3D Photoelectron Velocity Map Imaging and Four-Wave Mixing of Cylindrical Vector Modes

Edward Scott Goudreau

Thesis submitted to the University of Ottawa in partial fulfillment of the requirements for the degree of Doctor of Philosophy, Physics

Department of Physics
Faculty of Science
University of Ottawa

Supervisor: Dr. Albert Stolow

Dedication

To my parents, Sandy and Terry Goudreau,
for many years of love, support, and encouragement.

Abstract

The 2D photoelectron velocity map imaging (VMI) technique is commonly employed in gas-phase molecular spectroscopy and dynamics investigations due to its ability to efficiently extract photoelectron spectra and angular distributions in a single experiment. However, the standard technique is restricted to specific light-source polarization geometries by the need to perform a mathematical inversion of the measured 2D detector-plane projections in order to recover the spherical 3D particle distribution. This has led to significant interest in the development of 3D VMI techniques which are capable of measuring, at a detector, the transverse position (x, y) and time-of-flight (TOF, t) of individual events in order to obtain a full set of 3D coordinates, thus avoiding the need for inversion and the associated constraints. These techniques employ curved velocity-mapping electric field lines, making the general transformation of (x, y, t)-data into initial 3D recoil momentum vectors (p_x, p_y, p_z) a challenging problem which, until now, was not fully addressed. Here I present and demonstrate a novel time-stretched, 13-lens 3D VMI photoelectron spectrometer which has sub-camera-pixel spatial resolution and 72 ps (σ) TOF resolution. This instrument employs a kHz CMOS camera to image a standard 40 mm diameter microchannel plate (MCP)/phosphor anode detector (providing x and y positions), combined with a digitizer pickoff from the phosphor to obtain the electron TOF. This thesis contains my work on testing and evaluating the performance of this spectrometer, as well as developing a complete data processing and analysis protocol to convert raw 3D VMI data (camera images and digitizer waveforms) into 3D charged particle recoil momentum vectors. I demonstrate the advantages of the 13-element design, showing that the greater spread in electron TOF permits an accurate time- and position-stamping of up to six electrons per laser shot at a 1 kHz repetition rate.

In a second project, I develop a theoretical description of the nonlinear optical process of four-wave mixing (FWM) as it applies to a type of structured light mode called cylindrical vector (CV) beams. The CV modes are eigenmodes of optical fibre and, as such, they have a broad range of application, such as telecommunications, quantum cryptography, and fundamental optics research. Despite this, their nonlinear optical properties are not yet well understood. Here I derive the selection rules which determine the allowed FWM processes involving CV modes in optical fibre.

Acknowledgements

Firstly, I would like to thank my current supervisor, Dr. Albert Stolow, for his expert guidance, advice, support, and encouragement throughout my time in Ottawa. I would also like to thank Dr. Jeff Lundeen, who was my supervisor for the first year of my PhD program, for assisting me through the process of getting into the University of Ottawa, for the many things I learned from him during my time in his group, and for his understanding and acceptance of my reasons for switching groups.

I thank Dr. Andrey Boguslavskiy for his constant help and instruction in the operation of the lab equipment, collection of data, and data processing and analysis. Thanks to Douglas Moffatt, for writing the instrument interfacing and data acquisition code used with the NRC 3D VMI spectrometer, for continuing to maintain, diagnose, and repair this code throughout the many changes and updates we made over the years, and for his help in signal analysis and data processing. Thanks to Dr. Aurelien Sanchez for collaboration on the development of the momentum vector reconstruction technique, polarization misalignment correction, and MCP gain correction discussed in Chapter 4.

Further thanks to Dr. Iain Wilkinson, for the design of the spectrometer used for most of the work contained herein; Dr. Rune Lausten and Denis Guay, for maintaining the lasers and other lab equipment; Dr. Connor Kupchak, for continuing and completing my work on four-wave mixing of cylindrical vector beams after I left the Lundeen group; and Michael Hemsworth, for his work on the design of the new 3D VMI spectrometer being constructed at the University of Ottawa, much of which is reproduced in Chapter 2. Additionally, I would like to thank all of the other co-authors on the three publications [1–3] from which much of this thesis is derived.

Finally, I would like to thank my friends and family for their continuing support throughout my life and education.

Author Contributions

I, Edward Scott Goudreau, state that the work contained herein is either solely my own or was done in collaboration between myself and others, except where explicitly mentioned. In the cases where the work was collaborative, I have done my best to give credit to the collaborators and make clear my role in the work. The work that makes up this thesis has been published in three peer-reviewed journal articles [1–3].

In particular, data acquisition (Section 3.6) was performed with direct assistance from Dr. Andrey Boguslavskiy and Douglas Moffatt; the development of the momentum vector reconstruction (Section 4.4), polarization misalignment correction (Section 4.5), and the MCP gain correction (Section 4.7) was done in collaboration with Dr. Aurelien Sanchez; and part of the four-wave mixing selection rule derivation (Chapter 5) was done in collaboration with Dr. Connor Kupchak.

Section 2.5 on the design of the new 3D VMI spectrometer being constructed at the University of Ottawa was mostly adapted from the University of Ottawa Master's thesis of Michael Hemsworth [4].

Contents

List of Figures

List of Tables

List of Abbreviations

ACT Absolute Correlation Threshold

ADC Analog-to-Digital Converter

ADS Artificial Data Set

AST Absolute Sensitivity Threshold

CCD Charge-Coupled Device

CFD Constant Fraction Discriminator

CMOS Complementary Metal-Oxide-Semiconductor

COLTRIMS Cold Target Recoil Ion Momentum Spectroscopy

CV Cylindrical Vector

EVT Emergent Vision Technologies

FFT Fast Fourier Transform

FWHM Full Width at Half Maximum

FWM Four-Wave Mixing

GSR Global Single-hit Response

ID Internal Diameter

LHS Left-Hand Side

LSR Local Single-hit Response

MCP Microchannel Plate

MFPAD Molecular-Frame Photoelectron Angular Distribution

NRC National Research Council

OAM Orbital Angular Momentum

PAD Photoelectron Angular Distribution

PEPICO Photoelectron-Photoion Coincidence

PES Photoelectron Spectrum

PHD Pulse Height Distribution

PMT Photomultiplier Tube

QE Quantum Efficiency

RCT Relative Correlation Threshold

RHS Right-Hand Side

ROI Region of Interest

RSS Residual Sum of Squares

SAM Spin Angular Momentum

SFM Sum Frequency Mixing

SFWM Spontaneous Four-Wave Mixing

SHAD Single-Hit Amplitude Distribution

SHG Second Harmonic Generation

SNR Signal-to-Noise Ratio

SPM Self-Phase Modulation

TAM Total Angular Momentum

THG Third Harmonic Generation

TOF Time-Of-Flight

TTS Turn-around Time Spread

VMI Velocity Map Imaging

XPM Cross-Phase Modulation

1

Introduction

This thesis contains my work on two separate projects, both of which were performed as part of my doctoral degree program at the University of Ottawa. The first of these was on 3D velocity map imaging (3D VMI), a new charged particle imaging technique which circumvents a major limitation that greatly restricts polarization geometries in the more conventional 2D VMI method. In this project, I worked with a 3D VMI spectrometer located at the National Research Council (NRC) lab (100 Sussex Drive, Ottawa, Ontario), henceforth called the NRC spectrometer. This instrument was designed for 3D VMI capability, specifically for use in photoelectron imaging, although it was not used for this purpose prior to my involvement. I worked on the testing and calibration of this instrument, evaluation of its performance, and the development of novel data processing and analysis techniques for 3D VMI data. Time-resolved photoelectron VMI is a powerful

technique for isolating and studying ultrafast non-adiabatic molecular processes such as photodissociation, internal conversion, intersystem crossing, and charge transfer, many of which are of significant chemical and biological interest. This is done by projecting the evolving state of a photoexcited molecule onto the energy and angular distribution of an emitted photoelectron.

The second project is my theoretical work on the four-wave mixing (FWM) of cylindrical vector (CV) modes in optical fibre. FWM is one of many nonlinear processes which can cause optical pulses to mix, converting between different wavelengths and modes. CV modes are a type of structured light which are also eigenmodes of optical fibre and are thus of significant interest in photonics and telecommunications. I derived a set of rules which determine the allowed FWM processes which may occur between CV modes in fibre. The discussion of the FWM project is contained within Chapter 5 which has its own introduction; as such, I will refrain from giving additional background information pertaining to that project in this chapter.

The 3D VMI project was performed in the group of my current supervisor, Dr. Albert Stolow. The FWM project was performed in the group of Dr. Jeff Lundeen, who was my supervisor at the time, although I transferred into Dr. Stolow's group partway through my program.

1.1 Motivation for 3D Velocity Map Imaging

Charged particle imaging techniques, such as VMI [5] and COLTRIMS [6] (COLd Target Recoil Ion Momentum Spectroscopy), are important tools for making energy- and angle-resolved measurements of atomic and molecular processes, including dissociation, scattering, and ionization [7–14]. These "half-scattering" processes produce outgoing spherical ion/electron momentum distributions on energy shells (Newton spheres) which are guided toward a position-sensitive detector using electric fields. VMI, the primary technique discussed in this thesis, was devised by Eppink and Parker [5] as an evolution of earlier ion imaging methods [15]. It uses a non-uniform electric guiding field (an electrostatic lens) to focus particles with the same velocity vector to the same point on the detector, even when the particles are born from different origin points within an extended ionization volume. This minimizes the spatial blurring which would otherwise occur and greatly improves image resolution.

In a conventional VMI experiment, a laser-induced process creates an ion or electron Newton sphere which is projected onto the 2D detector. Long-exposure integrated images are acquired over a large number of laser shots, then transformed using a mathematical inversion algorithm such as Abel inversion to recover the full 3D distribution from its 2D projection. Such inversion algorithms require the Newton sphere to possess cylindrical symmetry about an axis in the plane of the detector and this, in turn, restricts the laser polarization geometry. However, there are important processes and techniques which pro-

duce Newton spheres lacking cylindrical symmetry, including circular dichroism [10, 16] and pump-probe experiments using crossed or magic-angle laser polarization geometries. The importance of 3D, non-inversion-reliant VMI for the case of photoelectron elliptical dichroism was recently noted [17].

Various approaches have been developed to address the polarization geometry limitation of the conventional inversion-reliant VMI design. These techniques for recovering the full three-dimensional Newton sphere, collectively called "3D VMI", are summarized in a recent review article [18]. For ion imaging, where the difference in arrival times between the front and back of the Newton spheres can be in the hundreds of nanoseconds, slice imaging can be employed through the use of fast electronic gating [19–21]. It is much more challenging to apply this method to photoelectron VMI, where the arrival time spread across the Newton sphere is often less than 10 ns.

One direct way to reconstruct the Newton spheres is by measuring, for each particle, its 3D coordinates, typically transverse position (x, y) and time-of-flight (TOF, t) at a detector. This can be done with the use of delay-line anode detectors [22], which allow accurate measurement of both position and arrival time for individual events. Delay-line anode detectors are limited in their ability to handle multiple events per laser shot [23], which in turn impedes the speed of data collection, particularly for electron imaging, owing to the aforementioned narrow arrival time spread. While delay-line anode detectors have been successfully applied to 3D electron VMI [14, 22], alternatives have also emerged which use a fast camera to image the spatial position of events. However, since conventional CCD/CMOS cameras lack the needed sub-nanosecond time resolution, the time-stamping

of events must be obtained by other means.

One such camera-based 3D VMI technique, developed by Li and co-workers, involves measuring spatial positions using a microchannel plate (MCP) detector (with phosphor screen) and fast CMOS camera while simultaneously time-stamping each event using either a photomultiplier tube (PMT) [24] or a capacitive pickoff from the MCP detector to a fast digitizer [25–30]. The latter camera-and-digitizer variation was shown to have sufficient time resolution for 3D electron VMI [25] and has been used successfully for electron-ion [27, 29] and electron-electron [26, 28] coincidence measurements. It has also been shown to have high multi-hit capability, time-stamping up to 27 ions per laser shot [30]. This technique benefits from the ability to make use of sufficiently fast CMOS cameras, reducing cost and enabling the freedom to customize the choice of camera to suit given experimental requirements. A similar approach was also developed by Picard and co-workers [31].

Recently, 3D VMI implementations [32–35] using the Timepix [36] series of event-driven camera sensors have shown promise as a way to circumvent the need for separate position and time measurements by time-stamping each event as it occurs. This approach is currently limited by the $\sim$250 ps time resolution and 256×256 pixel spatial resolution of the third generation Timepix sensor (TPX3CAM), although the recent fourth generation of Timepix [37] has a time resolution below $\sim$200 ps with a 448×512 pixel array. An improved time resolution has been achieved by combining the (comparatively) rough time stamps of the TPX3CAM sensor with more precise timing using a detector-pickoff and digitizer setup [33, 34]. Furthermore, using a coaxial 5-element electron lens and 6-element ion lens, 3D photoelectron VMI was combined with coincident (or covariant) ion detection,

allowing for high resolution photoelectron images to be correlated with their associated cation distributions [35].

Importantly, applying these charged particle imaging techniques to problems of physical and chemical interest requires transformation of the spatial and temporal coordinates (x, y, t) of the detected ion/electron events (the direct observables) into initial ion/electron recoil momentum vectors (p_x, p_y, p_z). In VMI, where curved focusing fields are used, this task is made more challenging for the momentum component along the time axis. A general method for reconstructing the initial momentum vectors in a 3D charged particle imaging experiment is essential in order to deconvolve the instrument-specific response and thus permit the comparison of experiment with theory [38].

1.2 Thesis Outline

In this thesis I make use of the NRC spectrometer, which was designed to build upon the camera-and-digitizer 3D VMI technique of Li *et al.* This spectrometer addresses one of the main challenges of 3D electron VMI: the short arrival time spread of the electron Newton sphere which limits momentum resolution along the TOF (z) axis. The NRC spectrometer uses a 13-element electrostatic lens (rather than the three-element lens of Eppink and Parker [5]) to allow the Newton sphere to expand significantly in time, stretching the arrival time distribution for better z-momentum resolution while also retaining good velocity mapping performance. This multi-electrode "thick lens" design concept originated in ion imaging [20, 39] and has since seen several successful applications [22, 35, 40–42].

In Chapter 2, I summarize aspects of conventional 2D VMI theory which are relevant to understanding the performance of the NRC spectrometer and the processing of VMI data, including a description of the velocity mapping fields, the relationship between image radius and energy resolution, and the particle arrival time distribution. I also describe how the photoelectron angular distribution (PAD) reveals properties of the parent atomic or molecular state in a photoionization experiment and how it can be represented in most cases by only a few asymmetry parameters. I include a description of the Abel transform and highlight a recent more advanced Abel-inversion algorithm [43] which I coded an implementation of in Python and have used to process 2D VMI data. I have also summarized the design considerations for a new 3D VMI spectrometer which improves on the NRC design and for which much of the analysis techniques described in this thesis will be later used.

In Chapter 3, I describe in detail the design of the current NRC spectrometer and the data acquisition setup used when performing 3D VMI experiments. I include the reasoning behind the camera and digitizer we have chosen to use with this instrument. I summarize the two laser setups used to collect 3D VMI data and the details of these data sets which I later use to characterize and evaluate the performance of the instrument. These data sets consist of photoionization experiments using atomic xenon, methyl iodide, and nitric oxide.

In Chapter 4, I describe in detail the steps required to process raw 3D VMI data in order to obtain initial ion/electron momentum vectors. This includes the process of assigning spatial positions (x, y) from the camera with their correlated time-stamps (t)

from the digitizer, which is a critical step when there are multiple hits per laser shot. The processing steps also include corrections for misalignment of the laser polarization vector and spatial inhomogeneity of the detector gain. I evaluated the performance of the NRC spectrometer and report here its multi-hit performance (up to six hits per laser shot), spatial resolution (0.2 pixels), and time resolution (< 80 ps). This chapter includes a general and novel methodology for recovering the initial momentum vectors, illustrated here for the NRC 3D photoelectron VMI instrument, which is applicable to any 3D charged particle imaging system (including those using curved velocity focusing fields) which time-stamps individual events. I conclude this chapter with an analysis of the 3D VMI data sets described in Chapter 3, including a comparison of the PADs with theoretical expectations and preexisting results and a demonstration of 3D imaging which would not be possible with conventional inversion-reliant 2D VMI.

Chapter 5 is a self-contained chapter covering a theoretical nonlinear optics project unrelated to my work on VMI. It details my work on developing a mathematical description of the process of FWM in optical fibre. Specifically, I derive the allowed FWM processes involving a type of structured light mode called cylindrical vector modes.

1.3 Photoelectron Spectroscopy

The design of the NRC spectrometer and my work on it was motivated primarily by the desire to investigate molecular dynamics using time-resolved photoelectron spectroscopy. Here I provide context for the benefits of VMI as an experimental technique for studying

molecular dynamics, explain how 3D VMI can reveal additional information not accessible with 2D imaging, and describe the types of experiments that are planned for this instrument and the similar University of Ottawa spectrometer in the future.

A general two-step scheme, called a pump-probe experiment, is often used to investigate the evolution of isolated molecular processes, typically in a gas-phase sample. The first step is an initial laser pulse (called the pump pulse) which prepares, in general, a superposition of excited states (called a wave packet) in an ensemble of sample molecules. The evolution of this wave packet represents the molecular processes of interest, for example: dissociation, internal conversion, or charge transfer. After a variable time delay Δt relative to the pump pulse, a second pulse (the probe pulse) ionizes the molecules in the ensemble. This projects information about the molecular wave packet, at that time delay Δt, onto the scattering state (energy and angular distribution) of the free photoelectron.

Photoelectron spectroscopy is an appealing technique for studying complex molecular wave packet dynamics using pump-probe experiments for a number of reasons. First, any state can be ionized, regardless of any symmetry or angular momentum constraints, so long as the probe laser photon energy exceeds the ionization potential. This ensures that the molecular dynamics can always be projected onto the photoelectron continuum at any time along the wave packet's evolution. Furthermore, the energy- and angle-resolved VMI measurement gives access to both the photoelectron spectrum (PES) and the PAD which provide complementary information about the wave packet dynamics. For example, the PAD can be used to monitor variations in molecular electronic symmetries, even between states which yield comparable photoelectron energies making such changes less discernible

from only the PES.

Typical pump-probe experiments involve a molecular ensemble which, without application of an alignment field, will be randomly oriented in space. This means that the PAD measured in the lab-frame will be averaged over all possible orientations of the molecule, resulting in a significant loss of information compared to the molecular-frame (MF) PAD. There has been significant recent interest in the development of techniques which can be used to access the MFPAD [14, 44–46]. The MFPAD is a much more information-rich and complex structure which does not in general possess the cylindrical symmetry required for 2D inversion-reliant VMI. As such, techniques capable of directly measuring the PAD, without inversion, such as 3D VMI, promise to be of significant use to this emerging field of MFPAD reconstruction.

2

Theoretical Background

2.1 Introduction

In this chapter, I discuss some of the theoretical concepts that are important to understand or contextualize my work on 3D VMI, as discussed later in this thesis. This includes a discussion of the theory underlying the conventional 2D VMI technique of Eppink and Parker [5], with a focus on the energy resolution and turn-around time spread of a 2D VMI spectrometer. Since this thesis primarily deals with photoelectron VMI, I discuss the nature of photoelectron angular distributions and how they can carry useful information about a molecular state prior to ionization. I describe the Abel transformation and its inversion, which is a necessary step for many 2D VMI experiments. In particular, I summarize the relatively new DAVIS algorithm [43] for performing the Abel inversion. I demonstrate

an implementation of this algorithm which I have written in Python and applied to the treatment of 2D VMI data. Finally, I summarize the design considerations which have lead to the currently under-construction next-generation 3D VMI spectrometer at the University of Ottawa. This section on the new design is mostly adapted from the work of Michael Hemsworth [4] and is included since this instrument will make extensive use of my work presented later in this thesis.

2.2 Velocity Map Imaging

The concept of VMI was devised in 1997 by Eppink and Parker [5] as an improvement to previous charged particle imaging techniques. This class of experimental technique seeks to measure the energy-, angle-, and (in the case of ion measurements) mass-resolved distribution of particles produced from an atomic or molecular sample undergoing a controlled ionization or fragmentation event. Particles produced in such an event will in general have kinetic energy, and thus a radial velocity directed outward from the origin point. For each energy, the spherical distribution of released particles (called a Newton sphere) will expand at a rate determined by the velocity. Newton spheres are captured by using electric fields to force the charged particles into a 2D imaging detector at the end of a vacuum-pumped flight tube.

Prior to VMI, such imaging instruments typically consisted of a single charged plate (called a repeller) on one end of the flight tube to direct the Newton sphere toward the detector on the other end. A second element, a fine mesh grid held at ground, was placed

between the particles' origin point and the detector to ensure a uniform field in this region while allowing transmission of the Newton spheres through the mesh. As a result of the uniform field, the radius r of the Newton sphere with initial particle release speed $|v|$ when it reaches the detector after a time-of-flight t is given by the simple relation $r = |v|t$.

This type of charged particle imaging presents two issues which are addressed by VMI. First, the ground mesh necessarily has a reduced transmission efficiency since some particles will either collide with the mesh or be deflected by it. Second, and most importantly, the particles being imaged are not all born from the same point origin; rather, they are generated from an extended region which has a shape typically determined by the size and orientation of the ionization volume. For UV laser-induced processes involving relatively long focal lengths ($f/\# \sim 75 - 100$), this interaction volume can be modelled as a cylinder along the propagation axis of the laser. Each point within this volume is the origin of its own Newton sphere. For a basic charged particle imaging system with a uniform guiding field, each Newton sphere is guided along a straight trajectory and arrives at the detector with a transverse offset equivalent to its origin within the interaction volume. That is, the image is blurred over the transverse area of the interaction volume. This effect encouraged experimental geometries which minimize the interaction volume, but this has the unwanted effect of reducing the yield of particles.

VMI addresses both of these issues by replacing the grid with an open ring-shaped ground electrode with an aperture for the Newton spheres to pass unimpeded. With this change, the field is no longer uniform, instead having some radial dependence from the central axis of the flight tube. However, by adding a third ring-shaped extractor electrode

between the interaction volume and the ground electrode, the non-uniform field can be adjusted to function as an electrostatic lens. Particles born away from the central axis of the flight tube now experience different guiding fields and no longer follow the same trajectory as those born on-axis. This is analogous to an optical lens – particles with the same initial velocity vector are focused (mapped) to the same spot within a focal plane; hence the name: velocity map imaging. By adjusting the ratio between the repeller and extractor voltages, V_R/V_E, one can move the focal plane to coincide with the detector plane. The magnification of the electrostatic lens is governed by the repeller voltage alone. Thus, VMI minimizes the effect of interaction region blurring, allowing experimentalists to make use of a larger interaction volume while still achieving better spatial resolution (and hence energy resolution) than was possible otherwise.

Due to the non-uniform fields in VMI, the simple relation $r = |v|t$ no longer holds. However, the proportionality $r \propto |v|t$ is still true to good approximation. With proper calibration for a given pair of voltages (V_R and V_E), the initial kinetic energy can be retrieved from the 2D detector-plane radius as described below in Section 2.2.2.

It is important to note that the focal length of the VMI electrostatic lens is dependent on the kinetic energy of the particles being imaged. Following the optical lens analogy, this equates to chromatic aberration. Typically, the focal length will be optimized for the maximum energy Newton spheres that can be imaged at a given magnification level. Newton spheres of any other energy will not be entirely in focus.

2.2.1 Definition of the Coordinate System

The coordinate system that will be used throughout this thesis (except where otherwise noted) is defined in Figure 2.1. The z and y axes are defined to be the axis normal to the detector face (the TOF axis) and the propagation axis of the laser ($\vec{k}$), respectively. A linear laser polarization is thus constrained to the XZ plane. For 2D inversion-reliant VMI, the laser must be polarized along the x axis to ensure the necessary cylindrical symmetry of the resulting particle angular distribution. The direction of the emission velocity vector ($\vec{v}$) for a particle is specified using two angles, θ and ϕ. I will refer to these angles as "detector-based angles" to distinguish them from another set of angles that I will introduce later in Chapter 4. The detector-based angles consist of a polar angle θ measured from the $+z$ axis with the range [0, π], and an azimuthal angle ϕ measured from the $+x$ axis toward the $+y$ axis with the range [0, 2π]. These are the usual spherical polar angles as defined in the standard physics convention.

2.2.2 Energy Resolution

The energy resolution of a VMI instrument is commonly described as the relative resolution $\Delta E/E$, given in percent. A VMI spectrometer can measure a charged particle angular distribution of any kinetic energy up to a maximum determined by the design and electrode voltages applied. Therefore, the energy resolution $\Delta E/E$ at a given voltage configuration is typically given with respect to the maximum energy particles for which the entire angular

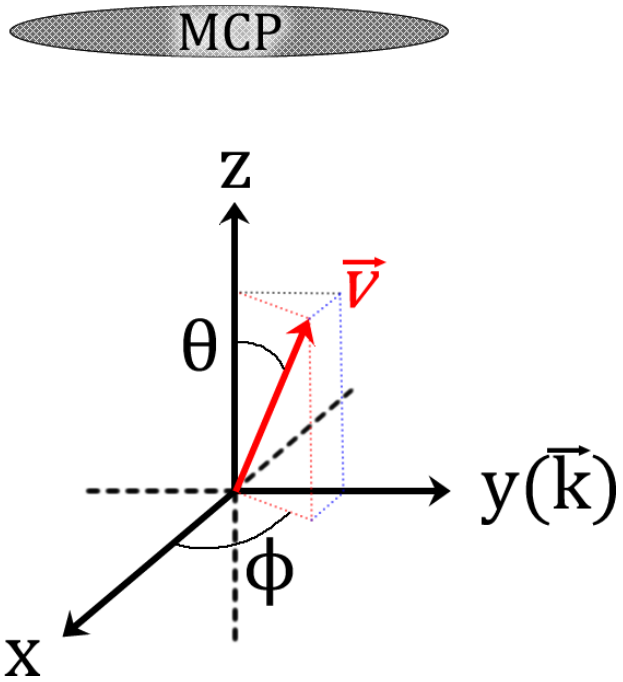

Figure 2.1: Lab frame coordinate system and definition of the detector-based angles. I define the z axis to be the axis normal to the plane of the MCP detector, with y being the axis of propagation of the laser ($\vec{k}$). The detector-based angles θ and ϕ follow the standard physics convention for polar coordinates.

distribution can be imaged. This corresponds to the maximum radius on the detector.

If we assume that the electric field experienced by the particles is strictly along the z (TOF) axis, and that they experience no transverse acceleration, their arrival position at the detector is given by $\mathbf{R} = \mathbf{v}_\perp t$, where $\mathbf{v}_\perp$ is the transverse component of the emission velocity and t is the TOF. The presence of the velocity mapping fields make this assumption untrue, particularly in the case of more complex thick-lens multi-electrode designs (see Sections 2.5 and 3.2.2), but it can still be used to give a rough evaluation of the relationship between initial kinetic energy and transverse detector position. For particles emitted in a transverse direction (perpendicular to the TOF axis), the kinetic energy E is proportional to $v_\perp^2$. Combined with the zero-transverse-field assumption ($\mathbf{R} \propto \mathbf{v}_\perp$), this implies $E \propto R^2$, or $E = kR^2$ for a proportionality constant k.

From this relation, we can derive a form for the absolute energy resolution ΔE in terms of the radial position resolution ΔR:

$$\begin{aligned} \Delta E = \frac{dE}{dR}\Delta R = 2(kR)\Delta R = 2(\frac{E}{R})\Delta R \\ \rightarrow \frac{\Delta E}{E} = 2\frac{\Delta R}{R}. \end{aligned} \tag{2.1}$$

As described in the previous section, VMI focusing can only be optimized for a single particle energy at one time. Typically this is the energy corresponding to the maximum radius on the detector, and the energy resolution is given at this energy. Any particles of lower energy than the maximum for the spectrometer's settings will necessarily have a lower resolution.

2.2.3 Turn-around Time Spread

The turn-around time spread (TTS) is defined as the difference in arrival time between a particle emitted directly toward the detector ($+z$) and one emitted directly away ($-z$). In the processing of 3D VMI data (see Chapter 4), where each event must be individually time-stamped, the TTS is of vital importance. Given an instrument-limited absolute arrival time resolution dt, the TTS Δt determines the number of meaningful time-slices n which can be taken through the PAD: $n = \Delta t/dt$. Since the arrival time is directly related to the momentum along the z axis, p_z, a longer TTS corresponds to a better p_z resolution.

A particle emitted from the centre of the interaction region directly away from the detector experiences a repulsive field causing it to slow down, reverse course, and arrive back at the point of origin. From there, the particle follows the same trajectory as an electron emitted from the same point directly toward the detector. Therefore, the TTS is determined by the field experienced by the particle during that turn-around time (hence the name turn-around time spread). It is the field near, and directly behind, the interaction region which contributes to the TTS. If the particle (of emission energy ϵ) does not travel far before reversing we can assume that the field in the turn-around region (E_{TTS}) is constant. In this case, the TTS is given by

$$\Delta t = C\frac{\sqrt{\epsilon}}{E_{TTS}}, \tag{2.2}$$

where $C = 2\sqrt{2m}/q$ contains the mass (m) and charge (q) of the particle. Clearly, the

TTS is inversely proportional to the turn-around field. To increase TTS, it is beneficial to have a lower E_{TTS}. However, the extent to which the field can be lowered is constrained by the size of the flight tube and detector; all particles in the distribution must still reach the detector without colliding with the walls of the instrument. This problem has been solved by employing a field gradient which starts weak near the interaction region and ramps up along the flight tube to ensure the distribution still reaches the detector (see Sections 2.5 and 3.2.2).

2.3 Photoelectron Angular Distributions

One of the main advantages of photoelectron VMI as a tool for studying molecular dynamics is having access to the photoelectron angular distribution in addition to the photoelectron kinetic energy spectrum. Once far enough away from its parent ion, the composition of the PAD can be described in terms of angular momentum scattering states of the continuum electron wave function in a Coulombic potential. This is the well-known partial wave expansion. Not all possible scattering states will be accessible for a given photoionization channel, and the allowed states depend on the angular momentum properties and symmetry of the molecular orbital being ionized, the final state of the ion, and the transition dipole moment. A detailed derivation of the photoionization cross-section [47] shows that the cross-section can only be non-zero if the direct product of the irreducible representation of the free electron, the ion electronic state, the dipole moment, and the neutral molecule's electronic state contains the totally symmetric representation of the molecular point group.

This condition demonstrates that the free electron state (and thus the PAD) is sensitive to changes in the electronic state of the neutral molecule. This makes the PAD a suitable observable with which to follow evolving molecular dynamics. The PAD complements measurement of the PES, allowing one to track changes in state symmetry without relying on a change in photoelectron emission energy.

For a pump-probe experiment, where a molecule is ionized after a time delay t relative to the pump pulse, the energy- and angle-resolved cross-section $\sigma(\epsilon, \hat{\mathbf{k}}_{\mathrm{L}}; t)$ is given by [47]

$$\sigma(\epsilon, \hat{\mathbf{k}}_{\mathrm{L}}; t) = \frac{\sigma_{\mathrm{total}}(\epsilon; t)}{4\pi} \sum_{LM} \beta_{LM}(\epsilon; t) Y_{LM}(\hat{\mathbf{k}}_{\mathrm{L}}), \tag{2.3}$$

where ϵ is the initial photoelectron kinetic energy and $\hat{\mathbf{k}}_{\mathrm{L}}$ is the initial lab-frame photoelectron momentum vector. The quantity σ_{total} (represented by the expansion coefficient β_{00}) is the energy-resolved total cross-section and represents the PES, and all other terms in the sum (besides the β_{00} term) contribute to the PAD. The sum in equation 2.3 describes the PAD as a partial wave expansion in the lab-frame spherical harmonics Y_{LM}. The expansion coefficients β_{LM} are referred to as asymmetry parameters (which are normalized by β_{00}) and fully characterize the PAD. Each term, or partial wave, in this expansion has a total angular momentum quantum number L and projection quantum number M along a lab-frame axis. For linear polarization of the ionizing laser, this axis is chosen to be the polarization axis. For a randomly oriented ensemble of molecules ionized by a linearly polarized laser, as is the case for the experiments described in this thesis, the resulting PAD will be cylindrically symmetric with respect to the polarization axis. Thus, for these

cases, only partial waves with $M = 0$ contribute. For simplicity, the $M = 0$ asymmetry parameters are often labelled without explicit inclusion of the M subscript: $\beta_L \equiv \beta_{L0}$. In addition, the random orientation of the molecular ensemble means that the PAD must posses $\pm$ symmetry about the polarization axis. This eliminates any contribution from partial waves with odd L. The maximum value of L in the partial wave expansion is determined by the number of photons involved in the ionization process: $L_{\max} = 2N$, for an N-photon ionization [48]. In this thesis, only up to two-photon ionization processes will be discussed. Therefore, the only non-zero asymmetry parameters that will be encountered are β_2 and β_4. The isotropic contribution β_0 represents the total ionization cross-section and does not describe the angular distribution.

2.4 The Forward and Inverse Abel Transform

In this section I discuss the conversion of a measured 2D VMI detector-plane photoelectron image into initial kinetic energy and angular distributions. This requires an inversion method, of which the Abel inversion is a prominent example. The Abel transform [49], in its most basic form, computes the one-dimensional projection $S(y)$ of a circularly symmetric two-dimensional function $f(\rho)$, where $\rho = \sqrt{y^2 + z^2}$ is the radial coordinate in the YZ plane:

$$S(y) = 2 \int_y^\infty \frac{f(\rho)\rho \mathrm{d}\rho}{\sqrt{\rho^2 - y^2}}. \tag{2.4}$$

I have defined the function $f(\rho)$ to be located in the YZ plane with the projection along the z-axis in preparation for the following generalization into 3D, and for consistency with the coordinate system used throughout this thesis. The integral in equation 2.4 shows that, at each value of y, the projection $S(y)$ contains contributions from $f(\rho)$ at each radial coordinate $\rho \geq y$ up to $\rho = \infty$ (or, practically, the maximum radius of the object represented by the function $f(\rho)$).

The 3D version of the Abel transform (given, for example, in reference [50]) describes the projection of a cylindrically symmetric 3D object onto a 2D plane parallel to the axis of symmetry:

$$S(R, \alpha) = 2 \int_R^\infty \frac{I(r, \theta) r \mathrm{d}r}{\sqrt{r^2 - R^2}}. \tag{2.5}$$

Here I define x to be the axis of cylindrical symmetry, with the projection being along the z-axis and onto the XY plane. The coordinate system used in the present discussion of the Abel transform is shown in Figure 2.2. The angles defined in Figure 2.2 are used instead of those defined in Figure 2.1 for this section only. The 3D object is described by the function $I(r, \theta)$, where $r = \sqrt{x^2 + y^2 + z^2}$ is the spherical radius and θ is the polar angle. The requirement of cylindrical symmetry has removed the dependence on the azimuthal angle ϕ. The projection $S(R, \alpha)$ is described using the 2D XY-plane radius $R = \sqrt{x^2 + y^2}$ and polar angle α.

The 3D Abel transform (equation 2.5) describes the relationship between a cylindrically symmetric charged particle momentum distribution (proportional to I) and the 2D

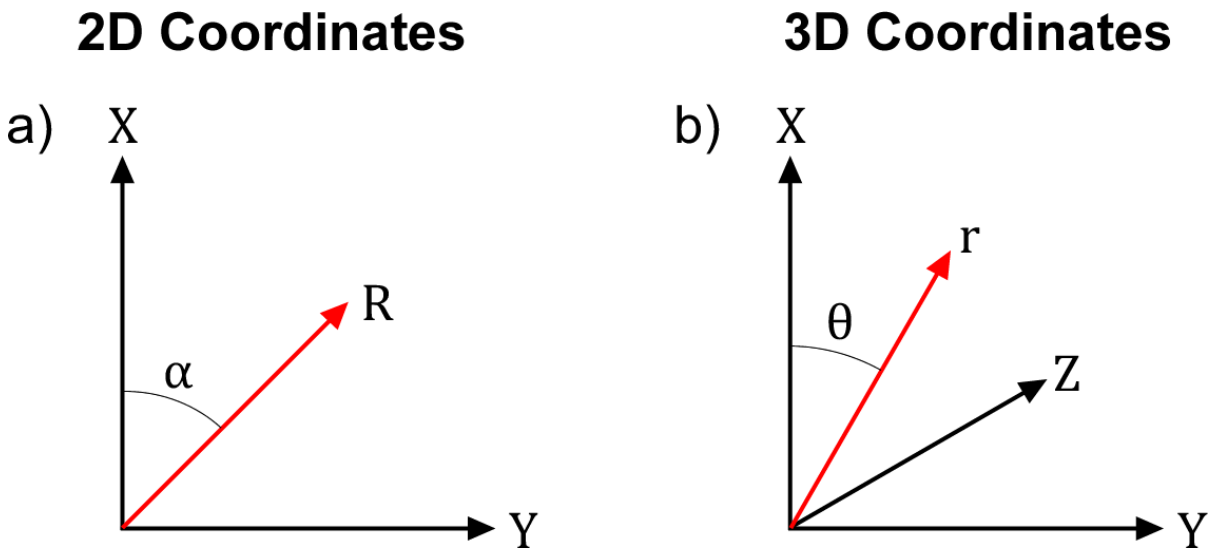

Figure 2.2: Coordinate systems used to define the 2D detector-plane projection (a) and 3D charged particle distribution (b) for the Abel transformation of a VMI image. The projection is taken along the z axis and is in the XY plane, and cylindrical symmetry about the x axis is assumed.

projected image (S) measured in a VMI experiment. As such, so long as the symmetry requirements are satisfied, the inversion of the Abel transform is a simple and convenient way of reconstructing the 3D distribution from a 2D VMI image.

2.4.1 The DAVIS Algorithm

In 2018, authors Harrison *et. al.* published a paper describing a new method for performing the inverse Abel transform, designed specifically for use in VMI applications [43]. The paper, titled "DAVIS: A direct algorithm for velocity-map imaging system", introduces the titular algorithm and demonstrates that it is simultaneously less sensitive to noise and less computationally demanding than the popular pBASEX inversion algorithm [50]. The DAVIS algorithm expands equation 2.5 as a sum of Legendre polynomials, making use of the physical constraints present in a typical VMI experiment, then derives from this a basis set which can be used to fit the experimental projection. This directly yields the asymmetry parameters which characterize the original 3D distribution. The pBASEX algorithm employs a similar strategy, but uses a numerical treatment of the integral in equation 2.5 instead of the analytical approach used in DAVIS.

In the following I reproduce the main steps of the derivation from the DAVIS paper [43] and outline my implementation of the algorithm in Python.

2.4.1.1 Description of the 2D Projected Image

The first step in the derivation is to express the form of a general 2D image in terms of a basis set which, when used to fit an experimental VMI image, directly yields the radial distribution and asymmetry parameters. The authors' derivation assumes the use of a linearly polarized laser in the model VMI experiment, although they note that the process can be generalized to other polarization states. Since the work in this thesis uses only linear polarizations, I will follow the derivation using the same assumption.

The 3D distribution $I(r, \theta)$ is first separated into a product of a radial part $F(r)$ and a radially-dependent angular part $G(r, \theta)$. The angular function $G(r, \theta)$ is then expanded as a sum of Legendre polynomials $P_n(\cos\theta)$:

$$\begin{aligned} I(r,\theta) &= F(r)G(r,\theta) \\ &= F(r)\sum_{n=0}^{2N}\beta_n(r)P_n(\cos\theta). \end{aligned} \tag{2.6}$$

The coefficients of the Legendre polynomials, $\beta_n(r)$, are the radially-dependent asymmetry parameters and contain all of the angular dependence of the 3D charged particle distribution. The Legendre polynomials here are in terms of θ, the polar angle measured from the $+x$ cylindrical axis. As discussed in Section 2.3, for linearly polarized light the Legendre expansion in equation 2.6 only includes terms up to $2N$, where N is the num-

ber of photons involved in the process. Note that in this thesis I only discuss processes involving up to two photons; for these cases, the sum ends at $2N = 4$.

Substituting equation 2.6 into equation 2.5 and pulling out the sum gives

$$S(R,\alpha) = \sum_{n=0}^{2N} 2 \int_R^\infty \frac{F(r)\beta_n(r)r}{\sqrt{r^2 - R^2}} P_n(\cos\theta)\mathrm{d}r. \tag{2.7}$$

Note that the 2D radius R and polar angle α are used here in the symbol $S(R,\alpha)$, as in equation 2.5; this is to emphasize the 2D nature of the projected image which is described using these coordinates. This continuous function is the projection of the 3D distribution $I(r,\theta)$ onto the plane of the detector (XY). When this projection is imaged by a camera in an actual VMI experiment, it is instead represented as a discrete array of pixels $\mathbf{S}_m$. Typically this is a regular rectangular array of Cartesian pixels. However, it can easily be transformed into a polar image, and in this case it can be expressed as the following integration of the continuous projection $S(R,\alpha)$:

$$(\mathbf{S}_m)_{ij} = \frac{1}{r^2} \int_{\check{R}_i}^{\hat{R}_i} \int_{\check{\alpha}_j}^{\hat{\alpha}_j} S(R,\alpha) R \mathrm{d}R \mathrm{d}\alpha, \tag{2.8}$$

where i and j index the radial and angular pixels, respectively. The symbols $\check{R}_i$, $\hat{R}_i$, $\check{\alpha}_j$, and $\hat{\alpha}_j$ are the upper and lower bounds of each pixel, as determined by the radial (ΔR) and angular ($\Delta\alpha$) pixel widths: $\hat{\check{R}}_i = R_i \pm \Delta R/2$ and $\hat{\check{\alpha}}_j = \alpha_j \pm \Delta\alpha/2$, where the upper symbol above R_i or α_j is taken with the upper symbol of $\pm$ and similarly for the lower symbols.

The Legendre expansion in equation 2.7 is in terms of θ, the polar angle of the 3D distribution. This is not a suitable basis for describing the 2D projection; it is more convenient to express equation 2.7 in terms of the 2D image-plane polar angle α. The authors make use of a Legendre polynomial identity [51] (derived by some of the same authors) which allows $P_n(\lambda x)$ to be expanded as a sum of $P_n(x)$:

$$P_n(\lambda x) = \sum_{k=0}^{\lfloor n/2 \rfloor} b_{\lambda,n,k} P_{n-2k}(x), \tag{2.9}$$

where the expansion coefficients $b_{\lambda,n,k}$ are given in reference [51]. This identity can be applied to equation 2.7 since the Legendre polynomial $P_n(\cos\theta)$ can be written as $P_n(\frac{R}{r}\cos\alpha)$. The immediate result, expressing $S(R,\alpha)$ as a sum of $P_k(\cos\alpha)$, can be found in reference [43]. For the sake of brevity I will not reproduce it here as it is long and represents an intermediate step in the full derivation. After substituting the full expression for $S(R,\alpha)$ into equation 2.8 for the discrete polar image $\mathbf{S}_m$ and rearranging the integrals and sums, $\mathbf{S}_m$ is expressed in a more concise matrix form:

$$\mathbf{S}_m = \sum_{n=0}^{2N} \sum_{k=0}^{\lfloor n/2 \rfloor} \mathbf{M}_{n,n-2k} \mathbf{F}_n \otimes \mathbf{P}_{n-2k}, \tag{2.10}$$

where $\mathbf{F}_n$ is a radial-pixel column vector with elements $F(r_i)\beta_n(r_i)$ and $\mathbf{P}_{n-2k}$ is an angular-pixel column vector with elements $P_{n-2k}(\cos\alpha_j)$. The vectors $\mathbf{F}_n$ contain all of the information necessary to reconstruct the 3D distribution $I(r,\theta)$, and extracting the $\mathbf{F}_n$ is the goal of the DAVIS algorithm. Equation 2.10 expresses the 2D image $\mathbf{S}_m$ explicitly in

terms of the 2D image-plane Legendre polynomials $\mathbf{P}$, the 3D distribution parameters $\mathbf{F}$, and the transformation matrices $\mathbf{M}$.

The general form of the matrices $\mathbf{M}$ are given in reference [43]. They also provide analytical forms for the first few matrices, up to $\mathbf{M}_{4,4}$, in Table I of that publication. Since I am only concerned with up to two-photon processes in this thesis, for which the maximum value of n is 4, the general form is unnecessary. I should note that Table I of reference [43] contains an error in the expression for $\mathbf{M}_{20}$; the + in the equation should be a − instead.

Equation 2.10 is further simplified by rearranging the sums and defining the vectors $\boldsymbol{\delta}_n$:

$$\boldsymbol{\delta}_n = \sum_{k=0}^{N-\lceil n/2 \rceil} \mathbf{M}_{n+2k}\mathbf{F}_{n+2k}. \tag{2.11}$$

The 2D image array $\mathbf{S}_m$ is then:

$$\mathbf{S}_m = \sum_{n=0}^{2N} \boldsymbol{\delta}_n \otimes \mathbf{P}_n. \tag{2.12}$$

2.4.1.2 Inverting the Image

The preceding formulation requires that the input 2D image be in polar form. For a Cartesian image the first step is the transformation into polar coordinates. This transformation is well-known and implemented in many program packages and will not be described in detail here.

The parameters that define the 3D distribution ($F(r)$ and $\beta_n(r)$) are contained within the vectors $\mathbf{F}_n$, which in turn are within the Legendre coefficient vectors $\boldsymbol{\delta}_n$. Therefore, the inversion begins by fitting the angular (α) dependence of the 2D polar image at each radial index to a sum of Legendre polynomials, directly giving the coefficients $\boldsymbol{\delta}_n$. For each Legendre polynomial order k (up to 4 for two-photon processes) and each radial pixel index i, the corresponding element $\delta_k(R_i)$:

$$\delta_k(R_i) = \frac{\Delta\alpha(2k+1)}{2} \sum_{j=1}^{N_\alpha} |\sin\alpha_j| \, P_k(\cos\alpha_j) I_{R_i}(\alpha_j), \tag{2.13}$$

where N_α is the total number of angular pixels in the polar image and $\alpha_j = 2\pi j/N_\alpha$.

The expression for the $\boldsymbol{\delta}_n$ vectors (equation 2.11) can be rearranged to give the 3D distribution vectors $\mathbf{F}_n$ as two independent recursively defined systems of equations (one even, one odd):

$$\mathbf{F}_k = \mathbf{M}_{k,k}^{-1} \left(\delta_k - \sum_{i=1}^{2N-\lceil k/2 \rceil} \mathbf{M}_{k+2i,k} \mathbf{F}_{k+2i} \right). \tag{2.14}$$

Solving these equations from highest to lowest k yields the $\mathbf{F}_n$ vectors containing the radial distribution and asymmetry parameters:

$$\begin{aligned} F(R_i) &= F_0(R_i), \\ \beta_k(R_i) &= \frac{F_k(R_i)}{F_0(R_i)}. \end{aligned} \tag{2.15}$$

2.4.1.3 Python Implementation

Data obtained using 3D VMI are not reliant on Abel inversion for their analysis. However, for 3D VMI data which can also be analyzed using an inversion-based approach (that is, data possessing cylindrical symmetry), it is beneficial to do so as a complement to the 3D analysis so that both methods can be directly compared. I have implemented the DAVIS inversion algorithm in Python and I will outline here the details of my code and demonstrate the results using experimental 3D VMI data.

The Legendre polynomial vectors $\mathbf{P}_n(\cos\alpha)$ and the transformation matrices $\mathbf{M}_{n,n-2k}$ are constant for any inversion, assuming the same radial and angular pixel widths are used. To save on computational resources, I generated and saved the $\mathbf{P}_n(\cos\alpha)$ and $\mathbf{M}_{n,n-2k}$ using widths of $\Delta r = 0.5$ pixels and $\Delta\alpha = 1^\circ$ (which give good performance with test data). I generated the $\mathbf{M}_{n,n-2k}$ with a maximum radius large enough to accommodate any image taken using the installed camera. Both the $\mathbf{P}_n(\cos\alpha)$ and $\mathbf{M}_{n,n-2k}$ are read in when the code is used, truncating the $\mathbf{M}_{n,n-2k}$ to the maximum radius of the image being inverted.

For my demonstration of the inversion process (shown in Figure 2.3) I use data from the one-photon ionization of the excited A-state of nitric oxide (NO) (outer ring) and two-photon ionization of the NO(X) ground state (inner ring), both at 266 nm. These data and the associated experimental details are described in detail in Section 3.6.3. The 2D image (Figure 2.3 (a)) must first be processed before it can be inverted. This involves determining the centre of the image and trimming it into a square array, rotating the image to align the cylindrical axis to the x axis, and symmetrizing by assigning to each quadrant the average

of all four. The last step is done since the images possess four-fold symmetry as a result of the limited allowed terms in the partial-wave expansion, as discussed above and in Section 2.3. The resulting prepared image is shown in Figure 2.3 (b).

The Cartesian image is converted into polar form (not shown) with the same pixel widths as the matrices. The transformation is performed in two steps: expanding the image at each radius as a sum of Legendre polynomials using equation 2.13 to get the coefficients $\boldsymbol{\delta}_n$, and solving equation 2.14 for the $\mathbf{F}_n$. The radial spectrum and asymmetry parameters are given directly by equation 2.15. The inverted image, which is the centre slice of the reconstructed 3D distribution (mathematically $r = R$) shown in Figure 2.3 (c), can be reconstructed using:

$$I(r = R, \theta) = F(r) \sum_{n=0}^{2N} \beta_n(r) P_n(\cos\theta). \tag{2.16}$$

The normalized radial spectrum $F(r) \times r$ and the asymmetry parameters β_2 and β_4 are shown in Figure 2.3 (d). As discussed in Section 3.6.3, the outer ring is expected to have a β_2 between 1.2 and 1.7 [52, 53]. Analyses of the same data set in Section 3.6.3 using a different inversion algorithm [54] and also fitting slices of the PAD directly gave values close to $\beta_2 = 1.7$. The value of β_4 should be zero by conservation of angular momentum for a one-photon process. The inner-ring photoionization dynamics are not as well-known, but being a two-photon process it is allowed to have a $\beta_4 > 0$. The value of β_2 for the outer ring is larger than expected; this could be explained by an improperly set polarization state resulting in a slightly non-cylindrically-symmetric distribution (see Section 4.5).

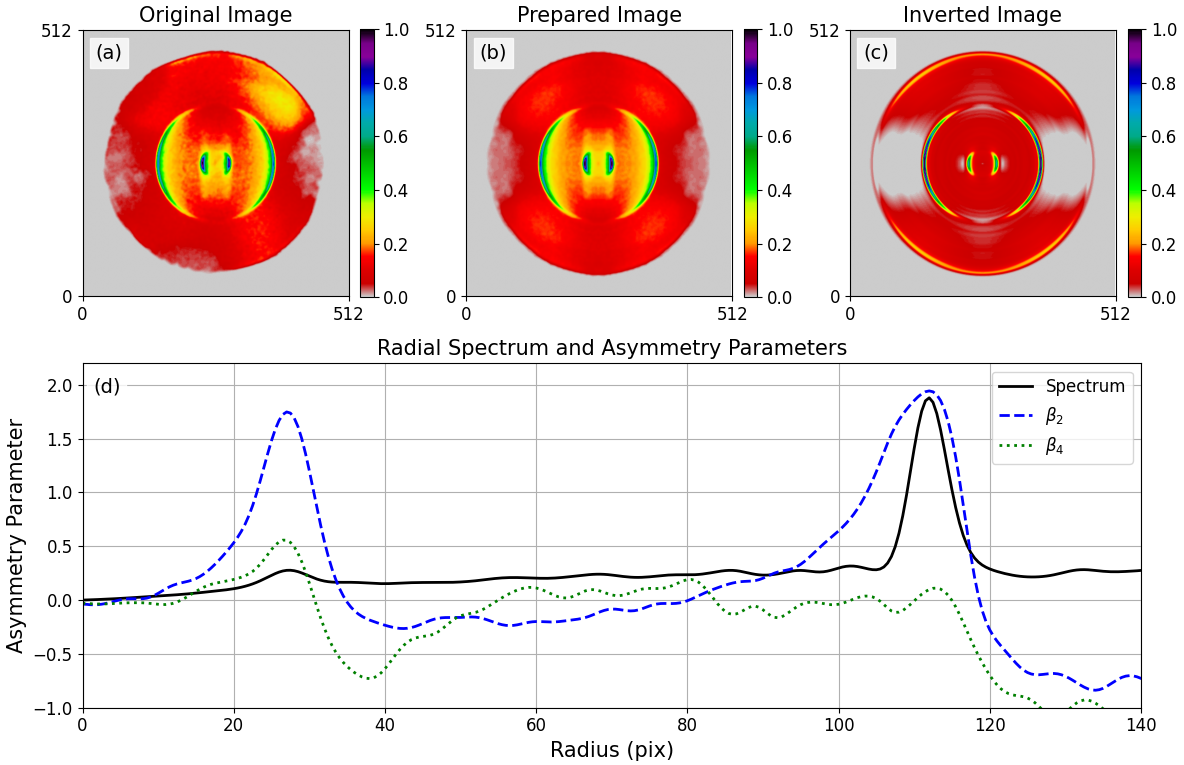

Figure 2.3: Abel inversion of experimental VMI data using the DAVIS algorithm. The data are from the one-photon ionization of the excited A-state of nitric oxide (NO) (outer ring) and two-photon ionization of the NO(X) ground state (inner ring), both at 266 nm (see Section 3.6.3 for additional details). The full VMI image (a) is centred, rotated, and symmetrized as described in the text to give the processed image (b) which is then inverted (c). The DAVIS algorithm directly yields the radial spectrum and asymmetry parameters (d). In (d), the radial spectrum is normalized to an arbitrary value and the vertical axis applies only to the asymmetry parameters (β_2, β_4).

