosions has come a long way since the 1930s and 1940s when the ideas were first introduced, but there is still much to learn. Nothing would more rapidly enhance our understanding of core collapse dynamics than a high-statistics multi-messenger observation of a galactic supernova.

A combined observation of electromagnetic and neutrino luminosity time profiles, neutrino energy spectra, and perhaps gravitational waves would constrain the list of feasible model scenarios and provide invaluable insight into which physical processes are most important in driving explosion dynamics, and which can be approximated.

We anxiously await the arrival of neutrinos from the next galactic core-collapse supernova here on Earth, but what can we do in the meantime? In the next chapter, I will describe what we know about how this signal will look and the global effort that is underway to ensure we don't miss it when the time finally comes.





## CHAPTER 4

# Supernova Neutrinos at Earth

We are fortunate enough to be located a safe distance from the majority of likely supernova progenitors. In the previous chapter, we explored what happens inside the core and surrounding region while it collapses, but when we observe the next core-collapse supernovae, we will do so from the vantage point of Earth, hundreds or thousands of light-years away. Humans have a long history of watching these cosmic cataclysms, but for the most part they had no idea what they were looking at. We presently stand on the precipice of being able to peer into the core of an exploding star and understand its inner workings. What will we see? How will the physical processes described in the previous chapter manifest themselves in the neutrinos we detect? What are we doing to prepare for such an historic observation? These questions are the focus of this chapter.

### Historical observations

The first known documented observation of a supernova occurred in the year 185 CE. Chinese astronomers identified a guest star<sup>1</sup> which stayed visible in the night sky for eight months. Its appearance intrigued these astronomers, who documented it for posterity in the Book of the Later Han<sup>2</sup>:

*“In the 2nd year of the epoch Zhongping, the 10th month, on the day Kwei Hae [December 7, Year 185], a ‘guest star’ appeared in the middle of Nan Mun [asterism containing Alpha Centauri], The size was half a bamboo mat. It displayed various colors, and gradually lessened. In the 6th month of Mod- the succeeding year it disappeared.”*

ern observations point to the remnant RCW 86 as the likely culprit. The supernova occurred 2.8 kpc away and was likely a type Ia [60].

SN 1006 is another notable supernova event which became visible on Earth in the year 1006 CE. This supernova was observed all around the world, including in Europe, Asia, Africa, and possibly North America. It is notable for possibly being the brightest astronomical event in recorded history. Ali ibn Ridwan, an Egyptian astrologer wrote that

*“[the] spectacle was a large circular body, 2 ½ to 3 times as large as Venus. The sky was shining because of its light. The intensity of its light was a little more than a quarter that of Moon light.”*

---

<sup>1</sup>a “guest star” is terminology used in ancient Chinese astronomical records to describe a bright star-like object that appeared in the sky where no star was previously present.

<sup>2</sup>The Book of the Later Han is a written history of the Han Dynasty covering the period of 6–189 CE.

The supernova may have left behind a signature on Earth. An analysis of Antarctic ice cores revealed a spike in the levels of nitrogen oxide in the atmosphere that year [61], which can be produced by a flux of gamma rays interacting with atoms in the atmosphere.

In October of 1604, Johannes Kepler recorded his observation of a supernova. SN 1604 was brighter than any star in the sky and remained visible during the day for three weeks. This is the most recent supernova in the Milky Way galaxy to be seen by the naked eye.

Since then there have been more notable supernova events. SN 2008D was the first supernova to be observed while it exploded [62], and SN 2015L is the most luminous supernova ever observed [63], but perhaps the most notable supernova of the modern age was SN 1987A, the first supernova with an observed neutrino counterpart.

#### 4.1 Supernova 1987A

On February 23, 1987 at 07:35:35 UTC, a tsunami of neutrinos washed over the the Kamiokande-II solar neutrino detector in Japan, the Baksan detector in Russia, and IMB in the United States. They went unnoticed in the moment, but those neutrinos would turn out to be a discovery of historic proportions, and they carried a message that more was to come.

Hours later, Ian Shelton at the Las Campanas observatory in Chile and Albert Jones of New Zealand independently discovered a bright object in the southern sky, the first supernova visible from Earth in 383 years. The International Astronomical Union Circular (IAUC) #4316, titled "Supernova 1987A in the Large Magellanic Cloud" was soon published announcing the observation<sup>3</sup>. In another first, the progenitor star was discovered to be Sanduleak -69 202, a blue supergiant located 50 kpc away with past measurements.

It was only well after news of a nearby supernova was disseminated that anyone thought to go back and look at data from the neutrino-sensitive detectors of the time; they found them. Of the  $\sim 10^{57}$  neutrinos produced in the explosion, only two dozen are known to have left any trace (Fig. 4.1).

#### 4.2 Supernova signal characteristics

Neutrinos are fermions, so one might expect their spectrum to be that of a thermal Fermi-Dirac distribution

$$f(\epsilon) \propto \frac{\epsilon^2}{1 + \exp(\epsilon/T - \eta)}. \quad (4.1)$$

where  $\epsilon$  is the neutrino energy,  $T$  is the average neutrino temperature, and  $\eta$  is a degeneracy parameter. However, neutrino interactions inside the collapsing core are strongly energy dependent, causing them to decouple from the surrounding medium at different radii depending on the neutrino energy. A better formulation [64] of the energy spectrum for supernova neutrinos is

$$f_\alpha(\epsilon) = \frac{(\alpha + 1)^{(\alpha+1)}}{\langle \epsilon \rangle \Gamma(\alpha + 1)} \left( \frac{\epsilon}{\langle \epsilon \rangle} \right)^\alpha e^{-(\alpha+1) \epsilon / \langle \epsilon \rangle} \quad (4.2)$$

where  $\langle \epsilon \rangle$  is defined as

$$\langle \epsilon \rangle \equiv \frac{\int_0^\infty d\epsilon \epsilon f_\alpha(\epsilon)}{\int_0^\infty d\epsilon f_\alpha(\epsilon)} \quad (4.3)$$

---

<sup>3</sup><http://www.cbat.eps.harvard.edu/iauc/04300/04316.html>

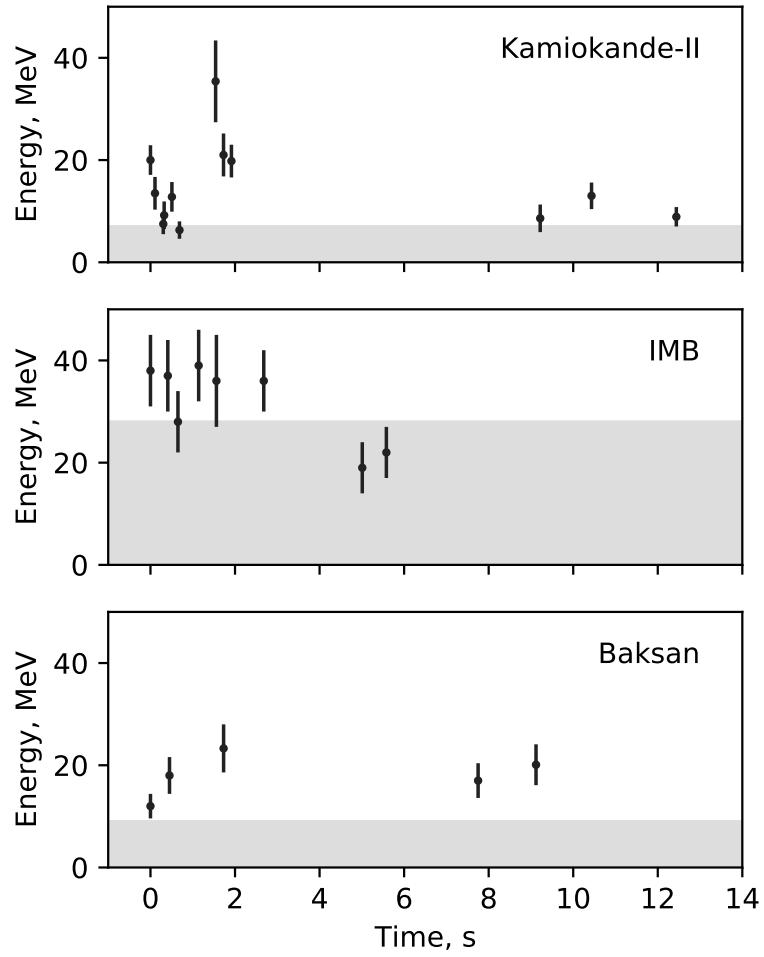


Figure 4.1: Neutrinos from Supernova 1987A recorded by the Kamiokande-II [5], IMB [6], and Baksan [7] detectors. The gray regions represent where the trigger efficiency is less than 30%. The time is relative to each detector, with the first event in each case occurring at  $t = 0$ .




Figure 4.2: Left: the triple-ring system of SN 1987A. Right: the inner ring. Credits: ALMA, Hubble Space Telescope, Chandra X-Ray Observatory




Figure 4.3: Neutrino luminosity (top), average neutrino energy (middle), and average square neutrino energy (bottom) versus time for the LS220-s27.0co (left) and SFHo-z9.6co (right) supernova neutrino flux models [55]. The values of these quantities at any one time parameterize the neutrino energy spectrum and flux.

and

$$\frac{\langle \epsilon^2 \rangle}{\langle \epsilon \rangle^2} = \frac{2 + \alpha}{1 + \alpha}. \quad (4.4)$$

The shape of Eq. 4.2 is similar to that of a standard Fermi-Dirac, but with some amount of pinching (narrowing) or anti-pinching (widening), depending on the value of  $\alpha$ . One difference between these is that we express Eq. 4.2 in terms of the average neutrino energy  $\langle \epsilon \rangle$  instead of the effective temperature  $T_\nu$ , although these quantities are related by  $T_\nu = 3 \langle \epsilon \rangle$ . It is important to note that—while not always explicitly stated—the quantities  $f_\alpha(\epsilon)$ ,  $\alpha$ , and  $\langle \epsilon \rangle$  are time-varying (see Figs. 4.3 and 4.4 as examples).

### 4.3 Flavor transitions

A neutrino produced in the collapsing core of a dying star propagates outward through a dense and turbulent environment before eventually freely streaming out of the stellar mantle, traveling through interstellar space for thousands of years, and passing through Earth’s atmosphere and possibly the interior of the planet itself before finally interacting in a supernova neutrino detector. During each of these




Figure 4.4: The supernova neutrino flux (LS220-s27.0co flux model [55]) at the point of production for various times based on the spectral shape in Eq. 4.2.

stages, neutrino oscillation phenomena imprint themselves on the detectable signal. Here, I will briefly touch on each major source of oscillation that may affect supernova neutrinos during their lifetime assuming a three-neutrino oscillation framework.

### Self-induced transitions

The neutrino density in the core is so high that the neutrino-neutrino interactions are no longer negligible. Supernovae are perhaps the only places where these interactions are possible, making them a unique laboratory for better understanding this phenomenon. The theoretical study of neutrino-neutrino interactions is still nascent, and the predictions about how they will imprint themselves on the detectable neutrino signal is not robust.

One possibly observable signature comes from neutrino pair conversions like  $\nu_e \bar{\nu}_e \rightarrow \nu_x \bar{\nu}_x$  [65]. These “collective effects” depend on the angles of the incoming neutrinos, and in cases where neutrino emission is anisotropic, the flux could be significantly affected [66]. A “spectral split” can occur when one flavor transforms into another above a certain energy threshold or a full spectral swap where one flavor entirely transforms into another. These would manifest themselves as discontinuities on the spectral shape [67].

### MSW in supernovae

The dense stellar environment allows for the MSW effect introduced in Sec. 2.2 to play a role in shaping the flavor distribution of neutrinos emitted from a supernova. In the region behind the shock front, the density is smoothly varying and neutrinos propagate adiabatically. At specific electron densities, neutrinos can experience resonant flavor transitions. There are two such resonances, the  $L$  resonance and the  $H$  resonance. The  $H$  resonance is the dominant of the two, and affects neutrinos for the normal mass ordering and antineutrinos for the inverted mass ordering. The less dominant  $L$  resonance affects neutrinos in both scenarios. Figure 4.5 depicts the neutrino energy eigenstates for both mass ordering scenarios as a function of electron density. The locations where two eigenstates approach each other corresponds to the resonances.




Figure 4.5: Neutrino level crossing diagrams in a supernova as a function of electron density  $n_e$  for the normal (left) and inverted (right) mass orderings. The solid vertical line labeled “vacuum” represents  $n_e = 0$ . Neutrinos propagate along the solid lines from regions of high density to zero density (right to left for neutrinos and left to right for antineutrinos). The dashed lines represent flavor eigenstate energies. From [54].

The neutrino fluxes transform in the following way after undergoing adiabatic MSW oscillations:

$$\text{Normal Mass Ordering} \begin{cases} F_{\nu_e} &= F_{\nu_x}^0 \\ F_{\bar{\nu}_e} &= \cos^2 \theta_{12} F_{\bar{\nu}_e}^0 + \sin^2 \theta_{12} F_{\bar{\nu}_x}^0 \end{cases} \quad (4.5)$$

and

$$\text{Inverted Mass Ordering} \begin{cases} F_{\nu_e} &= \sin^2 \theta_{12} F_{\nu_e}^0 + \cos^2 \theta_{12} F_{\nu_x}^0 \\ F_{\bar{\nu}_e} &= F_{\bar{\nu}_x}^0, \end{cases} \quad (4.6)$$

where the flux  $F_f^0$  and  $F_f$  are the initial neutrino fluxes of flavor  $f$  before and after oscillations.  $\theta_{12}$  is the neutrino mixing angle [54].

Non-adiabatic MSW transitions are also possible in regions where the matter density is not smoothly-varying, like the shock front [54].

## Decoherence

As supernova neutrinos begin their long interstellar journey to Earth, their mass eigenstate wave packets of width  $\sigma_x$  gradually begin to spread out and separate. The difference in group velocities of the packets causes a relative shift  $\Delta x_{\text{shift}}$  between the mass eigenstates. The coherence of these eigenstates is determined by a coherence length  $L_{\text{coh}}$ , the distance at which  $\Delta x_{\text{shift}} \sim \sigma_x$ . Beyond this distance, the neutrino mass states have separated enough so that they no longer interfere with each other, and the oscillation probability no longer depends on the baseline [68].

## MSW in Earth

If supernova neutrinos pass through the interior of the Earth, flavor transformations via the MSW effect can occur depending on the density profile traversed by the neutrinos. The detectable signature of these oscillations will appear as small modulations of the energy spectrum of a few percent at a frequency of a  $\sim 3\text{--}10$  MeV [54]. These modulations would be present in the antineutrino spectrum for the normal mass ordering case and in the neutrino spectrum for inverted mass ordering. Good energy resolution and statistics would be required to be able measure these wiggles.

Table 4.1: Expected interaction rates for a 10 kpc supernova as calculated by the SNOwGLoBES software package [69] using the s27.0c model from [55]. Event counts are summed over all channels (excluding elastic proton scattering in liquid scintillator) and assumes adiabatic MSW oscillations for both mass orderings. The effective mass for string detectors is derived from the event count for a 27.0  $M_{\odot}$  supernova and the normal mass ordering. Table adapted from Ref. [8].

	Experiment	Type		Mass (kt)	Location	27.0 $M_{\odot}$	
						NMO	IMO
<i>current</i>	Super-K	H <sub>2</sub> O	$\bar{\nu}_e$	32	Japan	7,800	7,600
	IceCube	H <sub>2</sub> O/String	$\bar{\nu}_e$	2500*	South Pole	660,000	660,000
	KM3NeT	H <sub>2</sub> O/String	$\bar{\nu}_e$	150*	Italy/France	37,000	38,000
	LVD	Liquid Scint.	$\bar{\nu}_e$	1	Italy	360	350
	KamLAND	Liquid Scint.	$\bar{\nu}_e$	1	Japan	360	350
	Borexino	Liquid Scint.	$\bar{\nu}_e$	0.278	Italy	100	97
	SNO+	Liquid Scint.	$\bar{\nu}_e$	0.78	Canada	280	270
	NOvA	Liquid Scint.	$\bar{\nu}_e$	14	USA	3,700	3,600
	Baksan	Liquid Scint.	$\bar{\nu}_e$	0.24	Russia	86	84
	HALO	Lead	$\nu_e$	0.079	Canada	9	8
	MicroBooNE	Argon	$\nu_e$	0.09	USA	12	11
	SBND	Argon	$\nu_e$	0.12	USA	16	15
<i>future</i>	Hyper-K	H <sub>2</sub> O	$\bar{\nu}_e$	220	Japan	53,000	52,000
	JUNO	Liquid Scint.	$\bar{\nu}_e$	20	China	7,200	7,000
	HALO-1kT	Lead	$\nu_e$	1	Italy	120	100
	DUNE	Argon	$\nu_e$	40	USA	5,500	5,200
	DarkSide-20k	Argon	any $\nu$	0.0386	Italy	250	250
	XENONnT	Xenon	any $\nu$	0.006	Italy	106	106
	LZ	Xenon	any $\nu$	0.007	USA	123	123
	PandaX-4T	Xenon	any $\nu$	0.004	China	70	70
Totals						776,092	775,234

#### 4.4 Supernova neutrino detection technologies

Today's neutrino detectors exhibit a variety of neutrino interaction media, geometries, materials, masses, and signal readouts. Every detector is unique, but each tends to fall into one of a few classes based on the interaction medium and therefore the neutrino flavor that is primarily sensitive to. I describe these detector classes below and Tab. 4.1 contains a list of current and future supernova-capable neutrino detectors, their detector type, and expected event rates for a canonical 10 kpc supernova.

##### Water Cherenkov

Water Cherenkov detectors instrument a volume of water with photo-multiplier tubes (PMTs). Charged particles produced from neutrino interactions travel faster than the speed of light in the medium, producing a Cherenkov cone that is then detected by the PMTs.

Because water is hydrogen-rich, these detectors are primarily sensitive to inverse beta decay interactions. Elastic scattering on electrons is another common interaction channel, which can be particularly useful for reconstructing the direction of neutrinos because the Cherenkov light cone indicates the direction of the scattered electron. Charged and neutral current interactions on Oxygen nuclei are present, but less common than the aforementioned channels.

Water Cherenkov detectors come in two varieties. The first involves storing water in a spherical or

cylindrical vessel and instrumenting the vessel walls with the PMTs. Examples of this type of detector include Super-K [70] and Hyper-K [71]. The second variety of Water Cherenkov detectors involves using a natural body of water as the interaction medium and instrumenting it with an array of long strings with optical modules attached. Examples include IceCube at the South Pole [72], KM3NeT in the Mediterranean Sea [73], and the Baikal Deep Underwater Neutrino Telescope in a lake in Russia [74].

## Scintillator

Scintillator detectors use organic hydrocarbons as an interaction medium, most commonly in liquid form. The passage of charged particles through a scintillating material excites its molecules which then isotropically emit photons. Liquid scintillator is known for producing higher light yields than Cherenkov radiation. This allows liquid scintillator detectors to have lower energy thresholds, but also leaves them vulnerable to backgrounds from radioactive decays below the MeV regime.

Liquid scintillator is rich in hydrocarbons, so much like Cherenkov detectors, inverse beta decay is the dominant interaction mode for supernova neutrinos. Examples of this type of detector include Baksan in Russia [75], LVD [76] and Borexino at the Gran Sasso laboratory in Italy [77], KamLAND in Japan, SNO+ in Canada [78], and eventually JUNO in China [79]. Both NOvA detectors, on which this thesis focuses, are liquid scintillator detectors. These detectors will be described in much greater detail in the next chapter.

## Noble element

Liquid argon time projection chambers (LArTPCs) are a relatively new detector technology that use strong electric fields to drift ionization charge generated by the passage of charged particles through the liquid argon to a collection plane. Combining the two-dimensional readout with drift time information and possibly photon detection from argon scintillation allows for precision reconstruction of particle tracks.

The water Cherenkov and liquid scintillator detectors were primarily sensitive to the  $\bar{\nu}_e$  component of the supernova signal through the dominant inverse beta decay channels, but for LArTPCs, the interaction  $\nu_e + {}^{40}\text{Ar} \rightarrow e^- + {}^{40}\text{K}^*$  provides sensitivity to the  $\nu_e$  channel. Current LArTPC detectors include MicroBooNE [80] and the Short Baseline Neutrino Detector (SBND) [81], both at Fermi National Accelerator Laboratory.

The Deep Underground Neutrino Experiment (DUNE) [82] is the next major international neutrino project. DUNE's 40 kt far detector in South Dakota has yet to be constructed, but once it begins taking data sometime after 2027, it will be a formidable supernova observatory. In fact, the ability to detect neutrinos from a galactic core-collapse supernova is one of DUNE's primary scientific goals.

Current and future dark matter TPC detectors searching for weakly-interacting massive particles (WIMPs) use liquid and sometimes gaseous noble elements as their detection medium, enabling their dual use as supernova observatories. These detectors are designed to be as radio-pure as possible to enable searches for extremely rare phenomena, like neutrinoless double beta decay. Future detectors such as LZ and XENONnT should be capable of detecting a supernova anywhere in the Milky Way galaxy at  $5\sigma$  significance [83].



## Lead-based

Currently, the only lead-based neutrino detector is the Helium and Lead Observatory (HALO) located at SNOLAB in Ontario, Canada [84, 85]. HALO's neutrino interaction medium is 79 t of lead with 128  $^3\text{He}$  neutron detectors embedded within. The primary channel is a charged-current interaction of  $\nu_e$  on protons in the lead nuclei. The  $\bar{\nu}_e$  interaction with neutrons is largely suppressed due to Pauli blocking. The  $^3\text{He}$  counters have relatively low backgrounds, which makes them useful supernova detectors. In fact, HALO is the world's first dedicated supernova neutrino detector. However, the low mass of the detector limits its sensitivity to supernovae to the nearest  $\sim 5$  kpc. An upgrade to HALO, HALO-1kT, has been proposed [86] which would consist of 1000 t of Pb from the OPERA experiment.

## 4.5 Physics opportunities

The dynamics of a supernova explosion and the properties of neutrinos are inextricably linked, and so those properties will be reflected in the characteristics of the detected neutrinos. In this section, I explore several of those properties and how they may or may not manifest in the signal.

### CP violation

The strength of CP violation in the neutrino sector, if there is any at all, is represented by the parameter  $\delta_{\text{CP}}$ . It is unlikely that a core-collapse supernova will reveal anything about this parameter because CP violation largely affects the  $\nu_\mu$  and  $\nu_\tau$  species. Current and future supernova neutrino detectors are primarily sensitive to  $\nu_e$ ,  $\bar{\nu}_e$ , or all neutrino flavors at once, and therefore have good way to discriminate between  $\nu_\mu$  and  $\nu_\tau$  events. However, it has been shown that CP violation can affect both electron neutrino species if the  $\nu_\mu$  and  $\nu_\tau$  fluxes are not equal at the neutrinosphere boundary [87].

### Absolute neutrino mass

Neutrinos travel at nearly the speed of light, but not quite. The time of arrival of a neutrino at a detector therefore depends on its energy  $E_\nu$  and mass  $m_\nu$ . After traveling a distance  $d$ , the time difference between a neutrino and a massless particle is

$$\Delta t \sim 5.14 \text{ ms} \left( \frac{m_\nu}{\text{eV}} \right)^2 \left( \frac{10 \text{ MeV}}{E_\nu} \right)^2 \frac{d}{10 \text{ kpc}}. \quad (4.7)$$

Currently, other methods of constraining the absolute neutrino mass, such as by measuring tritium beta decay endpoint energies or cosmological observations, provide better constraints than the 20 eV limit set by the neutrino data from SN 1987A [88], but it's possible that the next supernova observation could push these limits. Detectors with low thresholds and good energy resolution could stand a chance at maximizing  $\Delta t$  and thus constraining  $m_\nu$ . Using the current best limit on the neutrino mass of  $m_\nu = 1.1$  eV from the KATRIN experiment [89], a 10 MeV neutrino from a 10 kpc supernova would have a  $\Delta t$  of about 6 ms. If the absolute neutrino mass were much lower, say  $m_\nu = 50$  meV,  $\Delta t$  would drop to about 13  $\mu\text{s}$ . A sharp feature in the time profile, such as black hole formation, could enhance sensitivity [90].

### Mass Ordering

There are a variety of ways to probe the neutrino mass ordering with supernova neutrinos. Flavor transition effects all depend on this property, so observations of any of the phenomena discussed in Sec. 4.3 could provide sensitivity.

The neutronization burst is perhaps the lowest hanging fruit in this area for a couple reasons. For one, the presence of such a burst is largely model-independent. Second, the neutrinos emitted during this period are thought to be emitted from regions where neutrino self-interactions play little to no role in shaping the flux [65]. In this case, it is the  $\nu_e$ -sensitive detectors that will be able to make the most meaningful measurement: for the normal mass ordering case, the neutronization peak should be entirely suppressed or nearly so. For the inverted mass ordering, the neutronization peak will still be suppressed relative to the case of no oscillations, but it should be clearly present [54].

After neutronization, the early time profile is also sensitive to the mass ordering. Oscillation effects in the first few hundred milliseconds are thought to be largely dominated by MSW transitions and uncomplicated by other phenomena. The shape of the time profile during the accretion period could therefore be used to discriminate between mass ordering scenarios.

Mass ordering signatures are not only imprinted on the neutrino time profile. The energy spectrum may also contain information about this sought after neutrino property. Neutrino-neutrino interaction spectral splits and swaps would produce non-thermal spectra that would depend on the mass ordering. However, the current understanding of these collective effects are not robust enough to provide a consistent picture of what to expect.

## 4.6 Supernova Early Warning System and global readiness

The Supernova Early Warning System (SNEWS)<sup>4</sup> is an open, public alert system designed to issue an advance warning of an imminent core-collapse supernova based on the observations of a community of neutrino observatories [91]. This is possible because neutrinos are able to escape the center of the star minutes or hours before the explosion itself makes it to the star’s outer surface. SNEWS began operating in 2005, and has been watching for supernovae ever since. As a coincidence network, member detectors report to SNEWS when their data-driven triggers spot activity consistent with a supernova neutrino burst. If multiple detectors report a potential supernova within a 10 s window, then SNEWS issues an alert to subscribers and the public.

When a SNEWS alert is issued, the idea is that observatories around the world will stop what they are doing and search for the imminent supernova to appear in the sky. Issuing false-positive alerts was therefore a serious concern, and the system was designed to limit the probability of such a false alarm to one per century. This set the false-positive limit for reporting detectors to one per week.

Today, the experimental landscape has changed. LIGO and other multi-messenger networks now commonly issue their own alerts about various types of astronomical phenomena across the visible universe, and observatories are more accustomed to receiving alerts of various levels of urgency. Computer processing power and expanded network bandwidth have also increased since SNEWS was developed, so receiving and processing more data from contributing detectors is more feasible today than it once was.

Over the past few years, SNEWS has begun to undergo a major upgrade. SNEWS 2.0 aims to evolve beyond a simple coincidence alert system to something that is capable of taking in more detailed information from contributing detectors, performing on-the-fly combined analyses, and distributing multiple types of alerts that are updated when new information arrives. This section briefly describes

---

<sup>4</sup><https://snews.bnl.gov/>

SNEWS 2.0 in the context of these new capabilities. A more thorough explanation can be found in the SNEWS 2.0 whitepaper [8].

### Pre-supernova neutrinos

In the final stage of nuclear fusion before the catastrophic collapse, the temperature in the core increases as does the thermal emission of neutrinos [47]. Their energies don't make it much higher than a few MeV, but once they exceed the 1.8 MeV threshold for inverse beta decay, it becomes possible to detect by existing neutrino experiments. However, the flux during this silicon-burning stage is still relatively low when compared to that of the collapse itself, which limits the sensitivity to these pre-supernova neutrinos to a distance of about 1 kpc. The KamLAND experiment, for example, is capable of detecting these neutrinos, and can provide a  $3\sigma$  detection 48 hours before core collapse for a  $25 M_{\odot}$  star at a distance of 690 pc [92]. Super-Kamiokande in Japan recently doped their water with gadolinium, which will enhance their ability to tag neutrons from inverse beta decays and give them sensitivity to these pre-supernova neutrinos. They estimate that they could detect a supernova up to ten hours before core collapse with a false-alarm rate of one per century up to 600 pc away [93]. Future large dark matter experiments may also be able to detect these neutrinos [94] through the recently-discovered coherent elastic neutrino-nucleus scattering (CE $\nu$ NS) interactions [95], which unlike inverse beta decays do not have a minimum energy threshold.

### Pointing & triangulation

SNEWS hopes to not only be able to alert the public about an imminent supernova but also to provide some information about where it might appear in the sky. Considering that the most likely scenario is one wherein the electromagnetic component of the supernova is attenuated by dust in the galactic plane, it may be difficult to spot an optical signature even after the explosion has disrupted the outer layers of the star.

One method for localizing the origin of the neutrino burst is by pointing with a large number of individual neutrino observations. Forward elastic scatterings of neutrinos on electrons allow water Cherenkov experiments to reconstruct the neutrino direction based on the direction of the Cherenkov light cone (Fig. 4.6). However, these interactions are sub-dominant compared with those of inverse beta decay, and water detectors have difficulty in differentiating these two interaction types. The introduction of a dopant such as gadolinium enhances the neutron capture cross section and makes it easier to tag inverse beta decay events. Super-K has already begun to do this; they are currently the world leader in terms of supernova pointing capability with an  $1\sigma$  angular accuracy of  $4.3\text{--}5.9^{\circ}$ , and the addition of gadolinium is expected to improve this to  $3.3\text{--}4.1^{\circ}$  [8].

Another method is that of triangulation, where timing information from multiple detectors around the world are combined to constrain a region on the sky where the neutrino burst originated. This is the same technique that LIGO/Virgo uses to localize gravitational wave events. To do this, the time difference between pairwise detector observations ( $i$  and  $j$ ) is defined as

$$\Delta t_{ij} = \vec{d}_{ij} \cdot \vec{n} / c \quad (4.8)$$

where  $\vec{d}_{ij}$  is the distance vector between the two observatories and  $\vec{n}$  is the unit vector that points towards the supernova in the geographic horizontal coordinate system characterized by a right ascension  $\alpha$  and declination  $\delta$ .

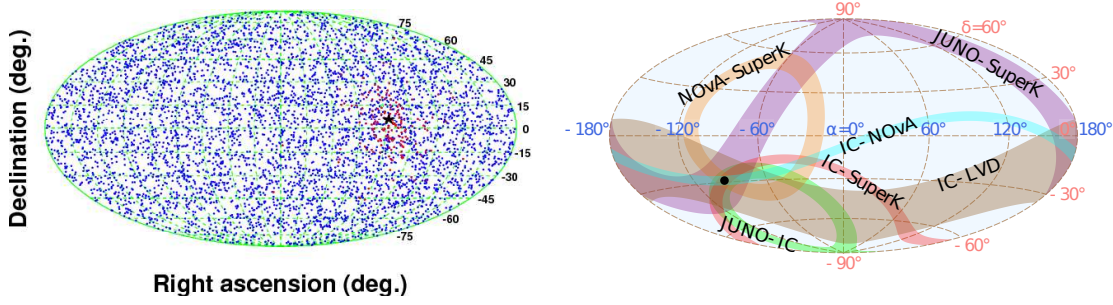


Figure 4.6: *Left*: Reconstructed event vectors for a simulated supernova in the Super-K detector. Blue points represent IBD events while red indicate elastic scatters on electrons. The star denotes the true location of the supernova. Plot from [96]. *Right*:  $1\sigma$  confidence regions for two-detector triangulation of a supernova that collapses into a neutron star. The black dot represents the true location of the supernova. Plot from [8] and adapted from [97].

Given a recorded time difference between two detectors  $\Delta t_{ij}^{\text{data}}$  and its associated uncertainty  $\delta(\Delta t_{ij})$ , which is taken to be the larger timing uncertainty of the two detectors, a  $\chi^2$  distribution is constructed as

$$\chi_{ij}^2(\alpha, \delta) = \left( \frac{\Delta t_{ij}(\alpha, \delta) - \Delta t_{ij}^{\text{data}}}{\delta(\Delta t_{ij})} \right)^2. \quad (4.9)$$

Additional detector pairs can be incorporated simply by summing their  $\chi^2$  contributions, which is then converted into a p-value to determine confidence regions on the sky (Fig. 4.6).

## SNEWS alerts

SNEWS 1.0 is a simple coincidence network; if two or more detectors report a supernova candidate within 10 s, SNEWS 1.0 issues an alert. SNEWS 2.0 aims to go beyond that by receiving a richer data payload from reporting detectors with which SNEWS can perform various “virtual experiments,” real-time analyses using certified data from a large number of experiments. Examples include those described above such as pointing and triangulation, a pre-supernova trigger, and the main supernova burst trigger. The purpose of these analyses is to generate more sophisticated and detailed alerts. Contributing experiments retain full ownership of the data they report.

These enhanced capabilities of SNEWS 2.0 requires more sophisticated cyberinfrastructure than its predecessor did. SNEWS must be able to aggregate incoming neutrino data from multiple experiments, including pointing information, neutrino light curves, and significance time-series. Once aggregated, these data need to be analyzed to determine the presence and properties of a supernova burst. Once an alert is issued, it must be tracked and updated as additional information comes in. Finally, these alerts must be combined and summarized in a way that is fit for consumption by the astronomy community. This should all happen with little-to-no human intervention. Figure 4.7 depicts a high-level architecture for the system.

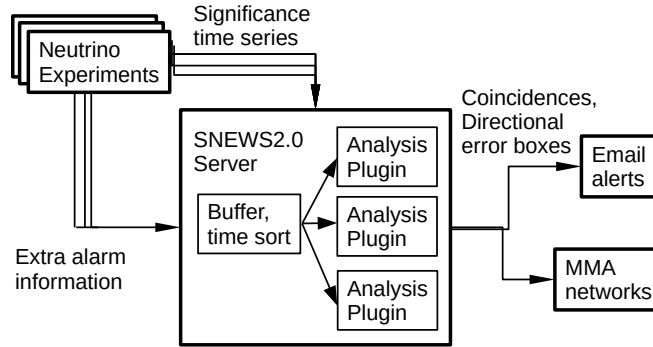


Figure 4.7: Example architecture of the SNEWS 2.0 network.

