PROBING ELECTRON HOLOGRAMS WITH STRONG FIELDS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Nicholas Werby
June 2024

Abstract

The ultrafast dynamics of strong-field tunnel ionization are recorded within intricate interference patterns in photoelectron momentum distributions. We categorize these dynamics with quantum electron trajectories which outline the possible paths an electron may take on its way to the detector. The pair-wise interference of these trajectories give rise to the holograms observable in strong-field spectra. These trajectories are determined by only three parameters: their ionization phases in the ionizing field, their initial transverse momenta, and the shape of the laser field they subsequently propagate within. The laser field tends to be a measurable quantity and the initial transverse momentum can often be inferred from the electron's position on the detector; however, experimentally measuring ionization phases is much more difficult. Existing techniques to measure these ionization phases are commonly stymied in the presence of multiple contributing trajectories, each with their own ionization phase, which reduces their viability when probing electron holograms. To more effectively study electron holograms in strong-field spectra, I developed two powerful techniques which help unlock the ionization phase of electron trajectories as an experimental observable. I methodically apply these techniques to strong-field ionized photoelectron momentum distributions of atomic argon to assess their effectiveness through comparisons with quantitative calculations of strong-field trajectory interference.

The first technique, the time-correlation filter, employs a differential Fourier analysis to reveal, isolate, or selectively remove individual holograms based on the ionization time separation of the pair of trajectories forming them. This has broad implications for experimentally measuring trajectories forming holographic structure. As one example, by removing unwanted intercycle interference structures I reveal a new hologram modulating the prominent spider-leg hologram. Furthermore, by isolating this new hologram in a defined window of time separations, I trace its origin to a specific class of electron trajectory which weakly interacts with the residual parent ion interfering with trajectories ionized half a laser cycle earlier. In order to test the validity of the technique, I make multiple comparisons of time-correlation filtered experimental data to calculations of specific pairs of interfering trajectories and show excellent agreement. Because the time-correlation filter is an analysis technique and does not require any experimental modification, it can be applied to all energy- and angularly-resolved datasets to measure pairwise interference in a broad range of strong-field ionization spectra. Additionally, the results of the time-correlation filter are independent of theory, and so can be used to directly measure time separations from an experiment, without relying on comparisons to supporting calculations from trajectory-based models of strong-field dynamics.

The second technique, three-color ionization phase extraction, involves adding a small, phase-controlled, second and third harmonic perturbation to the ionizing laser field to directly measure the ionization phases of multiple trajectories arriving at each momentum bin on the detector. To examine the benefits of a three-laser-color system, I present the predecessor technique of two-color phase-of-the-phase analysis on a high-resolution experimental dataset. I reveal how it effectively measures the ionization phase of trajectories arriving at regions of the detector where they are the heavily dominant or only trajectory, yet fails to extract even a single relevant phase in regions with more than one trajectory of comparable amplitude. With three laser colors, I outline how to extract the ionization phases of the two highest amplitude electron trajectories and compare to calculation to show quantitative agreement to within 5° of phase, which corresponds to 37 attoseconds for an 800nm laser. The success of this technique is a persuasive argument in favor of employing three-color laser probes in future studies of strong-field dynamics.

Acknowledgments

First and foremost I'd like to extend my most heartfelt gratitude to my advisor, Professor Philip Bucksbaum. When I first arrived at Stanford University, Phil persuaded me to join his lab using beautiful images of strong-field photoelectron momentum distributions. Restoring the original experimental setup and producing my own versions of these beautiful datasets was my first ten-week laboratory rotation which blossomed into the full PhD research project examined in this dissertation. Phil taught me the value of exploratory science, where discoveries are derived from novel measurement techniques on data-rich systems. Most of all, Phil has supported me throughout my graduate career, regularly meeting with me one-on-one to assess my progress and guide me towards insightful new avenues to examine experimental spectra.

I would also like to thank James Cryan for all of his assistance within the laboratory. Although I never worked for him directly, he was instrumental in helping me find my footing in a laser lab. Without his assistance, it is certain I would not have been able to make most of the measurements which are featured in this thesis.

I would like to thank Andrew Maxwell of Aarhus University and Carla Figueira de Morisson Faria of University College of London for the Coulomb Quantum Orbit Strong Field Approximation framework used extensively throughout this thesis. Our very fruitful years-long collaboration has resulted in several jointly published papers and innumerable stimulating conversations. For this thesis, I employ their calculations of transition matrices for argon at a wide-range of laser intensities for the individual quantum orbits predicted by the formalism they developed. Additionally, Andrew generated the Qprop calculations which are used to calibrate the experimental laser intensities as is described in Chapter 5.

My fellow members of the Bucksbaum group have been a source of much support, both as colleagues and as friends. When I first began working in the laser lab in which I'd spend the next five years I worked alongside Anna Li Wang on parallel research projects. With her guidance I learned the basics of ultrafast laser wrangling, which was critically important in my lab perhaps more so than most. In a similar vein I must express my most sincere appreciation for another graduate student, Jun Wang, who took it upon himself to become the foremost expert on whipping our laser back into shape when it (frequently) misbehaved. His laboratory acumen is unrivaled, and despite being the more senior graduate student, I often found myself relying on his insight and expertise when operating within the lab. I'd also like to thank my office-mate Emily Thierstein, with whom I've spent many productive afternoons discussing research hypotheses and techniques to analyze data, alongside an equal number of "less-productive" afternoons gossiping and trading work-related war-stories. As a whole, the Bucksbaum group has been the source of much inspiration. They are an extremely talented collection of graduate students and post-docs, and it was a pleasure to work alongside them for the past six years.

I could not have made it to where I am today without the tremendous support of my parents, Christopher and Olga Werby. After the birth of my daughter in the middle of my graduate career, my parents actively helped me manage both my research and my parenting obligations, without letting me neglect one or the other. My success in both has only been possible with their continuous effort.

Lastly I must acknowledge my beautiful daughter Suha. Even if she would much rather that I play with her all day every day rather than do my work, she is very excited that soon I will be *Doctor* Werby. She doesn't yet understand what I do, but knows that it is "as difficult as catching one hundred butterflies." Well within these pages are my hundred butterflies, and I can't wait until she's old enough to see them for herself.

My research is supported by the U.S. Department of Energy, Office of Science, Basic Energy Sciences (BES), Chemical Sciences, Geosciences, and Biosciences Division, AMOS Program.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Electrons govern the physics of electricity. Within ordinary conductors electrons play the role of the charge carriers, responding to external electric fields to generate electric currents. Electrons interact with electromagnetic fields by absorbing or emitting photons in accordance to their energy levels within atoms and molecules. The motion (or lack thereof) of electrons determines the properties of most novel materials such as semiconductors, superconductors, and topological insulators. Electrons also dominate the fundamental processes of chemistry, with the majority of properties of atoms and molecules being determined by the electron energy levels, orbitals, and bonds.

Because of its critical importance in so many facets of science and technology, there is a massive amount of research across the various scientific disciplines probing electron properties, dynamics, and control. However, because the natural time-scale of electron motion occurs on the order of hundreds of attoseconds (10×10^{-18} s), it is extremely difficult to attempt to resolve electronic motion directly. Instead, many disciplines focus on bulk properties of materials, statistical models, electric fields, and observable chemical phenomena to infer the broad behavior of electrons. Only relatively recently has technology advanced to the point where more direct methods are possible.

Within the last half century, there are few areas of science which have seen as explosive a growth in interest and scientific discovery as ultrafast physics. Scientists within this field take aim at the shortest time scales of chemistry, and attempt to image and describe the fundamental dynamics of molecules, atoms, and electrons. Currently, ultrafast physics explores phenomena occurring in timescales ranging from hundreds of femtoseconds (10×10^{-15} s) to tens of attoseconds [1, 3–13]. Examples of these processes on the "slower" end are atomic and molecular motion, photon absorption, and electron scattering. Even for ultrafast physicists, electron ionization dynamics in atoms and molecules are on the fastest edge of what is currently observable.

More than most other fields, the development of ultrafast physics has been driven by a marriage of technological advancement, experimental discovery, and theoretical study. Oftentimes, the experimental side of a field is forced to play catch-up to validate the predictions of a well-developed theory, or a field is stalled due to limitations of current technology. The history of ultrafast physics, on the other hand, has been punctuated by technological leaps which led directly to experimental discovery of fully unexpected and unexplained phenomena, inviting further study to develop new frameworks of explanation. Relatively recently, the development of chirped-pulse amplification has allowed for the production of high-intensity tabletop lasers with pulse durations of just a few femtoseconds [14], allowing physicists to study light-matter interactions on this ultrafast timescale.

This thesis focuses on using a commercially-available tabletop ultrafast laser to probe the dynamics of electrons undergoing strong-field tunnel ionization and subsequent motion in the laser field. Developing the tools to measure and understand these processes is a critical advancement for physics as a whole. As a quantum mechanical phenomenon, tunnel ionization is not yet fully understood, and so direct measurements of tunneling dynamics would help validate quantum mechanical theories. These measurement tools are also the gateway to attosecond quantum control, which would have major implications for steering chemical reactions and directing light-matter interactions. A complete toolbox for imaging electron dynamics on this timescale would pave the way for new types of quantum information processing and medical imaging technologies at the finest of resolutions.

1.1 Experimental Design Philosophy

In my first year as a graduate student, I was told: “A good experiment has exactly one hard component.” This is an experimental design philosophy revolving around streamlining experiments to target a specific problem that needs solving. Proper science builds on the work of previous scientists. Through discovery and innovation, techniques are developed to probe novel aspects of systems of interest. As experimentalists, we strive to design our experiment around an unknown. The results of the body of previous work provide a framework of what is currently known and achievable, and our insight and ingenuity are used to push beyond and probe the unknown.

In my studies of strong-field ionization, I take the role of the cinematographer. My focus in my research and thus in this dissertation is not to explore some novel phenomenon or observe some intricate dynamics of broad interest to the scientific community, but rather to develop the tools and imaging techniques that would enable a fellow scientist to build on my work to make those observations. To that end, the “hard” portion of the experiments I have performed has been my infrared strong-field probe laser. When ionizing a gas sample with a longer wavelength strong-field laser, electrons experience much more complicated dynamics in the oscillating laser field, resulting in broad energy distributions with many overlapping interference features. This probe is to me what a powerful camera is to a cinematographer, and these electron energy distributions are my photographs. It is my job to manipulate this complicated camera and develop the tools to interpret the images it produces in order to make ultrafast measurements that have never been made before.

Since developing the tools to capture and interpret these images is my primary research goal, my experimental design philosophy suggests that the other critical elements of strong-field ionization be simplified as much as possible. To that end, I focus on the strong-field spectrum of one of the simplest possible gas-phase targets: atomic argon. Atomic argon does not contain a complicated electronic structure, nor does it present any asymmetries which may confound analysis. Ionizing atomic argon allows me to focus nearly entirely on the dynamics of electrons in strong fields, which makes the techniques I have developed and present in the upcoming chapters widely applicable to strong-field experiments.

1.2 Thesis Structure

The following seven chapters present an in-depth demonstration of ultrafast quantum strong-field tunneling ionization dynamics by disentangling and analyzing photoelectron holograms collected through strong-field ionization of gas phase samples in a velocity map imaging spectrometer. Chapter two provides a brief overview of strong-field ionization and photoelectron holography. It also describes the primary theory models which is employed throughout the thesis. Chapter three describes the methodology of the experiments analyzed in later chapters. It pays particular focus to the optomechanical construction of the experiments and the techniques of data collection. Chapter four describes how to interpret photoelectron momentum distributions, which are the main type of data collected in these experiments. It explains how to recreate the cross section of the electron's three-dimensional momentum distribution, identifies the rough momentum boundaries for electron rescattering processes, and presents the primary two-path interference pathways which give rise to the bulk of visible holographic features. Chapter five introduces the time-correlation filter, which is a powerful tool developed to extract time separations between pairs of interfering electron trajectories given an angle-resolved electron energy distribution. It examines in detail the most prevalent holograms visible in strong-field spectra and isolate individual interferences for careful characterization. Chapter six discusses how to use shaped strong fields to extract ionization phases from multiple trajectories contributing to a hologram. It carefully describes two approaches to tackle this problem, one involving the overlap of two laser colors and the other involving the overlap of three. Chapter seven pivots to a speculative position, presenting other thought-provoking features visible in strong-field spectra which do not as of yet have a concrete explanation. Finally, Chapter eight concludes the thesis and presents parting thoughts and a direction for future application of the experimental techniques and analyses explored in the previous chapters.

1.2.1 One Final Note

One of the purposes of a physics dissertation is to provide specialized guidance for fellow scientists attempting to perform similar experiments. In an ordinary scientific article, there is little justification for providing long-winded explanations of the minutiae of an experimental procedure, and so important but small details about how to actually recreate an experiment are omitted. In this dissertation I attempt to provide sufficient explanation of the technical components of my experimental designs, data collection systems, hardware setups, and analysis considerations; however, these technical components may not be of broad interest to most readers.

> Where reasonable, I include these detailed explanations in indented paragraphs with a vertical bar at left such as this one. If the details being described are of no immediate interest to the reader, all such sections may be skipped, and no text outside of sections like these refer to material only described in these sections.

Let me also provide a quick comment on figure interpretation throughout this thesis. There are many different types of data analysis figures in this thesis, but I maintain a consistent style across all of them. A large portion of these figures show angle-resolved two-dimensional images which may be presented with logarithmic or linear colorscales depending on the type of data being shown. In general, all grey-scale images are logarithmic and all full-color colorscales are linear. All such figures are generated in Matlab, and for reference use the “bone” and “jet” color-map designations respectively. Certain figures use a cyclic colorscale to denote phase, and these are plotted on the uniform scientific color-map “viko” [15].

Except where otherwise indicated, all quantities and equations are written in atomic units.

Chapter 2

Electron Dynamics During Strong-Field Ionization

When a sufficiently intense, ultrashort laser pulse is focused onto a target, electrons are liberated and may undergo a variety of dynamics in the presence of the electric field of the laser. The interplay between these electrons and the electric field results in many possible interactions, during which electrons may be accelerated, driven back to their parent ion, and rescattered, all the while self-interfering [2, 5, 16–21]. Currently, experiments seeking to measure these processes cannot do so locally, and thus must rely on indirect methods to interrogate the system and interpret the results. Thus the first task of a photoelectron cinematographer is to develop the models of strong-field ionization that allow us to trace back the features measurable in an experiment to reveal the ultrafast phenomena of interest.

This chapter outlines the basic mechanics of strong-field ionization and explains the primary models and approximations which are applied throughout the rest of this thesis. There are many competing models of strong-field ionization [1, 3, 22–24], so in order to rein in the scope of this chapter, I exclusively focus on topics and descriptions which are directly relevant to the experiments presented in later chapters. Wherever possible, I use the specific parameters of these experiments to simplify equations from their general form to a form suitable for immediate application.

2.1 Overview of Strong-Field Ionization

In this section, I provide a brief outline of the basic mechanics of strong-field ionization, describe the approximations used throughout this thesis, outline the leading model of tunnel ionization, and introduce the framework of electron trajectories and photoelectron holography.

2.1.1 Single Active Electron Approximation

The single active electron (SAE) approximation is a model of electron dynamics commonly used in the study of light-matter interactions, particularly in the context of strong-field ionization. This approximation restricts the consideration of the interacting system to a single electron (the active electron), the laser field driving the interaction, and the combined background potential of the leftover nucleus and electrons of the target. Typically this electron is the most weakly bound electron, due to the sharply reduced rate of ionization for more tightly bound electrons with higher ionization energies [1].

There are a number of advantages to this approximation, with the most obvious being the reduction of complexity in calculation and explanation. In exchange, calculations relying on the SAE approximation are unable to predict any phenomena resulting from electron-electron interactions in many electron systems. Under most circumstances in strong-field ionization, it is reasonable to use the SAE approximation when considering the effect of laser pulses whose photon energy is much less than the ionization potential of the ionization target. Since this thesis exclusively describes experiments which employ 800 nm laser light, which has an individual photon energy of 1.55 eV, as our principle source of ionization while our principle target is argon gas which has an ionization potential (I_p) of 15.76 eV, the SAE approximation is implicitly assumed for all calculations when not otherwise indicated.

2.1.2 Basics of Strong-Field Ionization

Except where otherwise specified, the electric fields of the laser pulses described in this thesis are linearly polarized, and are described by

$$\vec{E}(t) = E_0 \exp\left(- \left(\frac{\omega t}{\pi N} \right)^2 \right) \cos(\omega t - \phi)\hat{z}. \tag{2.1}$$

Here E_0 is the peak amplitude of the pulse, N is the number of cycles contained within the full width at half maximum of the Gaussian envelope, ω is the central angular frequency of the laser, and ϕ is some phase offset. For the applications within this thesis, N is sufficiently high (~ 20) that the effect of the Gaussian envelope and the phase can be ignored for simplicity without disrupting the conclusions drawn in the text. In this case, when examining the dynamics occurring within the bulk of the pulse, the electric field can be modeled as a continuous wave:

$$\vec{E}(t) = E_0 \cos(\omega t)\hat{z}. \tag{2.2}$$

For simplicity and relevance, the following discussion in this chapter does not go out of scope of these electric fields. While there is much interest in the use of elliptically polarized electric fields (See, e.g., References [25–28]), I do not present any results from elliptically polarized lasers, and so

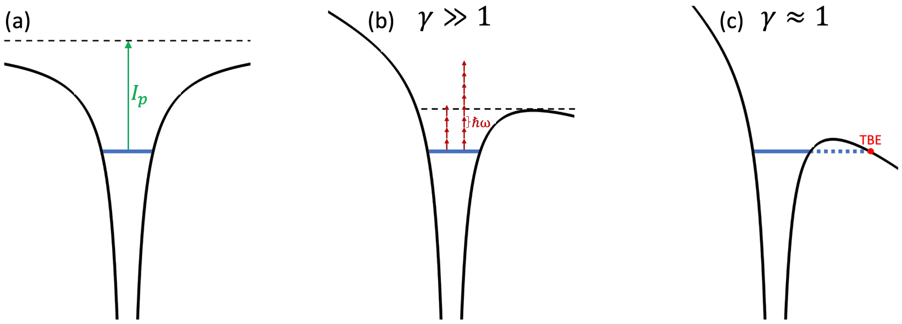

Figure 2.1: Schematic of the total electrical potential experienced by a bound electron with and without the presence of a laser field. For each panel the solid blue line indicates the weakest bound electron. (a) This drawing shows the field-free potential well, with the field-free ionization potential I_p shown as a green arrow. (b) This drawing shows a distorted Coulomb potential caused by a relatively weaker laser field for which $\gamma \gg 1$. The SFI in this regime is predominantly multi-photon. The red arrows correspond to a single laser photon of energy. (c) This drawing shows a more distorted Coulomb potential caused by a relatively stronger laser field for which $\gamma \approx 1$. The SFI in this regime is a mix of both multi-photon and tunnel ionization. The tunnel barrier exit (TBE) is shown as a red dot. Figure partially adapted from Reference [1]

.

restrict all equations and conclusions in this chapter to the linearly polarized case.