2.5 Summary of Design Considerations and Simulations for a New 3D VMI Spectrometer

The work described in Chapters 3 and 4 was performed using the NRC 3D VMI spectrometer, located at 100 Sussex Drive, Ottawa, Ontario. The Stolow research group is currently constructing a new 3D VMI spectrometer on the University of Ottawa campus (the uOttawa spectrometer) which improves upon the NRC spectrometer design. Since much of the methods described in this thesis will be applied to this new instrument, I will summarize here the main considerations that led to the final design for the uOttawa spectrometer. The contents of this section are adapted from the work described in the University of Ottawa physics Master's thesis of Michael Hemsworth [4].

The stated design goals for the new instrument were the following: (i) High 2D VMI energy resolution, which corresponds to high spatial resolution in the detector (XY) plane. (ii) High z-axis momentum resolution in 3D VMI operation, which corresponds to high arrival-time resolution and a longer TTS. (iii) Capability for ion TOF mass spectrometry and electron-ion covariance measurements. In addition to these three goals the usual velocity mapping conditions had to be retained.

2.5.1 Design Candidates

Three different designs were considered, and each was tested using the SIMION 8.0 charged particle optics simulation package [55]. Simulations were also performed on a basic three-electrode VMI design to provide a baseline for comparison. These four spectrometer designs were:

Three-Electrode VMI Design This is a basic three-electrode (repeller, extractor, and ground) spectrometer similar in design to the original VMI of Eppink and Parker [5]. To better compare with the designs under consideration for the new spectrometer, the flight distance (interaction region to detector) of the three-electrode simulations was made roughly the same as in the current NRC spectrometer (130 mm). This design uses a 40 mm diameter MCP detector and the electrodes have an internal diameter (ID) of 20 mm.

13-Electrode NRC Spectrometer Design The design of the existing NRC spectrometer has been proven to be effective and, as such, a duplicate of this instrument was considered for the uOttawa spectrometer. The design specifications for the NRC spectrometer are described in detail in Section 3.2.2. In brief, it uses eight electrodes with a linear voltage ramp in place of the single extractor electrode in the basic design. This allows the charged particle distribution to experience a weaker extraction field in the interaction region, resulting in a larger TTS (see Section 2.2.3), while still being accelerated rapidly enough towards the detector in order to avoid collisions with the flight tube. It

also uses three electrodes in place of the single repeller, and an extra (laser port) electrode, so that the distribution has room to expand in the negative TOF direction (away from the detector). As above, this design uses a 40 mm detector and has a flight distance of ~ 130 mm. The electrode ID is 50 mm.

25-Electrode Thick-Lens Design This design expands on the NRC design by using a 20-electrode thick-lens extractor. Combined with the three-electrode repeller, laser port electrode, and ground electrode, this is a total of 25. These extra electrodes would give a finer control of the voltage ramp. Furthermore, this design would have a longer (240 mm) and wider (84 mm ID) flight tube, and a larger (80 mm diameter) MCP detector. The benefits of this design are an improved z-momentum resolution (TTS), 2D energy resolution, and ion mass resolution, all resulting from the increase in dimensions of the flight tube. However, the increased scale of the instrument, particularly the 80 mm MCP detector, would drastically increase the price.

PEPICO Design This design (Figure 2.4) also expands on the NRC design, but in this case an additional coaxial ion flight tube is added in the opposite direction from the electron flight tube, thus enabling PhotoElectron-PhotoIon COincidence (PEPICO) measurements. The electron flight tube mimics the NRC design, but is slightly wider (60 mm ID). This increase serves to reduce distortions which can affect slow electrons passing close to the electrode surfaces (see Section 3.6.1). The ion side of the spectrometer consists of a second repeller grid 11 mm below the first, followed by a 560 mm flight tube (interaction region

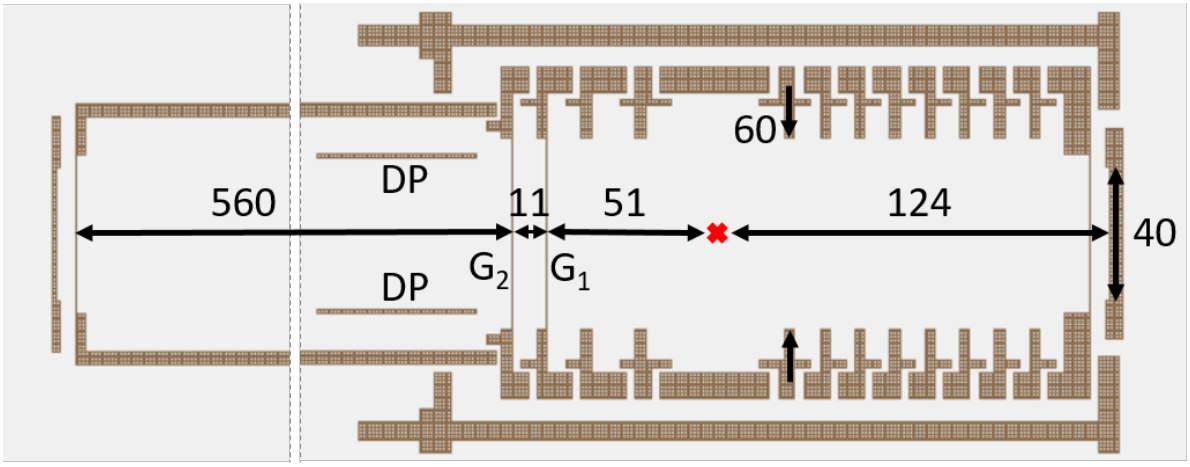

Figure 2.4: Layout of the PEPICO design for the new 3D VMI spectrometer, as used for the SIMION [55] simulations described in the text. Relevant dimensions are provided in mm. The labels G_1 and G_2 refer to the grids and DP refers to the deflector plates; these elements are described in the text. The red $\times$ shows the location of the interaction region. This figure was copied with permission from the Master's thesis of Michael Hemsworth [4].

to detector). The presence of the second flight tube means that the molecular beam can only enter from a port perpendicular to the flight axis. This means that the centre-of-mass velocity of the ion Newton sphere must be accounted for so that the ions do not collide with the side of the flight tube. Two deflector plates (labelled DP in the figure) are positioned after the ion-side repeller grid to compensate for this transverse velocity. While this design does require a second 40 mm MCP, it would still be considerably cheaper than the 25-electrode design. This, combined with the ability to perform PEPICO measurements, made this design an attractive option.

With the exception of the PEPICO design, as mentioned above, each of these designs can be implemented with the molecular beam entering either along the TOF (z) axis (to-

ward the detector) or perpendicular to the TOF axis. As described in Section 3.2.2, the NRC spectrometer uses the z-axis-oriented configuration. There are two main considerations which favour a perpendicular orientation for the molecular beam. First, the molecular beam possesses a velocity distribution which is carried over to the ions (electrons, being light in comparison, are not significantly affected). This distribution has a blurring effect on the ion mass resolution. Second, aiming the molecular beam along the z direction toward the detector makes it more difficult to pump out the excess gas, increasing background gas density. With the perpendicular molecular beam arrangement, a much higher target density can be achieved, as most of the gas passes through the interaction region without collisions. Over time, the excess gas particles impacting the detector can cause more rapid wear and damage the MCP. The velocity of a perpendicular molecular beam does affect the trajectory of the ions; however, as explained above, this can be compensated for with ion-side deflector plates. For these reasons, it was decided that the perpendicular orientation was preferred.

2.5.2 Simulations

The four designs described above were tested using SIMION simulations to evaluate their performance in three categories: 2D energy resolution, TTS (3D z-momentum resolution), and ion mass resolution.

The interaction region was modelled as a 2 mm long, 0.1 mm diameter cylinder along the laser axis. Particles were generated from nine origin points within this region and with

a range of emission angles. A second type of simulation was also performed, using a Monte Carlo sampling ($\sim 10^5$ particles) over a range of emission angles and origin points within the interaction region cylinder. Both of these simulation types are described in detail in Section 3.2.2. The Monte Carlo simulations were less clear than the discrete sampling method, and either supported the discrete results or were inconclusive (in the case of the electron energy resolution test). Thus, I will not distinguish between the two simulation types in the following.

For each of the designs, electrons were generated at 1, 2, 3, 4, and 5 eV, and the fields were optimized to minimize the energy resolution ($\Delta E/E$) at the maximum energy of 5 eV.

2.5.2.1 2D VMI Energy Resolution and Linearity

The energy resolution $\Delta E/E$ was calculated from the discrete simulations by taking ΔE to be the difference of the energies corresponding to the maximum and minimum radii from the discrete trajectories sampled. These are plotted in Figure 2.5. The simulations showed that the energy resolution was best for the 25-electrode design, followed in order by the PEPICO, NRC, and three-electrode design. This trend was true for each of the energies sampled. For each design, the energy resolution became monotonically worse for energies further from 5 eV (the energy that the fields were optimized for). Importantly, the 25-electrode design was not significantly worse at any energy than the PEPICO design. For example, at 1 eV (the worst case among the energies sampled) the difference between these two designs was 0.6%.

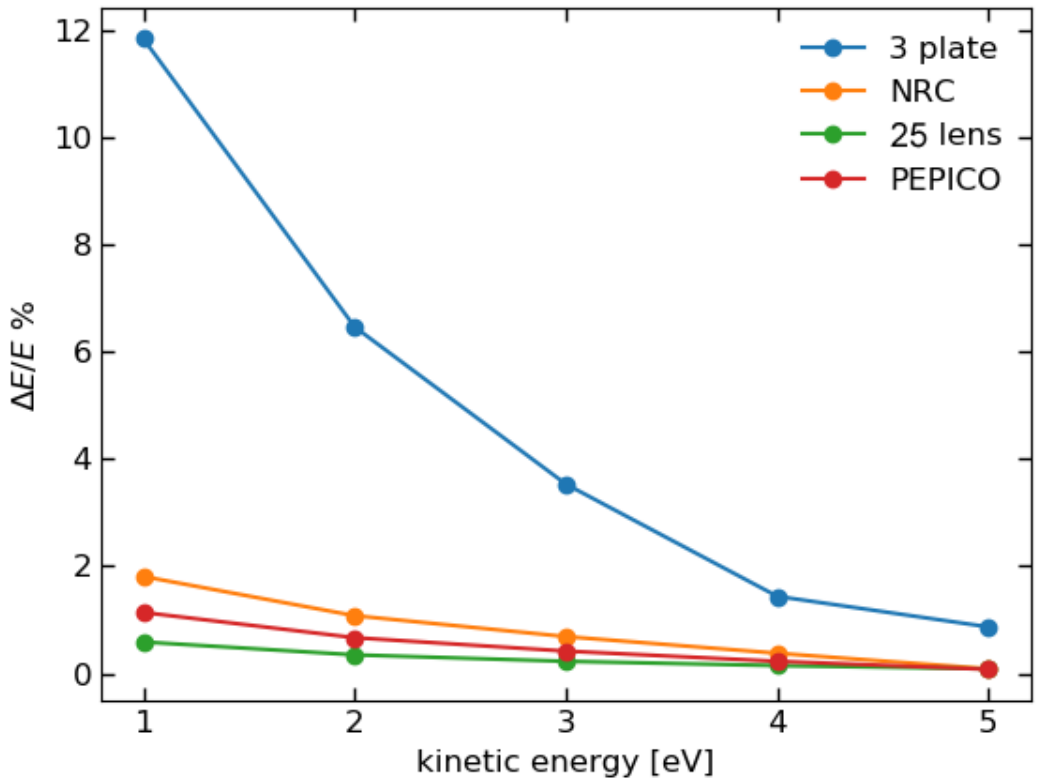

Figure 2.5: Relative energy resolutions ($\Delta E/E$) for each of the four spectrometer designs described in the text, derived from SIMION [55] simulations. The electrostatic lens voltages in the simulations were optimized for 5 eV electrons. This figure was copied with permission from the Master's thesis of Michael Hemsworth [4].

The relationship between the 2D detector-plane radius R and a particle's kinetic energy E is usually described by $E = kR^2$ for calibration constant k (see Section 2.2). The simulations highlight the relative validity of this simple linear relationship between E and R^2 for each of the designs. In his thesis [4], Hemsworth showed that each of the designs displayed clearly quadratic residuals when fitted to the form $E = kR^2$. For the NRC design currently being used, the maximum residuals (at 1 and 5 eV) were ~ 14 meV. This is larger than the 4.5 meV absolute energy resolution (ΔE) that the NRC spectrometer has (according to the simulations) at 5 eV. Instead, the relationship between E and R could be better described by adding a fourth-order term:

$$E = k_1 R^2 + k_2 R^4. \tag{2.17}$$

Fitting the simulated radii using equation 2.17 gives better residuals (< 2 meV for all except for the three-electrode design), but requires a more complex calibration source consisting of at least three energies.

2.5.2.2 Turn-around Time Spread

The simulations provide an estimate of the electron TTS that can be achieved using each of the spectrometer designs. These are: 1.8 ns (three-electrode), 8.0 ns (PEPICO), 8.6 ns (NRC), and 10 ns (25-electrode). While the 25-electrode design provides the largest TTS, the NRC and PEPICO designs are not far behind.

For any of the designs under consideration, it is possible to further increase the TTS by employing a nonlinear potential ramp instead of the linear ramp currently used in the NRC spectrometer. Simulations using a nonlinear ramp in conjunction with the PEPICO design show that a TTS of 15.5 ns should be possible. Furthermore, the nonlinear PEPICO configuration does not exhibit as severe of a drop in resolution when moving away from the optimal energy (5 eV in this case) compared to the linear ramp configurations. With individual control of the voltage applied to each electrode (possible for any of the designs), many such nonlinear configurations can be tested and optimized after the instrument is constructed. For this reason, the advantages of a nonlinear ramp were not directly taken into consideration when choosing the design for the new instrument.

The problem of transforming the TOF into the z-axis momentum component is described in detail in Section 4.4.

2.5.2.3 Ion Mass Resolution

In general, assuming a constant electric field and no initial velocity along the z axis, the TOF for an ion is proportional to $\sqrt{m/q}$, where m/q is the mass-to-charge ratio for the ion. The mass resolution is defined as the maximum mass n (in amu) such that masses of n and $n+1$ can be resolved. If the molecular beam were directed along the z-axis, the initial z-velocity would not be zero and this would result in a blurring of the mass spectrum and a lower resolution. To avoid this, and for the reasons described above, it is preferable to have the molecular beam perpendicular to the flight axis. The

simulations support this reasoning, showing that a much higher mass resolution can be achieved in the perpendicular configuration (> 500 amu) than in the parallel configuration (< 300 amu), regardless of the spectrometer design. These mass resolution estimates use a very conservative condition of approximately zero overlap between mass peaks, so the actual resolution will be considerably higher.

2.5.3 Final Decision

Considering the above results, the PEPICO design (Figure 2.4) was chosen for the new VMI instrument. Although the 25-electrode design did show some advantages, it was judged that the improvements were not enough to justify the increased cost. In addition, the PEPICO design brings the ability to perform electron-ion covariance measurements.

It is worth noting that, when operating in covariance mode, the electron kinetic energy and ion mass ranges are coupled. Optimizing for low-energy electron detection will also lower the range of detectable ion mass since the lower fields will allow the ions to collide with the inside surface of the spectrometer before entering the ion-side flight tube. Also, the required fields for covariance operation will result in a reduced electron energy resolution, having performance between that of the three-electrode VMI and NRC designs.

Construction of this new instrument is nearing completion and it is expected to be operational by December 2025.

3

Experimental Design

3.1 Introduction

In this chapter, I describe the experimental setup used for the measurements discussed throughout this thesis. The work described in this chapter has been published in two peer-reviewed journal articles [2, 3]. The NRC 3D VMI spectrometer is a unique instrument designed specifically to perform 3D measurements using the camera-and-digitizer method [25], while improving upon previous designs by maximizing effective momentum resolution with a "thick-lens" time-stretching electrode stack. While I joined this project after the NRC spectrometer construction was complete and thus had no involvement in the design and construction process, I include some of the design considerations here as necessary background information and to contextualize my work on the instrument. The camera

and digitizer that are currently used with the spectrometer were purchased during my time working on this project, and with my direct involvement. I discuss in detail the considerations which informed these purchases. We used two separate laser systems for the experiments discussed here. These produced femtosecond- and picosecond-length pulses and I refer to them as the femtosecond laser and the picosecond laser. I describe the optical setup used with both laser systems. Finally, I summarize each of the 3D VMI data sets that will be discussed in this thesis. These consist of photoionization experiments with atomic xenon, methyl iodide (CH_3I) and nitric oxide (NO).

The design of the NRC spectrometer is mostly credited to Dr. Iain Wilkinson. The 3D VMI data described in Section 3.6 were collected with the assistance of Dr. Albert Stolow, Dr. Andrey Boguslavskiy, Dr. Aurelien Sanchez, and Douglas Moffatt.

3.2 The NRC 3D VMI Spectrometer

In 3D electron VMI, the full 3D photoemission distribution can be determined directly, without the need for inversion, by simultaneously measuring the 2D position and arrival time (or TOF) of each photoelectron at an imaging detector. Typically, for the camera-and-digitizer approach used here, a high gain MCP detector is used in combination with a phosphor anode screen and high-speed camera to determine the 2D position, while a capacitive anode pickoff from the phosphor yields electron pulses which determine the electron TOF [30]. The 2D position yields (x, y), whereas the TOF is directly related to z, the coordinate along the flight axis. A key advantage of 3D VMI over other 3D recoil

momentum vector imaging techniques, such as those based on delay-line anode detectors, is the ability to readily characterize multiple electron events per laser shot, as demonstrated in Chapter 4. In this multi-event case, it is necessary to correctly match each 2D position with its corresponding TOF. This is achieved via correlations in the output charge variation of the MCP for single charged particle input, represented by its pulse height distribution (PHD). The broad PHD, typically $\sim$ 100% FWHM (Full Width at Half Maximum) for a two-plate chevron MCP, ensures that the amplified output charge from single events will sample a wide range of values. If an integrated 2D position signal and an integrated TOF signal correspond to the same total charge, they likely belong to the same event and are assigned to each other, thus yielding the 3D coordinates (x, y, t) of that event (this process is described in detail in Section 4.3). Resolving the 2D position of each event is straightforward since two electrons from the same laser shot rarely arrive at the same position on the MCP detector. Therefore, a key challenge in this multi-event approach is to resolve, in time, the individual electrons (which typically arrive within a few nanoseconds of each other) while maintaining the 2D velocity mapping conditions. In this section, I present the design and implementation considerations of the NRC 3D VMI spectrometer, which is based on a 13-plate electrode stack which stretches the photoelectron TOF distribution while retaining good 2D VMI focusing conditions.

3.2.1 Spectrometer and Data Acquisition Overview

In Figure 3.1 I present a schematic diagram of our experimental data acquisition setup and spectrometer. I have used circled numbers to label sections of the figure and these are described sequentially below:

(1) — A skimmed molecular beam (vertical blue dashed lines) from an Even-Lavie pulsed supersonic valve [56] intersects a weakly focused ($f = 75$ cm) laser beam (horizontal purple dashed line) inside a turbo-pumped (Pfeiffer TMU 521 YP) vacuum chamber (grey box), producing photoelectrons. The pulsed valve was operated at 1 kHz to match the repetition rate of the laser. We typically used a pulse duration of 16–19 μs and heated the valve to a temperature of 60°. We were able to switch between two different laser systems for different experimental requirements; these are described in detail in Sections 3.4 and 3.5. The pulsed valve is located in a lower turbo-pumped (Edwards STP-A2203CS) source chamber (not pictured) which is separated from the main chamber by a Beam Dynamics, Model 1A, 1 mm skimmer. The upper and lower chambers can be isolated from each other using a gate valve with an accessible exterior control. When the pulsed valve is not running, the pressures of the upper and lower chambers are kept below 5×10^{-9} and 5×10^{-8} torr, respectively. During operation of the valve, the pressures are typically around 2×10^{-7} torr (upper chamber) and 2×10^{-4} torr (lower chamber).

(2) — The 3D PAD, represented by the black dashed oval, is accelerated upward (toward the detector) by the electric fields generated from the 13-element electrode stack (black). Three example photoelectrons (red circles) and their trajectories (solid black lines) are

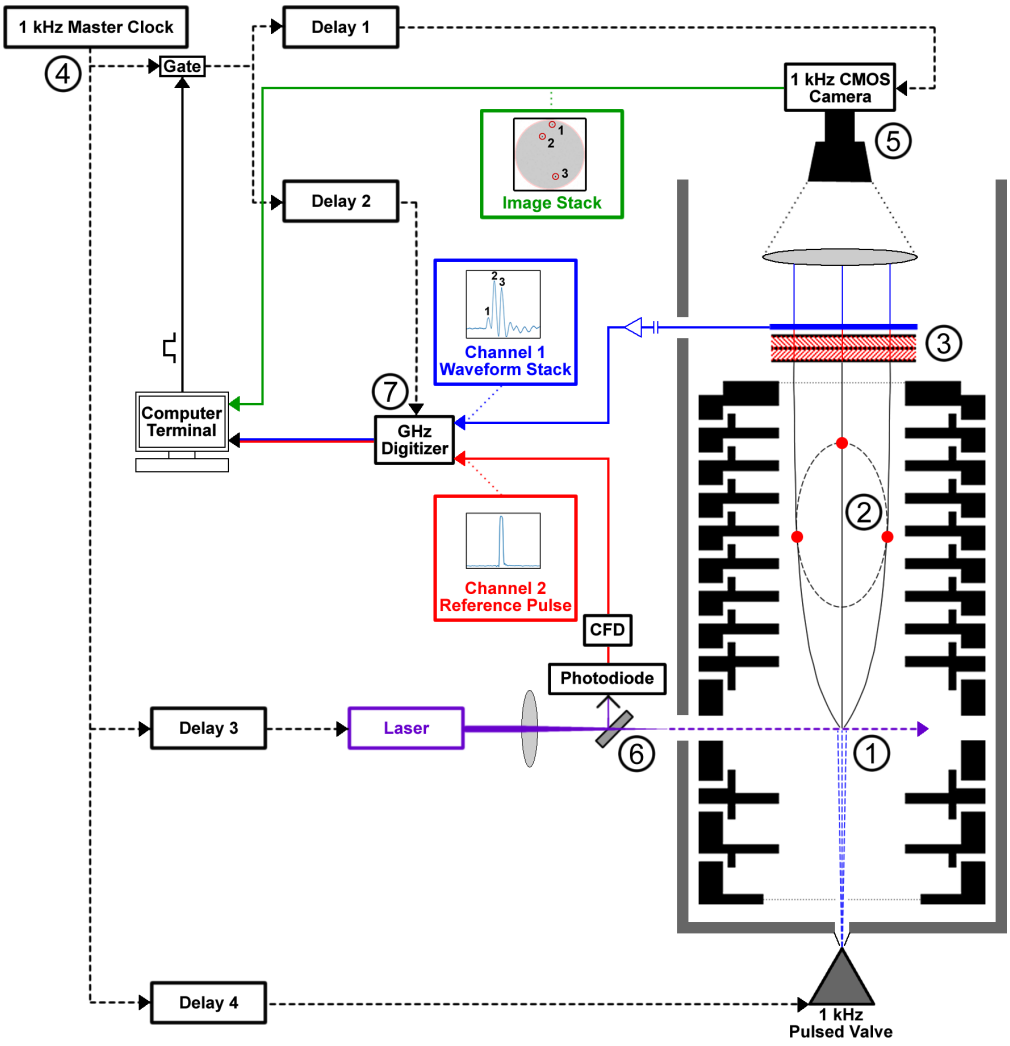

Figure 3.1: Schematic diagram of the triggering and data acquisition setup for the NRC 3D VMI spectrometer. Sections are labelled by circled numbers and each is described in detail in the text. (1) The interaction region between the laser pulses and skimmed molecular beam from the pulsed valve. (2) The photoelectron angular distribution being focused toward the detector by the 13-element electrostatic lens. (3) The MCP/phosphor detector assembly with a capacitive pickoff to the digitizer. (4) The 1 kHz master clock which triggers (after delays 1–4) the camera, digitizer, laser, and pulsed valve. (5) The CMOS camera which images the phosphor screen. (6) The photodiode timing reference to remove trigger jitter. (7) The GHz digitizer which records the pickoff waveforms and reference pulses.

shown for illustration purposes. The rectangular electrode segments above and below the laser axis both represent a single electrode with laser entrance and exit ports in the plane depicted in the figure. There are also identical unused ports in the axis perpendicular to the image. I describe the electrostatic lens system in more detail below in Section 3.2.2.

(3) — The electrons reach the chevron dual plate MCP (double-stack of red lines) and accompanying P47 phosphor (thick blue line), where they generate flashes of blue light (peak wavelength at 430 nm) which are imaged by the camera. At the same time, electrical impulses are created when the amplified electrons hit the phosphor screen, and these are decoupled through a 100 pF pickoff capacitor, propagated along an *in vacuo* SMA cable and doubled-ended SMA feedthrough, amplified (10× gain, 1.3 GHz analog bandwidth) and sent to the digitizer (blue arrow). The MCP/phosphor assembly (Photonis APD 2 PS 40/12/10/8 I 60:1 EDR P47) is discussed in more detail below in Section 3.2.3.

(4) — The precise timing of the laser and pulsed valve necessary to ensure that the laser pulse and molecular beam overlap in the interaction region, as well as collection and transfer of the data which must occur along with each laser pulse, are controlled by a 1 kHz master clock. The clock is used to generate four trigger channels, each with a different delay. The laser's switch-out Pockels cell (delay 3) is triggered through an analog delay setup which is not typically changed. The delay for the pulsed valve (delay 4) is set using a digital delay generator (Stanford Research Systems DG535), which allows us to easily adjust the timing for different gas sources. To allow time for the gas to reach the interaction region, the pulsed valve is triggered using the previous clock pulse. The triggers for the CMOS camera (delay 1) and GHz digitizer (delay 2) are set using a second DG535, which is gated

by the data acquisition computer so as to toggle between data acquisition and standby states. We collected data in batches of N laser pulses then paused for block data transfer to the hard drive. Typically N = 64,000 (64 seconds) in recent experiments, although some early data sets were collected with N = 4000. The laser and pulsed valve ran uninterrupted at 1 kHz throughout the block transfer cycles so as to maintain stability. While this downtime in acquisition means that data collection is not truly "real-time", the large value of N nonetheless results in a duty cycle greater than 90%.

(5) — The phosphor emissions were focused by a lens located inside the vacuum chamber (Edmund Optics MgF_2 AR-coated, 75 mm diameter, 300 mm focal length, N-BK7/N-SF5 achromatic doublet) and imaged by a 25 GigE monochrome CMOS camera (Emergent Vision Technologies, HB-2800-S-M). While the camera has a full resolution of 1936×1464 pixels, we run it with a reduced output resolution of 512×512 pixels and an 8-bit analog-to-digital converter (ADC) resolution so as to accommodate continuous acquisition which must match the laser repetition rate of 1 kHz. Acquisition of image data begins when the camera is initialized and instructed to capture a total of N images, after which it waits to be triggered by the master clock through delay 1. During the standby state after acquiring a batch of data, the stack of images (typically 64,000) is transferred to the main computer (green arrow). An example image with three electron events is shown in the green-bordered box. Writing the image stack to disk is slow and inconvenient due to the large data volume. To reduce this to manageable levels, the images are processed in real time (see Section 4.2.1.1) in order to identify features above a set threshold and only a 31×31 pixel region of interest (ROI) enclosing each feature is saved to disk. These

saved ROIs are subsequently treated by more time-consuming post-processing in order to yield accurate 2D positions and amplitudes. It should be noted that even with the high sensitivity afforded by our chosen camera (see Section 3.3.1), we are still operating at the low end of the camera's dynamic range. That is, with our MCP detector, a single hit on the sensor will typically register a peak (single pixel) value of less than 10 on an 8-bit (0–255) ADC. However, this is not a major impediment due to a combination of low readout noise and our ability to make use of the integrated full width of the single hit image spot (typically about 10 pixels FWHM). This allows us to accurately determine the position and amplitude (total charge) of single electron hits, as detailed in Section 4.2.1.2.

(6) — A photodiode and constant fraction discriminator (CFD) were used to detect the photoionization laser pulse, producing a precisely timed and highly stable electrical signal which is independent of any electronic trigger jitter. An example reference pulse is shown in the red-bordered box. This reference pulse is sent to the second (or first) channel on the digitizer (red arrow) where it is stored and processed as described in Section 4.2.2.2. Prior to the implementation of this photodiode-based timing reference, we estimated our jitter-limited TOF resolution to be ~ 210 ps (Gaussian σ). This agrees with the expected jitter from our digitizer, which has a sampling interval of 250 ps (4 GHz), since any trigger pulse arriving between intervals would be subjected to a delay of up to this amount before the first sample is taken. After the photodiode reference was added, we estimated our electron TOF resolution to be < 80 ps.

(7) — The pulses from the phosphor pickoff and the reference signal from the photodiode are both recorded by a 1.7 GHz analog bandwidth, 12-bit 4 GS/s digitizer card (AlazarTech,

ATS9373). An example pickoff waveform containing three electron events is shown in the blue-bordered box. We can operate the digitizer in either one-channel mode (where the reference pulse is co-added with the phosphor pickoff using a microwave splitter) or two-channel mode (where the reference pulse is recorded on its own channel). In Figure 3.1 the two-channel mode is illustrated with the reference pulse in channel two. In two-channel mode, the sampling rate of the digitizer is reduced to 2 GS/s on each channel; if the full 4 GS/s sampling rate is required the one-channel mode can be used instead. However, care must be taken to ensure that the reference pulse does not overlap in time with the photoelectron events. During each standby period, the stack of digitized waveforms were transferred to the data acquisition computer where post-processing (see Section 4.2.2.2) was later performed to determine the precise arrival time and amplitude of each event.

3.2.2 Thick-Lens Design and Simulations

The TOF resolution for a single electron event is limited by electronics and the transit time across the anode. Thus, the ability of this 3D VMI technique to resolve the arrival times of multiple events is limited by the TTS, which we define as the difference in arrival time between the first- and last-arriving particles in the spherical momentum distribution. Our spectrometer uses a 13-electrode "thick-lens" design to stretch the distribution of particles in time and improve multi-hit capability compared to a standard three-plate VMI design.

Much of the functionality of the spectrometer is determined by the conditions in and near the interaction region, which is defined, in our experiments, by the intersection of the

pulsed molecular beam, which travels along the TOF (z-axis) direction, and the laser beam with its propagation vector, $\vec{k}$, aligned along the y-axis. During the conceptualization stage, the SIMION 8.0 charged particle optics simulation package [55] was used to investigate photoelectron energy resolution and TOF spread for several different 3D VMI designs. In these simulations, the volume of the interaction region was modelled as a cylinder of length 2 mm (along the y-axis) and diameter 0.1 mm. The interaction region was Monte Carlo sampled as a 2D Gaussian distribution with 0.1 mm $1/e^2$ widths along the x and z (radial) axes of this cylinder, and as a uniform distribution over the 2 mm length of its y-axis (a uniform gas density approximation).

The simulations sampled thousands of electron trajectories from isotropic PADs with discrete kinetic energy values ranging from 0 to 5 eV. For each electron trajectory, the 2D position (x, y) at the MCP and the TOF (related to z) were recorded. The VMI energy resolution was assessed using the following standard criterion employed in 2D VMI. For electrons with kinetic energy E, which focus to a ring of radius r and width Δr at the MCP detector, the 2D positions are binned into a digital image, inverted using a standard Abel inversion algorithm, and radially integrated. The resulting radial spectrum is fitted with a Gaussian function to determine the peak width and position; the centroid of the Gaussian is taken as the 3D radius R, with ΔR being the width (FWHM). The photoelectron energy resolution is then given by $\Delta E/E = 2\Delta R/R$ (as described in Section 2.2.2).

The TTS is largely determined by the electric field strength in the interaction region (the extraction field). In order to stretch the electron distribution in time, the extraction field strength must be reduced. However, other constraints apply, namely: (i) the extraction

field must be strong enough that electrons emitted perpendicular to the TOF axis can still be collected and imaged on the MCP detector; and (ii) the VMI focusing condition must be maintained. The final design of our 3D VMI spectrometer is shown in Figure 3.2. Additional details of the types of SIMION simulations which led to this design are presented elsewhere [4]. An important advantage of the 13-electrode design is in its significantly improved TTS.

Although the full simulations involved thousands of electron trajectories, a discussion of selected SIMION results helps to illustrate the features of this design. In Figure 3.2 (a), I present an expanded view of the 3D VMI spectrometer, showing the results of simulations (for 5 eV photoelectrons) discretely sampling the interaction region defined above. Simulated electron trajectories (blue) are shown with illustrative electrons in transit (red dots), helping to visualize the improved electron TOF spread Δt. These trajectories were produced from nine electron origin points within this interaction volume (as shown in the lower-left inset). Three sets of three origin points, horizontally displaced along the laser axis, were located ± 1 mm from, and exactly on, the cylindrically symmetric TOF (z) axis. Within each set, the three origin points were placed ± 50 μm from, as well as on, the laser axis. In this illustrative figure, each origin point spawned eight electron trajectories (blue) with initial velocities of equal magnitude sampling, in 45° steps, the polar angle θ. The upper-left inset shows an expanded view of electron trajectories at the MCP, demonstrating the VMI focusing conditions achieved with the 13-electrode design. At 5 eV, the SIMION results for our 13-electrode spectrometer indicate a $\Delta E/E$ of 0.85%, compared with 1.3% in a simulation of the standard three-plate VMI design. The 13-electrode design also shows

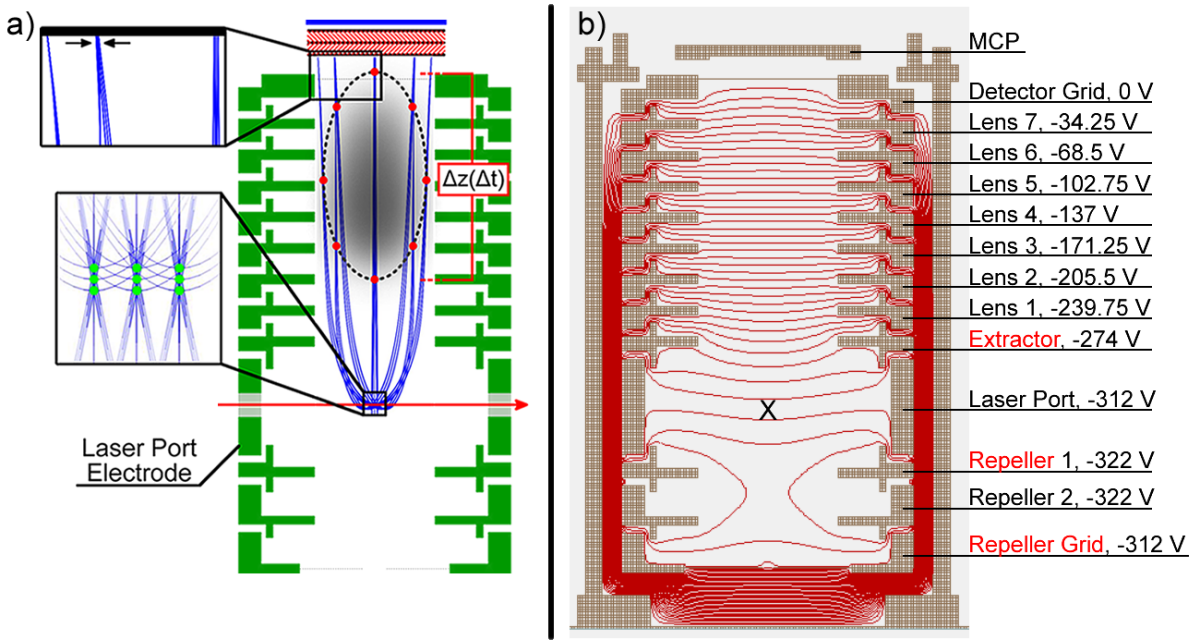

Figure 3.2: Expanded view of the 13-electrode NRC 3D VMI spectrometer design illustrating simulated (using SIMION) electron trajectories (a) and equipotential lines (b). In (a), trajectories (blue) with examples of time stamps (red dots) are shown to visualize the elongated photoelectron distribution Δz corresponding to a TOF spread Δt. The upper inset shows a close-up of electron trajectories at the MCP detector, demonstrating that VMI focusing conditions are retained with the 13-electrode design. The lower inset shows the origins (green dots) of the simulated trajectories (for details, see the text). Note that the cross-section depicts the laser axis (red) passing through the laser port electrode; the segments directly above and below the laser axis together comprise this single electrode, as shown in (b). Each of the labelled electrodes are described in detail in the text. Electrodes labelled with red text (repeller grid, repeller 1, extractor) are the only ones with exterior voltage control. The example voltages shown in (b) are for a maximum photoelectron kinetic energy of ~ 2.2 eV and are the conditions used in the nitric oxide experiments (see Section 3.6.3). The interaction region in (b) is marked by the **X**.

significantly reduced chromatic aberration relative to the standard three-plate VMI. It is important to note that this design is not restricted to 3D VMI measurements; the fields can easily be configured for 2D VMI, with performance comparable to that of a dedicated 2D instrument [4].

SIMION results show that, for 5 eV photoelectrons, our 13-electrode design has a TTS of $\Delta t = 8.6$ ns while maintaining VMI focusing conditions. The analogous standard three-plate VMI would have a TTS of 1.8 ns. Given that the electron TOF resolution is limited, an increased TTS results in improved z-coordinate resolution (corresponding to the polar angle θ of the photoelectron distribution). Furthermore, for typical MCP detectors, the output electron pulse duration is on the order of 2 ns FWHM, suggesting a pulse-pair resolution (the minimum resolvable time difference between two pulses) of roughly half of this value. I note, however, that deconvolution using the instrumental single-hit time response function can lead to improved single-pulse and pulse-pair resolution (see Section 4.2.2.1). Most importantly, an extended TTS allows for the detection and time-stamping of several electron events per laser shot.

As shown in Figure 3.2, the laser axis passes through a (large) 2 cm diameter hole in the laser port electrode, thus minimizing scattered UV light which can produce background photoelectrons from metal surfaces. A set of optical baffles helps to reduce the UV scatter. The 40 mm diameter MCP is mounted at a distance of 125 mm from the laser axis. Figure 3.2 (b) shows the same set of electrodes as in (a), but with SIMION-generated equipotential lines and labels for each of the 13 electrodes. As an example of typical settings, the voltages used in the nitric oxide experiments (described in Section 3.6.3) are included next to the

electrode labels. With these settings, the average electron TOF is ∼30 ns. The electrodes labelled with red text (repeller grid, repeller 1, and extractor) are the only ones for which the voltage can be set externally; the other electrodes are resistively coupled to maintain the same voltage ratio. There is also an adjustable offset voltage, which is applied to all of the electrodes (including the detector grid, normally at 0 V). This is used when imaging low-energy photoelectrons such as in the xenon experiments (described in Section 3.6.1) to ensure that the electrons are accelerated into the MCP with the amount of energy required for efficient detection.

Above the laser port electrode sits a stack of eight electrodes, the extractor and lens electrodes (Lens 1–7), followed by the top (mesh-covered) detector grid electrode defining the end of the TOF region. Finally, a flat-field electron acceleration or deceleration region before the MCP permits optimally efficient photoelectron detection. These eight lens electrodes are 3 mm thick at their inner edges and have an internal aperture diameter of 50 mm. The extractor and repeller 1 electrodes on either side of the thicker laser port electrode are separated by 44 mm. The distance from the bottom (mesh-covered) repeller grid electrode to the laser axis is 54 mm.

The detector grid and repeller grid electrodes (top and bottom) are made from CuBe in order to match their thermal expansion coefficients with that of the attached OFHC copper meshes. All other electrodes, the mesh electrode holders, and the lens mounting assembly are manufactured from minimally-magnetic and chemical-reaction-resistant 316 stainless steel. The repeller grid (bottom) electrode is covered by a flat, gold-plated electroformed copper mesh which has a laser-machined 3 mm diameter central hole for skimmed molecular

beam transmission. The molecular beam was verified to pass through the mesh without interference, as confirmed by the observation of cold clusters in the interaction region. The detector grid (top) electrode is covered by a 92% transmission gold-plated copper mesh. These meshes are gold-plated for electrical uniformity and chemical inertness, as well as to increase their work functions. The voltage varies linearly from the extractor to the detector grid electrode, with the latter set to ensure that the electrons transmitted though the detector grid are accelerated to 300–500 eV for maximum detection efficiency by the MCP. The three repeller electrodes below the laser port serve to mimic a flat repeller plate while simultaneously displacing the bottom grid from the laser axis, thus reducing spurious electron generation by scattered UV light. As an additional measure to reduce the effects of scatter, the repeller grid electrode is set at a slightly higher (more positive) voltage so that scattered electrons are pulled toward the electrode and thus away from the flight region. The TOF tube is enclosed by single-layer mu-metal magnetic shielding, reducing external magnetic fields by a factor of 1000.

3.2.3 MCP/Phosphor Detector

Our ability to interpret data containing multiple electron events per laser shot is contingent upon a correct determination of the 2D (x, y) position (from the CMOS camera image) and matching TOF (from the transient digitizer waveform) for each event in the shot. To accomplish this, we make use of the nature of MCP amplification as an avalanche process described by a gain PHD. Each single electron event will sample the broad PHD

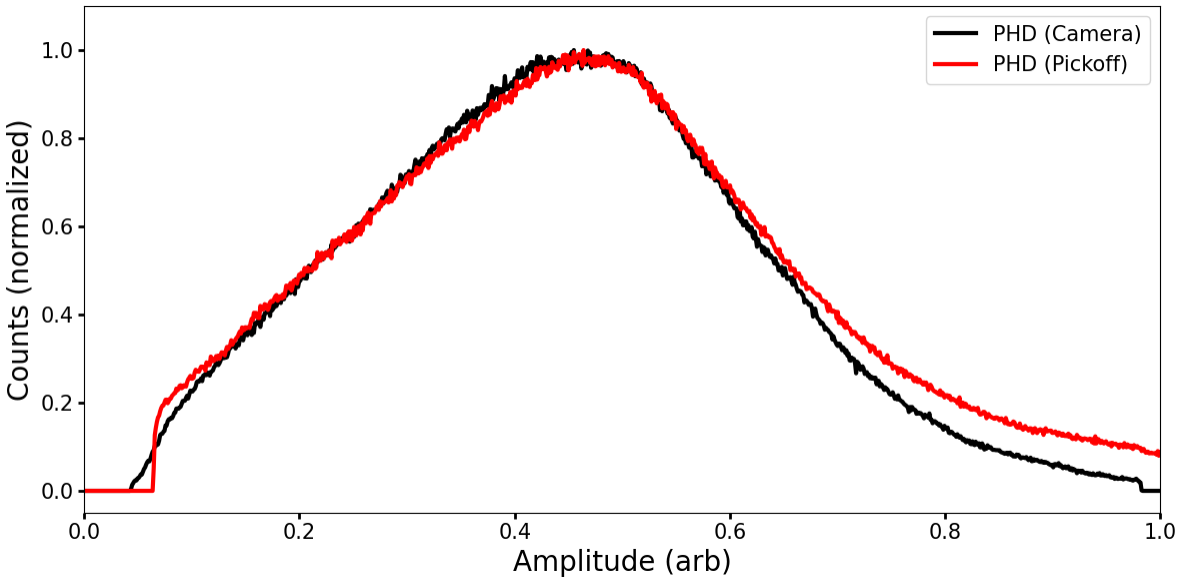

Figure 3.3: Pulse height distribution (PHD) for the MCP detector used in the NRC 3D VMI spectrometer. The singular PHD is reflected in both the integrated amplitudes of events detected in the camera images (black) and the fitted heights of events detected in the pickoff traces (red). Both represent the same PHD and are individually normalized.

and the resulting total charge will be reflected in the integrated amplitudes of both the camera and pickoff (digitizer) signals. As an example, for a laser shot resulting in three electron events, the event that experiences the highest MCP gain will produce the highest amplitude pulse on the digitizer waveform and, concomitantly, the most intense spot on the camera image. Similar arguments apply for the next and lowest gain electron events. For this reason, 3D VMI benefits from an MCP with high gain but a broad PHD, the latter being considered undesirable in most event counting applications. The PHD of our MCP detector, as independently derived from each of the data channels (camera and pickoff), is shown in Figure 3.3.

The unique way in which our 3D VMI application depends upon the precise performance characteristics of the MCP means it is valuable to discuss the specifications of our detector in detail. Our MCP/phosphor detector assembly was purchased from Photonis (now Exosens) and has the advanced performance detector (APD) product code: APD 2 PS 40/12/10/8 I 60:1 EDR P47. Note that this code may not be the same format currently used for Exosens detectors. The code indicates that the MCP has a chevron configuration (two stacked MCPs with opposing bias angles), the extended dynamic range option, and an attached P47 phosphor screen. The diameter, pitch (centre-to-centre pore distance), pore diameter, and bias angle (pore angle relative to the surface) are 40 mm, 12 μm, 10 μm, and 8°, respectively. The MCP thickness to pore diameter ratio is 60:1. For a 60:1 chevron MCP, the Photonis documentation recommends a maximum voltage of 2400 V (1200 V per plate); in our experiments we used a voltage of 2200 V. We also used a voltage of 5000 V (maximum) for the phosphor screen to achieve maximum brightness on the camera.

To improve detector performance, we have acquired a custom-engineered MCP with characteristics designed to produce a higher gain (with a corresponding loss of spatial resolution) and broader PHD. The loss in spatial resolution is acceptable since we accurately centroid each hit in the images to obtain the 2D peak position (see Section 4.2.1.2 for details). End-spoiling was performed to widen the cones of exit from the back of the pores and the spacing between the two plates in the chevron configuration were increased from 50 μm to 100 μm. Both of these changes have the effect of increasing gain and should broaden the PHD, but also significantly reduces the spatial resolution. The pitch and pore diameter were reduced to 6 and 5 μm to improve resolution and partially offset the effect of

the end-spoiling. Additionally, gold plating was added to potentially narrow the width of the pickoff pulses by reducing the transit time of the pulses across the face of the detector. Testing of this custom MCP confirmed an increase in PHD from $\sim 100\%$ to $\sim 170\%$ and a $\sim 10\times$ increase in gain.

3.3 Choice of Camera and Digitizer

The raw data from a 3D VMI experiment are collected from two primary sources: a camera imaging the flashes of light from the phosphor-backed MCP detector (containing the spatial information), and a digitizer recording the electrical impulses picked off from the phosphor (containing the temporal information). For each of these essential pieces of equipment (camera and digitizer), there are many choices available to suit different experimental needs. It was necessary to compare and contrast these options when deciding which instruments to purchase for our 3D VMI experiments. In this section, I will summarize the considerations and key factors which informed our decision. It should be noted that the camera was purchased in November of 2019 and the digitizer was purchased in April of 2021. Thus, the following represents the state of the market and technology at those respective times.

3.3.1 Camera Considerations

For a camera to be suitable for our 3D VMI application it must satisfy certain requirements. The phosphor screen backing the MCP detector is P47 (Y_2SiO_5:Ce), which has a short

emission lifetime of 70 ns and a central emission wavelength of 430 nm (above 50% of maximum intensity between 390 and 475 nm) [57]. Many parameters of camera sensors are wavelength-dependent, and it is important to evaluate these near the peak emission region whenever possible. The laser system operates at a repetition rate of 1 kHz so the camera must be capable of acquiring, transmitting, and storing data at this speed. The P47 phosphor has a low emission intensity; since we are imaging a very weak light source, our primary concern regarding the camera's performance was in its low-light sensitivity and noise characteristics. Following this, the next concern was spatial resolution, for which we prioritized the largest chip size that can still meet the above requirements. This camera was intended to replace a previous 3D VMI camera which was determined to not have the necessary level of sensitivity. The old camera was an IO Industries Flare 12M180 high-speed CMOS camera that uses the CMOSIS CMV12000 sensor. This sensor was used as a baseline for comparisons and the new camera was required to show considerable improvements in sensitivity over the CMOSIS sensor.

3.3.1.1 Camera Performance Parameters

There exist many parameters to characterize the performance of camera sensors and different manufacturers report their specifications differently. The EMVA1288 standard [58], established by the European Machine Vision Association, provides a set of precisely defined parameters for sensor performance, and instructions for how to measure and present them in standardized datasheets. However, EMVA1288 datasheets are not easily accessible

for all sensors, particularly newer high-end chips. Below is a summary of common sensor parameters that were considered, with an emphasis on the EMVA1288 definitions where applicable.

Quantum Efficiency, QE — $\eta(\lambda)$ **(%)** The quantum efficiency is the fraction of incoming photons that are converted into electrons by the sensor. It is defined as the ratio (usually expressed as a percent) $\eta(\lambda) = \mu_e/\mu_p$, where an average number of photons μ_p of wavelength λ hit the sensor during its exposure time, producing an average number of electrons μ_e. QE is highly wavelength-dependent and is often reported as a plot, or as a single value at the wavelength where the QE is maximum. High QE is important for our application since it represents a fundamental limitation on the low-light sensitivity. In the EMVA1288 standard, QE accounts for loss due to fill factor (non-sensing area per pixel) and is the only directly wavelength-dependent parameter.

Overall System Gain — K **(DN/e$^-$)** The overall system gain is the average digital output value (in DN) per electron in the pixel well. The dimensionless quantity DN is the discrete value assigned to the output by the ADC. There are two sources of electrons which are subjected to the gain, the electrons produced from the conversion of incident photons (μ_e) and the electrons present with no light (the dark signal, μ_d). Thus, the gain is described by the relation $\mu_y = K(\mu_e + \mu_d)$, where μ_y is the average digital output value. The average digital output value with no light is $\mu_{y.dark} = K\mu_d$. It is important to note that the meaning of DN depends on the ADC resolution (*i.e.* bit depth). That is, 8-bit

DN/e^- and 12-bit DN/e^- are not the same and they should be converted before any direct comparisons are made. To emphasize this, I will add a subscript to DN (DN_8, DN_{12}, *etc.*) when the distinction is important.

Temporal Dark Noise, TDN — σ_d (e^-) or $\sigma_{y.dark}$ (DN) In the EMVA1288 standard, there are three different contributions to the total output noise: shot noise, dark noise, and quantization noise. Shot noise (σ_e) results from the quantum mechanical fluctuations in electron count described by Poisson statistics and is equal to $\sqrt{\mu_e}$ for all sensors. Quantization noise (σ_q) depends only on the ADC resolution and arises from the difference between the analog value after application of the gain and the nearest discrete digital value to which it is assigned. Dark noise, or temporal dark noise, includes all other sources of noise in the sensor and is reported as two related forms: σ_d, which is the noise in the number of electrons before the ADC, and $\sigma_{y.dark}$, which is the noise in the digital value after the ADC and also includes the quantization noise. Neither form includes the shot noise. The total noise (σ_y) is given by:

$$\sigma_y^2 = K^2(\sigma_d^2 + \sigma_e^2) + \sigma_q^2. \tag{3.1}$$

As with gain, when comparing dark noise in DN the ADC resolution must be the same.

Saturation Capacity — $\mu_{p.sat}$ (p) or $\mu_{s.sat}$ (e^-) Saturation refers to the state where the pixel well is filled to its maximum capacity. There are multiple similar but distinct parameters which characterize the saturation limit. In the EMVA1288 standard the full-

well capacity is the maximum number of electrons the well can hold. A lower limit called the saturation capacity ($\mu_{s.sat}$) is also defined, which is the maximum number of electrons for which the fluctuations caused by noise will not significantly reach the full-well capacity. The saturation exposure $\mu_{p.sat}$ is the number of photons (abbreviated as p) which corresponds to the saturation capacity (accounting for the QE). This parameter is not important in our low-light application since we are not operating near saturation.