### Amateur astronomer engagement

After getting the word out through a SNEWS alert, the next most important task following a supernova neutrino burst will be to visually identify the explosion itself. Pointing information from individual neutrino observatories, SNEWS, and LIGO may or may not be available. The visible explosion could lag the neutrino burst by ten minutes, ten hours, or more. The explosion may or may not be visible from Earth, depending on how much dust is in between. There will be ample uncertainty in the moments after the neutrino burst washes over our detectors, but one thing is clear: we cannot miss the opportunity to witness a nearby supernova as the explosion bursts through the stellar surface.

We will need eyes on the skies. A lot of them. The amateur astronomer community will play a pivotal role in this search. I have always thought of the term “amateur astronomy” as something of a misnomer. This global community is home to members with a wide range of expertise, from enthusiasts to experts. Some are new observers starting out with nothing more than a pair of binoculars while others use sophisticated equipment and have decades of experience. Many are somewhere in between. Some of them will likely be the first humans to lay eyes on the next galactic core-collapse supernova, assuming it can be seen with eyes.

Until then, it is important that the SNEWS community engages with the amateur astronomy community. During my time as a member of the SNEWS collaboration, I stood up the Education and Public Outreach Working Group, whose mission is to facilitate this engagement. I developed a three-tiered strategy to guide our efforts:

1. **Awareness** — Put simply, the community cannot react to SNEWS alerts if they are not aware that they exist. This line of effort seeks to expand outside awareness of SNEWS by reaching out to organized clubs, attending conferences, and engaging with individuals on social media. Continuously cultivating community relationships in this way will help develop and maintain observational readiness.
2. **Preparedness** — This line of effort is largely an educational one. We must ensure that the public has the training and guidance they need to receive and interpret SNEWS alerts. Another component of this involves expanding the distribution of public alerts to additional media beyond email to maximize reach.
3. **Follow-up** — When the day finally arrives, execution will be key. There will be little time to

coordinate a global observing campaign. SNEWS will develop and distribute training materials containing information about existing reporting mechanisms and best practices so that observers can familiarize themselves ahead of and during a galactic supernova event.

One way to test the effectiveness of our engagement efforts is to conduct what SNEWS call a “fire drill”. This exercise would involve sending a test alert to SNEWS subscribers which is meant to mimic a true alert and directs them to begin a search for a transient object in the sky. The alert would be clearly labeled as a test, and subscribers would be notified ahead of time that a test would occur, but the exact time of the alert and the nature of the transient would be kept secret. This would allow SNEWS to test its technical infrastructure, benchmark the speed and accuracy with which the global community can identify a transient object with limited information, and evaluate which areas of community engagement could use more attention.

## **Summary**

Humanity’s understanding of supernovae has come a long way since the first-documented guest stars nearly two thousand years ago, but mysteries remain. It’s not entirely clear how they even explode. We now believe that Pauli’s “desperate remedy”, the neutrino, plays a central role. But we won’t know the details until we have the opportunity to observe a core-collapse supernova in or near our galaxy. We caught a glimpse in the late 1980s, as three neutrino-capable detectors identified a couple dozen neutrinos. Today’s neutrino detectors will collectively observe tens of thousands of neutrinos, and every last one of them will be precious. That is why the Supernova Early Warning System stands ready to coordinate the response of neutrino detectors around the world.

The remainder of this thesis will focus on two of these detectors, a pair of liquid scintillator calorimeters in the midwestern United States: the NOvA detectors.

## PART II

# NOVA AS A SUPERNOVA OBSERVATORY





## CHAPTER 5

# The NOvA Experiment

NOvA [98] is a long-baseline neutrino oscillation experiment which makes precision measurements of the  $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$  and  $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$  oscillation probabilities. The 290 t near detector is located at Fermi National Accelerator Laboratory (Fermilab) in Batavia, IL. It samples the NuMI neutrino beam at a baseline of 1 km from the beam target and sits 100 m underground. The 14 kt far detector is situated near the Canadian border in northern Minnesota (Fig. 5.1). The far detector is 810 km downstream from the beam target and sits near the Earth's surface with an overburden of 1.25 m of concrete and 16 cm of barite gravel for cosmic ray shielding. By comparing event rates and energy spectra measured at the near and far detectors, NOvA is making precise measurements of the mass splitting  $\Delta m_{32}^2$  and angle  $\theta_{23}$  and its data are sensitive to the neutrino mass ordering, the octant of  $\theta_{23}$ , and the CP violating phase  $\delta_{\text{CP}}$  [99–103]. In 2019, NOvA reported the first measurement of neutrino oscillation parameters using both neutrinos and antineutrinos [103].

In its efforts to measure these neutrino properties, NOvA has made significant contributions to neutrino event reconstruction through the development of deep learning algorithms. In 2016, NOvA developed convolutional neural network called CVN to identify neutrino interactions by their topologies without the need for detailed reconstruction [104]. This algorithm was the first of its kind to be applied to identifying interactions in sampling calorimeters and outperformed the traditional particle identification techniques that were in use before [105]. In 2019, this method was improved upon by incorporating contextual information [106] and expanded to include energy reconstruction in addition to particle identification [107].

The study of three-flavor neutrino oscillations is the primary physics objective of the NOvA experiment, but the full physics program extends well beyond that, including searches for sterile neutrino oscillations [108] and neutrino-nucleus cross section measurements [109, 110]. NOvA also searches for exotic and astrophysical phenomena, such as seasonal variations in the rate of cosmic ray muons [111], slow monopoles [112], multi-messenger signals coincident with LIGO/Virgo detections [113], and galactic core-collapse supernovae. Despite the fact that the NOvA detectors are designed to reconstruct neutrino interactions with energies one or two orders of magnitude larger than those expected from a supernova, NOvA is capable of detecting supernova neutrinos and already has a data-driven trigger in place to do exactly that [114]. In this chapter, I will describe the design of the NOvA detectors and the process of data acquisition.



Figure 5.1: Locations of the NOvA near and far detectors, separated by a distance of 809 km.

## 5.1 Detector technology and design

Although the NOvA experiment primarily uses a two-detector design, there are a total of four NOvA detectors. In addition to the near and far detectors, NDOS (Near Detector On the Surface) was a 210 t prototype detector at Fermilab which has since been decommissioned. More recently, NOvA has constructed a 30 t detector in the Fermilab test beam. The purpose of the NOvA Test Beam program is to study the detectors' response to electrons, muons, pions, kaons, and protons [115]. This will help reduce some of the largest systematic uncertainties facing NOvA analyses, such as detector response, calibration, and energy resolution. The work presented in this thesis relies only on the near and far detectors, so I will focus solely on those from this point on.

### Detector construction

The NOvA near and far detectors are segmented tracking calorimeters with liquid scintillator as an interaction medium. The detector structure itself is made of extruded hollow cells of PVC that are glued together to form planes. Each cell is  $5.59 \text{ cm} \times 3.56 \text{ cm}$  in cross-sectional area and spans a length of roughly 1600 cm at the far detector and 400 cm at the near detector. Planes of detector cells are placed together, altering their orientations by  $\sim 90^\circ$ . This provides two 2D projections of detector activity, which can be used to reconstruct particle tracks in three dimensions (Fig. 5.2). The far detector consists of 896 planes while the near detector has 214. In total, the far detector has 344,064 cells and the near detector has 20,192.

The detectors are constructed similarly, but differ in size; the 290 t near detector is about  $1/60$  the




Figure 5.2: Comparison of the NOvA near and far detectors (left) and a depiction of how individual planes of cells are alternated and staggered to allow for three-dimensional reconstruction (right).

volume of the 14 kt far detector (Fig. 5.3). The only other notable difference is that the last 22 planes in the near detector are part of a muon range stack. In this region of the detector, 10 cm thick steel plates are inserted between readout planes to provide extra stopping power for muons that might otherwise exit the back of the detector and complicate their energy reconstruction. Apart from the size and the presence of the muon range stack at the near detector, the detectors are designed identically in terms of physical construction, electronics, and data acquisition. Both detectors have an active mass fraction of 63%.

### Liquid scintillator and fiber

Each hollow cell in the NOvA detectors is filled with liquid scintillator that produces light when charged particles pass through it. The composition of the liquid scintillator is predominantly mineral oil, with pseudocumene (1,2,4-Trimethylbenzene) as the scintillant. Small amounts of PPO (2,5-diphenyloxazole) and bis-MSB (1,4-di(methylstyryl)benzene) are added to shift the light spectrum produced by the scintillant from 360 nm–390 nm to 400 nm–450 nm, which is more readily absorbed by the wavelength-shifting fiber. An anti-static agent was added to prevent the build-up of charge while filling the detector cells. Reference [116] describes the composition and production of NOvA’s liquid scintillator in greater detail.

A single wavelength-shifting fiber is inserted into each cell. The two ends of the fiber are attached to a single pixel on an avalanche photodiode (APD) at one end of the cell (Fig. 5.4) This allows the fiber to form a loop along the length of the cell. Light produced by the scintillator reflects off of the surface of the PVC cell walls until it is eventually captured by the fiber, which shifts the wavelength of the scintillation light from blue to green (490 nm–550 nm). Shifting the wavelength in this way is done to take advantage of the high quantum efficiency of the APDs in this wavelength range. One NOvA cell corresponds to one pixel on an APD, and each APD contains 32 pixels.

### Charge collection and readout

NOvA’s APDs are manufactured by Hamamatsu and have a quantum efficiency near 85% for wavelengths between 520 nm–550 nm (Fig. 5.5). They operate at a gain of 100 at the near detector and 140 at the far detector. APDs are subject to thermally-generated currents that can mimic signals, so keeping them at a low temperature is important for reducing the level of thermal noise in our data. To



Figure 5.3: Photographs of the NOvA near (left) and far (right) detectors. Credit: Fermilab/Reidar Hahn.




Figure 5.4: A NOvA FD cell filled with liquid scintillator is traversed by a charged particle (dashed line). The light produced by the scintillator (blue) reflects several times within the cell before being absorbed by the wavelength-shifting fiber (green). The light is internally reflected in the fiber to its ends which both terminate on the same APD pixel.




Figure 5.5: Liquid scintillator and fiber emission spectra and quantum efficiencies (QE) for avalanche photodiodes (APDs, pink) and photomultiplier tubes (PMTs, navy blue) as a function of wavelength in nanometers. The APDs outperform PMTs in quantum efficiency at the wavelengths emitted by the fiber.

address this, APDs are cooled using a thermoelectric cooler (TEC) which can bring them from room temperature to their nominal operating procedure of  $-15^{\circ}\text{C}$  within several minutes without significant mechanic stress. A chilled water system circulates 2 mL/s of  $15^{\circ}\text{C}$  water over the TECs to remove heat. A dry gas distribution system is used to prevent humidity in the ambient air from condensing on detector electronics.

Each APD is connected to a front-end board (FEB) (Fig. 5.6) for signal processing. An APD pulse is fed into a custom application specific integrated circuit (ASIC), which contains a preamp and shaper circuit which produce a pulse with a fast rise time and a comparatively slow fall time. An analog-to-digital converter (ADC) is used for digitization. The FEBs run with a clock speed of 16 MHz and data from the 32 channels are read out by the ASIC through a 8:1 multiplexer at the far detector, for a per-channel readout rate of 2 MHz. The near detector uses 2:1 multiplexing (8 MHz) to allow for better separation between hits due to the higher intensity of neutrino activity which creates a pileup effect.

Zero-suppression is achieved by applying a dual-correlated sampling (DCS) algorithm to the digitized signal stream. The algorithm samples a series of ADC values separated by 500 ns for every APD pulse. The ADC value at a point  $s_i$  is compared to ADC values in the recent past to determine whether the pulse exceeds a predetermined threshold value. Pulses that exceed the threshold are read out as a “hit”. Until October 2014, the far detector DCS algorithm operated in single-point mode, which uses the ADC values at points  $s_i$  and  $s_{i-3}$  to compute the difference. Since then, both detectors have been operating in multi-point mode, which uses a larger series of sample points  $\{s_{i-3}, s_{i-2}, s_{i-1}, s_i\}$  to better fit the pulse shape and more accurately determine the pulse time  $T_0$ . Figure 5.7 depicts these two DCS operating modes. Aside from thresholding, NOvA’s front-end electronics do not employ any hardware or software-based triggering; signals are continuously digitized, time-stamped, pedestal-subtracted, and zero-suppressed.






Figure 5.6: NOvA readout electronics schematic (left, viewed from side) and photo (right, viewed front-on).




Figure 5.7: Illustration of the advantage of multi-point over single-point readout. The red and black lines represent two similar pulse shapes which are sampled very 500 ns. In single-point mode, the ADC and TDC values at point  $s_i$  are taken to be the ADC and TDC values of the hit. These two pulses might be read out with the same values, even though the red pulse is delayed relative to the black one by about 500 ns. In multi-point mode,  $s_i$  still represents the hit ADC value, but all points in the set  $\{s_{i-3}, s_{i-2}, s_{i-1}, s_i\}$  are used to fit the pulse shape and more accurately determine the TDC value of the pulse.

### APD baseline sag

Although the channels operate independently, they can sometimes affect one another. Each channel on a given APD shares the same voltage source which supplies the baseline and when a pulse is produced on one APD channel, the baseline for all channels on that APD momentarily decreases slightly. This sag in the baseline depends on the amount of energy deposited on the APD, and when energy depositions are very large, the sag can be large enough that it causes the DCS values to exceed threshold as the voltage returns to baseline, mimicking a real pulse. This can create a ringing effect, which we call an “FEB flasher”, where some or all of the channels connected to a given FEB are repeatedly re-triggered at low ADC values for an extended period of time as the instigator pulse decays (Fig. 5.8).

NOvA test bench studies have revealed the conditions under which this FEB flasher phenomenon is




Figure 5.8: A far detector event display of an FEB flasher in the YZ-view. Pixel color and size represent hit ADC (darker and larger correspond to higher ADC). A high-energy shower deposited large amounts of energy into the cells along the track, causing APD baselines to sag and false hits (flashers) to occur in the affected FEBs.






Figure 5.9: A test bench digital scanning oscilloscope (DSO) scan showing a large charge deposition in channel 1 (gray) which induces an APD baseline sag in channel 2 (red) on the same FEB. The left panel shows the raw ADC values associated with the deposition and the right panel shows the DCS view which is used to read out the data if the value exceeds a threshold. In this example, the induced baseline sag causes channel 2 to take on a positive DCS value which exceeds threshold as the voltage is returning to baseline.

likely to occur. At the far detector, an FEB flash is expected if an instigator signal is 10,000–20,000 ADC counts ( $\sim 250$  times threshold). The near detector uses a different set of constants for signal shaping and thus these values are slightly different; about 30,000 ADC counts are needed to produce FEB flashing at the near detector, which represents a signal about 700 times threshold. Figure 5.9 depicts how one channel with a large charge deposition can influence a neighboring channel to produce this effect.

Hits from FEB flashers do not represent real detector activity and are therefore a background to analyses, especially a supernova analysis in which the ADC spectra of the signal and the flashers would overlap. Luckily, the fact that FEB flashers can be identified by high-energy events, are constrained to a limited set of FEBs, and have a predicable duration means that they can easily be identified and removed by software. I will revisit this topic in Sec. 7.2 when I discuss background rejection methods for this analysis.

## 5.2 Data acquisition

The NOvA data acquisition (DAQ) system is responsible for aggregating data streams from thousands of FEBs, making them available to software-based triggering algorithms, and writing them to disk for long-term storage. Data within the DAQ are organized in a hierarchical way. The most atomic data structure is called a nanoblock, which represents a single hit in the detector and contains information about the ADC pulse height, timestamp (more on timing in next section), and FEB and APD pixel addresses.

Data Concentrator Modules (DCMs) aggregate the data read out from up to 64 FEBs and services a geographical region of the detector. The near detector has 14 DCMs and the far detector has 168. DCMs concatenate streaming data from the FEBs into 5  $\mu$ s chunks called microslices. The Event Builder then groups multiple microslices together to form a microblock, which is 50  $\mu$ s in duration and represents a contiguous time period of data that has been read out by a single DCM. These microblocks are the basic

Table 5.1: Data aggregation layers from APD readout to disk storage for near and far detectors.

Layer	Input Streams (FD top, ND bottom)	Output Streams (FD top, ND bottom)	Channels Per Stream
APD Readout → FEB	344,064	10,752	32
	20,192	631	32
FEB → DCM	10,752	168	2,048
	631	14	~1,442
DCM → Buffer Node	168	200 (max)	344,064
	14	200 (max)	20,192
Buffer Node → Data Logger	200 (max)	1	344,064
	200 (max)	1	20,192
Data Logger → Disk Storage	1	21	344,064
	1	11	20,291

atom of data in the circular buffer system and are used as input to trigger algorithms that will decide which data are associated with physics events of interest to be stored and those which are not and to be discarded.

The circular buffer system that NOvA employs is made of a farm of computers (known as buffer nodes) which are connected by a shared memory segment. The depth of the buffer depends on the number of buffer nodes in the system; in the current configuration, the near detector buffer depth is approximately 30 min and that of the far detector is about 20 min. Triggering algorithms—which will be described in greater detail in Sec. 5.4—examine these buffered data in detail. When a trigger is issued for some time period, those data must be read from the circular buffer and written to disk. This buffering architecture is also useful because it takes longer to write data to disk than it does to stream it out of the detector. The data must also be packaged with important metadata, like run and subrun numbers. Table 5.1 summarizes the way in which data is aggregated into many APD channels to several files on disk.

### 5.3 Timing system

NOvA’s performance as a neutrino detector relies on meeting certain timing requirements. One requirement is intra-detector synchronization. This is particularly important at the near detector where the neutrino flux is large and the spatial spread of the beam is small. A single NuMI beam trigger often has multiple neutrino events present which overlap each other. To mitigate this pile-up, the detector components must be synchronized to one another to within 10 ns.

We also require synchronization of both detectors with an external reference clock. Both detectors rely on spill signals sent from Fermilab to decide when to record neutrino event data. If the detectors are not in-sync with each other, then one or both may fail to read out data at the right time and miss neutrino interactions. From a supernova perspective, detector synchronization is also very important. We need both detectors to read out a massive amount of data in the event of a supernova, and if they are not properly synced to an external reference, we may lose some or all of the data associated with the neutrino shock breakout.

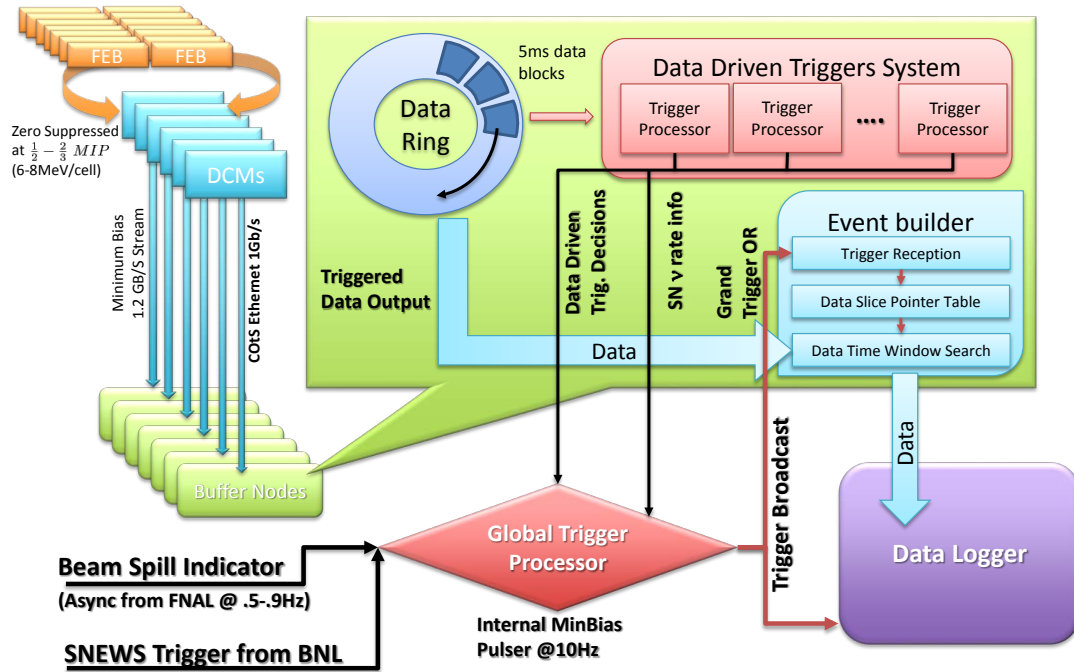


Figure 5.10: The NOvA DAQ and its subsystems.

## NOvA time base & clocks

The NOvA DAQ uses a custom time base for all clocks and timestamp counters. This time base is a 56-bit wide counter which represents the number of 64 MHz clock ticks since January 1, 2010 06:00:00 UTC. This NOvA epoch has a duration of 35.7 years before rolling over, and will end in 2045, well after the NOvA experiment has finished operating.

A 32 MHz clock is distributed across all components of each detector and serves as the base clock that drives data acquisition. This clock is derived from a phase lock loop (PLL) circuit which is driven by a 10 MHz Conor Winfield disciplined oscillator module. DCMs also contain undisciplined oscillators which can be used to derive 32 MHz and 16 MHz frequencies during diagnostic work when the full detector is decoupled from the central timing system.

## Timing chain and hardware

The timing system is made up of timing distribution units (TDUs) and DCMs. There are two types of TDUs: a main TDU (MTDU) which is responsible for originating timing commands and a secondary TDU (STDU) which is a repeater used to propagate signals coming from the MTDU to the DCMs. The timing system has a daisy chain design for distributing signals throughout the detector. The MTDU and several STDUs for a given detector are connected in series, forming a timing chain “backbone”. Each STDU also connects to two separate chains of DCMs. The final STDUs and DCMs in each chain have loopback connections which allows propagating signals to reflect back up the chain once they reach the end. Each DCM fans out to up to 64 FEBs. Redundancy is built into the system to mitigate any component failures along the timing chain. Each DCM on both detectors is physically connected to two physically distinct timing chains, only one of which is active at any one time. If one timing chain

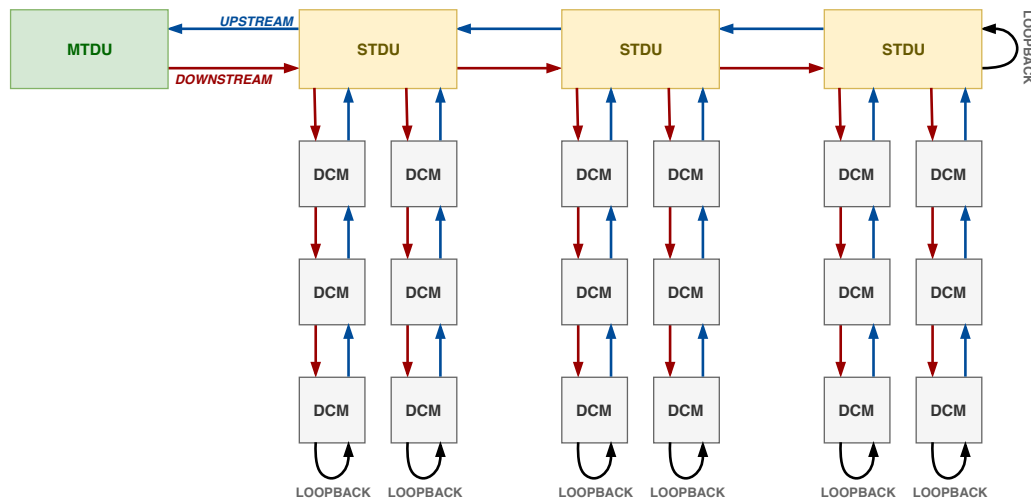


Figure 5.11: A NOvA timing chain. The main TDU (MTDU) and secondary TDUs (STDUs) form the backbone, and each STDU connects to two branches of DCMs. The end of the backbone and each DCM branch contains a loopback connection so that delay-learn signals can be echoed back through the chain. The far detector has 14 STDUs and 168 DCMs while the near detector has 2 STDUs and 14 DCMs.

fails, the second can be deployed by remote intervention by experts. In addition to this redundancy, a third MTDU exists at each detector to be used as a cold spare, but is not physically connected to the detectors' timing chains. A diagram of the timing chain topology is shown in Fig. 5.11.

### Delays and system synchronization

The TDUs and DCMs are connected by copper cables, and cable delays prevent the timing signals from reaching all parts of the detector simultaneously (Fig. 5.12). That's a problem if we're trying to keep everything in sync. We address this by calculating these delays before sending timing signals. For this procedure, the MTDU sends a signal through the system that is then reflected back in the opposite direction at the loopbacks. The time between a TDU or DCM seeing the signal initially and seeing it on its return tells that component how far away in cable-delay time it is from the end of the chain. Once that delay is learned, the component will not act on any future signals that it receives until the signal has reached the farthest component in the chain. This delay-learn procedure does not need to be repeated often; it is only necessary when components or cables are swapped out.

MTDUs are equipped with a GPS antenna and receiver for external synchronization. To synchronize all detector components, the system uses a "at the tone, the time will be" approach. The MTDU first sends a signal with a future timestamp for all components to set their clocks to. The MTDU will then send a signal that tells the components to start their preset clocks. Upon receiving the signal, a component will wait to start their clock according to their pre-learned delay. The result is that every component in the system starts their clocks at exactly the same moment, and the clocks are preset to be synchronized with a highly accurate GPS time.

### Timing Calibration

The MTDUs use their GPS receivers to periodically synchronize their clock. In the event that satellite lock is ever lost, the disciplined oscillator in the MTDUs takes over and may drift at a rate of about 2




Figure 5.12: Typical timing signal delays along the length and sides of the far detector. Each small square on the side and top of the detector represents one DCM.

parts in a billion per day. To monitor for such a drift, we use a timing calibration reference (TCR) with each TDU to check that the TDU times are accurate and not drifting too much. The TCR is a modified TDU with its own GPS receiver that is connected to a separate GPS antenna than the MTDU. The TCR is physically connected to the TDU that it is associated with and generates a 1 Hz timestamped signal that the TDU also timestamps. By comparing the offset between these timestamps, we can determine whether the MTDU has drifted from the current GPS time too much and thus warrants a resync with the GPS and of all other timing system components. These time offsets between the MTDUs and their TCRs are displayed on a website that is monitored by NOvA scientists. If a sufficiently large drift is detected, high-visibility alerts are displayed on the screen so that an expert can be notified.

Each TDU stores the computed timing delays between itself and the end of the chain, but that is not the case for the DCMs. DCMs therefore do not remember how far away from the end of the chain they are and so they don't have an appropriate delay value to use when syncing with the rest of the timing system elements. Although the DCMs are not synced in this way at a hardware level, we can derive their delay values offline. Part of the standard offline calibration process at NOvA involves computing and applying these delays. Cosmic-ray muon tracks provide many straight tracks in the detectors that often cross one or more DCM boundaries. By using the timing information at the borders between adjacent DCMs we can adjust the timing to make them match [117]. Because the DCMs and their cables are rarely changed or moved, we only occasionally have to apply this procedure.

### Spill signal distribution

A “spill” occurs when the NuMI accelerator extracts a bunch of protons and they strike a graphite target. When the accelerator extraction kicker fires, a signal is sent from the accelerator complex to a near detector MTDU which is specifically designated to decode and interpret these spill signals. This special TDU is referred to as the “NOvA spill server” (NSS), and differs from the other MTDUs in that it is not responsible for managing a timing chain. Each signal from the accelerator is decoded, timestamped, and placed into a queue on the spill server TDU.

A piece of software called the NssTDUApp, which runs on the spill server MTDU at the near detector, consumes spill messages from the queue, creates a XML-RPC packet, and distributes that packet to other spill server elements over the internet. The packets are initially sent to an application called a Spill Forwarder (NssSpillForwarder). Its purpose is to serve as a repeater; it receives packets from the spill server TDU and forwards them on to any number of remote hosts. Those remote hosts run an application called a Spill Receiver (NssSpillReceiver). The Spill Receiver listens for and reads the XML-RPC messages that it receives from a Spill Forwarder and converts it into a DDS<sup>1</sup> message that can be passed to the Global Trigger process. The Global Trigger uses information about the timing of the spill to decide which chunks of data should be pulled from the detector's data buffer and written to disk. These various components—The NssTDUApp and one instance of both NssSpillForwarder and NssSpillReceiver for each detector—must all work in concert for each detector to trigger at the right time. What begins as an electrical signal from the Fermilab accelerator complex is interpreted and timestamped by a GPS-synchronized TDU, repacked as an XML-RPC packet, and sent to each detector over the internet, and ultimately delivered to each detector's triggering system to save the relevant data.

This architecture takes advantage of the way the detectors buffer data. As long as the latency of spill packet delivery is much smaller than the depth of the data buffer—and it nearly always is—the spill signal can arrive well after the beam spill has been delivered to both detectors without the risk of losing the data. The architecture of the spill server system for each detector is detailed in Figs. 5.13 and 5.14.

## Monitoring and maintenance

There are several points of failure for this system. One is that packets are sent to the detectors' global trigger systems through local and wide area networks. If a critical router or switch goes down, spill information may not make it to one or either of the detectors. Another potential failure point is at the software level. Between both detectors, there are five pieces of software that need to be running and communicating with each other to allow accelerator signals to propagate to their intended destinations. Failures are fairly rare, but they do happen on occasion. Network problems are largely out of our control, but we have several monitoring tools and shifting protocols in place to identify problems as they arise.

One monitoring method involves watching for the presence or lack of heartbeat triggers. Along with signals from the beam extraction kicker, Fermilab also provides a 1 Hz heartbeat signal that we trigger on at NOvA. This heartbeat occurs even when the beam is not operating. If neither detector is registering a 1 Hz trigger from Fermilab, there may be a problem with the NssTDUApp, and shifters are trained to recognize this and contact an expert. If the heartbeat is seen at one detector and not the other, then perhaps one of the NssSpillReceiver or NssSpillForwarder applications is the problem.

We also have software that continuously monitors the machines that these applications run on. If one of these processes dies, the spill server monitor will notice and change the color of an indicator on the shift screen from green to red. This makes it easy for our shifters to quickly detect problems and contact the relevant expert. Even if the NssTDUApp stops running, the spill server TDU will continue to accumulate spill signals in its queue. Once the software is running again, it will consume the backlog of spill signals like normal. As long as the relevant data associated with those spills is still

---

<sup>1</sup> DDS, or Data Distribution Service, is a real-time messaging system that is used by many components of the NOvA DAQ.

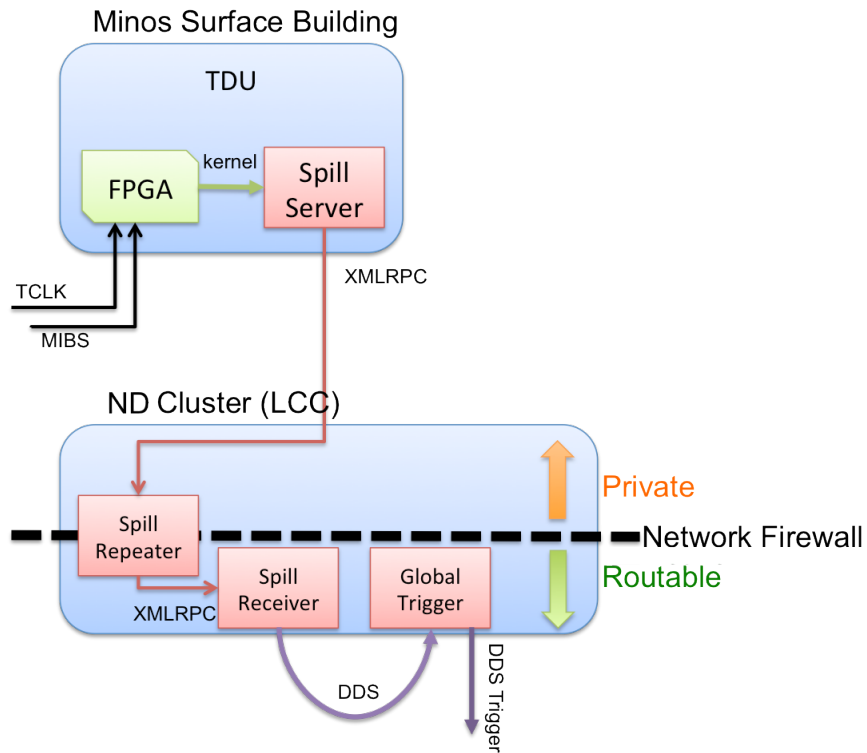


Figure 5.13: Spill server architecture at the near detector.

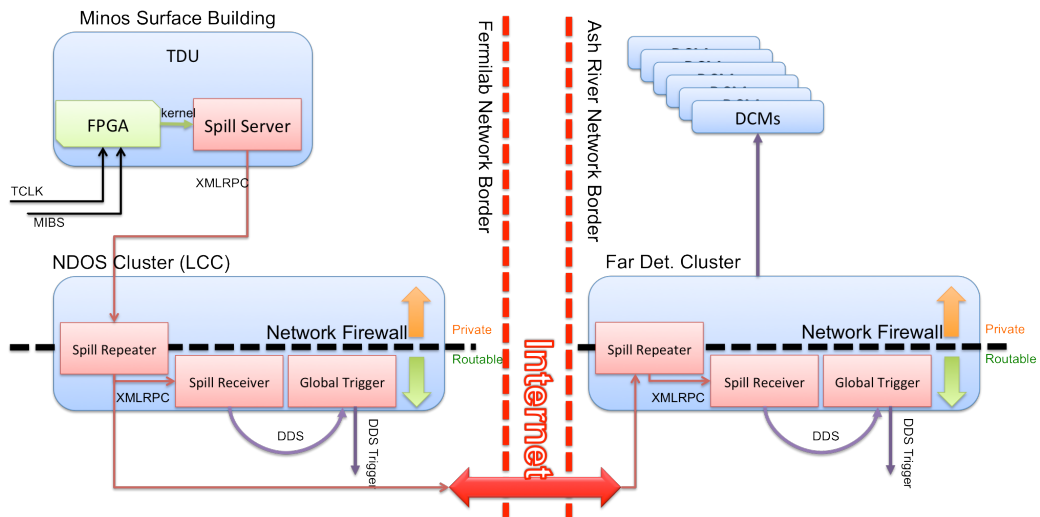


Figure 5.14: Spill server architecture at the far detector.

in a detectors' data buffers, it will be saved. Because of this design, a spill server downtime that lasts for several minutes will result in no lost beam data.