Electrons bound within an atom or molecule in its ground state will occupy distinct orbitals, with the weakest bound electron occupying the highest energy orbital. We represent this electron as an indeterminate position at a specific energy within a Coulomb potential well, as shown in Figure 2.1(a). When the atom or molecule is hit by a strong laser whose field amplitude is comparable to the magnitude of the Coulomb attraction of the nucleus, the potential well distorts. The oscillatory nature of the laser field causes the potential well to “rock” back and forth increasing the height on one side and decreasing it on the other.

By lowering the potential well barrier, the oscillatory field offers a number of ways for the electron to be ionized. First, the electron could absorb a sufficient number of photons whose combined energy exceeds the lowered barrier height above the electron’s bound state energy. This is called multi-photon ionization, and is depicted in Figure 2.1(b). When many more photons than required for ionization are absorbed by the electron, it is called above-threshold ionization (ATI), and the resulting freed electron attains distinct kinetic energies corresponding to the number of additional photons absorbed. Second, the barrier can be lowered sufficiently that the probability of tunneling through it increases to a measurable amount. This is called strong-field tunnel ionization, and can be seen in Figure 2.1(c) [1]. For certain extremely intense laser fields, the height of the barrier may fall below the energy of the electron’s orbital, causing the electron to ionize without tunneling. This is called over-the-barrier ionization, but practically speaking this form of ionization is extremely difficult to achieve using conventional laser setups, because essentially all of the electrons will either multi-photon ionize or tunnel ionize before the laser has a chance to reach the amplitude require to over-the-barrier ionize. This phenomenon is called electron depletion, which is undesirable for any of the results discussed in this thesis, and no presented experiment comes close to this laser intensity regime.

We can roughly characterize how a given laser pulse will ionize a given target by defining the Keldysh parameter of the system:

$$\gamma = \frac{\omega\sqrt{2I_p}}{E_0}. \tag{2.3}$$

Systems with a higher Keldysh parameter have a relatively higher proportion of ionization events from multi-photon ionization, while systems with a lower Keldysh parameter have relatively higher proportions of tunnel ionization events. Theorists commonly define “ionization regimes” at specific Keldysh parameter values in which they model all ionization events as coming from one type of ionization. Typically these regimes are defined as (1) Multi-photon ionization regime: $\gamma \gg 1$, (2) Tunnel ionization regime: $\gamma < 1$, (3) Over-the-barrier ionization regime: $\gamma \ll 1$ [1]. In my experience, these regimes are not easily distinguishable or realistically achievable in a practical laboratory setting using conventional ultrafast lasers. In this thesis, as in many other experimental studies of strong-field ionization [2,21,29–31], the experiments take place at $\gamma \approx 1$, such that there is a mix of multi-photon ionization and tunnel ionization processes. Despite this mixture, many of the more complicated interference structures visible in strong-field ionization experiments can be traced to sub-cycle interference which must arise from tunnel ionization.

2.1.3 Tunnel Ionization

The most widely used model of tunnel ionization in strong laser fields is Ammosov-Delone-Krainov (ADK) theory. The most important result from ADK for the purposes of this thesis is the tunnel ionization rate given by [1,32]

$$W_{ADK} = |C_{n^*l^*}|^2 \sqrt{\frac{6}{\pi}} f_{lm} I_p \left(\frac{2(2I_p)^{3/2}}{E_0}\right)^{2n^*-|m|-3/2} \exp\left(\frac{-2(2I_p)^{\frac{3}{2}}}{3E_0}\right). \tag{2.4}$$

Here, $n^* = Z/\sqrt{2I_p}$ is the effective principal quantum number where Z is the charge of the remaining ion, $l^* = n^* - 1$, and l and m are the initially bound electron's orbital and magnetic quantum numbers. The other terms in the formula are given by

$$|C_{n^*l^*}|^2 = \frac{2^{2n^*}}{n^*\Gamma(n^*+l^*+1)\Gamma(n^*-l^*)}, \tag{2.5}$$

where $\Gamma(x)$ is the Gamma function, and

$$f_{lm} = \frac{(2l+1)(l+|m|)!}{2^{|m|}|m|!(l-|m|)!}. \tag{2.6}$$

This rate equation is a model of the probability of tunnel ionization per unit time, and can be used directly to determine the probability of ionization at a given time in the laser field, with the addition of some key assumptions. In particular, we assume that the laser pulse envelope is sufficiently short, the laser frequency is sufficiently low, and the laser intensity is not overly intense such that the rate can adjust adiabatically in time and there is insignificant electron depletion between laser cycles.

All of the results in this thesis incorporate laser pulses which satisfy the above assumptions, which allows us to substitute $E_0 = |E(t)|$ in Equation 2.4. Furthermore, everywhere this rate equation comes into play in this thesis is for discussions of atomic argon, which allows us to replace the target specific quantities such as the ionization potential and quantum number terms, and achieve an ionization rate only dependent on the laser field. For singly ionized atomic argon, $Z = 1$, $I_p = 0.579$ (in atomic units), the weakest bound electron is in the $3p$ orbital and so $l = 1$, and by observation of the form of Equation 2.4 the ionization rate is maximized when $m = 0$. For atomic argon the ionization rate equation can thus be simplified numerically to

$$W_0(t) = 9.8791\left(\frac{2.4923}{|E(t)|}\right)^{0.3586} \exp\left(\frac{-0.8308}{|E(t)|}\right). \tag{2.7}$$

In later chapters, I use this equation for all quantitative results regarding the tunnel ionization rate of atomic argon.

2.1.4 Electron Trajectories

Strong-field tunnel ionization is an inherently quantum mechanical phenomenon, and thus in principle a complete picture of its resultant ultrafast electron dynamics cannot be achieved without considering electron wave packet dynamics, quantum tunneling, quantum diffusion, and quantum scattering. However, one compelling and intuitive approximation to describe these phenomena has been the electron trajectory model which bridges classical and quantum descriptions of strong-field dynamics. Succinctly stated, this model employs a quantum mechanical description for tunnel ionized electrons only during the ionization step as summarized above before tracing a classical trajectory in the laser field to describe all subsequent dynamics.

At first glance, this may seem to be a bad approximation. After all, electrons are certainly not classical objects, and it may seem unlikely that modeling their motion in the laser field classically could result in predictive insight of real strong-field processes. However, this model has been used ubiquitously in studies of strong-field ionization [33–36], and has been arguably the single most impactful insight in the history of studying strong-field ionization.

The electron trajectory model arose to prominence after Paul Corkum's 1993 formulation of the three-step model (also known as the Simple Man's Model, though many similar models also claim to share that name) [37]. In this seminal paper, Corkum proposes that experimental measurements of the features formed during high harmonic generation (HHG) experiments can be modeled very simply as quantum tunnel ionization, followed by classical acceleration in the laser field, followed by recollision and recombination with the parent ion. He assigns amplitudes to each trajectory by weighting them by the ADK ionization rate in Equation 2.4 according to the trajectory's ionization time within the laser field. Despite the simple framework, this model accurately reproduces the most prominent HHG features that are observed experimentally. In Chapter 4 I present a more detailed calculation of a similar Simple Man's Model alongside experimental data to introduce electron rescattering features clearly visible in experimental strong-field spectra, and so I do not go into further detail on this model in this chapter.

Since Corkum's work, dozens of more complicated electron trajectory models have been developed (See, for example, References [3, 18, 22–24] as a non-representative sample). In this thesis I fully embrace the electron trajectory paradigm of strong-field tunnel ionization, and express almost all results in the context of measuring the fundamental properties of observed electron trajectories. The model of electron trajectories I employ is more sophisticated than the purely classical trajectories of the Simple Man's Model, and have the following corrections:

1. Electron trajectories are semi-classical by construction and have a defined amplitude and phase advancement in the laser field. The phase advancement is given by the classical action along the trajectory:
$$\Delta\phi = S = \int_{t_0}^{t_f} dt\left(\frac{1}{2}v(t)^2 + I_p\right). \tag{2.8}$$

2. Electrons are ionized with zero velocity in the direction of tunneling and a non-zero transverse velocity, v_0, resulting from the dispersion of the electron wave packet while tunneling. This is modeled by weighting trajectories using the following transverse velocity distribution also derived from ADK [38]:
$$W_\perp(t_0, v_0) = \frac{\sqrt{2I_p}}{|E(t_0)|}\exp\left(-\frac{v_0^2\sqrt{2I_p}}{|E(t_0)|}\right). \tag{2.9}$$

3. The total amplitude, W, of any individual trajectory with given ionization time, t_0, and initial transverse velocity, v_0, is then determined solely by the ionization rate and transverse velocity weight depending on these parameters:
$$W(t_0, v_0) = W_0(t_0)W_\perp(t_0, v_0). \tag{2.10}$$

4. Using the short-range potential approximation [32, 39], electrons are ionized a distance away

from the ion at the adiabatic tunnel barrier exit:

$$z_{\mathrm{TBE}}(t_0) = \frac{I_p}{E(t_0)}, \tag{2.11}$$

which is commonly used when ADK theory is applied.

5. When multiple electron trajectories arrive at the same final position on a detector, their amplitudes sum coherently according to their phase advancement until that point.

The self-interferences caused by the semi-classical electron trajectories acquiring phase and summing coherently result in intricate interference fringes when measured. Within these interference patterns the full range of ultrafast electron dynamics are encoded, along with the structure of the residual ion and the laser field itself. The study of these interference patterns in order to extract useful information of the system is called photoelectron holography [3], which is the primary focus of this thesis.

2.1.5 Photoelectron Holography

In order to describe the motivation for photoelectron holography, it is useful to understand the principle behind holography, from which its name is derived. Ordinary light-based holography is the imaging technique centered around creating three-dimensional reconstructed images of an object from measurements of the amplitude and phase of coherent light waves reflected from that object. It roughly follows the following procedure:

1. Split a coherent light source into two beams.
2. Send one beam directly at a detector (the reference beam), and other towards the target (the signal beam).
3. Record the interference pattern generated from the phase difference arising from the path-length difference of the two beams. This is the desired hologram.
4. Using the hologram and known properties of the reference beam, the wavefront of the signal beam can be reconstructed, retrieving the three-dimensional information stored in the hologram's interference structure.

Depending on context, it is correct to label both the recorded interference pattern and the three-dimensional reconstructed image as a hologram. To avoid any ambiguity, in the following I use the term “hologram” only in reference to an interference pattern.

The theory behind photoelectron holography follows a similar outline to ordinary holography, except it uses the naturally diverse strong-field ionized electron trajectories as a substitute for the splitting of the coherent light source. Because electron trajectories are coherent and have different degrees of interaction with the ion subsequent to ionization, the interference structures characteristic of strong-field ionization are similar to the recorded holograms in holographic imaging, and can be used to probe the dynamics of the ionization system. It is for this reason that these strong-field interference patterns are commonly referred to as electron holograms.

When photoelectron holography was first conceived, there was a clear distinction between which trajectories could be related to the reference beam and which could be treated as the signal beam [3]. Electrons ionized when the laser electric field is increasing in magnitude tend to be accelerated directly away from its parent ion, and thus re-interact with the ion only minimally. Collecting these electrons is then akin to measuring the reference beam that is sent directly at a detector in ordinary holography. On the other hand, electrons ionized when the laser field is decreasing in magnitude tend to be driven back to the parent ion, and undergo many possible types of rescattering dynamics depending on the specific trajectories they are following. These electrons are akin to the reflected beam of light in ordinary holography, and thus were labeled the signal beam. It can still be useful to think about strong-field holograms using this language, but I and many other physicists studying photoelectron holography embrace a more relaxed definition which treat on equal footing the interference between any two trajectories, regardless of their level of interaction with the target [3].

In strong-field photoelectron holography, holograms can be divided into two types: intercycle and subcycle. Intercycle holograms form from the periodic interference of electron trajectories whose ionization times are separated by one or more full cycle of the driving laser. Generally speaking, intercycle holograms are of little interest, because they involve repeated trajectories, which do not separately encode information of the target and so can only be used to measure properties of the laser field itself. Furthermore, intercycle holograms tend to obfuscate more interesting features, and so can be actively detrimental for holographic studies. In contrast, subcycle holograms form from the interference of electron trajectories whose ionization times are all contained within a single laser cycle. These are of broad interest to the holography community, and great effort has been taken to isolate and characterize them [2, 3, 17, 21, 36, 40].

2.2 Calculations of Strong-Field Ionization from Theory

In the coming chapters, I make frequent direct comparisons of experimentally measured strong-field ionization spectra to calculations derived from theory. Depending on the context and the desired insight to be derived from the comparison I use one of two models. For the most faithful imitations of experimental spectra I employ calculations derived from the time-dependent Shrödinger equation (TDSE) solver Qprop [41]. When more physical insight is required, I employ calculations derived from the Coulomb Quantum Orbit Strong Field Approximation (CQSFA) which individually simulates predefined classifications of electron trajectories [18, 22].

2.2.1 Qprop

Qprop is a software package which performs the direct integration of the TDSE for arbitrary initial Hamiltonians. In particular, it has been optimized for spherically symmetric initial systems perturbed by a high-intensity laser [41], making it particularly suitable for studies of strong-field ionized argon. Direct integration of the Shrödinger equation yields calculations with very high numerical agreement to experimental results; however, they lack the interpretability that many other models provide.

Because of the lack of insight provided by the model, going into further details of the operation and calculative techniques of this software is not particularly useful for the interpretation of comparisons between experiment and Qprop calculations. If further details are desired, see the original documentation for Qprop in Reference [41]. However, these calculations serve a useful purpose for the conclusions drawn in this thesis. Strong agreement between an experimentally measured spectrum and a Qprop calculation is a way to directly measure properties of the laser field. In particular, the intensity of a laser, which is a critical quantity determining all of the quantities expressed thus far in this chapter, is very difficult to accurately determine using measurable quantities in the lab. Because generating a Qprop calculation requires one to explicitly record parameters of the laser within the defined Hamiltonian, directly comparing many calculations initialized with a range of input parameters and finding the best match to the experiment is a decent way to determine these quantities to reasonable accuracy. In Chapter 5 I show how Qprop calculations can be analyzed to relatively simply determine the intensity of a strong-field ionizing laser.

2.2.2 Coulomb Quantum Orbit Strong Field Approximation

The CQSFA is referred to much more frequently than Qprop and can be used to derive significant insight into the underlying dynamics of strong-field ionization, and so I go into much more detail of its formulation.

The CQSFA is an electron trajectory model of strong-field ionization which fully accounts for the Coulomb potential of the residual ion when calculating both the acquired phase and the dynamics of specified classes of trajectories. The version of CQSFA used in future chapters defines four unique classes of trajectories [33], which are called quantum orbits in the formalism. These are clearly defined and analyzed alongside an experimental measurement in Chapter 4.

We initialize our system with the full Hamiltonian

$$H(\mathbf{r}(t), \mathbf{p}(t), t) = \frac{1}{2}\left[\mathbf{p}(t) + \mathbf{A}(t)\right]^2 + V\left(\mathbf{r}(t)\right), \tag{2.12}$$

where $\mathbf{A}(t)$ is the vector potential of the laser and $V\left(\mathbf{r}(t)\right)$ is the value of the Coulomb potential of the target at coordinate $\mathbf{r}$ from the center of the target. In this thesis, $V\left(\mathbf{r}\right)$ is modeled by an effective potential provided in References [42,43].

The CQSFA begins from the transition amplitude $M(\mathbf{p}_f) = \langle \psi_{\mathbf{p}_f} | U(t, t_0) | \Psi_0(t_0) \rangle$, which details the propagation of the initial single active electron wave function $|\Psi_0(t_0)\rangle$ from its ionization time, t_0, to a later time, t, and eventually arriving a final state which has momentum $\mathbf{p}_f$. This quantity can be calculated exactly [18, 22] from Feynman path integrals [44] and the saddle point approximation to yield [45]

$$M_s(\mathbf{p}_f) \propto -i \lim_{t \to \infty} \det \left[\frac{\partial \mathbf{p}_s(t)}{\partial \mathbf{r}_s(t_s)} \right]^{-1/2} \mathcal{C}(t_s) e^{iS(\mathbf{p}_s, \mathbf{r}_s, t, t_s)}, \tag{2.13}$$

where the semi-classical action $S(\mathbf{p}, \mathbf{r}, t, t')$ is given by

$$S(\mathbf{p}, \mathbf{r}, t, t') = I_p t' - \int_{t'}^{t} [\dot{\mathbf{p}}(\tau) \cdot \mathbf{r}(\tau) + H(\mathbf{r}(\tau), \mathbf{p}(\tau), \tau)] d\tau \tag{2.14}$$

and the prefactor $\mathcal{C}(t_s)$ is given by

$$\mathcal{C}(t_s) = \sqrt{\frac{2\pi i}{\partial^2 S(\mathbf{p}_s, \mathbf{r}_s, t, t_s)/\partial t_s^2}} \langle \mathbf{p} + \mathbf{A}(t_s) | \hat{H}_I(t_s) | \Psi_0 \rangle, \tag{2.15}$$

where $\hat{H}_I(t) = \hat{\mathbf{r}} \cdot \mathbf{E}(t)$ is the interaction Hamiltonian. An additional factor of $-\pi/2$ must be added to the semi-classical action in Equation 2.14 to account for Maslov phase shifts whenever the momentum transverse to the polarization direction changes sign [46]. The proportionality in Equation 2.13 indicates an overall normalization must be computed at the end of the calculation to arrive at a proper probability distribution.

In these equations there is an index s which denotes each quantum orbit predefined in these calculations. These orbits are determined by solving the saddle point equations, which are given by

$$[\mathbf{p}_s(t_s) + \mathbf{A}(t_s)]^2/2 + I_p = 0, \tag{2.16}$$

which determines the ionization times of each orbit, and the pair of equations

$$\dot{\mathbf{p}}_s(t) = -\nabla_{\mathbf{r}} V[\mathbf{r}_s(t)] \quad \text{and} \quad \dot{\mathbf{r}}_s(t) = \mathbf{p}_s(t) + \mathbf{A}(t), \tag{2.17}$$

which describe how the orbit propagates after ionization [3, 17–20, 22, 42, 47, 48].

When calculating a strong-field spectrum using the CQSFA, a grid of final momenta $\mathbf{p}_f$ is selected, and the transition amplitude $M_s(\mathbf{p}_f)$ is calculated at each grid point for each quantum orbit. In a strong-field experiment, a detector measures the probability distribution of electrons arriving at these final momenta grid points, which in the CQSFA is given by the squared amplitude of transition amplitudes:

$$\text{Prob}(\mathbf{p}_f) = \sum_{i,j} M_i(\mathbf{p}_f) \overline{M_j(\mathbf{p}_f)}. \tag{2.18}$$

Here we are indexing any pair of trajectories with the sum over i and j. Importantly, the transition amplitudes are still separated into distinct terms for each orbit at this point, allowing us to selectively turn them off and on as desired by omitting any combination of orbits from the final summation. The basic CQSFA formalism ends here, but further corrections can be made to Equation 2.18 to better model strong-field experiments or to perform selective analysis.