Signal-to-Noise Ratio — SNR (dB) The signal-to-noise ratio is the ratio of the mean photo-induced signal ($\mu_y - \mu_{y.dark}$) to the total noise σ_y (given by equation 3.1). It is a dimensionless quantity and is often reported in dB ($20 \log_{10}$ SNR). In the EMVA1288 standard the reported quantity is SNR_{max}, the maximum achievable SNR, which corresponds to a signal at the saturation capacity. While SNR is of high importance in our low-light application, the SNR_{max} parameter is not directly relevant since we do not operate near the saturation limit.

Absolute Sensitivity Threshold, AST — $\mu_{p.min}$ (p), $\mu_{p.min.area}$ (p/μm^2), $\mu_{e.min}$ (e$^-$), or $\mu_{e.min.area}$ (e$^-$/μm^2) The absolute sensitivity threshold, or minimum detectable exposure, is defined in the EMVA1288 standard as the mean number of photons (per pixel) that result in an SNR of one. This is the minimum incoming signal that can be detected above the noise level and directly represents the low-light sensitivity of a sensor, accounting for the effects of QE and noise. For this reason, it is the most important parameter for our application. The AST is typically given as a number of photons ($\mu_{p.min}$) or as p/μm^2

($\mu_{p.min.area}$) which is the previous quantity divided by pixel area. Care must be taken when comparing sensors with different pixel size; for example, if the image will be focused onto the same number of pixels regardless of the size of the sensor, $\mu_{p.min}$ will be the important quantity. The AST can also be given in terms of electron rather than photon number, and this is simply the corresponding photon value multiplied by the QE.

Dynamic Range — DR (dB) The dynamic range is a measure of the full range of light levels detectable by the sensor. In the EMVA1288 standard, it is defined as the ratio of the saturation capacity to the AST and is often reported in dB ($20 \log_{10}$ DB). As with the saturation capacity, this is not an important parameter for low-light applications.

Fixed Pattern Noise (FPN) Fixed pattern noise is a term used to describe variation in pixel response between regions of the sensor. This is not noise in the temporally-varying sense, and in the EMVA1288 standard this variation is described instead by two nonuniformity parameters: DSNU (dark signal nonuniformity) and PRNU (photo-response nonuniformity). This type of "noise" is not a major concern since our image processing steps (described in Section 4.2.1.2) involve subtracting an average background level and average column and row projections from each image.

3.3.1.2 Comparing Camera Sensors

Our primary requirements of high sensitivity coupled with continuous acquisition of high-framerate data led us to consider cameras from Emergent Vision Technologies (EVT). EVT

offered a unique combination of the new Sony Pregius series of CMOS sensors and a fast 25 GigE (Gigabit Ethernet) interface. The third generation of the Pregius series, which follow the IMX4## naming scheme, were advertised as having higher QE and lower noise compared to previous designs. The available data support this claim, although the lack of EMVA1288 datasheets for all of the sensors made direct comparisons difficult. Of the products offered by EVT at the time, there were four cameras using sensors small enough (and thus fast enough) to be considered. Some of the properties of these cameras are summarized in Table 3.1. To aid in comparing sensors, I have defined a quantity ROI_{1000} which is the maximum square resolution that the camera can support at a framerate of 1 kHz (only a single dimension is given since they are the same). Because of the way in which these sensors are read (row-by-row), only the second dimension of the resolution will affect readout speed. This is the reason why a smaller sensor will be faster even if we restrict the readout region. Additionally, although the first dimension of the resolution can be expanded up to the maximum for that sensor without compromising speed, we are only interested in imaging a square region and will likely restrict it for convenience.

The four sensors in Table 3.1 are different sizes and the full-resolution framerate decreases with size, as one would expect. ROI_{1000} tends to be larger for cameras with a higher full-resolution framerate, since more of the full resolution can be used without lowering the framerate below 1000 Hz. An exception to this is the HB-500-S-M (IMX426), which can already run at 1000 Hz using its full resolution. Based solely on sensor size, the HB-1800-S-M (IMX425) would be the optimal choice.

I was only able to find EMVA1288 datasheets for two of the sensors in Table 3.1: the

Table 3.1: Four cameras from Emergent Vision Technologies (EVT) which were considered for use in the 3D VMI system. For each EVT camera, the Sony Pregius sensor (IMX4##), full resolution, pixel dimensions, framerate at full resolution, and maximum square resolution at 1000 Hz (ROI_{1000}) are listed.

Camera	Sensor	Resolution	Pixel Size (μm)	Framerate (Hz)	ROI_{1000}
HB-500-S-M	IMX426	812×620	9×9	1586	620
HB-1800-S-M	IMX425	1604×1100	9×9	660	660
HB-2000-S-M	IMX422	1624×1240	4.5×4.5	485	560
HB-2800-S-M	IMX421	1936×1464	4.5×4.5	410	540

IMX425 (EVT model HB-1800-S-M) and the IMX421 (EVT model HB-2800-S-M). In Table 3.2 I have summarized the relevant EMVA1288 data for these sensors, as well as for the CMOSIS CMV12000 sensor used in the old camera. Additionally, I have included data on the CMOSIS CMV4000 (which should be similar to the CMV12000) since the available EMVA1288 data for the latter is incomplete. Note that these parameters were measured at different wavelengths (ranging from 530–555 nm) as specified in the QE column. EMVA1288 data are not available near 430 nm (the wavelength of the phosphor emissions) since this is far from the peak of the QE curve. The Sony Pregius sensors should all have a similar wavelength dependence, so comparisons made at the QE peak should also be valid at other wavelengths. The CMOSIS sensors do have a different QE curve, but since they have a lower QE over all wavelengths of interest, direct comparisons are unnecessary.

The IMX421 is the only sensor for which I could find explicit mention of, and data for, two different gain settings: high conversion gain (HCG) and low conversion gain (LCG).

Table 3.2: Summary of EMVA1288 data for camera sensors that were considered for use in our 3D VMI instrument. The IMX425 and IMX421 (part of the Sony Pregius series) are the only two sensors of the ones under consideration (see Table 3.1) for which I could find EMVA1288 data. The other two sensors listed in Table 3.1 likely have similar properties to the corresponding sensors in this table with the same pixel size; the IMX426 should be similar to the IMX425 and the IMX422 should be similar to the IMX421. This table also includes the old camera sensor (CMV12000) for comparison, and the similar CMV4000 sensor which should have similar properties and for which more complete data was available. FLIR [59] provides two sets of data for the IMX421 sensor, corresponding to high (HCG) and low (LCG) conversion gain settings. The gain (K) values reported by FLIR were stated as being in units of e^- which is likely an error; I treated these values with the assumption that they were in units of DN_{16}/e^- as indicated elsewhere in the cited document. Values marked with an * were converted from the reported ADC resolution into DN_{12} for ease of comparison. The dark noise ($\sigma_{y.dark}$) for the CMV12000 sensor (marked by **) was obtained by multiplying $K \times \sigma_d$ since it was not given explicitly. According to their respective sources, the IMX421 and CMV4000 sensors were tested at 530 nm, the IMX425 was tested at 536 nm, and the CMV12000 was tested at 555 nm.

Sensor	Pixel Size (μm)	QE, η (%)	Gain, K (DN_{12}/e^-)	σ_d (e^-)	$\sigma_{y.dark}$ (DN_{12})	$\mu_{p.min.area}$ p/μm^2	$\mu_{p.min}$ p/pixel
IMX425 [60]	9×9	69	0.041	22.16	0.95	0.42	34.3
IMX421, LCG [59]	4.5×4.5	71	0.157*	5.93	0.93*	0.45*	9.11
IMX421, HCG [59]	4.5×4.5	69	0.363*	2.67	0.97*	0.23*	4.58
CMV12000 [61]	5.5×5.5	50	0.44*	13	5.72**	—	—
CMV4000 [59]	5.5×5.5	49	0.392*	16.53	6.48*	1.16*	35.11

It is likely that other members of the same series (IMX4##) would also have this feature but I could not find any sources which confirm this. For the data from FLIR [59] the gain (K) was given in e^- although I believe this is an error and the values are actually in DN_{16}/e^-. This is supported by statements made elsewhere in the document, and by the fact that the values are similar to those of other sensors when converted to DN_{12}/e^- under this assumption. I have therefore treated these values as DN_{16}/e^- when constructing Table 3.2.

From Table 3.2 it can be seen that the Sony sensors have peak QE values near 70%, which is a significant improvement over the old camera (CMV12000) which has a peak QE of 50%. The differences between the two Sony sensors (without considering the HCG setting on the IMX421) are mostly a result of their different pixel sizes. Since the IMX425 pixels are four times larger than those of the IMX421, the gain and σ_d electron noise are respectively decreased and increased by a factor of four. The $\sigma_{y.dark}$ digital noise combines the gain and electron noise and is similar for both. The AST, which is the most important parameter for our application, is similar for both sensors (again, considering only the LCG setting for the IMX421) when considering photons/area. However, due to the larger pixel size, the IMX425 has a much worse AST per pixel, requiring $\sim 4\times$ more photons per pixel to equal the noise level. This means that if we focus an image onto the same area of both sensors, the sensitivity will be approximately the same, but with the spatial resolution of the IMX425 reduced by a factor of four. If instead the same number of pixels are used, the sensitivity for the IMX425 will be four times worse. For this reason I did not recommend the IMX425, or other sensors with a large pixel size, for our application.

After the above considerations, the remaining options were the similar IMX421 and IMX422 sensors, used in the HB-2800-S-M and HB-2000-S-M cameras from EVT, respectively. The main advantage of the IMX422 is that it has a larger ROI_{1000} (560×560) compared to the IMX421 (540×540). This is, however, a small difference, even more so since we smooth and fit the peak positions (*i.e.* centroiding) to obtain sub-pixel resolution (see Section 4.2.1.2). When the purchasing decision was being made, EMVA1288 data were not directly available for the IMX422 sensor, and I could not confirm that it had a high conversion gain (HCG) setting. The IMX421 was the safer option since EMVA1288 data were available and the HCG setting was confirmed. With the only advantage of the IMX422 being a small improvement in resolution, I decided to recommend the IMX421 (EVT HB-2800-S-M) for purchase.

3.3.2 Digitizer Considerations

Prior to obtaining our current digitizer card (AlazarTech, ATS9373), we used an external Tektronix DPO7354 oscilloscope. This oscilloscope was capable of sampling at a rate of 10 GS/second, which was a faster sampling rate than we required for our application. However, there were problems with data storage and transfer which prompted us to look into a more appropriate PCIe digitizer card to install in the lab computer. When assessing potential cards to purchase, we needed to know the minimum bandwidth required. I investigated this by applying Fourier low-pass filters at different cutoff frequencies to determine the point at which there was a noticeable drop in signal quality. This process is shown in

Figure 3.4.

A typical trace taken from the oscilloscope is shown in (a), sampled at a rate of 10 GS/s. The Fast Fourier Transform (FFT) of this trace is shown in (b) along with two fourth-order Butterworth filters (at 0.25 GHz and 0.50 GHz). Although I performed this test with more than these two filters, I have only shown two here for illustration purposes. The original and filtered traces are shown in (c), with a zoomed-in time scale to focus on the peaks. The corresponding FFTs are also shown in (d). After applying the filters, it is clear from both the time- and frequency-domain plots that the 0.25 GHz filter (red) results in a severe loss in resolution, while the 0.5 GHz filter (blue) does not have any significant negative effect. I concluded from this that we could safely purchase a digitizer with a bandwidth of 0.5 GHz or greater. This lead us to consider and acquire the current 1.7 GHz analog bandwidth, 4 GS/s AlazarTech card described in Section 3.2.1 above.

3.4 Femtosecond Laser Optical Setup

The femtosecond laser source, used for the xenon (Section 3.6.1) and CH_3I (Section 3.6.2) experiments, was a Coherent Legend Elite Duo Ti:Sapphire laser which produced 800 nm, short ($\sim$ 40 fs) pulses with a repetition rate of 1 kHz. We used $\sim$ 0.7 mJ of the laser output to produce 200 nm pulses (fourth harmonic) by the nonlinear optical scheme shown in Figure 3.5. The beam diameter was narrowed using a 2:1 reflecting telescope and sent into the fourth harmonic generation setup. The fourth harmonic was produced by second harmonic generation (SHG), followed by two successive sum frequency mixing (SFM) steps

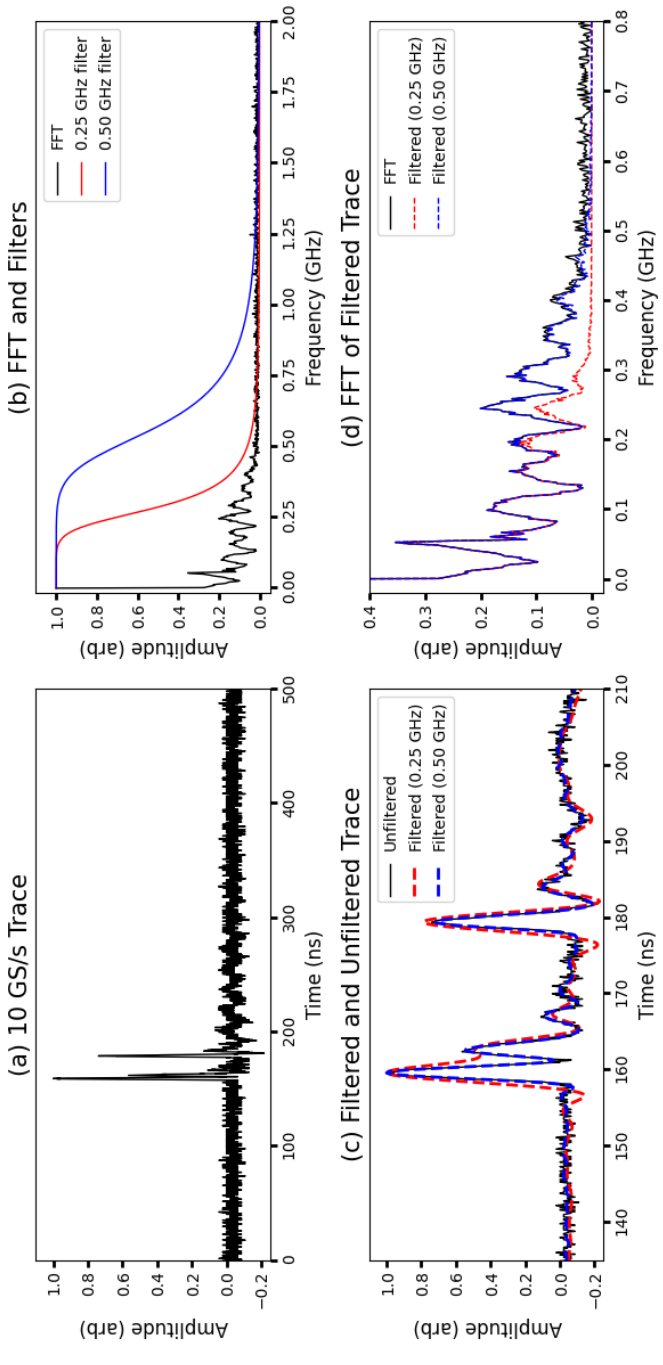

Figure 3.4: Determining the required bandwidth for the 3D VMI digitizer card. A typical experimental trace sampled at 10 GS/s (a) obtained using a Tektronix DPO7354 oscilloscope was filtered using several low-pass Butterworth filters to assess the minimum cutoff frequency which would not reduce signal quality. Two of these filters (fourth order, 0.25 and 0.50 GHz) are shown as examples (b) along with the Fourier transformed (FFT) trace. The original and filtered traces using both filters (c) and their corresponding FFTs (d) show that the 0.50 GHz filter does not reduce data quality and the 0.25 GHz filter does.

using the fundamental.

In detail, the fundamental was split into two branches, one of which was used to generate the third harmonic (THG) as shown in the lower box in Figure 3.5, and the other was sent through an adjustable delay stage to control pulse overlap during the final SFM step. The THG setup started with an initial BBO (β-BaB_2O_4) crystal (Type-1, 0.2 mm, 29.2°) for SHG (producing 400 nm pulses). This was followed by a calcite crystal (a-cut, 0.75 mm) to compensate for dispersion and regain pulse overlap and a half waveplate at 800 nm to rotate the polarization of the fundamental for phase-matching with the second harmonic. A second BBO crystal (Type-1, 0.1 mm, 44.3°) was used for SFM to produce the third harmonic (267 nm). Finally, SFM of the third harmonic with the delay-branch fundamental using another BBO crystal (Type-1, 0.05 mm, 64.8°) yielded fourth harmonic pulses (200 nm).

We also had a manual flip mirror mounted before the VMI chamber. The purpose of this was to direct the beam toward a pinhole positioned outside the spectrometer that was accurately measured to be at the same position relative to the mirror as the interaction region is within the spectrometer. This allowed us to easily align the beam against the pinhole before switching it into the VMI chamber. The laser entrance and exit ports of the VMI chamber were mounted with thin (1 mm) CaF_2 windows.

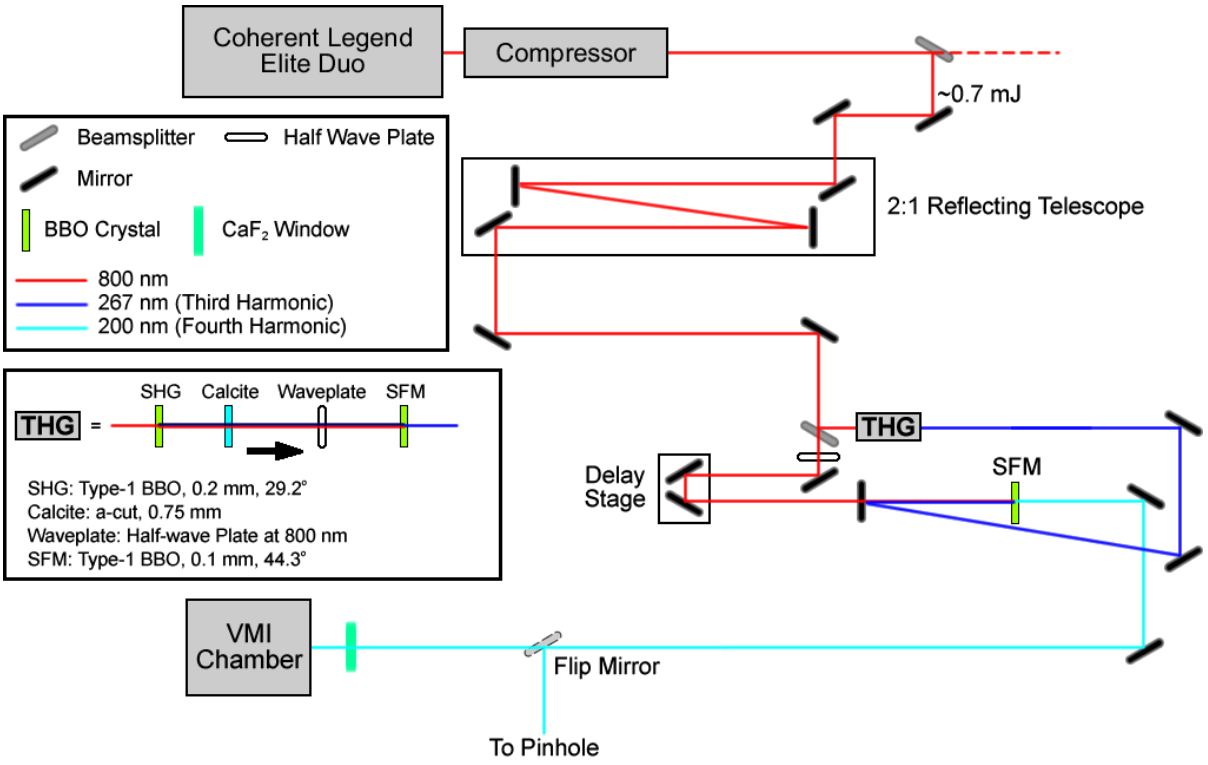

Figure 3.5: Layout of the optical setup used for femtosecond laser experiments. We generated $\sim$ 40 fs, 800 nm short pulses at 1 kHz using our Coherent Legend Elite Duo laser system. We narrow the beam diameter with a 2:1 reflecting telescope before sending it to the fourth harmonic generation setup. The fourth harmonic is generated by SFM of the (delay-controlled) 800 nm fundamental with the third harmonic (267 nm) as described in detail in the text. The lower inset shows how the THG optics are set up. A flip mirror allows the beam to be redirected to a pinhole for alignment.

3.5 Picosecond Laser Optical Setup

The laser system used for the nitric oxide (NO) experiments (Section 3.6.3) was a home-built Nd:YVO_4 Grazing Incidence Slab Laser, seeded by a 150 ps Nd:YAG SESAM mode-locked oscillator (Time-Bandwidth Products), yielding amplified ($\sim$1 mJ) 1064.7 nm, 150 ps transform-limited pulses at a repetition rate of 1 kHz. The fourth (266.2 nm) harmonic of this output was generated by two successive SHG steps and the fifth harmonic (212.9 nm) was generated by SFM of the fundamental with the fourth harmonic. Similarly to the femtosecond laser setup, we used a flip mirror to an external pinhole for alignment. The 1 mm thick entrance window into the VMI chamber, which is the exit window used by the femtosecond laser, is VUV grade CaF_2.

For the main feature (at 0.88 eV) studied in the NO experiments, the fifth harmonic ($\sim$ 213 nm) was used as the pump pulse (which dissociated the $(NO)_2$ dimers) and the fourth harmonic ($\sim$ 266 nm) was used as the probe pulse (which single-photon ionized the long-lived NO(A, 3s) photofragment products). All three NO photoionization processes used for the picosecond laser experiments are described in detail in Section 3.6.3. This laser was used for the NO experiments because the combination of narrow bandwidth ($\sim$ 0.2 cm^{-1}) and ps pulse duration ($\sim$ 75 ps) ensures that the measured 3D momentum distribution reflects the instrumental response rather than laser parameters. A 266 nm waveplate was used to freely rotate the linear polarization of the probe pulse relative to the 3D VMI TOF axis. This allowed us to generate non-cylindrically-symmetric photoelectron

angular distributions with which to test our 3D VMI capabilities. The time delay between the 213 nm photodissociation and 266 nm probe pulses was 2.1 ns. This ensured that the second (photoionization) pulse arrived well after the first, by which time the asymptotic photoproduct distribution had fully formed.

3.5.1 Echo Signal

One notable complication that was introduced by the picosecond laser system is the presence of an echo signal observed at a time delay of $\sim$ 1.8 ns after the main signal. An example of this is shown in Figure 3.6 (a) for an NO measurement with photoionization rings at 0.05 eV (inner) and 0.88 eV (outer). Figure 3.6 presents the 3D VMI data as a 2D histogram of photoelectron coordinates (radius and arrival time) after the data processing steps described in Chapter 4. Here, radius refers to the 2D detector-plane radius and the zero of the time axis is arbitrary. Short times indicate electrons emitted toward the detector (the front of the distribution).

The echo is highly consistent in its timing relative to the main signal and is always weaker in intensity. After investigating potential causes, we isolated the problem to the laser output. The picosecond laser system switches out a series of three pulses each time it fires, rather than one. These pulses are separated by the cavity round trip time of 1.8 ns, matching the observed delay of the echo signal. The primary (middle) output pulse is stronger than the two surrounding pulses. We were not sensitive to the early pulse since it was the weakest of the three. While we attempted to reduce the amplitude of the late

pulse, we could not make it weak enough to remove its effects from our data.

When analyzing data with the echo signal present, we removed it by simply subtracting a time-shifted and amplitude-scaled copy of the photoelectron distribution as shown in Figure 3.6 (b). The time shift of $\sim$ 1.8 ns was consistent, but the amplitude of the subtracted copy was adjusted manually by inspection during analysis to account for shifts in the relative heights of the main and secondary laser output pulses. As can be seen in panel (b), subtracting the scaled copy, which naturally also includes the echo, results in a negative impression of the echo displaced by an additional $\sim$ 1.8 ns. This is not an optimal treatment and could be further improved, but practically we did not encounter any negative effects from using this method.

3.6 Summary of 3D VMI Data

We collected and analyzed a number of experimental data sets using the 3D VMI spectrometer and laser systems described above. These were used to test, characterize, and demonstrate our implementation of 3D VMI. To construct a robust body of data for these purposes, we performed measurements using three different sample sources: atomic xenon, methyl iodide, and nitric oxide. In this section, I have summarized the data sets collected from each of these sources.

The first set of measurements, consisting of six data sets, was made with the following three goals: to test and demonstrate (1) the benefits of time-stretching on our ability to

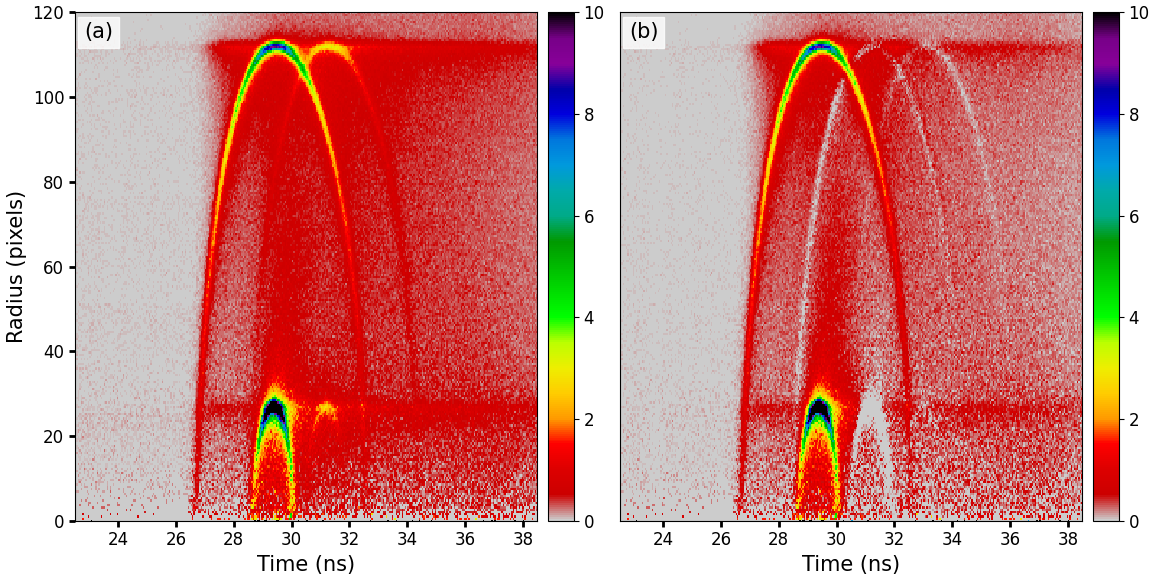

Figure 3.6: Treatment of the echo signal present in data taken with the picosecond laser system, arising from a secondary laser output pulse. The data shown here are from photoionization of nitric oxide (NO) and contain two photoelectron features (0.88 and 0.05 eV). (a) A 2D histogram of processed data showing the photoelectron distribution in terms of arrival time and 2D detector-plane radius. The echo is displaced by ~ 1.8 ns from the main signal (the laser cavity round trip time) and has a lower amplitude. (b) We remove the echo before analysis by subtracting a time-shifted and scaled copy of the distribution. The amplitude of the subtracted copy is determined by inspection.

efficiently process multi-hit data, (2) the extent to which our data processing and analysis methods can handle high hit-count data (more than three hits/shot), and (3) the ability of our 3D VMI spectrometer to acquire high-quality data using arbitrary (non-cylindrically-symmetric) polarization geometries. Goal (1) was addressed by collecting data at four different levels of time-stretching: 6.2 ns (NO, three data sets), 6.9 ns (CH_3I), 8.8 ns (Xe), and 31 ns (Xe), where the given times are the TTS of the photoelectron distributions. Goal (2) was addressed by collecting data that included hit counts up to $\sim$ 15 hits/shot (the 31 ns Xe and CH_3I data). Goal (3) was addressed by imaging a well-studied NO photoionization feature at three different polarization geometries, two of which are not cylindrically symmetric.

The second set of measurements were made with the goal of developing a robust transformation from the raw photoelectron detector coordinates (x, y, t) into the initial recoil momentum vectors (p_x, p_y, p_z). The transformation was to be tested by fitting the PADs and comparing the results with previously published data. We again used the well-known photoionization of nitric oxide, this time measuring three features with peak photoelectron kinetic energies of 0.05, 0.88, and 2.05 eV.

For each of the data sets in the first set of measurements, I present a visualization of the resulting PADs in the form of XY slices along the z (TOF) axis (Figures 3.7–3.10). The data were processed as described in Chapter 4. The slices in these figures are not constructed from the image data directly. Rather, they are made by binning the processed 3D (x, y, t) coordinate data by arrival time and plotting the (x, y) positions within each time bin as a 2D histogram. Note that with the full 3D coordinates available, I also have

the ability to produce 2D slices in the XZ (XT), YZ (YT), or any other arbitrary plane. In each case, the time width of the slices is 160 ps. The slices shown here are primarily for visualization purposes and have been smoothed to enhance their qualitative features. As such, the rings associated with specific electron kinetic energy features are artificially broadened and their width does not reflect our actual energy resolution. The rings all appear to be highly circular as expected for a VMI image, with the small exception of the 31 ns xenon data set, as described below. Since the time spread of these experiments is much longer than the width of the slices, I have the ability to produce many more slices than the few shown in the figures presented here. Each of the full time-binned data sets have been animated as videos and will be made available as supplementary material.

Panel (a) in Figures 3.7–3.9 and panels (a), (g), and (m) in Figure 3.10 are histograms of the photoelectron yield over time. Each of these show a significantly higher background level at the back end of the distribution (later arrival times) compared to the front end. This background signal is caused by photoelectrons produced from scattered light hitting the inner walls of the electrostatic lens system and is challenging to remove [62] when using UV light, where the photon energies exceed the work functions of the electrostatic lens components.

To better visualize and compare the shape and level of time-stretching for each of the data sets in the first set of measurements, I have plotted them as stacks of XY-plane slices in Figure 3.11. In Figure 3.11, the PADs are oriented so that the flight direction is up. It is clear, particularly for the xenon experiments, that the raw (x, y, t)-data are not spherical. This is a result of the nonlinear VMI focusing fields used for time-stretching.

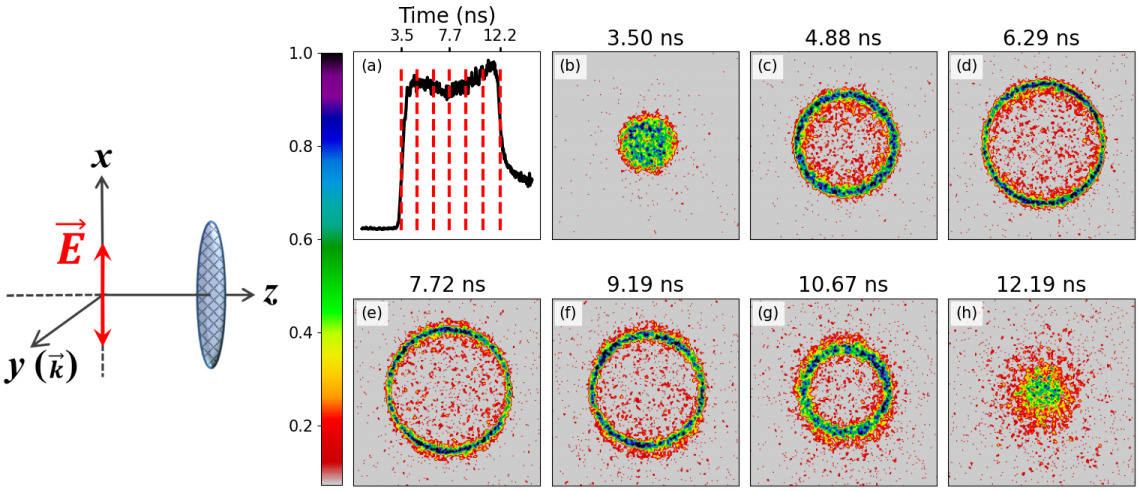

Figure 3.7: Time-stamped 3D electron VMI data recorded following two-photon ionization of xenon at 200 nm, resulting in a photoelectron distribution with a peak signal intensity at 0.27 eV. A ∼4.3 V/cm extraction field was implemented, resulting in an electron time spread of 8.8 ns. The coordinate system on the left highlights the experimental geometry, with the polarization vector of the linearly-polarized light aligned along the x-axis, the propagation axis of the laser ($\vec{k}$) along the y-axis, and the time-of-flight axis (z) directed toward the MCP detector. (a) Photoelectron yield over time at the detector, with vertical lines indicating the time bins used in the following slices. (b)–(h) A few exemplary time slices of the 3D photoelectron distribution through the XY plane, constructed by binning the processed (x, y, t)-data by arrival time and smoothing to highlight the qualitative features. Panels (b)–(d) show electrons emitted toward the detector and panels (f)–(h) shows electrons emitted away from the detector. Each slice has a temporal width of 160 ps.

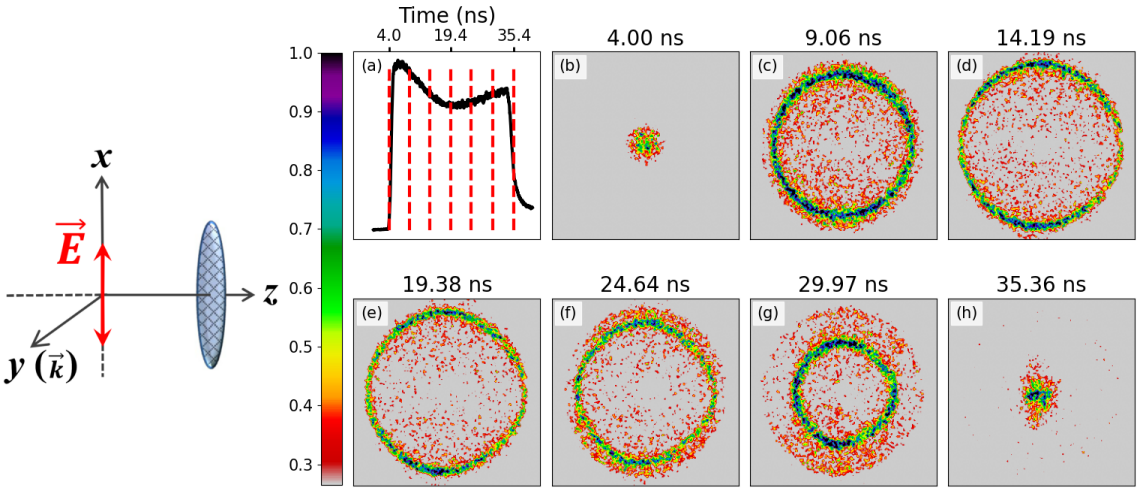

Figure 3.8: Time-stamped 3D electron VMI data recorded following two-photon ionization of xenon at 200 nm, resulting in a photoelectron distribution with a peak signal intensity at 0.27 eV. A ∼1.6 V/cm extraction field was implemented here, resulting in an electron time spread of 31 ns. The coordinate system on the left highlights the experimental geometry, as in Figure 3.7. (a) Photoelectron yield over time at the detector, with vertical lines indicating the time bins used in the following slices. (b)–(h) A few exemplary time slices of the 3D photoelectron distribution through the XY plane, constructed by binning the processed (x, y, t)-data by arrival time and smoothing to highlight the qualitative features. Each slice has a temporal width of 160 ps.

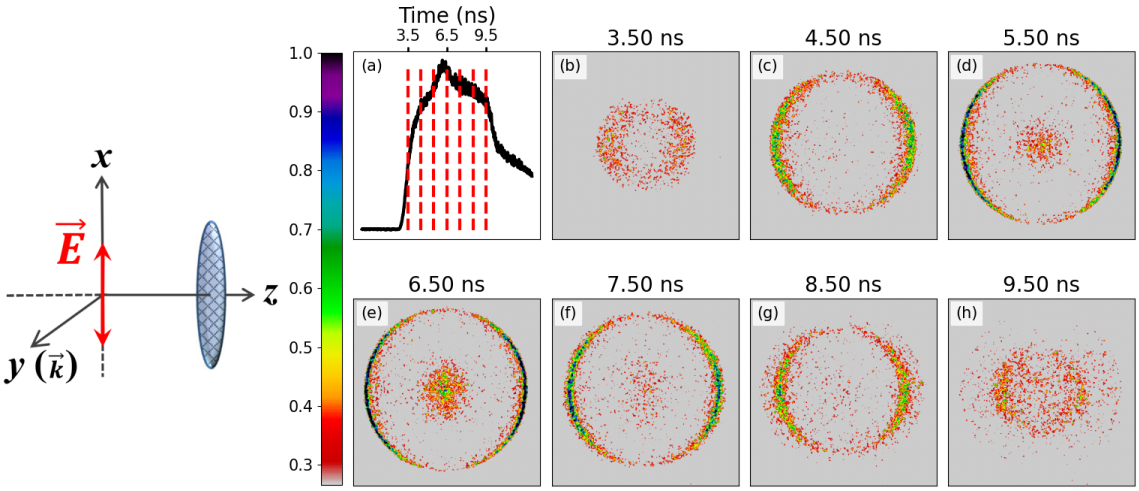

Figure 3.9: Time-stamped 3D electron VMI data recorded following two-photon ionization of CH_3I at 200 nm, resulting in a photoelectron distribution with a peak signal intensity at 2.78 eV. A ∼16 V/cm extraction field was implemented here, resulting in an electron time spread of 6.9 ns. The coordinate system on the left highlights the experimental geometry, as in Figure 3.7. (a) Photoelectron yield over time at the detector, with vertical lines indicating the time bins used in the following slices. (b)–(h) A few exemplary time slices of the 3D photoelectron distribution through the XY plane, constructed by binning the processed (x, y, t)-data by arrival time and smoothing to highlight the qualitative features. Each slice has a temporal width of 160 ps.

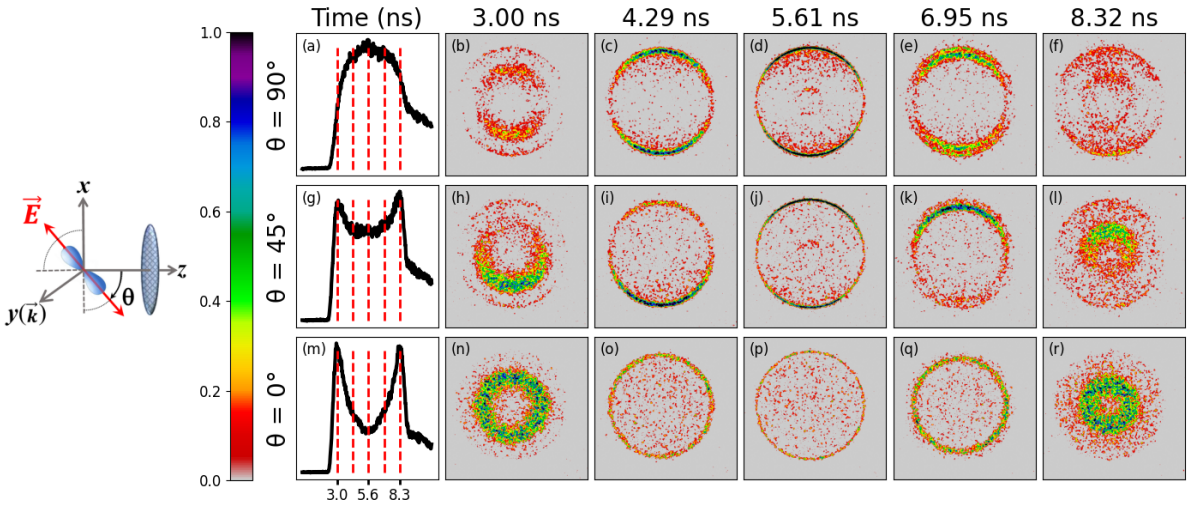

Figure 3.10: Time-stamped 3D electron VMI data for photodissociation of $(NO)_2$ at 213 nm and subsequent photoionization of the NO(A) photoproducts at 266 nm, resulting in a photoelectron distribution with a peak signal intensity at 0.88 eV. An ∼11 V/cm extraction field was implemented here, resulting in an electron time spread of 6.2 ns. This ionization scheme produces a nearly-pure p-wave photoelectron distribution aligned with the probe laser polarization $\overrightarrow{E}$ (at angle θ with the z-axis), as shown in the coordinate system to the left, where y ($\vec{k}$) is the laser propagation direction and z is the time-of-flight axis. To demonstrate our 3D imaging capability, we used three linear polarization cases: $\theta = 0°$ (bottom row), 45° (middle row), and 90° (top row). The 90° case is the only one accessible to conventional inversion-reliant 2D VMI. The first image on each row shows the photoelectron yield plotted as a function of time for each case. The images to the right of these plots show a few exemplary time slices of the 3D photoelectron distribution through the XY plane, constructed by binning the processed (x, y, t)-data by arrival time and smoothing to highlight the qualitative features. Each slice has a temporal width of 160 ps and is centred at the times indicated by the red vertical lines in the left panels (and noted at the top of each column).

The distortion from a spherical distribution is most extreme for the low-energy (0.27 eV) xenon data, where the electrons spend more time exposed to the focusing fields. One consequence of this distortion is that reconstruction of the initial photoelectron momentum vectors (described in Section 4.4) becomes more challenging.

For the xenon and CH_3I experiments, we collected image and waveform data for 800,000 laser shots (200 stacks of 4000 shots each). This corresponds to about 13.3 minutes of real-time data at a laser repetition rate of 1 kHz. However, due to the data transfer overhead between each stack of 4000 shots, the actual data collection time was about 37 minutes. As mentioned above in Section 3.2.1, we have greatly reduced this overhead since these early measurements were made. For the NO experiment, each of the three data sets consisted of 200,000 laser shots (50 stacks of 4000 shots each), corresponding to approximately 3.3 minutes of data and 9 minutes of collection time.

3.6.1 Xenon Data Sets

The first test case for our 3D VMI apparatus and data processing was the two-photon ionization of xenon at 200 nm:

$$\mathrm{Xe}\, 5p^6(^1S_0) + 2h\nu \rightarrow \mathrm{Xe}^+\, 5p^5(^2P_{3/2}) + e^-. \tag{3.2}$$

As mentioned above, we performed two experiments using this photoionization process, with the only difference being the implemented electrostatic lens voltages and the

Figure 3.11: Visualizations of the data sets shown in Figures 3.7–3.10, represented as stacked XY-plane slices (0.5 ns width) through each distribution. The xenon data sets (a)–(b), CH_3I data set (c), and NO data sets (d)–(f) have peak photoelectron kinetic energies of 0.27, 2.78, and 0.88 eV, respectively. Each PAD is oriented with the TOF axis in the vertical direction and the front of the distribution (electrons emitted toward the detector) on the top. The transverse (x, y) and time axis are to scale across all PADs, although the colour map is normalized separately for each. For data sets with a large amount of time-stretching (a) and/or with low energy (a)–(b), the distortion of the PADs from a spherical shape is clearly seen.

resulting degree of electron-distribution time-stretching (8.8 ns and 31 ns). These experiments allowed us to directly compare the effects of time spread on our results and analysis efficiency.

In these experiments, we only see ionization into the lower ($^2P_{3/2}$) of the two spin-orbit states of the ion [63], since two photons of 200 nm (12.40 eV) is less than the $^2P_{1/2}$ excited spin-orbit state ionization energy of xenon (13.44 eV). The ionization energy of xenon to the $^2P_{3/2}$ state is 12.13 eV and this results in a low, dominant photoelectron kinetic energy of 0.27 eV.

We used a linear laser polarization perpendicular to the TOF axis in these measurements, which is the polarization condition necessary for performing conventional 2D VMI using Abel inversion. The average hit count per laser shot was 1.0 and 8.7 hits/shot for the 8.8 ns and 31 ns experiments, respectively. The gas mixture used for the 8.8 ns TTS xenon experiment was 2% xenon in helium at a stagnation pressure of 5 bar, and for the 31 ns TTS xenon experiment we used 5% xenon at a stagnation pressure of 2 bar.

The analyzed and time-stamped xenon data are presented in Figure 3.7 (8.8 ns TTS experiment) and Figure 3.8 (31 ns TTS experiment) in the form of slices through the 3D photoelectron distribution where the processed (x, y, t)-data are binned by their arrival time. The TTS for each can be seen in the photoelectron yield over time, as shown in Figure 3.7 (a) and Figure 3.8 (a). For these experiments, the linear light polarization (x) axis is aligned in the vertical direction, which can be seen in the photoelectron distribution (most visible in panels (d) and (e) of each figure) where there is slight visible anisotropy

with higher photoelectron intensity along the $\pm x$-axis.

3.6.1.1 Symmetry-Breaking Distortion

Our SIMION simulations predict that with optimized electrode settings we would expect a TTS greater than 40 ns for this xenon photoionization process. However, when we attempted such an experiment, we found that stretching beyond 31 ns produced significant distortions in the shape of the photoelectron distribution, shown in Figure 3.12. This manifested primarily as a loss of circular symmetry in the electron rings, which became more square in shape and drifted away from the centre of the spectrometer. This effect was most severe for the later-arriving electrons (bottom of panel (a)), which spend more time in the extraction region of the VMI spectrometer. The full temporal extent of this photoelectron distribution is actually much longer than predicted by the simulations: $\sim$ 75 ns instead of 40 ns. This is further evidence of a departure from the expected behaviour. In Figure 3.12 (a), the 3D representation of the PAD shows a highly distorted shape which grows in radius near the bottom and drifts from the central axis. The integrated 2D VMI image in panel (b), viewed along the z (flight) axis, shows the square-shaped profile of the photoelectron ring. The distortions are also present, albeit to a lesser degree, in the 31 ns results (Figure 3.8).

I believe that the reason for this distortion effect is the broken symmetry caused by the entrance and exit apertures in the laser port electrode. This is supported by the direction of the corners of the distortion, which line up with these ports. Panel (b) of Figure

3.12 shows the laser propagation axis (labelled k) and the perpendicular polarization axis (labelled E), both of which coincide with apertures in the instrument (the ports located along the polarization axis being unused in this experiment). Furthermore, the severity of the distortion increases with the degree of time-stretching; as the electron-projection field gradient is reduced to allow the slow electron distribution to expand along the flight axis and detector face, the electrons move more slowly through the electrostatic lens, come closer to the laser port electrode, and are more susceptible to stray and inhomogeneous fields in the instrument. Thus, these trajectories are more sensitive to the presence of the electrode openings. Note that such electron distribution distortions are absent in higher electron kinetic energy data, such as that presented in Figures 3.9 and 3.10, for which the instrument was designed. We plan to address this low-electron-kinetic-energy issue in the next version of our 3D VMI spectrometer, which is currently under construction.

3.6.2 Methyl Iodide Data Set

The second test system that we used is the two-photon ionization of methyl iodide (CH_3I) at 200 nm [64], which results in 2.78 eV electrons:

$$CH_3I\,(X\,^1A_1) + 2h\nu \rightarrow CH_3I^+\,(X\,^2E_{3/2}) + e^-. \tag{3.3}$$

This process allowed us to demonstrate 3D imaging of photoelectrons with a significantly higher energy than that of the xenon measurements described above. In this experiment we were able to achieve a time spread of 6.9 ns and an average hit count of 3.2

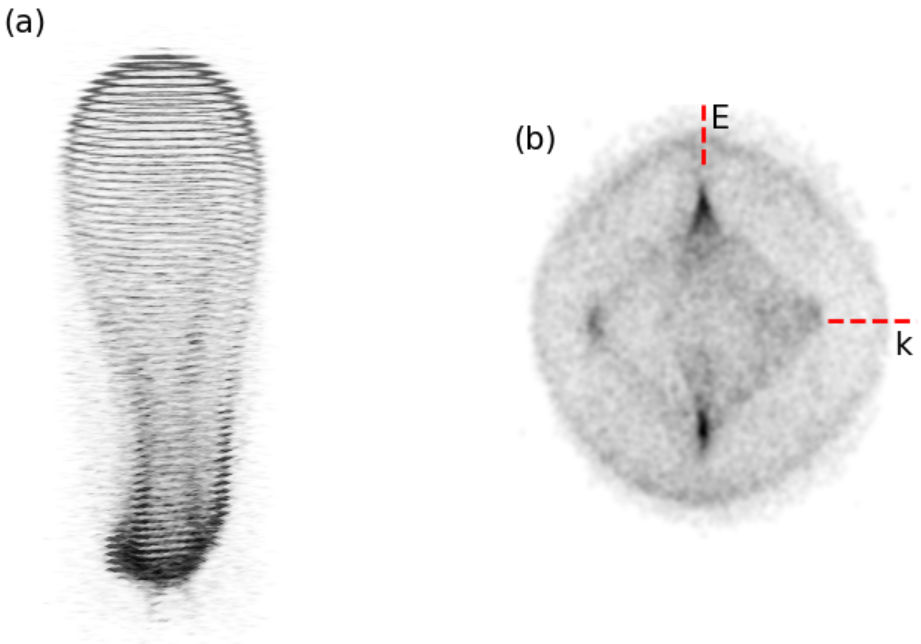

Figure 3.12: Distortions in the PAD of xenon when imaged with maximum time-stretching. The predicted TTS for this experiment was 40 ns, although the measured TTS was ~ 75 ns. (a) The full PAD, represented as a stack of XY-plane slices (1 ns width), as in Figure 3.11. The distortion can be clearly seen at the back end (bottom) of the distribution, where the distribution drifts away from the central axis. (b) An integrated 2D image of the entire data set, viewed along the z (flight) axis. The xenon photoelectron signal at 0.27 eV, which should be circular, has been distorted into the shape of a square. The corners of the square line up with the laser entrance and exit ports (along the axis labelled k) and the unused perpendicular ports (along the axis labelled E).

hits per laser shot. As with the xenon experiments, we used a linear laser polarization perpendicular to the flight axis, allowing us to also perform an inversion-based analysis of the data. We used a gas mix of 2% CH_3I in helium at a stagnation pressure of 1 bar.

The analyzed and time-stamped CH_3I data are presented in Figure 3.9 in the form of slices through the 3D photoelectron distribution. The 6.9 ns TTS can be seen in the photoelectron yield over time (Figure 3.9 (a)). The linear light polarization (x) axis is aligned in the vertical direction, as evidenced by the clear anisotropy in the photoelectron distribution (Fig 3.9 (b)–(h)) with higher photoelectron intensity along the $\pm y$-axis.

3.6.3 Nitric Oxide Data Sets

The third and most important test system that we used was nitric oxide (NO), with which we performed two sets of measurements using several photoionization processes. The first set of measurements used the one-photon photodissociation of the nitric oxide dimer $(NO)_2$ at 213 nm followed by subsequent one-photon ionization of the resulting NO(A, 3s) excited-state photoproducts at 266 nm, yielding 0.88 eV photoelectrons [65]. The second set of measurements used the above photodissociation process, followed by the one-photon ionization of the NO(A, 3s) photoproducts at 266 nm (again yielding 0.88 eV photoelectrons) and 213 nm (yielding 2.05 eV photoelectrons), as well as two-photon ionization of the NO(X) ground state monomer at 266 nm (yielding 0.05 eV photoelectrons). The photodissociation step is described as follows:

$$(\mathrm{NO})_2 + \mathrm{h}\nu_{213nm} \rightarrow \mathrm{NO}\,(\mathrm{X}\,^2\Pi) + \mathrm{NO}\,(\mathrm{A}\,^2\Sigma^+), \tag{3.4}$$

while the photoionization steps for each of the three processes are described by the following equations:

$$\mathrm{NO}\,(\mathrm{X}\,^2\Pi) + 2\mathrm{h}\nu_{266nm} \rightarrow \mathrm{NO}^+\,(\mathrm{X}\,^1\Sigma^+) + e^-_{0.05eV}, \tag{3.5}$$

$$\mathrm{NO}\,(\mathrm{A}\,^2\Sigma^+) + \mathrm{h}\nu_{266nm} \rightarrow \mathrm{NO}^+\,(\mathrm{X}\,^1\Sigma^+) + e^-_{0.88eV}, \tag{3.6}$$

$$\mathrm{NO}\,(\mathrm{A}\,^2\Sigma^+) + \mathrm{h}\nu_{213nm} \rightarrow \mathrm{NO}^+\,(\mathrm{X}\,^1\Sigma^+) + e^-_{2.05eV}. \tag{3.7}$$

All three of these processes occur in the presence of both laser pulses (213 nm and 266 nm). However, by adjusting their relative power, we can tailor the experiment to favour one process over the others. In particular, the 0.05 eV process (equation 3.5) only requires the 266 nm beam since NO(X) monomers exist in the gas sample even without dissociating the $(\mathrm{NO})_2$ dimer, and the 2.05 eV process (equation 3.7) only requires the 213 nm beam. For all NO measurements discussed here, the gas mixture was 20% NO in helium at a stagnation pressure of 5 bar.