## 5.4 Triggering system

The NOvA detectors are continuously reading out any signals that exceed the FEB readout threshold. A large portion of these data have no significant scientific value and are discarded as they reach the end of the circular data buffer. This is particularly true for the far detector which sits on the Earth's surface with only a modest barite overburden for shielding, and is subjected to a cosmic-ray muon rate of  $\sim 150$  kHz. This amounts to approximately 1.2 GB/s of raw hit data being streamed from the detector's 10,752 front-end boards, only a fraction of which represents activity associated with the physics processes that we are interested in.

The NOvA triggering system exists to determine which data in the buffer are worth storing and which are not. Triggering decisions can be data-driven or externally-driven. Data-driven triggers (DDTs) are made by examining the data in the buffer and checking whether they meet any number of predetermined criteria. Externally-driven triggers occur when some outside entity reports the occurrence of an interesting physics event. This most commonly occurs when Fermilab sends us a beam-spill signal. Another example is that of multi-messenger astronomy networks that report gravitational wave or supernova events. There is also a pulsed trigger, which we use to take regular triggers of minimum-bias cosmic data at an approximate rate of 10 Hz.

### Global Trigger and Data Logger

Every trigger, regardless of its source, is issued by a multi-threaded process called Global Trigger (GT). DDTs examining the buffered data or external events can send trigger requests to the GT. Trigger request consists of a start time, trigger duration, and the trigger type. GT sends the request to all the data buffer nodes. Upon receiving a trigger request, the buffer nodes package up the data requested and send it to a process called Data Logger, which is responsible for receiving data from the buffer nodes, collating the data with trigger information, and writing those data to disk. The data are written as binary files, with each file corresponding to a given run, subrun, and trigger stream. This keeps data that were tagged by different DDT algorithms or external events separate from one another, and makes it future analyses easier.



## CHAPTER 6

# Simulation and Supernova Triggering

The NOvA detectors are designed to reconstruct events from a 2 GeV neutrino beam. Supernova neutrinos have energies several orders of magnitude below this which presents several problems. For example, the spatial extents of supernova hit clusters are small and difficult to reconstruct, supernova neutrino energies are near—and sometimes below—the FEB triggering threshold, and most of NOvA’s instrumented mass sits near the Earth’s surface and is subject to a significantly larger cosmic flux than most other supernova neutrino detectors (Fig. 6.1).

NOvA’s design does have some inherent advantages though. The timing resolution of the detectors means that cosmic activity at the far detector is relatively quiet for small time windows. The data buffer is minutes deep, allowing us to read out a supernova event even when the trigger is delayed. The NOvA DAQ is capable of streaming continuous data from the detectors onto disk without incurring downtime.

In the face of these advantages and disadvantages, the NOvA detectors are capable supernova neutrino detectors. But the stakes are high; NOvA needs to be ready to trigger on the next galactic core-collapse supernova and save the precious data to disk without corrupting it. In this chapter, I will describe the NOvA supernova triggering infrastructure and data quality checks that ensure we are capable of writing the data to disk.

### 6.1 Supernova neutrino interactions in NOvA

Core-collapse supernovae emit neutrinos of all flavors, but the NOvA detectors are more sensitive to some than others. The detector materials are largely hydrocarbon-based, meaning the most prominent interactions are those of inverse beta decay, elastic scattering on electrons, and neutral current interactions on carbon (Fig. 6.2).

Inverse beta decay (IBD) is a charged current interaction between a proton and an electron antineutrino which produces a positron and a neutron. This is the dominant supernova neutrino interaction type for NOvA due to the abundance of hydrogen atoms. The final-state neutron scatters throughout the detector, traversing up to a couple dozen cells before thermalizing and capturing on a nucleus, usually that of chlorine. The nuclear de-excitation often proceeds by gamma emission, which can pair-produce detectable electrons. The positron is charged and therefore directly observable in the NOvA detectors. They are much lower in mass than the neutrons and as a result rarely make it farther than a few cells away from the point of production. The energy of the positron and that of the parent neutrino are related by a simple approximate kinematic relation:  $E_{e^+} \simeq E_{\bar{\nu}_e} - 1.29 \text{ MeV}$  [118].




Figure 6.1: Depth of various current and future supernova-capable detectors. For string detectors that span a range of depth, a gray bar is used to indicate the distance between the minimum and maximum depths. The red bars are the NOvA detectors.




Figure 6.2: Cross sections for relevant neutrino interactions in liquid scintillator.

Table 6.1: Predicted event rates by channel for both detectors given a 10 kpc supernova under the LS220-s27.0co and SFHo-z9.6co models.

Channel	Interaction	Far Detector		Near Detector	
		9.6 M <sub>⊙</sub>	27 M <sub>⊙</sub>	9.6 M <sub>⊙</sub>	27 M <sub>⊙</sub>
Inverse beta decay	$\bar{\nu}_e + p \rightarrow e^+ + n$	1,593	3,439	24	51
Elastic scattering on $e^-$	$\nu + e^- \rightarrow \nu + e^-$	143	259	3	5
Neutral current on $^{12}\text{C}$	$\nu + ^{12}\text{C} \rightarrow \nu + ^{12}\text{C}^*$	67	166	1	3

Elastic scattering on electrons occurs when an incoming neutrino of any flavor scatters elastically with either a free or an atomic electron. The incident electron is at rest, but the outgoing one is not, producing a detectable signature in the detectors. While this process can occur for neutrinos of all species, it is most pronounced for the electron flavor. This process, unlike IBD, preserves the momentum of the incident neutrino and therefore contains directionality information which could be exploited to point back to the location on the sky where the neutrino originated. This preservation of directionality is advantageous for water Cherenkov detectors where a light cone is projected along the axis of the electron's momentum, but liquid scintillator detectors require good spatial resolution to be able to reconstruct the path of the electron. In the NOvA detectors, these path lengths only span a couple of cells, which makes directionality reconstruction extremely difficult.

The third-most-dominant interaction class in the NOvA detectors is the neutral current interactions on carbon atoms. Here all neutrino flavors are represented, and the result is a nuclear excitation of the carbon atom which eventually de-excites by emitting a 15.1 MeV gamma [119].

## 6.2 Flux models

The supernova simulations described in Sec. 3.3 produce many outputs, but the one we're most interested in for our own simulation purposes is the predicted neutrino flux over time by flavor. There are many models to choose from, I will primarily focus on two 1-D models from the Garching group [55]. The LS220-s27.0co model is for a 27 M<sub>⊙</sub> progenitor star with the LS220 nuclear equation of state [120]. The SFHo-z9.6co model is for a 9.6 M<sub>⊙</sub> progenitor with the SFHo equation of state [121].

The model fluxes are organized as tab-delimited text files with four columns: time since core bounce in s, luminosity in  $10^{51} \text{ erg s}^{-1}$ , average neutrino energy in MeV, and average square neutrino energy in  $\text{MeV}^2$ . From this information we reconstruct the flux at each time step  $t$  using the pinched-thermal spectral shape  $f_{\alpha,t}(\epsilon)$  (Eq. 4.2) introduced in Sec. 4.2:

$$\Phi_t(\epsilon) = \frac{1}{4\pi d^2} \left( \frac{L_t}{\langle \epsilon \rangle_t} \right) f_{\alpha,t}(\epsilon). \quad (6.1)$$

where  $d$  is the distance to the supernova. The neutrino luminosity profiles and energy fluences for the two models are shown in Fig. 6.3. Table 6.1 lists calculated event rates for the three most dominant interaction channels for both detectors and neutrino flux models.

## 6.3 Simulating a supernova in the NOvA detectors

Simulating a supernova in the NOvA detectors happens in several stages. The first stage is event generation where the expected supernova neutrino flux is convolved with the detector material cross sections




Figure 6.3: Simulated neutrino luminosity evolution over time (left) and energy fluence (right) for the SFHo-z9.6co model (top) and LS220-s27.0co model (bottom) from the Garching group [55]. Taken from [122].

and geometry information to produce neutrino interaction vertices complete with final-state particle kinematic information. The next stage simulates the passage of those final-state particles through the detector geometry. Those energy depositions are then simulated to account for the expected response of the fiber and detector electronics. Finally, the expected background activity is overlaid with this signal simulation to produce a final result that is formatted the same way as any real data file. In the rest of this section, I will describe each of these steps in greater detail.

## Event Generation

The purpose of event generation is to determine which neutrinos from a given flux profile will interact inside of the detector volume, and where and when those interactions occur. The output is a list of outgoing particles and their kinematic properties. To do this, we use the GENIE event generation software [123] with some custom modifications, which we call GenieSNova [122]. For one, we supply our own flux driver to handle the time-varying flux profile that is typical for supernovae. This is not a common workflow for the average GENIE user who is typically interested in accelerator or atmospheric neutrino experiments. One of our inputs to this stage of the simulation therefore is the flux models described in the previous section.

We also provide cross sections as inputs. In particular, a subset of those described in Fig. 6.2. The IBD and elastic scattering cross sections are included, but the neutral current on carbon is not. These cross sections as well as the nuclear de-excitations that accompany this channel are not included out-of-the-box in the version of GENIE that we use. Future versions of GenieSNova will include this interaction channel, as well as other sub-dominant channels.

One limitation to the version of GENIE that we use—and thus to the GenieSNova software—is that neutrinos with energies less than 10 MeV are not supported. While we would prefer to simulate the full range of supernova neutrino energies, neutrinos in this regime often produce final state particles that are below the detection threshold of the detectors, and therefore do not appreciably contribute to the overall signal anyway.

## Detector Simulation

The output from GenieSNova is a list of outgoing particles from the neutrino interaction vertices with their associated positions and kinematics. The next stage is to take these outgoing particles and propagate them through the geometry of the detectors, simulated the interactions and energy depositions these particles have as they traverse the detector materials. We use GEANT v4.10 [124] to do this, without any additional modifications.

The output from GEANT is a list of energy depositions that are then fed to our in-house detector response simulation. At this stage, the response of the scintillation light is modeled, as well as its propagation as it bounces around the cell and gets absorbed by the wavelength-shifting fiber. The attenuation properties of the fiber is also simulated as the light signal travels through it towards the APD. The APD response is simulated as is the FEB shaping of the signal. The final result of this stage is a file that includes raw hit information that is identical to that in a real data file. At this point, we can treat the output of the simulation as if it had actually streamed out of the detector.

## Minimum-bias overlays

The event generation and detector simulation stages produce files that are formatted identically to real data files from the detectors, but they only contain activity related to the supernova neutrino signal. Background activity is not simulated. Instead, we rely on combining the simulated signal activity with that of real backgrounds from data.

The background data come from a daily readout that also serves as a test of the long-readout supernova triggering infrastructure that will be described in greater detail in the next section. The event that triggers the readout is a heartbeat signal from SNEWS. The data generated from these tests provide a large library of minimum-bias data that we overlay on top of simulated supernova data to produce realistic predictions of what we would expect to see in a real supernova event.

The full simulation chain, from event generation to minimum-bias overlays, is depicted in Fig. 6.4.

## 6.4 Data-driven supernova trigger

NOvA is subscribed to direct alerts from SNEWS and LIGO, and when an alert is received, the NOvA DAQ is configured to immediately read out 45 s of continuous data to disk starting 5 s before the time of alert issuance. These are important triggering mechanisms, but they have an external locus. We have also developed an internal trigger, which relies on NOvA detector data alone to make supernova triggering decisions. This allows NOvA to stay sensitive to a potential supernova when the connection to external sources is unavailable, and it also enables NOvA to become a contributing member of the SNEWS network. This section describes the operation of our data-driven supernova trigger and its sensitivity to galactic supernovae. A more detailed description of this system can be found in Ref. [122].

## Background Rejection

Cosmic-ray particles are the most obvious source of background activity, particularly at the far detector which has an overburden of only 3.6 meters water equivalent and observes an average cosmic ray rate of 148 kHz. The near detector, on the other hand, is located 100 m underground (225 meters water equivalent) and only experiences a cosmic ray rate of 37 Hz. These particles can produce a variety of secondary particles, such as neutrons, deltas, spallation products, and Michel electrons. Michel electrons

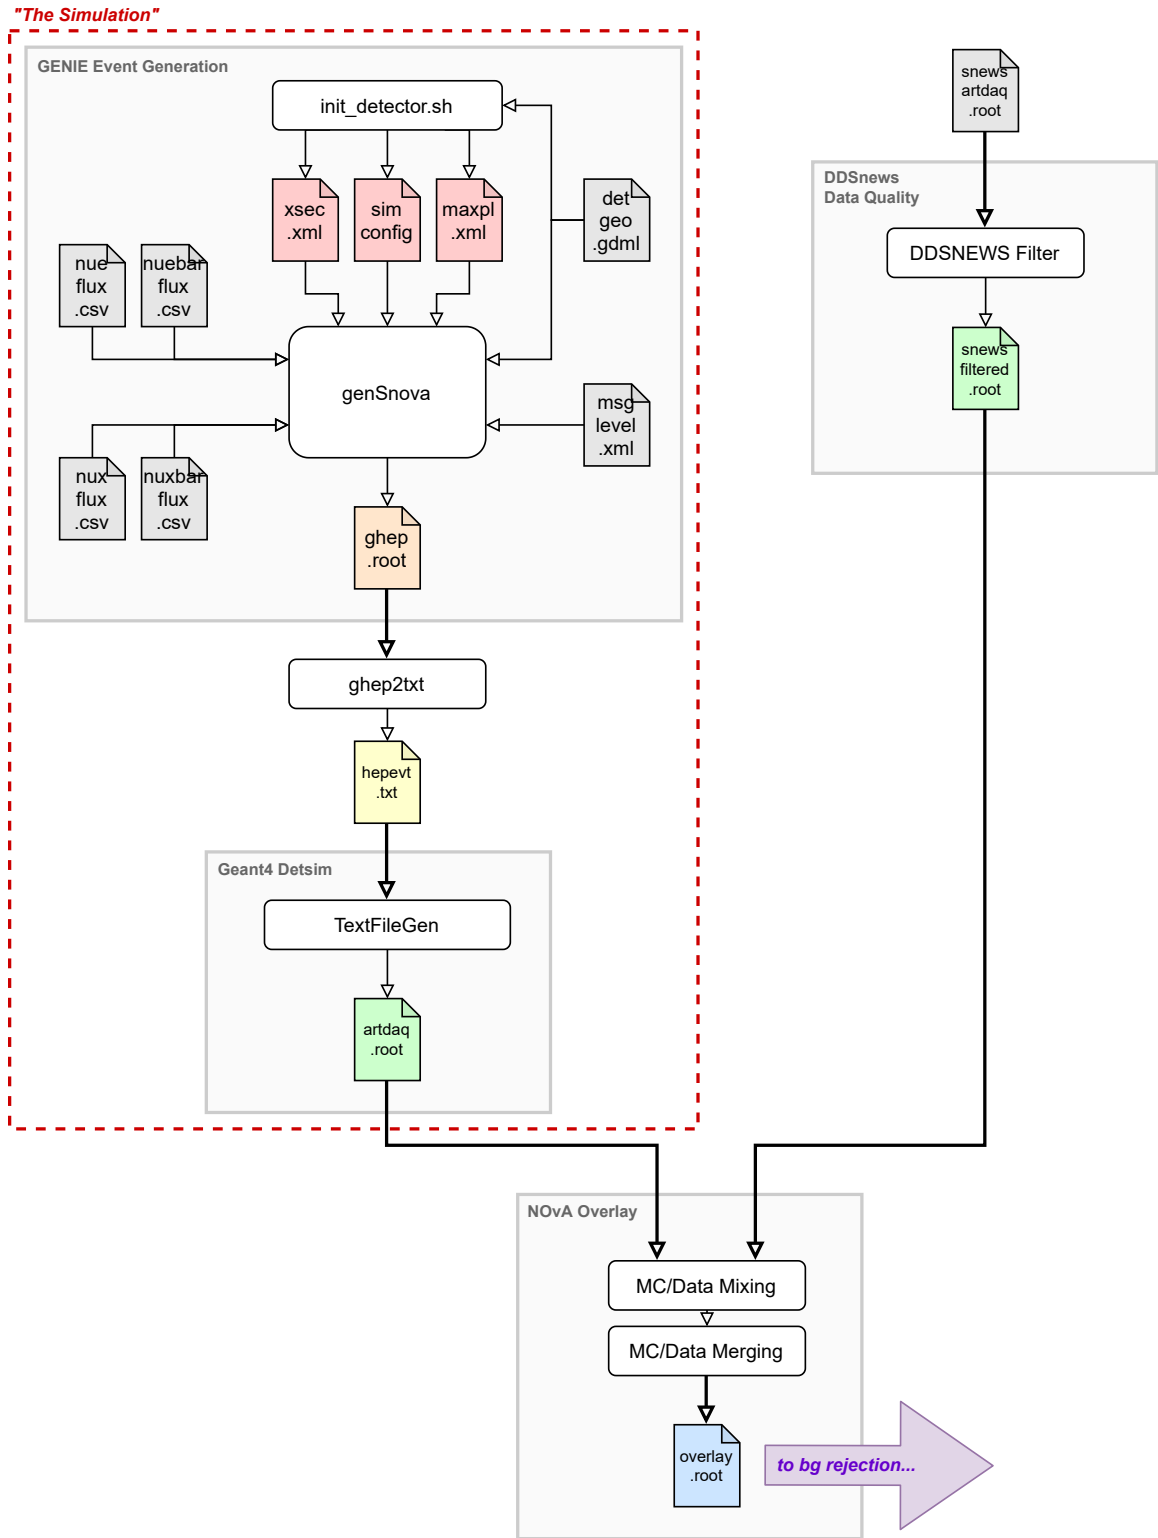


Figure 6.4: NOvA supernova simulation chain.

are particularly troublesome because their energies are similar to those of supernova neutrinos.

To reject Michel electrons, muon tracks are identified using a Hough transform-based algorithm [125]. Any trajectory end-point that occurs within the detector volume is treated as a potential stopping location of the muon, and hits that occur within 32 cm and 10  $\mu\text{s}$  of the point are rejected. All hits that are associated with the muon track itself are also removed.

High-energy air showers are produced by interactions of cosmic rays with the materials surrounding the detectors and produce high hit multiplicities. These bursts of activity can overwhelm reconstruction algorithms and drown out any signal we may hope to recover. For this reason, a simple time window veto is applied to all activity that occurs during one of these events. A rolling window of 1.1  $\mu\text{s}$  is applied to the data, and a high-energy shower is identified whenever the summed ADC of hits in a window exceeds 1,000,000 at the far detector or 300,000 at the near detector. When this condition is identified, all detector activity in the following 350  $\mu\text{s}$  is vetoed.

Dark current and leakage current in the detectors' circuitry can generate false hits that resemble those from supernova neutrinos. This effect can become more pronounced when FEBs or APDs are in need of maintenance or replacement. In these cases, hit rates can be suppressed or enhanced, and in either case we do not want to trust the hits coming from that component. We compute an hourly hit rate map to gauge which channels should be considered too hot or too cold to be included in the analysis. A cold channel is one that has an inactivity fraction exceeding 90% and a hot channel is one that has an average hit rate exceeding 1 kHz. A channel that is masked off by this map can become unmasked if it returns to nominal operation for a period of twenty-four hours.

## Clustering

Any hits that survive the background rejection process are passed through to a clustering algorithm. First, hits are grouped into 1  $\mu\text{s}$  clusters. Each temporal cluster is then decomposed into additional clusters based on their spatial proximity. The clusters formed this way have a large variation in topologies because no constraints have been applied to their formation based on the expected properties of the supernova signal. The only criteria that clusters are required to satisfy are that (1) hits from the same cluster must not be separated by more than one cell in the same plane, and (2) clusters must contain hits from both planes. These criteria make it less likely that we will cluster correlated noise activity between several cells that share the same front-end board electronics. Criterion (2) also ensures that we can deduce the cluster position in three dimensions, which is required for computing the time of the cluster.

Before the time of the cluster can be computed, the times associated with each cluster hit must be corrected. The timing system described in the previous chapter measures the time delay between timing units along the detector and applies those delay corrections at the hardware level. However, there are similar delays associated with signal travel time between the timing units and data concentrator modules (DCMs) that run along the sides of the detectors. These delays are not accounted for at the hardware level, so corrections need to be applied in the software. With the 3-D position information of the cluster, each hit's location within a cell can be deduced. A correction is applied based on the distance to the end of the cell, the distance from the end of the cell to the APD, and the distance from the DCM to the backbone of the timing chain. These distances are shown in Fig. 6.5.






Figure 6.5: Timing corrections applied to a hit based on its location in the cell, the distance between the cell readout and the APD, and the distance of the DCM from the rest of the timing chain. Taken from Ref. [122].




Figure 6.6: Time profile of supernova candidate clusters at the far detector before and after selection. Taken from Ref. [122].

## Selection

We apply a series of selection cuts to the clusters to exclude those which are not consistent with supernova activity. First, we apply a fiducial volume cut. At the far detector, any clusters that are positioned within 1.5 radiation lengths (16 cell widths) of a detector wall is excluded, with the exception of the top face of the detector where the veto region extends to 2.4 radiation lengths (24 cell widths) to further reject cosmic activity from above. At the near detector, all activity within 8 cell widths of any detector wall is excluded.

Next, we apply a cut on the total ADC of all hits in the cluster. The upper and lower cut boundaries are those that maximize the figure of merit  $N_{\text{sig}} / \sqrt{N_{\text{bkg}}}$ .

We reject any clusters which occur within 250 ns of another cluster. Supernova-related activity should not be correlated on this timescale, but cosmic-induced electromagnetic showers certainly can be. This cut helps to filter out the last of the clusters that are easily associated with cosmic activity. Finally, we exclude any clusters that have more than four hits.

After applying these selection cuts, most activity correlated with cosmic-ray interactions have been removed and the rate of supernova candidate clusters over time roughly follows a Poisson distribution (Fig. 6.6). The signal-to-noise ratios for a simulated  $9.6 M_{\odot}$  supernova at 10 kpc are reduced to 1:29 at the far detector and 2.5:1 at the near detector.

## Triggering infrastructure

The NOvA data-driven triggering system examines data in 5 ms chunks (called milliblocks) stored in the circular buffer. For most data-driven trigger (DDT) algorithms, trigger decisions are made for each individual milliblock. A DDT for identifying a supernova must operate beyond the standard DDT paradigm since the characteristic timescale of a supernova (several seconds) far exceeds the duration of a single milliblock. A large number of milliblocks must be considered simultaneously in order to “see” the supernova signal and issue the appropriate triggers.

Multiple instances of the supernova DDT process run on all buffer nodes and process milliblocks in parallel. Each one counts the number of candidate supernova clusters in its time window and reports back to the central process on the global trigger node of the DAQ. This central process is responsible for aggregating the candidate counts from each of the many supernova DDT processes and constructing a time series of the counts  $\vec{n} = \{t_i, n_i\}$ .

This time series is used to compare a background-only hypothesis ( $H_0 = B$ ) with a signal-plus-background hypothesis ( $H_1 = B + S$ ) given the data  $\vec{n}$ . The background level  $B$  is assumed to be constant across all time bins and is estimated by the triggering system every 10 min. The signal  $S$  is expected to vary according to a template that mimics the shape of a supernova neutrino burst. A log-likelihood ratio

$$l(\vec{n}) \equiv \log \prod_i \frac{P(n_i|H_1)}{P(n_i|H_0)} = \sum_i n_i \cdot \log(1 + S_i/B) \quad (6.2)$$

is used as a test statistic. The p-value is  $p(\vec{n}) = P(l > l(\vec{n})|H_0)$  which can be converted into a z-score by  $z(\vec{n}) = \text{erf}^{-1}(1 - 2p(\vec{n}))$ .

A supernova trigger is issued when the significance exceeds a predetermined threshold, which we chose to be  $z_0 = 5.645 \sigma$ . This choice of threshold will produce on average one false trigger per week, which is the target false-positive rate for contributing to SNEWS. Figure 6.7 shows a time series of supernova cluster candidates with a simulated 5 kpc supernova and the associated significance time series for three signal template choices. Each choice produces a similar significance time profile, so the trigger efficiency is not very sensitive to the choice of template, as long as the template mimics the very broad characteristics of a supernova event.

## Sensitivity

With the trigger significance threshold set to the value stated above, the sensitivity of the supernova DDT as a function of distance can be determined. We convolve the efficiency of the trigger with the supernova candidate density shown in Fig. 3.1 and find that for a supernova of a  $9.6 M_\odot$  progenitor, we are sensitive to 22.6% of supernova candidates. For the supernova of a  $27 M_\odot$ , we are sensitive to 49.2% (Fig. 6.8).

## 6.5 Long readout monitoring and data quality

Streaming 45 s of continuous data from the DAQ buffers to disk puts enormous strain on the DAQ system. To prevent the system from crashing, we limit the speed at which those data are written to disk. The only clock we are up against is the depth of the buffer. If we read out the data too slowly, it may reach the end of the buffer and disappear before we can record it. If we read out too quickly, the DAQ may crash.




Figure 6.7: Time series of supernova candidate clusters for a simulated  $9.6 M_{\odot}$  supernova at 5 kpc (top) and associated significance time series for an expected uniform signal (third from bottom),  $9.6 M_{\odot}$  signal (second from bottom), and  $27 M_{\odot}$  signal (bottom). Taken from Ref. [122].




Figure 6.8: NOvA's supernova DDT sensitivity versus distance for  $9.6 M_{\odot}$  and  $27 M_{\odot}$  progenitors convolved with the distribution of probable supernova candidates as shown in Fig. 3.1. Taken from Ref. [122].

There are two common failure modes that we see when something goes wrong during this process:

1. the readout takes too long and data begins to fall off of the buffer. When this happens, the triggers still appear in the raw data file on disk, but there is no data associated with it. It looks as if the detector registered absolutely no activity during that time; and
2. the readout speed overwhelms the global trigger process and individual triggers never get issued. In this case, those individual triggers are entirely missing, as if they were never meant to be readout at all.

It is important to identify these failure modes when they happen. Detector noise and operating conditions can change over time, meaning that a supernova trigger that performed well in the past may behave unstably in the future. Continuous monitoring is important for maintaining a functional trigger and ensuring that the NOvA DAQ is prepared to record the full detectable signal when the next galactic core-collapse supernova occurs.

I designed a monitoring tool for this purpose which I call the Long Readout Monitoring System (LRMS). A long readout is any readout for which multiple triggers are associated with a single triggering event. In this case, we are considering the internal and external supernova triggers, but this system can be extended to handle any number of long readout triggers.

The LRMS watches for new files associated with long readout trigger streams to appear on the raw data disk, and when a new file is detected, it is processed by the trigger analyzer (described below). Once all files associated with a given long readout have been processed, a final determination about the overall quality of the readout is made.

This system is composed of five components:

1. an analyzer to examine individual raw data files,
2. a scheduler to monitor the data disk and coordinate the execution of the processing script,
3. a summarizer script to compute the final data quality metrics once all files from a long readout have been processed, and
4. a web dashboard to help experts review the data quality of many recent long readouts at once.
5. A SQLite3 database for each component to store and share information.

These components are described in detail below, but first let me explain how a long readout is structured and what we are looking for to assess their integrity. Data readouts occur when triggering decision is made either by an external event (such as an accelerator signal or a SNEWS alert) or by a data-driven trigger process. When a trigger decision is issued, the global trigger requests a single contiguous chunk of data spanning a time less than or equal to 5 ms be written to disk from the data buffer. This chunk of data is called a “trigger”, not to be confused with a trigger process or a trigger decision. Each trigger in a run is given a unique trigger number for accounting purposes.

In the case of a long readout, a single 5 ms trigger is not enough to cover the full duration of the event in question (in our case, a supernova), so multiple contiguous triggers must be requested. In addition to the trigger number, each trigger is also assigned a parent trigger number. For simple readouts, the trigger number and the parent trigger number are identical. For long readouts, the parent trigger number for every trigger is equal to the first trigger number in the readout. This numbering scheme

enables the bookkeeping necessary to assess not only the data quality of a single trigger, but that of a group of triggers that are associated with one another. We can also look for instances where individual triggers are missing.

To assess the integrity of a long readout, the LRMS must do the following:

- identify errors or abnormalities for each individual trigger,
- count the number and types of all trigger-level errors or abnormalities associated with a long readout, and
- identify instances where individual triggers are missing from a long readout

## LRMS components

### Database

The database uses SQLite3, a serverless SQL database that is stored as a single file on disk. The data involved with this system are structured and are connected by simple relationships, making a relational database a natural choice.

The database is composed of several tables:

- `triggers` — all relevant information about an individual trigger, including the trigger number, parent trigger number, run number, subrun number, trigger start time, trigger duration, and any error modes present.
- `parent_triggers` — primarily derived information about long readouts, including the parent trigger number, number of triggers in the readout, number of missing triggers, start time of the first trigger, run number, number of subruns spanned by triggers, and processing status.
- `files` — information about files on the raw data disk that have already been analyzed by the system, including filename, trigger type, detector, run number, and subrun number.
- `configuration` — LRMS configuration values updated each time the system initializes.
- `log` — timestamped log messages of LRMS activity for troubleshooting.

### Trigger analyzer

The trigger analyzer is a module written for the art event-processing framework<sup>1</sup> [126]. Its purpose is to examine the individual triggers present in each data file (i.e. each subrun for particular detector and trigger stream) and store relevant information about them in the `triggers` table of the database.

This is where trigger-level data quality checks are performed. There are two potential error modes that the analyzer checks for. The first is whether the `isEventIncomplete` flag is set in the raw event header. When this flag is set to `True`, it indicates that at least one buffer node did not report to the Global Trigger when the trigger was requested. Only one buffer node contains the requested data, and it's unlikely that the one which did not report is the same one holding that data. Therefore, this rarely leads to data loss, but it is still worth tracking; a sudden uptick in these could indicate a more serious problem.

The trigger analyzer also counts the number of microslices. A microslice is a 50  $\mu$ s block of data from a single DCM. A 5 ms trigger should contain at least 16,800 microslices at the far detector and

---

<sup>1</sup>art is the framework that the NOvA collaboration uses to process its data files.

Table 6.2: Long readout individual trigger error register. Bits 2–7 are reserved for future error modes.

Bit	Name	Description
0	kNHitsOutOfRange	The number of microslices in the trigger is out of range.
1	kIncompleteEvent	The <code>isEventIncomplete</code> flag is set in the raw event header.
2	RESERVED	–
3	RESERVED	–
4	RESERVED	–
5	RESERVED	–
6	RESERVED	–
7	RESERVED	–

1,400 at the near detector. If there are fewer microslices than this present in the trigger, then it is likely that some of those data fell off the circular buffer before we managed to read it out to disk. Identifying and addressing this failure mode is critically important for ensuring that we do not lose any data during a real supernova event.

The presence or absence of these two failure modes are stored in an 8-bit error register (Tab. 6.2) which is stored in the database with the rest of the trigger metadata. Only two of the bits are currently used, the remaining six are free to be assigned to new data quality checks if needed.

### Long readout summarizer

The purpose of the summarizer script is to collate information about triggers processed by the trigger analyzer from across multiple files to characterize each long readout. Before this can happen, the summarizer checks any new long readouts for completeness. A long readout is considered complete once all files that contain triggers from the readout have been processed by the trigger analyzer program. To verify completeness, the summarizer checks for the two boundaries where a long readout begins and ends. If those boundaries occur in the same file, then the long readout is considered complete. If the boundaries occur in separate files, then the summarizer must verify that all intermediate files have also been processed. Figure 6.9 depicts several examples of this.

Once a new long readout is considered complete, the summarizer assesses its integrity. First, it looks for missing triggers. A trigger  $i$  has associated with it a start time  $t_i$  and a duration  $\delta t_i$ . From these two pieces of information, the end time of each trigger can be computed and compared against the start time of the following trigger. There should be no gap. The summerizer orders the triggers by time and computes the quantity  $\Delta = t_{i+1} - (t_i + \delta t_i)$ . If this quantity is greater than zero, there are triggers missing. The number of missing triggers can be deduced because all triggers in the long readout have the same  $\delta t$ :  $N_{\text{missing}} = \Delta / \delta t_i$ .

Finally, the summarizer indicates in the `parent_triggers` table that the long readout is “finalized”. This simply means that the summarizer has finished analyzing the long readout and has successfully stored its information in the database. This flag prevents the summarizer from attempting to summarize this long readout again in the future. Once a long readout is finalized, all error-free triggers associated with it are dropped from the database to save space and to reduce performance issues. They are no longer needed once the summary information has been computed. Triggers with errors are kept in the database as diagnostic information for experts.

Table 6.3 shows an example list of triggers for a long readout at the far detector with some instances




Figure 6.9: Examples of complete and incomplete long readouts. Each square represents an individual file and the shaded areas represent a series of triggers associated with a long readout. In scenario (a), the terminating boundary the long readout has been identified, but the initiating boundary is in a file that has yet to be processed. Scenario (b) is the same as (a), only reversed. Scenario (c) depicts a case where both boundaries of the long readout have been found, but an intermediate file in between has not yet been processed, and therefore the long readout is considered incomplete. Scenarios (d) and (e) show long readouts that are considered complete. In (d) the long readout is spread out across multiple files whereas in (e) it is entirely contained within a single file.

of the failure modes that the trigger analyzer and long readout summarizer look for.