Unit-Cell Averaging

In order to avoid the problem of intercycle holograms rapidly overwhelming more subtle holographic features, CQSFA calculations commonly restrict the range of allowed ionization times in Equation 2.16 to a single laser cycle, which is called a unit cell. This is an effective tool to eliminate intercycle holograms, but has a major shortcoming. If the end points of the restricted unit cell are arbitrarily selected, then non-physical asymmetries will develop in the final probabilities due to improperly defining an ordering for electron trajectories.

Any given electron trajectory has a specific ionization phase within a single unit cell; however all trajectories are periodic given a 2π ionization phase shift in either direction. When the ionization phase of a trajectory is shifted by 2π, it acquires an additional phase given by

$$\Delta S = \frac{2\pi}{\omega}\left(I_p + U_p + \frac{1}{2}\mathbf{p}_f^2\right), \tag{2.19}$$

where $U_p = \left(\frac{E_0}{2\omega}\right)^2$ is the ponderomotive energy of the electron in the laser field. All strong-field holograms are formed from the pair-wise interference of two or more electron trajectories. However, when looking at any given hologram formed from trajectories A and B, there is an uncertainty over whether A precedes or follows B. Of course in a real system the answer is both, as either trajectory's ionization phase can be shifted by $2\pi k$ radians in both directions for any integer k. Thus every pair of trajectories forms *two* distinct holograms which overlap at each final momentum.

We can account for this in the CQSFA formalism by incorporating the extra phase ΔS from Equation 2.19 in an integral over all possible unit cells with boundaries $[0+\phi, 2\pi+\phi]$:

$$I_\phi = \frac{1}{2\pi}\int_0^{2\pi} d\phi \exp\left[i\Delta H_{ij}(\phi)\Delta S\right], \tag{2.20}$$

where we have defined a quantity ΔH_{ij} given by

$$\Delta H_{ij}(\phi) = H(\phi - \omega t_i^{\mathrm{Re}}) - H(\phi - \omega t_j^{\mathrm{Re}}), \tag{2.21}$$

where $H(x)$ is the Heaviside function and ωt_i^{Re} and ωt_j^{Re} are the ionization phases for trajectories indexed by i and j respectively. This integral appropriately adds or subtracts the contribution ΔS to the phase as the unit-cycle boundaries shift over the ionization phases of the trajectories indexed by i and j.

With this phase correction in hand, we modify Equation 2.18:

$$\mathrm{Prob}(\mathbf{p}_f) = \sum_{i,j} M_i(\mathbf{p}_f)\overline{M_j(\mathbf{p}_f)} I_\phi, \tag{2.22}$$

which now correctly accounts for the time ordering of trajectories within any arbitrary unit cell. This correction is called unit-cell averaging, and it is applied to every CQSFA calculation shown in this thesis.

Time-windowing

The time-windowing modification described in this section is a very specific application of CQSFA which allows us to gate holographic interference formed from the coherent sum of trajectories whose ionization times are separated by more or less than a defined parameter, $\tau \in (0, 1)$. In spite of the previous discussion of unit-cell averaging in which we corrected for improper time ordering of trajectories, selecting a value of $\tau \neq 0.5$ allows us to analytically restrict the calculation to a specific ordering of trajectories, which are useful when we discuss time correlations in Chapter 5.

We apply selective use of Heaviside functions to add yet another correction to the integral in Equation 2.20:

$$I_\phi(\tau) = \frac{1}{2\pi}\int_0^{2\pi} d\phi \exp\left[i\Delta H_{ij}(\phi)\Delta S\right] H(2\pi\tau - |\omega\Delta t_{ij} + 2\pi\Delta H_{ij}(\phi)|). \tag{2.23}$$

Here we introduce $\Delta t_{ij} = t_i^{\mathrm{Re}} - t_j^{\mathrm{Re}}$ which is the time separation between the ionization of the two trajectories and τ is the time separation parameter in units of laser cycles. This correction allows us to select any value for τ and examine the set of all CQSFA orbits with launch time separations lower than τ cycles individually for each selected grid point $\mathbf{p}_f$. To modify the correction to observe the set of orbits with separations higher than τ cycles, simply multiply the argument of the Heaviside function by -1. This integral can then be used in Equation 2.22 to determine the probability at each grid point.

Chapter 3

Experiment Designs and Methods

This chapter focuses primarily on the tools required to replicate the experiments analyzed further in the thesis. I focus most heavily on the custom optical setups and the software I wrote to perform each of the experiments. Aspects of the experiments which are not custom built and are heavily documented elsewhere receive only a brief overview, and I encourage the reader to seek further explanation of these components from the referenced sources.

3.1 Overview

The experiments performed in this thesis all follow the block diagram shown in Figure 3.1. This is a tabletop setup, composed primarily of commercial laser hardware and optomechanics. The optomechanics in the block labeled "Variable Optical Setup" is further examined in Section 3.4 below.

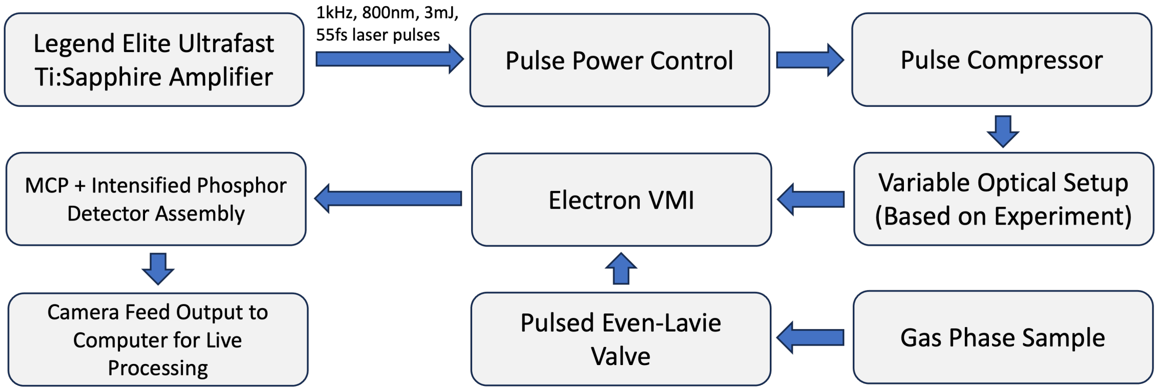

Figure 3.1: The general block diagram outline of the experiments performed in this thesis.

3.2 The Laser

The laser pulses are generated from a commercial ultrafast laser system: Coherent Legend Elite Ultrafast Ti:Sapphire Amplifier. This system outputs linearly polarized 800 nm laser pulses at a 1 kHz repetition rate. Typically, we can achieve a maximum of 3 mJ of pulse energy with a pulse duration of 55 fs. These pulses are throttled down to values suitable for strong-field experiments by a polarizing power throttle as described in Section 3.4.1. The typical range of acceptable pulse energies in a strong-field ionization experiment is between 30 μJ and 300 μJ. After the pulse throttle, the pulses are re-compressed to a final pulse duration of approximately 40 fs. A focusing mirror within the final vacuum chamber (see Figure 3.2) focuses the beam to approximately 50 microns in diameter, yielding final ionizing intensities between approximately 40 TW/cm^2 and 400 TW/cm^2.

3.3 The Velocity Map Imaging Spectrometer

The most important piece of experimental equipment I use is a velocity map imaging spectrometer (VMI) [49]. Succinctly, a VMI is a tool which extracts charged particles within an interaction region and uses a carefully calibrated set of electrostatic lenses to focus these particles onto a detector such that the particles' cross-sectional momenta in the interaction region are mapped one-to-one with their final position on the detector.

VMI spectroscopy is used ubiquitously throughout ultrafast AMO physics [2, 21, 36, 40, 49–55]. As a tool it is extremely versatile. It has been used to image molecular dissociation [13, 55], to time laser pulses with attosecond precision using attosecond streaking [56], and to image ultrafast electron tunneling dynamics [2, 21, 36], to name just a few. As there are many references detailing the construction [49], operation [52], and uses of VMIs [51–53], I mostly focus on the specifics of the VMI system that I used to generate the data examined in this thesis.

A cartoon schematic of the VMI spectrometer system is shown in Figure 3.2. This VMI consists of three electrostatically charged plates: the repeller, extractor, and an additional lens. The incoming laser is focused into the interaction region midway between the repeller and the extractor, and intersects an ultra-cold pulsed gas jet emitted from an Even-Lavie [57] (EL) valve (see below). The repeller and extractor are set to generate an electric field which pushes photoionized electrons downwards toward the detector assembly. The electrostatic simulation software Simion [58] was used to calibrate the voltages of the repeller and the extractor to determine the VMI focusing condition [49] for this spectrometer. The third electrostatic lens contracts the resulting spread of photoelectron momenta, allowing higher energy electrons to be captured on the surface of the detector. The voltage setting on this lens has no impact on whether the VMI condition is maintained, assuming it is set to a value between the voltage on the extractor and the voltage on the front of the detector. All three of these voltages are freely tunable to measure different ranges of final electron energy, but for all of the experiments shown in this thesis, the voltages are as labeled in Figure 3.2. With these voltages

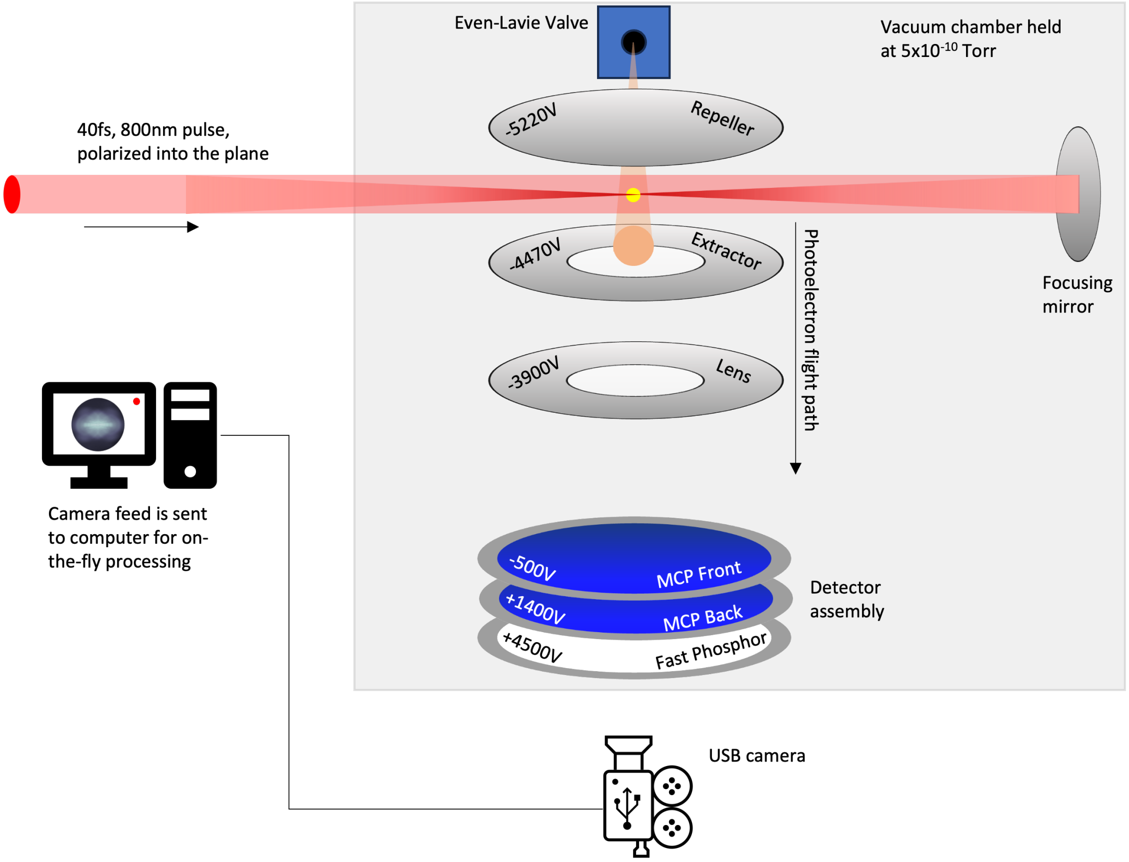

Figure 3.2: A schematic of the VMI spectrometer during the collection of data.

simulated in Simion, electrons with cross-sectional momenta up to approximately 2.3 a.u. can be measured on the detector.

The entire VMI spectrometer is contained within an ultra-high vacuum chamber held at approximately 5×10^{-10} Torr. A turbomolecular vacuum pump is located directly opposite to the path of the pulsed gas jet, which allows the jet to exit the chamber without significantly increasing the background pressure.

3.3.1 Detector Assembly

The detector assembly consists of a pair of chevron microchannel plates (MCPs) and a fast phosphor (P47) screen [59]. The MCP stack serves as an electron multiplier, converting individual electron impacts into electron cascades with orders of magnitude more electrons, which can illuminate the phosphor screen. Electrons impacting the front MCP plate enter one of the densely packed small channels which are bored at an angle with respect to the normal of the plate. When the electrons impact the side of the channel, they cause additional electrons to be knocked off which are accelerated by the voltage differential between the two MCP plates. These additional electrons impact the channel sides at other locations as they travel through them, each causing additional electrons to be released. After leaving the front plate, the electron cascade enters several channels on the back plate, each of which cause additional cascades. The chevron alignment of the plates help maintain the position of the final electron cascade relative to the impact position of the original electron. Together the two plates have a combined gain of approximately 1×10^4.

The final amplified electron signal is emitted from the back of the back plate onto the phosphor screen, which emits a flash of light. Increasing the voltage gradient between the back MCP plate and the phosphor screen increases the brightness of the flashes of light, but also increases the size of the flash, which may adversely affect the precision of a data collection if the incoming photoelectrons are not sufficiently sparse. A camera operating at 200 Hz with a 2 ms exposure time continuously records the underside of the phosphor screen, and transmits the live feed via a USB cable to a nearby computer. The live feed is processed in real time, as described below in Section 3.5.

3.3.2 Sample Delivery

Pulsed jets of ultra-cold gas-phase atoms and molecules are emitted by an EL valve [57] housed in a separate ultra-high vacuum chamber (the source chamber) and diffused through a skimmer connecting the two vacuum chambers before reaching the interaction region. An EL valve contains a piezoelectric actuator which can rapidly open and close the nozzle of the valve in accordance with a programmable incoming voltage signal. The maximum repetition rate of the EL valve is 500 Hz. A digital delay generator (Stanford Research Systems DG645) is used to precisely time the arrival of the gas jet with the laser pulses in the interaction region, which is determined by stepping the delay in 1 μs increments until the maximum electron count rate is measured on the detector. Because

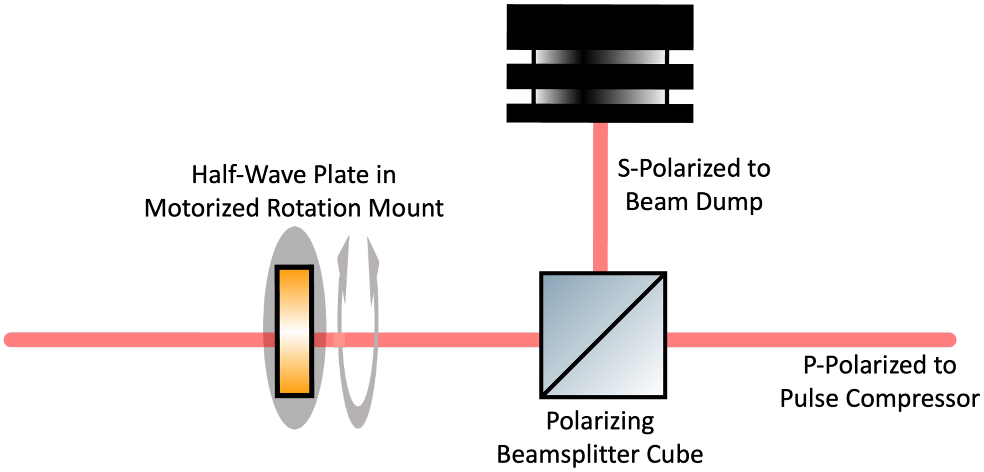

Figure 3.3: Schematic of the simple power throttle for controlling the intensity of laser pulses.

the laser repetition rate is 1 kHz, the EL valve is triggered such that it intersects every other laser shot. Due to the vacuum set-up of the chamber, the laser shots which miss the pulsed gas jet do not generate a significant amount of background electrons, which is verified by detuning the digital delay generator. The gas prior to the valve is held at a high backing pressure (125 psi) which causes it to enter the source chamber at a high velocity, before rapidly expanding and cooling in the high vacuum environment. In a typical experiment, the opening time of the EL valve is set to 20 μs, but may be adjusted to change the electron count rate without changing the laser intensity in the interaction region.

3.4 Experiment Design

In this section I go into significant detail on the optomechanical components of the experiment designs I discuss in later chapters of this thesis. These descriptions are written to be of specific interest to someone attempting to replicate these experiments, and where appropriate contain instructional information on how to arrange and manipulate optics on the laser table to facilitate the experiment.

3.4.1 Pulse Power

In order to adjust the intensity of the laser after the output of the Legend Elite system, I employ a simple polarizing power throttle, shown schematically in Figure 3.3. The full power, uncompressed P-polarized beam is sent through a half-wave plate mounted in a motorized rotation mount that can be programmatically set to any rotation angle to output any linear polarization. After the half-wave

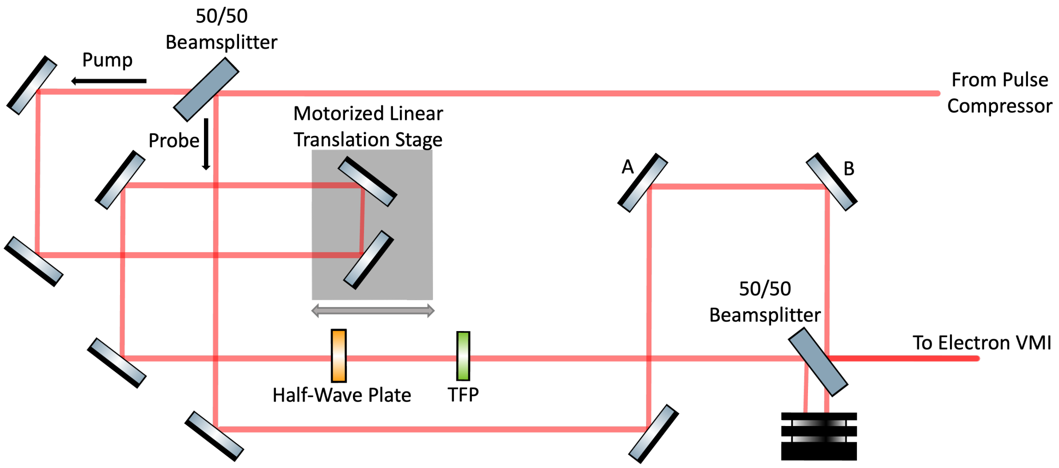

Figure 3.4: Schematic of the optical delay line used for pump-probe experiments requiring two 800 nm pulses.

plate, a permanent polarizing beam cube reflects components of the beam with S-polarization into a beam dump, and transmits the remaining P-polarized beam onwards through the pulse compressor. This throttle must be placed prior to the pulse compressor, as the polarizing beam cube adds a significant amount of dispersion to the beam, which must be corrected for before passing the beam to the VMI chamber.