3.6.3.1 First NO Experiment – 3D Imaging Demonstration

The first NO experiment used only the 0.88 eV process (equation 3.6) and was intended as a powerful qualitative demonstration of our implementation of 3D photoelectron VMI. The NO(A) state is long-lived (∼100 ns) and has almost perfect Rydberg 3s character [65–67]. Single-photon ionization of the NO(A) product with a 266 nm picosecond pulse therefore produces a nearly-pure p-wave photoelectron distribution with a lab-frame orientation that is entirely determined by the probe laser polarization, since the initially-formed photoproduct Rydberg state is spherically symmetric [53].

To demonstrate that 3D VMI measurements do not constrain the probe laser polarization to lie in the detector plane, as required for inversion-based 2D VMI, we collected data sets at three different probe polarization geometries. The polarization angles we used were $\theta = 0°$, 45° and 90°, where θ (defined in Figure 3.10) is measured from the z-axis. The 0° and 45° measurements represent a type of electron imaging experiment that would be impossible using the conventional inversion-based 2D VMI method. The average hit counts for these data sets were 1.9 ($\theta = 0°$), 2.1 ($\theta = 45°$), and 3.1 ($\theta = 90°$) hits per laser shot. These experiments were performed with an intermediate level of time-stretching, giving a TTS of 6.2 ns. Note that the same 3D photoelectron distribution was measured in each of these experiments, with each directly yielding a full 3D data set for that distribution. This should be contrasted with tomographic 3D photoelectron distribution reconstruction approaches, which extract similar information via a series of experiments performed with multiple laser polarization geometries [68, 69], thus unavoidably and significantly increasing

the duration of the measurement.

In Figure 3.10, I present the analyzed and time-stamped NO data as three rows of time slices through the 3D photoelectron distribution, with each row corresponding to the result of a different experiment performed with a different polarization geometry. The $\theta = 90^\circ$ configuration (top row) corresponds to the same ionizing polarization geometry used in the Xe and CH_3I photoionization experiments. This is also one of the polarization geometries required in conventional inversion-based 2D VMI experiments. The lobes of the photoelectron angular distribution are directed along the x-axis, and for this case we expect to see an electron yield which peaks at the centre (in time) of the distribution and is less intense at either end. We also expect the electron yield to be highly localized along the $\pm x$-axis where the lobes are located. These properties are clearly demonstrated in Figure 3.10. The "outer ring" visible in some of the slices (for example, in panel (f)) is actually data from the central (maximum radius) region of the same ring being incorrectly time stamped and appearing in the wrong slice. I discuss this misassignment background further in Section 4.3.3.

The $\theta = 0^\circ$ geometry (bottom row), with the laser polarization along the z (TOF) axis, is one that is not analyzable using conventional inversion-based 2D VMI. With this polarization geometry, we expect the photoelectron yield to be peaked on either end of the distribution and very weak near the centre where the node of the nearly-pure p-wave distribution is located. We also expect the distribution to be cylindrically symmetric about the z-axis with no anisotropy visible for the presented XY slices. These expectations are satisfied. The presence of some photoelectron density in Figure 3.10 (p) is an expected

result of the imperfect nature of the p-wave distribution [53] and the width of the slices (160 ps).

With the $\theta = 45°$ data set (middle row), we demonstrate our ability to image 3D photoelectron distributions produced using arbitrary polarization geometries in a single experiment. We set the polarization to be at 45° between the x and z axes, so that the photoelectron distribution is concentrated in the negative x-axis for the front of the distribution (the earliest arriving electrons) while shifting gradually toward the positive x-axis at the back end (late electron arrival times). As shown in Figure 3.10, we observe the expected result: at the front end of the distribution (panels (h) and (i)) the electrons are concentrated in the lower part of the images (the negative x-axis), and at the back end (panels (k) and (l)) the electrons are concentrated in the top of the image (positive x-axis). At the centre of the distribution (panel (j)), the electron yield is evenly split between the top and bottom. The total electron yield over time (panel (g)) also agrees with expectations, as its qualitative structure lies roughly between the $\theta = 0°$ and $\theta = 90°$ cases.

In Figure 3.10, a lower-energy photoelectron signal is also present, most visible in the middle slices (panels (d), (j), and (p)). This is the 0.05 eV ring described in equation 3.5, and is not used in the analysis of the data sets shown in Figure 3.10.

The polarization geometry of the $\theta = 90°$ data set allows for both inversion- and slice-based analyses. Since this set of NO data was collected and analyzed before some of the data processing steps were finalized, and was intended primarily as a qualitative demon-

stration, the following treatment does not make use of the momentum vector reconstruction and fitting techniques described in Chapter 4. Nevertheless, it did provide an intermediate evaluation of quantitative results from our 3D VMI spectrometer and helped build toward the analysis of the second set of NO data, which does use the techniques developed in Chapter 4.

From the 3D data, I first extracted centre slices in the XY, XZ, and YZ planes. The XY slices had a time width of 160 ps and the XZ and YZ slices had widths of 16 pixels in the axis normal to the slice plane. Analysis and fitting of the data requires that the PAD be spherical, so as to reflect the symmetry possessed by the initial momentum distribution for photoelectrons with a single, sharp kinetic energy. The detector-plane XY slices are already circular, as expected for a well-designed VMI instrument. The XZ and YZ slices, which include the time axis, are distorted and non-spherical (see Figure 3.11). Preparation of these slices required a transformation from t to z, such that the transformed coordinates form a sphere. This was accomplished by fitting each half of the data (forward- and backward-emitted electrons) to an ellipse and transforming each using the fit. In Figure 3.13 (a) I compare the photoelectron spectra derived from the Abel-inverted data [54] and the XY slice. Significantly, we achieve a better energy resolution ($\Delta E/E$) using the XY slice (2.2%) than with the standard 2D VMI inversion (3.8%). In Figure 3.13 (b) I show the PADs of the inverted data and all three slices (XY, XZ, YZ). Each PAD has been fitted to obtain β_2 asymmetry parameters, as described in Section 4.8. The results for the inverted, XY slice, and XZ slice are consistent with previous work [52, 53], which reported values between 1.2 and 1.7. The YZ slice for the $\theta = 90^\circ$ case being considered

here is expected to be angle-invariant as a result of the cylindrical symmetry of the distribution. Thus, we expect a β_2 of zero, which nearly agrees with our experimental result of 0.13 ± 0.06. This deviation is primarily caused by the dip in intensity near 90° in the PAD, which is an expected result of gain degradation at the centre region of our MCP detector. I describe a method to characterize and correct for this detector wear in Section 4.7.

3.6.3.2 Second NO Experiment – Momentum Vector Reconstruction

The second set of NO data was collected with the intention of developing a transformation between the raw (x, y, t)-data and the initial photoelectron momentum vectors. For this reason, we used all three photoionization events in equations 3.5–3.7, which provide a range of kinetic energies from 0.05 eV to 2.05 eV with which to characterize and test the transformation. Processes originating from the well-understood electronically excited A-state of NO (equations 3.6 and 3.7) also provide an opportunity to check the fitted PADs against the published values. We collected data with only 266 nm (allowing only the 0.05 eV process, equation 3.5), only 213 nm (allowing only the 2.05 eV process, equation 3.7) and with both wavelengths (allowing all three processes). When these data were collected, the improvements to the acquisition efficiency had been implemented and the acquisition was nearly real-time (very little data transfer overhead). The photodiode timing reference had also been added (as described in Section 3.2.1) and the time resolution for these measurements ($\sim$ 70 ps) is better than for the first set of NO experiments ($\sim$ 210 ps). Since the ability to

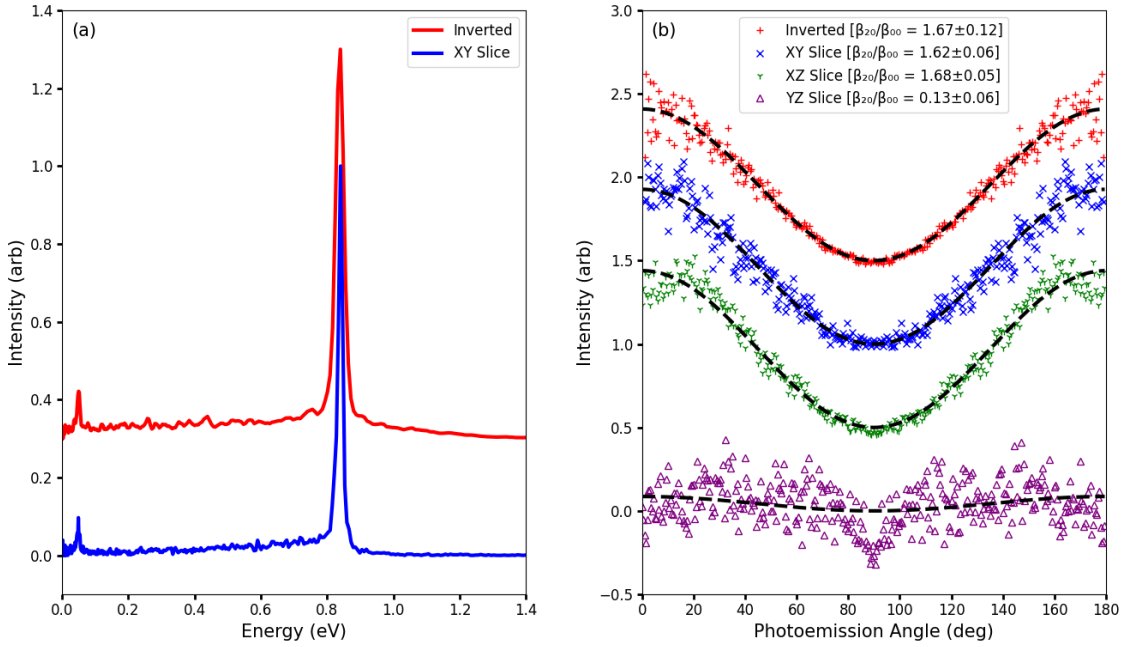

Figure 3.13: Photoelectron spectra (a) and angular distributions (b) for the single-photon photodissociation of $(NO)_2$ at 213 nm and subsequent single-photon photoionization of the NO(A) and two-photon ionization of the NO(X) photoproducts at 266 nm (high-intensity and low-intensity peaks, respectively). The data shown here correspond to the inversion-compatible 90° (detector plane) probe polarization case (as defined in Figure 3.10). In (a), I present the spectra of both the Abel-inverted image (red) and the central XY slice (blue), which yield energy resolutions ($\Delta E/E$) of 3.7% and 2.1%, respectively. In (b), I present the angular distributions and fitted asymmetry parameters (β_{20}/β_{00}) for the Abel-inverted image (1.67±0.12), XY slice (1.62±0.06), XZ slice (1.68±0.05), and YZ slice (0.13±0.06). In these plots, the emission angle is measured from the positive x-axis for the inverted, XY, and XZ slices, and from the positive y-axis for the YZ slice. The stated errors represent 95% confidence intervals. The dip in intensity near 90° in the YZ slice is a result of gain degradation at the centre of the detector. In both (a) and (b), the plots are displaced vertically for visibility proposes.

measure non-cylindrically-symmetric PADs was not the focus of these measurements, they were all performed using horizontal laser polarization. The VMI electrostatic lens voltages were kept the same throughout so that the data sets could be directly compared. These settings were the same as those used in the set of NO measurements described above.

For the 266 nm only data, we collected a total of 640,000 laser shots (10 runs of 64,000 shots each), representing 10.7 minutes of real-time data. The average hit count was 0.5 hits per laser shot and the TTS was $\sim$ 1.4 ns. For the 213 nm only data, we collected a total of 960,000 laser shots (15 runs of 64,000 shots each), representing 16 minutes of real-time data. The average hit count was 5 hits per laser shot and the TTS was $\sim$ 9.5 ns. For the data using both wavelengths, we collected a total of 2,240,000 laser shots (35 runs of 64,000 shots each), representing 37.3 minutes of real-time data. The average hit count was 4.2 hits per laser shot and the TTS of the 0.88 eV feature was $\sim$ 5.9 ns.

I have represented all three of these data sets as stacks of XY-plane slices in Figure 3.14. The data are all on the same x, y, and t scale and the PADs are oriented so that the TOF direction is upward. The echo signal described in Section 3.5.1 is clearly visible in the 266 nm (a) and two-colour (b) data; this echo was removed by subtraction before any fitting was done with these data. The 2.05 eV signal is not visible in panel (b) even though it is allowed by the presence of the 213 nm laser. This is because the power of the 213 nm beam in this experiment was set to be very low: only high enough to produce a population of NO(A) which would then be ionized with the 266 nm beam (equation 3.6). There is also a significant amount of noise interfering with the back end (bottom) of the 213 nm only data (c). This is a result of the high power of 213 nm necessary to see the 2.05 eV

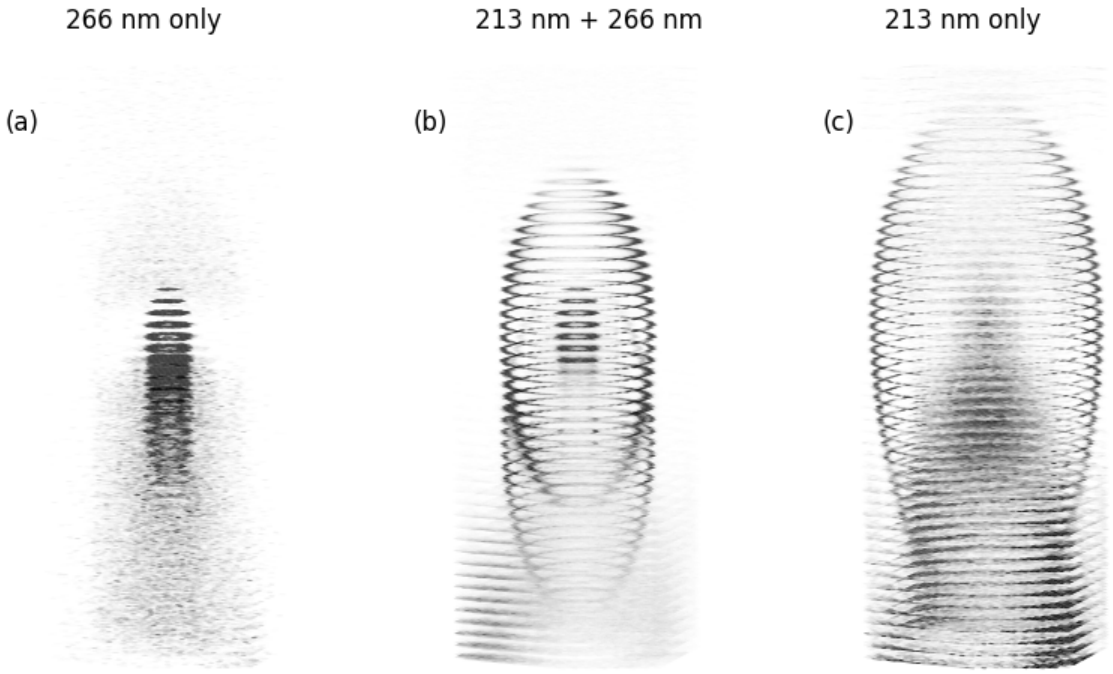

Figure 3.14: Visualizations of three NO data sets as stacked XY-plane slices (0.2 ns width). The data are all on the same x, y, and t scale and the PADs are oriented so that the flight direction is upward. The 0.05, 0.88, and 2.05 eV processes (equations 3.5–3.7) are visible in (a) and (b), (b), and (c), respectively. The echo clearly visible in (a) and (b) is described in Section 3.5.1.

signal, which produces UV scatter and photoemission from the inner metal surfaces of the spectrometer. As described in Section 4.4.2, the back half of the distribution is removed during analysis because of the presence of this noise.

3.7 Conclusion

In this chapter, I described the design of the time-stretching 13-electrode thick-lens NRC 3D VMI spectrometer. I presented the experimental details of the data collected to test and evaluate the performance of the instrument, including the optical configurations of both laser systems that were used. The data processing and analysis of these data sets are presented in the next chapter.

4

3D VMI Data Processing and Momentum Vector Reconstruction

4.1 Introduction

This chapter describes the steps required to fully process the raw 3D VMI output data from a photoionization experiment and reconstruct the photoelectron energy and angular distributions in the form of initial 3D recoil momentum vectors (p_x, p_y, p_z). The work described in this chapter has been published in two peer-reviewed journal articles [2, 3].

In brief, the 3D impact coordinates consisting of position (x, y) and arrival time (t) at the MCP/phosphor detector assembly are first extracted from the raw data. Since the

position and time coordinates are obtained from different sources, in the general case of multiple events at once (from the same laser shot) it is necessary to correctly associate each set of position coordinates with its corresponding time coordinate. The full set of impact coordinates (x, y, t) for each event are then converted into initial recoil momentum vectors using a nonlinear transformation function derived from a known, sharp calibration source. In this chapter, I also describe how the full 3D electron coordinates can be used to correct for errors in the experimentally prepared polarization state of the ionization laser and any misalignment of the laser propagation axis ($\vec{k}$). In addition, I show how to correct for a dip in gain near the centre of the MCP detector often caused by usage wear. Finally, I evaluate and quantify the ability of our 3D VMI spectrometer and data processing algorithms to correctly assign the spatial and temporal coordinates in multi-hit data sets.

The real-time data acquisition, instrument interfacing, and spatial image centroiding code was written by Douglas Moffatt. The peak-fitting treatment for the waveform data, the momentum vector reconstruction method, and the angular and MCP gain corrections were co-developed by myself and Dr. Aurelien Sanchez.

4.2 Extracting 3D Event Coordinates

A raw 3D VMI data set consists of a set of 2D camera images and a correlated set of digitized GHz waveforms, wherein each image-waveform pair comprises the electron detection events for a single laser shot. For a given laser shot, each electron event at the MCP detector is recorded as both a phosphorescence spot in the 2D image and a short

electrical pulse in the corresponding waveform. The image data contain the spatial (x, y) information for each event, whereas the digitized waveform data contain the arrival time (t) information. In the following, I will refer to an event detected in the image data as a position-hit (p-hit) and an event detected in the waveform data as a time-hit (t-hit). A p-hit and t-hit corresponding to the same event comprise the full 3D (x, y, t) coordinates for that event. The images and waveforms are first analyzed independently, determining the p-hit (x, y) and t-hit (t) coordinates and their associated amplitudes. These p-hit and t-hit amplitudes, both representing the same total charge sampling the MCP detector's broad PHD, are essential for correctly matching the spatial and temporal coordinates for each event in a laser shot containing multiple events.

4.2.1 Image Data Processing

The camera images must be processed to determine the x and y position of each p-hit as well as its amplitude. As described in Chapter 3, the full 512×512 pixel images are not saved to disk. Rather, features above a set threshold are identified during data collection and only 31×31 pixel ROIs centred on each feature are saved. Accurate fitting of the position and amplitude of these events requires post-processing since the necessary computation is not currently fast enough to be done in real time. We do, however, employ a faster but less accurate algorithm to estimate the positions and amplitudes in real time for use as a diagnostic for live monitoring during data acquisition.

4.2.1.1 Real-Time Processing of Image Data

The real-time processing of image data is performed in batches of 100 images and consists of two steps: smoothing and feature detection. We employ a simple and fast 2D smoothing routine consisting of convolution with a 2×2 array of ones. The smoothing step is required since the phosphor emissions being imaged only register at the lower end of the camera's dynamic range (*i.e.*, below 10 on the 0-255 scale of the 8-bit ADC), even when using the camera's high conversion gain setting. Smoothing allows us to lower the event detection threshold during the next step. This is particularly important for the camera-and-digitizer method of 3D VMI which relies on correlating the amplitudes of signals over the entire range (including the lower end) of the MCP detector's PHD. In addition, smoothing using a 2×2 convolution kernel allows us to then undersample the image by a factor of two, resulting in a smaller 256×256 image for the event detection step.

In the second and final real-time processing step, we identify features above a set threshold and save the surrounding regions for post-processing. This is done by raster scanning a small window (typically 10×10) through the data set. For each location along the scan, if the maximum value inside the window exceeds the threshold, the position of that value is flagged. At the end of the scan, the 31×31 region surrounding each flagged position is saved to disk. The real-time coordinates (x, y) are each calculated by taking the amplitude-weighted average position from all pixels above a set threshold. For the real-time amplitude (z), we use the average amplitude of all pixels above the same threshold.

The computer used for the data collection and real-time processing has an Intel Core

i5-4570 CPU with a clock speed of 3.2 GHz and four single-thread cores. Using this system (with all four threads), we are able to process 4000 images (4000 ms of data) in 907 ms (2960 ms if single-threaded). The code is written in C++ and compiled using Visual Studio Professional 2015 with level two (O2) optimization. We also tested the performance using a more powerful AMD Ryzen 7 2700X system with a clock speed of 3 GHz and eight dual cores. This system was able to process 4000 images in 360 ms with all 16 threads (403 ms with eight threads, and 2200 ms with one).

4.2.1.2 Post-Processing of Image Data

After data collection is finished, the saved ROIs are processed to determine more accurate values for the coordinates (x, y) and amplitude (z) of each p-hit. It is these values which are later used in the analysis of the data. From each ROI, we first subtract the average image background level and the average x and y projections, determined earlier from a large number of sample images. We then smooth the ROI using a 2D Gaussian filter, typically 5 pixels FWHM.

A typical 31×31 pixel ROI, with these initial steps already applied, is shown in Figure 4.1 (a). A larger 200×200 pixel region from the same 512×512 image is shown in Figure 4.1 (b), which contains two p-hits together. Note that this type of larger image is not stored during normal data acquisition, and is only shown here so that two p-hits can be directly compared and to demonstrate the size and structure of typical p-hits. The x and y coordinates are extracted by interpolating lineouts through the maximum along each axis

and precisely fitting the peak positions. For the two p-hits in Figure 4.1 (b), the x and y lineouts are shown in panels (c) and (d) respectively. It can be seen from the lineouts and the Gaussian fits (panels (e) and (f)) that the structure of both the p-hits and the noise is the same along both axes. The shape of a typical p-hit is, to good approximation, Gaussian along both axes, and the widths (FWHM) of both fits are consistent with each other. We are able to achieve a sub-pixel resolution of 0.2 pixels in the x and y p-hit coordinates using this data processing routine. For the associated p-hit amplitude, we simply take the maximum value of the smoothed ROI.

4.2.2 Waveform Data Processing

The signature of a single t-hit as seen in the waveform data has a complex but highly stable pulse shape caused by the effects of capacitive decoupling and imperfect high frequency impedance matching. In the following, I characterize this pulse shape and justify the use of an average of many of these pulses sampled from all regions of the detector as an instrument response function with which to fit the experimental waveform data. The data set used here is from the methyl iodide experiment described in Section 3.6.2

4.2.2.1 Pulse Shape Characterization

In Figure 4.2 (a) I show an average single-hit instrument response function, constructed using $\sim$ 20,000 single-hit waveforms. Single-hit waveforms were identified by applying a set of three thresholds of increasing height. To be accepted, a waveform must have exactly

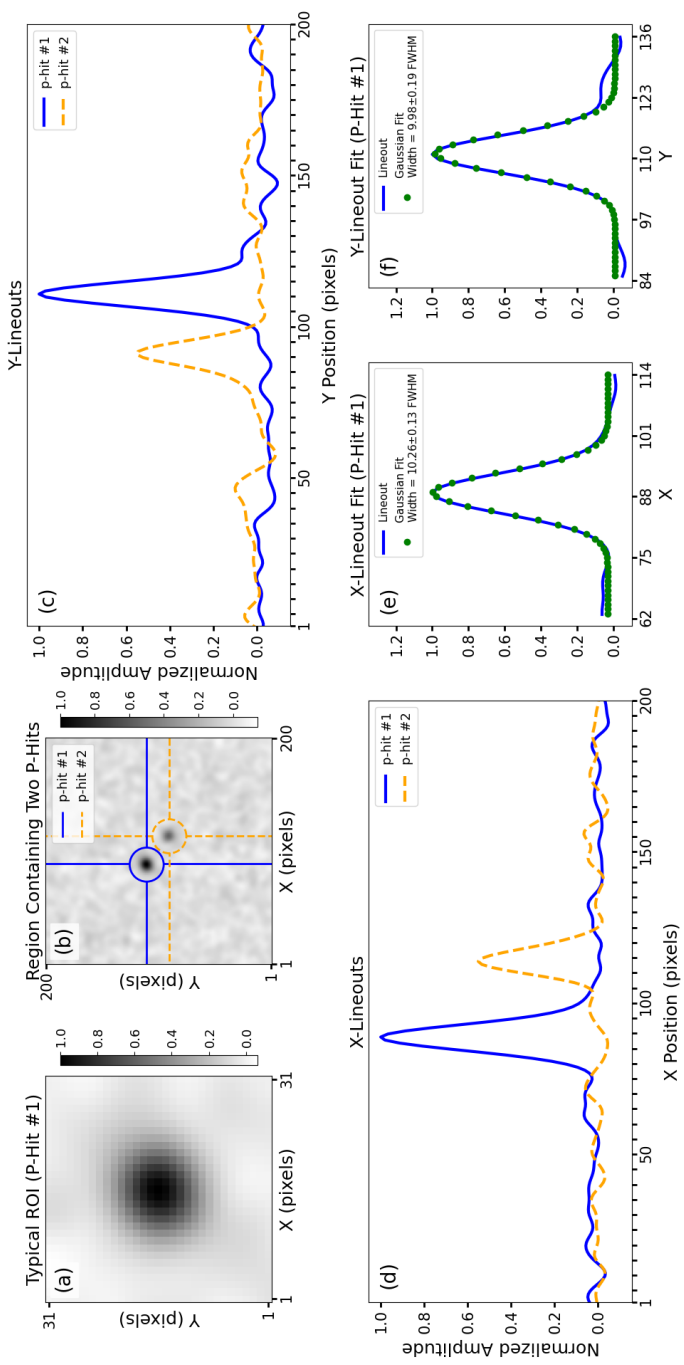

Figure 4.1: Post-processing of a typical experimental camera image containing two p-hits. (a) A 31×31 pixel region containing a single p-hit, taken from the full 512×512 image, representing the type of raw image data saved during a 3D VMI experiment. (b) A 200×200 pixel region of the same image, showing both p-hits for direct comparison. Both (a) and (b) have been processed as described in the text. Two p-hits of different heights are visible, marked by the vertical and horizontal lines. A typical p-hit has a FWHM of about 10 pixels along each axis. Lineouts through both p-hits (blue, p-hit #1, and yellow, p-hit #2) along the y-axis (c) and x-axis (d) show their relative height compared to the experimental noise. The pulse shape and noise characteristics are the same along each axis. Gaussian fits (green) to the x and y lineouts (blue) of p-hit #1 show good agreement with the expected Gaussian shape and their fitted widths are in agreement.

two local maxima above the first threshold (the main and secondary maxima in the pulse shape), a single local maximum above the second threshold (to check that the peak meets a minimum amplitude) and zero maxima above the last threshold (to avoid accepting saturated or overlapping peaks). Each of these pulses was fit to a Gaussian function to find the time coordinate of the largest peak in the pulse structure and shifted to the same time before averaging.

I will refer to this average as the Global Single-hit Response (GSR) since it represents a global average of single-hit waveforms including events sampling all spatial regions of the detector. A closer analysis of the data shows that there is a small variation in the single-hit pulse shape which depends on the spatial location of hits on the detector. This could be caused by differences in how the pulse propagates across the detector before reaching the pickoff line. An important question therefore arises: with what degree of accuracy can the GSR be used as a representative instrument response function for the global analysis of multi-hit waveforms?

To facilitate comparison of these local variations with the GSR, I also prepared a set of Local Single-hit Response (LSR) functions representing an average of 100 single-hit waveforms with the constituent events for each being constrained to a different spatial region. Specifically, the LSRs were obtained by binning the spatial region of the detector into a regular 20×20 grid, as shown in Figure 4.2 (c). Within each grid square, an LSR analogous to the GSR in Figure 4.2 (a) was generated using only single-hit waveforms where the corresponding p-hit falls within that square. Each LSR was deconvolved using the GSR as a kernel and the resulting peak was fitted to a Gaussian function to determine the peak

position (time) and amplitude. An example LSR corresponding to a region near the centre of the detector is shown in Figure 4.2 (b) along with its deconvolution. All deconvolutions were performed using a Wiener Deconvolution algorithm followed by smoothing with a $\sigma = 0.425$ ns (1.0 ns FWHM) Gaussian filter.

Deconvolving the GSR with itself and fitting gives a peak position and amplitude that the LSR fits can be compared to. In Figure 4.2 (d) and (e) I plot the time difference $t_{LSR} - t_{GSR}$ and amplitude ratio A_{LSR}/A_{GSR} between the GSR and LSRs, respectively. These 2D spatial maps use the same grid as was used to construct the LSRs and represents the systematic errors resulting from the spatially varying pulse shape. The dashed circles represent the region containing the methyl iodide photoionization feature of interest. The signal outside the circle arises from noise and is unimportant.

The range of the variation in time (Figure 4.2 (d)) can be seen to be about 50 ps within the circled region, which I consider to be acceptable and negligible on the basis that it is less than our experimental time resolution (see Section 4.4.3). The amplitude variation is about $\pm 10\%$ over the circled region, which would be enough to potentially reduce the efficiency with which we use the amplitudes to assign multi-hit data. However, as discussed later in Section 4.3.1, the spatially-varying ratio of p-hit to t-hit amplitudes is corrected in a way that accounts for both the pulse shape difference and the effects of phosphor inhomogeneity. In both cases (time and amplitude), the structure of the variation is a gradient roughly following the horizontal axis of the detector (in the orientation shown here). While it is difficult to verify the exact location of the pickoff line relative to the orientation of the camera, it seems likely that this location corresponds with one end of the

gradient. This would allow the variation to be explained as resulting from the distance from the hit location to the pickoff line, which would have an effect on the complex "splashing" of the wave-like electric impulse on the output plane of the MCP.

As an additional test of the effects of the spatial variation of the pulse shape, I constructed an artificial, precisely known double-hit test trace using each LSR. In each grid square, the double-hit test trace was constructed by summing two copies of the corresponding LSR with their peak times separated by exactly 10 ns and having equal amplitudes. An example test trace is shown in Figure 4.3 (a) for a grid square near the centre of the detector. As indicated in the figure, I will refer to the earlier peak as peak 1 and the later peak as peak 2. The purpose of this test trace is to quantify the errors associated with the relative time and amplitude of two t-hits within the same spatial bin. By deconvolving each test trace with the GSR and calculating the time difference $\Delta t = t_2 - t_1$ and amplitude ratio $A2/A1$, I can evaluate the spatial variation of these intra-bin errors. Figure 4.3 (b) and (c) show the variation of the time difference and amplitude ratio of the test traces, respectively. As in Figure 4.2, the dashed circles indicate the region containing the methyl iodide photoionization feature of interest. The range of the variation in Δt can be seen to be about 10 ps within the circled region, which is negligible compared to our time resolution. The amplitude ratios $A2/A1$ vary by less than 2% over the circled region, which is small enough to neglect. It is important that the intra-bin amplitude errors are negligible since the amplitude correction performed in Section 4.3.1 does not account for this type of variation. From this analysis, I can conclude that we are justified in using the GSR in the subsequent analysis of the waveform data.

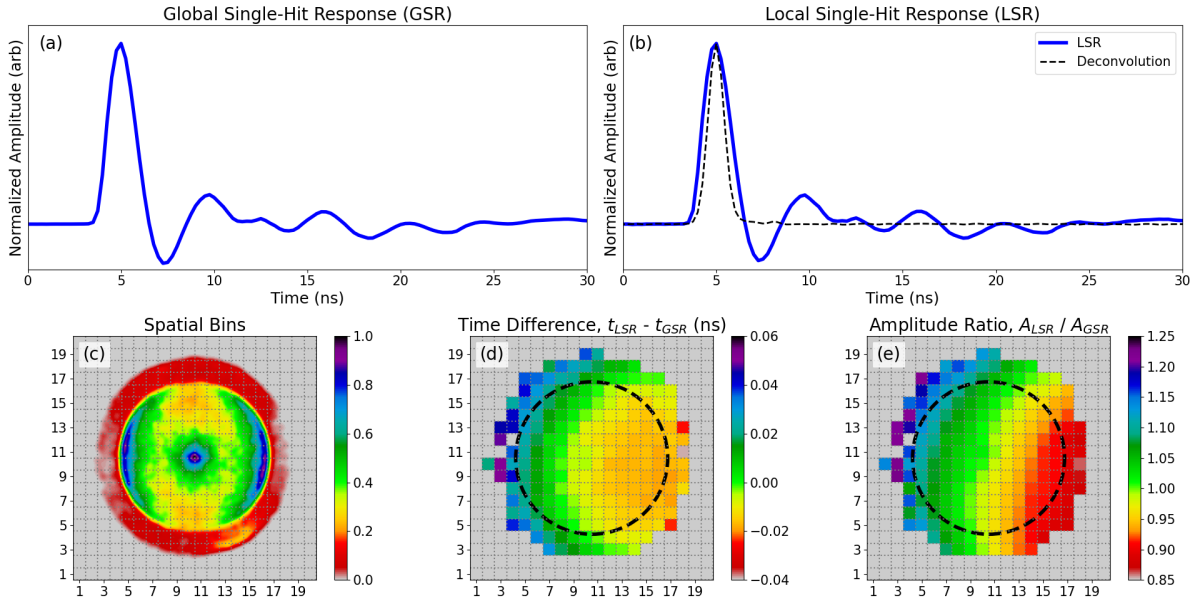

Figure 4.2: Characterization of the variation of the single-hit pulse shape with the location of the hit on the detector. (a) The GSR, which is an average of many single-hit waveforms sampled across all spatial regions of the MCP detector. The structure of the GSR is highly stable over time and specific to our VMI instrument. (b) A single LSR, which is an average of 100 single-hit waveforms sampled from a small spatial region represented by one of the grid squares near the centre of (c). The dashed line shows the deconvolution with the GSR, smoothed with a $\sigma = 1.7$ ns Gaussian filter. (c) A 20×20 grid representing the regions sampled to form the LSRs. Each of the squares which contained 100 suitable single-hit waveforms was used to construct an LSR. The colour in (c) represents the normalized yield and the outer low-yield region (red) contains hits not associated with the feature of interest (methyl iodide photoionization). The sharp cutoff indicates the edge of the MCP detector. (d) and (e) are spatial maps of the fitted peak time and amplitude of each LSR after deconvolution. Both the fitted time and amplitude are compared with the equivalent values resulting from deconvolving the GSR with itself. If the LSRs were all identical to the GSR, the time difference $t_{LSR} - t_{GSR}$ would be zero and the amplitude ratio A_{LSR}/A_{GSR} would be one.

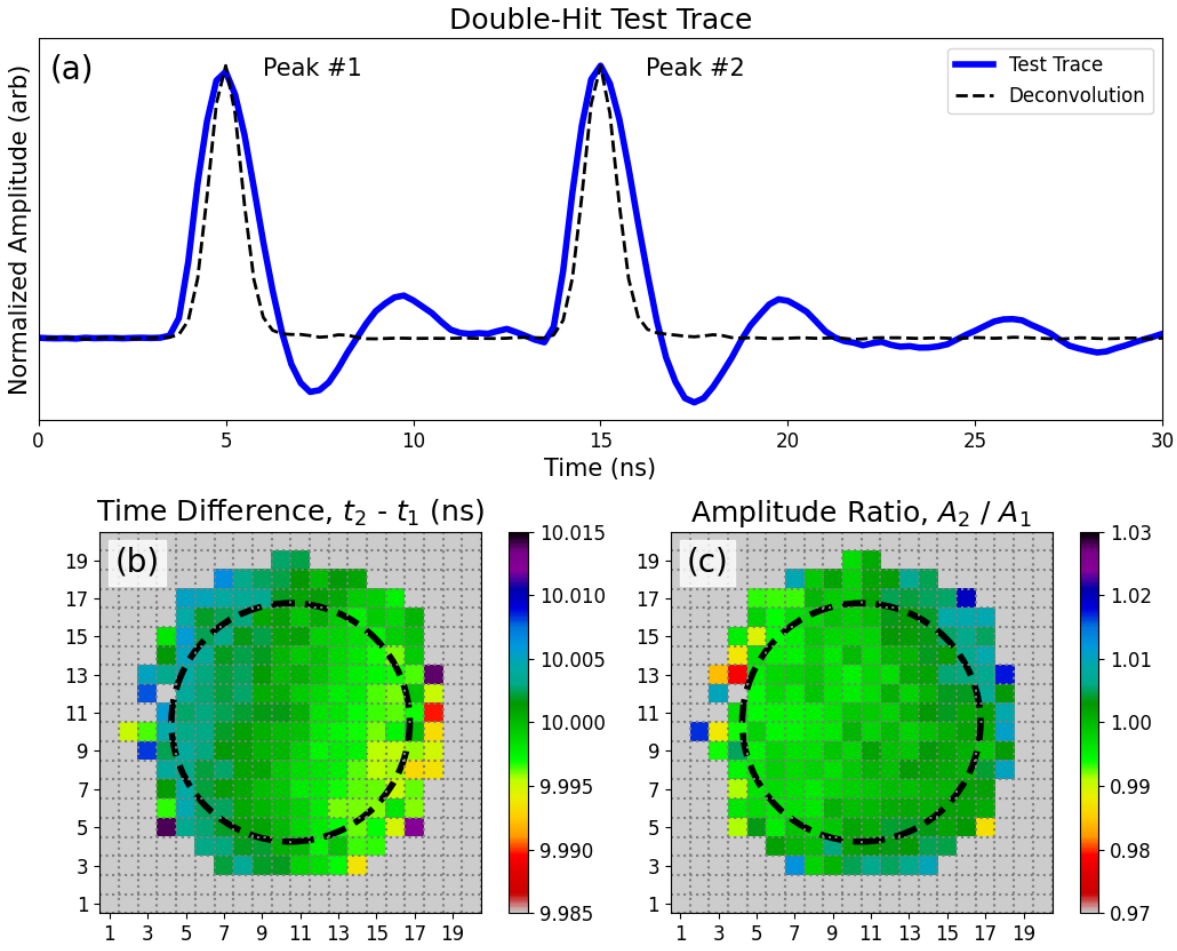

Figure 4.3: Spatial variation of the time difference and amplitude ratio of double-hit test traces. A double-hit test trace was created for each of the spatial bins in Figure 4.2 (c) by summing two equal-amplitude copies of the corresponding LSR with a time difference of exactly 10 ns. (a) The test trace, with the peaks labelled as peak #1 and peak #2. The dashed line shows the deconvolution with the GSR, smoothed with a $\sigma = 1.7$ ns Gaussian filter. (b–c) Spatial maps of the time difference Δt (b) and amplitude ratio $A2/A1$ (c) obtained by fitting the deconvolved trace in (a). The dashed circles indicate the region containing the methyl iodide photoionization feature of interest.

4.2.2.2 Peak Fitting

The GSR, derived above, represents the average structure of a single pulse. It is also necessary to characterize the average zero-hit trace, representing the background signal. The majority of the structure present in the background is Pockels cell noise produced from the laser amplifier picked up by the data collection electronics, which is highly stable and thus easy to subtract. This Pockels cell interference is substantially larger when using the picosecond laser system, which is poorly electrically shielded. For each data set, we create a background trace by averaging many zero-hit traces (determined by threshold). We have also implemented a photodiode reference pulse (see Section 3.2.1) which is accurately timed to the laser and is subjected to the same trigger jitter as the data. A reference pulse is collected along with each trace, either co-added into the data channel (one-channel mode) or on its own channel (two-channel mode), and the effect of trigger jitter can be removed by shifting the waveform data in time so that the reference pulses all align. An average of many reference pulses is shown in Figure 4.4 (a) along with an average background signal (using picosecond-laser data) and a typical single hit for comparison. I also show a zoomed in ($\times 80$) view of the average background so that the structure of the Pockels cell interference can be seen. Figure 4.4 (b) shows the detailed structure of the average photodiode reference pulse on a shorter time scale. This average is used to precisely fit the reference pulse in each experimental trace.

Two different peak-fitting algorithms were used in the analysis of the data presented in this thesis. For completeness, I will describe both here.

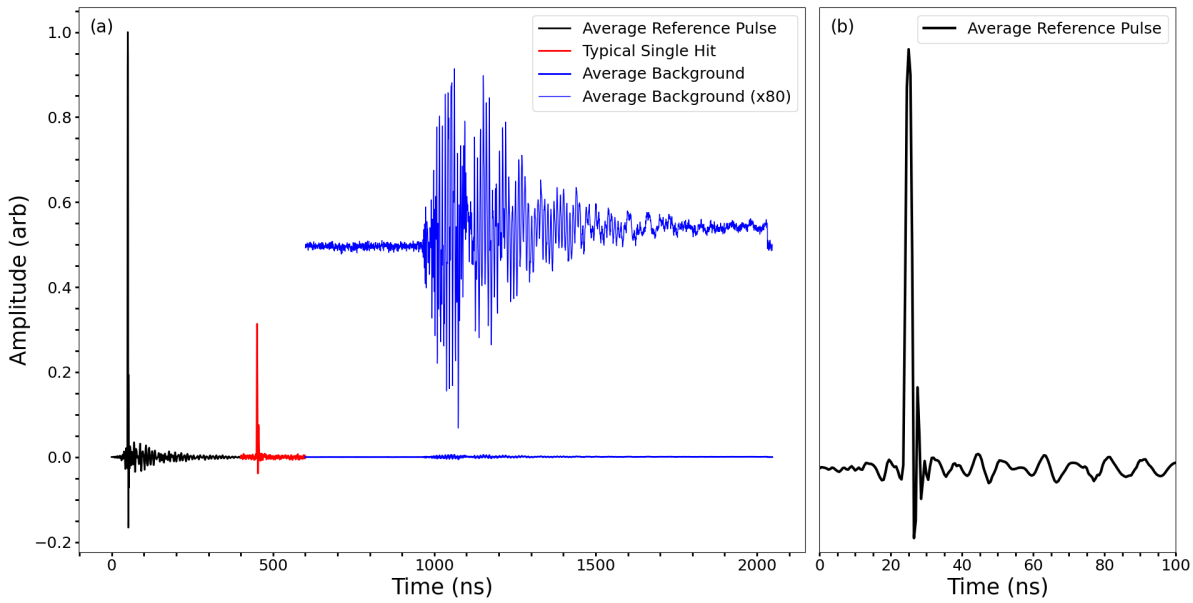

Figure 4.4: The structure of the average photodiode reference pulse and the average background signal, with a typical single hit for comparison. All three of these are shown together in (a) on the scale of the experimental traces (2048 ns). The average background is only plotted starting from 600 ns since the structure is constant for all times before the Pockels cell interference starting at 960 ns. In (b), the average reference pulse is shown on its own with a shorter time scale so that the structure of the pulse can be more easily seen.

In the first approach, the average reference pulse was used as a kernel to deconvolve the waveform data channel which contains the reference pulse. Each deconvolved trace was smoothed and fitted to a Gaussian function to determine the position (time) of the reference pulse, then (in one-channel mode only) the reference pulse was subtracted from the trace. The average background (Figure 4.4 (a)) was subtracted from the raw waveforms, which were then deconvolved using the GSR (Figure 4.2 (a)) to remove the complex instrument response function.

Typical experimental single-hit and multi-hit waveforms and their deconvolutions are shown in Figure 4.5 (a) and (b), respectively. Deconvolution enhances our resolution of individual hits in a multi-hit waveform, as demonstrated in panel (b), where it can be seen that the three hits are clearly resolved after deconvolution. Note that the amplitudes of the two highest peaks after deconvolution (the true amplitudes) are nearly identical, whereas the amplitudes appear to be different in the raw data due to the single-hit pulses being overlapped in time. Peak-picking following deconvolution was accomplished in two steps. For each trace, we first identify all local maxima above a set threshold. A region surrounding each maximum is then fit to a Gaussian function, yielding the peak position (time) and amplitude. For cases with N overlapping local maxima, the region surrounding this set of maxima is fit to a sum of N Gaussian functions, yielding a position and amplitude for each.

The second approach to peak-fitting involves first creating a filter function containing the average photodiode reference peak and the average Pockels cell background. Since the Pockels cell interference is generated by the laser, and is thus timed relative to the reference

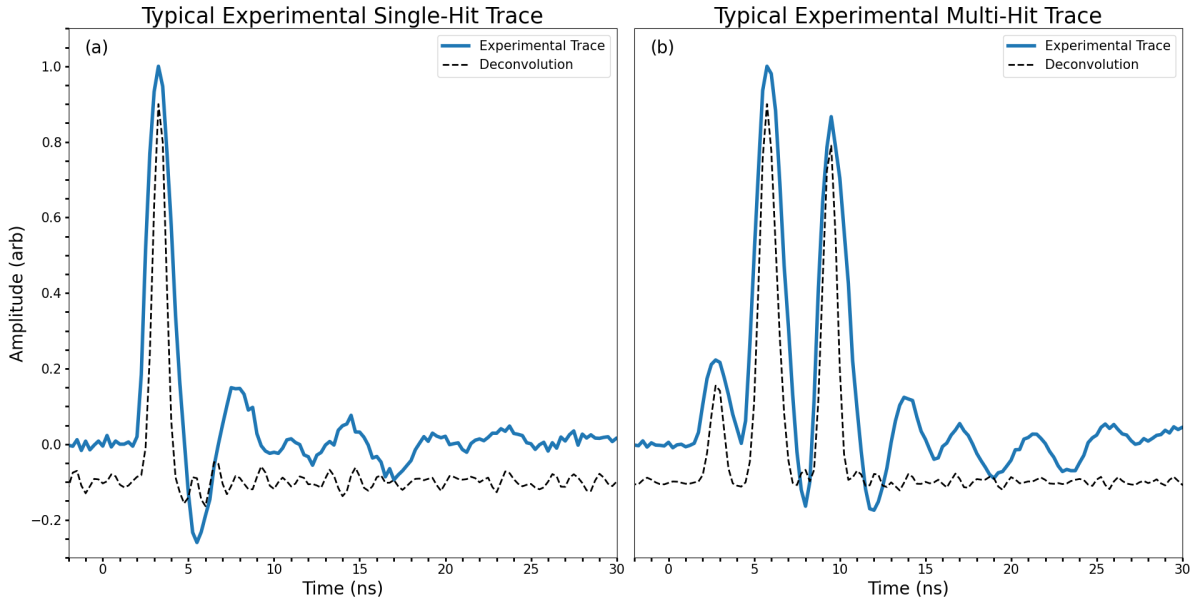

Figure 4.5: Deconvolution of typical single-hit and multi-hit traces using the average single-hit pulse shape as a kernel. (a) Deconvolution of a sample experimental single-hit trace using the kernel shown in Figure 4.2 (a). (b) Deconvolution of a sample experimental multi-hit trace using the same kernel as in (a). Both deconvolved traces have been smoothed using a $\sigma = 0.425$ ns Gaussian filter.

pulse, aligning the reference pulses for one-channel waveform data also aligns the Pockels cell background. For two-channel data, both channels were combined in sequence to make a double-length trace which was treated in the same way. This average reference-and-background trace was deconvolved using a Gaussian function ($\sigma = 0.72$ ns for one-channel data or $\sigma = 0.85$ ns for two-channel data) to create the filter function S0. Deconvolving an experimental trace with S0 has the effect of isolating the reference and background signals to a single peak at the start of the trace and applying the Gaussian smoothing. The average single-hit pulse shape was calculated similarly to how it was done in the first method; single-hit traces were identified, deconvolved with S0, and shifted to the same peak position to be averaged.

After applying the deconvolution, all local maxima above a set threshold are identified and the entire trace is fit to a sum of N functions for N maxima. These functions are comprised of one Gaussian which is fit to the peak representing the deconvolved reference pulse and background, and $N - 1$ copies of the average single-hit pulse shape for the remaining maxima. The fit parameters used are the position, amplitude, and width of the Gaussian, and the position and amplitude of the single-hit functions. In many traces there are pulses which overlap, where only one maximum was detected even though multiple hits were present. To account for this, we subtract the sum of the fit functions from the deconvolved trace and check again for local maxima above threshold. If any additional maxima are detected, an additional single-hit fit function is added and the fit is repeated. This step is repeated once more, to allow for a total of two additional fit functions beyond what was identified in the initial pass. Finally, the sum of residuals is calculated and

compared to a set minimum. If the sum is too high, indicating a poor fit, the entire trace is discarded.

The result of fitting a five-hit trace is shown in Figure 4.6. The high-amplitude peak on the left of the plot represents both the reference pulse and the Pockels cell interference. The five single-hit fit functions pictured above the full trace in the plot, along with the Gaussian fit to the reference-and-background peak, together make up the full fit (dashed black line) to the experimental trace (solid blue line). As can be seen, this results in a good fit to the trace. Since the position of the reference peak is fitted in each trace, the trigger jitter can be removed by recording the position of each hit relative to the reference.

4.3 Assigning Time-Stamps

Generally, our data sets contain multiple electron events per laser shot. In order to make use of this multi-hit data, we must determine the best way of grouping the p-hits and t-hits from each laser shot into correlated pairs (one p-hit and one t-hit) so that the members of each pair can be confidently assigned to the same event. In the following, a pair will refer to a set of one p-hit and one t-hit, whereas an assignment will refer to a set of pairs belonging to the same laser shot. Assigning the multi-hit data is accomplished by making use of the fact that each event is randomly subjected to a different gain when amplified by the MCP, as described by the MCP's PHD. Thus, the highest-amplitude hit of each type (p and t) are likely to form a correctly matched pair, as are the next highest-amplitude hits of each type and so on. However, this is not so simple in practice, since the measured

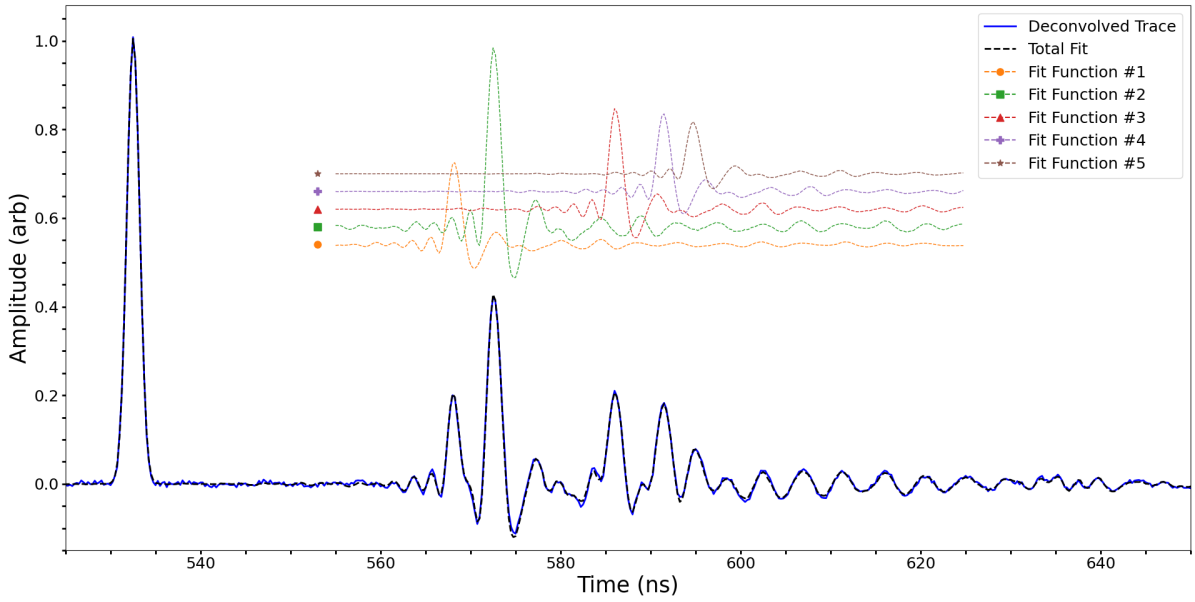

Figure 4.6: Example of peak-fitting for an experimental five-hit trace. The plotted experimental trace (solid blue line) has been deconvolved with the filter function S0 as described in the text and the time axis has been restricted to a 125 ns region of the full 1024 ns trace. This filter has the effect of smoothing the traces with a $\sigma = 0.72$ ns Gaussian function and combining the reference peak and Pockels cell background into the high-amplitude peak on the left of the plot. The full fit to the experimental trace (dashed black line) is a sum of the Gaussian fit to the reference-and-background peak (not shown) and five single-hit functions (shown above the full trace). The combination of these six functions shows good agreement with the experimental trace.

amplitude of each hit is subjected to instrument noise, which is different for the camera (p-hit) and digitizer (t-hit). This leads to some variation in the ratio of the p-hit and t-hit amplitudes associated with the same event, despite the equivalent MCP gain. As such, it is sometimes ambiguous which hits form a correctly matched pair, especially when there are multiple reasonable ways of assigning the hits from one laser shot.