## Scheduler

The scheduler coordinates the execution of the trigger analyzer on new files on the raw data disk. A separate instance for each detector is started every 15 minutes if it is not already running. First, the scheduler compiles a list of files on the data disk whose filenames match a pattern as specified in the configuration. This pattern forces the system to only consider files that are associated with long readout trigger streams. It also allows the user to extend this to any arbitrary trigger stream types that may be added in the future. The list of files on the disk is compared against the list of files that have already been processed by the system, and the difference is taken to generate a list of new files that have not yet been analyzed. Of these new, unprocessed files, only those that have a file modification time greater than 120 s from the current time are queued for analysis. This protects against attempting to analyze a new file that is currently being written to by the DAQ.

The scheduler then executes the trigger analyzer program on each new file one-by-one, starting with the most recent files and working backwards in time. After all new files have been processed by the trigger analyzer, the summarizer script is run to summarize any new and complete long readouts. Finally, the updated database is copied to a location that is accessible by the web interface.

## Web dashboard

A web dashboard which has access to the shared database presents information about recent long readouts in a way that is intended to provide supernova trigger experts with a bird's-eye view of stability over time, and also allows experts to drill down into specific long readouts.

Table 6.3: Example list of triggers for a run at the far detector with a long readout contained within it. This mock readout presents the pathologies that indicate the DAQ is struggling to continuously stream supernova data from the detector to storage. SNEWS HB triggers are minutely heartbeats from SNEWS and DDSN is the NOvA data-driven supernova trigger.

	Run / subrun	Parent trigger	Trigger	Type	Length ( $\mu$ s)	Number of microslices
	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
	13405 / 02	178	178	SNEWS HB	50	168
First long	13405 / 02	224	224	SNEWS HB	50	168
readout trigger $\rightarrow$	13405 / 02	236	236	DDSN	5,000	16,800
	13405 / 02	236	237	DDSN	5,000	16,800
Missing microslices $\rightarrow$	13405 / 02	236	238	DDSN	5,000	1,530
	13405 / 02	236	239	DDSN	5,000	16,800
Missing trigger $\rightarrow$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
Missing trigger $\rightarrow$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
	13405 / 02	236	242	DDSN	5,000	16,800
	13405 / 02	236	243	DDSN	5,000	16,800
Subrun rollover $\rightarrow$	13405 / 03	236	244	DDSN	5,000	16,800
(new file)	13405 / 03	236	245	DDSN	5,000	16,800
	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Final long	13405 / 19	236	9035	DDSN	5,000	16,800
readout trigger $\rightarrow$	13405 / 19	9062	9062	SNEWS HB	50	168
	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$



Parent Trigger ID	Start time	Type	Run/Subruns	Triggers	Duration (ms)	Trigger Errors	Missing Triggers
#80527008c8a	2020-03-26 08:29:02 CDT	SNEWS BEAT SLOW	35978/56 – 35979/32	9000	45000	0	0
#198e008c7a	2020-03-24 11:02:45 CDT	SN X-TRIG	35962/19 – 35962/59	9000	45000	0	0
#57c5008c70	2020-03-23 17:14:05 CDT	SN X-TRIG	35952/29	3	15	3	0
#56f0008c70	2020-03-23 17:14:04 CDT	SN X-TRIG	35952/29	173	865	173	0
#338b008c70	2020-03-23 17:13:19 CDT	SN X-TRIG	35952/26 – 35952/29	9000	45000	9000	0
#104a008c70	2020-03-23 17:12:34 CDT	SN X-TRIG	35952/14 – 35952/26	9000	45000	7014	0
#95664008c70	2020-03-23 17:12:05 CDT	SN X-TRIG	35952/14	4	20	0	0
#93306008c6f	2020-03-23 17:11:20 CDT	SN X-TRIG	35951/37 – 35952/14	9000	45000	2	0
#2ae3a008c1b	2020-03-14 08:29:02 CDT	SNEWS BEAT SLOW	35867/17 – 35867/57	9000	45000	0	0
#5e85008c01	2020-03-11 08:29:02 CDT	SNEWS BEAT SLOW	35841/02 – 35841/38	9000	45000	9000	0

Figure 6.10: Web dashboard home page for the far detector. Each row in the table represents a long readout.

Start time	Run/Subrun	ART Event No.	Trigger Number	Length (ms)	Microslices	Errors	Filename
2020-03-23 17:14:05.3700 CDT	35952 / 29	11517	22469	5	0	• Microslices are missing	faridet_r00035952_s29_DDsn.raw
2020-03-23 17:14:05.3750 CDT	35952 / 29	11520	22470	5	0	• Microslices are missing	faridet_r00035952_s29_DDsn.raw
2020-03-23 17:14:05.3800 CDT	35952 / 29	11523	22471	5	0	• Microslices are missing	faridet_r00035952_s29_DDsn.raw

Figure 6.11: Web dashboard page for a long readout at far detector that had some trigger-level errors.

The home page displays a list of the most recent readouts in a paginated table (Fig. 6.10). Basic information about the long readout is presented along with details about any errors that were encountered. The parent trigger ID is color-coded based on whether it had any errors or not. Clicking on the parent trigger ID takes the user to a page with more detailed information about that particular long readout. The search box in the upper right corner of the table allows the user to filter the results in the table in real-time. This is useful for when an expert wishes to examine the readouts from a particular trigger stream or is searching for a readout from a particular run or on a specific date.

The detailed long readout page (Fig. 6.11) provides more information about any one readout and can be useful for identifying specific triggers that experienced an error. Experts can use this information to find the exact file that contained problematic triggers and investigated further.

## Summary

The NOvA DAQ is capable of reading out 45 s of continuous data to disk to capture the full detectable period of activity from a galactic supernova event. These long readouts are initiated either from an external source, like SNEWS or LIGO, or from our data-driven supernova trigger process. The supernova DDT performs basic hit rejection, clustering, timing correction, and likelihood tests in real time to determine whether a triggering decision should be issued. When one detector's supernova DDT process issues a trigger decision, a cross-trigger is issued at the other detector so that both detectors will read out simultaneously.

These long readouts put considerable strain on the DAQ. The Long Readout Monitoring System I developed is designed to continuously monitor that data for signs that the DAQ is experiencing instability and dropping data from the readout. This monitoring is an important for ensuring that the NOvA detectors are ready to capture every moment of activity if and when the burst of supernova neutrinos wash over them.

One useful by-product of the daily SNEWS trigger tests is the collection of minimum-bias data that can be used to characterize expected backgrounds for both the development of the supernova DDT and that of an offline analysis. This chapter focused on these triggers. Next, I will turn my attention to the question of what we can do with supernova data once it has safely made it to disk.

## CHAPTER 7

# Background Rejection and Hit Clustering

Identifying hits from supernova neutrino interactions in the NOvA detectors is a challenge. The typical supernova neutrino will produce only a few hits due to its relatively low energy and thresholding effects of the detectors. Meanwhile, each cosmic-ray muon or beam neutrino will typically produce tens or hundreds of hits. The NOvA detectors were simply not designed for reconstructing events at supernova energies, and as a result, their topological signatures more closely resemble those of detector noise than those of neutrino activity.

For a supernova neutrino analysis at NOvA, nearly everything is a background. To stand a chance at teasing out the supernova signal, I must first remove as much non-supernova activity from the data as possible and then attempt to select the activity that most closely resembles the expected signal. In this chapter, I will describe the background rejection procedure I developed for this analysis and how the hits that remain after the procedure are grouped into clusters of candidate supernova neutrino interactions.

### 7.1 Calibration and initial reconstruction

The traditional reconstruction tools that the NOvA Collaboration has developed to identify activity of interest to other analyses are largely useless for identifying supernova neutrino activity. However, that makes them particularly useful for identifying detector activity that must be rejected before a search for the supernova neutrino signal can proceed. This section describes those pre-existing tools and algorithms that I use to reject activity clearly not associated with supernova neutrinos.

#### Hit calibration

A “raw” hit is characterized by its ADC and TDC values, but these quantities lack a clear physical meaning. The ADC can be rescaled to estimate the number of photoelectrons (PE) produced by the APD, but this too is an uncalibrated unit. To express the energy that is deposited into a cell in more meaningful units, these hits must be calibrated. There are several corrections that must be applied to compensate for phenomena such as detector thresholding, detector self-shielding, light attenuation in the fiber, and an overall energy scale.

Thresholding plays a role for lower-energy depositions that occur far away from the readout and whose light signal falls below detection threshold before arriving at the APD. Additionally, the top of the detector can act as shielding for hits deposited at the bottom of the detector. A scaling factor is applied based on the distance of a hit to the readout end of a cell to correct for these effects.

Another effect that must be corrected for is the attenuation of scintillation light as it travels through the wavelength-shifting fiber. Because of this, a hit that occurs closer to the readout end of a cell will have a higher ADC than an identical hit on the opposite end. This part of the calibration involves selecting through-going cosmic-ray muons which produce hits in three contiguous cells in a single plane, called “tricells”. For each tricell, the path-length traversed by the muon through the center cell is computed based on the direction of the track. The result is a collection of uncalibrated photoelectrons per path-length as a function of distance from the readout ( $W$ ). The attenuation length of each cell is determined by fitting the function

$$y = C + A \left( \exp \left( \frac{W}{X} \right) + \exp \left( -\frac{L + W}{X} \right) \right) \quad (7.1)$$

to the PE/cm profile, where  $y$  is the response,  $L$  is the cell length,  $X$  is the attenuation length of the cell, and  $A$  and  $C$  are additional free parameters [127]. These cell-by-cell attenuation fits allow us to correct a PE value for a given hit on a given cell. This ensures that a particular hit will have the same calibrated PE value regardless of where it was deposited in the detector.

An additional step is necessary to convert these calibrated PE values into energy. This step of the calibration relies on looking at the calibrated PE values for stopping muons, whose energy depositions are well-described by the Bethe-Bloch equation [35]. The calorimetric energy scale is set by performing this measurement for both real data and simulation, and taking the ratio:

$$\text{calorimetric energy scale} = \frac{(\text{MeV} / \text{cm})_{\text{sim}}}{(\text{PE}_{\text{calib}} / \text{cm})_{\text{data}}}. \quad (7.2)$$

This allows for the conversion of calibrated PE values into units of energy.

Throughout the rest of this thesis, the term “hit” will refer to a calibrated hit, unless otherwise specified.

## Clustering

After the raw hits are calibrated, the next step in the chain of nearly every NOvA reconstruction workflow is clustering hits that occur close together in time and space. The goal is for each cluster to contain the hits that were produced by a single charged particle in the detector volume.

NOvA has a clustering procedure called Slicer4D [128] that is based on the popular DBSCAN algorithm [129] developed in the 1990s. An advantage that DBSCAN has over other common clustering algorithms is that it does not require apriori knowledge of the shape, size, or distribution of clusters. DBSCAN is a density-based algorithm, which identifies clusters based on the proximity of points in a parameter space. It can also recognize that some points in the space may not be associated with any clusters and are instead treated as noise. These advantages are useful for NOvA data because we look for localized clusters of correlated activity on top of a background of noise hits. We do not need to specify the shape of these clusters nor how many we expect to have ahead of time.

The Slicer4D algorithm examines hits as points in a four-dimensional time-position parameter space. The density for each point is determined by counting the number of nearby neighboring hits in the space. Pairwise combinations of hits are each given a neighbor score  $\epsilon$  defined as

$$\epsilon = \left( \frac{\Delta T - \Delta \vec{r}/c}{T_{\text{res}}} \right)^2 + \left( \frac{\Delta Z}{D_{\text{pen}}} \right)^2 + \left( \frac{\Delta XY}{D_{\text{pen}}} \right)^2 \quad (7.3)$$

where  $T_{\text{res}}$  is the combined timing resolution of the hits,  $PE$  is the number of photoelectrons from both hits added in quadrature,  $\Delta Z$  is the plane distance between the hits,  $\Delta XY$  is the cell distance between the hits if they are in the same view and zero otherwise,  $D_{\text{pen}}$  is a distance penalty, and  $\Delta \vec{r}$  is the spatial distance between the two points, and is either two-dimensional or one-dimensional, depending on whether the hits are in the same or different views. Pairs of hits that satisfy causality and are close together in terms of cell and plane positions have a lower neighbor score than those that are not. Two points are considered neighbors if  $\epsilon \leq 5$  at the near detector and  $\epsilon \leq 2$  at the far detector. The specific values of these parameters were chosen for each detector by maximizing a figure of merit in a grid search, and can be found in Ref. [128].

Each point is classified as either a core point or a border point. A core point is any point which has 4 or more neighboring points, while border points have fewer. Each cluster is initially seeded by a single core point and continues to grow as more core points are identified and added to the cluster. Border points are also added to the clusters so long as those border points are neighbors with a core point. Once the cluster growth has terminated on a number of border points, the cluster is complete. This iterative process continues until all points have been designated as a core or border point in some cluster, or as a noise point with too few neighboring points to be considered for membership in any cluster. The noise points are typically ignored by the DBSCAN algorithm, but Slicer4D places them all into a single cluster that we call the “noise cluster”. In this way, every hit is either a member of one of the many “physics clusters” or a member of the single noise cluster.

In most cases, all supernova neutrino hits will end up in the noise cluster. This algorithm is therefore not effective at building clusters of individual supernova neutrino events, but it is useful in other ways. For one, the track-finding algorithm that I will describe next requires a vector of clusters as input, so this clustering is prerequisite. Also, clustering the many hits that are certainly not associated with supernova neutrinos (i.e. any hits in a physics cluster) is an effective way to directly reject activity that is not relevant to this analysis.

## Track-finding

The NOvA far detector sits on the Earth’s surface with only six inches of barite overburden, and is subject to a cosmic-ray hit-rate of about 150 kHz. These particles tend to produce long and straight tracks as they traverse the detector volume. NOvA needs to be able to identify these tracks for both the reconstruction of  $\nu_\mu$  events in the oscillation analyses and to identify cosmic ray activity. For this supernova analysis, identifying cosmic-ray muons and the path they take as they traverse the detectors is essential for being able to veto the hits they and their byproducts produce.

NOvA has developed a track-finding algorithm for cosmic-ray muons called WindowTrack. Given a list of non-noise clusters, WindowTrack attempts to fit a series of straight lines to the hits in that cluster. Cosmic-ray muons are relativistic and they traverse the detector quickly and with trajectories that resemble straight lines. Of course, Coulomb interactions between the muons and atoms in the detector can affect their trajectory, so a single straight line is not always sufficient to describe the muon’s path through the detector. Splitting the trajectory into segments and fitting a straight line to each segment allows for a more accurate reconstruction in situations where the muon undergoes multiple scatterings.

A cluster which spans  $n$  planes is split up into  $n - 1$  windows, and a straight line is fit to the points

in each window. For a track to be formed, a cluster must have hits that span at least four planes and at least half of the planes in that cluster must contain hits. In each window, the algorithm uses a signal-weighted average of hits in a given plane for the fit. Any hits which are considered to be outliers in a given plane are neglected during the fit. These fits are performed separately for each 2D view of the detector. Once the tracks for each view are fit, the algorithm attempts to match the two views together into full 3D tracks. Cosmic rays are expected to enter the detector from above, so an additional criterion requiring a downward direction is enforced by the algorithm. The output of this algorithm is a vector of tracks, each of which is composed of a vector of trajectory points which form the piece-wise linear trajectory of the track.

## 7.2 Background rejection

These clustering and track-finding algorithms are critically useful for identifying neutrino and cosmic-ray muon activity of interest in most NOvA analyses, but they are used in this analysis simply as a means of rejecting activity that is likely not supernova-related; one's signal is another's background.

While cosmic-ray muons are perhaps the most obvious source of background activity, they're not the only source. Detector electronics phenomena such as FEB flasher events produce correlated low-ADC activity, for example. High-energy interactions between cosmic-ray muons and the surrounding rock can produce massive shower events that saturate a large proportion of the detector for a period of time. Some detector channels may be in need of replacement or repair, and can produce abnormally high or low hit rates.

Radiological backgrounds are a nuisance for many particle physics experiments, but not for NOvA. Nuclear decays of trace radioactive materials inside and surrounding the detectors typically produce particles in the keV or single MeV range of energies. Due to thresholding, the detection lower limit at NOvA is near 8 MeV, depending on where along the cell the deposition occurs, and most radiological background activity is therefore not energetic enough to produce hits. For this reason, I make no attempt to identify and reject hits of radiological origin for this analysis.

All of these sources of relevant background activity must be enumerated and addressed before the supernova neutrino signal can shine through. In this section, I will describe the hit-level background rejection procedures that I developed for this analysis.

### Duplicated hits

The NOvA DAQ buffer stores data in 50  $\mu$ s blocks called microsllices. When a trigger is issued, it specifies a start time and a duration. All microsllices that span the time period requested by the trigger are written to disk. Because the data are discretized in time in this way, when a trigger specifies a start time that falls within a given microslice, the entire microslice is read out. For long triggers, a series of individual triggers are read out continuously with no gaps in between. This means that the first microslice of trigger  $N$  and the last microslice of trigger  $N - 1$  will be the same microslice. That is, there will be 50  $\mu$ s of overlapping data among those two adjacent triggers since both of them had a claim to a portion of that microslice (Fig. 7.1).

These duplicated data must be removed so that they are not double-counted. The procedure is straight-forward: I simply remove all hits that occur in the first 50  $\mu$ s of each 5 ms trigger. While this is the first step in a series of algorithms meant to reject hits that would be a background to the




Figure 7.1: 2D locations of hits in the final microslice of a trigger (left) and the first microslice in the following trigger (right) for the far detector. The same hits appear in both triggers when the trigger boundary does not fall exactly on a microslice boundary.

supernova analysis, this algorithm differs from the others in that it does not filter out hits because they are suspected to be background; they are removed because they are duplicated. The removal of these hits therefore does not truly remove any hits at all<sup>1</sup>, it simply corrects for a feature of the DAQ architecture that would otherwise lead to an over-counting of some hits.

### Hot and cold channels

Sometimes channels malfunction and are in need of replacement or maintenance. In some cases, these channels will register hits at a particularly high rate; these are considered to be “hot” channels. “Cold” channels, on the other hand, are those that produce hits at an abnormally low rate. In either case, we ought not to trust the activity registered by those channels, and any hits that they produce should be rejected.

I calculate the average channel rates by accumulating per-channel hit statistics for each subrun and divide by the total trigger duration of that subrun. Any hits that come from a channel whose average rate is too high or too low will be rejected.

But what hit rates should be considered acceptable? To determine this, I adjusted cut values on the high and low side independently and computed a figure of merit for each possibility in the case of a simulated supernova overlaid with real background data. While varying either the cold or hot cut variable, I fixed the non-varying parameter. I used  $S/\sqrt{S+B}$  as the figure of merit where  $S$  and  $B$  are the numbers of signal and background hits which survive the cut.

Figure 7.2 depicts the percentages of signal and background hits rejected under these scenarios as well as the figures of merit for both detectors. An appropriate choice of cut value is one which exactly or nearly maximizes the figure of merit without being too close to a region where the figure of merit varies rapidly. The location of any fast fluctuations of the figure of merit may shift under different detector conditions; choosing a cut value too close to any such shoulder may perform well today but cut far too harshly under tomorrow’s detector conditions. For this reason, I chose cut values by eye from the figure of merit plots in Fig. 7.2. The acceptable channel hit rates at the near detector are 3 Hz–40 Hz and those at the far detector are 85 Hz–400 Hz (Fig. 7.3). Any hits originating from channels with subrun-averaged hit rates outside of this range are rejected.

<sup>1</sup>Except in the case of the very first trigger of the readout, in which case hits that occur in the first 50  $\mu$ s of the trigger truly

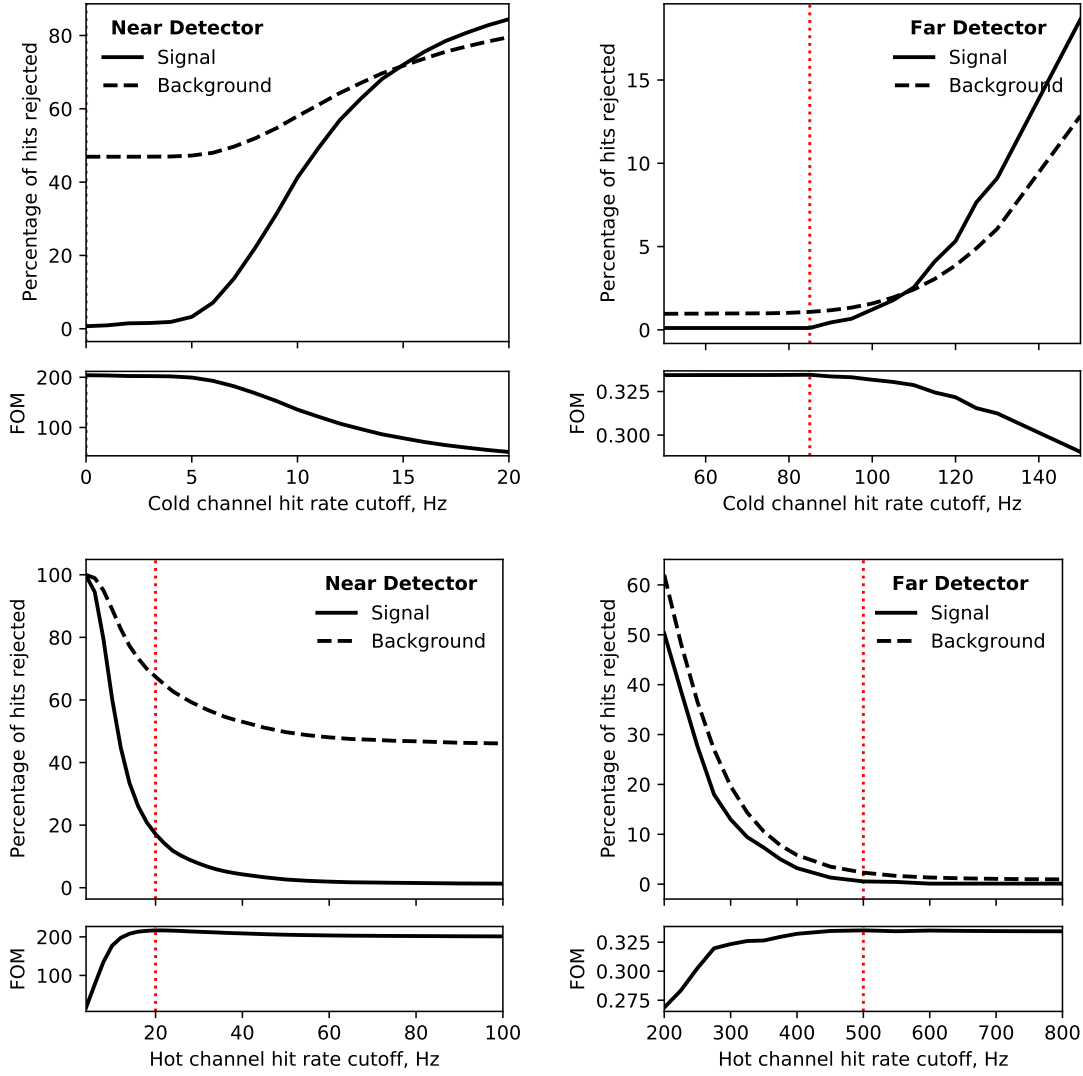


Figure 7.2: Percentage of hits rejected for different choices of cold (top) and hot (bottom) channel cut values for each detector. Figures of merit (FOM) are plotted and defined as  $S/\sqrt{S+B}$  where  $S$  is the number of signal hits surviving the cut and  $B$  is the number of background hits surviving the cut. The red dotted line indicates where the figure of merit is maximized.






Figure 7.3: Typical channel rates for the near detector (left) and far detector (right) with cut values for background rejection indicated by the dashed red lines.

### High- and low-ADC hits

Supernova neutrinos are lower in energy than much of the activity in the NOvA detectors and therefore tend to have lower ADC values. Rejecting hits whose ADC values are lower or higher than expected for supernova neutrino hits is an effective way to reduce backgrounds.

To choose the optimal ADC cut values, I followed the same procedure as the hot and cold channels rejection, where one variable is fixed while the other is adjusted and the cut value is chosen based on where the figure of merit is maximized. Like before, this procedure is performed independently for both detectors because the ADC spectra of the two detectors differ from one another. The fact that the near detector is underground is one reason for the difference, but the size of the detectors also plays a role in shaping the ADC spectra of both signal and background hits. This is because ADC values are proportional to the amount of charge collected by the APD and do not incorporate any corrections for attenuation of the scintillation light as it traverses the cell and fiber. Two hits at opposite ends of a cell produced by otherwise identical energy depositions will register different ADC values. The cells in the near detector are about one-quarter the length of those in the far detector, so attenuation of the scintillation is less pronounced.

Figure 7.4 shows the percentages of signal and background hits rejected under these scenarios as well as the figures of merit for both detectors. In these cases, the cut values that maximized the figure of merit did not sit too close to any steep shoulders, so I simply chose those values. The acceptable ADC ranges for hits at the near detector are 65–1900 and those at the far detector are 70–600 (Fig. 7.5). Any hit whose ADC is outside of these ranges is rejected.

### FEB flashers

In Sec. 5.1, I described how large charge depositions on an APD can cause the baseline voltage for all channels on a given FEB to sag, causing false hits to be triggered as the voltage returns to baseline. These false hits, known as “FEB flashers”, are a source of low-ADC hits that could resemble those produced by supernova neutrino interactions, but luckily they are easily identified by the conditions that produce them.

---

are rejected.

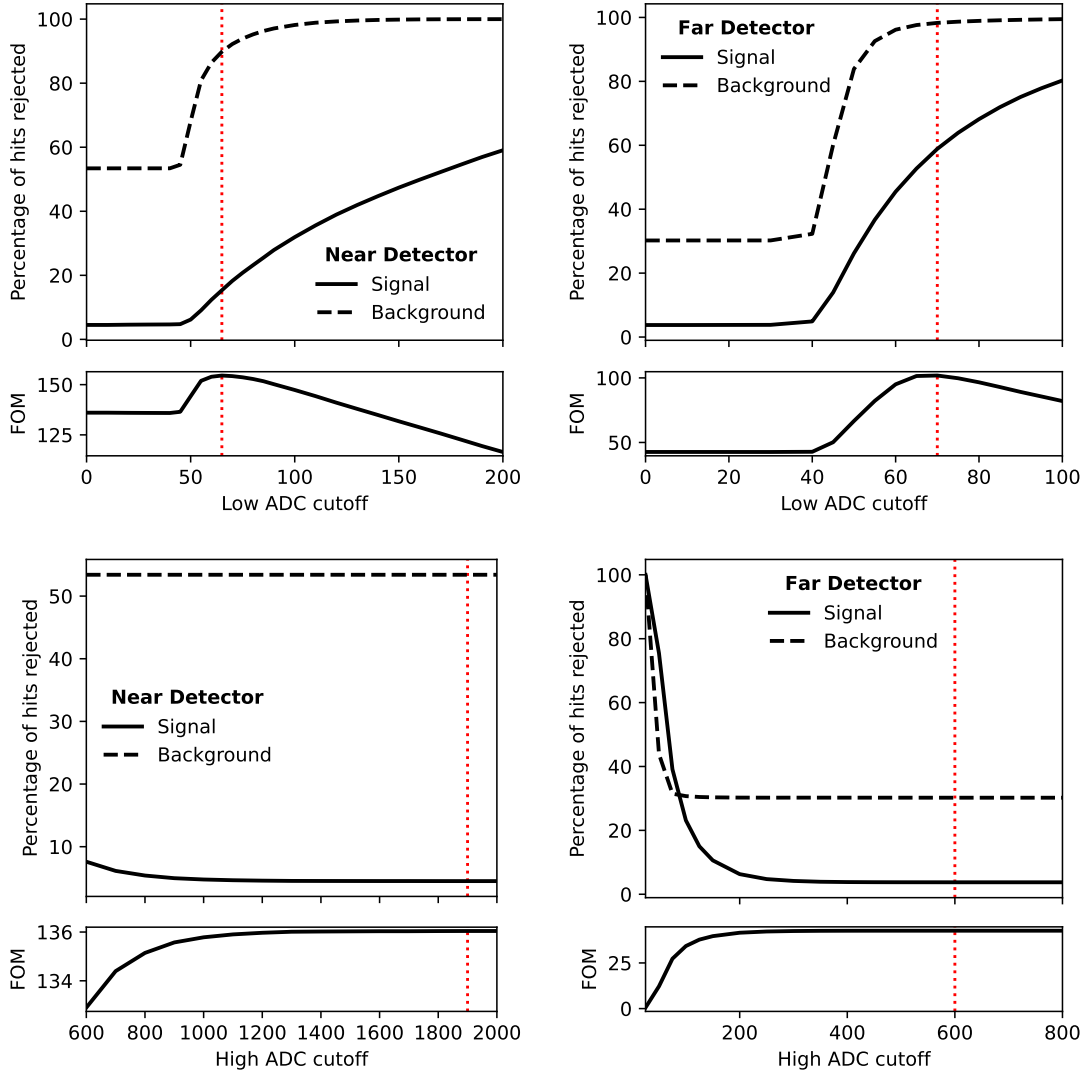


Figure 7.4: Percentage of hits rejected for different choices of low (top) and high (bottom) ADC cut values for each detector. Figures of merit (FOM) are plotted and defined as  $S/\sqrt{S+B}$  where  $S$  is the number of signal hits surviving the cut and  $B$  is the number of background hits surviving the cut. The red dotted line indicates where the figure of merit is maximized.




Figure 7.5: Typical ADC spectra for background and a supernova signal simulated at 1 kpc in the near detector (left) and far detector (right) with cut values for background rejection indicated by the dashed red lines.

The key to identifying FEB flasher events is spotting the high-ADC instigator hits that trigger the phenomenon. Large ADC depositions in a single FEB within a short period of time are marked as instigator hits. In both detectors, this timescale is  $4 \mu\text{s}$ . When this condition occurs, the FEB in question is considered to be “flashing”. Figure 7.6 depicts a typical FEB flasher event and its characteristic features. The hits to reject are those that occur on a flashing FEB within  $30 \mu\text{s}$  of when the instigator hits occurred and with an ADC value less than 50.

The FEB flasher rejection algorithm relies on these four parameters to decide which hits to reject: the total ADC contributed on one FEB from instigator hits, the period of time over which those hits occurred, the upper limit on the ADC of suspected flasher hits, and the time period over which flasher hits should be rejected. Out of these parameters, the total instigator ADC is the one which will have the greatest impact on background rejection, since it determines the threshold at which point combined activity on a given FEB is considered to be a flasher event.

Flashers like the one in Fig. 7.6 are the more obvious cases of this phenomenon, and setting a relatively large instigator ADC cut like 3,500 is enough to identify and reject most of the flasher hits. However, there are more subtle cases where only one or two FEBs are affected, and often those cases are only identified if the cut on instigator ADC is set lower.

To optimize this instigator ADC cut, I again turned to the method of testing different possibilities and comparing the figures of merit for each scenario. The result is shown in Fig. 7.7, and it shows that setting the instigator ADC cut to a very low value like 750 is more effective at rejecting background hits, while leaving the signal hits largely unaffected.

Though it may seem to perform well according to the figure of merit, choosing a low value for this parameter is a bad idea for several reasons. For one, setting the cut value too low ensures that this procedure will reject a lot of activity that is not flasher-related. In many cases, hits will be rejected simply because the FEB recently saw moderate ADC depositions. There is already a rejection procedure in place for both over-active channels and for hits whose ADC is too large or small to have been produced by a supernova neutrino interaction. Another issue is that as the instigator cut value decreases, the computation time required by the algorithm increases sharply.

For these reasons, I chose an instigator ADC cut value of 1,500 in both detectors. This keeps the cut

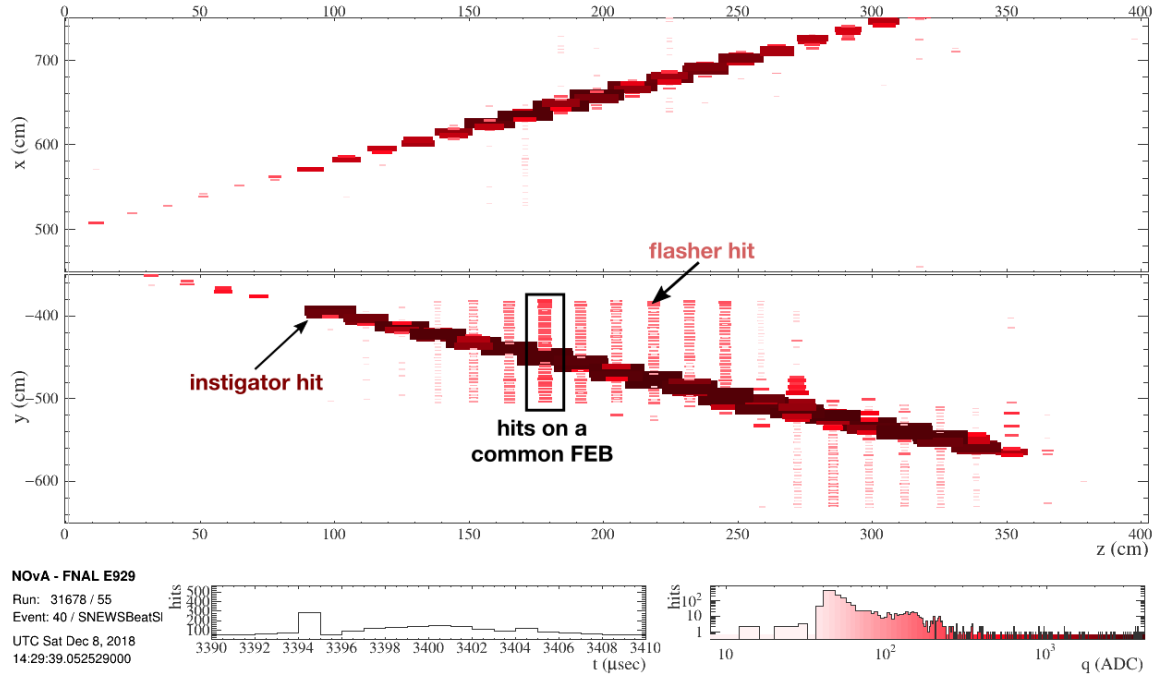


Figure 7.6: An annotated far detector event display of a typical FEB flasher event. Color represents hit ADC.