3.4.2 Two-Pulse Delay

When performing experiments that require two 800 nm pulses in a pump-probe type setup, I have constructed an interferometer with an optical delay line on one of the two paths, as shown in Figure 3.4. The setup actually used is faithfully recreated here; however, this design uses more mirrors than required in order to fit into a constrained area on the optical table while maintaining roughly equal path lengths in each arm of the interferometer. The motorized translation stage with a mounted retroreflector in the pump arm sets the delay between the pulse in each arm. This stage has a very wide range of approximately 500 ps of optical delay. This level of delay is useful for probing certain rotational motions, such as impulsive molecular alignment. These types of experiments tend to have a relatively less intense pump pulse and a very intense probe pulse, so a manual power throttle similar to the overall pulse power control consisting of a half-wave plate and a thin film polarizer (TFP) is placed in the pump arm to reduce the pump laser's pulse intensity. This is also why the probe arm contains no transmissive optics, to maximize its intensity, while the pump arm contains several. Since the incoming pulse durations are upwards of 40 fs, the extra dispersion from passing through glass optics can be mostly ignored.

The overlap of the two beams in the VMI is controlled using the mirrors labeled "A" and "B" and the final beamsplitter. First the pump is directed through the final beamsplitter and aligned to maximize the resulting electron count rate when directed into the chamber with the probe blocked. Then the mirrors "A" and "B" are used to steer the probe roughly collinear to the pump after the beamsplitter, with particular emphasis on good overlap immediately after the beamsplitter. The motorized stage is translated until the two beams generate an interference pattern on fluorescent viewing paper, which indicates rough temporal overlap has been achieved. The thickness of the interference fringes indicates how close to true spatial overlap the two beams are, with thicker fringes indicating more overlap. The orientation of the fringes also indicates which direction the two beams are offset in. The final beamsplitter can be very finely adjusted to maximize the thickness of the fringes until spatial overlap is achieved far down the optical table. After this point, iterating between finely adjusting mirror "B" for spatial overlap immediately after the beamsplitter, and finely adjusting the beamsplitter for far spatial overlap yields close collinearity. Lastly the true temporal overlap position of the translation stage is determined by maximizing the count rate of electrons in the VMI chamber with fine stage-translation steps.

3.4.3 Relative Phase between Fundamental and Second Harmonic

Some of the rather more interesting experiments to probe electron holograms involve using multiple laser frequencies. Overlapping two or more different frequency pulses yields a summed pulse with both frequency components, which changes the resulting ionization pattern. For experiments involving a combined fundamental 800 nm (ω) and second harmonic 400 nm (2ω) perturbation, I use the collinear optical setup shown in the upper arm of Figure 3.5. Here, the 2ω light is generated in a beta barium borate (BBO) crystal, which has been cut such that type I phase matching occurs at normal incidence. The conversion efficiency of this process is around 2% (meaning that the output 2ω pulse power is 2% of the incoming fundamental pulse power). As is standard for type I phase matching processes, the 2ω pulse is polarized perpendicular to and delayed with respect to the fundamental, and so birefringent calcite crystals are used as time delay compensators (TDCs) to achieve temporal overlap in the VMI chamber. As shown in Figure 3.5, there are two TDCs in this optical setup. The first (labeled "1" in Figure 3.5) controls the $\omega/2\omega$ temporal overlap at the third harmonic crystal in the lower arm, and is discussed in the next section, and the second (labeled "2") controls the final temporal overlap between the 800 nm light and the 400 nm light in the interaction region of the chamber.

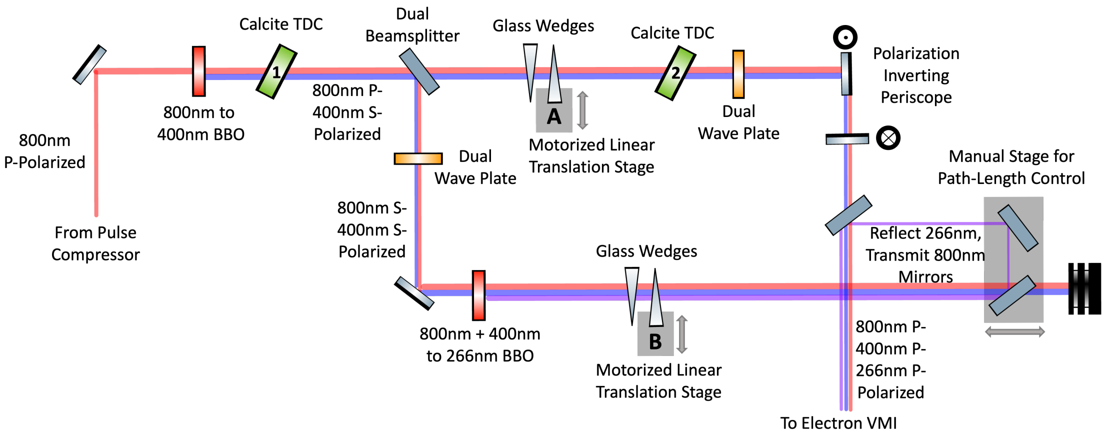

Figure 3.5: Schematic of the optical beamline used for experiments involving overlapped second and third harmonics of the 800 nm fundamental laser. The 800 nm, 400 nm, and 266 nm laser pulses are shown in red, blue, and violet respectively. These pulses are spatially overlapped and the offset of the colors here is for illustration only. Experiments involving only the second harmonic use only the upper path after the dual beamsplitter, and block the lower path after the beamsplitter.

Because the usage of the calcite crystals as TDCs was not obvious to me when I began this work, and I could not find a simple written explanation for how to mount and manipulate the optic on the laser table to achieve the delays labeled on the optics when they were purchased, I include an explanation here.

Birefringence is the phenomenon in which light traveling through a crystal experiences a different index of refraction depending on its direction of propagation and its polarization. The simplest birefringence is uniaxial, in which a three dimensional crystal has exactly one axis (the optical axis) such that light traveling parallel to this axis has the same index of refraction (n_o) regardless of its polarization. If we send in a ray of light perpendicular to the optical axis, then there is one linear polarization of the light perpendicular to the optical axis which will also have index of refraction n_o, whereas the polarization parallel to the optical axis will have a different index of refraction, n_e. A ray of light propagating through this crystal is labeled as ordinary if it experiences the index n_o, and is labeled as extraordinary if it experiences the index n_e. If a light ray has a polarization between these extremes, it will experience an index of refraction between n_o and n_e.

For a uniaxial time delay compensator, such as calcite, we can use these differing indices of refraction to change the relative timing of laser pulses with polarizations perpendicular to each other. The first step in doing this is to identify which laser pulse we want to be the ordinary ray, and which we want to be the extraordinary ray. Calcite is a negative uniaxial crystal, which means that $n_e < n_o$, and so the ordinary ray is delayed more than the extraordinary ray. For second harmonic generation (SHG) the 2ω pulse is delayed with respect to the fundamental within the BBO, and so we want to delay the fundamental in the TDC, and thus the fundamental pulse should be our ordinary ray. For birefringent crystals that are positive uniaxial and thus $n_o < n_e$, one would want to set the 2ω pulse as the ordinary ray.

In order to control the amount of delay the TDC introduces to the pulses, the angle between the extraordinary ray and the extraordinary axis can be changed, smoothly transforming the index of refraction it experiences between n_e and n_o. The ordinary ray must be unaffected by this rotation, and so the optic must be turned about the axis defined by the ordinary polarization, which may involve assembling a slightly more complicated optical mount that allows rotations into the plane of the optical table. The delay values quoted by the manufacturer typically cover a rotation angle of around ±12°, but greater or lesser delays can be achieved with larger rotations if the beam width is small enough to clear the optic.

Polarization control of the laser pulses is achieved using a zero-order dual-wavelength wave plate, which acts as a full-wave plate for the 400 nm pulse and a half-wave plate for the 800 nm pulse. This optic leaves the polarization of the 400 nm pulse unchanged at any orientation, while providing control of the 800 nm polarization. Since the 800 nm pulse entering the SHG BBO is P-polarized, the emitted 400 nm pulse is S-polarized. The two useful orientations of the dual-wave plate either leaves these polarizations unchanged (labelled state 1) or rotates the 800 nm pulse to S-polarization (labelled state 2). After the dual wave plate, a polarization inverting periscope rotates the polarizations of both pulses by 90° resulting in S-polarized 800 nm light and P-polarized 400 nm light in state 1, or both pulses P-polarized in state 2. The polarization is not further affected by any downstream optics.

In the interaction region of the VMI chamber, both pulses must be polarized parallel to the optical table (P-polarized) to collect electron spectra; however, alternative polarization states are necessary for calibrating experimental parameters. In particular, in order to verify the time overlap of the 800 nm and 400 nm, I insert another SHG crystal at the output of the optical setup to generate another 400 nm pulse. Additionally, a spare VMI chamber window is inserted before the second SHG crystal to account for the extra layer of glass the pulses pass through when entering the vacuum chamber on the way to the interaction region. Two overlapping laser pulses of the same frequency and polarization generates an interference pattern when they are temporally overlapped. Thus, we require that the 800 nm pulse entering the second SHG crystal be S-polarized (state 1), in order to generate a second P-polarized

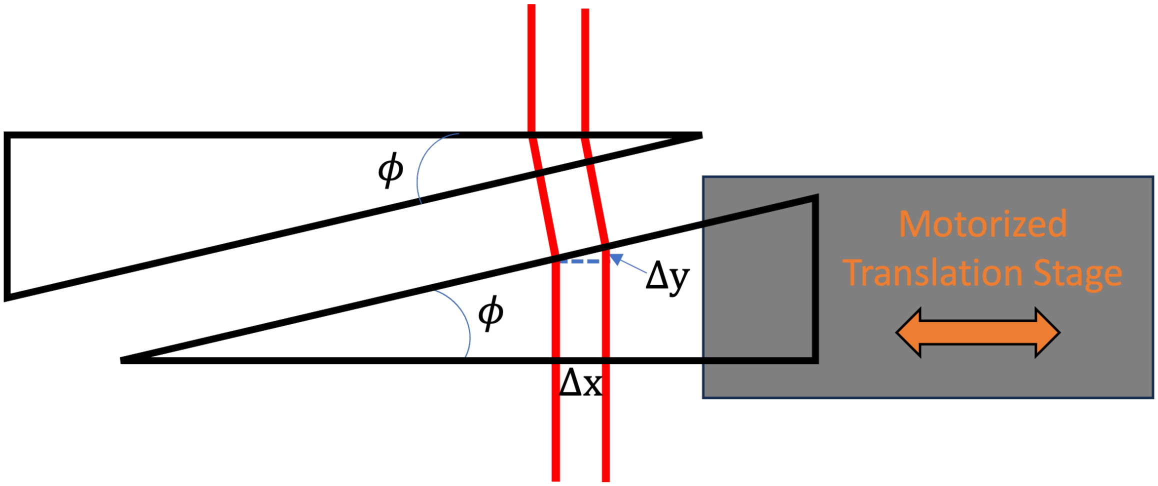

Figure 3.6: Schematic of the pair of glass wedges used to adjust the relative timing of the laser pulses. A high precision motorized translation stage moves one of the wedges while the other is stationary and mounted to the optical table. When the stage inserts the wedge a distance Δx, the optical thickness of the wedge increases by $\Delta y = \Delta x \tan(\phi)$. For the wedges used in this setup $\phi = 2.8°$. The schematic shown uses a much larger angle ϕ for illustrative purposes, and the deviation of the beam is negligible as the glass insertion varies.

> 400 nm pulse to match the first one. The overlapping spot is then observed on fluorescent paper and the second TDC is rotated until an interference pattern is observed. The TDC is then tuned to the angle which causes the most contrast in the resulting interference pattern. This temporal overlap between the two 400 nm pulses guarantees temporal overlap between the 800 nm pulse and the first 400 nm pulse, and so once the second TDC is set to the correct value the second SHG crystal and the spare chamber window is removed. The dual waveplate is then rotated to the position in state 2, such that both pulses are overlapping with P-polarization in the interaction region and electron spectra can be produced.

One final optical element is used in this design to control the relative phase between the 800 nm pulse and the 400 nm pulse: an adjustable pair of thin fused silica wedges. A pair of thin wedges can be used to change the delay between two different frequencies of light, without otherwise altering the Poynting vector of the incoming beam. The geometry of these wedges is shown in Figure 3.6, with interior wedge angle $\phi = 2.8°$. One of the two wedges is mounted stationary on the table, and the other is mounted on a high precision motorized translation stage (labeled "A" in Figure 3.5). In the figure, the insertion of the stage by Δx is indicated by a second offset laser beam, which acquires an additional $\Delta y = \Delta x \tan\phi$ more glass on its path. The additional glass in the laser's path results in the laser pulse being delayed in accordance with its wavelength dependent index of refraction in glass. During an experiment, the electron spectrum is collected at varying glass wedge insertions amounts, and thus the relative phase between the two colors are modified accordingly. We can calculate the laser pulse phase change with respect to glass insertion according to the following:

$$\frac{\Delta\theta}{\Delta x} = \tan(\phi)\frac{n_\lambda - 1}{\lambda}, \tag{3.1}$$

where $\Delta\theta$ is the phase change in units of cycles, Δx is the amount of glass inserted in the path of the laser, ϕ is the inner wedge angle, n_λ is the wavelength-dependent index of refraction of the wedge, and λ is the wavelength. These values are shown in Table 3.1.

λ **(nm)**	$\mathbf{n}_\lambda$	$\boldsymbol{\Delta\theta/\Delta}$**x (cycles/**$\mu$**m)**
800	1.4533	0.0277
400	1.4701	0.0575
266	1.4997	0.0919

Table 3.1: This table summarizes the effect of inserting a glass wedge with the geometry shown in Figure 3.6 in the path of a laser pulse. The final column indicates the number of cycles the laser pulse advances per micrometer of inserted glass, as calculated from Equation 3.1.

Since this glass wedge is inserted into the path of both the 800 nm pulse and the 400 nm one, it is more useful to examine the relative phase between the pulses. This is especially true because the absolute timing of the laser pulses is not recorded at any step in the experiment, and so the arrival of the 800 nm pulse is our reference clock. The relative phase advancement can be calculated from Equation 3.1, but it is clearer if we first convert from units of cycles, which are wavelength-dependent, to units of femtoseconds, which are not. Then equation 3.1 reads:

$$\frac{\Delta t}{\Delta x} = \tan(\phi)\frac{n_\lambda - 1}{c}. \tag{3.2}$$

where $c = f\lambda = \lambda/T$ is the speed of light, where T is the period of the pulse. If we treat our 800 nm pulse as stationary and want to measure the relative phase advancement of the 400 nm pulse, we subtract the time advancement of the 800 nm pulse from the time advancement of the 400 nm pulse and then divide by the period of the 400 nm pulse:

$$\frac{\Delta\theta_{\lambda_1,\lambda_2}}{\Delta x} = \tan(\phi)\frac{n_{\lambda_1} - n_{\lambda_2}}{cT_1} = \tan(\phi)\frac{n_{\lambda_1} - n_{\lambda_2}}{\lambda_1}. \tag{3.3}$$

Here, $\Delta\theta_{\lambda_1,\lambda_2} = \Delta\theta_{\lambda_1} - \Delta\theta_{\lambda_2}$ is the relative phase between pulses indexed by subscripts 1 and 2, measured in units of cycles of pulse 1.

In a typical relative phase experiment, I want to advance over a range of at least one full cycle of the 400 nm perturbation. To determine the glass insertion needed for this we simply compute the reciprocal of the right side of Equation 3.3, which is the glass insertion per single cycle of relative phase. Calculating this quantity using the values in Table 3.1 for $\lambda_1 = 400$ nm and $\lambda_2 = 800$ nm yields

$$\frac{\Delta x}{\Delta\theta_{\lambda_1,\lambda_2}} = 486.8\ \mu\text{m/cycle}. \tag{3.4}$$

Thus to achieve a relative phase of at least one cycle of the 400 nm perturbation the glass wedge must be inserted a minimum of 486.8 μm. It's important to recognize that this value (and others similarly calculated in this section) is the calculated value only, and in an actual experiment the true value of relative phase advancement is determined empirically from the data as accurately as possible. However, the calculated values are useful for understanding the approximate range to scan the linear translation stages and verify that the resulting signal's periodicity is roughly consistent.

I'd like to make a note on sign conventions here. A positive phase change, $\Delta\theta$, refers to the laser pulse being delayed, such that it takes longer for a selected peak within the wave to arrive at the interaction region. When glass is inserted in the path of multiple laser pulses of different wavelength, the lower wavelength one is delayed more than the higher wavelength one, and thus $\frac{\Delta\theta_{\lambda_1,\lambda_2}}{\Delta x} > 0$ if $\lambda_1 < \lambda_2$. This means that in plots where the time in laser cycles is plotted on the horizontal axis, increasing the insertion of glass results in the shorter wavelength pulse shifting to the right (in the direction of positive time) with respect to the longer wavelength one.

In order to correct for the optical Gouy phase shift acquired when focusing the final beam into the interaction region, a final iris is placed just before the entrance to the VMI chamber. By irising the 800 nm and 400 nm pulse down to the same width, it is ensured that both pulses experience the same shift. See References [46] and [60] for more details on this process.

3.4.4 Relative Phase between Fundamental and Third Harmonic

Controlling the relative phase of the third harmonic in an overlapped fundamental plus third harmonic experiment is quite a bit trickier than controlling the phase of just the second harmonic. For this setup the lower arm of the optical setup in Figure 3.5 is used. A beamsplitter with two dichroic coatings for 400 nm and 800 nm is used to split the incoming two pulses into each arm. 50% of the 800 nm and 80% of the 400 nm light is sent into the lower arm. A second BBO crystal cut for type I third harmonic generation (THG) is used to generate 266 nm light from the sum of an 800 nm S-polarized pulse and a 400 nm S-polarized pulse. Since the 800 nm pulse is P-polarized after the beamsplitter, another dual waveplate is used just prior to the THG crystal to advance the 800 nm pulse by $\lambda/2$ and the 400 nm pulse by λ, resulting in both being S-polarized as desired. This generates a pulse of 266 nm light which is P-polarized as desired for the experiment, so this polarization is not be changed.

In order to generate third harmonic light, the 400 nm and the 800 nm pulse must be overlapping, which is done using the first TDC after the SHG BBO. The overlap is fine-tuned by splitting the output colors of the third harmonic crystal with a prism, measuring the power of the 266 nm spot, and adjusting the rotation of the TDC to maximize the reading. Because second and third harmonic generation are nonlinear optical processes heavily dependent on the intensity of the incoming laser pulses, adjusting the compression of the incoming beams to maximize the intensity at the location of the crystal (rather than in the interaction region of the VMI chamber) can yield better results. Thus I found that I could generate more signal contrast in the resulting measurement by also adjusting the pulse compression in this step to maximize the third harmonic power. After the THG crystal, a set of glass wedges on a motorized stage (labeled “B”) is set to adjust the third harmonic relative phase, as in the previous setup. The final set of mirrors in the lower arm act as a retroreflector and sit on a manual stage, which can be adjusted for coarse path-length corrections to set the overlap of the 266 nm pulse from the lower arm and the 800 nm pulse from the upper arm. The position of this stage is roughly set such that the center of the stage ensures the measured path length of the upper and lower arm are approximately equal.