There are also many shots in which the number of p-hits and the number of t-hits are not equal. This can occur if an event results in a gain on the lower end of the MCP's PHD and the signal falls below the noise level on one instrument but not the other. Since the camera images have a lower SNR than the digitizer, this is most likely to result in t-hits without a matching p-hit, rather than the converse. There are also many cases where two or more pulses on the waveform overlap and the t-hits cannot be resolved, thus appearing as a single pulse. This results in the opposite effect: one or more p-hits appear without matching t-hits. This effect will also be seen if events resulting from laser scatter occur outside the time window of the digitizer but are still captured in the camera image. The effect of pulse overlap is mitigated by the application of time-stretching, since the same pulses arrive within a wider time window and are therefore easier to resolve.

It is possible to discard any shots which have ambiguous assignments or unequal numbers of p-hits and t-hits, but this would mean discarding a large amount of potentially usable data. Even if it is not feasible to confidently assign all of the hits in a multi-hit laser shot, it is still possible to make use of a substantial fraction of them.

4.3.1 Single-Hit Amplitude Distribution

In order to quantify the level of agreement between the correlated amplitudes of a pair, and thus properly compare multiple possible assignments, it is necessary to develop a better understanding of the amplitude distribution of correctly matched data. This distribution is represented as a plot of t-hit amplitude versus p-hit amplitude for single-hit data, as shown in Figure 4.7 (a) for the two-photon ionization of atomic xenon (equation 3.2) acquired with a TTS of 8.8 ns. The phrase "single-hit data" in this case refers to data containing only shots which have exactly one p-hit and one t-hit, and "multi-hit data" refers to data containing shots which have more than one p-hit or t-hit. Since this characterization is a prerequisite of the multi-hit treatment, it is necessary to consider only single-hit data for which it can be assumed that the p-hit and t-hit belong to the same event. I will abbreviate this Single-Hit Amplitude Distribution as SHAD in the following.

For multi-hit data, the likelihood that a pair is correctly assigned (that is, that the p-hit and t-hit belong to the same event) depends on the degree to which the amplitudes of the pair agree with the SHAD. For this reason, it is important that the SHAD is as narrow as possible so that pairs can be confidently assigned. The SHAD associated with the aforementioned 8.8 ns TTS xenon data is shown in Figure 4.7 (a). This SHAD (which is representative of other SHADs collected under similar conditions) is fairly broad, which impedes its use in assigning multi-hit data. One of the reasons for this is that the amplitudes of the p-hits vary significantly depending on the spatial location of the p-hit on the detector. Since the t-hits do not display the same spatial variation, this cannot be a

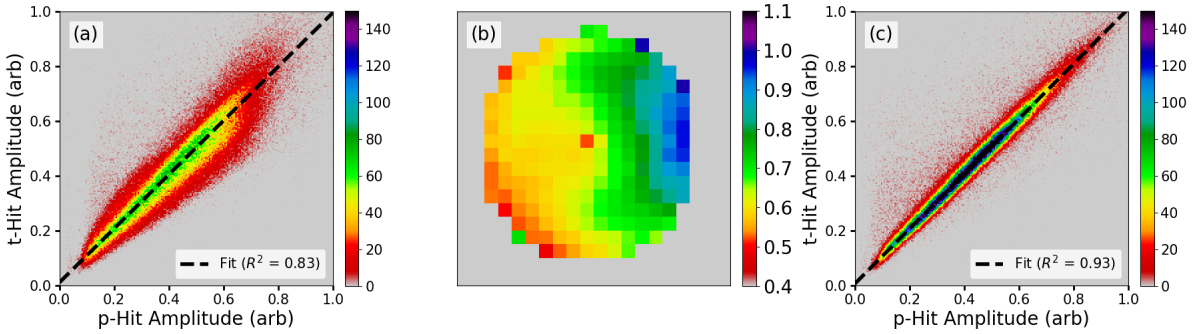

Figure 4.7: The SHAD from two-photon ionization of atomic xenon acquired with a TTS of 8.8 ns. The SHAD is shown before (a) and after (c) correction for the spatial variation (b) of the p-hit/t-hit amplitude ratio over the phosphor screen of the detector. The SHAD shows the distribution of correlated p-hit (from the camera image) and t-hit (from the digitizer waveform) amplitudes using only single-hit data. (b) The spatial region of the MCP detector imaged by the camera, divided into a regular 20×20 grid. The single-hit data are binned according to where they fall within this grid. The colour represents the normalized slope of the SHAD, using only hits from each bin, and the sharp cutoff indicates the edge of the MCP detector. We observe that the slope of the SHAD exhibits a clear gradient from the right to the left of the detector in the orientation shown, and a minimum at the centre. (c) To correct for the spatial variation of the slope, we scale the amplitude of the position data so that the average slope in each bin is one. This results in a more narrow distribution which allows us to better assign p-hits and t-hits in multi-hit data.

result of a spatially dependent MCP gain (which would affect the t-hit and p-hit amplitudes equally) but is rather caused by spatial variation in the phosphor efficiency. The variation of the single-hit pulse shape (see Section 4.2.2.1) also likely makes a contribution to the width of the SHAD. A spatially dependent phosphor efficiency could be a result of manufacturing imperfections in the phosphor screen and/or due to experimentally-induced degradation of the phosphor. This spatial variation manifests as a broadening of the SHAD as a result of its spatially dependent slope.

We correct for the spatial variation of the SHAD by first binning the single-hit data (p-hit/t-hit pairs) into an equally spaced 20×20 grid based on the location of the p-hit on the detector. This is presented in Figure 4.7 (b), where the value assigned to each square is the slope of the SHAD using only pairs where the p-hit falls within the corresponding spatial bin. I have used a 20×20 grid since it is both fine enough to show the structure of the spatial variation of the slope as well as coarse enough so that each bin has a large enough sample size to give a sufficiently accurate value for the slope. For the $\sim$200,000 single-hit shots in this data set, this grid size results in bins containing 100 to 5000 pairs. Figure 4.7 (b) highlights a clear gradient in slope from the right to the left of the detector, as the image is oriented here. It also shows a minimum in the slope located at the centre of the detector, likely representing degradation of the phosphor screen. To correct the data, we scale the p-hit amplitudes in each bin by the slope of the SHAD associated with that bin. This ensures that the average slope in each bin will be one. After correction, the SHAD is narrowed significantly, as shown in Figure 4.7 (c).

4.3.2 Assignment Algorithm

For a general laser shot with N p-hits and M t-hits, there are $N * M$ possible pairs that can be formed. To determine the best assignment (the set of pairs which are most likely to be correctly matched), we must quantify the degree to which each pair agrees with the SHAD. This value, which I will call the quality of the pair, is determined by evaluating the smoothed SHAD at the coordinates corresponding to the amplitudes of the pair. Since the SHAD is the observed density of single hits as a function of their amplitudes, its value at a certain pair of amplitudes represents the likelihood of observing that pair. This assumes that the hits in a multi-hit shot can be considered independent events, which is justified since the hits typically occur in different spatial regions of the detector, especially for hit counts below 10 hits/shot as is typical, and thus should not interact. I will represent the qualities of the $N*M$ possible pairs as an $N*M$ matrix which I will call the quality matrix $\boldsymbol{Q}$.

The rows and columns of $\boldsymbol{Q}$ represent pairs which contain the same p-hit (rows) or t-hit (columns), so accepting one pair as part of an assignment will necessarily involve rejecting all other pairs in the same row and column. Consequently, it is better to evaluate the quality of a multi-hit assignment as a whole rather than only considering the quality of the individual pairs that belong to it. I will define the quality of an assignment as the product of the qualities of the pairs it contains. Since the quality value for a pair represents the likelihood of observing that pair, the product of the qualities of pairs belonging to an assignment will represent the likelihood of observing that assignment. To make the most

probable assignment, a set of all possible n-pair assignments is constructed, where n is the minimum dimension of $\boldsymbol{Q}$. The best assignment is then the one with the greatest quality. The task of finding the best assignment from the matrix $\boldsymbol{Q}$ is a variation of a mathematical optimization problem called the assignment problem [70], for which the solution is known and has been implemented in the Python package *scipy.optimize* [71]. Currently, this assignment is accepted without further consideration, although improvements could be made using a more advanced statistical treatment.

4.3.2.1 Assignment Thresholds

The algorithm described above simply accepts the best assignment of each multi-hit laser shot, entirely disregarding any potential ambiguity between assignments with similar quality. This approach involves the obvious risk of accepting incorrect pairs along with the correct ones in an assignment. In this section, I describe a more complex algorithm for attempting to remove ambiguity from the accepted assignments. Ultimately, this version of the assignment algorithm was not used in the data analysis due to a lack of statistical justification. However, I include it here because I believe it has value as a starting point for a more advanced treatment of the multi-hit assignment problem.

The first difference from the version of the assignment algorithm described above occurs immediately after the creation of the quality matrix $\boldsymbol{Q}$. I reasoned that there should be a minimum quality below which pairs should not be accepted even if there are no competing assignments using the same p-hit or t-hit. To account for this, I introduced a threshold

to be applied to $\boldsymbol{Q}$, which I called the absolute correlation threshold (ACT). Any element of $\boldsymbol{Q}$ below the ACT is set to zero and any hits (p or t) which only contribute to pairs with a quality of zero (corresponding to a row or column of all zeros) are removed from the analysis. That is, the rows and/or columns of these non-contributing hits are removed and and a new reduced $\boldsymbol{Q}$ matrix is constructed to be used in the following steps. This algorithm was tested using an ACT value which is expected to reject $\sim 5\%$ of the lowest quality correct pairs from the observed experimental amplitude distributions.

The second and most important change in this version of the algorithm occurs after evaluating the qualities of each possible assignment of the hits in a laser shot. If there are multiple assignments with non-zero quality then, rather than accepting the highest-quality assignment immediately, it is compared with the other assignments to check for ambiguity. This comparison is done using another threshold, which I call the relative correlation threshold (RCT). The RCT is a value between zero and one and the product of the RCT with the quality of the best assignment (q_1) is compared with the quality of the second best assignment (q_2). The best assignment is considered unambiguous if $q_1 - (q_1 * RCT) > q_2$ and any assignments with qualities greater than $q_1 - (q_1 * RCT)$ are considered to be ambiguous with the best assignment. In the event of ambiguity between two or more assignments, the pairs that are common to all of the ambiguous assignments are accepted and the remaining are discarded.

There are some cases where it may be preferable to accept a subset of the pairs belonging to an assignment. This could occur if an assignment can be made with high qualities for all pairs except for one, which has a quality below the ACT. However, any assignment

containing a pair with a quality below the ACT (set to zero due to the ACT thresholding) would have an assignment quality product of zero and would thus be discarded even if the other qualities belonging to the assignment are high. To avoid this situation, an exception is made when calculating the qualities for assignments containing a zero-quality pair. In these cases, the zero-quality pair will instead contribute a value equal to the ACT when forming the product. This is done so that such an incomplete assignment will still be considered, albeit with a low weighting as a result of the presence of a low quality (the ACT) in the product. Of course, when the best assignment is chosen, any pair with a quality of zero will still be rejected.

I tested this assignment algorithm with various values of the RCT using an artificial data set (described in detail in the next section) where the correct assignments are known with certainty. The results showed a substantial reduction in the acceptance of both correctly and incorrectly matched pairs for any significant value of the RCT. As discussed in the following section, the efficiency of the current assignment algorithm is sufficient even without the inclusion of this more advanced thresholding. Should the need for more efficient data processing become a priority, I believe that additional testing and further statistical considerations could make this approach, or a similar implementation, a viable starting point. Currently, without a compelling justification for how to properly set the ACT and RCT, both thresholds were set to zero for analysis of the data discussed in this thesis.

4.3.3 Multi-Hit Performance

One way to quantify the efficiency of the multi-hit data analysis described above is by the average number of correctly assigned pairs per laser shot. Since we are primarily concerned with characterizing our ability to handle higher hit-count data, it is instructive to consider how the number of correctly assigned pairs varies with the total number of hits in the shot. The number of hits in an experimental laser shot can be ambiguous when there are unequal numbers of p-hits and t-hits. When grouping the data by hit count to perform the following analysis, I chose to use the number of p-hits (rather than the number of t-hits) to define the number of hits in the shot. This is a more reliable value than the number of t-hits since many waveforms suffer from unresolved time overlap of peaks. In the following, I compare the multi-hit assignment efficiency of four different experimental data sets with a range of photoelectron energies (E) and turn-around time spreads: xenon ($E = 0.27$ eV, TTS = 31 ns), xenon ($E = 0.27$ eV, TTS = 8.8 ns), CH_3I ($E = 2.78$ eV, TTS = 6.9 ns), and NO ($E = 0.88$ eV, TTS = 6.2 ns). In this section, only the energy and TTS of the data sets are important; additional experimental and spectroscopic details are provided in Section 3.6. Note that the distribution of hit counts was lower for some of the experiments (8.8 ns TTS Xe and 6.2 ns TTS NO) so the higher hit counts are not well represented in these data sets.

It should be noted that not all of the pairs that are assigned will be assigned correctly. In order to quantify the number of correctly assigned pairs, it is necessary to estimate the number of pairs that were assigned incorrectly. For an experimental data set, it is

impossible to know with certainty which assignment is the correct one. For this reason, I constructed an artificial data set (ADS), for which the correct assignments are known, corresponding to, and mimicking the statistical properties of, each of the experimental data sets listed above. Since the experimental data sets are large, I constructed the ADS by sampling the experimental values themselves. This should be equivalent to (and easier than) modelling the distribution and sampling the model. Previously processed and confidently assigned single-hit amplitude pairs from the corresponding experimental data were randomly sampled and grouped into sets of 1–20 to form shots (4000 shots for each hit count). These were analyzed in an identical way to the experimental data, as described above, and the resulting assignments were compared to the known correct assignments. The ratio of correct to incorrect assignments was then applied to the experimental results to produce an estimate of the number of correctly assigned pairs.

Incorrect assignments, where the position and time each belong to different events, will contribute a background signal distributed over the cylinder defined by the spatial and temporal coordinates sampled in the data set. It is important to note the features of interest in an experiment will typically occupy only a small volume of this cylindrical region and, as such, the presence of incorrect assignments has a minimal negative effect on the quality of the data. Furthermore, it possible to model the distribution of incorrectly assigned pairs by randomly assigning p-hits and t-hits from different shots over the entire data set. With an accurate estimate for the ratio of incorrectly to correctly assigned pairs, this misassignment background can be subtracted from the angular distribution later in the analysis.

The correct and incorrect pair assignments for each ADS are plotted in Figure 4.8 as a function of hit count (up to 20 hits per shot). Since the ADS was constructed with only shots where the number of p-hits and t-hits are equal, all hits were assigned. Therefore, the sum of the correctly (green) and incorrectly (red) assigned pairs is equal to the total number of hits in the shot (black dashed line). The intersection point where the number of correct and incorrect assignments are equal can be taken as an indication of the multi-hit performance. For the NO (d) and both xenon data sets (a–b), the intersection point increases with the TTS as would be expected since a more broad distribution of arrival times should make higher hit-count data easier to resolve. However, this trend is not followed by the CH_3I data (c), which, although the intersection point is not visible on the 0–20 hit scale, can be assumed to intersect at the highest hit count of all four data sets. This behaviour is unexpected and at the time of writing I currently have no explanation for why this might be the case.

Figure 4.9 contains similar plots for the experimental data. Since the experimental data sets contain many shots with unequal numbers of p-hits and t-hits, some of the p-hits will not be assigned. In these plots, the sum of the correctly (green) and incorrectly assigned (red) pairs is represented by an additional line (blue), which is less than the total number of hits in the shot (black dashed line). Note that the horizontal scale is different for each of the plots in Figure 4.9 since the higher hit counts are not represented in all of the experimental data. In Figure 4.10, I present a combined plot showing the correctly assigned pairs for each data set together. The lines shown in Figure 4.10 are the same as the green lines from each of the plots in Figure 4.9. It is clear from Figure 4.10 that the multi-hit assignment

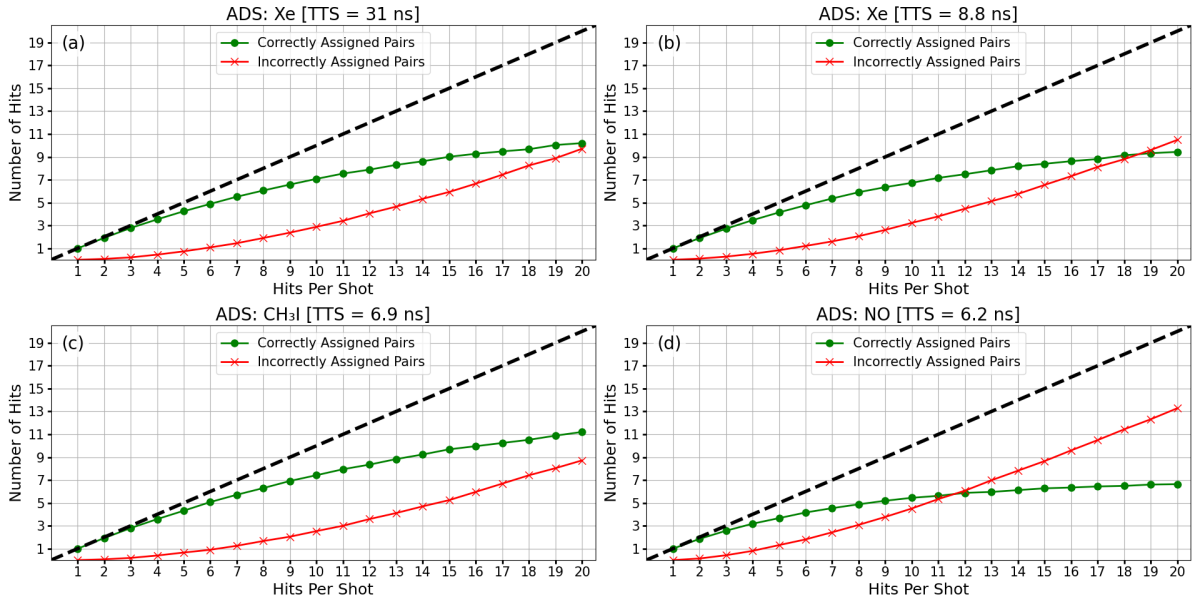

Figure 4.8: Multi-hit assignment efficiency for the ADSs, represented by the number of correctly assigned pairs for simulated laser shots with hit counts ranging from 1–20. An ADS was constructed by sampling each of the following experimental data sets with different TTS: (a) xenon (TTS = 31 ns), (b) xenon (TTS = 8.8 ns), (c) CH_3I (TTS = 6.9 ns), and (d) NO (TTS = 6.2 ns). For each ADS, the sum of the correctly (green) and incorrectly assigned (red) pairs is equal to the total number of hits in the shot (black dashed line). The ratio of incorrect to correct assignments is used to estimate the number of incorrect assignments in the experimental data corresponding to each ADS.

efficiency increases with the time spread of the photoelectron distribution. This trend is present even when considering photoelectrons of different kinetic energy (0.27–2.78 eV). We can correctly assign up to six hits per shot for the case of the 31 ns TTS xenon experiment, and up to four hits per shot for the higher-energy CH_3I case. Importantly, we have not seen a substantial drop in assignment efficiency even with higher hit count, up to 15 hits per shot (the highest hit count sampled in the experimental data). This means that having a wide range of hit counts in the data is not harmful to our data collection rate. This, in turn, demonstrates that 3D multi-hit VMI can operate at data rates over an order of magnitude greater than that of true coincidence techniques such as COLTRIMS.

There are two major limiting factors for assignment efficiency. The first is the ability to resolve individual t-hits, especially when the hit count is high. It is this factor which is addressed by the application of time-stretching, as is evident in Figure 4.10. The other factor is the inability to resolve the difference in amplitude between hits. This issue can be addressed by using a high-gain MCP detector with a broader PHD, and we are currently working with modified MCP geometries as a method for improving this aspect (see Section 3.2.3).

4.4 Momentum Vector Reconstruction

The preceding analysis steps yield a set of (x, y, t) coordinates for the observed electron events in a 3D VMI experiment. These coordinates are necessarily dependent on experimental details such as the mechanical specifications of our spectrometer and the electrode

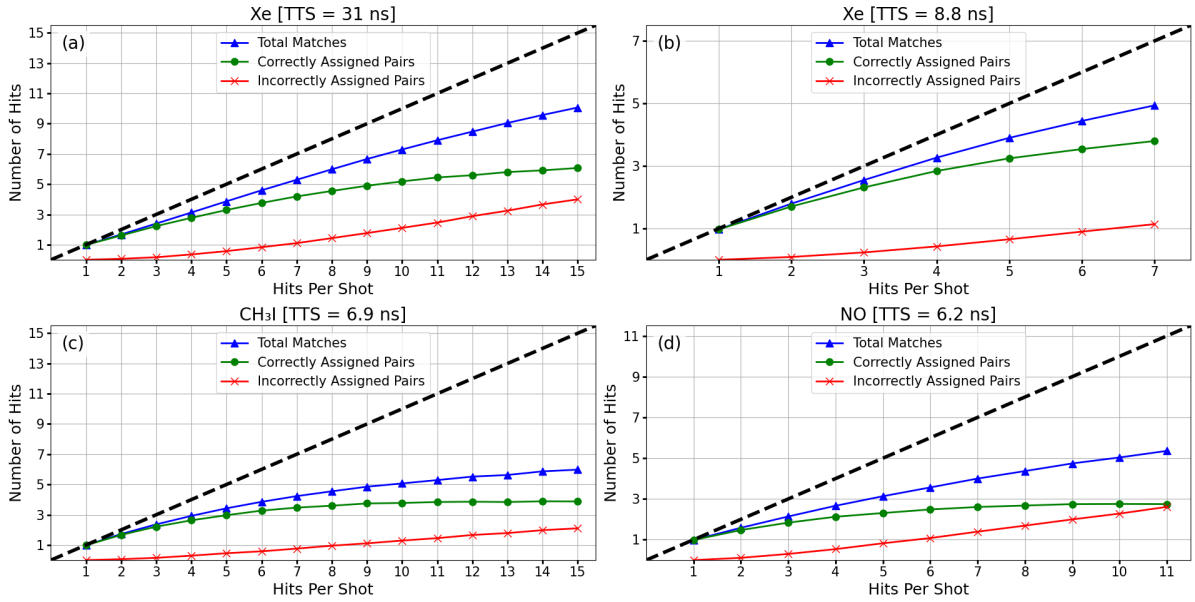

Figure 4.9: Multi-hit assignment efficiency for experimental data sets with different TTS: (a) xenon (TTS = 31 ns), (b) xenon (TTS = 8.8 ns), (c) CH_3I (TTS = 6.9 ns), and (d) NO (TTS = 6.2 ns). The total number of assigned pairs (blue) for laser shots with hit counts ranging from 1–15 (depending on the data set) are obtained directly from the processing of the experimental data. This is less than the total number of hits in the shot (black dashed line) since some p-hits were not assigned. Estimates for the number of correctly (green) and incorrectly assigned (red) pairs are derived by assuming the same ratio of incorrect assignments as in the ADSs (Figure 4.8).

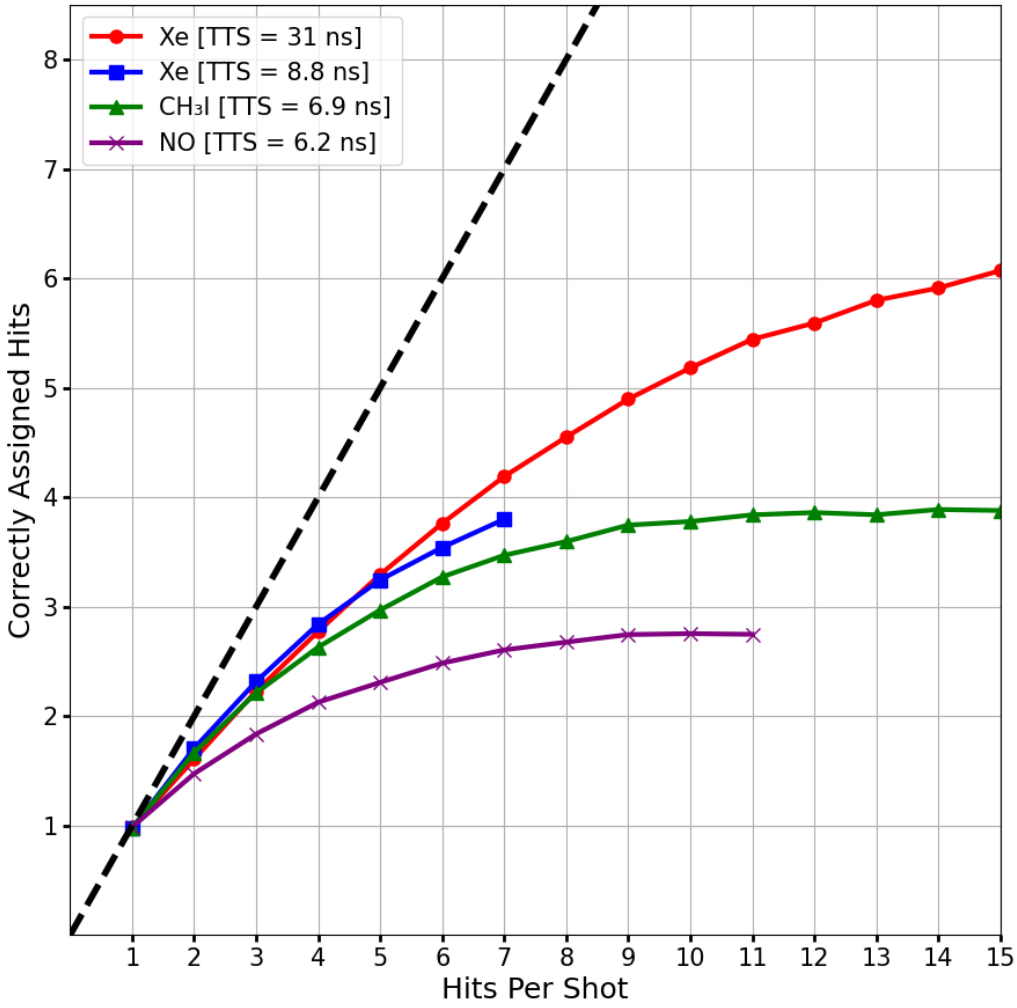

Figure 4.10: Comparison of the multi-hit assignment efficiency for data sets with different TTS. Four different experiments using photoelectron distributions with varying kinetic energies and TTS are shown, with hit counts up to 15 hits/shot (except for the 8.8 ns xenon and 6.2 ns NO experiments where the experimental hit count was lower). The dashed line represents the maximum possible efficiency (*i.e.* all hits in the shot are correctly assigned). These results show that a longer time spread results in a greater efficiency, and for the case of the 31 ns xenon experiment we can correctly assign up to six hits/shot. The 8.8 ns xenon, 6.9 ns CH_3I, and 6.2 ns NO experiments permitted correct assignments of up to four, four, and three hits/shot, respectively.

voltage settings. Therefore, these data must be transformed into initial photoelectron momentum vectors (p_x, p_y, p_z) which reflect atomic or molecular properties and are, as such, instrument independent. Below I present a method to develop such a transformation using a calibration source that produces sharp velocity-map-image features associated with well known initial recoil momentum vectors. The transformation from the raw 3D coordinates (x, y, t) to the initial momentum coordinates (p_x, p_y, p_z) is derived in two steps. The 2D detector-plane coordinates (x, y) are converted into their respective momenta first, then the time-stamp is converted into the z-axis momentum. In the following treatment, we used the three nitric oxide photoionization features described in Section 3.6.3. We know the precise kinetic energy E of the photoelectron signal associated with each ionization channel, thus we know the total momentum $p = \sqrt{2E}$ (in atomic units).

4.4.1 Transverse Coordinates

First we treat the 2D coordinates in the plane of the detector. This step is well-understood since the same transformation is performed in conventional 2D VMI. For a well-designed VMI spectrometer operating with velocity-focusing electrode voltage ratios, this is simply a linear transformation:

$$p_x = Kx;\ p_y = Ky. \tag{4.1}$$

The conversion factor K between the 2D coordinate (in pixels) and the corresponding momentum (in atomic units) must be determined from the calibration source. We obtain

K by first taking a 2D (x, y) slice through the centre of the photoelectron distribution, as shown in Figure 4.11 (a–b) for the case of the 0.88 eV photoionization signal of nitric oxide (equation 3.6). For the centre of the distribution, we use the arrival time coordinate that maximizes the observed radius in the plane of the detector. This coordinate can be accurately determined by inspection of the data. We then use a Gaussian fit to find the radial maximum r_{max} of the resulting ring in the centre slice, as shown in Figure 4.11. The conversion factor K is then given by:

$$K = p/r_{max}, \tag{4.2}$$

where p is the known photoelectron momentum of the calibration source. For the demonstration case in Figure 4.11, the value of K is 0.00226 a.u./pixel.

4.4.2 Time-to-Momentum Transformation

The transformation of the time coordinate (t) to p_z is nonlinear and requires a function $p_z = F(t)$ which we must derive from the calibration source. To accomplish this, we make use of p_x and p_y from equation 4.1 and calculate the transverse, detector-plane momentum:

$$p_\perp = \sqrt{p_x^2 + p_y^2}. \tag{4.3}$$

For each photoelectron feature in the data, plotting $p_\perp$ versus t gives a half-ring, as shown in Figure 4.12 (a) for all three NO photoionization processes. These three rings

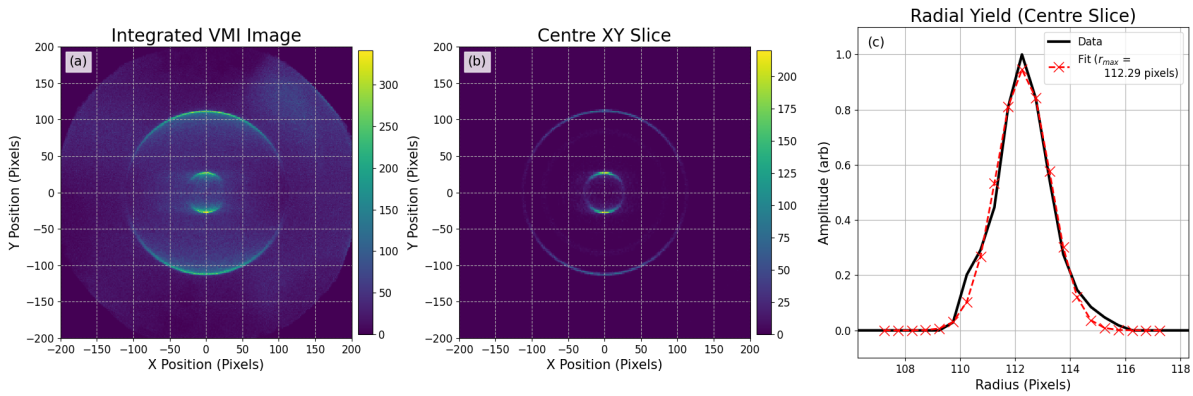

Figure 4.11: Demonstration of the process for finding the maximum 2D detector-plane radius of a 3D VMI data set. Combined with a known photoelectron momentum corresponding to the observed feature, this allows one to obtain the calibration constant K between radius (pixels) and momentum (atomic units). The pictured data are from photoionization of nitric oxide and contain signals with peak energy at 0.05 and 0.88 eV (equations 3.5 and 3.6, respectively). Only the outer 0.88 eV signal is used here. We start from a set of event coordinates (x, y, t), represented in (a) by a 2D VMI image integrated in time over the entire distribution. We then take a central time-slice (b) where the 2D ring is at its maximum radius (determined by inspection). In this case, the width of the slice was 200 ps. The radial spectrum from the central slice (c) is fit to a Gaussian to find the maximum radius r_{max} of the 0.88 eV ring. This analysis gives an r_{max} of 112.29 pixels and using equation 4.2 results in a K of 0.00226 a.u./pixel. Note that the inner ring is not pictured in (c) but would be located at lower radius.

originate from three different photoionization processes and, as such, they are displaced in time relative to each other. In the following fit, these small (< 1 ns) time differences have been subtracted from the data to facilitate proper comparison. We divide each ring into narrow vertical slices in time (using widths of 68 ps, 96 ps, and 172 ps for the 0.05 eV, 0.88 eV, and 2.05 eV rings, respectively), then perform Gaussian fits to locate the maximum values of $p_\perp$ at each discrete t-value. We further perform such fits to horizontal slices of the data shown in Figure 4.12 (a) to obtain $p_\perp$-resolved errors (σ of these Gaussian fits) in the time coordinate. This results in a set of (t, $p_\perp$) coordinates describing each photoelectron-feature ring (coloured circles in Figure 4.12 (a)). We transform the $p_\perp$ values into p_z via the following relation:

$$p_z = S\sqrt{2E - p_\perp^2}, \tag{4.4}$$

where the sign S is -1 for the negative-going half of the distribution ($t < 0$) and +1 for the positive-going half ($t \geq 0$). We then plot the (t, p_z) coordinates as shown in Figure 4.12 (b) and fit them to a second-order polynomial function F. After testing with various polynomial fits, we found that a second-order polynomial well-describes the simulated and measured dependence of p_z on t. The inclusion of the quadratic term in the fit accounts for the variation in momentum resolution along the z-axis ($\partial p_z/\partial t$), as discussed below and shown in Figure 4.12 (d). The horizontal errors bars in Figure 4.12 (b) are the widths (σ) of the Gaussian fits to the $p_\perp$ slices. The vertical p_z error bars are derived by propagating the error in $p_\perp$ (from the time slices) using the following transformation:

$$\sigma_{p_z} = \left| \frac{p_\perp}{p_z} \right| \sigma_{p_\perp}. \tag{4.5}$$

The errors near $p_z = 0$ become large as a result of the vanishing denominator. With this polynomial fit, the complete set of transformations to recover the initial momentum vector are:

$$p_x = Kx, \quad p_y = Ky, \quad p_z = F(t). \tag{4.6}$$

We observe excellent agreement between all three rings in Figure 4.12, which range in kinetic energy from 0.05 eV to 2.05 eV. This supports the applicability of this method as a robust calibration for this 3D VMI system. The fit shown in Figure 4.12 (b) uses the 0.88 eV data only and deviations of all three data sets from the fit are shown in (c). The agreement with the 2.05 eV data suggests that one could even confidently extrapolate such a fit to energies higher than that of the calibration source.

As can be seen in Figure 4.12 (d), the momentum resolution along the z-axis, $\partial p_z/\partial t$, is better for negative p_z (the electrons initially emitted away from the detector) and becomes worse as p_z increases. This is because the electrons emitted away from the detector spend more time exposed to the weak field gradient near the interaction region and therefore experience more time-stretching as a result. In practice, however, the negative-going half of the photoelectron distribution is often more difficult to analyze due to the presence of small stray fields and spurious signals, such as those associated with the scattered UV light back-

ground which are difficult to completely remove [62]. These negative-going electrons can be less well-behaved compared to the forward end of the photoelectron distributions and, for these reasons, are often discarded [72]. The difficulty in analyzing the negative-going electrons could be a result of them spending more time in the spectrometer, expanding closer to the electrodes, and passing through $p_z = 0$ as they turn around and therefore being more sensitive to weak stray fields.

4.4.3 Momentum Errors and Resolution

In the previous section, I performed a series of Gaussian fits to obtain the transverse momentum uncertainty ($\sigma_{p_\perp}$) along the z (t) axis. Here I use these measured errors to estimate the spatial and temporal resolution of this experiment, as shown in Figure 4.13. In the following I use only the nitric oxide photoionization feature at 2.05 eV and exclude the high-noise region below ~ -0.25 a.u. The momentum distribution for this data set, transformed from the raw (x, y, t)-data as described in the previous section, is shown in Figure 4.13 (a) along with the previously derived 1σ transverse momentum errors ($\sigma_{p_\perp}$).

The 1σ confidence region between the error bars is used to produce a binary mask (Figure 4.13 (b)), which is assigned a value of one inside the region and zero elsewhere. By propagating the spatial (σ_r) and temporal (σ_t) resolution, we can derive the following expression for the error in the total momentum:

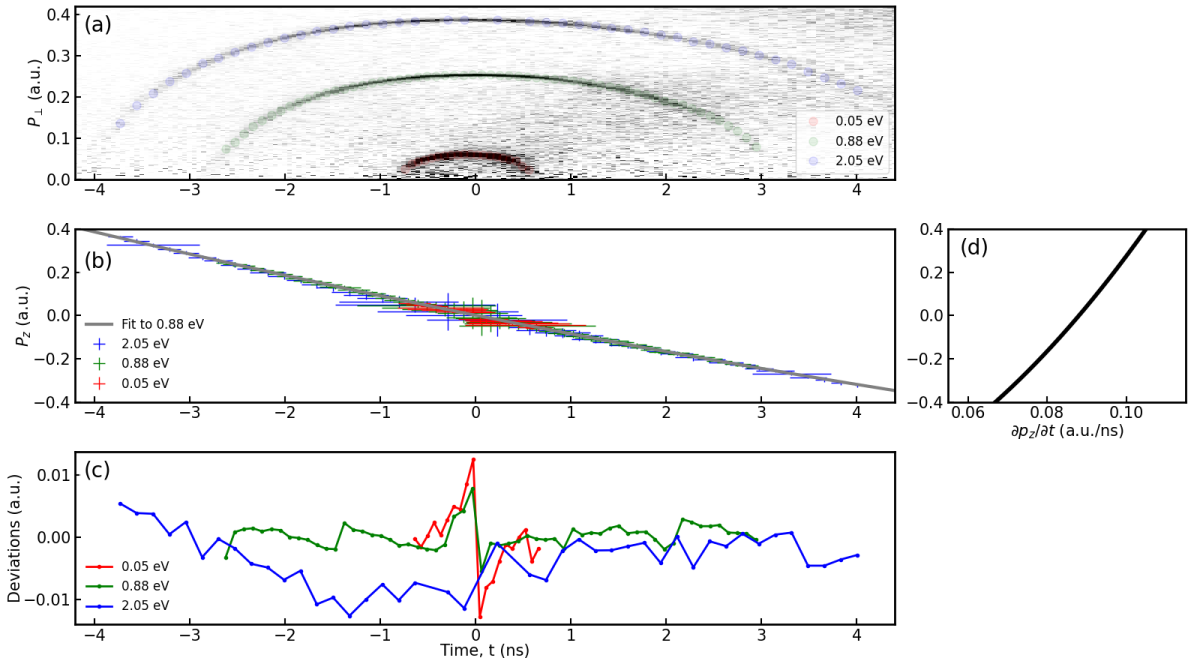

Figure 4.12: Demonstration of the process for fitting a calibration function to transform measured 3D VMI photoelectron time coordinates (t) into longitudinal momentum (p_z). The plots depict three sharp velocity-mapped nitric oxide photoionization features (described in Section 3.6.3) with known photoelectron energies of 0.05, 0.88, and 2.05 eV, for the inner, middle, and outer rings, respectively. The transverse momentum $p_\perp$ is obtained from equations 4.1 and 4.3 and is plotted against the corresponding time-stamp data in (a). We obtain (t, $p_\perp$) coordinates characterizing each photoelectron ring (the coloured markers in panel (a)) by taking narrow slices in time and performing a Gaussian fit to find the maximum. In (b), we convert these coordinates into (t, p_z) using equation 4.4 and fit a second-order polynomial to obtain a functional form for the transformation. The polynomial fit in (b) uses only the 0.88 eV data, and demonstrates a strong agreement with the other (0.05 and 2.05 eV) data sets. The deviations of all three data sets from the 0.88 eV fit are shown in (c). The derivative of this polynomial function (d) shows the dependence of the error ($\partial p_z/\partial t$) on p_z. This demonstrates that the errors are not uniform along the z (t) axis, but rather are smaller for the photoelectrons that were initially ejected away from the detector, where the time-stretching is stronger.

$$\begin{aligned}\sigma_p &= \frac{1}{p}\sqrt{(p_\perp \sigma_{p_\perp})^2 + (p_z \sigma_{p_z})^2}, \\ &= \frac{1}{p}\sqrt{\left(p_\perp \frac{\partial p_\perp}{\partial r}\sigma_r\right)^2 + \left(p_z \frac{\partial p_z}{\partial t}\sigma_t\right)^2}, \\ &= \frac{1}{p}\sqrt{(p_\perp K \sigma_r)^2 + \left(p_z \frac{\partial F(t)}{\partial t}\sigma_t\right)^2}. \end{aligned} \tag{4.7}$$

We can then construct an analogous simulated mask by filling in the region $p \pm \sigma_p$ as shown in Figure 4.13 (c) using approximate values of $\sigma_t = 100$ ps and $\sigma_r = 1$ pixel. Here I have superimposed the experimental mask for comparison. We fit the simulated mask M_S to the experimental mask M_E by minimizing $\sum(M_S - M_E)^2$, where $(M_S - M_E)^2$ is zero where the masks overlap and one elsewhere. The result of the fit is shown in Figure 4.13 (d) and we obtain values of $\sigma_t = 72$ ps and $\sigma_r = 1.6$ pixels for our temporal and spatial resolution, respectively. The Residual Sum of Squares (RSS) (plotted with a logarithmic scale in Figure 4.14) shows that the minimum is well-localized.

As previously noted in Section 4.2.1.2, our camera setup is capable of sub-pixel spatial resolution for individual events. The slightly poorer resolution emerging from this analysis is likely (supported by SIMION trajectory simulations) due to a broadening effect associated with photoelectrons being produced over the extended laser/molecular beam interaction volume (a $\sim$100 μm diameter by $\sim$2 mm long cylinder aligned parallel to the detector face). We could improve the spatial resolution by optimizing the VMI electrode voltages to improve focusing, or by reducing the extent of the interaction volume by skim-

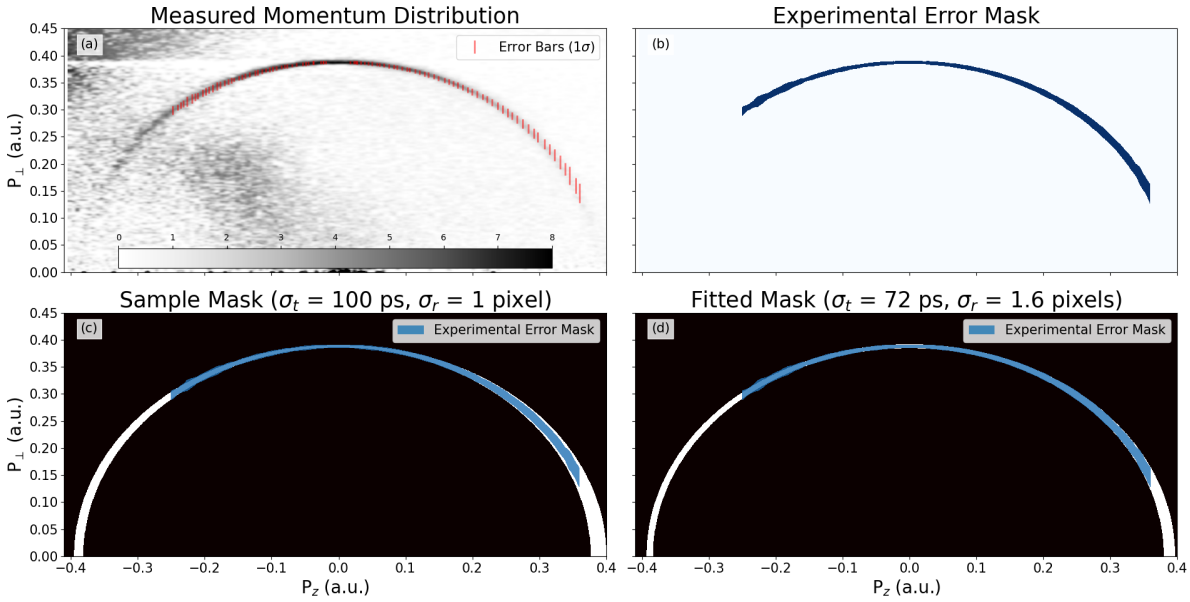

Figure 4.13: The least-squares fitting process used to characterize the spatial and temporal resolution of a 3D VMI experiment from the measured momentum uncertainties. All four panels show the same range of $p_\perp$ (momentum in the plane of the MCP detector) and p_z (momentum along the flight axis). In (a), we show the experimental momentum distribution obtained from the 2.05 eV nitric oxide photoionization data set (equation 3.7). The vertical lines (red) are 1σ error bars obtained by taking slices of the data in the $p_\perp$ direction and performing a series of Gaussian fits. The 1σ confidence region defined by the error bars is shown in (b). Note that I have excluded values of p_z below ~ -0.25 a.u. to avoid the noise present in this region. In (c) I show a simulated confidence region (white) with example error values of $\sigma_t = 100$ ps and $\sigma_r = 1$ pixel. The experimental region from (b) is superimposed for comparison (blue). After performing a least-squares fit to minimize the difference between the simulated and measured regions (d), we obtain values of $\sigma_t = 72$ ps and $\sigma_r = 1.6$ pixels for our experimental resolution.

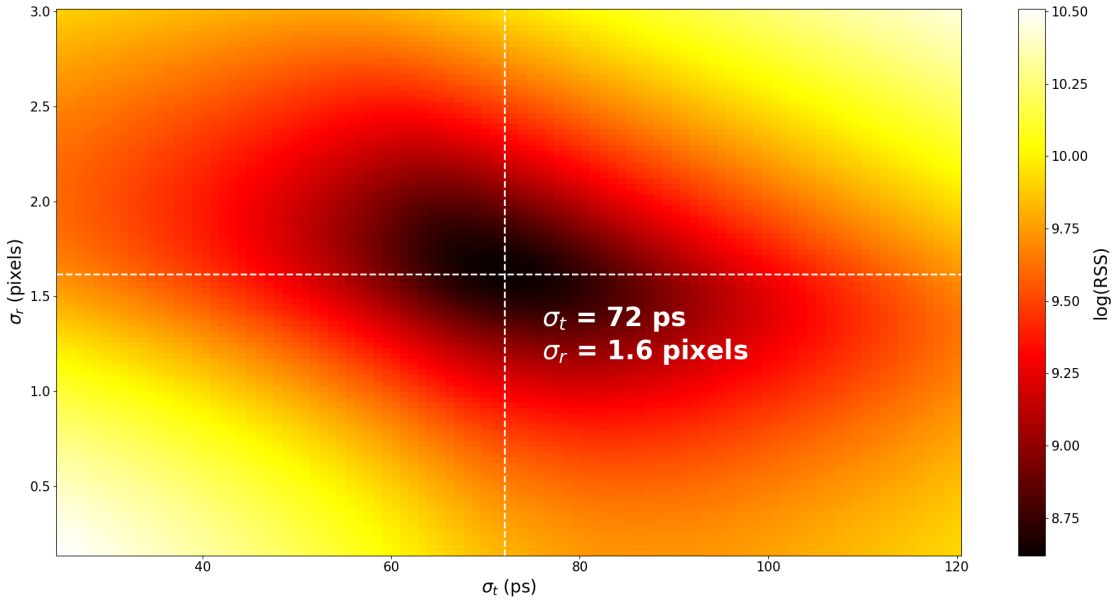

Figure 4.14: The RSS for the two-parameter (σ_t and σ_r) least-squares fit to the experimental momentum errors (shown in Figure 4.13). The minimum RSS is well-localized and found at values of $\sigma_t = 72$ ps and $\sigma_r = 1.6$ pixels.

ming the molecular beam, although this would be accompanied by a drop in yield.

We also estimated our experimental time resolution by measuring the arrival-time distribution of scattered UV photons which directly hit the MCP. With the ∼100 ps time resolution and short flight distances in this experiment, UV photons scattered from the walls of the VMI spectrometer can be assumed to arrive at the MCP detector simultaneously, with zero time delay. The observed time spread in UV photon events therefore represents our experimental time resolution. We obtained a sample of photon scatter by collecting data without the molecular beam (laser only) and by reversing the polarity of the VMI electrodes so as to prevent any photoelectron signals from being detected. This photon scatter signal is shown in Figure 4.15 along with a Gaussian fit to determine its width. We fit a time resolution of $\sigma = 99.8$ ps using this method. Importantly, we observed signal (shoulders) on either side of the central peak which does not fit the expected Gaussian distribution. This represents scatter from other surfaces (for example, input and exit windows) within the spectrometer. As such, I have excluded these shoulders from the fit in order to better characterize the central peak. Despite this, the presence of multiple scattering sources contributing to the width of the fit likely make this an overestimation of our actual time resolution. For this reason, I believe that our actual time resolution is closer to the value obtained from the photoelectron distribution (72 ps), as described above.

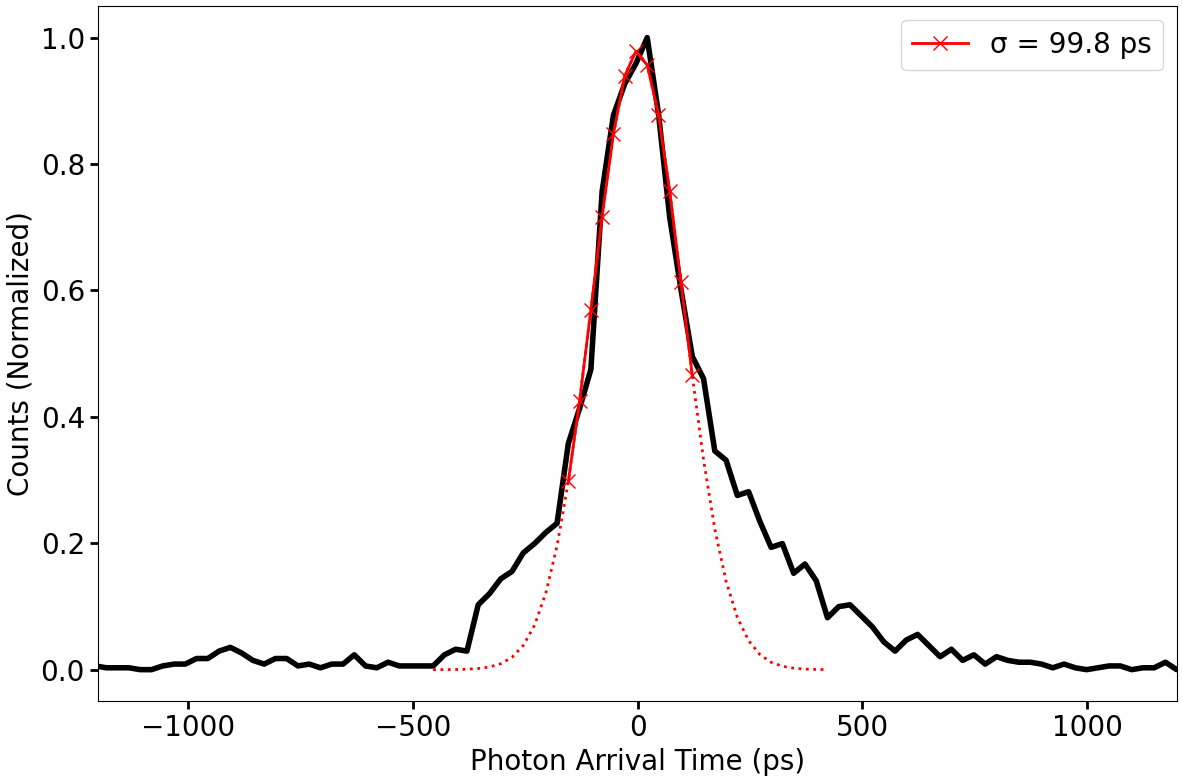

Figure 4.15: Estimation of the temporal resolution of a 3D VMI experiment using the arrival-time spread of photons scattered within the spectrometer. This approach relies on the assumption that, on the time scale of the experiment, all scattered photons can be considered to arrive at the detector simultaneously. Thus, the observed spread in the photon scatter signal is a measurement of the timing resolution of the experiment. The signal shown here contains structure on either side of the main peak, which is attributed to photon scatter from multiple surfaces within the instrument. To obtain an accurate fit to the main peak, I have included only the central data points in the fit (solid line with markers) and excluded the outer points (dotted line). Nevertheless, the resulting fitted width of $\sigma_t = 99.8$ ps may still be an overestimation of our actual experimental time resolution, as detailed in the text.