Figure 7.7: Percentage of hits rejected for different choices of instigator ADC cut values for the far detector. The figures of merit (FOM) is plotted and defined as  $S/\sqrt{S+B}$  where  $S$  is the number of signal hits surviving the cut and  $B$  is the number of background hits surviving the cut. The red dotted line indicates where the figure of merit is maximized.




Figure 7.8: The same event displays as shown in Fig. 7.6 after FEB flasher hits have been removed. Color represents hit ADC.

aggressive in terms of flasher hit rejection while not straying too far into non-flasher rejection territory. This choice also maintains a reasonable algorithm computation time. Figure 7.8 shows an event display of an FEB flasher event after this removal procedure is applied.

### Cosmic-ray muons

Both NOvA detectors are exposed to a continuous flux of cosmic-ray muons. The rate is relatively low at the near detector where the detector sits 100 m underground, but most of the far detector is above ground with only 16 cm of barite as shielding. The rate of cosmic-ray muons traversing the far detector has been observed to be about 150 kHz. The hits produced by these tracks must be removed, but cosmic-ray muons can induce secondary activity such as bremsstrahlung and spallation which must also be removed. Finally, muons which stop within the detector are likely to decay into a pair of neutrinos and an electron (or positron). These Michel electrons have an end-point energy of  $m_\mu/2 = 52.8$  MeV, which is the same energy regime as supernova neutrinos, making them a potentially significant background for this analysis. They too must be removed.

I apply the Window Track algorithm described earlier in this chapter to identify track-like features and the hits that make them up. These hits are rejected. For each 2D view, I reject all hits that occur within 20 cm and 5  $\mu$ s of the track body. Finally, I reject any hits that occur within 40 cm and 30  $\mu$ s of the end of any stopping tracks.

### Burst events

Large energetic showers can occur when cosmic-ray muons interact with the rock just above the detector. In these cases, the detector is flooded with particles and the channels in large regions of the

detector can become saturated. During such an event, the entire detector—or a significant portion of it—is rendered useless. Figure 7.9 shows an event display of a high-energy shower inside the NOvA far detector. The rejection algorithms mentioned so far will reject many of the hits produced in these situations because their ADC values are too high, or because they exhibit FEB flasher instigator behavior, but some artifacts will often remain (Fig. 7.10). To handle this, the best course of action is to simply veto all hits that occur within the period of time during which the detector is reeling from the high-energy burst event.

To identify these events, I look for peaks in the time-distribution of hits for each trigger after all aforementioned rejection algorithms have been applied. I bin the distribution in units of  $1\ \mu\text{s}$ , and use ROOT's [130] one-dimensional peak-finding algorithm (`TSpectrum::Search`) to identify any instances where a gaussian can be fit with a width greater than or equal to two bins and a height that rises at least 50% above the local background level.

I also fit a constant to the entire distribution, save for the first two and final two bins, to estimate the background level  $B$ . Out of all the identified peaks, I only keep those that exceed  $5 \times \sqrt{B}$  in height. Each of these selected peaks represent a window of time during which all hits should be vetoed. To determine the duration, I start at the time of the peak and walk bin-by-bin to either side and identify the times at which the distribution crosses below  $B$ . These times represent the beginning and end of the burst activity, but to be sure that most activity is rejected, a buffer equal to 25% of the veto window is added to either end of the veto window. In other words, for a peak which falls below the background level at times  $t_1$  and  $t_2$ , the width will be  $\delta t = t_2 - t_1$  and the veto window will be  $[t_1 - 0.25\ \delta t, t_2 + 0.25\ \delta t]$ . All hits that occur during this time will be rejected. Figure 7.11 shows where the algorithm identifies the peak, background level, and veto window for the time distribution of a real burst event in the far detector and Fig. 7.12 shows the event display after all background rejection algorithms, including this burst rejection, have been applied. This algorithm is only applied to far detector data, because the near detector rarely sees such high-energy activity.

Because an entire block of time is being masked off, the effective live-time of the readout is reduced. This reduction in effective live-time is on average 0.5% of the total readout time, and is written into the data files by algorithm, so that it can be taken into account by any analyses that might be dependent on the live-time.

## Traditional physics clusters

The Slicer4D algorithm described earlier in this chapter clusters hits that occur close together in time and space which then are used as inputs to the track-finding algorithm. The vast majority of hits produced from supernova neutrino interactions end up in the noise cluster, so the final step in the chain of background rejection algorithms is to reject all hits that appear in any physics cluster. This step is used to reject any hits that survived all of the other rejection steps.

The tuned parameters for each of the background rejection algorithms described in this section are in Tab. 7.1.

## Background rejection performance

The background rejection algorithms are applied in the order presented in this section. Most algorithms are passed the full, unfiltered vector of hits, so their performance is independent of the order in which

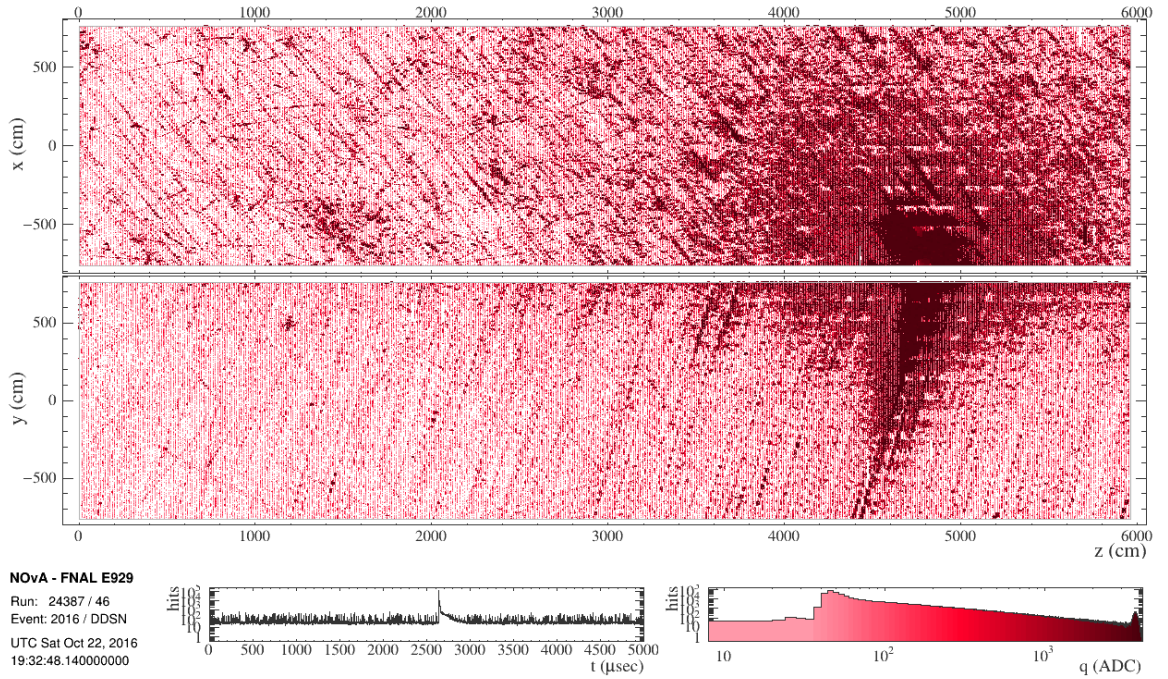


Figure 7.9: Far detector event display of a high-energy shower event. Color represents hit ADC.




Figure 7.10: Far detector event display of a high-energy shower event after background rejection, but without specifically rejecting burst activity. Color represents hit ADC.






Figure 7.11: Peak-finding for burst event rejection. The black curve represents the number of hits per  $1\ \mu\text{s}$  for a high-energy event at the far detector after applying other rejection algorithms. The triangular markers indicate where the algorithm identified prominent peaks. The horizontal teal line is a constant fit to the data. The two vertical red lines represent the boundaries of the veto region, including the 25% buffer on either side.




Figure 7.12: Far detector event display of a high-energy shower event after all background rejection has been applied, including burst event rejection. Color represents hit ADC.



Table 7.1: Tuned background rejection algorithm parameters for both detectors. The algorithms for removing hits from the first microslice of each event and for removing hits associated with physics clusters are not listed here, because they do not take any parameters.

Algorithm				
Parameter	Units	Near Detector	Far Detector	
<i>Cold and hot channels</i>				
Cold rate threshold	Hz	3	85	
Hot rate threshold	Hz	40	400	
<i>Low and high ADC</i>				
Low threshold	ADC	65	70	
High threshold	ADC	1,900	600	
<i>FEB flashers</i>				
Instigator charge	ADC	1,500	1,500	
Instigator time window	TDC	256	256	
Veto charge	ADC	500	500	
Veto time window	TDC	1920	1920	
<i>Track activity</i>				
Muon body veto radius	cm	20	20	
Muon body veto time	ns	5,000	5,000	
Muon end veto radius	cm	40	40	
Muon end veto time	ns	30,000	30,000	
<i>High-energy bursts</i>				
Peak threshold		—	5 $\sigma$	
Veto window padding		—	25%	

they occur. If this weren't the case, they would be less effective. For example, one can imagine a situation where an FEB flasher event is not correctly identified because a previously-applied rejection procedure had already removed a critical number of instigator hits. The one exception to this is the high-energy burst rejection algorithm, which relies on identifying peaks in an otherwise stable hit-time distribution. That algorithm therefore only considers hits which have not otherwise been marked for removal by the other algorithms.

When any background rejection algorithm identifies a hit for removal, that hit is added to a list of hits to be masked. In this way, multiple algorithms may flag the same hit for removal; a hit may be a flasher hit and a low-ADC hit and a physics slice hit all at once. After all algorithms have run, the list of hits that survive is the difference between the unfiltered list of hits and the list of hits to veto.

Figure 7.13 shows the percentage of surviving signal and background hits after each stage of background rejection for both detectors. At the near detector, 90% of the background hits and 15% of the signal hits are removed. The peak signal-to-noise ratio (SNR) for a  $27 M_{\odot}$  supernova at 1 kpc before and after rejection is 1:29.1 and 1:3.3 respectively. At the far detector, 98% of the background hits and 50% of the signal hits are removed. The SNR for a  $27 M_{\odot}$  supernova at 1 kpc before and after rejection is 1:453 and 1:21.3 respectively. Figure 7.14 shows example event displays before and after all background rejection procedures have been applied for both detectors.

For many supernova distances, the background hit rate exceeds that of the signal even after the background rejection has been applied, sometimes by several orders of magnitude at the far detector. This may hardly seem ideal, but even when the signal-to-noise ratio is significantly smaller than unity, two important things have been achieved by applying these background rejection algorithms: (1) the time distribution of surviving hits more closely follows Poisson statistics and remains relatively stable for the duration of the long-readout, and (2) the small percentage of remaining hits form a clean sample to be used as input for an algorithm that will attempt to form supernova-like clusters out of the surviving hits.

### 7.3 Clustering supernova neutrino interactions

Supernova neutrino interactions produce very few hits in the NOvA detectors. Despite the fact that their ADC spectrum resembles that of electronics noise and low-energy cosmic ray byproducts, the few hits that they do produce tend to be clustered close together in time and space, and this correlation is useful for identifying these interactions. Before attempting to design a clustering algorithm, I created the distributions in Fig. 7.15 by clustering Monte Carlo signal hits that survive background rejection together based on their parent neutrino. These clusters represent the ideal case: they each consist of all hits from the simulated particle that produced them and contain no background hit contamination. The properties of these ideal clusters reveal several things: (1) most interactions produce one or two hits and rarely more than 6–8; (2) the total ADC of the hits is generally less than 1,000 at the far detector and less than 1,500 at the near detector; and (3) the hits from a given interaction span a time period of less than 50 ns at the near detector and less than about 100 ns at the far detector, and for both detectors the hits can span several tens of planes or cells, although the majority are confined to the first few.





Figure 7.13: Percentage of surviving signal and background hits after each stage of the background rejection for each detector.

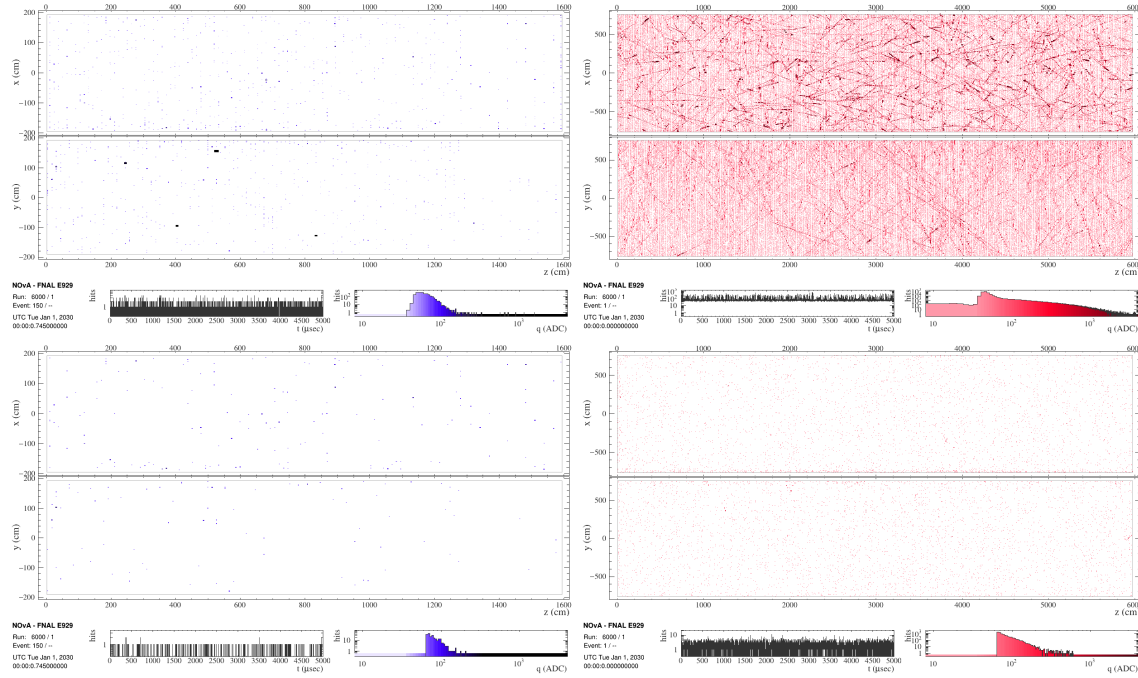


Figure 7.14: Event displays of the near (left) and far (right) detectors before (top) and after (bottom) background rejection algorithms have been applied to a 5 ms trigger. Color represents hit ADC.

## Algorithm

The algorithm I developed to identify supernova-like clusters of hits is called “SNSlicer.” I begin by sorting all calibrated hits which survived the background rejection process by time, from earliest in the trigger to latest. The earliest-occurring non-clustered hit is used as the seed hit for a new cluster. Hits that occur later than the seed hit are considered for inclusion into the cluster.

To be included, a hit must:

1. occur within (ND: 40 ns, FD: 100 ns) of the seed hit;
2. occur within (ND: 40 planes, FD: 15 planes) of the seed hit;
3. occur within (ND: 40 cells, FD: 15 cells) of the seed hit if in the same view;
4. not cause the total ADC of the cluster to exceed (ND: 2,400, FD: 1,000); and
5. not cause the cluster to exceed 8 hits.

Once a candidate hit is encountered which occurs too late after the current cluster seed hit as defined by criterion 1 above, no more hits will be considered for inclusion. At that point, the next cluster will be constructed using the earliest-occurring unclustered hit as the seed. This process continues until there are no remaining hits to be clustered. All clusters then undergo an initial selection. For a cluster to be selected, it must contain two or more hits. In the case of the far detector, the cluster must also have at least one hit in each view (XZ or YZ). Clusters which fail these criteria are discarded; those which satisfy them are passed to the next stage of this analysis.

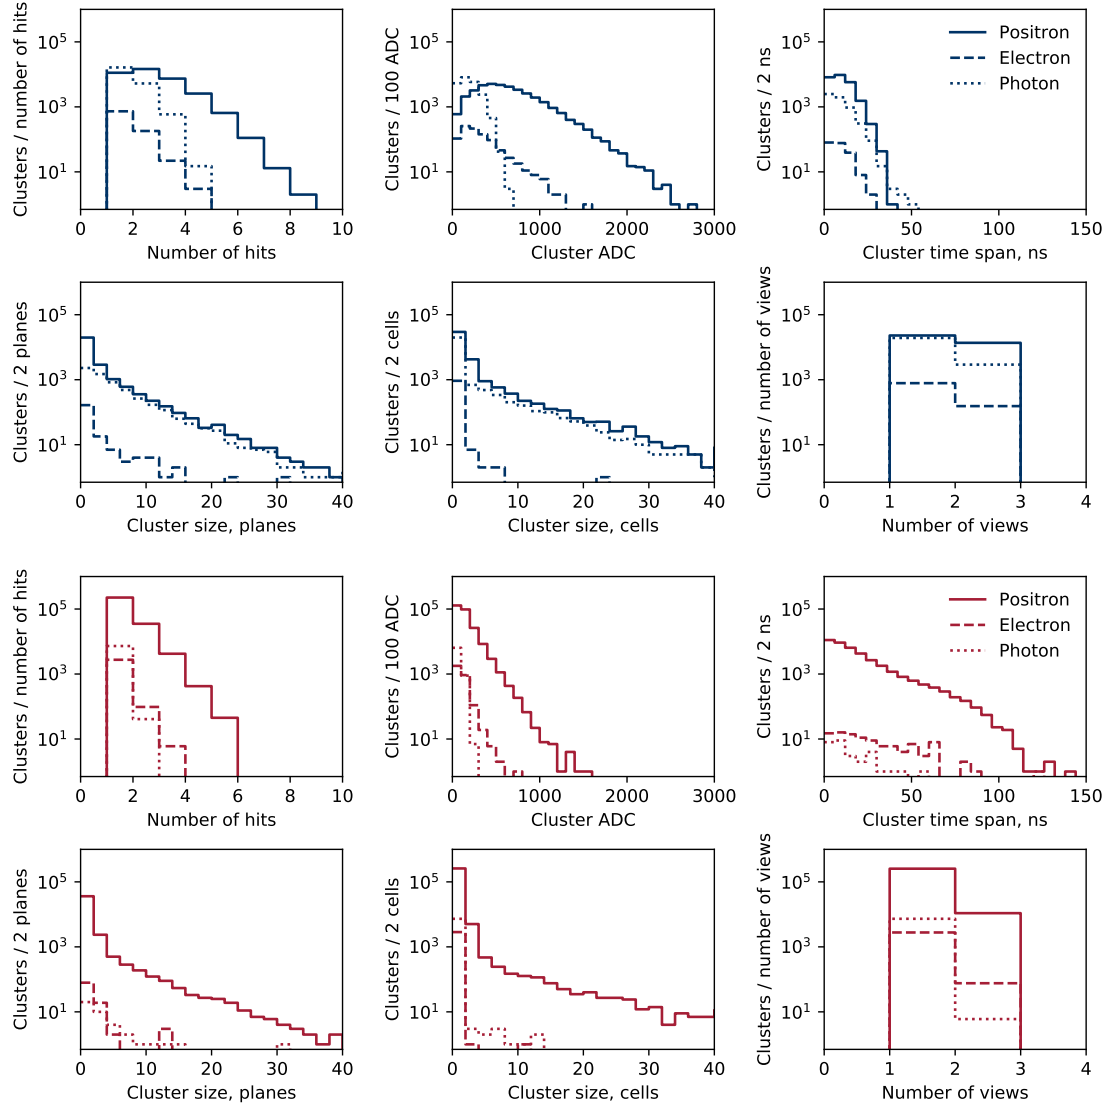


Figure 7.15: Cluster properties if the hits from each simulated supernova neutrino daughter particle were perfectly clustered for the near (top two rows) and far (bottom two rows) detectors: (a) the total number of hits, (b) the summed ADC of all hits, (c) the largest time difference between hits, (d) the largest plane difference between hits, (e) the largest cell difference between hits in either view, and (f) the number of views which contain hits.




Figure 7.16: SNSlicer cluster purity for both detectors.

## SNSlicer Performance

A key performance metric for many particle physics clustering algorithms is efficiency, defined as the fraction of available signal hits that get clustered. This metric is not as useful for the purposes of clustering supernova neutrino hits in the NOvA detectors. For one, 47% of interactions at the near detector and 86% at the far detector produce only one hit. These hits are indistinguishable from background hits in the detector and are therefore not clustered. Aside from one-hit interactions, signal clustering efficiency can be increased by loosening the criteria for cluster formation as described above. However, this also increases background cluster efficiency. For the clustering criteria described above, SNSlicer clustered 26.5% of the available signal hits at the far detector and 73.5% at the near detector.

Another common metric for quantifying the performance of a clustering algorithm is cluster purity. It is expected that any attempt to cluster supernova neutrino hits in the NOvA detectors will produce many clusters that contain background hits. It is difficult to cluster only signal-like activity while leaving background-like activity behind. However, if one hopes at a later time to distinguish signal-dominant clusters from those dominated by background hits—and I do—it is important that the clusters are as pure as they can be. Cluster purity is defined as the fraction of signal hits in a cluster, so a purity of zero means that all hits in the cluster are background. A purity of one means all hits are signal. Any value between zero and one describes a cluster that contains both signal and background hits. Figure 7.16 shows the purity of SNSlicer clusters for each detector. The near detector is dominated by signal-pure clusters while the far detector is dominated by those that are background-pure, and for both detectors the total fraction of clusters that are neither pure signal nor background is less than 1%.

The final performance metric I consider is that of interaction purity. That is, the number of neutrinos which contributed hits to a single cluster. Ideally there should be a one-to-one correspondence between neutrinos and clusters; each hit-producing neutrino interaction should produce one cluster, and each cluster should contain hits from only one parent neutrino. If multiple neutrinos are represented in a single cluster, it could adversely affect the ability to reconstruct an energy spectrum, for example. This becomes more of an issue for supernovae that are closer to Earth as the flux increases

Table 7.2: Clusters produced by the SNSlicer algorithm grouped by the number of neutrinos contributing hits to the cluster.

Neutrinos contributing hits	near detector	far detector
0	2,960 (08.60%)	447,904 (91.58%)
1	31,446 (91.34%)	41,112 (08.41%)
2	19 (00.06%)	58 (00.01%)

according to the inverse square law and neutrino-induced hits become more tightly packed in time and space throughout the detectors. Using Monte Carlo truth information from the simulation, I match each hit to the particle that produced it and the parent neutrino of that particle. Table 7.2 shows a breakdown of how many neutrinos are present in SNSlicer clusters for both detectors. These numbers are based on a supernova distance of 350 pc for the near detector and 500 pc for the far detector. Despite these nearby supernovae exposing the detectors to a large flux, SNSlicer produced highly interaction-pure clusters with more than 99.9% containing hits from either zero (background hits only) or one parent neutrino.

Despite the relative simplicity of the algorithm and its seemingly low efficiencies, SNSlicer performs well for clustering supernova neutrino activity in the NOvA detectors; the clusters it produces are largely pure signal or pure background and are not prone to pile-up effects for nearby supernovae. The SNSlicer signal clusters—clusters that contain one or more signal hits—resemble those from MC truth information. This can be seen in Fig. 7.17 which compares the cluster properties shown in Fig. 7.15 for both SNSlicer and MC truth. The SNSlicer clusters are separated into signal and background distributions to show that the signal portion is in good agreement with the truth clusters.

This chapter began with a problem to be solved: the background hits overwhelm the signal hits. It now ends on a similar note as the background clusters overwhelm the signal clusters, particularly at the far detector. But much has been achieved here; the proverbial haystack is significantly reduced and it has been partitioned into manageable hay bales that will enable a systematic search for needles. That search is the topic of the next chapter, and is motivated by a single question: what distinguishes the typical background cluster from the typical signal cluster?

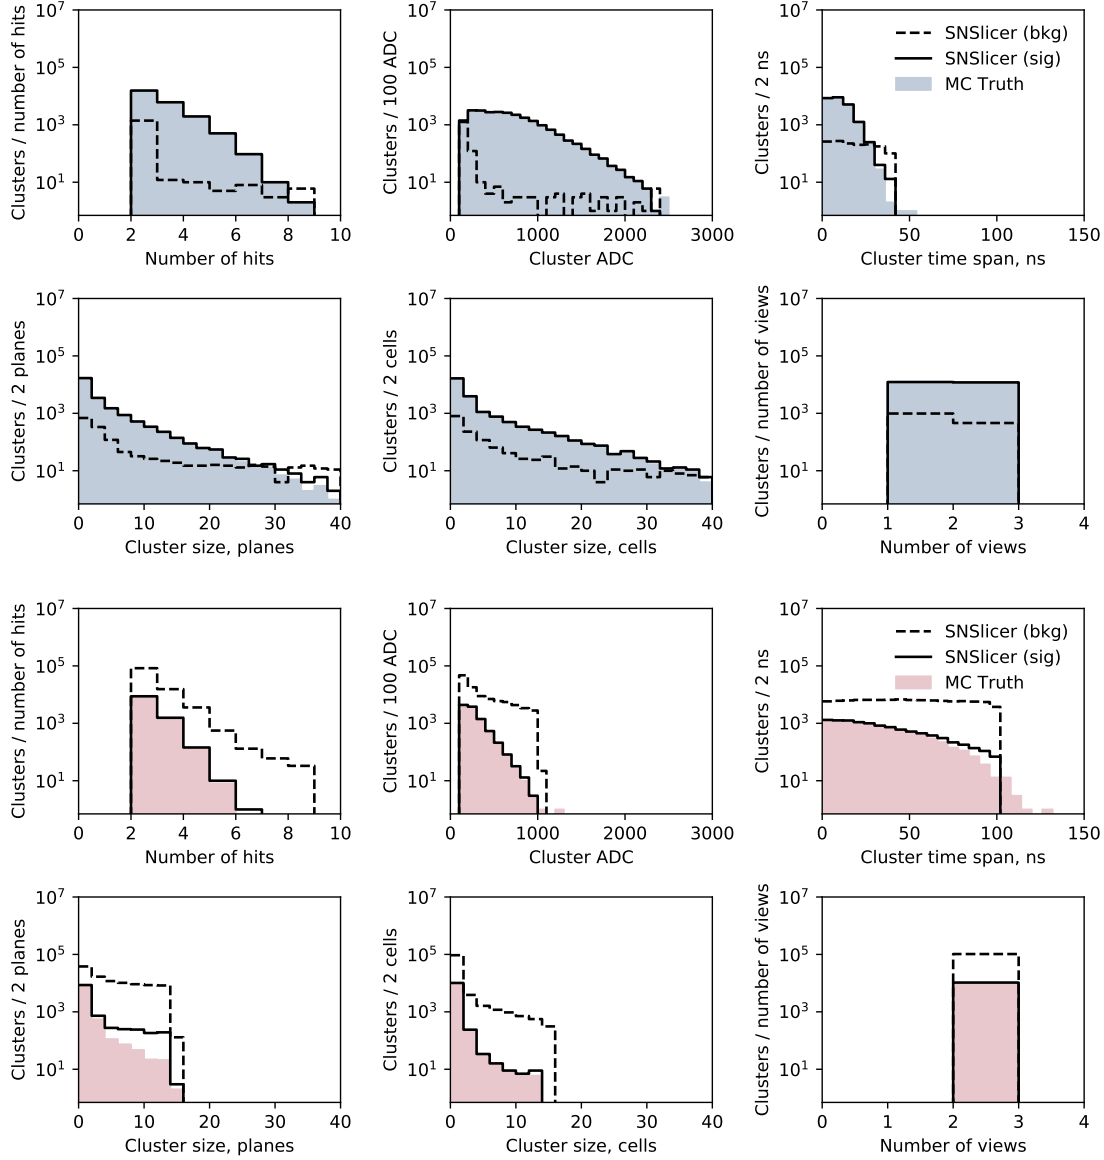


Figure 7.17: Cluster properties for clustering based on MC Truth information versus clustering with the SNSlicer algorithm for the near (top two rows) and far (bottom two rows) detectors: (a) the total number of hits, (b) the summed ADC of all hits, (c) the largest time difference between hits, (d) the largest plane difference between hits, (e) the largest cell difference between hits in either view, and (f) the number of views which contain hits. The MC Truth clusters have been filtered to match the same criteria required by the SNSlicer algorithm.



## CHAPTER 8

# Cluster Selection and Reweighting

After applying the background rejection algorithms described in the previous chapter and clustering the surviving hits, mostly signal-dominant clusters at the near detector and background-dominant clusters at the far detector are left behind. They are well sorted; clusters are composed entirely of signal hits or of background hits, with very few containing a mixture of both. The goal now is to discern one class from the other and produce a refined sub-sample of clusters which will serve as supernova candidate clusters for an analysis.

In addition to that candidate selection, I wish to replicate any analysis under different supernova model scenarios. But running the neutrino simulation, detector simulation, background rejection, and SNSlicer for a large variety of supernova distances and models is computationally expensive, especially for the far detector. This was not feasible with the computational resources I had available to me; another method was needed.

In this chapter, I will present a solution to the computing problem, which involves assigning a weight to every signal cluster and adjusting those weights to produce distributions for different progenitor distances or models. I will also describe the process by which clusters produced by the SNSlicer algorithm are selected as supernova candidate clusters for further analysis.

### 8.1 Cluster reweighting

To avoid the computing problem, I chose an approach that uses knowledge of the various model fluxes to convert signal cluster distributions from one model to another. This method scales the simulated time-energy cluster distribution by a target-to-simulated model flux ratio:

$$n_{\text{target model}}(E_\nu, t_\nu) = \left( \frac{F_{\text{target model}}(E_\nu, t_\nu)}{F_{\text{simulated}}(E_\nu, t_\nu)} \right) \times n_{\text{simulated}}(E_\nu, t_\nu), \quad (8.1)$$

where each flux  $F$  in the ratio is a weighted sum of the partial fluxes for each flavor:

$$F(E_\nu, t_\nu) = \sum_{\alpha \in \{\text{flavors}\}} w_\alpha \times F_\alpha(E_{\nu_\alpha}, t_{\nu_\alpha}). \quad (8.2)$$

The weights  $\{w_\alpha\}$  are determined by the relative rates of interaction for that flavor. Since inverse beta decay is the dominant interaction channel in the NOvA detectors, the  $\bar{\nu}_e$  flavor is the most heavily-weighted. The energy-time distributions  $n(E_\nu, t_\nu)$  correspond to the signal clusters. The energy and time here is not the calorimetric energy and the mean time of the cluster, but rather the true energy of

the neutrino that produced the cluster hits and the time at which that neutrino interacted. This works because the vast majority of signal clusters contains hits from exactly one neutrino.

The one simulation that I have run for each detector is the same one that has been used so far in this thesis: the LS220-s27.0co model from the Garching group [55]. The near detector was simulated at a distance of 350 pc and the far detector at 500 pc. These distances were chosen to be as close to Earth as computationally feasible to maximize the statistics of each simulation. Scaling the distributions to a different distance is simpler than scaling to a different model: instead of computing a separate weight for each energy-time bin, a constant weight of

$$w = \left( \frac{d_{\text{simulated}}}{d} \right)^2 \quad (8.3)$$

is applied to every signal cluster.

To demonstrate this procedure, I generated two supernova neutrino simulations in the near detector at a distance of 2 kpc. One simulation used the LS220-s27.0co model flux while the other used SFHo-z9.6co. The neutrino fluxes for the  $\nu_e$  and  $\bar{\nu}_e$  flavors in both models is shown in Fig. 8.1, as well as the partial flux ratios as described in Eqs. 8.1 and 8.2.

Figure 8.2 shows this method applied to data from the two simulations. The LS220-s27.0co is in black and the SFHo-z9.6co in gray. The red curves are the LS220-s27.0co data after being reweighted to match the SFHo-z9.6co model. The plots in the top row show this procedure applied to the model fluxes themselves, and the true SFHo-z9.6co curve and reweighted curve match each other perfectly, as expected. In the bottom row, this procedure is applied to the data which shows good agreement between the true SFHo-z9.6co data and the reweighted LS220-s27.0co data.

There are advantages and disadvantages to this reweighting method. The obvious advantage is what inspired this method in the first place: computational ease. The full simulation-to-clusters pipeline described in previous chapters takes several days to run for a nearby supernova at the far detector, and that would need to be repeated for every desired model scenario. In comparison, the reweighting procedure described here takes under two seconds to convert the distribution of signal clusters from one model and distance to another, and only requires a one-time investment of computational resources. Another advantage of producing one nearby supernova and scaling it to farther distances is that of increased statistics, especially when a supernova is so far away that bins might otherwise end up with no counts due to statistical fluctuations.

One disadvantage of this technique is that it assumes that background rejection and supernova clustering would have done the exact same thing; the background clusters are not altered in any way. If instead I re-ran the full simulation, signal hits would appear in different locations and times throughout the detector. The clustering algorithm might end up grouping a signal hit into what used to be a background cluster, for example. However, this effect is thought to be negligible.