After the wedges, there are three pulses of different colors in the lower arm, but only the 266 nm is desired. The two mirrors on the manual stage and the recombining mirror are coated to reflect 266 nm and transmit 800 nm in order to discard the remaining colors. This worked well to get rid of the 800 nm light, but I later discovered evidence that the not-fully-extinguished residual 400 nm light from the lower arm had a noticeable effect on the final collected spectrum. If I were to redesign this portion to improve the experiment I would replace the second mirror on the retroreflector with a mirror coated to reflect 266 nm and transmit 400 nm, and additionally replace the recombining mirror with a laser-line filter set to reflect 266 nm and transmit other colors.

> The time overlap of the 266 nm pulse and the 800 nm pulse is predominantly controlled through manipulating the horizontal position of the manual stage to increase or decrease the path length of the lower arm. To determine if time overlap is achieved, I employ a difference frequency generation (DFG) scheme in another Type I THG BBO crystal. Simply stated, by placing a tripling crystal after the recombining optic along with a prism afterwards, if the 800 nm pulse from the upper arm is co-timed with the 266 nm pulse from the lower arm, their interaction in the tripling crystal will yield additional 400 nm light. This requires a few considerations to work correctly. First, the 800 nm light must be cross-polarized to the 266 nm light, and so the dual waveplate in the upper arm must be set to make the 800 nm light S-polarized at the recombining optic. Because there is still the 400 nm light from the upper arm and some residual 400 nm light in the lower arm, there is still a bright spot corresponding to 400 nm light after the prism. This bright spot obscures the small effect from the DFG, and so I place glass filters in each of the arms to transmit the desired color and block the 400 nm light. These change the true relative timing between the pulses and have to be compensated for at the end, but are useful for coarse alignment. Lastly the duplicate chamber window is

> inserted before the DFG crystal to account for the relative time delay caused by the extra glass upon entering the VMI chamber. Once all of these are accounted for, it's a simple process of carefully adjusting the manual stage position and observing the prism output to see a sudden flash where the 400 nm light should be once time overlap is achieved. This flash can be extremely subtle, depending on the relative intensities of light entering the DFG crystal, so very close observation is required.
>
> After this time overlap is achieved, the filters should be removed, and time overlap should be looked for again around the calculated stage position adjustment due to the thicknesses and indices of refraction of the filters used. By tweaking the horizontal angle of the recombining optic, and walking off spatial overlap of the 800 nm and the 266 nm pulse, the resulting DFG 400 nm pulse can be slightly moved away from either of the two 400 nm spots from the other arms. This can allow one to see the formation of the DFG spot by manipulating the manual stage, but it is very difficult to see. If the incoming laser pulses have enough power stability and are of sufficient intensity to generate a significant DFG pulse, a power meter can be used at the location of the 400 nm spot to measure when time overlap has been achieved; however, I did not have access to that level of stability and could not observe that effect with the tools available to me. After this rough time overlap has been attained, the duplicate vacuum window, DFG crystal, and prism are removed, and the polarization of the 800 nm light is reset to horizontal.

Because the splitting and recombination scheme does not maintain the collinearity of the 800 nm and the 266 nm pulse, significant attention must be devoted to ensuring proper spatial overlap of the two pulses after the recombining optic. I followed two main procedures for ensuring this overlap. First, an iris was inserted and centered in each of the two arms prior to the recombining optic and contracted to form the minimum spot size. The position of the laser pulses from each arm is measured just after the recombining optic as close to the optic as possible using a laser-viewing card. These spots are overlapped by eye as closely as possible by adjusting the position of the recombining mirror on the table and subtly tweaking the vertical position of the first 266 nm mirror. Next, a set of mirrors is inserted into the output of the recombining mirror to reflect the output far across the laser table (in my case about 3 meters of additional path length) onto laser viewing paper. The recombining optic is manipulated to align the far spots on top of each other by eye. These two steps are iterated as many times as necessary to ensure good rough alignment.

Once rough alignment has been achieved, the two beams can be sent into the VMI chamber and focused onto the interaction region. Again, a final iris is used to reduce each of the beam widths to the same size. At rough alignment, there should be an enhancement of the electron count rate exceeding the contribution from each of the two beams individually. This is because the strong-field tunnel ionization rate (Equation 2.7) scales non-linearly with field strength, and when roughly overlapped the field amplitudes from each laser color should sum. From here, optimal temporal overlap can be achieved by subtly adjusting the position of the relative phase wedges until the maximum count rate is achieved. Optimal spatial overlap is achieved by tweaking the position of the recombining mirror to further maximize the count rate.

Just as for the case where we consider the relative phase of the 800 nm pulse and the 400 nm perturbation when the glass wedge is inserted into the path, we should examine the relative phase of the 800 nm pulse and the 266 nm perturbation. The major distinction between these two cases is that the wedge mounted on motorized stage B only affects the delay of the 266 nm pulse. Therefore we no longer need to subtract the effect of the delay of the 800 nm pulse to determine the relative phase, and we can directly use the reciprocal of Equation 3.1 which is presented in Table 3.1 for 266 nm to determine the amount of glass insertion needed per cycle of relative phase:

$$\frac{\Delta x}{\Delta \theta_{266\text{nm}}} = 10.88\ \mu\text{m/cycle}. \tag{3.5}$$

Note that this is a much smaller amount of glass insertion than the amount needed for the 400 nm case. In order to accurately sweep over a full cycle of relative phase, it's important to ensure that the linear stage being used has significant accuracy to move with sub-micron precision.

3.4.5 Simultaneous Control of Three Laser Colors

The most complicated of these experiments involves overlapping and controlling all three of these laser colors simultaneously. The spatiotemporal overlap of three laser pulses is difficult to achieve; however, I found the most success through the following procedure. First, optimize the output of the third harmonic BBO crystal using the first TDC; second, optimize the temporal overlap of the 800 nm and 400 nm pulses through the detailed procedure above and by using the electron count rate; third, follow the second detailed procedure above to get spatiotemporal overlap of the 800 nm and 266 nm pulses. While these three steps are rather simple when presented in this form, in practice this is a quite finicky procedure that requires significant time and attention to achieve a stable three-color overlap.

After initial three-color overlap is achieved, the real challenge begins: maintaining overlap over the course of the experiment while adjusting the relative phase of each color. When motorized stage B moves only the timing of the 266 nm pulse changes. This is ideal, and it allows for uncoupled manipulation of that experimental parameter using the one stage. However, when the other motorized stage (A) in the combined 800 nm and 400 nm arm moves, both the relative phase between the 800 nm and the 400 nm pulses and the relative phase between the 800 nm and 266 nm pulse change. This is not a small effect either, as I will demonstrate.

Let's examine the effect of inserting the glass wedge mounted on stage A on the relative phase between the 800 nm and 266 nm pulse. Ignoring the effect on the 400 nm pulse for now, we see that only the 800 nm pulse shifts from this wedge insertion. Using Equation 3.3 with the values from Table 3.1 and $\Delta\theta_{266\text{nm}} = 0$ we have

$$\frac{\Delta\theta_{266\text{nm},800\text{nm}}}{\Delta x} = \tan(\phi)\frac{(1 - n_{800\text{nm}})}{266\ \text{nm}} = -0.0833\ \text{cycles}/\mu\text{m}. \tag{3.6}$$

Thus for every micron of stage A glass insertion, the 266 nm pulse gets ahead of the 800 nm pulse by about a twelfth of a cycle. Recall the result of Equation 3.4: a typical experiment varying the phase of the 400 nm perturbation at a minimum involves nearly 500 microns of stage A insertion. This sweep without any correction results in the relative phase of the 266 nm pulse to advance ahead of the 800 nm pulse by about 40 cycles! This effect dominates the relative phase of the 266 nm pulse, and so a correction factor must be applied which changes the insertion of the glass wedge on stage B whenever stage A moves in order to preserve the relative phase between the 800 nm and 266 nm pulse. Using the values we currently have access to, we can estimate this correction factor simply by using the results of Equations 3.5 and 3.6 as follows:

$$\frac{\Delta x_B}{\Delta x_A} = -\left(\frac{\Delta\theta_{266\text{nm},800\text{nm}}}{\Delta x_A}\right)\left(\frac{\Delta x_B}{\Delta\theta_{266\text{nm}}}\right) = 0.907. \tag{3.7}$$

Here I added the subscripts A and B to denote the translation of stage A and stage B respectively. The ratio $\Delta x_B/\Delta x_A$ indicates the amount stage B must be translated to preserve the relative phase of the 800 nm pulse and the 400 nm pulse. The sign of the result indicates that the two stages should be moved in the same direction whenever an adjustment is made.

For a reference on the magnitude of this correction factor let's consider a typical three-color experiment which sweeps over approximately one cycle of the second harmonic and approximately one cycle of the third harmonic. As per our previous discussion, to achieve this single cycle of third harmonic phase, we would advance stage B by about 11 microns; however, with the correction factor on top of approximately 500 microns of stage A insertion, we end up needing to translate stage B approximately 460 microns in total, while still maintaining sub-micron precision.

As mentioned earlier, the values presented in this section are all calculations. In a real experiment the true relative phase advancements and the correction factor must be calibrated empirically. This calibration is discussed towards the end of the following data collection section.

3.5 Data Collection and Experimental Control

In order to make successful measurements using the previously described experimental setups, I have written a comprehensive software application which interfaces with the detector, the experimental control devices mounted on the optical table, and various monitoring tools. In this section I outline the data collection techniques used to perform the experiments described previously, and explain how the tools available to me are incorporated into the exported measurements.

3.5.1 Control Devices

For a single data acquisition with fixed parameters, no experimental control is necessary. One would simply set up the experiment, affix the static parameters to their proper values, and allow electrons to accumulate on the detector. These types of experiments have a lot of value, and indeed in subsequent chapters I discuss at length a static scan of photoionized argon gas. However, most of the experiments I have performed involve actively controlling the optics on the table in order to sweep through parameters of interest during a single data collection. To do this, I use motorized translation stages and rotation mounts manufactured by Newport which are controllable programmatically from the lab computer.

Currently, there are four motorized optomechanical components in regular operation which allow me to control several parameters of interest. Firstly, a motorized rotation mount containing a half-wave plate is used to finely control the laser pulse power. Secondly, a translation stage upon which a retroreflector is mounted is used to control the delay between two laser pulses for pump-probe type experiments. Thirdly, a translation stage, upon which one wedge from a matched thin glass wedge pair is mounted, is used to control the relative phase between the laser and its second harmonic in experiments requiring $\omega - 2\omega$ control. Lastly, a similar translation stage is used to control the relative phase between the laser and its third harmonic while preserving the relative phase of other harmonic components for experiments requiring $\omega - 2\omega - 3\omega$ control.

3.5.2 Data Collection

The entire following section goes into more detail about the mechanisms of data collection than is required to understand the the following chapters outlining the results of these measurements. A reader not overly interested in replicating these experiments for themselves can quickly skim this section for key items before advancing to the next chapter.

On-the-Fly Peak-Finding

Every data acquisition requires that the video stream from the camera recording the intensified phosphor of the detector stack be converted into a format that records the locations of all electron impacts. There are several methods, in principle, to achieve this.

The most simple method would be to simply save the entire video over the course of the data collection, and do analysis after the fact to extract electron counts frame-by-frame. This method is prohibitively storage intensive for all but the shortest of experiments. As each experiment I discuss in this thesis was collected on the timescale of days, a different solution was required.

Another commonly employed method is to continuously sum each frame of the video feed, allowing the electron impacts to build up over time into a final summed spectrum. This is fast and uses little computing power, but has a number of disadvantages. Firstly, the MCP stack and intensified phosphor do not have equal gain across their effective area due to ordinary wear and tear, which results in individual electron impacts appearing brighter or dimmer in the video feed. An ordinary sum over-counts the brighter impacts and under-counts the dimmer ones, resulting in a final spectrum that does not accurately reflect an electron probability distribution. Secondly, each electron impact illuminates multiple pixels on the camera, so the final spectrum would blur the location of each electron. This blurring can be very significant and fully obscure the electron holograms that are the focus of this thesis.

The method I employ for all data collections is called "on-the-fly peak-finding." Simply summarized: while a data collection is ongoing, every single frame of the video feed is analyzed to extract out the single-pixel locations of every electron impact, and those locations are summed to a master distribution of all previous single-pixel electron impacts. This method is computationally expensive and must satisfy the following constraints. The entire peak-finding algorithm must operate faster than the frame rate of the camera, for which mine is 200 Hz. Each electron impact must be counted a single time only, regardless of their relative brightness or spread over the camera pixels. No electrons may be missed by the algorithm and thus omitted from the final summed distribution.

In order to optimize the speed of the algorithm, all included functions and variables are hardware accelerated in parallel using a graphics processing unit (GPU). With two Nvidia Titan RTX graphics cards, the following algorithm runs in 0.7 ms in Matlab and accurately produces single pixel precision for each electron impact for experiments in which there are approximately 100 electron impacts per frame.

1. Receive as input the 1024 by 1024 matrix of unsigned 8-bit integers from the live camera feed.
2. Convert the input to a GPU Array (Matlab's data format which can be processed on a computer's GPU) to allow hardware acceleration.
3. Apply a threshold filter to eliminate spurious signal below a defined value.
4. Apply a two-dimensional median filter to eliminate white noise by assigning each pixel the median value from its 3 by 3 neighborhood.
5. Convert the image to single-precision floating point format and then convolve it with a 7 by 7 Gaussian function in order to emphasize the Gaussian shape of electron impacts while suppressing spurious signal. Single-precision floating point format is necessary to avoid multiple pixels near the peak acquiring the same value and thus disrupt steps 7 and 8 below.

6. Stretch the dynamic range of the image by dividing by the maximum value of the image and multiplying by 255, before converting back to an unsigned 8-bit integer. Store this image.
7. Apply a dilation filter by assigning each pixel the maximum value from its 3 by 3 neighborhood.
8. Generate and output the 1024 by 1024 binary matrix of every pixel whose value after dilation equals its value before dilation. These are the single-pixel locations of every electron impact.
9. Sum the binary matrix to the master matrix recording electron impacts over the course of the data collection.
10. Repeat all of the above steps for the next input video frame.

After a full data collection, the summed master matrix output of this algorithm accurately reflects the electron momentum distribution from the measured ionization experiment.

Pulse Power Control

The power control is calibrated by taking the readout from a power monitor and sweeping through the rotation positions of the motorized mount. This yields a sinusoidal function of rotation angle versus measured power, which is programmed into the software suite for calibration. After the power control, many optomechanical components may lower the pulse power, so the calibration is stored in the suite as a percentage of maximum input power. When controlling the pulse power, the software suite predominantly operates in three modes: percentage of max power, desired pulse power, or histogrammed intensities.

The percentage of max power is the simplest form of operation, in which the software rotates the motorized mount to the position determined by the calibration to yield a defined percentage of the incoming laser power. This is most reliable in circumstances in which you are not too concerned about gradual variation in the power entering the power throttle. For the laser system I used during my graduate career, there were periods in which the power was extremely stable and this mode was more than sufficient, but there were also periods in which the power was very unstable and I observed fluctuations of 10-20% or more over the course of a day. During these unstable periods, different modes of operation are needed.

When running a single experiment at one particular pulse intensity, it can be useful to adjust the power controller at set intervals such that a photodiode measuring the transmitted light behind a mirror reads a near constant value over the duration of the experiment. In this mode of operation, the experiment is initialized at the desired laser intensity, and the averaged photodiode readout is recorded along with the current set percentage. After a specified number of collected laser shots, the data collection is paused for approximately two seconds while the pulse power is adjusted to match the desired value. During this time, a new photodiode readout is taken, averaging over one second. The ratio of the new power to the desired power is taken and multiplied by the initial (at the start of the data collection) percentage to yield the adjusted pulse power percentage which the power control is then set to:

$$P_f = \frac{(V_f)(P_i)}{V_i}. \tag{3.8}$$

Here V is the measured photodiode voltage, P is the pulse power percentage the power control is set to, and i and f index the initial and final states respectively. This mode of operation typically yields more stable single intensity measurements.

When running an experiment at multiple intensities in order to determine the intensity response of various spectral features, rather than focusing on a single pulse intensity it can be more useful to bin the final spectra according to the photodiode readout. In this mode, a set of initial nominal pulse power percentages is specified that are scanned over during the course of the experiment. At the beginning of each new sweep over these positions, the average photodiode readout is measured, and the positions are sorted into new bins in accordance with Equation 3.8. Importantly for this mode, the positions of the rotation mount are fixed throughout the experiment and all that changes is how the collected detector image is sorted. The final output spectrum is three dimensional, with the first two dimensions for the detector geometry and the third for the intensity bin index.

Two Pulse Delay

Pump-probe experiments using two IR pulses follow a very straightforward data collection procedure. The stage position of overlap of the two beams is determined from the pattern of the two beams as described previously and recorded as the reference point x_0 for this procedure. The time delay between the pulses is then simply calculated from double the linear translation stage's translation:

$$t_D = \frac{2\Delta x}{c}, \tag{3.9}$$

where $\Delta x = x - x_0$ is the displacement of the stage from the reference point (with $\Delta x > 0$ defined as the stage moving to the right with respect to the schematic shown in Figure 3.4) such that $2\Delta x$ is the change in path length of the pump pulse.

When collecting a two-pulse measurement, the desired delays in picoseconds are specified as a range or a defined vector. The required stage positions corresponding to these delays are extracted from the inverse of Equation 3.9, and stored as an output vector. During the collection, the stage sweeps over a random permutation of these stage positions, with the collected counts stored in the third index of the output spectrum, just as in the intensity binning case discussed previously.

Relative Phase of Fundamental and Second/Third Harmonic

The data collection scheme for $\omega/2\omega$ experiments follows a similar structure to the two pulse delay data collection approach. A range or defined vector of stage positions is provided, using the result of Equation 3.4 to guide the selection. During the collection, the stage sweeps over a random permutation of these stage positions just as in the two pulse case, and the final output is again three-dimensional, with the third index corresponding to the index of the stage positions. The case for $\omega/3\omega$ is nearly identical, except that it uses Equation 3.5 to guide the selection of stage positions.

During these data collections it is important to empirically determine the values calculated in Equations 3.4 and 3.5 for future calibration of the three-color correction factor. Since the 2ω signal is periodic and anti-symmetric with respect to the sign of the electric field of the fundamental, I examine the change in total counts of a region of interest on the detector compared to the equivalent region of interest on the opposite side of the detector. This yields two sinusoidal curves which are π out of phase with each other. With sufficient statistics (usually accumulating in about an hour at a count rate of 100 ionized electrons per laser shot), an accurate measurement of the period of this oscillation in units of stage position is determined, which is used in part to adjust the value of Equation 3.6. The 3ω signal is also periodic but is symmetric with respect to the sign of the electric field of the fundamental, and so the total counts in the regions of interest on opposite sides of the detector form two in-phase sinusoidal oscillations. Accurately measuring the period of these oscillations provides the other half of the correction needed to modify Equation 3.6.