4.5 Angular Misalignment Corrections

In many experiments, the laser polarization state is set outside the spectrometer before passing through an entrance window into a vacuum-pumped chamber. This window may be subject to some degree of stress-induced birefringence as a result of the pressure differential. Practically, it is difficult to accurately measure the polarization vectors inside the interaction region of the spectrometer and thus the actual polarization state used in the experiment can be slightly different than the one prepared. With access to the full 3D photoelectron momentum vectors, we can extract the actual internal polarization state and apply the appropriate angular correction to align the (linear) polarization axis of the data with the intended direction.

In the case of the nitric oxide photoionization data used here, the polarization axis ($\vec{E}$) was set outside the spectrometer to be in the plane of the detector (horizontal polarization). In addition, no effort was made to align the mounting angle of our camera with the axes of our coordinate system as it is simple to perform this correction during analysis. We wish to align the polarization axis $\vec{E}$ with the x-axis, which is the axis in the detector plane perpendicular to the laser propagation vector $\vec{k}$. In general two rotations are required to accomplish this, and I define these as χ_z and χ_y for rotations about the z and y axes, respectively. Rotation about the polarization axis itself (which might be caused by $\vec{k}$ being slightly out of plane with the detector) is not being considered here since the PAD possesses symmetry about this axis and a rotational correction would be meaningless.

The process used to determine these correction angles is shown in Figure 4.16. The angular deviation χ_z arises from the mounting angle of our camera and is large as expected ($27.5°\pm0.6°$ for the data in Figure 4.16). The deviation χ_y is the difference between the actual and intended polarization vector angle of the laser. This angle is small ($7.3°\pm0.5°$ for the data in Figure 4.16) but not zero despite our efforts to correctly set this laser polarization. This shows that significant deviations in the polarization vector can be present in such an experimental setup and must be accounted for in any polarization-sensitive work.

The values of χ_z and χ_y shown in Figure 4.16 were obtained by respectively extracting the central XY- and XZ-plane angular slices (width 10°) from the overall photoelectron momentum distribution and integrating the yield radially over the extent of the signal (in this case, a ∼20 meV region around 0.88 eV). I have plotted the yield in each slice along the appropriate polar angle (θ_{XY} for the XY slice (a); θ_{XZ} for the XZ slice (b)) and fit this to an even Legendre polynomial expansion with an angular offset:

$$Y(\alpha) = A(1 + \frac{\beta_2}{2}(3x(\alpha)^2 - 1) + \frac{\beta_4}{8}(35x(\alpha)^4 - 30x(\alpha)^2 + 3)), \tag{4.8}$$

where $x(\alpha) = \cos(\alpha + \chi_z)$ or $x(\alpha) = \cos(\alpha + \chi_y)$ depending on the slice. Although we obtain the asymmetry parameters β_2 and β_4 from these fits, only the values of χ_z and χ_y are needed during this step; final fitting of the asymmetry parameters is performed only after additional processing steps are applied to the data.

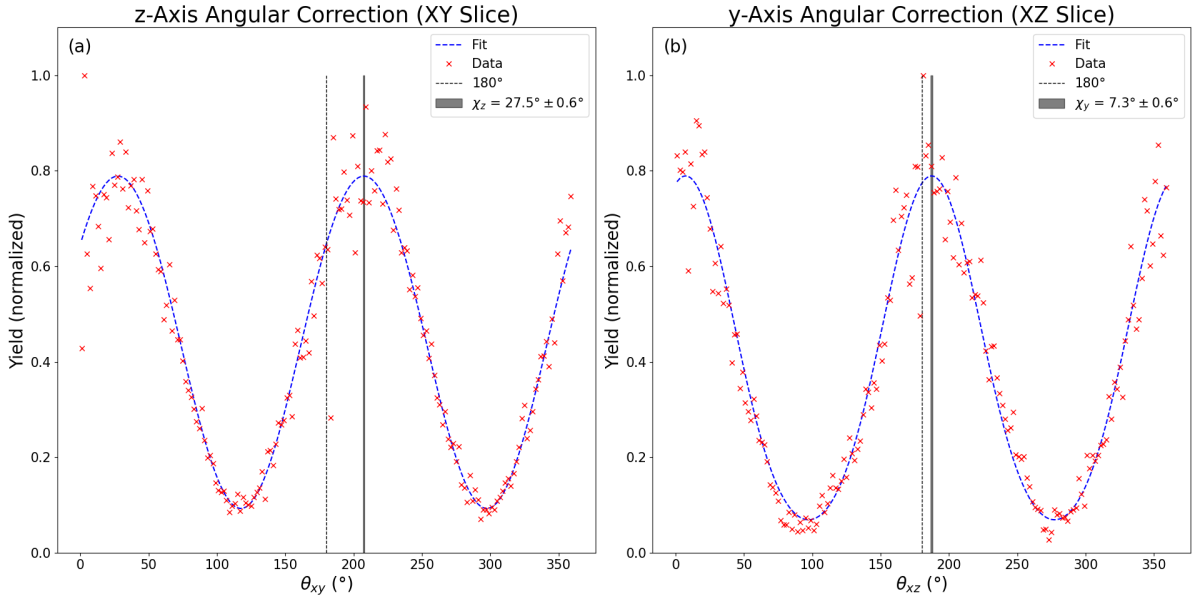

Figure 4.16: 3D angular corrections about the (a) z and (b) y axes relative to the coordinate system defined in Figure 2.1. The angular correction χ_z about the z axis represents a rotation of the camera's XY plane, and the angular correction χ_y about the y axis represents a rotation of the XZ plane due to a small deviation of the laser polarization vector away from the intended detector-plane alignment. The single-photon ionization of NO(A) at 266 nm (equation 3.6) results in a nearly-pure p-wave distribution which is peaked along the polarization axis (details in Section 3.6.3). We therefore expect the maxima in the yield in the XY (a) and XZ (b) planes to each align with the x axis. Taking slices in these planes, I have plotted the variation in yield with the 2D polar angle (θ_{xy} for the XY slice and θ_{xz} for the XZ slice). We obtain the angular corrections by fitting the angular signal variation in each slice to an even Legendre expansion (equation 4.8). For this data set, the angular corrections were found to be $\chi_z = 27.5° \pm 0.6°$ and $\chi_y = 7.3° \pm 0.6°$.

4.6 Polarization-Based Angles

The data processing steps in the following sections are best described using a new set of angles defined relative to the polarization axis, which I will call "polarization-based angles", shown in Figure 4.17. These angles use the same coordinate system defined in Figure 2.1. The remaining sections of this chapter involve only the NO data sets (see Section 3.6.3) where the laser was polarized along the x-axis. As such, these angles are defined with the assumption of an x-aligned laser polarization. These polarization-based angles consist of a polar angle ϕ' measured from the $+x$ axis with the range [0, π], and an azimuthal angle θ' measured from the $-z$ axis toward the $-y$ axis with the range [0, 2π].

4.7 MCP Gain Correction

An inhomogeneous gain response over the spatial extent of an MCP detector, particularly a dip in gain near the central region, is a well-known and common result of usage wear. As our current MCP detector has been in use for many years, this effect is noticeable in any PAD derived from our measurements, particularly for slices which include the z axis, such as the XZ and YZ slices (for example, in Figure 3.13 in Section 3.6.3). For many 2D VMI experiments, this is not a major concern since the data are collected and analyzed as an integrated 2D image and the central region of the MCP makes only a small contribution to this. For 3D VMI, we want to make use of our ability to separately analyze and compare arbitrary slices through the PAD. For such an analysis to be valid, we must characterize

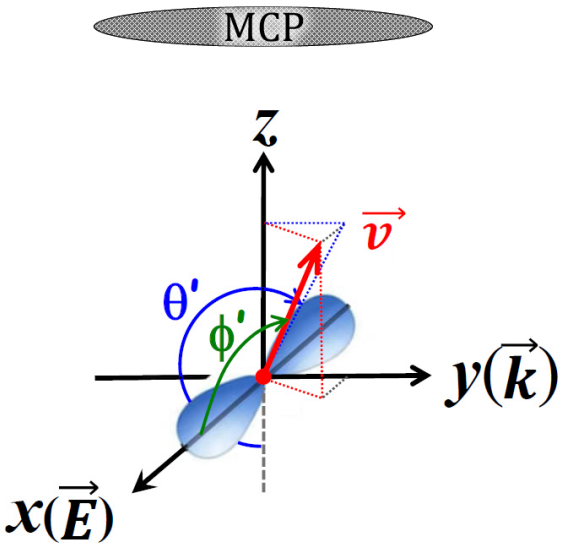

Figure 4.17: Lab frame coordinate system and definition of the polarization-based angles. I define the z axis to be the axis normal to the plane of the MCP detector, with y being the axis of propagation of the laser ($\vec{k}$). For the experiments discussed here, the laser is polarized along the x axis. I define the polarization-based angles θ' and ϕ' as shown, where the polar angle ϕ' is measured from the $+x$ (polarization) axis, and the azimuthal angle θ' is measured from the $-z$ axis toward the $-y$ axis.

and correct these gain inhomogeneities.

In Figure 4.18 I demonstrate our process for obtaining a functional form for the radial variation of the gain from the centre of the detector, which is subsequently used to correct our data. The PADs shown in Figure 4.18 are all prepared by binning the (p_x, p_y, p_z)-data into spherical coordinates, integrating radially over the feature of interest, and smoothing to enhance the main features. Since this gain behaviour is primarily a result of usage wear on the detector, this procedure should be repeated periodically to account for continuing wear over time. In Figure 4.18 (a) I show an experimental PAD (single-photon ionization of NO(A) at 213 nm) integrated over the radius corresponding to the maximum of the photoelectron signal (2.05 eV in this case), using the previously defined polarization-based angles. The PAD shown here is not corrected for the y-axis rotational misalignment (χ_y) and the associated curvature can be seen in the image. This is because it is better to derive each correction independently from the other corrections, then apply all of them at once during the final analysis.

For the data used in this gain inhomogeneity characterization, we found that the signal in the region from $\theta' = 0°$ to 90° and 250° to 360° (the lower half, or $-z$ region, and 20° on the $+y$ side of the upper half) was corrupted by the presence of background (likely laser scatter) and had to be rejected. Only the region between the black lines in Figure 4.18 (a) was used, and the regions above and below the black lines are mirrored from the $\theta' = 90°$ to 180° region. Despite the corruption of part of the PAD, this 2.05 eV data set was used here because it fills the largest region of the detector and thus provides the most complete sampling of the MCP area. In addition, only the front half of the PAD is used

in this characterization (see below) so the mirroring has minimal effects on the resulting correction.

The effect of the gain inhomogeneity can be seen clearly in the variation in yield with θ', which has minima at 0°, 180°, and 360°, all corresponding, upon mirroring, to the $+z$ axis (the centre of the detector). In Figure 4.18 (b) we normalize the yield by dividing the PAD (a) by the average signal at $\theta' = 90°$, where the gain degradation is smallest (between the dashed lines). This removes the variation associated with the anisotropic nature of the PAD, isolating the MCP gain variation (and the effect of the χ_y angular offset). We compare this with an ideal simulated PAD (c), generated as a set of (p_x, p_y, p_z) data points possessing the same χ_y (4.1°) and approximate β_2 (1.17) as the experimental data and processed in the same way. The normalized simulated PAD (also divided by the signal at $\theta' = 90°$) is shown in panel (d), in which the alternating structure is a result of the simulated angular misalignment. The ratio of normalized experimental and simulated PADs (e) represents the deviation from the ideal behaviour and forms the basis of the correction.

In Figure 4.18 (f) I plot the θ' and ϕ' dependence of the gain, taken from the following regions enclosed by dashed lines in panel (e). For θ', I plot only from 90° to 270° (the $+z$ half of the data), and integrate from $\phi' = 75°$ to 105° (a 30° slice centred at 90°). For ϕ', I plot the full range (0° to 180°) and integrate from $\theta' = 165°$ to 195° (a 30° slice centred at 180°). Using the maximum radius r_{max} of the experimental photoelectron signal (that is, the radial position of the maximum yield in a centre slice through the XY plane), we

convert θ' and ϕ' both into the 2D radial position r on the detector:

$$r = r_{max}|\sin(\theta' - 180°)| = r_{max}|\sin(\phi' - 90°)|. \tag{4.9}$$

The average of these gives the radial dependence of the MCP gain, shown in panel (g). We found that the gain curve in (g) is well-represented by a damping function of the form

$$G(r) = Ae^{-\lambda r^{3/2}}, \tag{4.10}$$

with amplitude A and damping constant λ, and we used a fit to this function (dashed line) to characterize and later correct for the radial gain dependence. Note that although this correction was developed using a specific data set, the ratio of experimental and simulated PADs (e) represents the data-set-independent detector gain. The correction is thus general to any data collected using this MCP detector with the current level of wear.

4.8 Asymmetry Parameter Fitting

Having corrected the experimental PADs for both the angular misalignment and MCP gain inhomogeneity, we can now perform fits to determine the photoionization asymmetry parameters (β_2 and β_4) for each of the measured NO photoionization signals (equations 3.5–3.7). I present these fits in figures 4.19–4.21. The NO(A) state is well known to be predominately of 3s Rydberg character [53, 65–67]; single-photon ionization of NO(A) results

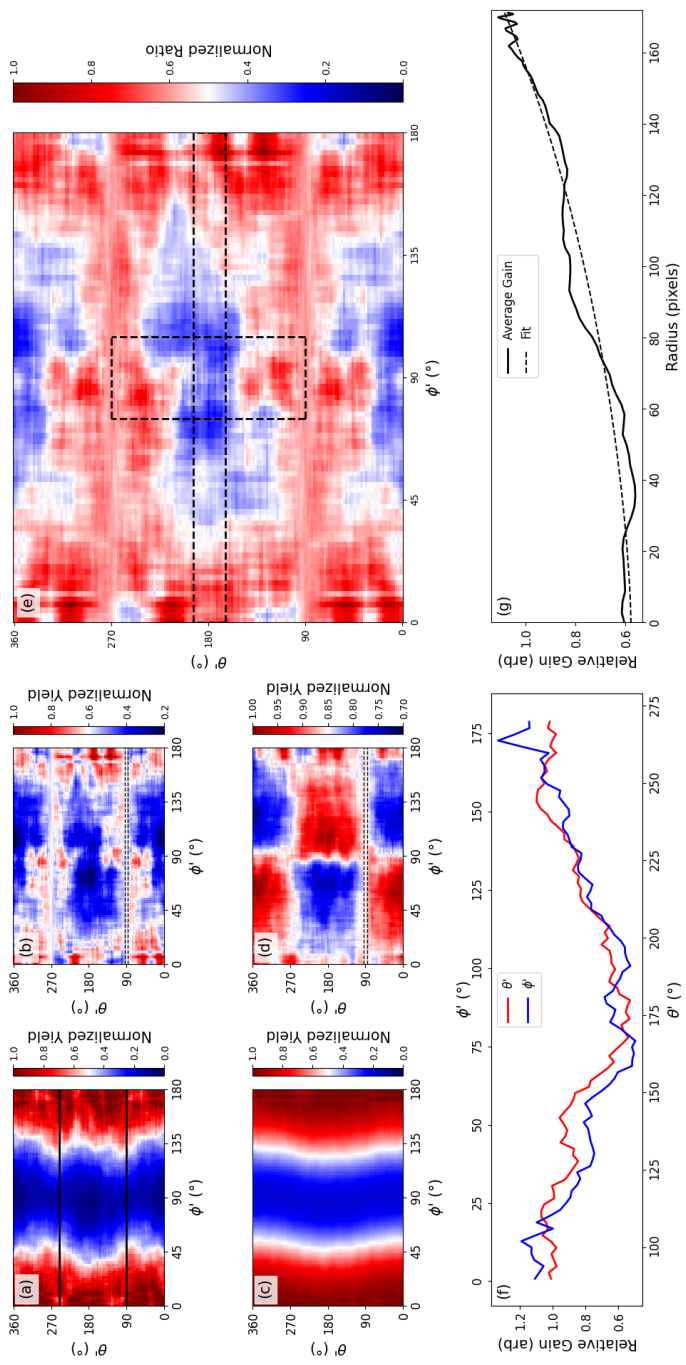

Figure 4.18: Illustration of the method we used to characterize the radial variation of the MCP detector gain in a 3D VMI experiment. We first took the angular yield of an experimental NO(A) photoionization signal (details in the text) at 2.05 eV (a) and divided it by the yield at $\theta' = 90°$ to give a normalized yield (b). The dashed lines in (b) show the region used for the normalization. For this data set, we only used the region between $\theta' = 90°$ and $\theta' = 250°$ (indicated by the solid black lines in (a)); the data outside this region were generated by mirroring as described in the text. In panels (c) and (d), I show the angular yield (c) and normalized (by the signal at $\theta' = 90°$) yield (d) for a simulated data set with the same angular distribution and y-axis rotation (χ_y) as the experimental data. In (e), I show the ratio of the experimental normalized yield and the simulated normalized yield, revealing the angular deviation between the two. We extracted the radial dependence by taking the region between the dashed lines (θ' from 90° to 270°, integrated between $\phi' = 75°$ and $\phi' = 105°$; ϕ' from 0° to 180°, integrated between $\theta' = 165°$ and $\theta' = 195°$) and plotting these in (f). The angular axes in (f) were then converted into a radial axis using the known peak radius of the experimental data (equation 4.9), then averaged to give the radial MCP gain curve (g). We fit this curve to a damping function (equation 4.10) to obtain a radial correction that can be applied to any data collected with this MCP detector.

in a nearly-pure p-wave photoelectron distribution, aligned along the laser polarization direction. Prior studies of this state have reported values between 1.2 and 1.7 [52, 53] for β_2. The value of β_4 is expected to be zero for a pure p-wave distribution. Although the polarization geometry of these data sets (laser polarized along x) yields a cylindrically symmetric PAD compatible with 2D inversion-based methods, it is important to note that this is a full 3D measurement applicable to arbitrary polarization states.

We fit the asymmetry parameters to a standard even Legendre polynomial expansion up to fourth order:

$$Y(\cos\alpha) = A(1 + \frac{\beta_2}{2}(3\cos^2(\alpha) - 1) + \frac{\beta_4}{8}(35\cos^4(\alpha) - 30\cos^2(\alpha) + 3)). \tag{4.11}$$

Owing to our instrument's 3D imaging capabilities, I am able to calculate and report asymmetry parameters for not only the full integrated data set (panel (a) in figures 4.19–4.21), but also for any arbitrary meridional slice through the 3D distribution. This represents a stringent test of our 3D recoil momentum vector imaging capability. For each of figures 4.19–4.21, I show the full fits to the XY (b) and XZ (c) slices, as well as the resulting β_2 (e) and β_4 (f) from fits to each 10°-width slice in θ' through the full 3D distribution (d). The shaded region in panels (e) and (f) represent 1σ uncertainty in the fits.

For the 0.88 eV (Figure 4.19) and 2.05 eV (Figure 4.20) NO(A) photoionization data sets, I find β_2 values of 1.488 ± 0.023 and 1.174 ± 0.032, respectively. These are in agreement

with the previously reported values, although not with each other. The values of β_4 were found to be 0.114 ± 0.019 (0.88 eV) and -0.007 ± 0.029 (2.05 eV). These values are small, as expected, and the latter is consistent with zero to within one standard deviation.

The fits to the XY and XZ slices for the 0.88 eV data set give 1.47 ± 0.03 and 1.54 ± 0.04 for β_2, respectively. The values for β_4, 0.07 ± 0.03 (XY slice) and -0.03 ± 0.04 (XZ slice) are consistent with zero to within 3σ and 1σ, respectively. The absolute difference between the β_2 values for the two slices is small (0.07) and they are in agreement with each other to within the uncertainty of the fits. The variation (with θ') of the asymmetry parameters over the distribution shows relatively stable behaviour over the full angular range.

For the 2.05 eV data set, the XY and XZ slices give β_2 values of 1.13 ± 0.03 and 1.19 ± 0.09, respectively. The values of β_4 are -0.07 ± 0.03 (XY slice) and -0.21 ± 0.08 (XZ slice), which are both consistent with zero to within 3σ. The β_2 values for the XY and XZ slices are in agreement with each other and have a small absolute difference of 0.06. As with the 0.88 eV data set, the variation of the asymmetry parameters with θ' is relatively stable.

The 0.05 eV data results from two-photon ionization of the NO(X) monomer ground state and, therefore, β_4 is no longer expected to be zero. For this data set, the integrated fit gives values of $\beta_2 = 1.585 \pm 0.043$ and $\beta_4 = 0.590 \pm 0.036$. The results for the XY ($\beta_2 = 1.65 \pm 0.10$, $\beta_4 = 0.69 \pm 0.08$) and XZ ($\beta_2 = 1.65 \pm 0.17$, $\beta_4 = 0.57 \pm 0.14$) slices are consistent with each other and with the integrated result. As above, the asymmetry parameters are relatively stable over the full range of θ'.

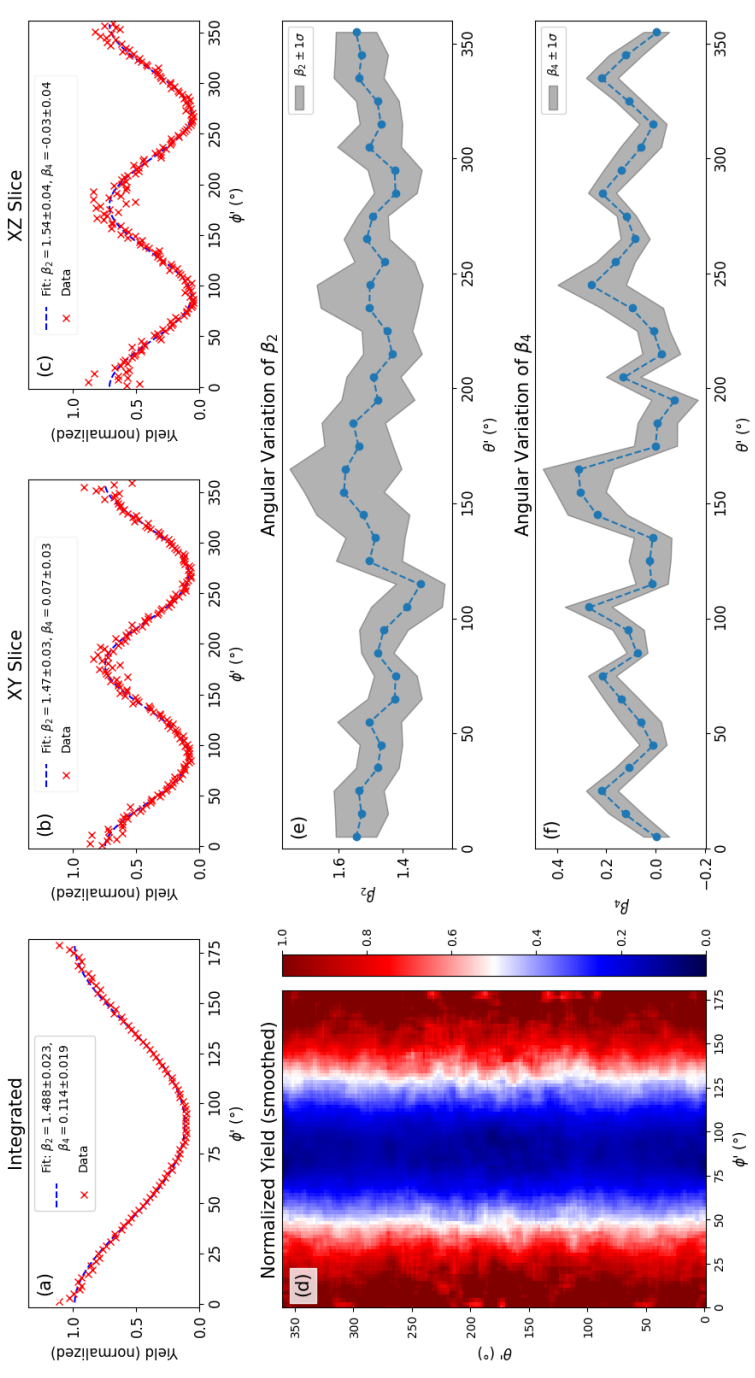

Figure 4.19: The PAD and asymmetry parameter fits (β_2 and β_4) for the photodissociation of $(NO)_2$ at 213 nm and subsequent photoionization of the NO(A) photoproducts at 266 nm, resulting in a photoelectron distribution with a peak signal intensity at 0.88 eV (equation 3.6). We obtain β_2 and β_4 by fitting the PAD to the Legendre expansion in equation 4.11. To demonstrate the 3D imaging capability of our instrument and technique, I present fits of the integrated 3D PAD (a), the XY slice (b), and the XZ slice (c). In addition to these selected slices, I show the full 3D PAD (d), from which I can extract and fit arbitrary slices. The angular variation of β_2 (e) and β_4 (f) are obtained by fitting 10° slices along θ', with 1σ error bars indicated by the grey shaded area. Note that mirroring has been applied to the data, as described in the text.

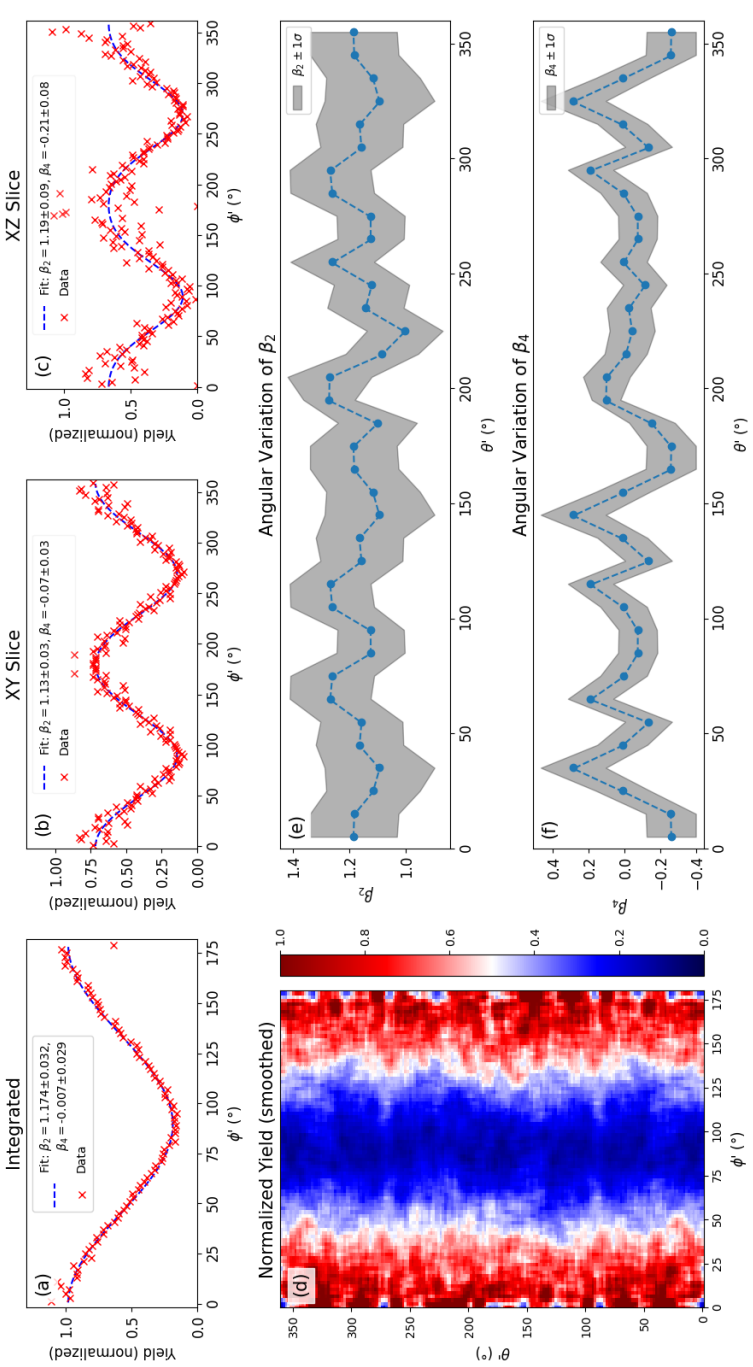

Figure 4.20: The PAD and asymmetry parameter fits (β_2 and β_4) for the photodissociation of $(NO)_2$ at 213 nm and subsequent photoionization of the NO(A) photoproducts at 213 nm, resulting in a photoelectron distribution with a peak signal intensity at 2.05 eV (equation 3.7). We obtain β_2 and β_4 by fitting the PAD to the Legendre expansion in equation 4.11. To demonstrate the 3D imaging capability of our instrument and technique, I present fits of the integrated 3D PAD (a), the XY slice (b), and the XZ slice (c). In addition to these selected slices, I show the full 3D PAD (d), from which I can extract and fit arbitrary slices. The angular variation of β_2 (e) and β_4 (f) are obtained by fitting 10° slices along θ', with 1σ error bars indicated by the grey shaded area. Note that mirroring has been applied to the data, as described in the text.

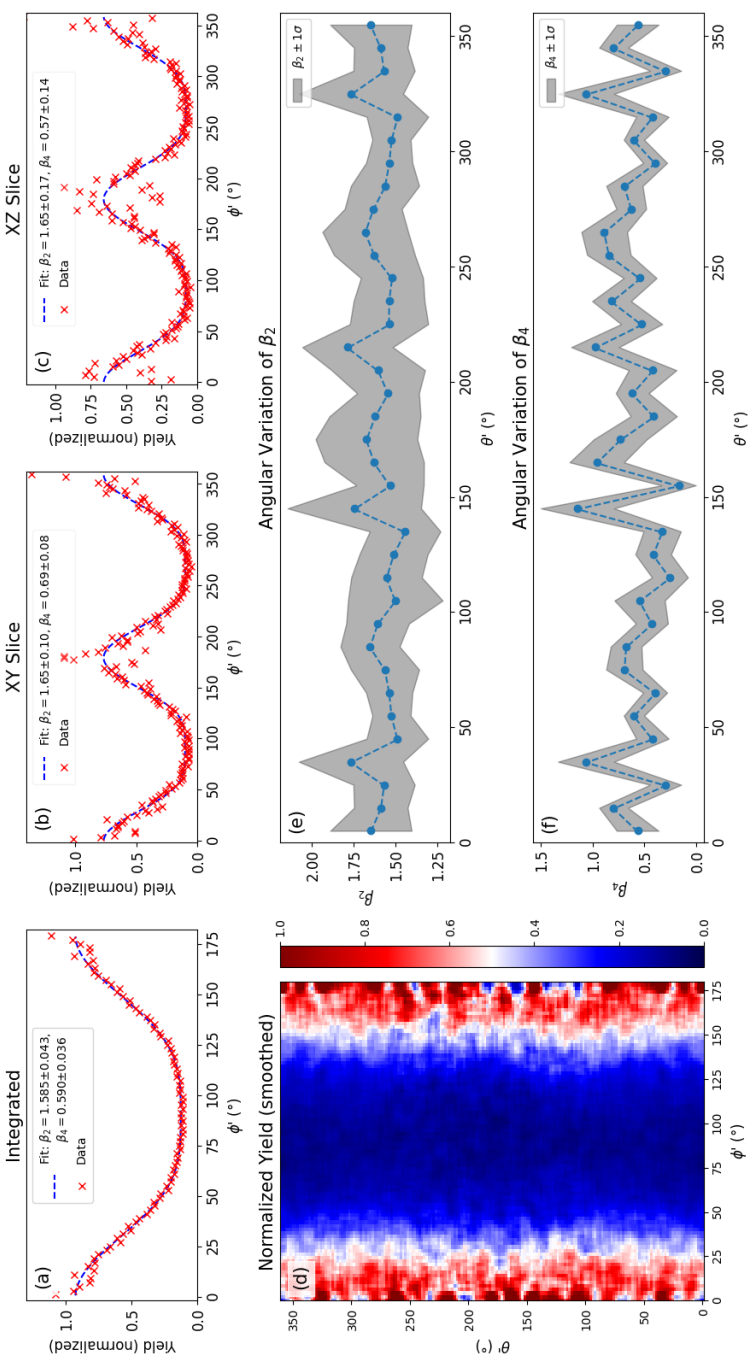

Figure 4.21: The PAD and asymmetry parameter fits (β_2 and β_4) for the two-photon photoionization of the NO(X) monomer at 266 nm, resulting in a photoelectron distribution with a peak signal intensity at 0.05 eV (equation 3.5). We obtain β_2 and β_4 by fitting the PAD to the Legendre expansion in equation 4.11. To demonstrate the 3D imaging capability of our instrument and technique, I present fits of the integrated 3D PAD (a), the XY slice (b), and the XZ slice (c). In addition to these selected slices, I show the full 3D PAD (d), from which I can extract and fit arbitrary slices. The angular variation of β_2 (e) and β_4 (f) are obtained by fitting 10° slices along θ', with 1σ error bars indicated by the grey shaded area. Note that mirroring has been applied to the data, as described in the text.

4.9 Conclusion

In this chapter, I described all of the data processing and analysis steps necessary to convert raw 3D VMI data (images and waveforms) into time-stamped particle coordinates (x, y, t) and then into the instrument-independent initial recoil momentum vectors (p_x, p_y, p_z) which may be compared with theory. This includes first processing the camera images to obtain the transverse spatial coordinates (x, y) and fitting the arrival time peaks in the waveforms to obtain the time-stamps (t). In the general case of multi-hit data, the spatial coordinates and time-stamps are then assigned to each other to form the 3D coordinates for each event. I developed and demonstrated a method for fitting a calibration function which can transform the time-stamps into the momentum component along the z-axis; this approach is general for any similar time-stamping charged particle imaging technique. This set of data processing steps also includes corrections for two common sources of error that may occur in 3D VMI data: misalignment of the laser polarization vector, and usage-induced gain degradation of the central region of the MCP detector.

Finally, I demonstrated that, after processing the data as described, we achieve agreement with previously published asymmetry parameter fits to the PAD for the well-understood photoionization of the excited A-state of nitric oxide. This confirms the precision and accuracy of our approach.

5

Four-Wave Mixing Theory for Cylindrical Vector Beams

5.1 Preface

The contents of this chapter are adapted from a paper [1] published in the Journal of the Optical Society of America B, titled "Theory of four-wave mixing of cylindrical vector beams in optical fibers", representing my work on a theoretical nonlinear optics project. As such, the contents of this chapter are independent and self-contained. My contributions to this publication include conceptualization of the general approach, most of the mathematical derivations, early drafts of the article, and feedback on the final version. Although most of the following is my own work, it naturally also contains the efforts of

the other authors, primarily Dr. Connor Kupchak, and these are included for context and completeness.

5.2 Introduction

The term "structured light" refers to optical beams whose intensity, phase, or polarization are non-uniform across the beam's transverse profile. Cylindrical vector (CV) beams are a type of structured light which exhibits intensity and polarization profiles that are spatially dependent but also displays cylindrical symmetry under discrete rotations about the beam axis. Specifically, the modes for such structured light beams, CV modes, are described by a superposition of product states between the spin angular momentum (SAM) and orbital angular momentum (OAM) degrees of freedom. CV beams are also the eigenmodes of optical fibre and, as such, are of wide-spread practical importance in photonics. The degrees of freedom available in the transverse profile of an optical beam can benefit many applications, including mode-division spatial multiplexing techniques to increase communications bandwidth. Such multiplexing techniques have already been demonstrated with scalar OAM modes [73, 74] and vector beam modes [75, 76]. Similarly, from a quantum optics viewpoint, CV modes can increase the information capacity available in single photons via the additional OAM and SAM degrees of freedom. In free-space quantum key distribution, the rotational symmetry of CV modes can be exploited to alleviate the need for rotational alignment between sending and receiving parties [77]. Outside of communications, the behaviour of optical beams in CV modes are of fundamental interest. For

example, radially polarized CV modes focus to smaller spot sizes than Gaussian modes of a comparable beam waist [78], and azimuthally polarized CV modes produce a longitudinal magnetic field at their focus.

Thorough understanding of nonlinear optical processes is vital to many practical applications. An example is in telecommunications, where unwanted nonlinear interactions between optical pulses in fibre present a roadblock for increasing signal power, and thus the bandwidth [79]. To date, the theory for four-wave mixing (FWM) in fibre systems has been developed for both uniformly polarized light [80, 81] and spatial modes [82–84]. Despite their potential for increasing communications bandwidth, the nonlinear optics of structured light and vector modes is in its infancy [85–87]. In order to address this, we have derived the coupled amplitude equations that describe FWM of CV and other complex modes. From this, we derived a set of general selection rules for the allowed mixing processes between CV modes. As shown below, FWM processes can convert photons between different CV modes and may therefore provide insight into conversion processes that involve structured light and the conservation of angular momentum of light. Moreover, these FWM transitions can potentially produce mode-entangled CV photons through spontaneous FWM (SFWM) photon pair generation.

5.3 Optical Mode Descriptions

I begin with a review of optical spatial modes that includes a succinct general mathematical description suitable for deriving selection rules. Starting with the nonlinear optical

wave equation, I derive general coupled-amplitude equations for FWM of fields that vary spatially in intensity, phase, and polarization. The derived FWM theory applies to both free-space environments such as bulk nonlinear media and to weakly-guiding cylindrically-symmetric fibres (examples are shown in Figure 5.1). I then focus on examples of these fields, particularly CV modes, but also circularly polarized OAM modes, and modes that are eigenstates of the total angular momentum (TAM) along z, the beam or fibre axis. For each of these examples, I derive and present selection rules. These spatial mode solutions are highly relevant in the development of photonic devices, since the CV modes are eigenmodes of all weakly-guiding cylindrically-symmetric waveguides, such as standard optical fibres [88, 89]. Furthermore, CV modes represent approximate solutions to the paraxial vector wave equation and, in free space, follow the intensity distribution of Laguerre-Gauss modes [90].

I begin by defining a general type of transverse mode which has a cylindrically symmetric intensity distribution:

$$\mathbf{M}^{[u]}(r,\phi) = R^{[m]}(r)\mathbf{\Phi}^{[n]}(\phi), \tag{5.1}$$

where r and ϕ are the standard radial and azimuthal cylindrical coordinates, respectively, and the full position vector is $\mathbf{r} = (r, \phi, z)$. Throughout, quantities in bold font are vectors. Here, $\mathbf{\Phi}^{[n]}(\phi)$ gives the vector azimuthal dependence of the polarization (which is strictly transverse) and phase for azimuthal mode index n. The scalar radial dependence $R^{[m]}(r)$ is usually implicitly dependent on n and will typically exhibit a number of rings

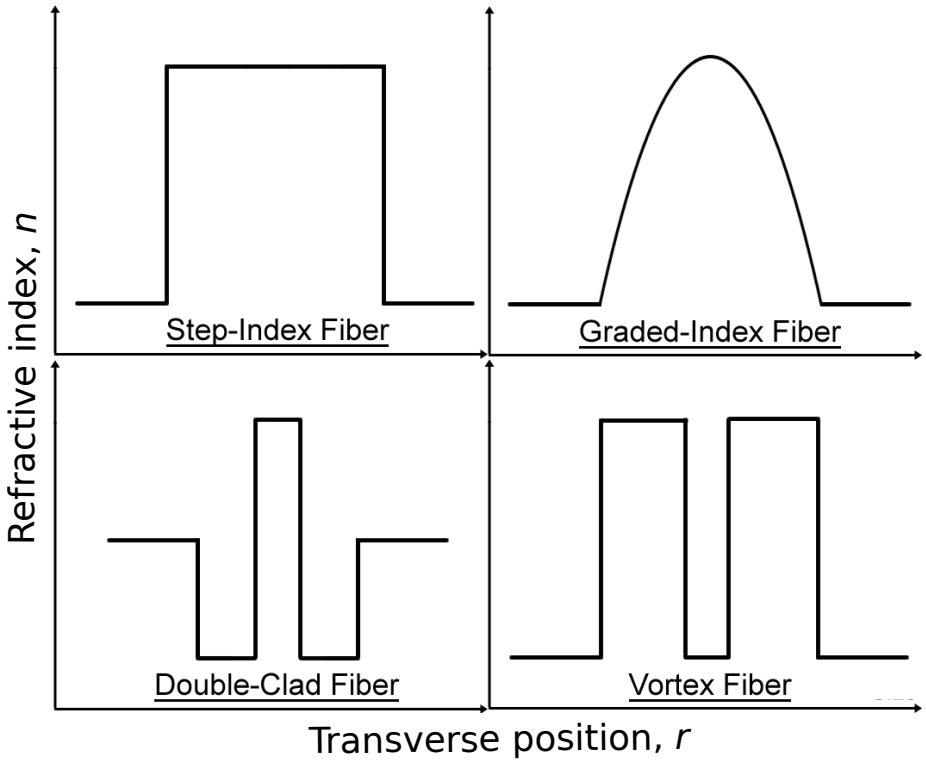

Figure 5.1: Four examples (step-index, graded-index, double-clad, and vortex) of fibre types which support CV modes and to which this work is directly applicable.

in the intensity profile that equal the radial mode index m. The combined mode index u is defined to be the set of radial and azimuthal mode indices: $u = (m, n)$. By defining $\mathbf{M}^{[u]}(r, \phi)$ in this way the intensity distribution is explicitly cylindrically symmetric.

Inside a cylindrically symmetric waveguide with weak guiding (small refractive index contrast), cylindrically symmetric modes are approximate paraxial solutions to the wave equation [90]:

$$\nabla^2 \mathbf{L}(\mathbf{r}, t) + \frac{n(r)^2}{c^2} \frac{\partial^2 \mathbf{L}(\mathbf{r}, t)}{\partial t^2} \approx 0, \tag{5.2}$$

where $n(r)$ is the cylindrically symmetric index profile and c is the speed of light in vacuum. Here, I have defined the normalized electric field solutions as

$$\mathbf{L}^{[u]}(\mathbf{r}, t) = \mathbf{M}^{[u]}(r, \phi) \exp(i\beta^{[u]} z) \exp(-i\omega t). \tag{5.3}$$

Each mode $\mathbf{M}^{[u]}$ has a potentially distinct effective wavevector $\beta^{[u]}$. Moreover, both the mode's radial dependence ($R^{[m]}$) and $\beta^{[u]}$ depend on the angular frequency of the field, ω. In contrast, given the cylindrical symmetry, $\mathbf{\Phi}^{[n]}(\phi)$ will be independent of frequency. Figure 5.1 shows a sample of common fibre structures to which this work is applicable.

The $\mathbf{M}^{[u]}(r, \phi)$ modes obey the following orthonormality relations [88] for transversely polarized modes:

$$\delta_{u,u'} = \int \mathbf{M}^{[u]^*}(r,\phi) \cdot \mathbf{M}^{[u']}(r,\phi) r \mathrm{d}r \mathrm{d}\phi, \tag{5.4}$$

$$\delta_{m,m'} = \int R^{[m]}(r) R^{[m']}(r) r \mathrm{d}r, \tag{5.5}$$

$$\delta_{n,n'} = \int \mathbf{\Phi}^{[n]^*}(\phi) \cdot \mathbf{\Phi}^{[n']}(\phi)(r,\phi) \mathrm{d}\phi, \tag{5.6}$$

where $\delta_{j,k}$ is the Kronecker delta. The integral in equation 5.4 is over the transverse plane and u and u' are combined mode indices as defined above. Equations 5.5 and 5.6 follow from equations 5.4 and 5.1. Since $R^{[m]}$ depends on frequency, equations 5.4 and 5.5 are only strictly true for fields of equal wavelength. The theory developed here focuses on the azimuthal mode function $\mathbf{\Phi}^{[n]}$, which describes the spatial polarization and phase variation of the modes. Since this variation is independent of wavelength, the $\mathbf{\Phi}$ orthonormality in equation 5.6 will also be wavelength independent. In any case, later derivations only rely on the orthonormality of modes at the same wavelength. In the next three subsections, I will introduce three types of $\mathbf{\Phi}$ modes, each of which forms a complete mode basis. Examples of the mode types are shown in Figure 5.2.

5.3.1 Definite Spin and Orbital Angular Momentum Modes

The azimuthal mode function $\mathbf{\Phi}^{[n]}(\phi)$ describes both SAM and OAM. This first type of mode has a definite value for both SAM and OAM along the system axis z (the beam or fibre axis). The SAM of a photon is given by its circular polarization,

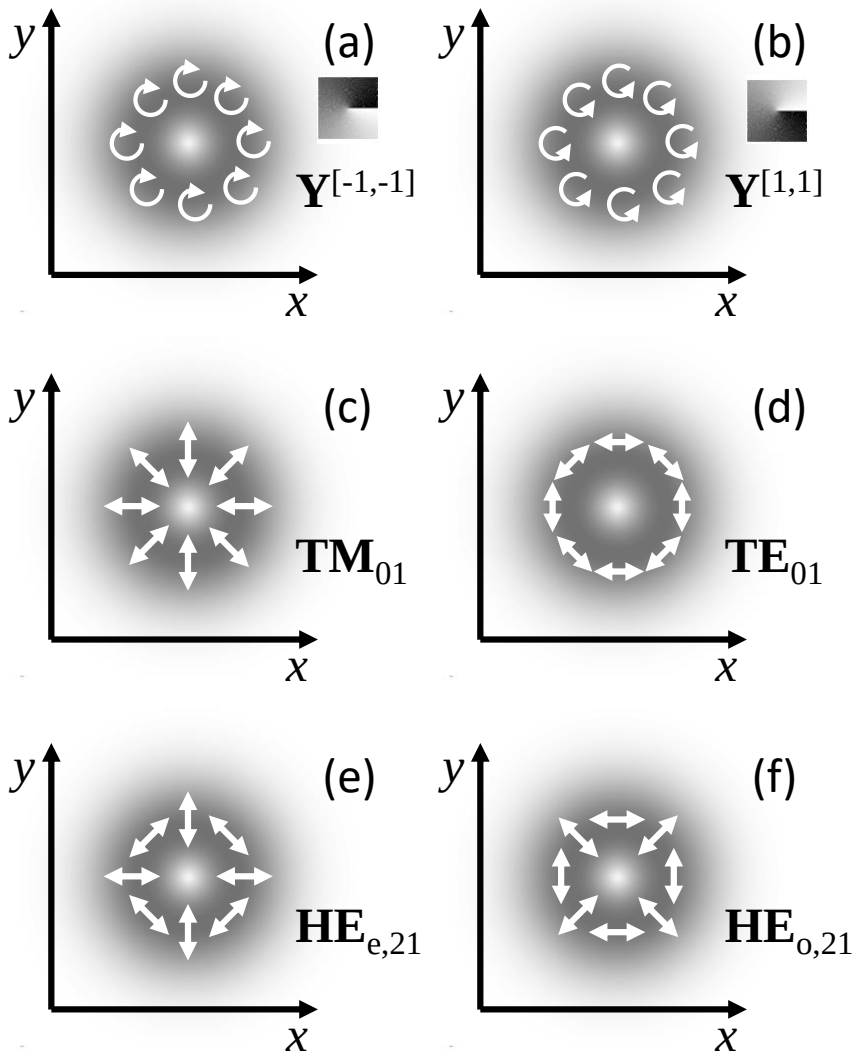

Figure 5.2: Examples of structured light modes. The greyscale images represent intensity, the arrows indicate the polarization, and phase is given by the greyscale inset for (a) and (b). All of the depicted modes have the same orbital angular momentum magnitude, $|l| = |L| = 1$. The top two plots are spin and orbital angular momentum eigenstate modes $\mathbf{Y}$ with (a) $s = -1$ and $l = -1$ and (b) $s = 1$ and $l = 1$. The bottom four plots are cylindrical vector modes: (c) the radial mode ($L = -1$, $S = 1$, TM_{01}), (d) the azimuthal mode ($L = 1$, $S = -1$, TE_{01}), and the hybrid modes (e) even ($L = 1$, $S = 1$, $\mathrm{HE}_{e,21}$) and (f) odd ($L = -1$, $S = -1$, $\mathrm{HE}_{o,21}$). The CV mode indices L and S are defined in Section 5.3.2.

$$\boldsymbol{\sigma}^{[s]} = \frac{(\mathbf{x} + is\mathbf{y})}{\sqrt{2}}, \tag{5.7}$$

with right or left circular modes yielding a spin projection along the system axis of $s\hbar$ ($s = \pm 1$). The OAM results from an azimuthal phase gradient of the field about the beam axis, $\exp(il\phi)$, and has a value of $l\hbar$ ($l = 0, \pm 1, \pm 2, \ldots$) projected along that same axis [91]. With these functions, I can define azimuthal modes $\mathbf{\Phi}^{[n]}(\phi)$ with definite values of SAM and OAM, which I call the Y modes:

$$\mathbf{Y}^{[l,s]} = e^{il\phi}\boldsymbol{\sigma}^{[s]}. \tag{5.8}$$

In free-space, these are paraxial vector solutions to the wave equation. In waveguides, these azimuthal modes are solutions only when $l \geq 2$ [88].

5.3.2 Cylindrical Vector Modes

The second type of azimuthal $\mathbf{\Phi}$ mode we define is the CV mode. These are the general set of solutions in cylindrically-symmetric weakly-guiding waveguides, valid for $l \geq 1$ (unlike the Y modes):

$$\mathbf{CV}^{[L,S]} = \frac{1}{\sqrt{2}}\left(\mathrm{e}^{i\frac{\pi}{4}(S-1)}\mathbf{Y}^{[L,S]} + \mathrm{e}^{-i\frac{\pi}{4}(S-1)}\mathbf{Y}^{[-L,-S]}\right), \tag{5.9}$$

where $S = \pm 1$ and $L = \pm|l|$ are mode indices, not angular momentum quantum

numbers. A unique feature of these modes is that in contrast to the $\mathbf{Y}^{[l,s]}$ modes, they are real at all transverse points. The CV modes cannot be factored into functions for OAM and SAM; hence, they are non-separable and are no longer eigenstates of OAM or spin (unlike the Y modes). Thus, it is not correct to refer to the OAM or spin of a particular CV mode. However, the magnitude of L is equal to the magnitude of l in each CV mode: $|L| = |l|$.

The CV modes can be divided into groups of four that have similar effective wavevector ($\beta^{[u]}$) values [88]. Each mode in a particular mode group has an identical radial mode function (index m) and identical OAM magnitude $|l|$:

$$\mathbf{CV}^{[L=-|l|,S=1]}(\phi) = \frac{1}{\sqrt{2}}\left(\mathbf{Y}^{[-|l|,1]} + \mathbf{Y}^{[|l|,-1]}\right); \\ (|l| = 1 : \mathrm{TM}_{0m}, |l| \geq 2 : \mathrm{EH}_{e,|l|-1,m}) \tag{5.10}$$

$$\mathbf{CV}^{[L=|l|,S=-1]}(\phi) = \frac{i}{\sqrt{2}}\left(\mathbf{Y}^{[-|l|,1]} - \mathbf{Y}^{[|l|,-1]}\right); \\ (|l| = 1 : \mathrm{TE}_{0m}, |l| \geq 2 : \mathrm{EH}_{o,|l|-1,m}) \tag{5.11}$$

$$\mathbf{CV}^{[L=|l|,S=1]}(\phi) = \frac{1}{\sqrt{2}}\left(\mathbf{Y}^{[-|l|,-1]} + \mathbf{Y}^{[|l|,1]}\right); \\ (|l| \geq 1 : \mathrm{HE}_{e,|l|+1,m}) \tag{5.12}$$

$$\mathbf{CV}^{[L=-|l|,S=-1]}(\phi) = \frac{-i}{\sqrt{2}}\left(\mathbf{Y}^{[-|l|,-1]} - \mathbf{Y}^{[|l|,1]}\right); \\ (|l| \geq 1 : \mathrm{HE}_{o,|l|+1,m}). \tag{5.13}$$

The parenthetical line below each mode CV mode definition indicates the standard waveguide mode label for each value of $|l|$. The four modes above have identical intensity distributions but different patterns of spatially varying polarization. The set of modes for $|L|= 1$ is shown in Figure 5.2. This ladder of modes is the mode solution to cylindrically-symmetric weakly-guiding waveguides.

To understand the relationship between the CV and Y modes as waveguide solutions, one must consider the degeneracy of the CV modes. Some of the CV modes within each group of four will be degenerate. That is, they will have equal $\beta^{[u]}$ within the weakly-guiding approximation in a waveguide. Notice that two of the four modes have OAM and SAM aligned in each term, whereas the other two modes have anti-aligned angular momenta. The aligned pair of modes in any weak-guiding cylindrically-symmetric fibres are degenerate, and likewise for the anti-aligned pair. Any superposition of two degenerate modes is also a solution and has the same $\beta^{[u]}$ as the original modes. Each Y mode is a superposition of two degenerate CV modes and, consequently, is a waveguide-mode solution. An exception is the case of $l = \pm 1$ where the anti-aligned degeneracy is broken [88, 92]; that is, $\mathbf{CV}^{[-1,1]}$ and $\mathbf{CV}^{[1,-1]}$ are not degenerate. In summary, whereas Y modes are solutions to cylindrically-symmetric weakly-guiding waveguides for $l \geq 2$, CV modes are solutions for $l \geq 1$.