## 8.2 Cluster selection

The first selection step is a pre-selection stage. To a certain degree, SNSlicer has already applied some pre-selection in that it discards any clusters that only contain one hit and for the far detector it discards any clusters that contain hits from only one detector view. Additionally, I apply a fiducial volume cut to ensure containment and also to discard clusters that are more prone to background contamination




Figure 8.1: Energy-time fluxes for  $\nu_e$  (left) and  $\bar{\nu}_e$  (right) flavors. *Top:* SFHo-z9.6co model neutrino flux. *Center:* LS220-s27.0co model neutrino flux. *Bottom:* Ratio of model neutrino fluxes (SFHo-z9.6co/LS220-s27.0co).

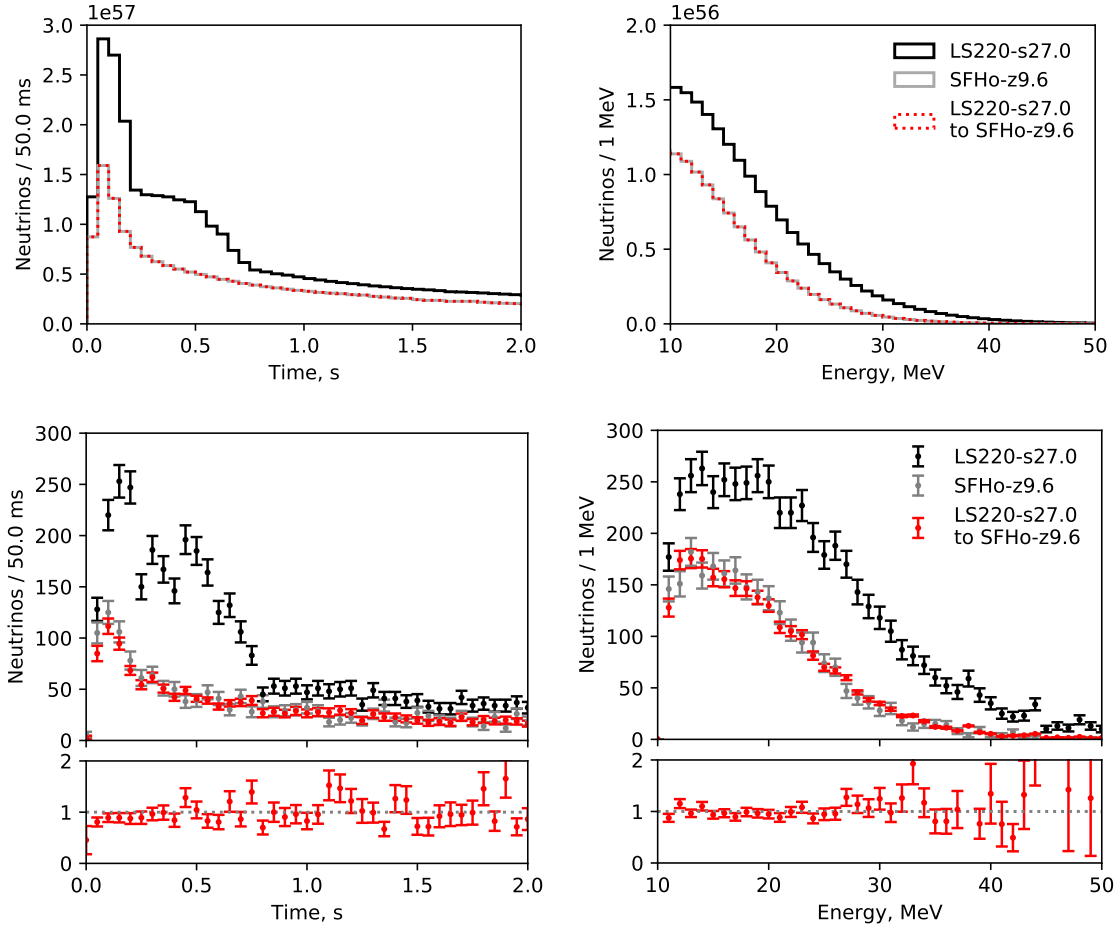


Figure 8.2: Reweighting the neutrino time profile (left) and fluence (right) to convert between supernova models. The black and gray solid lines are the true distributions for the two models simulated at the same distance. The red distributions are the results of reweighting the LS220-s27.0co model to SFHo-z9.6co. *Top*: the model fluxes used as input to the simulation as a cross-check that this process works as expected. *Bottom*: the neutrinos generated by the simulation.

due to interactions in the surrounding rock and overburden. For the near detector, I discard all clusters whose mean position is within 30 cm of any detector wall. The same selection applies to the far detector except for the top face of the detector, where the distance-to-wall cut is 45 cm due to increased cosmogenic activity in this region of the detector.

After pre-selection, the aim is to distinguish signal clusters from background clusters based on their macroscopic properties. Descriptions of these properties are listed in Tab. 8.1 and are plotted in Figs. 8.3 and 8.4. A visual inspection of these distributions for signal and background clusters reveals that there is a lot of overlap between the two classes. A common tactic in particle physics for signal selection involves performing a series of rectangular cuts on the variables. This can be a simple and effective method when the signal and background distributions have good separation, but that is not the case here. It is difficult to tell by eye which variables provide the best discriminating power or in what order cuts should be applied. Rectangular cuts are therefore not an attractive choice for supernova cluster selection.

But another visualization technique suggests that there is hope. t-SNE<sup>1</sup> [131] is a method of projecting high-dimensional data onto two or three dimensions for visualization. It preserves the separation between points in the higher-dimensional space and can be useful for identifying the presence of clusters in the data. Figure 8.5 shows a 2D t-SNE plot for each detector and reveals that there are indeed some groupings. The near detector consists of mostly signal, but the few background clusters that do exist are largely located together near the top of the plot. Likewise the far detector exhibits a grouping of signal clusters in the upper-right portion of the plot, although they are not tightly packed and many are scattered throughout the plot among the background clusters. This suggests that it will be more difficult to separate the signal and background clusters at the far detector than at the near detector, but not nearly as difficult as the feature distributions indicated. To address this challenge, I turn to machine learning.

There are many details to consider when choosing a machine learning strategy. First, there is the question about how the data need to be cleaned and processed to make them suitable for training. Then there is the choice of the algorithm itself and optimal choices for its hyperparameters. Finally, one must decide how to evaluate the output. How do we quantify performance? How do we verify that the machine learning model has not been under-trained or over-trained and that it will generalize well to new data? There is rarely a “one size fits all” approach to most of these considerations because the details often depend on the nature of the problem one is trying to solve, the quality and quantity of data available for training, and the desired outcomes. In the following sections, I will lay out the machine learning strategy that I employ to classify clusters as either “signal” or “background.”

## Machine learning classifiers

Training data are partitioned into two components: a set of features  $X$ , which in this case correspond to the cluster properties, and a vector of class labels  $Y$ :

$$X = \begin{bmatrix} x_{11} & x_{12} & \cdots & x_{1m} \\ x_{21} & x_{22} & \cdots & x_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \cdots & x_{nm} \end{bmatrix}; \quad Y = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}. \quad (8.4)$$

---

<sup>1</sup>t-SNE is short for “t-distributed Stochastic Neighbor Embedding.”

Table 8.1: Cluster properties.

Feature	Variable name	Description
Number of hits	NHits	Number of hits contained in the cluster
Number of views	Views	Number of detector views that contributed hits to the cluster (1 or 2)
Largest time difference	TimeGapLarge	Largest time difference (ns) between any two hits in the cluster
Smallest time difference	TimeGapSmall	Smallest time difference (ns) between any two hits in the cluster
Largest plane difference	PlaneGapLarge	Largest plane difference between any two hits in the cluster
Smallest plane difference	PlaneGapSmall	Smallest plane difference between any two hits in the cluster
Largest cell difference	CellGapLarge	Largest cell difference between any two same-view hits in the cluster
Smallest cell difference	CellGapSmall	Smallest cell difference between any two same-view hits in the cluster
Mean X-position	XMean	Mean detector position (cm) of the cluster in the detector X-coordinate
Mean Y-position	YMean	Mean detector position (cm) of the cluster in the detector Y-coordinate
Mean Z-position	ZMean	Mean detector position (cm) of the cluster in the detector Z-coordinate
Total ADC	AdcTotal	Total ADC of all hits in the cluster
Mean ADC	AdcMean	Mean hit ADC in the cluster
Smallest ADC difference	AdcGapSmall	Smallest ADC difference between any two hits in the cluster
Largest ADC difference	AdcGapLarge	Largest ADC difference between any two hits in the cluster
Photoelectrons	PE	Total photoelectrons of all hits in the cluster
Total energy	TotalEnergy	Total calorimetric energy (MeV) of the cluster
Time to track	TimeToTrack	Time difference (ns) between cluster and most recent track

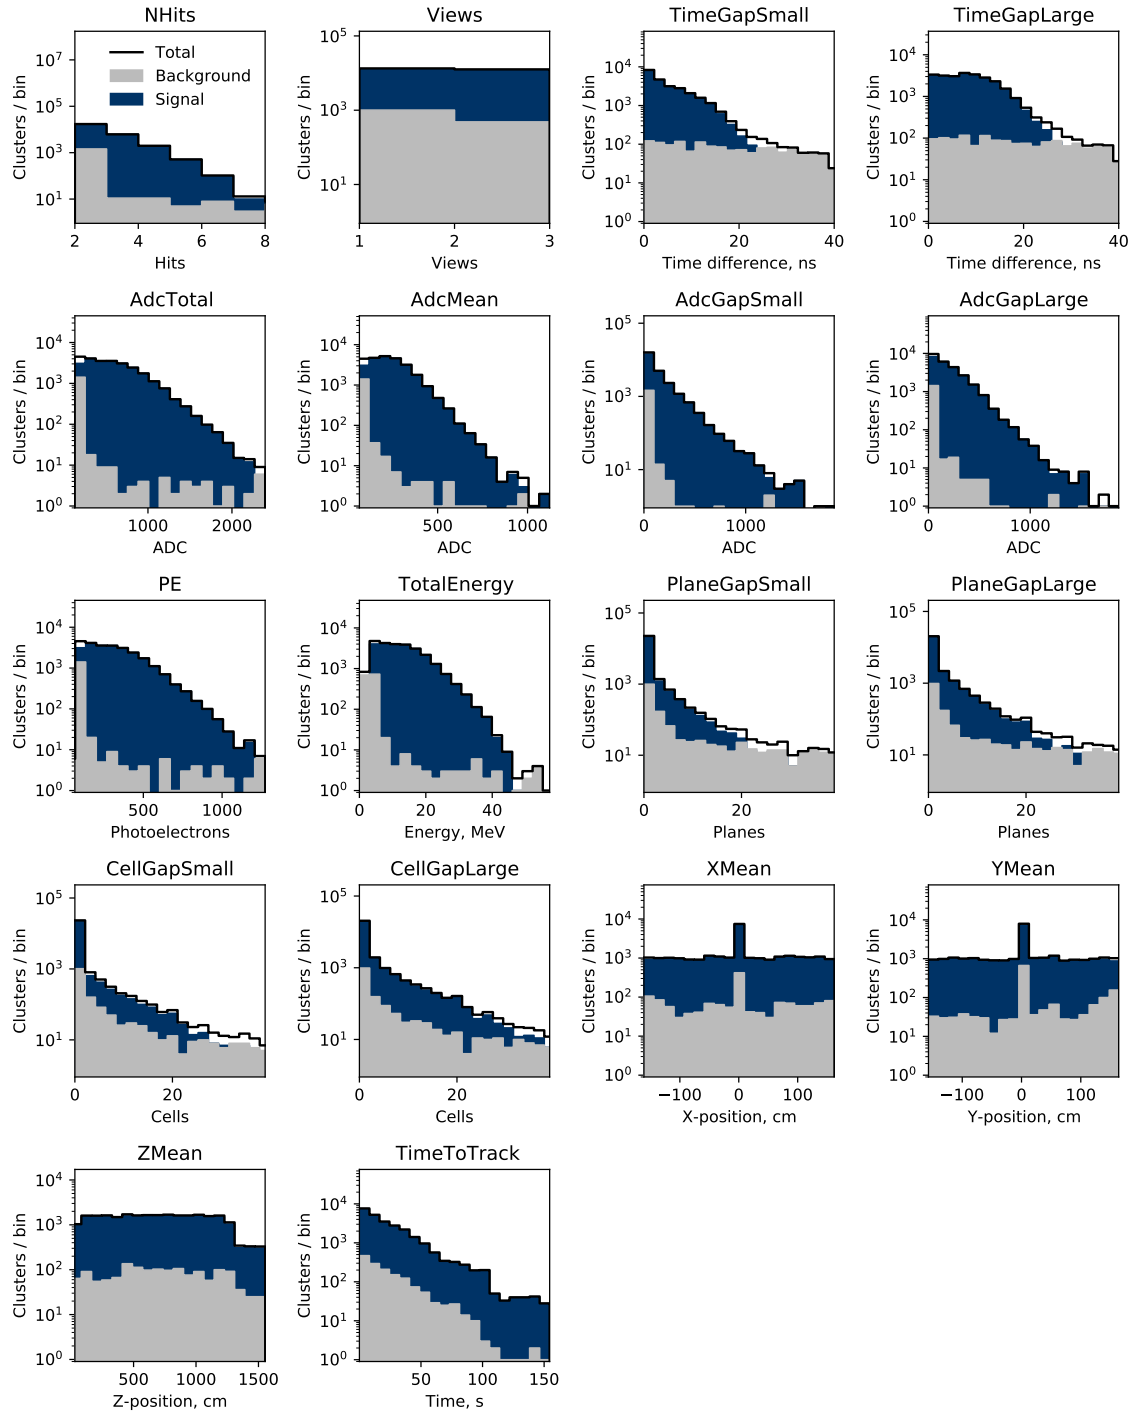


Figure 8.3: Near detector cluster properties.

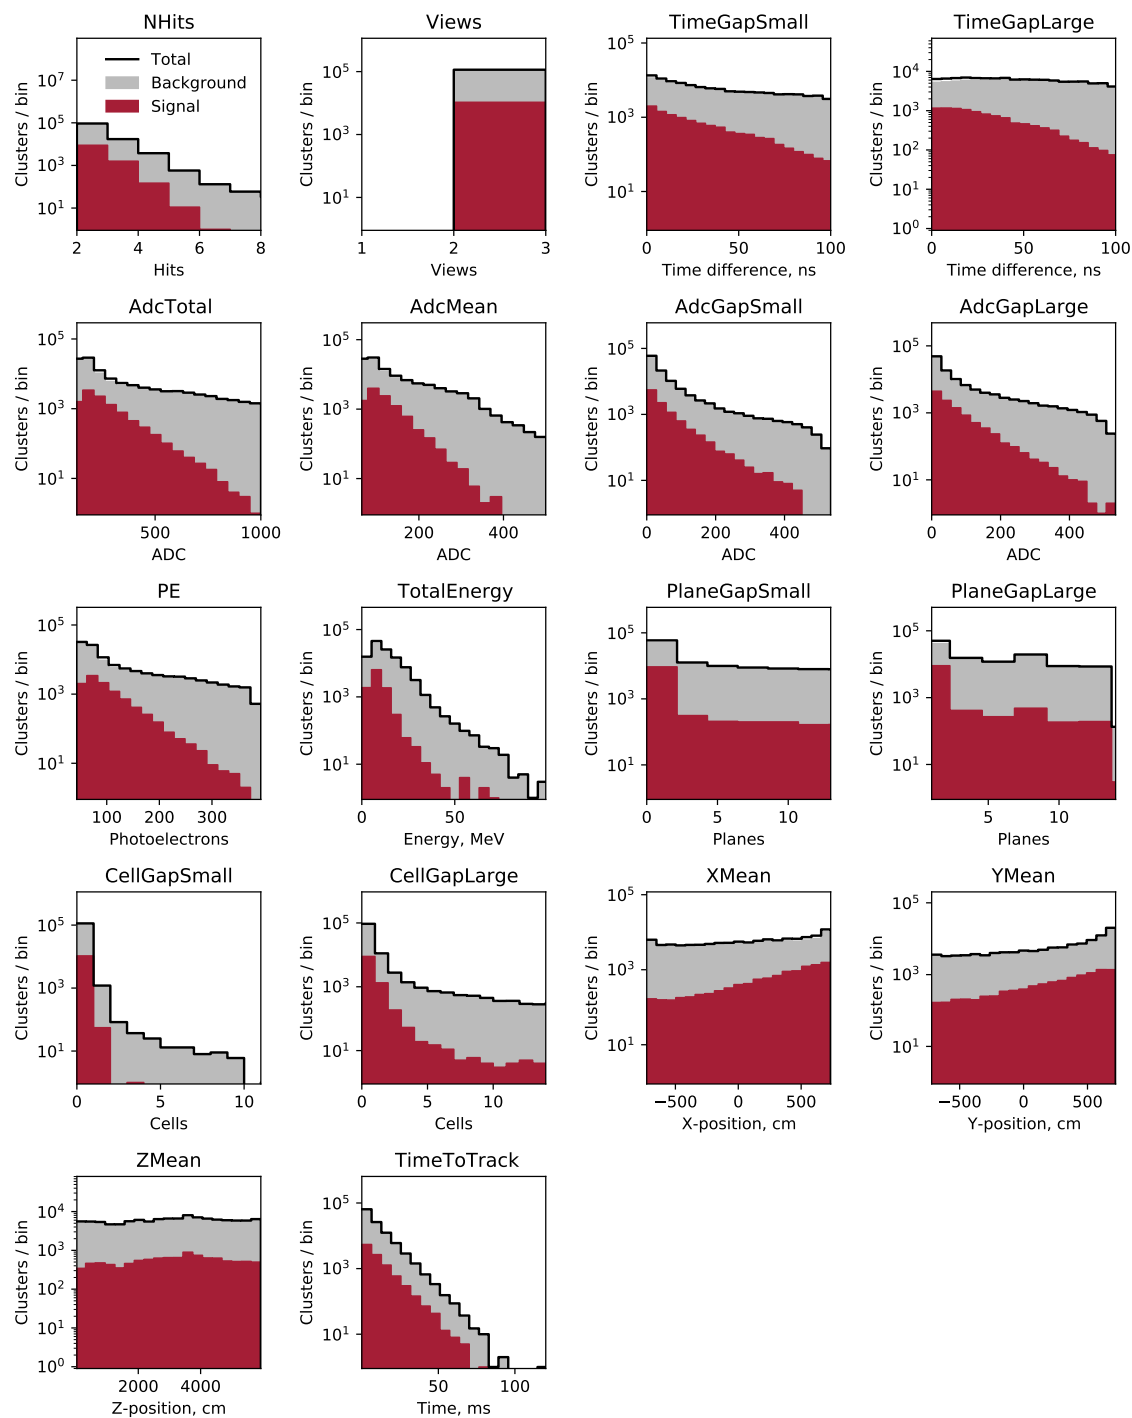


Figure 8.4: Far detector cluster properties.






Figure 8.5: t-SNE plots for the near (left) and far (right) detectors. Blue or red points indicate signal clusters and gray points indicate background clusters. The axes represent a projection of a high-dimensional space onto two dimensions and therefore lack a physical meaning, so axis and tick labels are not shown.

The class labels for binary classification are 0 and 1, indicating whether the true class of that cluster is the “positive” class (signal) or the “negative” class (background). This differs from regression problems, where the goal is to estimate the value of a continuous variable. Focusing solely on binary classification limits the options of suitable algorithms, but not by much; there are still many options and no way to choose the best one without trying each one and comparing their performance for this particular data set. I decided to test four classifiers from different families:

**Random forest** [132]. An ensemble method that consists of a number of individual decision tree (DT) classifiers. Each DT classifier is trained on a random subset of the full data set and with a random subset of the full list of features to reduce the correlation between individual DTs. The random forest classifier makes its predictions based on the collective predictions made by its component DTs.

**AdaBoost** [133]. A sequence of “weak learners”—simple models which only slightly outperform randomly guessing. The weak learners in this case are decision stumps, which are just very small decision trees. The final prediction is made by taking a weighted average of predictions from the component decision stumps, where the weights are optimized during the model training process.

**Logistic regression** [134]. A sigmoid function is used to model the probability distribution of being in the positive class and a threshold somewhere in that distribution enables binary classification. This requires the  $N$ -dimensional feature-space to be reduced down to a single dimension on the domain  $[0, 1]$ , and it is this reduction in dimensionality that is optimized during training.

**Gaussian naive Bayes** [135]. A probabilistic model based on Bayes’ theorem [136]. This method is described as naive because the features are assumed to be independent of one another (correlations are not considered) and each feature is considered to have an equal effect on the outcome. Given these assumptions, a gaussian posterior probability distribution is used to make classification decisions.




Figure 8.6: Class populations for both detectors.

### Imbalanced classes

When training any machine learning model, care must be taken when dealing with data sets that exhibit large imbalances between classes. This is typical in searches for rare events. For example, spam filtering, credit card fraud detection, and cancer screening are all situations in which positive cases are largely outnumbered by the negative ones, but identifying positive cases is critically important. Such is the case for this supernova search.

Both detectors exhibit imbalanced classes, but in opposite ways. Background clusters for the far detector outnumber signal clusters by about 9-to-1. At the near detector, the signal outweighs the background clusters 16-to-1 (Fig. 8.6). If this imbalance is not addressed, the models trained for both detectors are likely to be over-trained on the data in the majority class and under-trained on the data in the minority class.

There are a variety of ways to approach the problem, which tend to fall broadly into the following categories:

**Under-sampling.** Under-sampling balances the training data set by selecting a subset of the majority class so that the majority and minority classes are equal in size. This can be done by simply removing majority class records randomly or through a more sophisticated method, such as removing majority class points which occur close to minority class points in the feature space (known as “tomek links”) and thus improving the separation between the classes. Regardless of the method, under-sampling has the disadvantage of throwing out good data that could be used to improve model performance. This may prove to be untenable for smaller data sets.

**Simple over-sampling.** Over-sampling adds more instances of the minority class by duplicating existing instances at random until the two classes are balanced. This can be useful when the available data set is fairly small and under-sampling is not an option. However, the duplication of data can lead to over-fitting the model, and it may not generalize well to new data.

**Synthetic over-sampling.** The synthetic minority over-sampling technique (SMOTE) [137] increases the size of the minority class by generating synthetic samples of the class. This process

		Predicted		
		Positive	Negative	
True	Positive	True positive $T_p$ (Correct)	False positive $F_p$ (Incorrect)	Precision $\frac{T_p}{T_p + F_p}$
	Negative	False negative $F_n$ (Incorrect)	True negative $T_n$ (Correct)	Negative Predictive Value (NPV) $\frac{T_n}{T_n + F_n}$
		Recall $\frac{T_p}{T_p + F_n}$	Specificity $\frac{T_n}{T_n + F_p}$	

Figure 8.7: Confusion matrix for binary classification.

involves selecting a minority class sample, choosing one of its  $k$  nearest neighbors, drawing a straight line in the feature space between these two minority class samples, and generating a new sample at a random point along that line. This produces a set of new minority class samples that are not simply duplicates of existing samples, and can reduce the effects of over-fitting that can plague duplicative over-sampling methods.

## Performance metrics

It is not possible to compare combinations of the aforementioned classifiers and class imbalance techniques without specifying how performance will be quantified. One commonly-used metric is the accuracy, the fraction of correct classifications. However, this becomes less insightful as the data set becomes more imbalanced. For example, if the minority class made up 1% of all the data, an algorithm that naively classifies everything as the majority class will have an accuracy of 99%.

More meaningful metrics can be derived by looking at the confusion matrix, which compares the number of class predictions from a given model with the true classes (Fig. 8.7). All predictions are decomposed into true positives  $T_p$ , false positives  $F_p$ , true negatives  $T_n$ , and false negatives  $F_n$ .

Two important metrics derived from the confusion matrix are “recall” ( $T_p / (T_p + F_n)$ ) and “precision” ( $T_p / (T_p + F_p)$ ). Recall is the percentage of positive class samples predicted correctly and precision is the percentage of the predicted positive classes that are actually positive. These are useful when identifying the positive cases is the most important aspect of classification. A third metric, the f1-score, is a harmonic average of precision and recall, and is often used to gauge the performance for situations with class imbalances.

Out of the four metrics described so far—accuracy, recall, precision, and f1-score—the best choice depends on the nature of the data and what outcomes are most important for the problem, but that can make it difficult to compare the performance of various class balancing scenarios. For this, we can derive a new metric from the receiver operator characteristic (ROC) curve.

A binary classifier generally assigns a numerical probability to each prediction that it makes, and if the probability is above some threshold then it is predicted to be the positive class. The ROC curve is




Figure 8.8: Receiver operator characteristic (ROC) curves. The dashed gray line corresponds to  $y = x$  and is the ROC curve for a classifier that simply guesses. The step function curve (solid black) represents the response of a perfect classifier. The other curves correspond to models that perform between those two extremes.

a plot of the true positive rate versus the false positive rate for various possible thresholds (Fig. 8.8). If the threshold is set to zero, then everything is labeled positive, meaning the true positive rate is 1 and false positive rate is 1. If the threshold is set to one, then everything is labeled negative, making the true positive rate 0 and the false positive rate 0. So  $(0, 0)$  and  $(1, 1)$  are two points that every ROC curve will pass through. What happens in between depends on the performance of the classifier. A perfect classifier would have a point at  $(0, 1)$  and the ROC curve would look like a step function. A classifier that was no better than random guessing would not have any preferred threshold, and would be represented by the line  $y = x$ .

By plotting the ROC curves for multiple binary classifiers together, we can compare the performance by noticing which one gets pulled closer to the point  $(0, 1)$ . Or, we can distill that information into a single numerical score by computing the area under the ROC curve, known as the ROC AUC. For this metric, 1 represents a perfect classifier while 0.5 is equivalent to a model that guesses at random.

Table 8.2 shows the ROC AUC scores for each combination of classifier and class-balancing method described above. I applied the class-balancing techniques using the Imbalanced-learn Python package [138] and trained the classifiers using the scikit-learn Python package [139]. I used the default hyperparameters for each algorithm. That is, I made no attempt to optimize the performance of any of the techniques and therefore these trials do not demonstrate the best achievable ROC AUC score for any one scenario, but it does allow for something of an apples-to-apples comparison.

For this comparison, I also used a method called “k-fold cross validation” (Fig. 8.9). In this method, the training sample is split evenly into  $k$  sub-samples, or “folds”. During each training iteration,  $k - 1$  sub-samples are used as the training data and the remaining one sub-sample is used as a validation set to test the performance. A different sub-sample is used as the test sample during each iteration. The results of each iteration are averaged together to produce the final result, which reduces the dependence

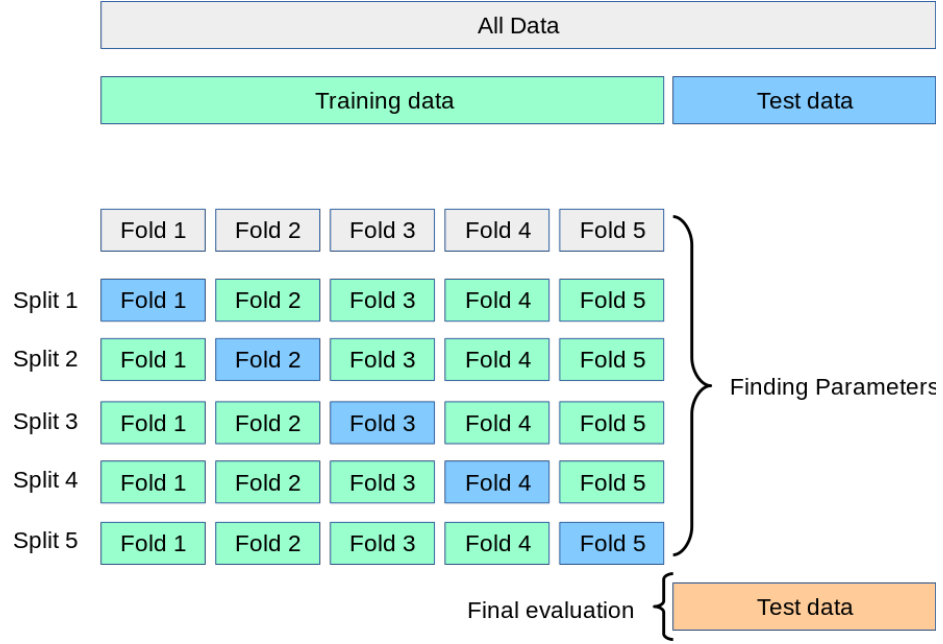


Figure 8.9:  $k$ -fold cross-validation example. Instead of training a model once and evaluating on a single test set, this procedure breaks the training data up into  $k$  folds to allow training the model  $k$  separate times. For each training iteration  $1/k$  of the data is used as a test data set, and  $(k - 1)/k$  is used as the training set. With this method, one can evaluate the average performance of a model on a single training data set. Diagram from the scikit-learn user guide: [https://scikit-learn.org/stable/modules/cross\\_validation.html](https://scikit-learn.org/stable/modules/cross_validation.html).

of the final result on the particular of how the training and test samples were split. For the comparisons shown in Tab. 8.2, I used  $k = 10$  folds for cross-validation.

The results show that the ROC AUC score for both detectors is maximized when using a random forest classifier in combination with the SMOTE technique to synthetically up-sample the minority class (signal at the far detector, background at the near detector).

### Decision trees and random forests

A decision tree is essentially a collection of yes/no questions called nodes. When a trained decision tree classifier is asked to predict the class of new data, it begins by passing the features to the root node. The root node checks if the value of a particular feature is less or greater in value than some pre-determined threshold, and passes the data to one of two sub-nodes depending on the outcome. Whichever node receives the data then repeats the process using a different feature and a different threshold, and this continues until the data reach a leaf node corresponding to a specific class, rendering the classification prediction.

Training a decision tree classifier is an iterative process that begins at the root node. The feature and the split threshold value for the node is chosen based on a cost function, a measure of misclassification. A popular choice for the cost function is the Gini impurity

$$G = \sum_{i=0}^{N_C-1} p(i) \cdot (1 - p(i)), \quad (8.5)$$

where  $N_C$  is the number of classes ( $C = 2$  for binary classification) and  $p(i)$  is the purity of class  $i$  in

Table 8.2: ROC AUC scores for different classifier models and class-balancing methods. The best score for each detector is shown in bold.

		Classifier types			
Resampling method		Random Forest	AdaBoost	Logistic Regression	Gaussian Naive Bayes
<i>Near detector</i>	None	0.9831	0.9840	0.9655	0.9504
	Under-sampling	0.9829	0.9791	0.9605	0.9498
	Tomek links	0.9848	0.9834	0.9651	0.9503
	Over-sampling	0.9837	0.9827	0.9616	0.9503
	Over/under-sampling	0.9852	0.9803	0.9640	0.9508
	SMOTE	<b>0.9857</b>	0.9811	0.9619	0.9549
<i>Far detector</i>	None	0.8706	0.8603	0.8382	0.7931
	Under-sampling	0.8681	0.8412	0.8031	0.7933
	Tomek links	0.8731	0.8569	0.8405	0.7924
	Over-sampling	0.8791	0.8563	0.8080	0.7958
	Over/under-sampling	0.8724	0.8492	0.8224	0.7896
	SMOTE	<b>0.8804</b>	0.8548	0.8078	0.7689

the sample. The Gini impurity is calculated for each branch of the split and a weighted sum of the two is taken as the total cost function for the node. The lower the value, the better the separation power. This process continues until a stopping criterion—such as reaching the maximum depth of the tree or failure to further improve the separation of the sample—has been satisfied. The final nodes of the tree are the leaf nodes which determine the prediction based on the majority class present in the node.

A random forest (RFC) classifier is an ensemble of individual decision tree classifiers. The random aspect of a RFC comes from two sources: (1) random sampling of the training data with replacement and (2) random selection of a subset of features used to train each tree. In this way, each decision tree sees different data and a different feature set during training, which reduces correlations between the decision trees in the forest. The final classification from a RFC is a majority vote from the classifications of all the component decision trees.

## Feature engineering

The performance of any machine learning algorithm strongly depends on both the quality and the quantity of the data that it is trained on. But more data is not necessarily better. For one, the computational cost of training a model generally increases with the number of features, so including highly-correlated features can increase the training time without producing any appreciable boost in performance. It is also often the case that some features are far more important to a trained model than others, so removing useless features can further decrease the time it takes to train a model on a full data set with little-to-no impact on the performance of the model. In some cases, the presence of non-important features can actually decrease the performance of a model, so even if training time is not significant, selecting the right set of features is still worthwhile.

My approach to this is straightforward. First, I look for highly-correlated features and remove all but one of them (Fig. 8.10). Next, I train an out-of-the-box random forest model with SMOTE oversampling on a subset of the full training data set and examine the importance of each feature. The importance of a feature is quantified by assessing the average relative amount that the feature

Table 8.3: Random forest hyperparameters and optimized values for each detector.  $N_{\text{feat}}$  is the total number of features.

Variable	Description	ND	FD
n_estimators	Number of trees in the forest	875	575
criterion	Function used to measure the quality of a split	Gini	Gini
min_samples_split	Minimum samples required to split an internal node	2	2
min_samples_leaf	Minimum samples required at any leaf node	1	1
max_features	Number of features to consider when looking for the best split	$\sqrt{N_{\text{feat}}}$	$\sqrt{N_{\text{feat}}}$
max_depth	Maximum depth of any one tree	$\infty$	90

contributes to impurity decrease. I discard any feature with an importance score less than or equal to 0.1 (Fig. 8.11).

### Hyperparameter tuning

With a model in hand and a reduced feature set, it's almost time to train the model and test its predictive power. But first, the hyperparameters that constrain the structure of the model during training must be chosen. Up to this point, I have been training the random forest model using the default hyperparameters provided by scikit-learn [139]. The defaults are fine for making comparisons, but not for maximizing performance. So what are the optimal hyperparameters for identifying supernova neutrino clusters at NOvA with a random forest binary classifier? There is no way to deduce that. Like most machine learning tasks, the best approach is to simply try various combinations and choose the one that performs best.