Simultaneous Control of Fundamental, Second Harmonic, and Third Harmonic

When controlling all three colors simultaneously, the correction factor proposed in Equation 3.7 must be incorporated, but is instead defined in terms of the empirically determined stage translations per full cycle advancement of each color. Using these translations, we calculate the empirical correction factor as

$$\frac{\Delta x_B}{\Delta x_A} = \frac{n_{800\text{nm}} - 1}{n_{400\text{nm}} - n_{800\text{nm}}} \frac{3\Delta y_B}{2\Delta y_A}, \tag{3.10}$$

where we have introduced the quantities Δy_A and Δy_B as the empirically determined stage translations required for stages A and B to respectively advance the relative phase of the 400 nm pulse and the relative phase of the 266 nm pulse by one full cycle.

When selecting the stage ranges over which to sweep the data, the positions for stage A, $\mathbf{x}_A$, are determined and stored exactly as before. The positions for stage B, however, now must incorporate the empirical correction factor determined from Equation 3.10. Thus when a single range or vector is defined for stage B, $\mathbf{x}_B$, a two-dimensional matrix, $\mathbf{X}_B$, is generated in accordance with

$$X_B(m,n) = x_B(m) + \frac{\Delta x_B}{\Delta x_A}(x_A(n) - x_A(1)). \tag{3.11}$$

Here, we have the indices m and n to index stage B positions and stage A positions respectively.

During the data collection, an index n is randomly selected (without replacement) to advance to stage position $x_A(n)$, at which point a random permutation of the indices m are swept over to determine the position of stage B in accordance with $X_B(m,n)$. After a full sweep of the m indices, the index n is advanced to the next randomly permuted value and this process repeats.

The ending outputs of this procedure are a four-dimensional matrix where the first two dimensions index the detector pixels, the third dimension indexes the position of Stage A in accordance with $x_A(n)$, and the fourth dimension indexes the position of Stage B in accordance with $X_B(m,n)$ which incorporates the index of the third dimension.

Chapter 4

Analyzing Photoelectron Momentum Distributions

This chapter focuses on the fundamental analysis techniques used ubiquitously throughout the rest of the thesis to process the collected data from the VMI. A well-studied but quintessential example of a single long-duration scan of argon gas is examined throughout, and particular attention is devoted to identifying and discussing the major visible holographic features of interest in this spectrum. After the detailed feature identifications and descriptions presented in this chapter, later discussions mostly refer to these features by name only.

4.1 Data

The data shown in this chapter originated from the following collection parameters. We are observing the PMD of argon gas atoms photoionized by single, linearly-polarized, 800 nm wavelength laser pulses. The argon gas was delivered to the VMI source chamber via an Even-Lavie valve operating at 500 Hz and triggered to every other laser shot with an opening time of 22 μs and a backing pressure of 125 psi. The laser pulses were focused to a peak intensity of approximately 125 ± 10 TW/cm^2, which corresponds to a Keldysh parameter of $\gamma \approx 1.03$. The determination of this peak intensity is discussed in Chapter 5 as it is not directly measurable during a data collection. Each interaction event between the laser and the pulsed argon gas beam generated approximately 50 photoelectrons, which were then captured by the VMI spectrometer with voltages tuned to acquire up to approximately 70 eV electrons. The camera recorded the intensified phosphor detector screen at a frame-rate of approximately 200 Hz with a 2 ms integration time, capturing one interaction event per frame. Each frame was processed using on-the-fly peak-finding to extract single pixel precision (without centroiding) electron positions, which were integrated to form the final PMD. In

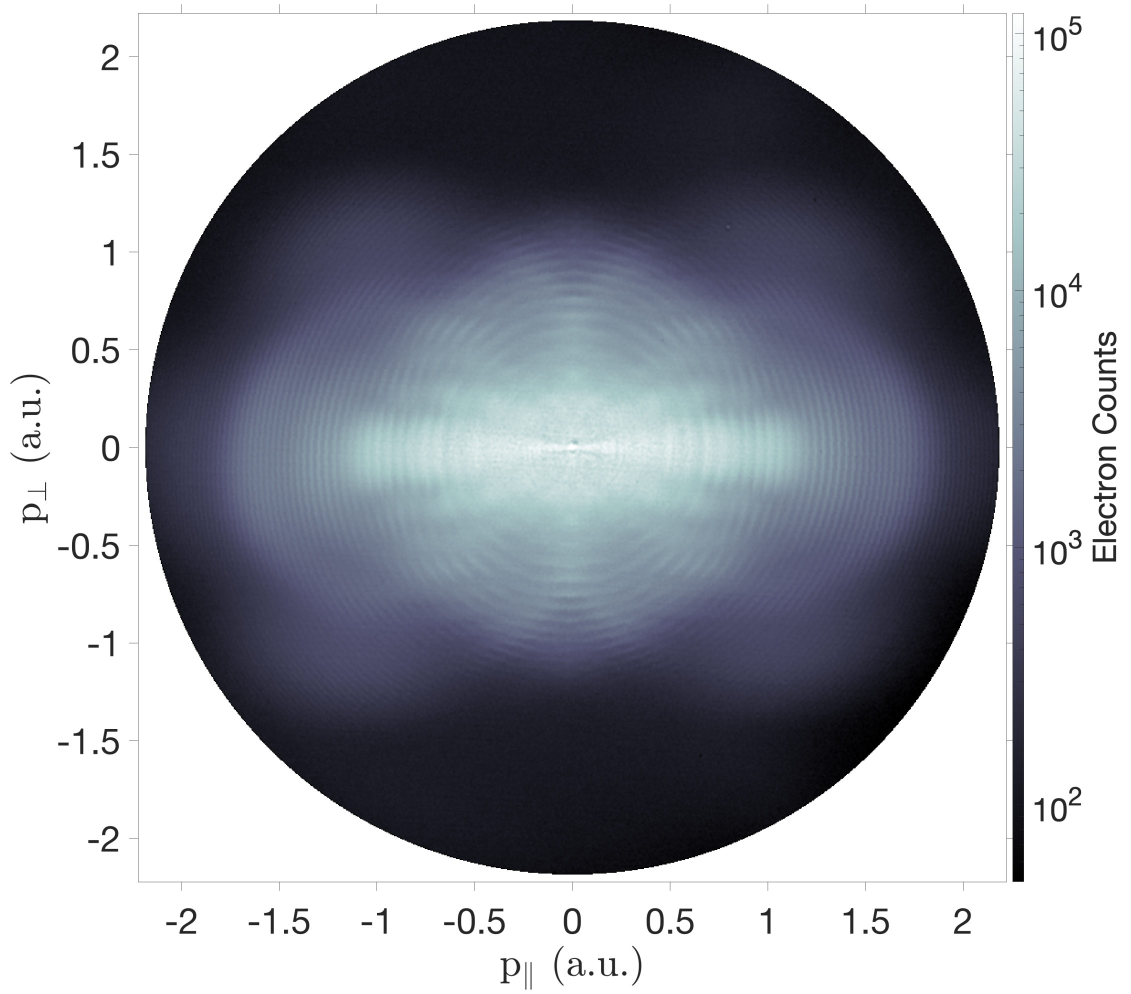

Figure 4.1: The raw PMD of the strong-field ionized argon dataset used throughout Chapters 4 and 5 of this thesis.

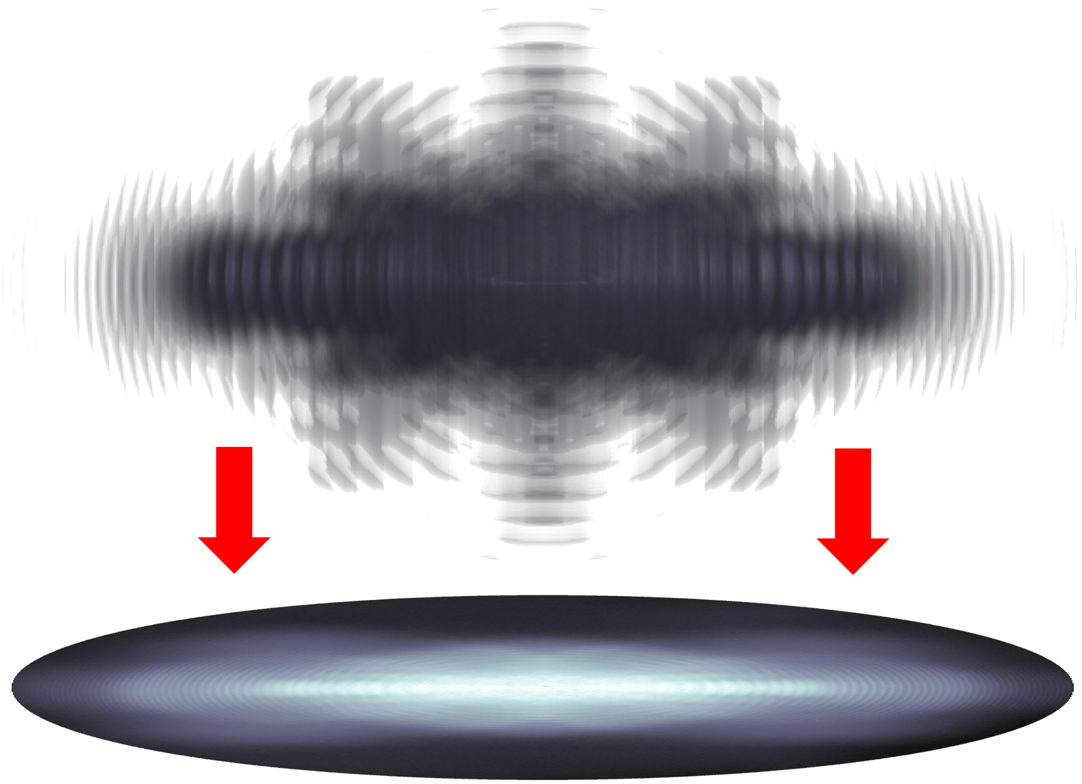

Figure 4.2: Representation of the three-dimensional Newton sphere of final electron momenta being projected onto the electron detector to form the raw PMD in Figure 4.1.

total, this collection ran for three days, and a total of 2.4 billion electron counts were recorded. The resulting raw VMI PMD (after centering and aligning) is presented in Figure 4.1.

4.2 Newton Spheres and Abel Inversion

An electron VMI focuses the photoelectrons onto specific pixels of the detector depending on their final momenta in the interaction region. Thus each pixel of the collected raw spectrum shown in Figure 4.1 presents the integrated yield of electrons which attained the final momentum indicated by the axes. Here, $p_{||}$ refers to momentum parallel to the direction of the laser polarization, while $p_\perp$ refers to the momentum component transverse to the laser polarization and the VMI spectrometer axis. However, this is an incomplete description of the system because the full 3-dimensional distribution of photoionized electron vector momenta ($p_{||}$, $p_\perp$, p_s) is being projected onto a 2-dimensional detector, see Figure 4.2. The 3-dimensional momentum distribution is referred to as a “Newton sphere” in this field of research for historical reasons, and I maintain that terminology.

One of the first steps in disentangling the raw PMD is to correct the spherical errors introduced by flattening the Newton sphere onto the detector. Because the laser pulses are linearly polarized, the entire ionization process maintains cylindrical symmetry about the laser polarization axis. There exists an integral transform which generates the 2-dimensional planar projection from a cylindrically symmetric 3-dimensional distribution, which is termed the 3-dimensional Abel Transform [53], and is given by the following:

$$F(y,z) = \int_y^\infty \frac{f(\rho,z)\rho d\rho}{\sqrt{\rho^2 - y^2}}. \tag{4.1}$$

Here, $\rho^2 = x^2 + y^2$ is the cylindrical radius of the three dimensional distribution, $f(\rho, z)$, which has cylindrical symmetry about the z-axis. This distribution is projected onto the yz plane forming $F(y, z)$.

In our PMD, we start with the projection $F(y, z)$ and wish to recover the 3-dimensional distribution $f(\rho, z)$. From now on, when describing the experimental data I describe these functions using our coordinate system: $F(p_\perp, p_\parallel)$ and $f(p_\rho, p_\parallel)$, where I'm defining $p_\rho^2 = p_\perp^2 + p_s^2$. To do this we need to invert the integral shown in Equation 4.1. This is called the 3-dimensional inverse Abel Transform [53] or colloquially Abel inversion.

Abel inversion is used ubiquitously in light-matter data analysis techniques; however, there is no simple closed-form formula for 3-dimensional Abel inversions, and so we rely on numerical approaches. The two most commonly employed techniques are "polar basis set expansions" (pBasex) [61] and "polar onion peeling" (POP) [53]. In my work, I exclusively employ POP-type approaches to Abel inversion, and I detail its use here.

4.2.1 Polar Onion Peeling

In order to understand the motivation for the POP algorithm, it is useful to understand the form of the forward Abel Transform shown in Equation 4.1. Firstly, we note from the fact that the integral ranges from y to ∞ that $F(y, z)$ is "constructed" from the sum of the projection of infinitesimally thin cylindrical shells concentric to the cylindrical axis z, beginning at a radius of $\rho = 0$. An inverse transform would then compose the 3-dimensional distribution out of contributions from infinitesimally thin pairs of parallel strips on the projection, beginning at $y = \pm\infty$. These would then have to be subtracted from the projection as the strips advanced toward the z-axis. As a mathematical transform, this cylindrical formulation is reasonable, but it has several limitations for real data. Most critically, it shifts statistical noise from the edges of the projection towards the central axis, which tends to be where the signal is maximal. Because the PMD in Figure 4.1 is contained within a circular region, a polar Abel transform is superior. Rather than adding concentric cylindrical shells from the central axis, a polar transform constructs spherical shells from the origin to infinity. Inverting this process requires subtracting the contribution of rings from the projection beginning outward from infinity and progressing towards the center. This is polar onion peeling.

The onion-peeling algorithm relies on an inductive process. First, assuming the full Newton sphere projection is recorded by the detector, the outermost ring of data in the original raw PMD exactly records the outermost equatorial ring of the Newton sphere, as no other portion of the Newton sphere is projected onto it. Then, subtracting the contribution of the spherical shell projection generated by the outermost ring yields a new partially peeled PMD in which the previous ring is uniformly zero, and the outermost non-zero ring now exactly records the proper value of the equatorial ring of the Newton sphere at that radius. This can be repeated until the entire raw spectrum is peeled. The complete algorithm of POP is summarized as follows:

1. Cast the Cartesian matrix of electron yields to polar coordinates, where each pixel is binned by radial distance from the origin in steps equal to the Cartesian pixel separation.
2. Construct the Legendre decomposition angular fit at the outermost unpeeled radius.
3. Convolve the angular fit to form a spherical shell, and project that shell back onto the detector plane.
4. Subtract the contribution of the projection from the remainder of the data. This is termed "peeling" the outermost radial contribution from the data.
5. Iterate to the next outermost unpeeled radius and repeat the above steps for all radii.
6. Take the union of all the constructed spherical shells to form the fully inverted 3-dimensional PMD.

While the 3-dimensional inverted PMD (shown in Figure 4.2) is the final output of the mathematics of Abel inversion, throughout this thesis we instead look at the equatorial cross section of the fully inverted PMD, which is formed from the union of all the angular fits. This cross section is shown in Figure 4.3. Because the system is cylindrically symmetric about the polarization axis, only two momentum parameters are needed to fully describe the complete momentum distribution, $p_{\|}$ and $p_{\perp}$. From the equatorial cross section of the Newton sphere, the full Newton sphere can be constructed simply by convolving the cross section about the polarization axis. We can describe this cross section using its Legendre decomposition using its polar coordinate representation as follows:

$$Y(p_r,\theta) = C_0(p_r)\left(1+\sum_{n=1}^{N}\beta_n(p_r)P_n(\cos\theta)\right) \tag{4.2}$$

Here, $Y(p_r,\theta)$ refers to the electron yield of the inverted PMD cross section (which I hereafter just refer to as either the "inverted PMD" or just the "PMD") at radial momentum coordinate p_r and azimuthal coordinate θ. It is congruent to $Y(p_{\|},p_{\perp})$ which is shown in Figure 4.3, and I use these interchangeably depending on which coordinate basis is more useful to discuss. For experiments which maintain symmetry between the positive and negative polarization axis (as is the

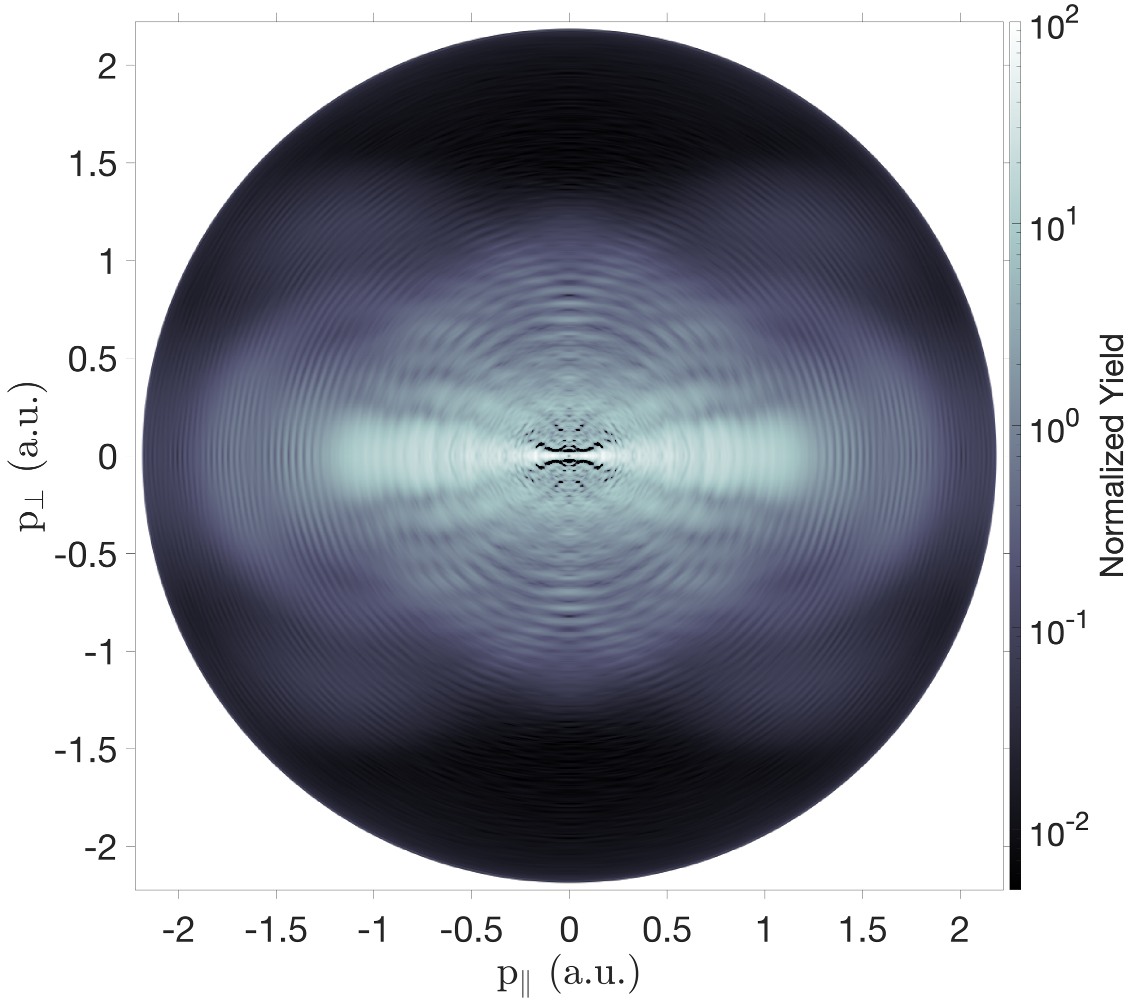

Figure 4.3: The normalized argon PMD after Abel inversion.

case for ionization due to all unshaped, linearly-polarized laser fields) only even orders of n need to be considered, and then N would be the highest even order to which the angular fits are numerically decomposed. $C_n(p_r)$ refers to the fitted nth order Legendre coefficient of the photoelectron angular distribution (PAD) at radius p_r. The parameters $\beta_n(p_r) \equiv \frac{C_n(p_r)}{C_0(p_r)}$ for $n \geq 1$ are defined as the nth order anisotropy parameters, whereas $C_0(p_r) \equiv \beta_0(p_r)$ is the isotropy parameter. For convenience, I refer to the complete set of the isotropy parameter and the nth order anisotropy parameters as simply the "anisotropy parameters." This Legendre decomposition representation in terms of the anisotropy parameters is very useful, and I frequently refer to them and manipulate them in future analysis. One immediate consequence of this representation is that rather than working with a two-dimensional grid of electron yields, we can instead analyze and manipulate a relatively small number of one-dimensional anisotropy parameters before reconstructing the modified PMD from Equation 4.2, which is how the majority of PMD transformations are performed in this thesis.