5.3.3 Total Angular Momentum Modes

In addition to the CV and Y mode bases, it is useful to define an azimuthal $\mathbf{\Phi}$ mode basis for the TAM projected along the beam or fibre axis. In this basis, each mode has a definite TAM of $j = l + s$, which will allow the TAM conservation to be tracked in the FWM processes. Since the FWM process occurs along the full length of the medium, this is only useful to do if these j eigenstates are also eigenmodes of the fibre, so that j is conserved during propagation. Note that an eigenstate refers to a state with a definite value of an observable, whereas an eigenmode refers to a mode which remains unchanged upon propagation. If the j eigenstates were not eigenmodes of the fibre any change in j could be caused by either FWM or linear propagation and the role of FWM would be ambiguous.

For $|l| \geq 2$, these definite TAM modes will be the Y modes. Since the Y modes have definite s and l, the TAM (j) will be definite as well. However, as explained above, for the $|l| = 1$ subspace, only two of the four Y modes are waveguide mode solutions: $\mathbf{Y}^{[1,1]}$ and $\mathbf{Y}^{[-1,-1]}$. For the other two Y modes, s and l will not be preserved during propagation. In their place, we use the anti-aligned pair of $|l| = 1$ CV modes. Note that both superposition terms in each of these CV modes, equations 5.10 and 5.11, have the same value for TAM: $l + s = j = 0$. Consequently, the anti-aligned CV modes are definite TAM states with $j = 0$. The TAM mode basis comprises these four modes, which are listed in Table 5.1 and are represented by Z in subsequent theory.

Table 5.1: Total angular momentum (TAM) modes.

$l+s=j$	Mode
$-1+1=0$	$\mathbf{Z}^{[+0]}$=$\mathbf{CV}^{[L=-1,S=1]}$
$1-1=0$	$\mathbf{Z}^{[-0]}=\mathbf{CV}^{[L=1,S=-1]}$
$1+1=2$	$\mathbf{Z}^{[2]}=\mathbf{Y}^{[1,1]}$
$-1-1=-2$	$\mathbf{Z}^{[-2]}=\mathbf{Y}^{[-1,-1]}$

5.4 Four-Wave Mixing Theory for Spatial Modes

The starting point for this derivation of FWM theory is the wave equation for a third-order nonlinear process. Unlike most other published theory, I will retain the transverse vector and spatial variation of the fields. As will be discussed later, I assume an isotropic material (such as silica) and that the nonlinearity is frequency independent. This permits the use of a simplified form for the third-order nonlinear polarization. In these systems (such as optical fibre), the coupled amplitude equations for the four fields are dependent on the vector eigenmodes. In this treatment, the effects of self-phase modulation (SPM) and cross-phase modulation (XPM) arising from the pump beams are also included. However, SPM and XPM arising from the signal and idler are neglected on the basis that they are far less intense than the pumps. Similarly, pump depletion is also neglected.

5.4.1 Definition of Fields

FWM converts light from two pump (p, p') fields to two outgoing fields, typically called signal (s) and idler (i): $p + p' \rightarrow s + i$. In the non-degenerate case, all four fields can have distinct frequencies and spatial modes. I define the total electric field vector in the system as the sum of the four fields,

$$\mathbf{E}(\mathbf{r}, t) = \sum_{v=p,p',s,i} \mathbf{E}_v(\mathbf{r}, t). \tag{5.14}$$

Each field is identified by the subscript v and is assumed to be monochromatic and exist in just one of the spatial modes discussed in Section 5.3. Thus, this subscript will be taken to implicitly represent the field identity (p, p', i, or s), the field frequency ω_v, and the spatial mode $\mathbf{M}^{[u]}$. When considering an individual field $\mathbf{E}_v$, I can factor out the slowly-varying field amplitude $A_v(z)$:

$$\mathbf{E}_v(\mathbf{r}, t) = A_v(z)\mathbf{L}^{[u]}(\mathbf{r}, t) + \text{c.c.}. \tag{5.15}$$

With this definition, I can later isolate the slow change in field amplitude $A_v(z)$ due to the nonlinear interaction.

5.4.2 The Nonlinear Wave Equation

The wave equation with a nonlinear source term $\mathbf{P}_{\text{NL}}$ is

$$\nabla^2\mathbf{E}(\mathbf{r},t) + \frac{n(r)^2}{c^2}\frac{\partial^2\mathbf{E}(\mathbf{r},t)}{\partial t^2} = -\mu_0\frac{\partial^2\mathbf{P}_{\mathrm{NL}}(\mathbf{r},t)}{\partial t^2}, \tag{5.16}$$

where μ_o is the permeability of free space. When considering all four fields pertinent to FWM, the left-hand side (LHS) of equation 5.16 is

$$\mathrm{LHS[Eq.\ 5.16]} = \sum_{v=p,p',s,i}\left[\nabla^2\mathbf{E}_v(\mathbf{r},t) + \frac{n(r)^2}{c^2}\frac{\partial^2\mathbf{E}_v(\mathbf{r},t)}{\partial t^2}\right]. \tag{5.17}$$

Carrying out the spatial differentiation in the Laplacian, equation 5.17 becomes

$$\begin{aligned}\mathrm{LHS[Eq.\ 5.16]} = \sum_{v=p,p',s,i}\Bigg[&\left(\frac{\partial^2 A_v(z)}{\partial z^2} + 2i\beta_v\frac{\partial A_v(z)}{\partial z}\right)\mathbf{L}_v(\mathbf{r},t)\\ &+ A_v(z)\left(\nabla^2\mathbf{L}_v(\mathbf{r},t) + \frac{n(r)^2}{c^2}\frac{\partial^2\mathbf{L}_v(\mathbf{r},t)}{\partial t^2}\right) + \mathrm{c.c.}\Bigg].\end{aligned} \tag{5.18}$$

In the right-hand side (RHS) of the first line of this equation, I make use of the fact that $A(z)$ is assumed to be slowly varying compared to the fast oscillation associated with the effective wavevector $\beta^{[u]}$ along z. This is the slowly-varying amplitude approximation, $\partial^2 A(z)/\partial z^2 = 0$. The second line is just the source-free wave equation, equation 5.2, which is zero for the solution $\mathbf{L}_v(\mathbf{r},t)$. With these simplifications, equation 5.16 becomes

$$\sum_{v=p,p',s,i}\left[2i\beta_v\frac{\partial A_v(z)}{\partial z}\mathbf{L}_v(\mathbf{r},t) + \mathrm{c.c.}\right] = -\mu_0\frac{\partial^2\mathbf{P}_{\mathrm{NL}}(\mathbf{r},t)}{\partial t^2}. \tag{5.19}$$

5.4.3 Nonlinear Polarization and the Coupled Amplitude Equations

I now determine the form of the nonlinear polarization $\mathbf{P}_{\mathrm{NL}}(\mathbf{r}, t)$. By omitting contributions from other third-order processes (such as third harmonic generation, *etc.*) and considering only contributions $\mathbf{P}_v(\mathbf{r})$ created at the pump, signal, and idler frequencies, the nonlinear polarization can be expressed as the sum:

$$\mathbf{P}_{NL}(\mathbf{r}, t) = \sum_{v=p,p',s,i} \mathbf{P}_v(\mathbf{r}) \exp(-i\omega_v t) + \text{c.c.}. \tag{5.20}$$

With this definition, the RHS of the nonlinear wave equation (equation 5.19) becomes

$$\text{RHS[Eq. 5.19]} = \mu_0 \sum_{v=p,p',s,i} \omega_v^2 \mathbf{P}_v(\mathbf{r}) \exp(-i\omega_v t) + \text{c.c.}. \tag{5.21}$$

In materials that are both isotropic and satisfy the Kleinman symmetry condition (that is, that the frequency dependence of $\chi^{(3)}$ can be neglected [93]), the third-order nonlinear susceptibility tensor can be expressed in terms of a single scalar component: $\chi^{(3)}_{xxxx}$. The nonlinear polarization at the signal frequency is then given by

$$\begin{aligned}\mathbf{P}_s(\mathbf{r}) = 2\epsilon_0\chi^{(3)}_{xxxx}\Big[&\left(\mathbf{E}^*_p\cdot\mathbf{E}_p\right)\mathbf{E}_s + \left(\mathbf{E}^*_p\cdot\mathbf{E}_s\right)\mathbf{E}_p + \left(\mathbf{E}_p\cdot\mathbf{E}_s\right)\mathbf{E}^*_p \\ &+ \left(\mathbf{E}^*_{p'}\cdot\mathbf{E}_{p'}\right)\mathbf{E}_s + \left(\mathbf{E}^*_{p'}\cdot\mathbf{E}_s\right)\mathbf{E}_{p'} + \left(\mathbf{E}_{p'}\cdot\mathbf{E}_s\right)\mathbf{E}^*_{p'} \\ &+ \left(\mathbf{E}^*_i\cdot\mathbf{E}_p\right)\mathbf{E}_{p'} + \left(\mathbf{E}^*_i\cdot\mathbf{E}_{p'}\right)\mathbf{E}_p + \left(\mathbf{E}_p\cdot\mathbf{E}_{p'}\right)\mathbf{E}^*_i\Big],\end{aligned} \tag{5.22}$$

where I have omitted the spatial-temporal coordinates for clarity. The derivation of equation 5.22 is nontrivial and I have moved it to Section 5.7.1 to avoid interrupting the flow of this section. A similar equation can be found for the nonlinear polarization at the idler frequency ($\mathbf{P}_i$) by exchanging s and i throughout equation 5.22. The first six terms in equation 5.22 represent XPM from the pump beams, and the last three terms are the contribution to FWM.

With this simplified form for the nonlinear polarization, I can now evaluate its action on the signal and idler fields. The nonlinear wave equation (equation 5.19) for just the signal field is

$$2i\beta_s\frac{\partial A_s(z)}{\partial z}\mathbf{M}_s(r,\phi)\exp(i\beta_s z) = \frac{\omega_s^2\mathbf{P}_s(\mathbf{r})}{\epsilon_0 c^2}. \tag{5.23}$$

Substituting the nonlinear polarization of the signal frequency ($\mathbf{P}_s$) from equation 5.22 into equation 5.23 gives

$$\frac{\partial A_s(z)}{\partial z}\mathbf{M}_s = \frac{-i\omega_s^2\chi^{(3)}_{xxxx}}{2\beta_s c^2}\Bigg(A_p^* A_p A_s\Big[\boldsymbol{\alpha}_{pps} + \boldsymbol{\alpha}_{psp} + \bar{\boldsymbol{\alpha}}_{psp}\Big] + A_{p'}^* A_{p'} A_s\Big[\boldsymbol{\alpha}_{p'p's} + \boldsymbol{\alpha}_{p'sp'} + \bar{\boldsymbol{\alpha}}_{p'sp'}\Big] + A_p A_{p'} A_i^*\Big[\boldsymbol{\alpha}_{ipp'} + \boldsymbol{\alpha}_{ip'p} + \bar{\boldsymbol{\alpha}}_{pp'i}\Big]\exp(i\Delta k z)\Bigg), \tag{5.24}$$

where $\Delta k = \beta_p + \beta_{p'} - \beta_s - \beta_i$ is the phasematching term and where I introduce the quantities $\boldsymbol{\alpha}_{ijk} = (\mathbf{M}_i^* \cdot \mathbf{M}_j)\,\mathbf{M}_k$ and $\bar{\boldsymbol{\alpha}}_{ijk} = (\mathbf{M}_i \cdot \mathbf{M}_j)\,\mathbf{M}_k^*$. As before, I have suppressed the dependence of A and $\mathbf{M}$ on the spatial coordinates.

I now isolate the behaviour of the scalar amplitude $A_s(z)$ by taking the dot product with $\mathbf{M}_s^*$ on both sides and integrating over the transverse plane. With this, the orthonormality condition in equation 5.4 yields the coupled amplitude equation for the signal (or idler, by exchanging s and i throughout) as

$$\frac{\partial A_s(z)}{\partial z} = \kappa_s\Big(\Big|A_p\Big|^2 A_s\left[\mathsf{O}_{p^*ps^*s} + \mathsf{O}_{p^*ss^*p} + \mathsf{O}_{pss^*p^*}\right] + \Big|A_{p'}\Big|^2 A_s\left[\mathsf{O}_{(p')^*p's^*s} + \mathsf{O}_{(p')^*ss^*p'} + \mathsf{O}_{p'ss^*(p')^*}\right] + A_p A_{p'} A_i^*\left[\mathsf{O}_{i^*ps^*p'} + \mathsf{O}_{i^*p's^*p} + \mathsf{O}_{pp's^*i^*}\right]\exp(i\Delta k z)\Big). \tag{5.25}$$

Here, the field coupling constant κ_s is given by

$$\kappa_s = \frac{-i\omega_s^2 \chi^{(3)}_{xxxx}}{2\beta_s c^2}, \tag{5.26}$$

and the quantity

$$\mathsf{O}_{a^{(*)}b^{(*)}c^{(*)}d^{(*)}} = \int \left(\mathbf{M}_a^{(*)} \cdot \mathbf{M}_b^{(*)}\right) \left(\mathbf{M}_c^{(*)} \cdot \mathbf{M}_d^{(*)}\right) r\mathrm{d}r\mathrm{d}\phi \tag{5.27}$$

is the mode overlap integral. In $\mathsf{O}_{a^{(*)}b^{(*)}c^{(*)}d^{(*)}}$, the subscripts a, b, c, and d represent the fields (p, p', s, or i) and an optional conjugate on a particular subscript applies to the corresponding mode in the integral. For example, $\mathsf{O}_{i^*ps^*p'} = \int \left(\mathbf{M}_i^* \cdot \mathbf{M}_p\right) \left(\mathbf{M}_s^* \cdot \mathbf{M}_{p'}\right) r\mathrm{d}r\mathrm{d}\phi$. So far, I have not made use of the cylindrical symmetry of the modes. This coupled wave equation is applicable in bulk media and waveguides since it only assumes that $\mathbf{M}_v$ are paraxial modes that are approximate solutions to the wave equation. I have also derived the analogous coupled amplitude equation for the pump field (see Section 5.7.1).

I now describe the conditions necessary for efficient FWM. The first two lines on the RHS of equation 5.25 describe XPM in the signal field induced by the pumps. The last line describes FWM. FWM requires perfect energy conservation and is most efficient when photon momentum is also conserved. These two conditions, which constitute perfect phase-matching, are $\Delta\omega = \omega_p + \omega_{p'} - \omega_i - \omega_s = 0$ and $\Delta k = 0$, and can occur in many different FWM processes since β_v varies with mode and frequency [93]. For example, different CV-mode combinations in the FWM process will be phasematched at different sets of frequencies. Since the energy and momentum conservation for FWM will depend on the specific

medium and waveguide of interest, the details of phasematching will not be considered further in this thesis.

The selection rules of interest here arise from a third condition for FWM. Specifically, the sum of overlap integrals in the final line of the coupled wave equation (equation 5.25), the total process amplitude U given by

$$\mathsf{U} \equiv \mathsf{O}_{i^*ps^*p'} + \mathsf{O}_{i^*p's^*p} + \mathsf{O}_{pp's^*i^*}, \tag{5.28}$$

must be non-zero.

5.5 Four-wave Mixing Selection Rules

So far, the coupled wave equations have been derived for any transverse modes satisfying the linear wave equation. By introducing cylindrically symmetric modes into the FWM integrals listed in equation 5.28, I now restrict the analysis to systems with cylindrical symmetry, such as optical fibres. This will reveal a set of selection rules which describe the allowed transitions between the modes. I begin by using the mode definition in equation 5.1 to separate each of the overlap integrals comprising the process amplitude U into a product of radial and azimuthal parts such that

$$\mathsf{O}_{a^{(*)}b^{(*)}c^{(*)}d^{(*)}} = F \int \left(\mathbf{\Phi}_a^{(*)} \cdot \mathbf{\Phi}_b^{(*)}\right) \left(\mathbf{\Phi}_c^{(*)} \cdot \mathbf{\Phi}_d^{(*)}\right) \mathrm{d}\phi, \tag{5.29}$$

where $F \equiv \int R_i^* R_p R_s^* R_{p'} r \mathrm{d}r$ is the integral over the radial components of the modes. Since the radial mode function $R_v = R^{[m(v)]}(r)$ depends on the specific radial index profile $n(r)$ of a chosen waveguide, I will focus only on the selection rules associated with the azimuthal modes. These selection rules will then be general to any medium with cylindrical symmetry. Specifically, I assume that F is non-zero, as would be true if all four fields were in the same radial mode (the same value of m). Henceforth, a "mode" will refer to only its azimuthal component $\mathbf{\Phi}_v(\phi)$. Note that for the CV mode basis only, the potential complex conjugates of $\mathbf{\Phi}_v$ in the integral in equation 5.29 can be dismissed since at all transverse positions the modes are real.

The total FWM intensity is proportional to the sum U of three overlap integrals (equation 5.28), each of which is defined by equation 5.29. By considering the mode combinations for the four fields that lead to a non-zero U, one can find the corresponding sets of selection rules for a given set of mode indices. I will use the notation $\mathbf{M}_p + \mathbf{M}_{p'} \rightarrow \mathbf{M}_i + \mathbf{M}_s$ to express allowed FWM processes. All allowed processes can also occur in reverse, that is, with input and output modes interchanged. If a process occurs for a particular set of modes in fields p and p' then it will also occur with the modes interchanged between p and p', and likewise for the output fields s and i.

Of particular interest is the scenario where all of the involved modes are taken from the family of four modes (within the Y, CV, and Z mode types) that have equal $|l|$. With this limited set of four, there is a finite number of potential processes.

5.5.1 Y Mode Selection Rules

For the sake of completeness and for use later, I derive here the selection rules for fields in the Y angular momentum modes. These results agree with similar FWM studies involving linearly polarized modes [81]. After inserting combinations of Y modes into the three O overlap integrals inside U (equation 5.28), each integral factors into OAM and SAM components. For example,

$$\begin{aligned}\mathsf{O}_{i^*ps^*p'} &= F\left(\boldsymbol{\sigma}^{[-s_i]}\cdot\boldsymbol{\sigma}^{[s_p]}\right)\left(\boldsymbol{\sigma}^{[-s_s]}\cdot\boldsymbol{\sigma}^{[s_{p'}]}\right)\\ &\quad\times\int e^{i\left(-l_i+l_p-l_s+l_{p'}\right)\phi}\mathrm{d}\phi\\ &= F\delta_{s_i,s_p}\delta_{s_s,s_{p'}}\delta_{l_p+l_{p'},l_s+l_i}.\end{aligned}\tag{5.30}$$

The other two O integrals give similar results. Given that s is limited to the values ±1, the Kronecker deltas for the s values in the three integrals can be combined to arrive at the selection rules:

$$\begin{aligned}s_p+s_{p'}&=s_i+s_s,\\ l_p+l_{p'}&=l_i+l_s,\end{aligned}\tag{5.31}$$

where l_v and s_v are the mode indices of fields ($v = p$, p', s, or i). When these rules are satisfied, the process will occur with amplitude $\mathsf{U} = 4\pi F$. I provide a full derivation of

equation 5.31 in Section 5.7.2.1.

These two rules show that SAM and OAM are independently conserved; there is no coupling between the SAM and OAM degrees of freedom. Consequently, the Y modes obey the following selection rules:

$$\begin{aligned} \Delta s &= s_p + s_{p'} - s_s - s_i = 0, \\ \Delta l &= l_p + l_{p'} - l_s - l_i = 0, \\ \Delta j &= j_p + j_{p'} - j_s - j_i = 0. \end{aligned} \tag{5.32}$$

In other words, in the weakly-guiding or paraxial approximation, TAM cannot be conserved by converting between SAM and OAM. This is a result of the fact that, in these approximations, the fields are transverse everywhere.

5.5.2 CV Mode Selection Rules

I now consider the situation in which each of the four fields is in a CV mode. In this treatment, each mode can have any value of the OAM magnitude $|L|$. For each field v ($v = p$, p', s, or i) the mode indices will be labelled as L_v and S_v. Table 5.2 gives the four selection rules, the full derivation of which is in Section 5.7.2.2. More than one rule can hold true simultaneously, and in this case their associated process amplitudes should be summed to find the total process amplitude U. Since S is limited to ± 1, the selection rules disallow processes where there are an odd number of like S values among the four fields.

Table 5.2: Selection rules for CV modes in terms of their mode indices L and S, and their associated process amplitudes.

Rule	L	S	Process Amplitude, U
1	$L_p + L_{p'} = L_s + L_i$	$S_p + S_{p'} = S_s + S_i$	$2\pi F$
2	$L_p + L_{p'} = -(L_s + L_i)$	$S_p + S_{p'} = -(S_s + S_i)$	$-2\pi F$
3	$L_p - L_{p'} = L_s - L_i$	$S_p - S_{p'} = S_s - S_i$	$2\pi F$
4	$L_p - L_{p'} = -(L_s - L_i)$	$S_p - S_{p'} = -(S_s - S_i)$	$2\pi F$

To better understand these selection rules, I first consider only the restricted case where all of the modes have the same value of $|L|$. The following will use the mode labels for $|L|= 1$ (TM, TE, *etc.*), but the results generalize to any value of $|L|$. All 100 distinct possible processes for $|L|= 1$ are compiled and tabulated in Figure 5.3 (the definition of an "individual process" is discussed with the full derivation in Section 5.7.2.2). By the same reasoning as above for S, the selection rules for $|L|= 1$ (where L is limited to $\pm|L|$) disallow processes where there are an odd number of like L values among the four fields. This is confirmed in Figure 5.3, where all processes in which the modes do not appear in pairs has an amplitude of zero. For example, (TM + TE $\rightarrow$ TM + TE), ($\mathrm{HE}_e + \mathrm{HE}_e \rightarrow \mathrm{HE}_e + \mathrm{HE}_e$), and (TM + TM $\rightarrow \mathrm{HE}_o + \mathrm{HE}_o$) are allowed but (TM + TE $\rightarrow$ TM + TM) is not. This single selection rule is similar to the selection rules for linear polarized fields in an isotropic medium [82] that involve only two orthogonal polarizations. Examples of allowed CV mode conversions with total process amplitudes ranging from $2\pi F$ to $6\pi F$ are shown in Figure 5.4.

When considering higher-order CV modes beyond the $|L|= 1$ manifold, one can find

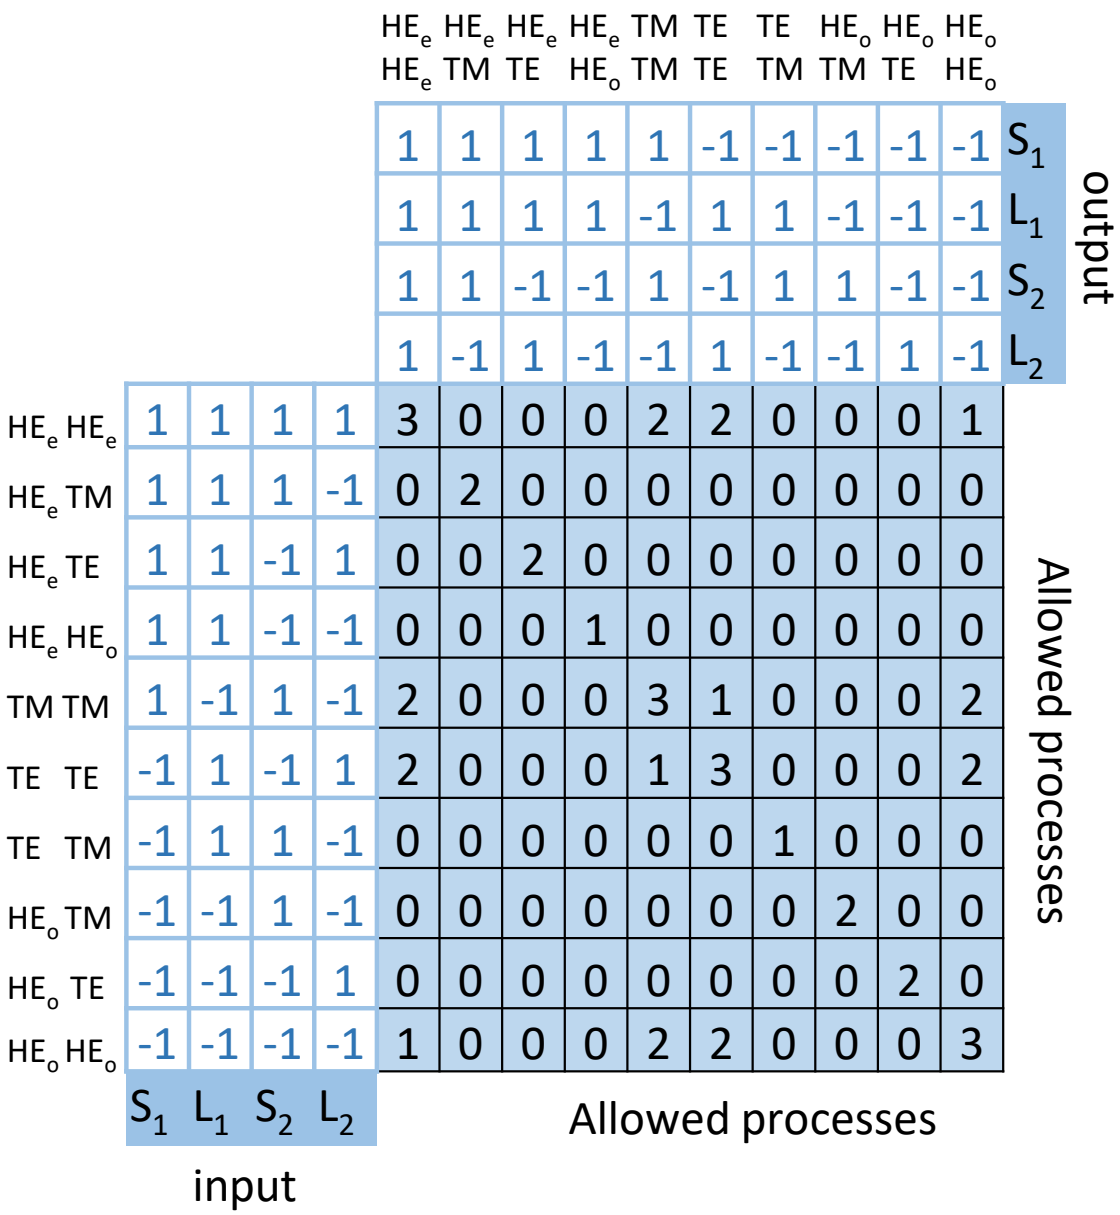

	S$_1$	L$_1$	S$_2$	L$_2$	HE$_e$HE$_e$	HE$_e$TM	HE$_e$TE	HE$_e$HE$_o$	TMTM	TETE	TETM	HE$_o$TM	HE$_o$TE	HE$_o$HE$_o$
HE$_e$HE$_e$	1	1	1	1	3	0	0	0	2	2	0	0	0	1
HE$_e$TM	1	1	1	-1	0	2	0	0	0	0	0	0	0	0
HE$_e$TE	1	1	-1	1	0	0	2	0	0	0	0	0	0	0
HE$_e$HE$_o$	1	1	-1	-1	0	0	0	1	0	0	0	0	0	0
TMTM	1	-1	1	-1	2	0	0	0	3	1	0	0	0	2
TETE	-1	1	-1	1	2	0	0	0	1	3	0	0	0	2
TETM	-1	1	1	-1	0	0	0	0	0	0	1	0	0	0
HE$_o$TM	-1	-1	1	-1	0	0	0	0	0	0	0	2	0	0
HE$_o$TE	-1	-1	-1	1	0	0	0	0	0	0	0	0	2	0
HE$_o$HE$_o$	-1	-1	-1	-1	1	0	0	0	2	2	0	0	0	3

Figure 5.3: Process amplitudes U for four-wave mixing between CV modes. All four modes have the same value of $|L|=|l|$. For simplicity, $|l|$ is taken to be 1 here, but the amplitudes for $|l|>1$ will be identical to these. The process amplitudes given here were calculated from equation 5.53 and are normalized as $\mathsf{U}/2\pi F$.

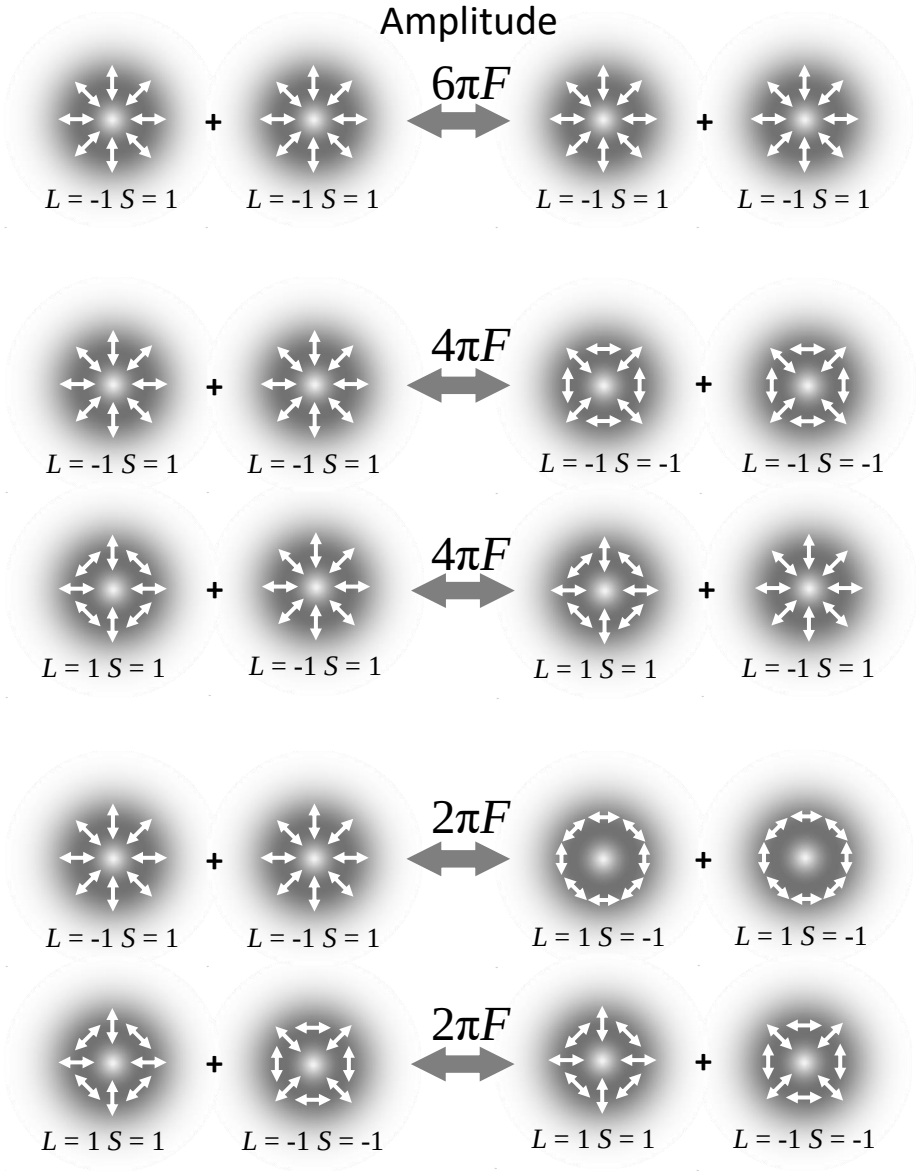

Figure 5.4: Examples of some of the allowed four-wave mixing processes between cylindrical vector modes in terms of their mode indices L and S, and their associated process amplitudes.

non-trivial allowed FWM processes such as ones with $L_p = L_{p'} = 2$, $L_s = 3$, and $L_i = 1$, among others. However, since the radial functions $R^{[m]}(r)$ typically implicitly depend on $|l|$, they will have different shapes with different magnitudes of L. This will decrease the radial overlap integral F and, thus, these processes will occur with a decreased efficiency.

5.5.3 TAM Mode Selection Rules

I now consider the situation in which each of the four fields is in a TAM mode. The set of TAM modes consists of the two $j = 0$ modes, $\mathbf{Z}^{[-0]} = \mathbf{CV}^{[-1,1]}$ and $\mathbf{Z}^{[+0]} = \mathbf{CV}^{[1,-1]}$ and the $j = \pm 2$ modes $\mathbf{Y}^{[1,1]}$ and $\mathbf{Y}^{[-1,-1]}$. I summarize the selection rules here (the full derivation of which is given in Section 5.7.2.3) and address their role in the conservation of TAM. Specifically, the question is whether the sum of the TAM of the two input photons is conserved in the FWM process; that is, does $j_i + j_s = j_p + j_{p'}$?

The TAM mode processes which involve only Y modes or only Z modes are already covered by the selection rules presented in the previous two sections. The selection rules for the Y modes conserve j since they separately conserve OAM and SAM (equation 5.32). However, the Z modes are not eigenstates of OAM or SAM and, thus, one cannot compare s and l values before and after the process; it is impossible to evaluate their conservation. Instead, the TAM j will be considered directly. A process containing only Z modes trivially conserves the two photons' TAM, given that each mode carries none. The remaining non-trivial processes which interconvert $j = 0$ and $j = \pm 2$ modes are summarized in Table 5.3.

Table 5.3: Selection rules for definite total angular momentum (TAM) modes, not including the trivial processes with only Z ($j = 0$) modes.

Process Type	Allowed Processes	Amplitude, $\mathsf{U}^{(\mathrm{TAM})}$
$\mathbf{Z}^{[w0]} + \mathbf{Z}^{[w0]} \rightarrow \mathbf{Y}^{[q,q]} + \mathbf{Y}^{[-q,-q]}$	$w = \pm, q = \pm 1$	$4\pi F$
$\mathbf{Y}^{[q,q]} + \mathbf{Y}^{[-q,-q]} \rightarrow \mathbf{Z}^{[w0]} + \mathbf{Z}^{[w0]}$	$w = \pm, q = \pm 1$	$4\pi F$
$\mathbf{Z}^{[w0]} + \mathbf{Y}^{[q,q]} \rightarrow \mathbf{Z}^{[w0]} + \mathbf{Y}^{[q,q]}$	$w = \pm, q = \pm 1$	$4\pi F$

The TAM conservation expression for the first allowed process type is $j_i + j_s - j_p - j_{p'} = \mp 2 \pm 2 - 0 - 0 = 0$. The second process type is the reverse of the first, and conserves TAM in a similar manner. The last process type gives $j_i + j_s - j_p - j_{p'} = 0 \pm 2 - 0 \mp 2 = 0$. Thus, all of the TAM mode interconversion processes conserve the TAM of the two photons.

One might ask whether conservation of TAM could be used as the sole selection rule – the answer is no. There are many processes that would conserve TAM but are not permitted by the other selection rules, such as $\mathbf{Z}^{[+0]} + \mathbf{Y}^{[\pm 1,\pm 1]} \rightarrow \mathbf{Z}^{[-0]} + \mathbf{Y}^{[\pm 1,\pm 1]}$. Allowed processes can occur with different amplitudes of $\mathsf{U} = 2\pi F$, $4\pi F$, and $6\pi F$. Note that the processes with amplitudes of $2\pi F$ and $6\pi F$ are the trivial cases involving only Z modes; these are covered later in Section 5.7.2.3.

5.5.4 Prospects for the Generation of Mode-entangled Photon Pairs

In spontaneous FWM (SFWM), vacuum-state fields in any chosen pair of input signal and idler modes (since all the possible input modes contain vacuum) undergo FWM with a

strong pump field. The mixing will occasionally produce a pair of photons, one in the output idler field and one in the output signal field. This will occur with the same relative amplitude as the corresponding FWM process. Consequently, for any given two input pump modes, the FWM selection rules identify signal and idler output mode combinations that SFWM will produce photon pairs in. That is, it identifies allowed SFWM processes.

Previous work has generated photon pairs that are entangled in their OAM modes [94]. In those cases, entanglement occurred naturally through the conservation of OAM. More generally, to generate entanglement through a spontaneous nonlinear process such as downconversion or FWM, one requires two simultaneous processes that each produce a different pair of signal and idler modes. In order to generate entanglement, the modes containing the signal and idler photons must be indistinguishable other than in the degree of freedom that is entangled [95]. For FWM in a fibre, this entails phasematching the two processes with identical signal wavelengths and identical idler wavelengths. Since each mode has a different effective wavevector β with different dependence on wavelength, Δk will typically vary differently with wavelength for the two processes. Consequently, such phasematching is nontrivial. Fortunately, control of the effective index β can be achieved by careful design of the fibre index profile and choice of material [96].

In order for these two processes together to create mode entanglement, they must produce different pairs of modes, $\mathbf{M}_i$ and $\mathbf{M}_s$, from one another. An example of two appropriate processes is $\left(\mathbf{Y}_p^{[1,1]} + \mathbf{Y}_{p'}^{[-1,-1]} \rightarrow \mathbf{Z}_i^{[+0]} + \mathbf{Z}_s^{[+0]}\right)$ and $\left(\mathbf{Y}_p^{[1,1]} + \mathbf{Y}_{p'}^{[-1,-1]} \rightarrow \mathbf{Z}_i^{[-0]} + \mathbf{Z}_s^{[-0]}\right)$. The entangled state of the two spontaneously generated photons would be $\mathbf{Z}_i^{[+0]}\mathbf{Z}_s^{[+0]} + \mathbf{Z}_i^{[-0]}\mathbf{Z}_s^{[-0]}$. The availability of many allowed FWM processes for cylindrical vector modes

means that there are many possibilities for the generation of photons entangled in their spatial-polarization modes.

5.6 Conclusion

In this chapter, I presented a theoretical investigation into the nonlinear optics of structured light. Specifically, I and the other co-authors of the paper this chapter was based on [1] have developed a theory describing FWM of CV modes in optical fibre or free space and the mode selection rules that follow. For comparison, I have also reviewed the analogous theory for modes with definite SAM and OAM. Given the cylindrical symmetry of a fibre and free space, one might expect that the only quantity that must be strictly conserved in FWM is the TAM along the system axis. Indeed, in some processes the SAM and OAM are independently conserved and, thus, so too will be the TAM. However, for OAM magnitude $|l|= 1$, the fundamental eigenmodes of an optical fibre are not eigenstates of either SAM or OAM and, consequently, these properties will change during propagation. Nonetheless, I demonstrated that for modes that are TAM eigenstates, the sum of two photons' TAM is conserved and emerges in the output photons. Future work on this project could investigate the link between FWM selection rules and angular momentum conservation laws outside of the paraxial and weakly-guiding regime.

Since CV modes are the eigenmodes of weakly-guiding fibres and of free-space propagation, the selection rules presented here are pertinent to a range of applications. These include telecommunications, where structured modes can increase bandwidth [73, 74]. Ad-

ditionally, these results could find use in all-optical switching by utilizing the optical nonlinearity of structured spatial modes as a way to route such beams. In quantum optics, CV modes are currently being investigated for use in quantum cryptography [97]. Using these results, the CV mode photons could be generated directly by FWM inside optical fibres. I have also demonstrated how to generate photon pairs entangled in their CV modes. Beyond fibre optics, these results have implications in free-space FWM, and could shed light on the debate over the roles of SAM and OAM in photons [98, 99].

5.7 Derivations

This section contains the full derivations of some of the equations used earlier in this chapter, but which were too lengthy to include above without interrupting the flow of the sections in which they appeared.

5.7.1 Derivation of the Nonlinear Polarization

Here I will derive the form of the third-order nonlinear polarization $\mathbf{P}$ for an isotropic material where the frequency dependence of the susceptibility can be neglected (the Kleinman symmetry condition). In general, each component P_i ($i = x$, y, or z) of the $\chi^{(3)}$ nonlinear polarization for a given process is [93]

$$P_i = \epsilon_0 D \sum_{jkl=x,y,z} \chi^{(3)}_{ijkl} E_j(\omega_o) E_k(\omega_n) E_l(\omega_m), \tag{5.33}$$

where the degeneracy factor D is equal to the number of distinct permutations of the fields. Except where otherwise noted, all subscripts in this section refer to the Cartesian coordinates (x, y, z), unlike in previous sections where subscripts p, p', s, and i are used to identify the field.

The non-zero components of the $\chi^{(3)}$ tensor are

$$\begin{aligned}
\chi^{(3)}_{xxyy} &= \chi^{(3)}_{xxzz} = \chi^{(3)}_{yyxx} = \chi^{(3)}_{yyzz} = \chi^{(3)}_{zzxx} = \chi^{(3)}_{zzyy} \\
&= \chi^{(3)}_{xyxy} = \chi^{(3)}_{xzxz} = \chi^{(3)}_{yzyz} = \chi^{(3)}_{yxyx} = \chi^{(3)}_{zxzx} \\
&= \chi^{(3)}_{zyzy} = \chi^{(3)}_{xyyx} = \chi^{(3)}_{xzzx} = \chi^{(3)}_{yxxy} = \chi^{(3)}_{yzzy} \\
&= \chi^{(3)}_{zxxz} = \chi^{(3)}_{zyyz} = \frac{1}{3}\chi^{(3)}_{xxxx} = \frac{1}{3}\chi^{(3)}_{yyyy} = \frac{1}{3}\chi^{(3)}_{zzzz}.
\end{aligned} \tag{5.34}$$

The 21 non-zero $\chi^{(3)}$ components in equation 5.34 contribute seven terms to each of the Cartesian components of the nonlinear polarization in equation 5.33. Each nonlinear polarization component P_i now contains only seven non-zero terms from the general 27-term sum in equation 5.33. The last line of equation 5.34 shows that each $jkl = iii$ term is a factor of three larger than the other terms. In order to make evident an upcoming vector product I will expand each of these larger iii terms as three separate (and equal) terms for a total of nine non-zero terms all together. The nonlinear polarization is then

$$
\begin{aligned}
P_i = \frac{\epsilon_0 D \chi^{(3)}_{xxxx}}{3} \Bigg[& \Big(E_i(\omega_o)E_i(\omega_n) + E_j(\omega_o)E_j(\omega_n) + E_k(\omega_o)E_k(\omega_n) \Big) E_i(\omega_m) \\
& + \Big(E_i(\omega_n)E_i(\omega_m) + E_j(\omega_n)E_j(\omega_m) + E_k(\omega_n)E_k(\omega_m) \Big) E_i(\omega_o) \\
& + \Big(E_i(\omega_o)E_i(\omega_m) + E_j(\omega_o)E_j(\omega_m) + E_k(\omega_o)E_k(\omega_m) \Big) E_i(\omega_n) \Bigg] \\
= & \frac{\epsilon_0 D \chi_{xxxx}}{3} \Bigg[\Big(\mathbf{E}(\omega_o) \cdot \mathbf{E}(\omega_n) \Big) E_i(\omega_m) \\
& + \Big(\mathbf{E}(\omega_m) \cdot \mathbf{E}(\omega_n) \Big) E_i(\omega_o) \\
& + \Big(\mathbf{E}(\omega_m) \cdot \mathbf{E}(\omega_o) \Big) E_i(\omega_n) \Bigg],
\end{aligned} \tag{5.35}
$$

where i, j, and k are any permutation of $\{x, y, z\}$. Combining the three Cartesian components of P_i gives the following vector form:

$$
\begin{aligned}
\mathbf{P} = \frac{\epsilon_0 D \chi^{(3)}_{xxxx}}{3} \Bigg[& \Big(\mathbf{E}(\omega_o) \cdot \mathbf{E}(\omega_n) \Big) \mathbf{E}(\omega_m) \\
& + \Big(\mathbf{E}(\omega_m) \cdot \mathbf{E}(\omega_n) \Big) \mathbf{E}(\omega_o) + \Big(\mathbf{E}(\omega_m) \cdot \mathbf{E}(\omega_o) \Big) \mathbf{E}(\omega_n) \Bigg].
\end{aligned} \tag{5.36}
$$

The nonlinear polarization in equation 5.22 follows from applying this to FWM and XPM (conjugating fields with negative frequency) and evaluating the degeneracy factor: $D = 6$ for three distinct fields and $D = 3$ for two distinct fields.

For completeness, the nonlinear polarization for the pump field p (or p', by exchanging

p and p' throughout) is:

$$\mathbf{P}_p(\mathbf{r}) = \epsilon_0 \chi^{(3)}_{xxxx} \Big[(\mathbf{E}_p \cdot \mathbf{E}_p)\, \mathbf{E}_p^* + 2\left(\mathbf{E}_p^* \cdot \mathbf{E}_p\right) \mathbf{E}_p + 2\left(\mathbf{E}_{p'}^* \cdot \mathbf{E}_{p'}\right) \mathbf{E}_p + 2\,(\mathbf{E}_{p'} \cdot \mathbf{E}_p)\, \mathbf{E}_{p'}^* + 2\left(\mathbf{E}_{p'}^* \cdot \mathbf{E}_p\right) \mathbf{E}_{p'} \Big]. \tag{5.37}$$

Here, the subscripts on the electric fields and polarization indicate the field identity, p or p'. By the same procedure used in Section 5.4.3, the coupled equations for the pump amplitude p (and p', by exchanging p and p' throughout) are found to be

$$\frac{\partial A_p(z)}{\partial z} = \frac{1}{2}\kappa_p \Big(\Big|A_p\Big|^2 A_p \left[\mathsf{O}_{ppp^*p^*} + 2\mathsf{O}_{p^*pp^*p}\right] + 2\Big|A_{p'}\Big|^2 A_p \left[\mathsf{O}_{(p')^*p'p^*p} + \mathsf{O}_{p'pp^*(p')^*} + \mathsf{O}_{(p')^*pp^*p'}\right] \Big). \tag{5.38}$$

5.7.2 Derivation of the FWM Process Amplitudes

Here I outline how to evaluate the overlap integral O defined in equation 5.29 for each of the mode bases introduced in Section 5.3. The sum of three O integrals gives the FWM process amplitude U as in equation 5.28. The mode indices for each field v ($v = p$, p', s, or i) will be labelled with the corresponding subscript – for example: s_p, L_i, or S_s. As explained in Section 5.4.3, a complex conjugate on the field v in the subscript of O indicates that the corresponding mode should be conjugated. In all three O integrals, two of the

indices are conjugated.

5.7.2.1 Process Amplitudes for the Definite Spin and Orbital Angular Momentum Modes

I start by considering FWM processes in which all four beams are in the SAM and OAM eigenstate modes, $\mathbf{Y}^{[l,s]} = e^{il\phi}\boldsymbol{\sigma}^{[s]}$. The terms comprising the FWM process amplitude are formed by the product of four modes, in this case these are the Y modes. I will first evaluate half of this product, the dot product of two general Y modes. For fields a and b this is

$$\boldsymbol{Y}_a^{(*)} \cdot \boldsymbol{Y}_b^{(*)} = \boldsymbol{\sigma}^{[\pm s_a]} \cdot \boldsymbol{\sigma}^{[\pm s_b]} \mathrm{e}^{i(\pm l_a \pm l_b)\phi}. \tag{5.39}$$

Here, the superscript $(*)$ indicates the optional presence of a complex conjugate on the Y mode, in which case each mode index of the v field (s_v, l_v) is preceded by the bottom symbol of $\pm$ or $\mp$. It follows from the definition of $\boldsymbol{\sigma}^{[s]}$ (equation 5.7) that, with $s = \pm 1$, $\boldsymbol{\sigma}^{[s_i]} \cdot \boldsymbol{\sigma}^{[s_j]} = \delta_{s_i,-s_j}$. The full argument of one overlap integral is then

$$\begin{aligned} \mathsf{O}_{a^{(*)}b^{(*)}c^{(*)}d^{(*)}} &= F \int \left(\boldsymbol{Y}_a^{(*)} \cdot \boldsymbol{Y}_b^{(*)}\right)\left(\boldsymbol{Y}_c^{(*)} \cdot \boldsymbol{Y}_d^{(*)}\right) \mathrm{d}\phi \\ &= \delta_{\pm s_a,\mp s_b}\delta_{\pm s_c,\mp s_d} F \int \mathrm{e}^{i(\pm l_a \pm l_b \pm l_c \pm l_d)\phi} \mathrm{d}\phi \\ &= 2\pi F \delta_{\pm s_a,\mp s_b}\delta_{\pm s_c,\mp s_d}\delta_{0,(\pm l_a \pm l_b \pm l_c \pm l_d)}, \end{aligned} \tag{5.40}$$

where the last line uses $\int_0^{2\pi} \exp(iq\phi)\mathrm{d}\phi = 2\pi\delta_{0,q}$. With this result, the three O integrals composing $\mathsf{U}^{(\mathrm{Y})}$ are

$$\mathsf{O}_{i^*ps^*p'} = 2\pi F \delta_{s_i,s_p}\delta_{s_s,s_{p'}}\delta_{l_p+l_{p'},l_i+l_s} \tag{5.41}$$

$$\mathsf{O}_{i^*p's^*p} = 2\pi F \delta_{s_i,s_{p'}}\delta_{s_s,s_p}\delta_{l_p+l_{p'},l_i+l_s} \tag{5.42}$$

$$\mathsf{O}_{pp's^*i^*} = 2\pi F \delta_{s_p,-s_{p'}}\delta_{-s_s,s_i}\delta_{l_p+l_{p'},l_i+l_s}. \tag{5.43}$$

The total process amplitude is then the summation of equations 5.41–5.43:

$$\begin{aligned}\mathsf{U}^{(\mathrm{Y})} &= 2\pi F \left(\delta_{s_i,s_p}\delta_{s_s,s_{p'}} + \delta_{s_i,s_{p'}}\delta_{s_s,s_p} + \delta_{s_p,-s_{p'}}\delta_{-s_s,s_i}\right)\delta_{l_p+l_{p'},l_i+l_s} \\ &= 4\pi F \delta_{s_i+s_s,s_p+s_{p'}}\delta_{l_i+l_s,l_p+l_{p'}},\end{aligned} \tag{5.44}$$

where the last line makes use of an identity presented below in equation 5.52. Each Kronecker delta provides a selection rule: $s_i + s_s = s_p + s_{p'}$ and $l_i + l_s = l_p + l_{p'}$, or more succinctly, $\Delta s = 0$ and $\Delta l = 0$. In summary, the SAM and OAM are independently conserved in this FWM process.

5.7.2.2 Process Amplitudes for the CV Modes

I now derive the process amplitude $\mathsf{U}^{(\mathrm{CV})}$ for a FWM process where all the beams are in CV modes. A CV mode consists of a superposition of two Y modes, as in equation 5.9. Since these modes are real, the complex conjugates in O are dropped. The mode overlap integral

(equation 5.29), contains two inner products of two CV modes each. When expanded using equation 5.9, O contains 16 terms, each containing a $\mathbf{Y}$ mode factor from all four fields ($v = a,\ b,\ c,\ d$). From equation 5.9, each of these four Y mode factors will either be of the form $\mathbf{Y}^{[+L_v,+S_v]}$ or $\mathbf{Y}^{[-L_v,-S_v]}$. Since the signs of S_v and L_v are equal for the Y modes appearing here, I introduce a new index $n_v = \pm 1$ which is the sign appearing in front of L_v and S_v in the Y mode label. Using this index, the these 16 terms can be represented as the following sum:

$$
\begin{aligned}
\mathsf{O}_{abcd} &= F \int \left(\mathbf{CV}^{[L_a,S_a]} \cdot \mathbf{CV}^{[L_b,S_b]}\right)\left(\mathbf{CV}^{[L_c,S_c]} \cdot \mathbf{CV}^{[L_d,S_d]}\right) \mathrm{d}\phi \\
&= \frac{F}{4} \sum_{n_a..n_d=\pm 1} \mathrm{e}^{i\frac{\pi}{4}(n_aS_a+n_bS_b+n_cS_c+n_dS_d-K)} \\
&\qquad \times \int \left(\mathbf{Y}^{[n_aL_a,n_aS_a]} \cdot \mathbf{Y}^{[n_bL_b,n_bS_b]}\right)\left(\mathbf{Y}^{[n_cL_c,n_cS_c]} \cdot \mathbf{Y}^{[n_dL_d,n_dS_d]}\right) \mathrm{d}\phi \\
&= \frac{\pi F}{2} \sum_{n_a..n_d=\pm 1} \mathrm{e}^{i\frac{\pi}{4}(n_aS_a+n_bS_b+n_cS_c+n_dS_d-K)} \\
&\qquad \times \delta_{n_aS_a,-n_bS_b}\delta_{n_cS_c,-n_dS_d}\delta_{0,n_aL+n_bL_b+n_cL_c+n_dL_d} \\
&= \frac{\pi F}{2} \sum_{n_a..n_d=\pm 1} \mathrm{e}^{-i\frac{\pi}{4}K}\delta_{n_aS_a,-n_bS_b}\delta_{n_cS_c,-n_dS_d} \\
&\qquad \times \delta_{0,n_aL+n_bL_b+n_cL_c+n_dL_d},
\end{aligned}
\tag{5.45}
$$

where the azimuthal integral was evaluated using the results from the last section. Each term differs by a constant K factor in the exponential: $K = n_a + n_b + n_c + n_d$. In the final expression, I have used the S Kronecker deltas to simplify the exponential factor.