This approach occurs in two stages. The first stage is a random search of 50 combinations of hyperparameters with each hyperparameter chosen at random based on user-defined bounds. Each combination is k-fold cross-validated with  $k = 3$ . The optimal combination is the one that maximizes the ROC AUC score. This random search identifies roughly optimized values for the hyperparameters, and the second stage is a grid search designed to systematically search a reduced space to produce a more finely-tuned optimization. Table 8.3 lists the hyperparameters included in the search and the optimized values for both detectors.

### Training and performance

For each detector, I trained a random forest classifier using the SMOTE technique to balance the class frequencies on features that are neither highly correlated nor unimportant to the classifier's performance.

The training data set is 80% of the full data set with the other 20% set aside for testing the performance of the classifier after training. These data were split randomly while preserving the relative class distribution of the full data set.

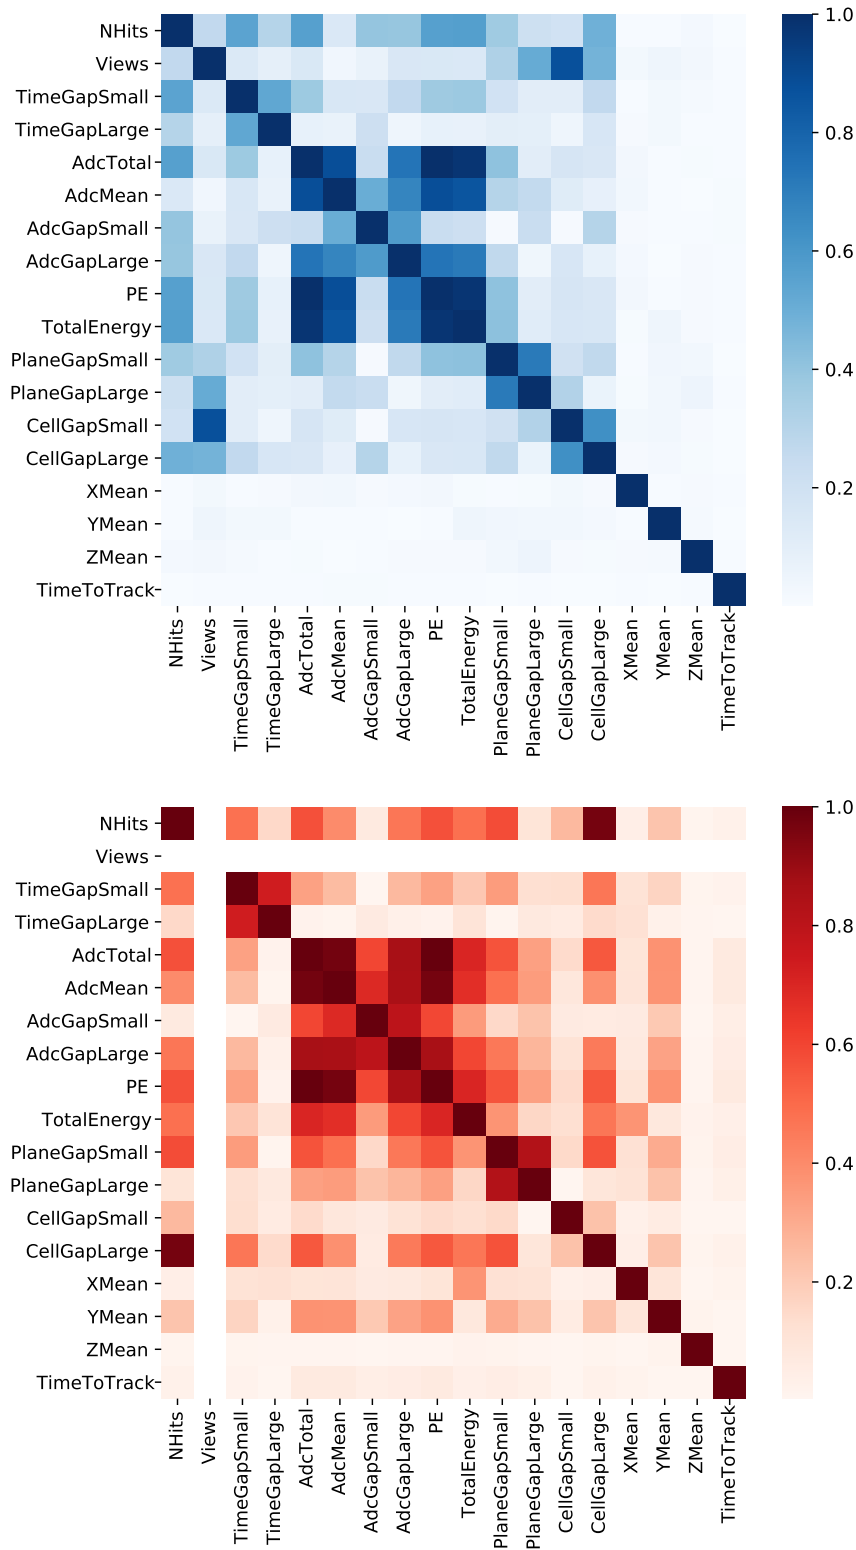


Figure 8.10: Feature correlations for the near detector (top) and far detector (bottom).






Figure 8.11: Importances for full (top) and reduced (bottom) feature sets for the near detector (left) and far detector (right).

The confusion matrices for each classifier are shown in Fig. 8.12, and the results are promising. Both the near and far detectors correctly classify nearly 95% of the true background clusters. The classifiers are able to correctly identify signal clusters 96% of the time for the near detector and 85% of the time for the far detector. As far as mistakes go, the near detector model misclassifies clusters 4% of the time. The far detector model misclassifies background clusters 6% of the time, but it misclassifies signal clusters at the higher rate of 15%. The performance comparison between these two classifiers is shown in the ROC curves in Fig. 8.13.

The values in the confusion matrix are based on classification decisions being a majority vote from all decision trees in the forest. If more than 50% of the trees predict a cluster as signal, then the RFC will predict signal, and vice-versa for background predictions. But this 50% threshold is somewhat arbitrary; why should a simple majority be the criterion that determines the classification? Different thresholds may be more suitable in specific situations. For that reason, instead of considering the classification decision for each cluster, I instead look at the fraction of trees that classified it as signal. Figure 8.14 shows the distribution of these scores and the selection efficiencies for true signal and background clusters. Additionally, I have plotted the figure of merit  $S/\sqrt{S+B}$  to demonstrate that a threshold of 0.5 is not the ideal choice in this case.

I do not use this figure of merit to choose the thresholds, however. Instead, I want to choose thresholds based on how they affect the sensitivity of the analyses in this thesis. For now, I have two optimized and trained random forest classifiers that are able to quantify how signal-like any cluster is with acceptable error rates. I also have in hand the capability to transform cluster time and energy




Figure 8.12: Confusion matrices for the near detector (left) and far detector (right). The numbers are the fraction of true signal or true background clusters in each quad.




Figure 8.13: ROC curves for near (dashed) and far (solid) detectors. The dotted line is  $y = x$  and represents a classifier that randomly guesses.

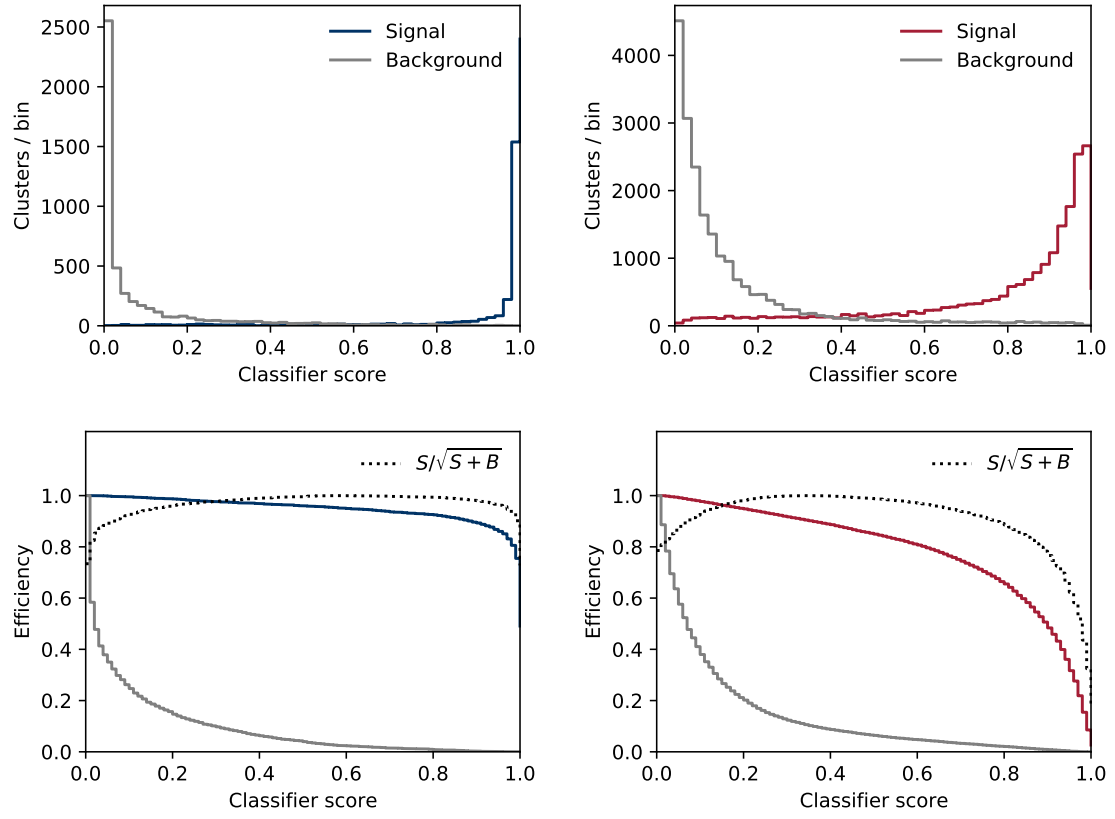


Figure 8.14: Clusters (top) and selection efficiencies (bottom) versus classifier score from the random forest classifiers trained for the near (left) and far (right) detectors.

distributions produced from the LS220-s27.0co supernova neutrino flux model into any other model on-the-fly and without the need for running additional expensive and time consuming simulations. In the next chapter, I will lay out several analyses that could be performed in the event of a real supernova detection with the NOvA detectors and compute the sensitivities of those measurements for a variety of model and progenitor distance scenarios.

## CHAPTER 9

# Analyses and Results

After ensuring the quality of the data, rejecting obvious background activity, clustering hits together based on topology, and determining how supernova-like those clusters are with a tuned random forest classifier, the next step is to analyze the data. There is only one problem: the most recent supernova with a detectable neutrino burst occurred twenty-six years before NOvA began taking data. But, this is okay. Nearly<sup>1</sup> every other current and planned neutrino detector is in the same proverbial boat.

For now, we are all limited to exploring capabilities and estimating sensitivities, but even this is a challenge. As I have discussed in previous chapters, we still lack much information about the mechanics of a supernova and how they influence neutrino emission. There are a wide range of supernova neutrino models available, each consisting of a distinct selection of relevant physics, nuclear equations-of-state, progenitor masses, and dimensionality, which begs the question: what is the “canonical” supernova? How can we benchmark the sensitivity of the NOvA detectors—or any neutrino detector—when there are so many viable models to choose from?

There is no clear answer to that question. We will not be able to seriously constrain neutrino emission models until the next galactic core-collapse supernova occurs. And even then, no one neutrino detector will be enough. Each detector is primarily sensitive to one or two interaction channels and therefore will have only a piece of the puzzle. Liquid scintillator and water Cherenkov detectors will mainly see the  $\bar{\nu}_e$  piece while noble element detectors will see  $\nu_e$ . High-Z detectors like HALO and future dark matter detectors will have sensitivity to all flavors through neutral current interactions and coherent elastic neutrino-nucleus scattering (CE $\nu$ NS). The world’s neutrino detectors will need to combine their observations as well as those produced by optical telescopes and gravitational wave observatories before a complete picture will emerge.

Give the absence of real neutrino data and the lack of any single canonical supernova neutrino model, the analysis I will present in this chapter will be that of a case study. I have chosen two models (LS220-s27.0co and SFHo-z9.6co) for which I will examine NOvA’s sensitivity to two things: (1) the ability to constrain the initial arrival time of the neutrino burst  $t_0$  and the distance to the supernova  $d_{\text{SN}}$ , and (2) the neutrino mass ordering under a specific set of assumptions.

---

<sup>1</sup>The Baksan Underground Scintillation Telescope (BUST) in the North Caucasus region of Russia reported an observation of neutrinos from SN 1987A and still operates today.




Figure 9.1: Far detector template cluster data shifted and scaled to various values of  $d_{\text{SN}}$  and  $t_0$ . Note that the background level is not affected by these transformations, as demonstrated by the overlap of every curve prior to core-bounce time of the earliest distribution at 1 s.

## 9.1 Modeling and fitting the time profile

For this analysis, I produced one simulation for each detector ( $d_{\text{SN}}$  of 350 pc for the near detector and 500 pc for the far detector) from the LS220-s27.0co model. Instead of fitting a function to the data, I use a time profile for a very close supernova overlaid with minimum-bias data as a template that can be adjusted for various values of progenitor distance and core-bounce time. The model  $F(d_{\text{SN}}, t_0, b)$  that describes this template for either detector depends on the supernova distance  $d_{\text{SN}}$ , the core-bounce time  $t_0$ , and the average background level  $b$ . This model can be decomposed into a signal portion and background portion:

$$F(d_{\text{SN}}, t_0, b) = S(d_{\text{SN}}, t_0) \oplus B(b). \quad (9.1)$$

The model can be produced for any core-bounce time  $t_0$  by shifting the template core-bounce time  $t_0^{\text{sim}}$  for each signal template cluster by some amount  $\delta t_0$  such that  $\delta t_0 = t_0 - t_0^{\text{sim}}$ . Similarly, models of any supernova distance  $d_{\text{SN}}$  can be produced by scaling the signal component of the model by a fraction  $w$  relative to the simulated distance  $d_{\text{SN}}^{\text{sim}}$ :  $w = (d_{\text{SN}}^{\text{sim}}/d_{\text{SN}})^2$ . In both of these cases, the background component of the model is not affected. With this model, I can replicate the expected time profile for a LS220-s27.0co model for any distance, core-bounce time, and average detector background level (Fig. 9.1).

In addition to these basic transformations, I can use the reweighting infrastructure described in Sec. 8.1 to transform the time profile from one supernova neutrino flux model to another. This same procedure can also be used to apply neutrino flavor transformations. The LS220-s27.0co model does not already account for any oscillations, though several effects would likely imprint themselves on the signal. Among them are MSW oscillations in the stellar material, collective oscillations in regions of high density, decoherence effects as the mass state propagate through space, and Earth-matter effects if the neutrinos traverse the interior of the planet before detection.

For simplicity, I only consider the case of adiabatic MSW oscillations in the stellar interior. In re-




Figure 9.2: Selected clusters over time under different mass ordering scenarios. Only adiabatic MSW effects within the collapsing core are considered.

regions of high electron density, neutrino mass eigenstates propagate adiabatically so long as the matter density profile is slowly-varying. The fluxes are transformed according to the Eqs. 4.5 and 4.6 introduced in Sec. 4.3. To reweight the selected clusters to either of these mass ordering scenarios, I take a ratio of the oscillated and unoscillated model fluxes in bins of time and energy and scale the weight of each cluster by that ratio based on its parent neutrino true time and energy (Fig. 9.2).

### Time Binning

The time evolution of neutrino emission from a core-collapse supernova does not play out on any one time scale. The intricacies of the neutronization burst are contained within the first few milliseconds, the turbulent instabilities that modulate the neutrino flux during accretion last for tenths of a second, the remaining neutrino emission spills slowly out of the core during the cooling phase which lasts for tens of seconds. So when it comes to fitting the time profile of a template to data, what is the appropriate binning? Should it be fine or coarse or something in between? Perhaps there is no correct answer, but for this thesis I chose a coarse binning of 0.5 s for several reasons.

For one, it is simply the natural choice. I am not considering only the very early supernova activity in my analysis; I am looking out as far as possible before the signal becomes buried in background. That timescale is on the order of seconds, and binning at neutronization timescales would require thousands of bins. This choice would have logistical consequences for supernova progenitors at larger distances where low statistics could drive the expected or measured counts in some bins to zero, complicating the computation of key statistical metrics such as  $\chi^2$ .

Another benefit of using coarse time binning is that model-dependent details on smaller timescales get washed out. The templates that I fit to the data—whether the data are those simulated for this thesis or those from a future real supernova—were produced from a simulation and therefore assume a particular supernova model. For the purposes of model discrimination in the presence of real data,

a fine binning might be desirable, but for the purposes of this thesis, I only seek to fit the general properties of a supernova event: a sharp peak early on followed by a brief shoulder during accretion and terminated by a period of decay during cooling. Coarse binning achieves this.

## Fitting algorithm

I chose to use Powell’s method [140] as the minimization algorithm to perform the fit. This algorithm is a good choice because it does not rely on estimating derivatives (the  $\chi^2$  surfaces in parameter space in this case have sharp boundaries that are not well described with functions) and it was more robust to time and distance shifts than many other popular algorithms. Powell’s method requires an initial guess  $X_0$  and two non-parallel direction vectors  $\vec{h}_1$  and  $\vec{h}_2$  as input. The point  $X_1$  is the minimum point on the line in the direction of  $\vec{h}_1$  that passes through  $X_0$ . The point  $X_2$  is the minimum point on the line in the direction of  $\vec{h}_2$  that passes through  $X_1$ . A conjugate direction is then chosen by drawing a line through the points  $X_0$  and  $X_2$ , and this process repeats until it converges near the optimal value  $X_{\text{opt}}$ .

For the initial guess  $X_0$ , I compute a coarse estimate for three of the four parameters in the following way:

- **Detector background levels ( $b_{\text{ND}}$  and  $b_{\text{FD}}$ )** — I fit a background-only model (a constant) to each detector separately using Powell’s method. The initial guess in each case is 1. I take the result of this fit  $b_{\text{guess}}$  to be the guess for the signal-and-background model with bounds of  $[0, 1.5 \times b_{\text{guess}}]$ .
- **Neutrino burst time ( $t_0$ )** — I use the `scipy.signal` library to identify any peaks with a prominence greater than zero in the near detector time profile. I take the largest peak to be the assumed neutrino burst, and use its time  $t_{\text{guess}}$  as the initial guess. I also constrain the possible range of this parameter to be  $[t_{\text{guess}} - 0.5 \text{ s}, t_{\text{guess}} + 0.5 \text{ s}]$ . I only use the near detector time profile for this because it is more likely to have a prominent peak for distant supernovae.

The initial guess for the remaining parameter  $d_{\text{SN}}$  is 1 kpc and is allowed to take on any values between 0 and 100 kpc during the joint-detector fit.

## Asimov data sets

I was only able to produce a single simulation per detector for analysis due to computational limitations, and therefore any sensitivity computation will depend on the particular minimum-bias background data that I overlaid on the simulated signal. With unlimited resources, I would produce many simulations and overlay them with different minimum-bias data and combine them to calculate the detection sensitivity to a “typical” supernova. But this is not possible.

To get around this, I make use of the so-called Asimov data set. In his 1955 sci-fi short story *Franchise* [141], Isaac Asimov envisions a future where the President of the United States is no longer elected by the people as a whole, but rather by a single individual<sup>2</sup> who is meticulously selected by a supercomputer before each election by virtue of being the most representative American.

This idea can be applied to characterizing the sensitivity of an experiment to an observed signal by replacing an ensemble of simulated scenarios with a single representative instance [142].

---

<sup>2</sup>The story covers the 2008 U.S. election, in which the “Voter of the Year” is Norman Muller of Bloomington, Indiana. I have to believe that somehow the coincidence that I am writing and defending this thesis in the same city is further evidence that the use of Asimov data sets for this thesis is the right choice, but only so much so that this thought is relegated to a tongue-in-cheek footnote.



## Joint-detector fits

Two fits are performed on both detectors simultaneously using the `lmfit` Python package [143]: (1) a joint-detector fit of the signal-plus-background model and (2) a joint-detector fit of the background-only model. The signal-plus-background model contains four parameters, two of which ( $t_0$  and  $d_{\text{SN}}$ ) are constrained to be equal for both detectors, and the background-only model contains two parameters,  $b_{\text{ND}}$  and  $b_{\text{FD}}$ . After performing the fit, `lmfit` computes the  $1\sigma$  uncertainties on the parameters by inverting a matrix of the second derivatives of fit quality for each parameter.

The fit minimizes the sum of bin-by-bin residuals across both detectors, with each residual being weighted by the total number of clusters in that bin. Minimizing the relative residuals instead of the absolute residuals in this way is necessary so that the detector with the higher background rate (far detector) does not dominate the calculation.

To evaluate the goodness-of-fit, I compute the  $\chi^2$  which is proportional to the log-likelihood under the assumption that the variables are Poisson distributed:

$$\chi^2 = -2 \ln \lambda(\theta) = 2 \sum_{i=1}^N \left[ \mu_i(\theta) - n_i + n_i \ln \frac{n_i}{\mu_i(\theta)} \right], \quad (9.2)$$

where the last term is zero when  $n_i$  is zero and the logarithm is undefined. The  $\chi^2$  is calculated for both the background-only and the signal-plus-background models, and the difference is taken as a measure of the significance:

$$\Delta\chi^2 = \chi_B^2 - \chi_{S+B}^2. \quad (9.3)$$

The signal-plus-background model has four parameters and the background-only has two, so  $\Delta\chi^2$  has two degrees of freedom, and this can be used to translate  $\Delta\chi^2$  into a significance measured in  $\sigma$ . It is this sensitivity that I use to optimize the classifier score cut by choosing the combination of near detector and far detector scores that maximize  $\Delta\chi^2$ . Figure 9.3 shows the  $\Delta\chi^2$  values for different choices of score threshold for a 1 kpc simulated supernova using the LS220-s27.0co model. The maximum  $\Delta\chi^2$  corresponds to a near detector score of 0.625 and a far detector score of 0.650.

With the classifier score thresholds chosen, cluster selection is complete. Figure 9.4 shows the selection efficiency of neutrinos that interacted in the fiducial volumes of each detector for a simulated supernova using the LS220-s27.0co model. The efficiency increases with energy out to 70 MeV, beyond which the uncertainty on the efficiency grows rapidly due to diminished statistics. The near detector has an overall neutrino selection efficiency of  $(47.7 \pm 0.5)\%$ . At the far detector, the overall efficiency of  $(0.748 \pm 0.008)\%$  is substantially lower, but it reaches a maximum value of  $(10.7 \pm 2.5)\%$  at 68 MeV.

Figure 9.5 shows both detector time profiles and joint-detector fit result for that same 1 kpc supernova using the optimized classifier score cuts for simulated clusters and a synthetic Asimov dataset.

## 9.2 Sensitivity to supernovae

By constructing an Asimov dataset for both detectors and conducting joint fits with a parameterized template for a given supernova flux model at a variety of distances, I determined NOvA's detection sensitivity at the  $3\sigma$  level to be 14 kpc for the SFHo-z9.6co model and 16.6 kpc for the LS220-s27.0co (Fig. 9.6).




Figure 9.3:  $\Delta\chi^2 = \chi_B^2 - \chi_{S+B}^2$  for a joint-detector fit of a 1 kpc supernova at different classifier score thresholds. The white marker indicates the cut values that maximize  $\Delta\chi^2$ : 0.650 for the far detector and 0.625 for the near detector.




Figure 9.4: All neutrino interactions and selected signal clusters for the near (left) and far (center) detectors and selection efficiency for both detectors (right) as a function of true neutrino energy. Only neutrino interactions within the fiducial used for cluster pre-selection. Neutrinos were simulated using the LS220-s27.0co model flux at a distance of 350 pc for the near detector and 500 pc at the far detector.

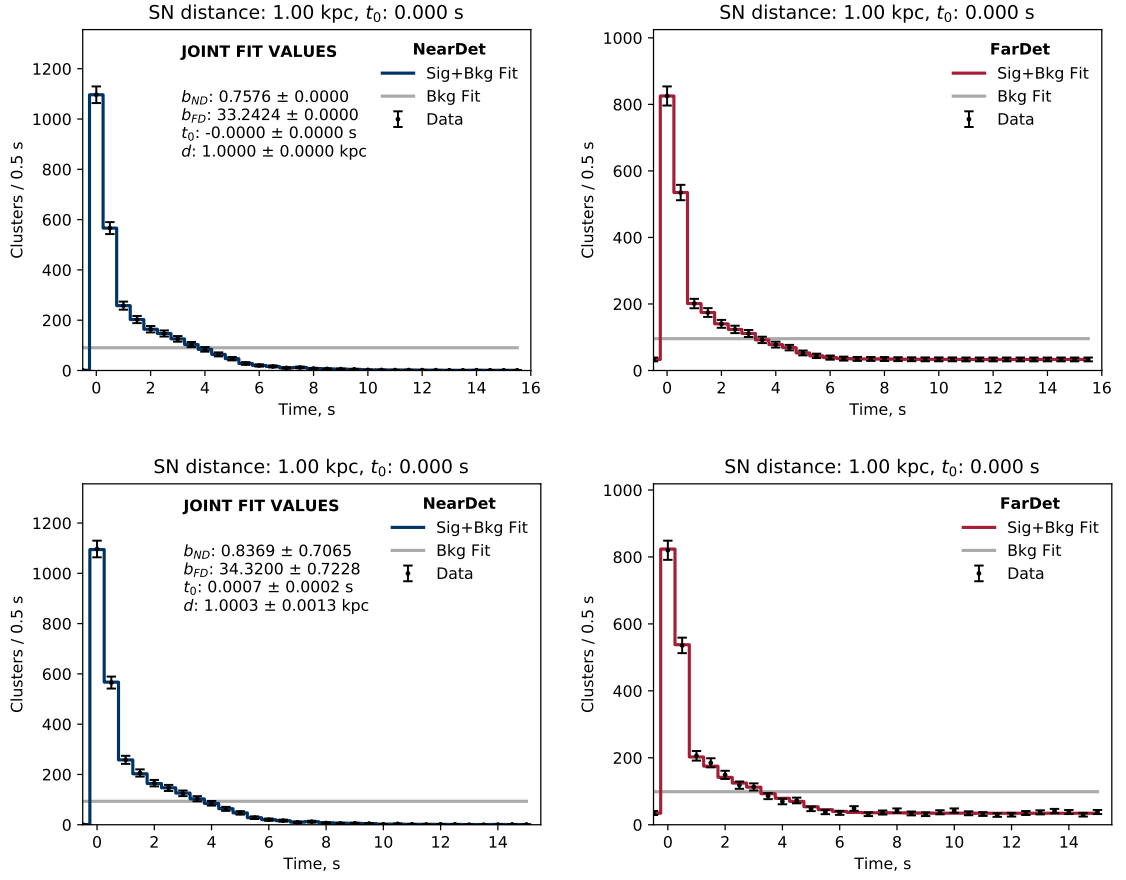


Figure 9.5: Joint-detector time profile fit for Asimov data sets (top) and simulated clusters (bottom) for the near detector (left) and far detector (right).




Figure 9.6: Joint-detector time profile fit sensitivity as a function of supernova distance.




Figure 9.7: Joint detector time profile fit to mass ordering sensitivity for LS220-s27.0co (left) and SFHo-z9.6co (right).

### 9.3 Mass ordering

As described in Sec. 4.3, there are a variety of flavor transformations that may affect the observable supernova signal, including MSW effects in the star and Earth, neutrino self-interactions, and decoherence. For this analysis, I consider only the MSW effects in the interior of the star on the neutrino time profile according to Eqs. 4.5 and 4.6. This assumes that self-interactions do not appreciably affect observed signal and that the neutrinos do not travel a significant distance in the Earth prior to detection.

I examine this effect by reweighting the full unoscillated time distribution according to Eqs. 4.5 and 4.6. I approach this using two separate methods, which are described in the following sub-sections.

#### Method 1: Joint-detector fit

This method is very similar to that of determining supernova sensitivity; I transform the time profile for a particular distance and a mass ordering scenario. Instead of comparing the signal-plus-background and background-only hypotheses, for each of the two model scenarios I considered, I compared the hypothesis of one mass ordering when fit to data for the other ordering.

Figure 9.7 shows the results. For both models, this method was better at disfavoring the normal ordering hypothesis when fitting to the inverted ordering. For all cases, a  $3\sigma$  sensitivity to the mass ordering extends beyond 1 kpc.

#### Method 2: Time asymmetry

The other method I use to determine the mass ordering is comparing the time profile at early times to later times. The MSW effect is largely considered to be robust during the neutronization burst period because the situation is uncomplicated by collective oscillations or other complicated physics (Fig. 9.8).

One way to quantify this is by computing an asymmetry between the neutronization and accretion phases. However, the event counts during the burst phase are suppressed at NOvA due to the lower energies involved and thresholding effects. To get around this, I expanded the search window to instead compare the burst and accretion periods to the initial cooling period. The assumption that the MSW effect plays a dominant role during this period isn't as robust as it would be for an earlier time window given our current understanding. But if it does play such a role, NOvA will have some sensitivity to it




Figure 9.8: Simulated event rates for a 10 kpc electron-capture supernova in a 20 kt liquid scintillator detector given different mass ordering scenarios. The event rates are only affected by MSW effects in the star. Figure from [54].

and thus the mass ordering.

Calculating the asymmetry requires counting the number of signal clusters in each time region and comparing those counts. The time profile consists of number of clusters from both detectors added together in each time bin. The number of clusters in the  $i$ -th time bin is  $N_i$ . The number of clusters  $N_R$  in any time window  $R$  is then

$$N_R = \sum_{i=i_{R,0}}^{i_{R,f}} N_i \quad (9.4)$$

where  $i_{R,0}$  and  $i_{R,f}$  indicate the first and final indices that define  $R$ .

$N_i$  can be split up into its component signal and background,  $N_i = S_i + B_i = S_i + b$  where  $b$  is the average number of background clusters per time bin. This background rate and its uncertainty is estimated by fitting the signal-plus-background time profile hypothesis to the data. The asymmetry in the signal between two time windows  $R1$  and  $R2$  is

$$A(R1, R2) \equiv \frac{S_{R2} - S_{R1}}{S_{R2} + S_{R1}} = \frac{(N_{R2} - b n_{R2}) - (N_{R1} - b n_{R1})}{(N_{R2} - b n_{R2}) + (N_{R1} - b n_{R1})} \quad (9.5)$$

$$= \frac{N_{R2} - N_{R1} - b(n_{R2} - n_{R1})}{N_{R2} + N_{R1} - b(n_{R2} + n_{R1})} \quad (9.6)$$

where  $n_{R1}$  and  $n_{R2}$  represent the number of bins in regions  $R1$  and  $R2$ . The uncertainty on the asymmetry is then

$$\sigma_A^2 = \left( \frac{\partial A}{\partial N_{R1}} \right)^2 \sigma_{N_{R1}}^2 + \left( \frac{\partial A}{\partial N_{R2}} \right)^2 \sigma_{N_{R2}}^2 + \left( \frac{\partial A}{\partial b} \right)^2 \sigma_b^2 \quad (9.7)$$

where the quantities  $N_{R1}$  and  $N_{R2}$  are assumed to obey gaussian statistics, and therefore  $\sigma_{N_x}^2 = N_x$ . Computing the partial derivatives and simplifying yields

$$\sigma_A^2 = \frac{4}{\alpha^2} \left[ (b n_{R2} - N_{R2})^2 N_{R1} + (b n_{R1} - N_{R1})^2 N_{R2} + (N_{R2} n_{R1} - N_{R1} n_{R2})^2 \sigma_b^2 \right] \quad (9.8)$$




Figure 9.9: MSW oscillations applied to Asimov data sets for both detectors for a supernova at 1 kpc under the LS220-s27.0co (left) and SFHo-z9.6co (right) flux models. The regions  $R1$  and  $R2$  are demarcated by the vertical red lines.  $R1$  ranges from 0 to 0.7 s and  $R2$  ranges from 0.7 s to 6 s. The asymmetry between these regions is calculated for each mass ordering scenario, as well as the asymmetry difference  $\Delta A = A_{\text{NO}} - A_{\text{IO}}$  and its associated significance  $\sigma$ .

where  $\alpha \equiv (N_{R1} + N_{R2} - b(n_{R1} + n_{R2}))^2$ .

Figure 9.9 illustrates this how this calculation is performed for a 1 kpc supernova under both flux models. Replicating this procedure for many distances results (Fig. 9.10) shows that the asymmetries remain constant in each mass ordering scenario, as one would expect, but the uncertainties grow as described in Eq. 9.8, and so the separation power between these two scenarios decreases with distance (Fig. 9.11).

## 9.4 Discussion and summary

The sensitivity of the NOvA detectors to the mass ordering shown in this thesis are promising. In each of two methods, a  $3\sigma$  measurement is possible for progenitors within 1 kpc depending on the model. However, this result depends the key assumptions that the adiabatic MSW effects in the stellar interior remain robust out to times of 700 ms after core bounce and that no other flavor transformation effects play an appreciable role.

That latter assumption is almost certainly wrong, but in what way we do not yet know. It's likely that neutrino self-interactions will have an effect on the observed signal and it's entirely possible that the neutrinos do not come from above, but from below after traveling a considerable distance through the Earth. There is also the question of what effect the expanding shock front has on the observed signal. At such a region of discontinuous density, the MSW flavor transformations are no longer adiabatic.