4.3 Understanding Photoelectron Momentum Distributions

The PMD shown in Figure 4.3 contains a wealth of information of strong-field quantum tunneling ionization encoded in the patterns of electron interference features visible throughout. These patterns overlap each other and form a tangled web of interconnected dynamics which confounds simple analysis. The vast majority of analytical effort for these types of experiments is spent disentangling these features and separately classifying individual holograms in order to extrapolate the ultrafast dynamics of strong-field ionization. Under the semi-classical electron trajectory model employed throughout this thesis, these holograms are the result of different electron trajectories which attain the same final momentum in the continuum after ionization. The motion and interference of these trajectories are determined by only three parameters: the shape of the ionizing strong field ($E(t, \vec{z})$), the launch phase at which the electron ionizes measured with respect to the ionizing strong field (ωt_0), and the transverse momentum the electron has at the tunnel ionization exit point (v_0). Different regions of the PMD have different types of these trajectory classes, or "orbits," which dominate the spectrum and yield unique holographic structures which can be individually studied.

4.3.1 Classical Evaluation of Quantum Orbits

While understanding the dynamics of these electron orbits clearly falls under the purview of quantum mechanics, it is useful to use ordinary classical mechanical models to estimate the approximate momentum domains different orbits may inhabit. For this section, I employ a Simple Man's Model (SMM) for the initial probe of strong-field electron dynamics to develop an intuition for how these electrons broadly behave and where they end up in our PMDs. The SMM is a simple classical framework for electron trajectories and rescattering which exclusively employs Newtonian kinematics after the initial tunnel ionization step. There are three kinematic assumptions of the SMM: the electron is born with zero momentum (both parallel and transverse to the laser polarization), it is born at the origin of our coordinate system before evolving in the laser field, and all rescattering behaves like isotropic hard-sphere elastic scattering.

Let us begin the SMM estimates. As before, we let the laser angular frequency be given by ω, it is polarized in the $\hat{z}$ direction, its amplitude is E_0, and we approximate it as a continuous wave such that

$$E(t, \vec{z}) = E_0 \cos(\omega t)\hat{z}. \tag{4.3}$$

In the dipole approximation the vector potential reads:

$$A(t, \vec{z}) = \frac{E_0}{\omega} \sin(\omega t)\hat{z} = A_0 \sin(\omega t)\hat{z}. \tag{4.4}$$

It is common practice to express all calculations in terms of the amplitude of the vector potential $A_0 = \frac{E_0}{\omega}$, such that conclusions can be drawn independently of the ionizing laser color. We can describe the momentum of an ionized electron using classical kinematics:

$$p_{\parallel} = \int_{t_0}^{t_f} dt(-A_0\omega \cos(\omega t)) = A_0 \sin(\omega t_0) - A_0 \sin(\omega t_f) \tag{4.5}$$

For direct electrons, this is the only equation we need, and we take $t_f \to \infty$ to determine the final momentum. Since obviously a sinusoidal function does not converge at infinity, we consider the average velocity of the oscillation, which is 0. Thus

$$\overline{p_{\parallel}} = \lim_{t_f \to \infty} \int_{t_0}^{t_f} dt(-A_0\omega \cos(\omega t)) = A_0 \sin(\omega t_0) - \overline{A_0 \sin(\omega t_f)} = A_0 \sin(\omega t_0) \tag{4.6}$$

The magnitude of $p_{\parallel}$ is maximized when $\sin(\omega t_0) = \pm 1$, yielding a maximal extent of $-A_0 \leq p_{\parallel} \leq A_0$. Rewriting in terms of energy: $0 \leq U_f \leq \frac{A_0^2}{2}$, where U_f is the final energy of the photoelectron. Recall that the ponderomotive potential is given by $U_p = \frac{A_0^2}{4}$, and so we can write $0 \leq U_f \leq 2U_p$, which is a famous and commonly cited result for direct electrons [62].

If we instead consider elastically rescattering electrons, we have to consider an intermediate time, t_r, at which the electron returns to the parent ion and changes direction. Since everything is cylindrically symmetric, we only need to consider planar scattering with the azimuthal angle θ, and can look at just the cross section coordinates $p_{\parallel}$ and $p_{\perp}$. Thus we write

$$\vec{p} = \begin{bmatrix} p_{\parallel} \\ p_{\perp} \end{bmatrix} = \int_{t_0}^{t_r} dt(-E(t)\hat{z}) \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \end{bmatrix} + \lim_{t_f \to \infty} \int_{t_r}^{t_f} dt(-E(t)\hat{z}) \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \tag{4.7}$$

Once again this yields

$$\begin{bmatrix} p_{\parallel} \\ p_{\perp} \end{bmatrix} = A_0(\sin(\omega t_0) - \sin(\omega t_r)) \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \end{bmatrix} + A_0 \sin(\omega t_r) \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \tag{4.8}$$

This equation has a very important consequence. The second term only having a component of momentum in the parallel direction indicates that in the final PMD, isotropic rescattering forms rings with radius $A_0(\sin(\omega t_0) - \sin(\omega t_r))$ centered on a point offset from the origin along the parallel axis by $A_0 \sin(\omega t_r)$. This is relevant for future discussions.

The last piece of this derivation is what the value of t_r ends up being for each value of t_0. We know in the SMM $z(t_r) = z(t_0) = 0$, and from kinematics $z(t) = \int_{t_0}^{t} dt' p(t')$. Thus we can write

$$0 = \int_{t_0}^{t_r} dt' p_{\parallel}(t') = \int_{t_0}^{t_r} dt' A_0(\sin(\omega t_0) - \sin(\omega t')) = A_0(t_r - t_0)\sin(\omega t_0) - \frac{A_0}{\omega}[\cos(\omega t_0) - \cos(\omega t_r)]. \tag{4.9}$$

This can be simplified to

$$\cos(\omega t_r) - \cos(\omega t_0) = (\omega t_0 - \omega t_r)\sin(\omega t_0). \tag{4.10}$$

This is a transcendental equation relating t_r and t_0, which can be solved numerically given either parameter or the final vector momentum alongside Equation 4.8. Importantly for future discussions, these solutions are not unique. Since the electron is oscillating back and forth in the field, there exist solutions where the electron does not rescatter on its first return to $z = 0$, but does rescatter at some later time when the electron once more returns to $z = 0$. It should be noted that given the physical assumptions and purely classical nature of the SMM, these multiple returns do not make much sense, but other models and experiments do rather clearly show the effects of multiple return rescattering. Interestingly, despite its shortcomings, the quantitative values given by SMM for multiple return rescattering are reasonably close to the values given by more complex models [63], and so these numbers are also useful to quote.

If we treat the scattering to be pure backscattering, where $\theta = \pi$, then we can use Equations 4.8 and 4.10 to numerically maximize the magnitude of $p_{\parallel}$. Assuming only a single return this comes out to $\omega t_0 = 0.2702$ radians (15.48 degrees) and $\omega t_r = 4.5239$ radians (259.20 degrees) for maximal momentum in the negative direction. Plugging these values into Equation 4.8:

$$\begin{bmatrix} p_{\parallel} \\ p_{\perp} \end{bmatrix} = A_0(\sin(0.2702) - \sin(4.5239)) \begin{bmatrix} -1 \\ 0 \end{bmatrix} + A_0 \sin(4.5239) \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} -2.2315A_0 \\ 0 \end{bmatrix}. \tag{4.11}$$

This value is symmetric with a π-shift for both ωt_0 and ωt_r, and so we get a momentum range of

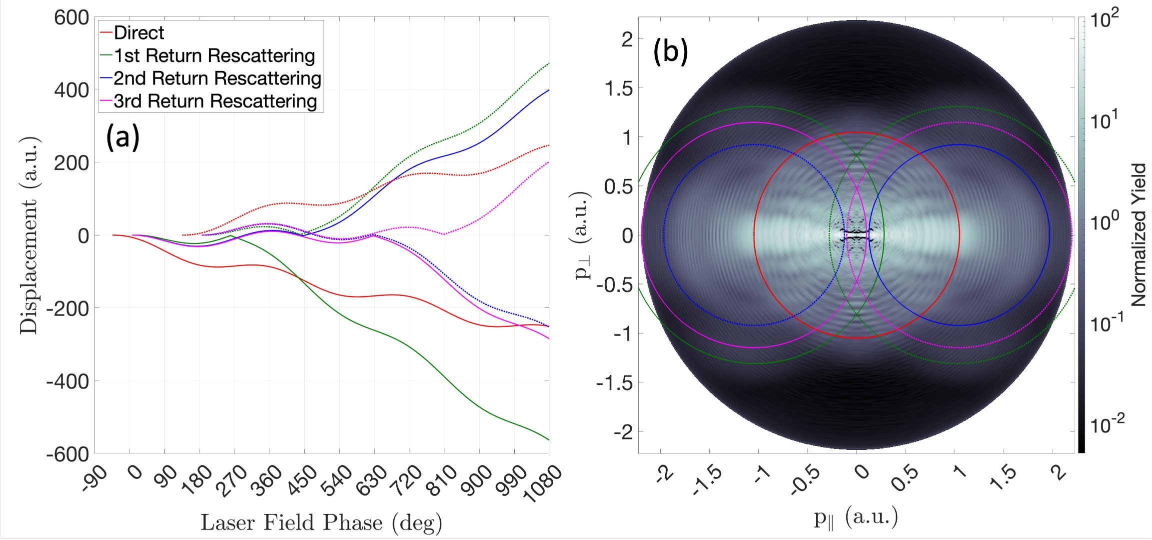

Figure 4.4: Calculations of the classical SMM rescattering trajectories and their maximal cutoffs. (a) This plot shows the one-dimensional SMM trajectories for electrons achieving the maximal final momentum for direct electrons and for rescattering electrons on their first three returns. Only pure backscattering is shown at the rescattering step. Two trajectories for each classification are shown, representing the same trajectory ionized from the opposite half cycle of the laser and resulting in opposite directions of the final momentum. (b) The final momenta from the classes of trajectories shown in (a) are plotted over the experimental data using the same line styles.

SMM Trajectory	**Launch Time Ranges (radians)**	**Return Time Ranges (radians)**	**Momentum Bounds**	**Energy Cutoff**
Direct	$\pi/2 \leq \omega t_0 \leq \pi$; $3\pi/2 \leq \omega t_0 \leq 2\pi$	N/A	$-A_0 \leq p_\parallel \leq A_0$	$2U_p$
Long Rescattering	$0 \leq \omega t_0 \leq 0.27$; $\pi \leq \omega t_0 \leq 3.41$	$4.52 \leq \omega t_r \leq 2\pi$; $7.67 \leq \omega t_r \leq 3\pi$	$-2.2A_0 \leq p_\parallel \leq 2.2A_0$	$10U_p$
Short Rescattering	$0.27 \leq \omega t_0 \leq \pi/2$; $3.41 \leq \omega t_0 \leq 3\pi/2$	$\pi/2 \leq \omega t_r \leq 4.52$; $3\pi/2 \leq \omega t_r \leq 7.67$	$-2.2A_0 \leq p_\parallel \leq 2.2A_0$	$10U_p$
Second Return Rescattering	$0 \leq \omega t_0 \leq 0.13$; $\pi \leq \omega t_0 \leq 3.27$	$2\pi \leq \omega t_r \leq 3\pi$; $3\pi \leq \omega t_r \leq 4\pi$	$-1.9A_0 \leq p_\parallel \leq 1.9A_0$	$7U_p$
Third Return Rescattering	$0 \leq \omega t_0 \leq 0.13$; $\pi \leq \omega t_0 \leq 3.27$	$3\pi \leq \omega t_r \leq 4\pi$; $4\pi \leq \omega t_r \leq 5\pi$	$-2.1A_0 \leq p_\parallel \leq 2.1A_0$	$8.8U_p$

Table 4.1: A summary of the trajectories and bounds predicted by the SMM. All results here can be determined from Equations 4.8 and 4.10. While not as quantitatively accurate as higher level models, the results shown here do roughly predict the structures and cutoffs found in experiment.

$-2.2315A_0 \leq p_\parallel \leq 2.2315A_0$ for first return perfectly backscattered electrons. Converting to energy again, we have $U_f \leq 2.4898A_0^2 = 9.9592U_p$, yielding the well known 10 U_p cutoff point in strong-field photoelectron spectra.

For completeness and to motivate future discussion, it is useful to examine the multiple return rescattering cutoffs as well. We can constrain the return times generated by Equation 4.10 according to the return number. For example, the second return to the parent ion occurs between 2π and 3π radians after the launch time, rather than between π and 2π radians for the first return. For n^{th} return rescattering electrons, I label the momentum and energy cutoffs as $p_{\parallel,n}$ and $U_{f,n}$ where appropriate.

We use these constraints to numerically generate the maximum magnitudes of $p_{\parallel,n}$ for second and third return rescattering electrons, which are the only multiple return rescattering electrons we examine further. For second return this yields $\omega t_0 = 0.1194$ radians (6.84 degrees) and $\omega t_r = 7.7911$ radians (446.40 degrees), yielding $p_{\parallel,2} = 1.877A_0$ using Equation 4.8. In energy this leads to $U_{f,2} \leq 1.7615A_0^2 = 7.0459U_p$, which I approximate as $7U_p$ in future discussion. For third return we have $\omega t_0 = 0.0967$ radians (5.54 degrees) and $\omega t_r = 10.9437$ radians (627.03 degrees), yielding $p_{\parallel,3} = 2.0939A_0$. In energy again this leads to $U_{f,3} \leq 2.1921A_0^2 = 8.7685U_p$, which I approximate as $8.8U_p$.

All of these cutoffs and classical trajectories are presented in Figure 4.4, and summarized in Table 4.1. Figure 4.4(a) shows the one-dimensional SMM trajectories for electrons which either do not scatter, or scatter on their first, second, or third return. The displacement from the parent ion is calculated by propagating Equation 4.7 in the case that $\theta = \pi$ using the numbers calculated above. A pair of trajectories is shown for each trajectory class, showing the same trajectory ionized half a

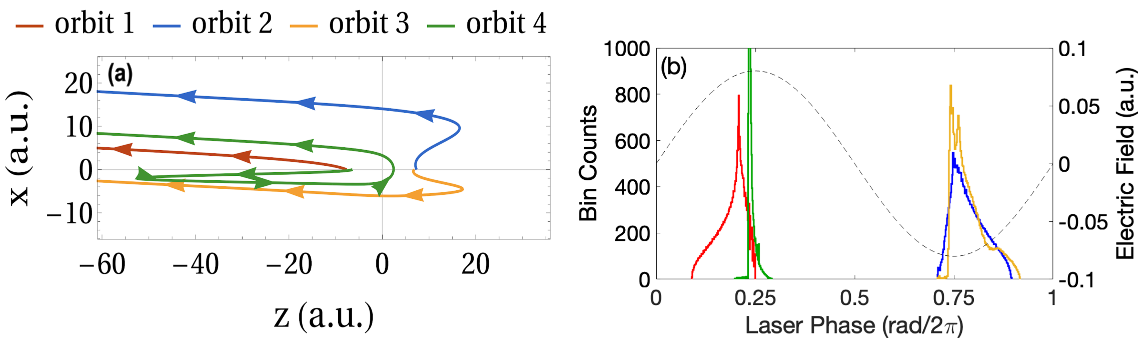

Figure 4.5: (a) Schematic example of the four orbits identified by CQSFA. The arrows on each trajectory denote the direction of travel, and each arrowhead is separated by 0.2 laser cycles. The position of the parent ion is located at the origin. The coordinate z denotes the position of the electron along the laser polarization axis with respect to the parent ion, while x denotes the position transverse to the polarization axis. Each of the shown trajectories has the same final momentum $[p_{\|}, p_{\perp}] = (-1.0, 0.13)$, and electrons following them would end up at the same position on the detector. (b) A histogram of the launch times of the orbits identified by the CQSFA. The simulated laser field is overlaid as the black dashed line. These launch times correspond only to the trajectories whose final momenta have a negative parallel component. The launch times for positive parallel momenta are identical with a half-cycle shift. Figure reproduced from Reference [2].

cycle later. Figure 4.4(b) shows these momentum cutoffs overlaid upon the PMD from Figure 4.3. Here $\theta \in [0, 2\pi]$ to show the two dimensional cross-sectional rescattering. Momentum rings from opposing half cycles of the laser are shown using the solid/dashed lines from Figure 4.4(a). We see the first return rescattering ring (green) closely aligns with the outermost lobes appearing in each quadrant at approximately $[p_{\|}, p_{\perp}] = [\pm1.25, \pm1.2]$. This outermost portion of the spectrum most closely shows the cross-section of scattering of the ionization target, and is a useful reference for target-specific measurements [34, 64]. Intriguingly the second return rescattering rings (blue) closely outline a prominent ring-like structure on each half of the spectrum. This partially suggests that this ring-like structure is related to second return rescattering, but we leave that discussion to Chapters 6 and 7 in which additional experimental data is analyzed to validate that claim.

4.3.2 Electron Holograms

While individual trajectories are interesting, and indeed significant work has been done to describe these trajectories [3, 18, 20, 22, 33], in experimental spectra they are difficult to observe directly. Rather, experimentalists have access to electron holograms generated by the interference of two or more electron trajectories. By examining the different holograms present in experimental PMDs and developing tools to extract the contributing electron trajectories, we gain insight into the ultrafast strong-field dynamics occurring during tunnel ionization. With robust analytical methods to disentangle strong-field holograms, electron VMI spectroscopy becomes a powerful camera to image ionizing electrons on their natural time scales.