This expression can be further simplified by pairing each of the 16 terms with its complex conjugate, where a term with a particular set of (n_a, n_b, n_c, n_d) values is conjugate to the ($-n_a$, $-n_b$, $-n_c$, $-n_d$) term. Combining each conjugate pair, the general form for each of the eight resulting terms is:

$$G_{n_a n_b n_c n_d} = \pi F \cos\left(\frac{\pi}{4}K\right) \delta_{n_a S_a, -n_b S_b} \delta_{n_c S_c, -n_d S_d} \times \delta_{0, n_a L + n_b L_b + n_c L_c + n_d L_d}, \tag{5.46}$$

where the notation for the $n_v = \pm 1$ in the subscripts has been simplified to just $n_v = \pm$. With this definition, $\mathsf{O}_{abcd} = G_{++++} + G_{+++-} + G_{++-+} + G_{++--} + G_{+-++} + G_{+-+-} + G_{+--+} + G_{-+++}$. Each of the eight terms has either zero, one, or two n_v values of -1; the other possible combinations of (n_a, n_b, n_c, n_d) were represented in the eight complex conjugate terms. Consequently, the value of K is 4, 2, and 0 for these cases, respectively. The four terms with $K = 2$ can be eliminated since, for this value of K, the cosine factor is zero. This leaves four terms: $\mathsf{O}_{abcd} = G_{++++} + G_{++--} + G_{+-+-} + G_{+--+}$. If $K = 0$ then the cosine factor equals 1 and if $K = 4$ it equals -1. Applying this to each term gives

$$\begin{aligned}\mathsf{O}_{abcd} = \pi F \Bigg[& - \delta_{S_a,-S_b}\delta_{S_c,-S_d}\delta_{0,L_a+L_b+L_c+L_d} \\ & + \delta_{S_a,-S_b}\delta_{S_c,-S_d}\delta_{0,L_a+L_b-L_c-L_d} \\ & + \delta_{S_a,S_b}\delta_{S_c,S_d}\delta_{0,L_a-L_b+L_c-L_d} \\ & + \delta_{S_a,S_b}\delta_{S_c,S_d}\delta_{0,L_a-L_b-L_c+L_d} \Bigg],\end{aligned} \tag{5.47}$$

where the identity $\delta_{n_jS_j,n_kS_k} = \delta_{-n_jS_j,-n_kS_k}$ was used to ensure that the first argument in all the S Kronecker deltas is positive.

With this, the three overlap integrals that compose $\mathsf{U}^{(\mathrm{CV})}$ can be written as follows:

$$\begin{aligned}\mathsf{O}_{i^*ps^*p'} = \pi F \Bigg[& - \delta_{S_i,-S_p}\delta_{S_s,-S_{p'}}\delta_{0,L_i+L_p+L_s+L_{p'}} \\ & + \delta_{S_i,-S_p}\delta_{S_s,-S_{p'}}\delta_{0,L_i+L_p-L_s-L_{p'}} \\ & + \delta_{S_i,S_p}\delta_{S_s,S_{p'}}\delta_{0,L_i-L_p+L_s-L_{p'}} \\ & + \delta_{S_i,S_p}\delta_{S_s,S_{p'}}\delta_{0,L_i-L_p-L_s+L_{p'}} \Bigg],\end{aligned} \tag{5.48}$$

$$\begin{aligned}\mathsf{O}_{i^*p's^*p} = \pi F \Big[&- \delta_{S_i,-S_{p'}}\delta_{S_s,-S_p}\delta_{0,L_i+L_{p'}+L_s+L_p} \\ &+ \delta_{S_i,-S_{p'}}\delta_{S_s,-S_p}\delta_{0,L_i+L_{p'}-L_s-L_p} \\ &+ \delta_{S_i,S_{p'}}\delta_{S_s,S_p}\delta_{0,L_i-L_{p'}+L_s-L_p} \\ &+ \delta_{S_i,S_{p'}}\delta_{S_s,S_p}\delta_{0,L_i-L_{p'}-L_s+L_p} \Big],\end{aligned} \tag{5.49}$$

$$\begin{aligned}\mathsf{O}_{pp's^*i^*} = \pi F \Big[&- \delta_{S_p,-S_{p'}}\delta_{S_s,-S_i}\delta_{0,L_p+L_{p'}+L_s+L_i} \\ &+ \delta_{S_p,-S_{p'}}\delta_{S_s,-S_i}\delta_{0,L_p+L_{p'}-L_s-L_i} \\ &+ \delta_{S_p,S_{p'}}\delta_{S_s,S_i}\delta_{0,L_p-L_{p'}+L_s-L_i} \\ &+ \delta_{S_p,S_{p'}}\delta_{S_s,S_i}\delta_{0,L_p-L_{p'}-L_s+L_i} \Big].\end{aligned} \tag{5.50}$$

Note that the same four L Kronecker deltas appear in each O, albeit paired with different S deltas in each. In the $\mathsf{U}^{(\mathrm{CV})}$ sum, I now group these L Kronecker delta terms together:

$$
\begin{aligned}
\mathsf{U}^{(\mathrm{CV})} &= \mathsf{O}_{i^*ps^*p'} + \mathsf{O}_{i^*p's^*p} + \mathsf{O}_{pp's^*i^*} \\
&= \pi F \Bigg[- \left(\delta_{S_i,-S_p} \delta_{S_s,-S_{p'}} + \delta_{S_i,-S_{p'}} \delta_{S_s,-S_p} + \delta_{S_p,-S_{p'}} \delta_{S_s,-S_i} \right) \delta_{0,L_i+L_p+L_s+L_{p'}} \\
&\quad + \left(\delta_{S_i,-S_p} \delta_{S_s,-S_{p'}} + \delta_{S_i,S_{p'}} \delta_{S_s,S_p} + \delta_{S_p,S_{p'}} \delta_{S_s,S_i} \right) \delta_{0,L_i+L_p-L_s-L_{p'}} \\
&\quad + \left(\delta_{S_i,S_p} \delta_{S_s,S_{p'}} + \delta_{S_i,S_{p'}} \delta_{S_s,S_p} + \delta_{S_p,-S_{p'}} \delta_{S_s,-S_i} \right) \delta_{0,L_i-L_p+L_s-L_{p'}} \\
&\quad + \left(\delta_{S_i,S_p} \delta_{S_s,S_{p'}} + \delta_{S_i,-S_{p'}} \delta_{S_s,-S_p} + \delta_{S_p,S_{p'}} \delta_{S_s,S_i} \right) \delta_{0,L_i-L_p-L_s+L_{p'}} \Bigg].
\end{aligned} \tag{5.51}
$$

The following identity can be used to considerably simplify the expression for $\mathsf{U}^{(\mathrm{CV})}$:

$$
\begin{aligned}
\delta_{S_a,q_1S_c} \delta_{S_b,q_1S_d} + \delta_{S_a,q_2S_d} \delta_{S_b,q_2S_c} + \delta_{S_a,-q_1q_2S_b} \delta_{S_c,-q_1q_2S_d} \\
= 2\delta_{S_b+q_1q_2S_a,q_2S_c+q_1S_d},
\end{aligned} \tag{5.52}
$$

where $q_k = \pm 1$. This identity's validity can be understood by considering the three possible values that can be taken by each of the arguments of the Kronecker delta on the RHS: -2, 0, or 2. For both delta arguments to be 2 or both arguments to be -2, all four terms in both arguments must be equal. The first two deltas on the LHS sum to 2 if all the terms in the arguments on the RHS are equal, thereby covering these cases. The third delta and one of the first two deltas on the LHS cover the remaining case, in which both RHS delta arguments are 0.

Applying this identity to $\mathsf{U}^{(\mathrm{CV})}$ four times and reordering the arguments of the Kro-

Table 5.4: Selection rules for the CV modes. Either the top or bottom symbol of $\pm$ and $\mp$ should be taken consistently throughout each rule.

Rule	L	S	Amplitude
1&2	$L_p + L_{p'} = \pm(L_s + L_i)$	$S_p + S_{p'} = \pm(S_s + S_i)$	$\pm 2\pi F$
3&4	$L_p - L_{p'} = \pm(L_s - L_i)$	$S_p - S_{p'} = \pm(S_s - S_i)$	$2\pi F$

necker deltas gives

$$\begin{aligned}\mathsf{U}^{(\mathrm{CV})} = 2\pi F\Big[&- \delta_{S_s+S_i,-(S_p+S_{p'})}\delta_{L_s+L_i,-(L_p+L_{p'})} + \delta_{S_s-S_i,S_p-S_{p'}}\delta_{L_s-L_i,L_p-L_{p'}} \\ &+ \delta_{S_s+S_i,S_p+S_{p'}}\delta_{L_s+L_i,L_p+L_{p'}} + \delta_{S_s-S_i,-(S_p-S_{p'})}\delta_{L_s-L_i,-(L_p-L_{p'})}\Big]. \end{aligned} \tag{5.53}$$

These four terms correspond to four selection rules, each of which has an L sub-rule and S sub-rule. These are summarized in Table 5.4. More than one rule can hold true simultaneously and in this case, their amplitudes should be added to find U, the total process amplitude.

In Figure 5.3 in Section 5.5, the amplitudes are compiled (normalized by $2\pi F$) for all distinguishable processes involving CV modes with the same value of $|L| = |l|$ (in this case, $|L| = 1$, but the results are general to any value of $|L|$). The labels p and p' are interchangeable without changing the physical process, and the same is true for the output fields s and i. In accordance with this, the input and output mode indices are labelled by the subscripts 1 and 2. This results in ten physically distinct mode combinations for

Table 5.5: Summary of allowed FWM processes between TAM modes and their associated amplitudes.

Case	Process Type	Allowed Processes	Amplitude, $\mathsf{U}^{(\mathrm{TAM})}$
1	$\mathbf{Z}^{[w0]} + \mathbf{Z}^{[w0]} \to \mathbf{Z}^{[q0]} + \mathbf{Z}^{[q0]}$	$w = \pm, q = \pm$	$\begin{cases} 2\pi F & \text{for } w \neq q \\ 6\pi F & \text{for } w = q \end{cases}$
2	$\mathbf{Y}^{[s_p,s_p]} + \mathbf{Y}^{[s_{p'},s_{p'}]} \to \mathbf{Y}^{[s_i,s_i]} + \mathbf{Y}^{[s_s,s_s]}$	$s_p + s_{p'} = s_i + s_s$	$4\pi F$
3a	$\mathbf{Z}^{[w0]} + \mathbf{Z}^{[w0]} \to \mathbf{Y}^{[q,q]} + \mathbf{Y}^{[-q,-q]}$	$w = \pm, q = \pm 1$	$4\pi F$
3b	$\mathbf{Y}^{[q,q]} + \mathbf{Y}^{[-q,-q]} \to \mathbf{Z}^{[w0]} + \mathbf{Z}^{[w0]}$	$w = \pm, q = \pm 1$	$4\pi F$
3c	$\mathbf{Z}^{[w0]} + \mathbf{Y}^{[q,q]} \to \mathbf{Z}^{[w0]} + \mathbf{Y}^{[q,q]}$	$w = \pm, q = \pm 1$	$4\pi F$

the pump modes and the same for the output modes. In combinatorial multiset notation, $\binom{n}{k} = 10$ for $n = 4$ modes and $k = 2$ input/output fields. It follows that there are $10 \times 10 = 100$ potential distinct CV mode FWM processes in total in each $|L|$ manifold.

5.7.2.3 Process Amplitudes for the Total Angular Momentum Modes

Here I use the results of the previous two sections to find selection rules for the TAM mode basis, that is, the modes with definite j that are simultaneously fibre eigenmodes. The potential processes can be divided into three cases, which I address separately below. The allowed processes and their amplitudes $\mathsf{U}^{(\mathrm{TAM})}$ are summarized in Table 5.5.

Case 1: Interconversion of $j = 0$ modes First I consider interconversion between the two $j = 0$ modes, $\mathbf{Z}^{[-0]}$ and $\mathbf{Z}^{[+0]}$, which are the TE and TM modes, respectively. All of

the potential processes will trivially conserve the TAM j since, for each mode, $j = 0$. Since these are CV modes, these cases will be governed by the corresponding selection rules and amplitudes from equation 5.53. As in the previous section, the only disallowed processes are ones in which there is a single unpaired mode among the four modes in the process (for example, TM + TE → TM + TM). Note that this process would be allowed by TAM conservation, so this principle by itself is not the sole selection rule for the TAM basis. If all of the modes are alike, $\mathsf{U}^{(\mathrm{TAM})} = 6\pi F$, otherwise $\mathsf{U}^{(\mathrm{TAM})} = 2\pi F$.

Case 2: Interconversion of $j = \pm 2$ modes Next I consider the similar case of conversion between the $j = \pm 2$ modes in the TAM basis. These are the $\mathbf{Z}^{[2]} = \mathbf{Y}^{[1,1]}$ and $\mathbf{Z}^{[-2]} = \mathbf{Y}^{[-1,-1]}$ modes. The relevant selection rules, given by equation 5.44, are simply those that separately conserve SAM (s) and OAM (l): $s_i + s_s = s_p + s_{p'}$ and $l_i + l_s = l_p + l_{p'}$. Since in each beam $s = l$, one type of angular momentum conservation is automatically accompanied by the other. The TAM will also be trivially conserved since $j = s + l$. Thus, the only way to conserve TAM for these modes is for SAM and OAM to be independently conserved.

Case 3: Conversion between $j = 0$ and $j = \pm 2$ modes Finally, I consider interconversion between the $j = \pm 2$ modes and the $j = 0$ modes. The two $j = 0$ modes $\mathbf{Z}^{[\pm 0]}$ can be described with a single expression that depends on one parameter, $m = \pm 1$:

$$\mathbf{Z}^{[m0]} = \mathbf{C}\mathbf{V}^{[-m,m]} = \frac{1}{\sqrt{2}} \mathrm{e}^{i\frac{\pi}{4}(m-1)} \left(\mathbf{Y}^{[1,-1]} + m\mathbf{Y}^{[-1,1]} \right). \tag{5.54}$$

Since the overlap integral O has a linear dependence on each beam's mode, the presence of each Z mode (which has two terms) in the product will double the number of terms in $\mathsf{U}^{(\mathrm{TAM})}$. Each of these terms will correspond to a process amplitude $\mathsf{U}^{(\mathrm{Y})}$ involving the Y modes in equation 5.54. Therefore, the selection rules for the TAM modes can be derived using the selection rules of the constituent Y modes. Importantly, each field q in a Z mode will have anti-aligned SAM and OAM, $s_q = -l_q$, whereas they will be aligned for each beam in a Y mode, $s_q = l_q$. The following will address the sub-cases in which one, two, or three of the beams are in Z modes.

Sub-case 3a: One beam in a $j = 0$ mode. First is the case of a sole beam (for example, the idler beam) in a Z mode. By linearity, the superposition of two Y modes in the Z-mode factor leads to a doubling of terms in the TAM process amplitude, which can be described using the Y-mode process amplitudes $\mathsf{U}^{(\mathrm{Y})}$:

$$\mathsf{U}^{(\mathrm{TAM})} = \frac{1}{\sqrt{2}} \mathrm{e}^{i\frac{\pi}{4}(m_i - 1)} \left(\mathsf{U}^{(\mathrm{Y})}_{+spp'} + m_i \mathsf{U}^{(\mathrm{Y})}_{-spp'} \right). \tag{5.55}$$

Here, $\mathsf{U}^{(\mathrm{Y})}_{ispp'}$ is the process amplitude for all four fields in Y modes (equation 5.44). A field subscript of $n_v = \pm 1$ in place of the usual field label v (i, s, p, or p') indicates that the field v is in mode $\mathbf{Y}^{[n_v, -n_v]}$. Each $\mathsf{U}^{(\mathrm{Y})}_{ispp'}$ term is only non-zero if it satisfies the standard Y-mode selection rules, in which case $\mathsf{U}^{(\mathrm{Y})}_{ispp'} = 4\pi F$ regardless of which four Y modes are involved. At least one term must be an allowed process for $\mathsf{U}^{(\mathrm{TAM})}$ to be non-zero. In the current case, the Y-mode selection rules describing both $\mathsf{U}^{(\mathrm{Y})}_{n_i spp'}$ terms are $n_i + s_s = s_p + s_{p'}$

and $-n_i + l_s = l_p + l_{p'}$ with the value of $n_i = \pm 1$ controlling which term the rules apply to. However, all of the other fields have $s_j = l_j$ (that is, modes $j = i$, s, and p are $\mathbf{Y}^{[\pm n_j, \pm n_j]}$). With this, the l selection rule becomes $-n_i + s_s = s_p + s_{p'}$, and when taken together with the s selection rule $n_i + s_s = s_p + s_{p'}$, the result is $n_i = 0$. This contradicts $n_i = \pm 1$, and thus neither term is allowed and $\mathsf{U}^{(\mathrm{TAM})}$ is always zero in the current case. In other words, a process involving a single $\mathbf{Z}$ mode is not allowed.

Sub-case 3b: Both input beams or both output beams in $j = 0$ modes. Next is the case with two input beams (or output beams, since the allowed FWM processes are the same in reverse) in Z modes. The derivation is similar but one must also consider the relative amplitudes of the terms, since they could cancel each other. Now that there are two Z modes, there are four $\mathsf{U}^{(\mathrm{Y})}$ terms in the process amplitude:

$$\begin{aligned}
\mathsf{U}^{(\mathrm{TAM})} &= \frac{1}{2} \mathrm{e}^{i\frac{\pi}{4}(m_i+m_s-2)} \left(\mathsf{U}^{(\mathrm{Y})}_{++pp'} + m_i \mathsf{U}^{(\mathrm{Y})}_{-+pp'} + m_s \mathsf{U}^{(\mathrm{Y})}_{+-pp'} + m_i m_s \mathsf{U}^{(\mathrm{Y})}_{--pp'} \right) \\
&= \frac{1}{2} \mathrm{e}^{i\frac{\pi}{4}(m_i+m_s-2)} \left(m_i \mathsf{U}^{(\mathrm{Y})}_{-+pp'} + m_s \mathsf{U}^{(\mathrm{Y})}_{+-pp'} \right) \\
&= \frac{1}{2} \mathrm{e}^{i\frac{\pi}{4}(m_i+m_s-2)} 4\pi F \left(m_i + m_s \right) . \qquad (5.56)
\end{aligned}$$

Analogous to the last case, each term must satisfy $n_i + n_s = s_p + s_{p'}$ and $-n_i - n_s = s_p + s_{p'}$. Adding and subtracting these gives $n_i = -n_s$ and $s_p = -s_{p'}$. The latter does not depend on the identity of the Z modes, and applies to any allowed process; the former

governs which terms are non-zero. The first and last terms on the first line of equation 5.56 do not satisfy this requirement and vanish. The second line contains the remaining terms, the cross-terms (such as $n_i = 1$, $n_s = -1$) between the Z-mode superpositions. The process amplitudes are equal, $\mathsf{U}^{(\mathrm{Y})}_{-+pp'} = \mathsf{U}^{(\mathrm{Y})}_{+-pp'} = 4\pi F$, and this gives the third line in equation 5.56. Consequently, in this sub-case, $\mathsf{U}^{(\mathrm{TAM})}$ is non-zero only if $m_i = m_s$. In other words, the two Z modes must be the same (for example, $m_i = m_s$) and the Y modes must be opposite (for example, $s_p = -s_{p'}$). The total process amplitude is $\mathsf{U}^{(\mathrm{TAM})} = \pm 4\pi F \mathrm{e}^{i\frac{\pi}{4}(\pm 2-2)} = 4\pi F$.

Sub-case 3c: One input beam and one output beam in $j = 0$ modes. The next case, where one input beam and one output beam are in a Z mode, follows similar reasoning to the previous cases. That is, $n_i + s_s = n_p + s_{p'}$ and $-n_i + s_s = -n_p + s_{p'}$ must be satisfied for a term to be non-zero. This leads to $s_s = s_{p'}$ and $n_i = n_{p'}$, and the process amplitude becomes

$$\mathsf{U}^{(\mathrm{TAM})} = \frac{1}{2}\mathrm{e}^{i\frac{\pi}{4}(m_i+m_s-2)} 4\pi F \left(1 + m_i m_s\right). \tag{5.57}$$

Thus, $\mathsf{U}^{(\mathrm{TAM})}$ is non-zero only if $m_i = m_s$, which means the input beam is in the same Z mode as the pump, and the remaining input and pump beams are in the same Y mode (for example, $s_s = s_{p'}$). Again, the total process amplitude is $\mathsf{U}^{(\mathrm{TAM})} = 4\pi F$.

Sub-case 3d: Three beams in $j = 0$ modes. The last case, which has three of the four beams in Z modes, is relatively simple. There will be six terms in $\mathsf{U}^{(\mathrm{TAM})}$, each of

which must satisfy $n_i + n_s = n_p + s_{p'}$ and $-n_i - n_s = -n_p + s_{p'}$. Combining these equations leads to the contradiction, $s_{p'} = 0$. Consequently, no such process is possible.

6

Conclusion

In this thesis I have described my work on the following two projects: time-stretched 3D velocity map imaging [2, 3] and four-wave mixing theory for cylindrical vector beams [1]. My work on both of these projects has been published in peer-reviewed journals, as cited above.

6.1 The 3D VMI Project

The 3D VMI project was motivated by the limitations of conventional 2D VMI techniques, which typically rely on 2D-projection inversion algorithms that restrict the variation of the laser polarization geometries. The main advantage of 3D VMI over 2D ion/electron imaging techniques is the capability to measure the full 3D spatial coordinates (x, y, and z)

of each particle as it reaches the detector, thus avoiding the need for inversion algorithms and the constraints they impose. The method we use to accomplish this [25] involves obtaining 2D transverse coordinates (x, y) from a fast CMOS camera imaging an MCP detector backed by a phosphor anode screen. Here the camera is run at the same repetition rate as the laser (that is, one image per laser pulse) and the exposure is much longer than the time spread of the Newton sphere. We simultaneously time-stamp all electron events produced by each laser pulse (TOF being proportional to the z coordinate) by capacitively picking off and digitizing the electrical pulses generated by each impact on the phosphor.

Multi-hit 3D VMI represents a high data rate alternative to 3D photoelectron imaging techniques using delay-line anodes. It accomplishes 3D imaging without relying on the performance of multiple experiments with different polarization geometries and tomographic reconstruction of the 3D electron distribution [69]. The camera-and-digitizer approach to 3D VMI used here does not require the use of a specialized camera such as the Timepix [36] series of event-driven cameras and, as demonstrated here, has a much better (< 80 versus 200 ps) TOF resolution. Indeed, any fast frame rate (1 kHz) CMOS camera with sufficient sensitivity (AST) can be used, which has the benefit of affordability and allows the freedom to select from a wide variety of cameras to suit the needs of a given experiment.

As part of this project, I tested, characterized, and performed several experiments demonstrating the 3D imaging capability of the NRC 3D VMI spectrometer. This instrument is a custom-built 13-electrode thick-lens spectrometer designed for multi-hit 3D photoelectron VMI which stretches the particle Newton spheres in time to maximize z-axis momentum resolution and multi-hit capability. It maintains good VMI focusing conditions

and, importantly, is capable of both 3D and 2D VMI operation. The work presented in this thesis represents the first publications [2, 3] demonstrating 3D imaging with this instrument.

I performed demonstrative photoelectron experiments using atomic xenon, methyl iodide, and nitric oxide, ranging in electron kinetic energy from 0.27 to 2.78 eV, although good performance up to 10 eV is expected for this instrument. From these experiments, I demonstrated a TTS (time-stretching) of 31 and 6.9 ns for 0.27 and 2.78 eV electrons, respectively. I showed that we can achieve a 72 ps TOF resolution over the full 40 mm diameter face of the microchannel plate, making use of a photodiode reference pulse to minimize the effects of electronic trigger jitter. This is a considerable improvement over the estimated time resolution without the photodiode reference, which is 210 ps. For a photoelectron time spread of 10 ns (a reasonable number for our system) this equates to effectively over 130 time-slices through the full distribution. Considering that the PADs are generally smoothly varying functions, this should be sufficient for many applications.

I have created a code package (written in Python) which contains a variety of tools for processing and analyzing 3D VMI data. This will be of great use for later users of the NRC spectrometer and the new spectrometer under construction at the University of Ottawa. This code will be made available and freely accessible through an online repository. Included in this package is an algorithm for assigning the spatial coordinates with their corresponding time-stamps in multi-hit data. From an analysis of the multi-hit capability of the instrument using this data processing code, I demonstrated that time-stretching provides a significant improvement, yielding correct (x, y, t) assignment of more than six

hits per laser shot with near-maximum stretching.

We have performed demonstrative measurements of the full 3D PADs for three photoionization processes in nitric oxide with photoelectron kinetic energies of 0.05 eV, 0.88 eV, and 2.05 eV. The latter two processes result from ionization of the well-known excited NO(A, 3s) state, and fits to the asymmetry parameters for these processes agree with previous results and expectations.

We developed a general and novel momentum vector reconstruction method which can be applied to current and future implementations of 3D VMI and similar 3D charged particle imaging techniques. We showed that the 3D charged-particle distribution can provide a measurement of the actual polarization state of light within the interaction region of the spectrometer, which will in general differ from the polarization state prepared outside the spectrometer. For polarization-sensitive work, access to this information is essential. The ability to measure the polarization state inside the spectrometer is one advantage offered by 3D charged particle imaging. Furthermore, we presented a method for correcting the radially varying gain in MCP detectors often caused by usage wear.

All of the above results and methods will be used in conjunction with the second-generation time-stretched 3D VMI instrument which is currently under construction. It incorporates a co-axial Wiley-McLaren time-of-flight mass spectrometer, as was previously demonstrated [35], allowing us to perform photoelectron-photoion image covariance VMI. We also have plans to improve the multi-hit assignment of the temporal (digitizer) and spatial (camera) coordinates by using an MCP detector having both a higher average gain

and a broader pulse height distribution. As noted in Section 3.6.1, for especially low photoelectron kinetic energies and implemented extraction field gradients, loss of circular symmetry is observed for large time-stretching, which is due to the four-fold broken symmetry of the laser port electrode. We anticipate approximate correction of this via the use of additional small electrodes centred on each port in order to permit local field corrections, thus removing the four-fold symmetry breaking.

The problem of assigning the multi-hit data, particularly when multiple ambiguous assignments are available, is a complex one which warrants a proper statistical treatment. I made some progress toward this, as described in Section 4.3.2.1, but additional work is necessary to develop a statistically justified and optimized solution.

The advantages afforded by 3D photoelectron VMI will permit a variety of novel time-resolved linear and circular dichroism experiments which rely on measuring non-cylindrically-symmetric charged particle distributions. Additionally, 3D VMI promises to enable many experiments in the field of MFPAD reconstruction, for which there has been significant recent interest [14, 44–46].

One experiment which is being planned, which will apply the technique of time-stretched 3D photoelectron VMI to an unanswered problem of chemical interest, involves the photodissociation dynamics of carbon disulfide (CS_2). This experiment involves preparing (using a 200 nm pump pulse) the dissociative $^1B_2(^1\Sigma_u^+)$ state of CS_2 which, upon dissociation, produces ground state CS(X) and atomic sulfur [100]. This dissociation process can proceed along two different pathways, producing either S(^1D) or S(^3P) product states.

It has been previously observed [101] that the apparent lifetime of CS_2 dissociation depends strongly on the polarization geometry of the pump and probe (ionization) pulses. In this paper, the authors propose that the two channels have a polarization-dependent photoionization cross section and that varying the polarization geometry switches which channel is being probed. If this suggestion is correct and if the two channels have different angular momentum properties, it is likely that the PADs also have a distinct polarization dependence. The need for control over the polarization geometry means that an experiment designed to test this theory requires an inversion-free technique. Thus, this is an excellent opportunity to apply 3D VMI to a problem which makes use of the advantages that it offers.

6.2 The Four-Wave Mixing Theory Project

In this project, I derived a set of general selection rules which determine the allowed FWM processes involving CV modes. These structured light modes are eigenmodes of optical fibre and, as such, they are of interest in various applications. These include telecommunications, where the OAM content of such structured light beams represents an additional degree of freedom for information encoding, and can thus be used to increase bandwidth [73, 74]. Furthermore, unwanted nonlinear interactions (such as FWM) and the mode mixing that can result, are a major limitation present in fibre-based signal transmission [79]. Interest in CV modes and their properties extends to quantum cryptography [97] and fundamental optics research [98, 99]. Despite this variety of applications, the fundamental theoretical

understanding of structured light and, in particular, CV modes, is still being developed. This work represents a significant step toward a complete understanding of the nonlinear optical properties of these important modes of light.

6.3 Velocity Map Imaging With Structured Light

Given the nature of the two projects I have summarized above, one might wonder if there is any advantage to be gained by applying structured light beams to a velocity map imaging experiment. In a photoionization experiment, the partial wave contribution to the PAD is determined by the angular momentum contained in the light-matter interaction. Unlike SAM (polarization), which is constrained to only values of ± 1, OAM can take any integer value. Therefore, light containing OAM in addition to SAM can possess higher values of angular momentum than light containing only SAM. This could result in more complex and information-rich PADs and allow access to different photoionization pathways not accessible with ordinary polarized light. Recent photoionization helical dichroism experiments have shown that OAM-containing light beams can be used to control photoelectron yield [102]. In general, the more complex PADs produced in this type of experiment will not possess the cylindrical symmetry required by inversion-reliant 2D VMI and such experiments will thus benefit from the 3D VMI techniques presented in this thesis. Furthermore, once the new coincidence-capable 3D VMI spectrometer at the University of Ottawa is completed, there will be opportunities to perform experiments which can explore the effect of photoionization with OAM-containing light by looking at the ion and electron states

simultaneously.

References

[1]E. S. Goudreau, C. Kupchak, B. J. Sussman, R. W. Boyd, and J. S. Lundeen, "Theory of four-wave mixing of cylindrical vector beams in optical fibers", Journal of the Optical Society of America B - Optical Physics **37**, 1670–1682 (2020).

[2]E. S. Goudreau, A. E. Boguslavskiy, D. J. Moffatt, V. Makhija, M. Hemsworth, R. Lausten, C. Marceau, I. Wilkinson, and A. Stolow, "Time-stretched multi-hit 3d velocity map imaging of photoelectrons", Rev. Sci. Instrum. **94**, 063002 (2023).

[3]A. Sanchez, E. S. Goudreau, A. E. Boguslavskiy, D. J. Moffatt, R. Lausten, I. Wilkinson, V. Makhija, and A. Stolow, "Complete 3d photoelectron momentum vector reconstruction from time-position charged particle imaging", Journal of Physics B: Atomic, Molecular and Optical Physics **58**, 035101 (2025).

[4]M. Hemsworth, MA thesis (University of Ottawa, 2020).

[5]A. T. J. B. Eppink and D. H. Parker, "Velocity map imaging of ions and electrons using electrostatic lenses: application in photoelectron and photofragment ion imaging of molecular oxygen", Review of Scientific Instruments **68**, 3477–3484 (1997).

[6]R. Dörner, V. Mergel, O. Jagutzki, L. Spielberger, J. Ullrich, R. Moshammer, and H. Schmidt-Böcking, "Cold target recoil ion momentum spectroscopy: a 'momentum microscope' to view atomic collision dynamics", Physics Reports **330**, 95–192 (2000).

[7]M. N. R. Ashfold, N. H. Nahler, A. J. Orr-Ewing, O. P. J. Vieuxmaire, R. L. Toomes, T. N. Kitsopoulos, I. A. Garcia, D. A. Chestakov, S.-M. Wu, and D. H. Parker, "Imaging the dynamics of gas phase reactions", Phys. Chem. Chem. Phys. **8**, 26 (2006).

[8]D. M. Neumark, "Slow electron velocity-map imaging of negative ions: applications to spectroscopy and dynamics", J. Phys. Chem. A **112**, 13287 (2008).

[9]R. Mabbs, E. R. Grumbling, K. Pichugin, and A. Sanov, "Photoelectron imaging: an experimental window into electronic structure", Chem. Soc. Rev. **38**, 2169 (2009).

[10]M. H. M. Janssen and I. Powis, "Detecting chirality in molecules by imaging photoelectron circular dichroism", Phys. Chem. Chem. Phys. **16**, 856 (2014).

[11]M. J. J. Vrakking, "Attosecond imaging", Phys. Chem. Chem. Phys. **16**, 2775 (2014).

[12]M. N. R. Ashfold and D. H. Parker, "Imaging molecular dynamics", PHYSICAL CHEMISTRY CHEMICAL PHYSICS **16**, 381–382 (2014).

[13]A. G. Suits, "Photofragment imaging", Rev. Sci. Instrum. **89**, 111101 (2018).

[14]D. Dowek and P. Decleva, "Trends in angle-resolved molecular photoelectron spectroscopy", Phys. Chem. Chem. Phys. **24**, 24614–24654 (2022).

[15]D. W. Chandler and P. L. Houston, "Two-dimensional imaging of state-selected photodissociation products detected by multiphoton ionization", The Journal of Chemical Physics **87**, 1445–1447 (1987).

[16]S. Beaulieu, A. Comby, A. Clergerie, J. Caillat, D. Descamps, N. Dudovich, B. Fabre, R. Géneaux, F. Légaré, S. Petit, B. Pons, G. Porat, T. Ruchon, R. Taïeb, V. Blanchet, and Y. Mairesse, "Attosecond-resolved photoionization of chiral molecules", Science **358**, 1288–1294 (2017).

[17]C. Sparling, A. Ruget, L. Ireland, N. Kotsina, O. Ghafur, J. Leach, and D. Townsend, "The importance of molecular axis alignment and symmetry-breaking in photoelectron elliptical dichroism", The Journal of Chemical Physics **159**, 214301 (2023).

[18]G. Basnayake, Y. Ranathunga, S. K. Lee, and W. Li, "Three-dimensional (3d) velocity map imaging: from technique to application", J. Phys. B **55**, 023001 (2022).

[19]D. Townsend, M. P. Minitti, and A. G. Suits, "Direct current slice imaging", Rev. Sci. Instrum. **74**, 2530 (2003).

[20]J. J. Lin, J. Zhou, W. Shiu, and K. Liu, "Application of time-sliced ion velocity imaging to crossed molecular beam experiments", Rev. Sci. Instrum. **74**, 2495 (2003).

[21]W. Li, S. D. Chambreau, S. A. Lahankar, and A. G. Suits, "Megapixel ion imaging with standard video", Rev. Sci. Instrum. **76**, 063106 (2005).

[22]S. Kauczok, N. Gödecke, A. I. Chichinin, M. Veckenstedt, C. Maul, and K.-H. Gericke, "Three-dimensional velocity map imaging: setup and resolution improvement compared to three-dimensional ion imaging", Rev. Sci. Instrum. **80**, 083301 (2009).

[23]O. Jagutzki, A. Cerezo, A. Czasch, R. Dorner, M. Hattas, M. Huang, V. Mergel, U. Spillmann, K. Ullmann-Pfleger, T. Weber, H. Schmidt-Bocking, and G. Smith, "Multiple hit readout of a microchannel plate detector with a three-layer delay-line anode", IEEE Transactions on Nuclear Science **49**, 2477–2483 (2002).

[24]S. K. Lee, F. Cudry, Y. F. Lin, S. Lingenfelter, A. H. Winney, L. Fan, and W. Li, "Coincidence ion imaging with a fast frame camera", Review of Scientific Instruments **85**, 123303 (2014).

[25]S. K. Lee, Y. F. Lin, S. Lingenfelter, L. Fan, A. H. Winney, and W. Li, "Communication: time- and space-sliced velocity map electron imaging", J. Chem. Phys. **141**, 221101 (2014).

[26]Y. F. Lin, S. K. Lee, P. Adhikari, T. Herath, S. Lingenfelter, A. H. Winney, and W. Li, "Note: an improved 3d imaging system for electron-electron coincidence measurements", Rev. Sci. Instrum. **86**, 096110 (2015).

[27]L. Fan, S. K. Lee, Y.-J. Tu, B. Mignolet, D. Couch, K. Dorney, Q. Nguyen, L. Wooldridge, M. Murnane, F. Remacle, H. B. Schlegel, and W. Li, "A new electron-ion coincidence 3d momentum-imaging method and its application in probing strong field dynamics of 2-phenylethyl-n, n-dimethylamine", J. Chem. Phys. **147**, 013920 (2017).

[28]A. H. Winney, S. K. Lee, Y. F. Lin, Q. Liao, P. Adhikari, G. Basnayake, H. B. Schlegel, and W. Li, "Attosecond electron correlation dynamics in double ionization of benzene probed with two-electron angular streaking", Phys. Rev. Lett. **119**, 123201 (2017).

[29]Q. Liao, A. H. Winney, S. K. Lee, Y. F. Lin, P. Adhikari, and W. Li, "Coulomb-repulsion-assisted double ionization from doubly excited states of argon", Phys. Rev. A **96**, 023401 (2017).

[30]C. Weeraratna, C. Amarasinghe, S. K. Lee, W. Li, and A. G. Suits, "Demonstration of multi-hit and multi-mass capability of 3d imaging in a conventional velocity map imaging experiment", J. Chem. Phys. **149**, 084202 (2018).

[31]X. Urbain, D. Bech, J.-P. Van Roy, M. Géléoc, S. J. Weber, A. Huetz, and Y. J. Picard, "A zero dead-time multi-particle time and position sensitive detector based on correlation between brightness and amplitude", Review of Scientific Instruments **86**, 023305 (2015).

[32]A. Zhao, M. van Beuzekom, B. Bouwens, D. Byelov, I. Chakaberia, C. Cheng, E. Maddox, A. Nomerotski, P. Svihra, J. Visser, V. Vrba, and T. Weinacht, "Coincidence velocity map imaging using tpx3cam, a time stamping optical camera with 1.5 ns timing resolution", Review of Scientific Instruments **88**, 113104 (2017).

[33]D. A. Debrah, G. A. Stewart, G. Basnayake, A. Nomerotski, P. Svihra, S. K. Lee, and W. Li, "Developing a camera-based 3d momentum imaging system capable of 1 mhits/s", Review of Scientific Instruments **91**, 023316 (2020).

[34]C. Cheng, G. Moğol, T. Weinacht, A. Nomerotski, and C. Trallero-Herrero, "3d velocity map imaging of electrons with tpx3cam", Review of Scientific Instruments **93**, 013003 (2022).

[35]M. Davino, E. McManus, N. G. Helming, C. Cheng, G. Mogol, Z. Rodnova, G. Harrison, K. Watson, T. Weinacht, G. N. Gibson, T. Saule, and C. A. Trallero-Herrero, "A plano-convex thick-lens velocity map imaging apparatus for direct, high resolution 3d momentum measurements of photoelectrons with ion time-of-flight coincidence", Rev. Sci. Instrum. **94**, 013303 (2023).

[36]R. Ballabriga, M. Campbell, and X. Llopart, "Asic developments for radiation imaging applications: the medipix and timepix family", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **878**, 10–23 (2018).

[37]X. Llopart, J. Alozy, R. Ballabriga, M. Campbell, R. Casanova, V. Gromov, E. Heijne, T. Poikela, E. Santin, V. Sriskaran, L. Tlustos, and A. Vitkovskiy, "Timepix4, a large area pixel detector readout chip which can be tiled on 4 sides providing sub-200 ps timestamp binning", Journal of Instrumentation **17**, C01044 (2022).

[38]R. Wu, B. Ding, Y. Feng, K. Wu, X. Jin, and X.-J. Liu, "Scaling analysis parameters in 3d-focusing velocity map imaging", MEASUREMENT SCIENCE AND TECHNOLOGY **34** (2023).

[39]N. Yonekura, C. Gebauer, H. Kohguchi, and T. Suzuki, "A crossed molecular beam apparatus using high-resolution ion imaging", Review of Scientific Instruments **70**, 3265–3270 (1999).

[40]S.-Y. Liu, K. Alnama, J. Matsumoto, K. Nishizawa, H. Kohguchi, Y.-P. Lee, and T. Suzuki, "He i ultraviolet photoelectron spectroscopy of benzene and pyridine in supersonic molecular beams using photoelectron imaging", J. Phys. Chem. A **115**, 2953 (2011).

[41]I. Wilkinson, PhD thesis (University of Leeds, 2010).

[42]N. G. Kling, D. Paul, A. Gura, G. Laurent, S. De, H. Li, Z. Wang, B. Ahn, C. H. Kim, T. K. Kim, I. V. Litvinyuk, C. L. Cocke, I. Ben-Itzhak, D. Kim, and M. F. Kling, "Thick-lens velocity-map imaging spectrometer with high resolution for high-energy charged particles", J. Instrum. **9**, P05005 (2014).

[43]G. R. Harrison, J. C. Vaughan, B. Hidle, and G. M. Laurent, "Davis: a direct algorithm for velocity-map imaging system", The Journal of Chemical Physics **148**, 194101 (2018).

[44]C. Marceau, V. Makhija, D. Platzer, A. Y. Naumov, P. B. Corkum, A. Stolow, D. M. Villeneuve, and P. Hockett, Phys. Rev. Lett. **119**, 083401 (2017).

[45]M. Gregory, P. Hockett, A. Stolow, and V. Makhija, J. Phys. B **54**, 145601 (2021).

[46]P. Hockett and V. Makhija, "Topical review: extracting molecular frame photoionization dynamics from experimental data", Journal of Physics B: Atomic, Molecular and Optical Physics **56**, 112001 (2023).

[47]A. Stolow and J. G. Underwood, "Time-resolved photoelectron spectroscopy of nonadiabatic dynamics in polyatomic molecules", in *Advances in chemical physics* (John Wiley & Sons, Inc., 2008), pp. 497–584.

[48]C. N. Yang, "On the angular distribution in nuclear reactions and coincidence measurements", Phys. Rev. **74**, 764–772 (1948).

[49]R. Bracewell, *The fourier transform and its applications*, 3rd (McGraw Hill, 2000).

[50]G. A. Garcia, L. Nahon, and I. Powis, "Two-dimensional charged particle image inversion using a polar basis function expansion", Review of Scientific Instruments **75**, 4989–4996 (2004).

[51]G. M. Laurent and G. R. Harrison, *The scaling properties and the multiple derivative of legendre polynomials*, 2017.

[52]T. Horio and T. Suzuki, "Multihit two-dimensional charged-particle imaging system with real-time image processing at 1000 frames/s", Rev. Sci. Instrum. **80**, 013706 (2009).

[53]M. Tsubouchi, C. de Lange, and T. Suzuki, "Ultrafast dissociation processes in the no dimer studied with time-resolved photoelectron imaging", J. Electron Spectrosc. Relat. Phenom. **142**, 193 (2005).

[54]Y. Cho and S. Na, "Application of abel inversion in real-time calculations for circularly and elliptically symmetric radiation sources", Measurement Science and Technology **16**, 878–884 (2005).

[55]IMI Adaptas, Palmer, MA, *SIMION*, version 8.0.8.1, May 14, 2012.

[56]U. Even, "The even-lavie valve as a source for high intensity supersonic beam", EPJ Tech. Instrum. **2**, 17 (2015).

[57]B. Winter, S. J. King, M. Brouard, and C. Vallance, "A fast microchannel plate-scintillator detector for velocity map imaging and imaging mass spectrometry", Review of Scientific Instruments **85**, 023306 (2014).

[58]European Machine Vision Association, *Emva standard 1288 - standard for characterization of image sensors and cameras, release 4.0 general*, tech. rep. (2020).

[59]FLIR Systems, Inc., *Sensor review: mono cameras*, tech. rep., 19-0709-B&CS (2019).

[60]MATRIX VISION GmbH, *Emva 1288 data sheet m0782*, tech. rep. (2017).

[61]ams AG, *Datasheet - cmv12000*, tech. rep., DS000603 (2019).

[62]O. J. Clarkin, PhD thesis (Queen's University, 2012).

[63]C. Heckenkamp, F. Schäfers, G. Schönhense, and U. Heinzmann, "Experimental characterization of the xe 5p photoionization by angle- and spin-resolved photoelectron spectroscopy", Zeitschrift für Physik D Atoms, Molecules and Clusters **2**, 257–274 (1986).

[64]N. Thire, R. Cireasa, V. Blanchet, and S. T. Pratt, "Time-resolved photoelectron spectroscopy of the ch3i (be)-e-1 6s [2] state", Physical Chemistry Chemical Physics **12**, 15644–15652 (2010).

[65]O. Geßner, A. M. D. Lee, J. P. Shaffer, H. Reisler, S. V. Levchenko, A. I. Krylov, J. G. Underwood, H. Shi, A. L. L. East, D. M. Wardlaw, E. t. H. Chrysostom, C. C. Hayden, and A. Stolow, "Femtosecond multidimensional imaging of a molecular dissociation", Science **311**, 219 (2006).

[66]A. V. Demyanenko, A. B. Potter, V. Dribinski, and H. Reisler, "No angular distributions in the photodissociation of (no)2 at 213 nm: deviations from axial recoil", J. Chem. Phys. **117**, 2568 (2002).

[67]V. Blanchet and A. Stolow, "Nonadiabatic dynamics in polyatomic systems studied by femtosecond time-resolved photoelectron spectroscopy", J. Chem. Phys. **108**, 4371 (1998).

[68]M. Wollenhaupt, M. Krug, J. Koehler, T. Bayer, C. Sarpe-Tudoran, and T. Baumert, "Three-dimensional tomographic reconstruction of ultrashort free electron wave packets", Applied Physics B - Lasers and Optica **95**, 647–651 (2009).

[69]C. Sparling and D. Townsend, "Tomographic reconstruction techniques optimized for velocity-map imaging applications", Journal of Chemical Physics **157**, 10.1063/5.0101789 (2022).

[70]D. F. Crouse, "On implementing 2d rectangular assignment algorithms", IEEE Transactions on Aerospace and Electronic Systems **52**, 1679–1696 (2016).

[71]P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python", Nature Methods **17**, 261–272 (2020).

[72]J. A. Davies, R. E. Continetti, D. W. Chandler, and C. C. Hayden, "Femtosecond time-resolved photoelectron angular distributions probed during photodissociation of NO_2", Phys. Rev. Lett. **84**, 5983–5986 (2000).

[73]N. Bozinovic, Y. Yue, Y. Ren, M. Tur, P. Kristensen, H. Huang, A. E. Willner, and S. Ramachandran, "Terabit-scale orbital angular momentum mode division multiplexing in fibers", Science **340**, 1545–1548 (2013).

[74]J. Liu, S.-M. Li, L. Zhu, A.-D. Wang, S. Chen, C. Klitis, C. Du, Q. Mo, M. Sorel, S.-Y. Yu, X.-L. Cai, and J. Wang, "Direct fiber vector eigenmode multiplexing transmission seeded by integrated optical vortex emitters", Light: Science & Applications **7**, 17148–17148 (2018).

[75]W. Qiao, T. Lei, Z. Wu, S. Gao, Z. Li, and X. Yuan, "Approach to multiplexing fiber communication with cylindrical vector beams", Opt. Lett. **42**, 2579–2582 (2017).

[76]J. Li, J. Zhang, F. Li, X. Huang, S. Gao, and Z. Li, "Dd-ofdm transmission over few-mode fiber based on direct vector mode multiplexing", Opt. Express **26**, 18749–18757 (2018).

[77]G. Vallone, V. D'Ambrosio, A. Sponselli, S. Slussarenko, L. Marrucci, F. Sciarrino, and P. Villoresi, "Free-space quantum key distribution by rotation-invariant twisted photons", Phys. Rev. Lett. **113**, 060503 (2014).

[78]R. Dorn, S. Quabis, and G. Leuchs, "Sharper focus for a radially polarized light beam", Phys. Rev. Lett. **91**, 233901 (2003).

[79]P. P. Mitra and J. B. Stark, "Nonlinear limits to the information capacity of optical fibre communications", Nature **411**, 1027–1030 (2001).

[80]Q. Lin and G. P. Agrawal, "Vector theory of four-wave mixing: polarization effects in fiber-optic parametric amplifiers", J. Opt. Soc. Am. B **21**, 1216–1224 (2004).

[81]K. Garay-Palmett, D. Cruz-Delgado, F. Dominguez-Serna, E. Ortiz-Ricardo, J. Monroy-Ruz, H. Cruz-Ramirez, R. Ramirez-Alarcon, and A. B. U'Ren, "Photon-pair generation by intermodal spontaneous four-wave mixing in birefringent, weakly guiding optical fibers", Phys. Rev. A **93**, 033810 (2016).

[82]G. Agrawal, *Nonlinear fiber optics*, Electronics & Electrical (Elsevier Science, 2007).

[83]H. Pourbeyram and A. Mafi, "Photon pair generation in multimode optical fibers via intermodal phase matching", Phys. Rev. A **94**, 023815 (2016).

[84]E. Nazemosadat, H. Pourbeyram, and A. Mafi, "Phase matching for spontaneous frequency conversion via four-wave mixing in graded-index multimode optical fibers", J. Opt. Soc. Am. B **33**, 144–150 (2016).

[85]R. Jáuregui and J. P. Torres, "On the use of structured light in nonlinear optics studies of the symmetry group of a crystal", Scientific Reports **6**, 20906 (2016).

[86]R. Y. Saaltink, L. Giner, R. W. Boyd, E. Karimi, and J. S. Lundeen, "Super-critical phasematching for photon pair generation in structured light modes", Opt. Express **24**, 24495–24508 (2016).

[87]J. Arlt, K. Dholakia, L. Allen, and M. J. Padgett, "Efficiency of second-harmonic generation with bessel beams", Phys. Rev. A **60**, 2438–2441 (1999).

[88]A. . Snyder and J. Love, "Optical waveguide theory", in, 1st ed. (Springer, 1983) Chap. 14, pp. 304–305.

[89]S. Chen and J. Wang, "Theoretical analyses on orbital angular momentum modes in conventional graded-index multimode fibre", Scientific Reports **7**, 3990 (2017).

[90]Q. Zhan, "Cylindrical vector beams: from mathematical concepts to applications", Adv. Opt. Photon. **1**, 1–57 (2009).

[91]L. Allen, M. W. Beijersbergen, R. J. C. Spreeuw, and J. P. Woerdman, "Orbital angular momentum of light and the transformation of laguerre-gaussian laser modes", Phys. Rev. A **45**, 8185–8189 (1992).

[92]P. Gregg, M. Mirhosseini, A. Rubano, L. Marrucci, E. Karimi, R. W. Boyd, and S. Ramachandran, "Q-plates as higher order polarization controllers for orbital angular momentum modes of fiber", Opt. Lett. **40**, 1729–1732 (2015).

[93]R. W. Boyd, *Nonlinear optics, third edition*, 3rd (Academic Press, 2008).

[94]R. Fickler, R. Lapkiewicz, W. N. Plick, M. Krenn, C. Schaeff, S. Ramelow, and A. Zeilinger, "Quantum entanglement of high angular momenta", Science **338**, 640–643 (2012).

[95]P. G. Kwiat, K. Mattle, H. Weinfurter, A. Zeilinger, A. V. Sergienko, and Y. Shih, "New high-intensity source of polarization-entangled photon pairs", Phys. Rev. Lett. **75**, 4337–4341 (1995).

[96]O. Cohen, J. S. Lundeen, B. J. Smith, G. Puentes, P. J. Mosley, and I. A. Walmsley, "Tailored photon-pair generation in optical fibers", Phys. Rev. Lett. **102**, 123603 (2009).

[97]B. Ndagano, I. Nape, M. A. Cox, C. Rosales-Guzman, and A. Forbes, "Creation and detection of vector vortex modes for classical and quantum communication", Journal of Lightwave Technology **36**, 292–301 (2018).

[98]L. Allen, S. Barnett, and M. Padgett, *Optical angular momentum*, Optics & Optoelectronics (Taylor & Francis, 2003).

[99]D. Andrews and M. Babiker, *The angular momentum of light* (Cambridge University Press, 2013).

[100]A. Mank, C. Starrs, M. Jego, and J. Hepburn, "A detailed study of the predissociation dynamics of the b-1(2)((1)sigma(+)(u)) state of cs2", JOURNAL OF CHEMICAL PHYSICS **104**, 3609–3619 (1996).

[101]D. Townsend, H. Satzger, T. Ejdrup, A. M. D. Lee, H. Stapelfeldt, and A. Stolow, "B21(Σu+1) excited state decay dynamics in cs2", The Journal of Chemical Physics **125**, 234302 (2006).

[102]J.-L. Bégin, PhD thesis (University of Ottawa, 2024).