So what does this result really tell us about NOvA's sensitivity to the mass ordering from the next galactic core-collapse supernova? NOvA has sensitivity under a specific set of conditions, but more generally that will depend on the dynamics of the explosion. We simply do not have enough knowledge of what to expect of the signal properties to make robust predictions. Every detector today has to grapple with those same uncertainties. I have made no attempt in this work to combine the expected observations of other detectors with that of NOvA, but when the supernova finally comes, information from other observatories will allow us to make more informed assumptions that will improve the

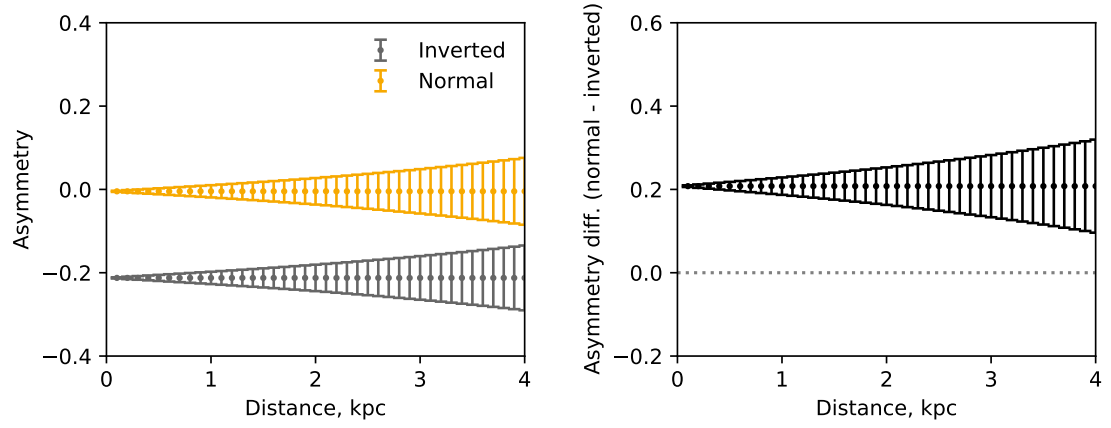


Figure 9.10: Calculated asymmetries for the normal and inverted mass orderings (left) and their difference (right) as a function of supernova distance.




Figure 9.11: Significance of the asymmetry difference between normal and inverted mass orderings versus distance for the LS220-s27.0co and SFHo-z9.6co models.

robustness of analysis like the one presented here.

Until then, the development of this and other supernova sensitivity studies at NOvA can add levels of sophistication to our assumptions and develop entirely new analysis pathways. For example, while it was a wish-list item for this thesis, I have not attempted to reconstruct the supernova neutrino energy spectrum, upon which mass ordering signatures and other interesting phenomena will be imprinted. Candidates clusters with hit multiplicities greater than three of four should allow for a decent estimate of the calibrated visible energy, and studying the many Michel electrons in NOvA's cosmic data can provide a good handle on the energy resolution and overall energy scaling in the MeV regime.

For now, the results presented here demonstrate that, despite NOvA's relatively large exposure to cosmic rays, it is a capable supernova observatory. This work also creates a foundation upon which additional sensitivity studies can be built and that a future analysis can draw from in the event of a real galactic core-collapse supernova.



## CHAPTER 10

# Summary

The next galactic core-collapse supernova will be an incredible opportunity to advance humanity's understanding of supernovae, neutrinos, and nuclear physics. Every capable neutrino detector must be ready for this once-in-a-career event which could come at any time. The Supernova Early Warning system connects a global network of supernova-sensitive detectors together with the goal of maximizing that readiness.

The NOvA detectors are standing by and will read out 45 s of continuous data from its buffers at both detectors in the event of a supernova. There are multiple avenues to triggering this readout, including our internal data-driven trigger, a SNEWS alert, or even a LIGO alert. The last several years have seen major developments in these capabilities, and soon NOvA will not only be a consumer of these alerts, but a contributor.

Both of NOvA's detectors sit closer to the surface than any other current or planned supernova-capable neutrino detectors, especially the far detector which is only partially underground and subjected to an onslaught of cosmic particles. NOvA's coarse spatial resolution is fine for reconstructing GeV neutrinos from an accelerator beam, but it presents an interesting challenge for identifying the MeV-scale supernova neutrino depositions.

Despite these challenges, I have shown in this thesis that with the right collection of background rejection techniques and machine-learning-assisted event selection, NOvA is capable of identifying a supernova event by performing a joint-detector fit of a flux model to a simulated supernova. For one particular model, a  $27 M_{\odot}$  progenitor using the LS220 nuclear equation of state, this method was capable of identifying the supernova at the  $3\sigma$  significance level out to 16.6 kpc.

Additionally, I showed that NOvA may have sensitivity to the mass ordering under the assumption that MSW effects in the star are appreciable during the first second of the explosion and that no other flavor transformations take place. In this simple scenario, a joint-detector fit of the time profile under both ordering hypotheses reveals a  $3\sigma$  sensitivity out to 1.51 kpc under the normal hypothesis and 1.06 kpc for the inverted case. A second method which considered the asymmetry in signal counts between two adjacent time windows yielded a  $3\sigma$  sensitivity at 2.8 kpc.



# Bibliography

- [1] Yoshiyuki Fukuda, T Hayakawa, E Ichihara, K Inoue, K Ishihara, Hirokazu Ishino, Y Itow, T Kajita, J Kameda, S Kasuga, et al. Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, 81(8):1562, 1998.
- [2] Q Retal Ahmad, RC Allen, TC Andersen, JD Anglin, JC Barton, EW Beier, M Bercovitch, J Bigu, SD Biller, RA Black, et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory. *Physical review letters*, 89(1):011301, 2002.
- [3] G. Gamow and M. Schoenberg. Neutrino Theory of Stellar Collapse. *Phys. Rev.*, 59(7):539, 1941.
- [4] Stirling A. Colgate and Richard H. White. The Hydrodynamic Behavior of Supernovae Explosions. *Astrophys. J.*, 143:626, 1966.
- [5] K. Hirata, T. Kajita, and M. Koshiba. Observation of a neutrino burst from the supernova SN1987A. *Physical Review ...*, 58(14):1490–1493, April 1987.
- [6] R. Bionta, G. Blewitt, C. Bratton, D. Casper, A. Ciocio, R. Claus, B. Cortez, M. Crouch, S. Dye, S. Errede, G. Foster, W. Gajewski, K. Ganezer, M. Goldhaber, T. Haines, T. Jones, D. Kielczewska, W. Kropp, J. Learned, J. LoSecco, J. Matthews, R. Miller, M. Mudan, H. Park, L. Price, F. Reines, J. Schultz, S. Seidel, E. Shumard, D. Sinclair, H. Sobel, J. Stone, L. Sulak, R. Svoboda, G. Thornton, J. van der Velde, and C. Wuest. Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud. *Physical Review Letters*, 58(14):1494–1496, April 1987.
- [7] E. N. Alekseev, L. N. Alekseeva, V. I. Volchenko, and I. V. Krivosheina. Possible detection of a neutrino signal on February 23, 1987 with the Baksan underground scintillation telescope of the Nuclear Research Institute of the Soviet Academy of Sciences. *Pisma v Zhurnal Eksperimentalnoi i Teoreticheskoi Fiziki*, 45:461–464, May 1987.
- [8] S Al Kharusi, SY BenZvi, JS Bobowski, W Bonivento, V Brdar, T Brunner, E Caden, M Clark, A Coleiro, M Colomer-Molla, et al. SNEWS 2.0: a next-generation supernova early warning system for multi-messenger astronomy. *New Journal of Physics*, 23(3):031201, 2021.
- [9] A H Bécquerel. Sur les radiations invisibles emises par les corps phosphorescents. *C. R. Acad. Sci. Paris*, 122:501, 1896.

- [10] J. Chadwick. Intensitätsverteilung im magnetischen Spectrum der  $\beta$ -Strahlen von radium B + C. *Verhandl. Dtsch. Phys. Ges.*, 16:383, 1914.
- [11] J. Chadwick. private communication, 1914.
- [12] D. E. Ellis and W. A. Wooster. The average energy of disintegration of radium E. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 117(776):109–123, December 1927.
- [13] Wolfgang Pauli. Pauli letter collection: letter to Niels Bohr. Handwritten copy. Transcript available, March 1929.
- [14] Wolfgang Pauli. Pauli letter collection: letter to Niels Bohr. Handwritten copy. Transcript available, July 1929.
- [15] W. Pauli. Dear radioactive ladies and gentlemen. *Phys. Today*, 31N9:27, 1978.
- [16] Enrico Fermi. An attempt of a theory of beta radiation. *Z. phys*, 88(161):19, 1934.
- [17] Hans Bethe and Rudolph Peierls. The neutrino. *Nature*, 133(3366):689–690, 1934.
- [18] ME Nahmias. An attempt to detect the neutrino. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 31, pages 99–107. Cambridge University Press, 1935.
- [19] C. Jarlskog. Birth of the neutrino, from Pauli to the Reines-Cowan experiment. In *History of the Neutrino 2018 conference proceedings*, sep 2018.
- [20] B. Pontecorvo. Inverse  $\beta$  Process. Technical report, National Research Council of Canada, Division of Atomic Energy, Chalk River, nov 1946. Report PD-205.
- [21] F. Reines and C. L. Cowan. Detection of the free neutrino. *Phys. Rev.*, 92:830–831, Nov 1953.
- [22] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: a confirmation. *Science*, 124(3212):103–104, 1956.
- [23] Gaillard Danby, Jean Maurice Gaillard, Konstantin Goulianos, Leon M Lederman, Nariman Mistry, Melvin Schwartz, and Jack Steinberger. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Physical Review Letters*, 9(1):36, 1962.
- [24] Martin L Perl, GS Abrams, AM Boyarski, Martin Breidenbach, DD Briggs, F Bulos, William Chinowsky, JT Dakin, GJ Feldman, CE Friedberg, et al. Evidence for anomalous lepton production in  $e^+e^-$  annihilation. *Physical Review Letters*, 35(22):1489, 1975.
- [25] K Kodama, N Ushida, C Andreopoulos, N Saoulidou, G Tzanakos, P Yager, B Baller, D Boehnlein, Walter Freeman, B Lundberg, et al. Observation of tau neutrino interactions. *Physics Letters B*, 504(3):218–224, 2001.

- [26] D Decamp, B Deschizeaux, J-P Lees, M-N Minard, JM Crespo, M Delfino, E Fernandez, M Martinez, R Miquel, ML Mir, et al. Determination of the number of light neutrino species. *Physics Letters B*, 231(4):519–529, 1989.
- [27] MZ Akrawy, G Alexander, J Allison, PP Allport, KJ Anderson, JC Armitage, GTJ Arnison, P Ashton, G Azuelos, JTM Baines, et al. Measurement of the  $Z^0$  Mass and Width with the OPAL Detector at LEP. *Physics Letters B*, 231(4):530–538, 1989.
- [28] P Aarnio, P Abreu, W Adam, P Adrianos, T Adye, G Akopdzhanov, G Alekseev, J Allaby, P Allen, P Allport, et al. Measurement of the mass and width of the  $Z^0$ -particle from multi-hadronic final states produced in  $e^+e^-$  annihilations. *Physics Letters B*, 231(4):539–547, 1989.
- [29] Hans Albrecht Bethe. Energy production in stars. *Physical Review*, 55(5):434, 1939.
- [30] John N Bahcall. Solar neutrinos. i. theoretical. *Physical Review Letters*, 12(11):300, 1964.
- [31] Raymond Davis Jr. Solar neutrinos. ii. experimental. *Physical Review Letters*, 12(11):303, 1964.
- [32] John N Bahcall, Neta A Bahcall, and Giora Shaviv. Present status of the theoretical predictions for the cl 37 solar-neutrino experiment. *Physical Review Letters*, 20(21):1209, 1968.
- [33] Raymond Davis Jr, Don S Harmer, and Kenneth C Hoffman. Search for neutrinos from the sun. *Physical Review Letters*, 20(21):1205, 1968.
- [34] John F Amsbaugh, JM Anaya, J Banar, TJ Bowles, MC Browne, TV Bullard, TH Burritt, GA Cox-Mobrand, X Dai, H Deng, et al. An array of low-background  $^3\text{He}$  proportional counters for the Sudbury Neutrino Observatory. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 579(3):1054–1080, 2007.
- [35] P.A. Zyla et al. Review of Particle Physics. *PTEP*, 2020(8):083C01, 2020.
- [36] Lincoln Wolfenstein. Neutrino oscillations in matter. *Physical Review D*, 17(9):2369, 1978.
- [37] SP Mikheev and A Iu Smirnov. Resonant amplification of  $\nu$  oscillations in matter and solar-neutrino spectroscopy. *Nuovo Cimento C Geophysics Space Physics C*, 9:17–26, 1986.
- [38] Glenn Ledrew. The real starry sky. *Journal of the Royal Astronomical Society of Canada*, 95:32, 2001.
- [39] W. Baade and F. Zwicky. On Super-Novae. *Proc. Nat. Acad. Sci.*, 20(5):254–259, 1934.
- [40] W. Baade and F. Zwicky. Cosmic rays from super-novae. *Proceedings of the National Academy of Sciences*, 20(5):259–263, 1934.
- [41] Hans-Thomas Janka. Neutrino-Driven Explosions. In Athem W. Alsabti and Paul Murdin, editors, *Handbook of Supernovae*, pages 1–56. Springer International Publishing, Cham, 2017.
- [42] Stan Woosley and H. T. Janka. The Physics of Core-Collapse Supernovae. *Nature Physics*, 1(3):17, January 2006.

- [43] W. Hillebrandt, M. Kromer, F. K. Röpkke, and A. J. Ruiter. Towards an understanding of Type Ia supernovae from a synthesis of theory and observations. *Front. Phys. (Beijing)*, 8:116–143, 2013.
- [44] Alessandro Mirizzi, G. G. Raffelt, and P. D. Serpico. Earth matter effects in supernova neutrinos: Optimal detector locations. *JCAP*, 05:012, 2006.
- [45] Mainak Mukhopadhyay, Cecilia Lunardini, F. X. Timmes, and Kai Zuber. Presupernova neutrinos: directional sensitivity and prospects for progenitor identification. *Astrophys. J.*, 899(2):153, 2020.
- [46] Karolina Rozwadowska, Francesco Vissani, and Enrico Cappellaro. On the rate of core collapse supernovae in the milky way. *New Astron.*, 83:101498, 2021.
- [47] Andrzej Odrzywolek and Alexander Heger. Neutrino signatures of dying massive stars: From main sequence to the neutron star. *Acta Phys. Pol. B*, 41:1611–28, 2010.
- [48] Adam Burrows. Perspectives on Core-Collapse Supernova Theory. *Reviews of Modern Physics*, 85(1), February 2013.
- [49] H. Th. Janka. Neutrino-driven Explosions. 2 2017.
- [50] H Janka, K Langanke, A Marek, G Martinezpinedo, and B Muller. Theory of core-collapse supernovae. *Physics Reports*, 442(1-6):38–74, April 2007.
- [51] H.-Th Janka. Neutrino Emission from Supernovae. *arXiv:1702.08713 [astro-ph, physics:hep-ph, physics:nucl-th]*, pages 1575–1604, 2017. arXiv: 1702.08713.
- [52] John M. Blondin, Anthony Mezzacappa, and Christine DeMarino. Stability of standing accretion shocks, with an eye toward core collapse supernovae. *Astrophys. J.*, 584:971–980, 2003.
- [53] M. Kachelrieß, R. Tomàs, R. Buras, H.-Th. Janka, A. Marek, and M. Rampp. Exploiting the neutronization burst of a galactic supernova. *Phys. Rev. D*, 71:063003, Mar 2005.
- [54] Kate Scholberg. Supernova Signatures of Neutrino Mass Ordering. *J. Phys. G*, 45(1):014002, 2018.
- [55] Alessandro Mirizzi, Irene Tamborra, Hans-Thomas Janka, Ninetta Saviano, Kate Scholberg, Robert Bollig, Lorenz Hudepohl, and Sovan Chakraborty. Supernova Neutrinos: Production, Oscillations and Detection. *La Rivista del Nuovo Cimento*, 39(102):1–112, February 2016. arXiv: 1508.00785.
- [56] Evan O’Connor, Robert Bollig, Adam Burrows, Sean Couch, Tobias Fischer, Hans-Thomas Janka, Kei Kotake, Eric J Lentz, Matthias Liebendörfer, OE Bronson Messer, et al. Global comparison of core-collapse supernova simulations in spherical symmetry. *Journal of Physics G: Nuclear and Particle Physics*, 45(10):104001, 2018.
- [57] Sean M Couch, MacKenzie L Warren, and Evan P O’Connor. Simulating Turbulence-aided Neutrino-driven Core-collapse Supernova Explosions in One Dimension. *The Astrophysical Journal*, 890(2):127, 2020.

- [58] Tobias Melson, Hans-Thomas Janka, and Andreas Marek. Neutrino-Driven Supernova of a Low-Mass Iron-Core Progenitor Boosted by Three-Dimensional Turbulent Convection. *The Astrophysical Journal*, 801(2):L24, March 2015.
- [59] J. Nordhaus, A. Burrows, A. Almgren, and J. Bell. Dimension as a Key to the Neutrino Mechanism of Core-Collapse Supernova Explosions. *Astrophys. J.*, 720:694–703, 2010.
- [60] Sjors Broersen, Alexandros Chiotellis, Jacco Vink, and Aya Bamba. The many sides of RCW 86: a Type Ia supernova remnant evolving in its progenitor’s wind bubble. *Monthly Notices of the Royal Astronomical Society*, 441(4):3040–3054, 2014.
- [61] Yuko Motizuki, Kazuya Takahashi, Kazuo Makishima, Aya Bamba, Yoichi Nakai, Yasushige Yano, Makoto Igarashi, Hideaki Motoyama, Kokichi Kamiyama, Keisuke Suzuki, et al. An Antarctic ice core recording both supernovae and solar cycles. *arXiv preprint arXiv:0902.3446*, 2009.
- [62] Alicia M Soderberg, E Berger, KL Page, P Schady, J Parrent, D Pooley, X-Y Wang, EO Ofek, A Cucchiara, A Rau, et al. An extremely luminous X-ray outburst at the birth of a supernova. *Nature*, 453(7194):469–474, 2008.
- [63] Subo Dong, BJ Shappee, JL Prieto, SW Jha, KZ Stanek, TW-S Holoien, CS Kochanek, TA Thompson, N Morrell, IB Thompson, et al. ASASSN-15lh: A highly super-luminous supernova. *Science*, 351(6270):257–260, 2016.
- [64] M. Keil, G. Raffelt, and H. Janka. Monte Carlo Study of Supernova Neutrino Spectra Formation. *The Astrophysical Journal*, 590(2):971–991, jun 2003.
- [65] Steen Hannestad, Georg G Raffelt, Günter Sigl, and Yvonne YY Wong. Self-induced conversion in dense neutrino gases: Pendulum in flavor space. *Physical Review D*, 74(10):105010, 2006.
- [66] Raymond F Sawyer. Multiangle instability in dense neutrino systems. *Physical Review D*, 79(10):105003, 2009.
- [67] Andreu Esteban-Pretel, Alessandro Mirizzi, Sergio Pastor, Ricard Tomàs, Georg G Raffelt, Pasquale D Serpico, and Günther Sigl. Role of dense matter in collective supernova neutrino transformations. *Physical Review D*, 78(8):085012, 2008.
- [68] Jörn Kersten and Alexei Yu. Smirnov. Decoherence and oscillations of supernova neutrinos. *The European Physical Journal C*, 76(6):339, June 2016.
- [69] A. Beck et al. SNOwGLOBES: SuperNova Observatories with GLOBES. <https://github.com/SNOwGLOBES/snowglobes>, 2016.
- [70] S Fukuda, Y Fukuda, T Hayakawa, E Ichihara, M Ishitsuka, Y Itow, T Kajita, J Kameda, K Kaneyuki, S Kasuga, et al. The Super-Kamiokande Detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 501(2-3):418–462, 2003.

- [71] Hyper-Kamiokande Proto-Collaboration. Hyper-Kamiokande Design Report, 2018.
- [72] Mark G Aartsen, M Ackermann, J Adams, JA Aguilar, M Ahlers, M Ahrens, D Altmann, K Andeen, T Anderson, I Ansseau, et al. The IceCube Neutrino Observatory: instrumentation and online systems. *Journal of Instrumentation*, 12(03):P03012, 2017.
- [73] Silvia Adrian-Martinez, M Ageron, F Aharonian, S Aiello, A Albert, F Ameli, E Anassontzis, M Andre, G Androulakis, M Anghinolfi, et al. Letter of intent for KM3NeT 2.0. *Journal of Physics G: Nuclear and Particle Physics*, 43(8):084001, 2016.
- [74] O Suvorova et al. Multi-messenger and real-time astrophysics with the Baikal-GVD telescope. In *Proceedings of 37th International Cosmic Ray Conference—PoS (ICRC2021)*, volume 395, page 946, 2021.
- [75] VB Petkov, AN Fazliakhmetov, AM Gangapshev, VN Gavrin, TV Ibragimova, MM Kochkarov, VV Kazalov, D Yu Kudrin, VV Kuzminov, BK Lubsandorzhev, et al. Baksan large volume scintillation telescope: a current status. In *Journal of Physics: Conference Series*, volume 1468, page 012244. IOP Publishing, 2020.
- [76] M Aglietta, B Alpat, ED Alyea, P Antonioli, G Anzivino, G Badino, Y Ban, G Bari, M Basile, A Benelli, et al. The most powerful scintillator supernovae detector: LVD. *Il Nuovo Cimento A (1965-1970)*, 105(12):1793–1804, 1992.
- [77] G Alimonti, C Arpesella, H Back, M Balata, D Bartolomei, A De Bellefon, G Bellini, J Benziger, A Bevilacqua, D Bondi, et al. The Borexino Detector at the Laboratori Nazionali del Gran Sasso. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 600(3):568–593, 2009.
- [78] S Andringa, E Arushanova, S Asahi, M Askins, DJ Auty, AR Back, Z Barnard, N Barros, EW Beier, A Bialek, et al. Current status and future prospects of the SNO+ experiment. *Advances in High Energy Physics*, 2016, 2016.
- [79] Fengpeng An, Guangpeng An, Qi An, Vito Antonelli, Eric Baussan, John Beacom, Leonid Bezrukov, Simon Blyth, Riccardo Brugnera, Margherita Buizza Avanzini, et al. Neutrino physics with JUNO. *Journal of Physics G: Nuclear and Particle Physics*, 43(3):030401, 2016.
- [80] R Acciarri, C Adams, R An, A Aparicio, S Aponte, J Asaadi, M Auger, N Ayoub, L Bagby, B Baller, et al. Design and construction of the MicroBooNE detector. *Journal of Instrumentation*, 12(02):P02017, 2017.
- [81] R. Acciarri et al. A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam, 2015.
- [82] The DUNE Collaboration. Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics, 2020.



- [83] Rafael F Lang, Christopher McCabe, Shayne Reichard, Marco Selvi, and Irene Tamborra. Supernova neutrino physics with xenon dark matter detectors: A timely perspective. *Physical Review D*, 94(10):103009, 2016.
- [84] Taylor C. Shantz. Design and construction of the Helium and Lead Observatory for supernova neutrinos. Master’s thesis, Laurentian University (Canada), December 2010.
- [85] Justin A. Vasel. Under the Hood: Preparing the Helium and Lead Observatory for Full Operation. Master’s thesis, University of Minnesota, January 2013.
- [86] Barry Pointon. HALO-1kT-Helium And Lead Observatory for Supernova Neutrinos with High Sensitivity to  $\nu_e$ . In *XXVIII International Conference on Neutrino Physics and Astrophysics*, page 474, 2018.
- [87] Akif Baha Balantekin, J. Gava, and C. Volpe. Possible CP-Violation effects in core-collapse Supernovae. *Phys. Lett. B*, 662:396–404, 2008.
- [88] David N Schramm and James W Truran. New physics from supernova 1987a. *Physics Reports*, 189(2):89–126, 1990.
- [89] Max Aker, K Altenmüller, M Arenz, M Babutzka, J Barrett, S Bauer, M Beck, A Beglarian, J Behrens, T Bergmann, et al. Improved upper limit on the neutrino mass from a direct kinematic method by katrin. *Physical review letters*, 123(22):221802, 2019.
- [90] Rasmus SL Hansen, Manfred Lindner, and Oliver Scholer. Timing the neutrino signal of a Galactic supernova. *Physical Review D*, 101(12):123018, 2020.
- [91] Pietro Antonioli, Richard Tresch Fienberg, Fabrice Fleurot, Yoshiyuki Fukuda, Walter Fulgione, Alec Habig, Jaret Heise, Arthur B McDonald, Corrinne Mills, Toshio Namba, et al. SNEWS: the supernova early warning system. *New Journal of Physics*, 6(1):114, 2004.
- [92] K Asakura, A Gando, Y Gando, T Hachiya, S Hayashida, H Ikeda, K Inoue, K Ishidoshiro, T Ishikawa, S Ishio, et al. KamLAND sensitivity to neutrinos from pre-supernova stars. *The Astrophysical Journal*, 818(1):91, 2016.
- [93] C. Simpson et al. Sensitivity of Super-Kamiokande with Gadolinium to Low Energy Anti-neutrinos from Pre-supernova Emission. *Astrophys. J.*, 885:133, 2019.
- [94] Nirmal Raj, Volodymyr Takhistov, and Samuel J Witte. Presupernova neutrinos in large dark matter direct detection experiments. *Physical Review D*, 101(4):043008, 2020.
- [95] D. Akimov et al. Observation of Coherent Elastic Neutrino-Nucleus Scattering. *Science*, 357(6356):1123–1126, 2017.
- [96] K Abe, Y Haga, Y Hayato, M Ikeda, K Iyogi, J Kameda, Y Kishimoto, M Miura, S Moriyama, M Nakahata, et al. Real-time supernova neutrino burst monitor at Super-Kamiokande. *Astroparticle Physics*, 81:39–48, 2016.

- [97] Vedran Brdar, Manfred Lindner, and Xun-Jie Xu. Neutrino astronomy with supernova neutrinos. *Journal of Cosmology and Astroparticle Physics*, 2018(04):025, 2018.
- [98] D S Ayres, G R Drake, M C Goodman, J J Grudzinski, V J Guarino, R L Talaga, A Zhao, /Argonne, P Stamoulis, E Stiliaris, G Tzanakos, M Zois, and Duluth /Minnesota U. /Athens U. /Caltech /UCLA /Fermilab /College de France /Harvard U. /Indiana U. /Lebedev Inst. /Michigan State U. /Minnesota U. The NOvA Technical Design Report. 10 2007.
- [99] NOvA Collaboration. First measurement of muon-neutrino disappearance in NOvA. *Physical Review D*, 93(5):051104, March 2016.
- [100] NOvA Collaboration. Measurement of the Neutrino Mixing Angle  $\theta_{23}$  in NOvA. *Physical Review Letters*, 118(15):151802, April 2017.
- [101] NOvA Collaboration. Constraints on Oscillation Parameters from  $\nu_e$  Appearance and  $\nu_\mu$  Disappearance in NOvA. *Physical Review Letters*, 118(23):231801, June 2017.
- [102] NOvA Collaboration. New constraints on oscillation parameters from  $\nu_e$  appearance and  $\nu_\mu$  disappearance in the NOvA experiment. *Physical Review D*, 98(3):032012, August 2018.
- [103] NOvA Collaboration. First measurement of neutrino oscillation parameters using neutrinos and antineutrinos by NOvA. *Physical Review Letters*, 123(15):151803, October 2019.
- [104] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M.D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle. A convolutional neural network neutrino event classifier. *Journal of Instrumentation*, 11(09):P09001–P09001, sep 2016.
- [105] Fernanda Psihas. *Measurement of Long Baseline Neutrino Oscillations and Improvements from Deep Learning*. PhD thesis, Indiana U., 2018.
- [106] F. Psihas, E. Niner, M. Groh, R. Murphy, A. Aurisano, A. Himmel, K. Lang, M. D. Messier, A. Radovic, and A. Sousa. Context-Enriched Identification of Particles with a Convolutional Network for Neutrino Events. *Phys. Rev. D*, 100(7):073005, 2019.
- [107] Pierre Baldi, Jianming Bian, Lars Hertel, and Lingge Li. Improved Energy Reconstruction in NOvA with Regression Convolutional Neural Networks. *Phys. Rev. D*, 99(1):012011, 2019.
- [108] NOvA Collaboration. Search for active-sterile neutrino mixing using neutral-current interactions in NOvA. *Physical Review D*, 96(7):072006, October 2017.
- [109] NOvA Collaboration. Measurement of neutrino-induced neutral-current coherent  $\pi^0$  production in the NOvA near detector. *Physical Review D*, 102(1):012004, July 2020.
- [110] NOvA Collaboration. Adjusting neutrino interaction models and evaluating uncertainties using NOvA near detector data. *The European Physical Journal C*, 80(12):1119, December 2020.
- [111] NOvA Collaboration. Observation of seasonal variation of atmospheric multiple-muon events in the NOvA Near Detector. *Physical Review D*, 99(12):122004, June 2019.

- [112] NOvA Collaboration. Search for slow magnetic monopoles with the NOvA detector on the surface. *Physical Review D*, 103(1):012007, January 2021.
- [113] NOvA Collaboration. Search for multimessenger signals in NOvA coincident with LIGO/Virgo detections. *Physical Review D*, 101(11):112006, June 2020.
- [114] NOvA Collaboration. Supernova neutrino detection in NOvA. *Journal of Cosmology and Astroparticle Physics*, 2020(10):014–014, October 2020.
- [115] Michael Wallbank. The NOvA Test Beam Program. *PoS, ICHEP2020*:188, 2021.
- [116] S. Mufson, B. Baugh, C. Bower, T. E. Coan, J. Cooper, L. Corwin, J. A. Karty, P. Mason, M. D. Messier, A. Pla-Dalmau, and M. Proudfoot. Liquid scintillator production for the NOvA experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 799:1–9, November 2015. arXiv: 1504.04035.
- [117] Evan David Niner. *Observation of Electron Neutrino Appearance in the NuMI Beam with the NOvA Experiment*. PhD thesis, January 2015.
- [118] P Vogel and John F Beacom. Angular distribution of neutron inverse beta decay,  $\nu e + p \rightarrow e++ n$ . *Physical Review D*, 60(5):053003, 1999.
- [119] B Armbruster, I Blair, BA Bodmann, NE Booth, G Drexlin, V Eberhard, JA Edgington, C Eichner, K Eitel, E Finckh, et al. Measurement of the weak neutral current excitation  $^{12}\text{C}(\nu\mu\nu\mu')$   $^{12}\text{C}[(1+, 1; 15.1 \text{ MeV})$  at  $E\nu\mu = 29.8 \text{ MeV}$ . *Physics letters B*, 423(1-2):15–20, 1998.
- [120] James M Lattimer and F Douglas Swesty. A generalized equation of state for hot, dense matter. *Nuclear Physics A*, 535(2):331–376, 1991.
- [121] Andrew W Steiner, Matthias Hempel, and Tobias Fischer. Core-collapse supernova equations of state based on neutron star observations. *The Astrophysical Journal*, 774(1):17, 2013.
- [122] M. A. Acero et al. Supernova neutrino detection in NOvA. *JCAP*, 10:014, 2020.
- [123] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. *Nucl. Instrum. Meth.*, A614:87–104, 2010.
- [124] Sea Agostinelli, John Allison, K al Amako, John Apostolakis, H Araujo, Pedro Arce, Makoto Asai, D Axen, Swagato Banerjee, GJNI Barrant, et al. GEANT4—a simulation toolkit. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.
- [125] M Baird, Jianming Bian, M Messier, E Niner, D Rocco, and K Sachdev. Event reconstruction techniques in NOvA. In *Journal of Physics: Conference Series*, volume 664, page 072035. IOP Publishing, 2015.
- [126] C. Green, J. Kowalkowski, M. Paterno, M. Fischler, L. Garren, and Q. Lu. The Art Framework. *J. Phys. Conf. Ser.*, 396:022020, 2012.

- [127] Bruce Lynn Howard Jr. *Toward a Precision Era of Neutrino Oscillation Physics: Liquid Argon Scintillation Detector Development for DUNE and Neutrino Oscillation Studies with NOvA*. PhD thesis, Indiana University, 2019.
- [128] Michael David Baird. *An analysis of muon neutrino disappearance from the NuMI beam using an optimal track fitter*. PhD thesis, Indiana University, 2015.
- [129] Martin Ester, Hans-Peter Kriegel, and Xiaowei Xu. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. *Kdd*, 96(34):226–231, 1996.
- [130] Fons Rademakers et al. ROOT - An Object Oriented Data Analysis Framework (v6-20-06), June 2020.
- [131] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11), 2008.
- [132] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [133] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- [134] Markus Svensén and Christopher M Bishop. *Pattern recognition and machine learning*, 2007.
- [135] Harry Zhang. The optimality of naive Bayes. *AA*, 1(2):3, 2004.
- [136] Thomas Bayes. LII. An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, FRS communicated by Mr. Price, in a letter to John Canton, AMFR S. *Philosophical transactions of the Royal Society of London*, (53):370–418, 1763.
- [137] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, 16:321–357, Jun 2002.
- [138] Guillaume Lemaitre, Fernando Nogueira, and Christos K. Aridas. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *Journal of Machine Learning Research*, 18(17):1–5, 2017.
- [139] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [140] Michael JD Powell. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The computer journal*, 7(2):155–162, 1964.
- [141] Isaac Asimov. Franchise. *If: Worlds of Science Fiction*, Aug 1955.
- [142] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur.Phys.J.C*, 71:1554, 2011. [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)].

- [143] Matt Newville, Renee Otten, Andrew Nelson, Antonino Ingargiola, Till Stensitzki, Dan Allan, Austin Fox, Faustin Carter, Michał, Dima Pustakhod, Ineuhaus, Sebastian Weigand, Ray Osborn, Glenn, Christoph Deil, Mark, Allan L. R. Hansen, Gustavo Pasquevich, Leon Foks, Nicholas Zobrist, Oliver Frost, Alexandre Beelen, Stuermer, kwertyops, Anthony Polloreno, Shane Caldwell, Anthony Almarza, Arun Persaud, Ben Gamari, and Benjamin F. Maier. `lmfit/lmfit-py` 1.0.2, February 2021.