To begin our analysis of electron holograms, we first must characterize the most prevalent holograms in strong-field spectra. To do this, we introduce the holograms predicted by the quantum-trajectory theory framework, the CQSFA, we compare to the majority of our experimental results.

The most important advantage when employing a quantum-trajectory theory such as the CQSFA is that we can isolate the holograms formed from any combination of orbits identified by the theory. The CQSFA distinguishes four specific electron orbits based on the qualitative motion of the electron about the parent ion after ionization. These orbits are labeled by both a number and a name, and are as follows: #1 the direct orbit, #2 the forward deflected orbit, #3 the forward rescattered orbit, and #4 the backward rescattered orbit. An example of each of these orbits ending at the final momentum $[p_{\parallel}, p_{\perp}] = (-1.0, 0.13)$ is shown schematically in Figure 4.5(a), and a histogram of their launch times for all such orbits ending with negative parallel momentum is shown in panel 4.5(b). These four orbits are not a complete set of all possible electron trajectories; however, for most of the features discussed in this chapter they are sufficient. An important omission for CQSFA are orbits with multiple returns to the ion, which are the subject of much discussion in Chapters 6 and 7. Table 4.2 overviews all the orbits which are discussed in this thesis, and orbits not included in the CQSFA description are further explained in later chapters where they are useful.

Orbit Name	**CQSFA Orbit?**	**Description**
Direct, #1	Yes	Electron travels directly to detector without reversing its transverse or parallel velocity direction.
Forward Deflected, #2	Yes	Electron ionizes on the opposite side of the ion as its eventual final parallel velocity, returns to the parent ion and does not reverse its transverse velocity.
Forward Rescattered, #3	Yes	Electron ionizes on the opposite side of the ion as its eventual final parallel velocity, returns to the parent ion and reverses its transverse velocity.
Backscattered (Long), #4	Yes	Electron ionizes on the same side of the ion as its final parallel velocity, returns to and loops around the parent ion.
Multiple Return Rescattered	No	After ionization, the electron changes directions multiple times in the field, looping more than one time around the parent ion.

Table 4.2: A concise list of the electron trajectories which are examined in many passages in this thesis. The orbits appearing in the CQSFA formalism are often referred to by number rather than by name in both the literature and in this thesis for brevity.

The holograms formed by interfering combinations of CQSFA orbits are presented in Figure 4.6. Panels 4.6(a)-(g) contain CQSFA calculations of the interference pattern of the orbits shown in the insets at the top center of each panel. These insets are reproductions of the example trajectories shown in Figure 4.5(a) with the axis labels removed for simplicity. As a reminder, these inset

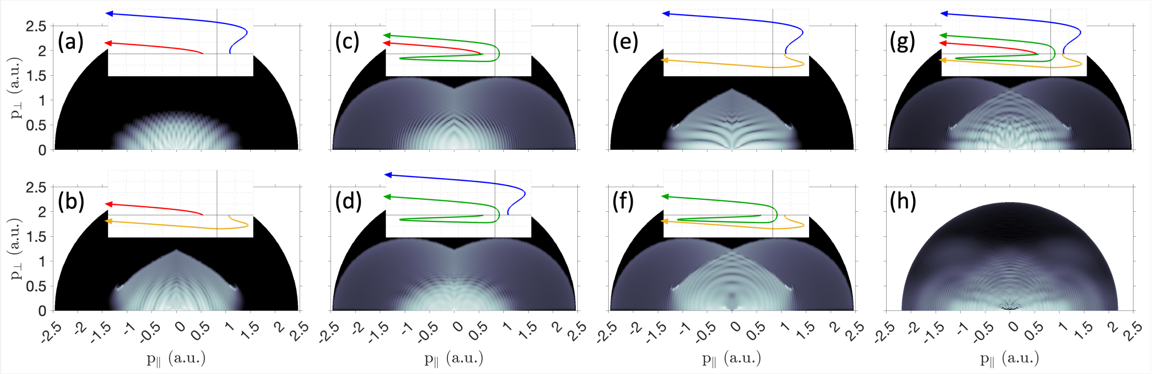

Figure 4.6: A presentation of calculated holograms from the CQSFA. Panels (a)-(f) present all the holograms formed from the interference of exactly two orbits identified by the CQSFA. Panel (g) presents the CQSFA calculation of the holographic structure formed from the interference of all trajectories simultaneously. Panel (h) reproduces the experimental Abel-inverted PMD from Figure 4.3 for comparison. Panels (a)-(g) contain insets which schematically show the orbit combination resulting in the calculated hologram. These insets are pulled from the trajectories shown in Figure 4.5(a) and have the axis labels removed for visual clarity.

trajectories are only accurate for a single pixel on the left side of the PMDs; however, the shapes of these trajectories are representative of the orbits forming the full PMD. Panels 4.6(a)-(f) display all the pairwise combinations of the CQSFA orbits, while panel 4.6(g) shows the full interference pattern formed by interfering all orbits at once. Panel 4.6(h) reproduces the inverted experimental PMD from Figure 4.3 for comparison.

Several of these holograms are well documented in the literature, and in particular I refer the reader to Reference [3] for further details on holograms involving the CQSFA orbits. For the purposes of this thesis, two of the pairwise holograms deserve special attention: the Spider and the Carpet (panels 4.6(e) and 4.6(f) respectively).

The Spider

It is a running tradition among scientists in this field to use forced acronyms to name apparatus, experiments, and analytical techniques after animals. For instance, there is RABBITT [65] ("reconstruction of attosecond beating by interference of two-photon transitions"), FROG [66] ("frequency-resolved optical gating"), or the endearingly egregious GRENOUILLE [67] ("grating-eliminated no-nonsense observation of ultrafast incident laser light e-fields"). Confusingly for our purposes there is even SPIDER [68] ("spectral phase interferometry for direct electric-field reconstruction"). Thankfully, the Spider holographic structure is so named because of its direct resemblance to the legs of a spider, as can be seen in Figure 4.6(e).

The Spider (or spider-leg structure, as I often refer to it) is the most prominent holographic structure visible in the majority of experimental PMDs. This is because it is formed from the interference of two orbits which both have high rates of occurrence and have similar amplitudes. These two orbits in the CQSFA formulation are the forward deflected orbit (#2) and the forward rescattered orbit (#3). From Figure 4.5(b), we note orbits #2 and #3 have ionization times at approximately the same phase in the laser field. As can be seen in the inset plot in Figure 4.6(e), the two contributing orbits are born at approximately the same tunnel exit point, but they have oppositely signed initial transverse momenta. Thus the Spider is the interference pattern formed from similar electron trajectories passing on opposite sides of the parent ion.

The Carpet

The Carpet is a subtle hologram which can be difficult to observe in experimental spectra without sufficient resolution. It is a two-trajectory interference formed in the CQSFA formulation from the forward rescattered orbit (#3) and the backscattered orbit (#4). Visible in Figure 4.6(f), the Carpet presents as a checkerboard-like pattern of nodes and antinodes along the transverse axis spaced by 2ω in both transverse and parallel momentum. The Carpet is the most studied hologram in this thesis for several reasons. First, it is one of the less completely studied holograms in the literature. Second, there have been very few experimental attempts to disentangle and study the Carpet directly. Third, the Carpet structure is formed by relatively rapidly changing phase differences between the contributing trajectories, and so is an excellent candidate for resolving phase information. Lastly, the Carpet structure is unique among the holograms for its distinctive qualitative appearance, which allows for unambiguous determination of the structure during disentanglement efforts. Another very useful feature of the Carpet is that it is formed from the interference of two trajectories launching from opposite sides of the parent ion on opposite half cycles of the laser field. This fact is not particularly relevant to our discussion thus far, but in the next chapter we discuss how to accurately retrieve time differences between interfering trajectories, for which the Carpet is a sterling example.

Above-Threshold Ionization Rings

By far the most studied feature of strong-field ionization are the distinctive peaks resulting from above-threshold ionization (ATI). In a photon-based model, these peaks result from the electron absorbing additional photons beyond what is required for ionization, causing the electron to attain relatively discrete energies. This results in rings of accentuated electron yield centered on the origin, with spacing set by the energy of a photon from the driving laser field. In an electron trajectory-focused model, these same rings are caused by identical trajectories whose launch times are spaced exactly one laser cycle apart. For ultra-short laser pulses such as the ones we work with, these are equivalent explanations. Due to the nature of the CQSFA which restricts all calculations to a single unit cell of the laser, the calculations from CQSFA do not contain any ATI interference. The majority of this thesis focuses on extracting useful information from sub-cycle interference patterns in the domain where the CQSFA is valid, meaning we examine holograms whose contributing trajectories have launch-time separations less than a laser cycle apart. In many ways the ATI structure is a hindrance to this goal, as it is an undesired feature which sits on top of all the rest of the holograms of interest. In the next chapter, one analysis avenue we examine involves removing the ATI rings from the data without disrupting the underlying spectrum in order to gain access to purely sub-cycle structure.

Chapter 5

The Time-Correlation Filter

My overall goal when generating PMDs such as the one shown in Figure 4.3 is to develop techniques to directly image ultrafast electron motion and the electronic structures of the parent. In the field of photoelectron imaging, as of this writing, the relationship between theorists and experimentalists is fairly one-sided, with theorists using calculations to characterize the features experimentalists see but cannot independently evaluate in their data [3]. Identifying the various holographic structures present in these spectra and determining (from calculation) the contributing electron orbits is broadly useful; however, far better would be the direct examination of electron trajectories from the data itself. This is the approach I take for the next two chapters of the thesis. Rather than rely on calculations to predict the features present in the data, I directly extract quantities of interest from the data, and compare those to theory to verify the predictions of the calculations.

Recall that holograms are formed from the interference of two electron trajectories. The path each of these trajectories follow are determined by three parameters: the launch time in the laser field, the initial transverse momentum of the launched electron, and the shape of the electric field during the electron's motion. In most experiments, the shape of the field can be determined to reasonably high precision prior to the laser entering the interaction regime and ionizing the sample, so we consider that to be a known parameter. In accordance with ADK theory, we treat the initial transverse momentum distribution as a Gaussian centered at zero, see Equation 2.9 [38]. Sometimes the specific transverse momentum of a trajectory arriving at a pixel in a PMD can be determined from the data directly, as the data is of the form of a momentum distribution to begin with. It is much more difficult to extract the launch times of interfering trajectories from experimental spectra. It is because of this difficulty that calculations have been so necessary for the study of these spectra, as trajectory-based calculations (such as the CQSFA) explicitly record the launch times for each trajectory, allowing experimentalists to roughly match features in data to their preferred calculation to attempt to derive insight into the observed ultrafast dynamics.

This chapter focuses on a Fourier analysis technique that can be used on all types of angle-resolved photoelectron spectra to directly extract timing information from experimental data. Rather than focus on the launch times of individual trajectories, we can instead determine the time correlations between *pairs* of interfering trajectories forming the bulk of the holograms present in strong-field spectra. This has three very important consequences that are expanded upon in the remainder of the chapter. First, time correlations can be used to determine the launch-time separation between interfering trajectories in quarter cycle steps, allowing for unambiguous verification of the predictions of theory for each observed hologram. Second, the data can be filtered using time correlations in order to remove specific holograms and reveal the underlying spectrum. Finally, the data can be truncated to extract a specific hologram and examine it independently of the rest of the spectrum which may otherwise obscure it. This technique is called the "Time-Correlation Filter" (TCF).

5.1 Constructing the Time-Correlation Filter

Since we are interested in extracting temporal information from our PMD, it seems logical to turn to Fourier analysis. The conjugate quantity to units of time has units of energy, so we begin by restructuring our data to have units of kinetic energy. Recall that after the Abel inversion step of data processing we have both an inverted two-dimensional PMD and a finite set of one-dimensional anisotropy parameters, which are the coefficients of the Legendre basis with respect to radial momentum:

$$Y(p_r, \theta) = C_0(p_r) \sum_{n=0}^{N} \beta_n(p_r) P_n(\cos\theta). \qquad \text{(4.2 revisited)}$$

It is a simple process to re-interpolate the quantities $\beta_n(p_r)$ to $\beta_n(E)$, where the electron kinetic energy E is uniformly sampled from 0 to the maximum value of $\frac{1}{2}p_r^2$. In this form, these $\beta_n(E)$ are the coefficients to an angle-resolved power spectrum, which when Fourier transformed yields a temporal autocorrelation function [2, 21, 69]:

$$\tilde{\beta}_n(t) \equiv \frac{1}{2\pi} \int_0^{\infty} \beta_n(E) e^{iEt} dE = \int y_n(t' - t) y_n(t') dt'. \qquad (5.1)$$

Here, $y_n(t')$ is the complex electron yield from all electron trajectories launched at time t', and t is the time difference between interfering pairs of trajectories. Thus $\tilde{\beta}_n(t)$ are a new set of one-dimensional parameters relating complex electron yield to time differences between interfering trajectory pairs for all possible ionization times. I refer to these parameters as the "time-correlation parameters" when discussed with words. This is not a trivial result and is not immediately self-evident from the form of Equation 5.1 alone; and so I provide further justification for this conclusion in the following sections.

If we, for now, accept this explanation for $\tilde{\beta}_n(t)$, then we can transform the units of both E and t from atomic units into more natural units. By dividing E by $\hbar\omega$ we can express E in units of laser's central photon energy, which when Fourier transformed yields t in units of laser cycles. $\beta_0(E)$ and $|\tilde{\beta}_0(t)|$ with these natural units are shown in blue in panels (a) and (b) respectively of Figure 5.1.

The original PMD can be reconstructed from the time-correlation parameters by inverse Fourier transforming them back into the energy-sampled anisotropy parameters, resampling those back to momentum space, and then using Equation 4.2. This indicates that we can manipulate the time-correlation parameters to selectively suppress or enhance some portion of the holographic structure present in the original spectrum. We can do this most effectively using digital Fourier filters operating on $\beta_n(E)$. By applying a low-pass filter with specified cutoff α to $\beta_n(E)$ before transforming back to a PMD, we remove the contributions of all interfering trajectories whose launch times are separated by more than α cycles.

For most of the applications in this chapter, I use digital low-pass filters constructed using Matlab's filter designer tool. There are three major considerations when designing the filter. First, the filter should be maximally flat in the pass-band in order to minimize the generation of artifacts in the features we want to observe. Second, the attenuation in the stop-band should be at minimum 20 decibels to sufficiently suppress the desired features. Third, the transition window width between the pass-band and the stop-band should be minimized as much as possible. To satisfy these considerations, we employ a Kaiser-windowed, finite-impulse response, low-pass filter with a transition width of 0.2 cycles to ensure 20 decibels of attenuation. The allowed pass-band ripple magnitude is constrained to 0.1 decibels. For such a filter, the cutoff of 6 decibels of attenuation occurs at the center of the transition window. These parameters were carefully tuned to satisfy the desired considerations. For details on how to interpret these parameters, I refer the reader to Reference [70], which is an excellent review of tailoring digital filters to satisfy most purposes. Due to the numerical resolution of the datasets in this thesis, the lowest cutoff is restricted to 0.1 cycles. For convenience, I refer to filters such as the one described above as TCF$^\alpha$ in accordance with its low-pass cutoff α in units of laser cycles. Additionally, in equations or when describing variables which have been filtered by such a low-pass filter, I treat the TCF$^\alpha$ as an operator and express its result with a similar superscript. For example, the time-correlation filtered anisotropy parameters with a cutoff of 0.7 cycles is defined as $\text{TCF}^{0.7}[\beta_n(p_r)] \equiv \beta_n^{0.7}(p_r)$.

5.2 Applying the Time-Correlation Filter

Figure 5.1 presents the result of TCF$^{0.85}$ acting upon our selected argon dataset. Panel 5.1(a) shows the filtered and unfiltered order 0 anisotropy parameter, while Panel 5.1(b) shows the time-correlation parameters corresponding to the unfiltered and filtered anisotropy parameter in 5.1(a). Panel 5.1(c) shows the magnitude response of TCF$^{0.85}$. There is a significant amount of structure visible in the unfiltered time-correlation parameter in (b). Most notably there is a tall peak at 1 field cycle of time separation. This peak corresponds to the prominent ATI ring structure that sits

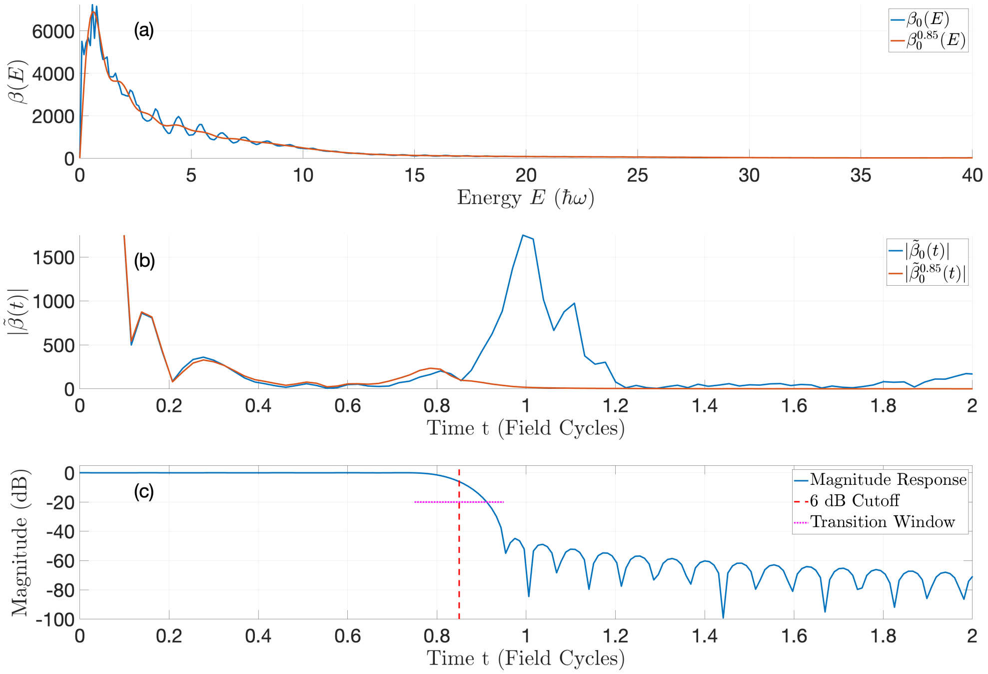

Figure 5.1: How TCF$^{0.85}$ operates on the anisotropy parameters. Panel (a) shows the energy-spaced anisotropy parameters $\beta_0(E)$ (blue) and $\beta_0^{0.85}(E)$ (orange) on the same set of axes. Panel (b) shows the time-correlation parameters $\tilde{\beta}_0(E)$ (blue) and $\tilde{\beta}_0^{0.85}(E)$ (orange) on the same set of axes. Panel (c) shows the magnitude response of the digital filter employed by TCF$^{0.85}$, using the same time axis as in panel (b). The 6 decibel cutoff is shown as a vertical dashed red line, and the extent of the filter's transition window is indicated by a horizontal dotted magenta line.

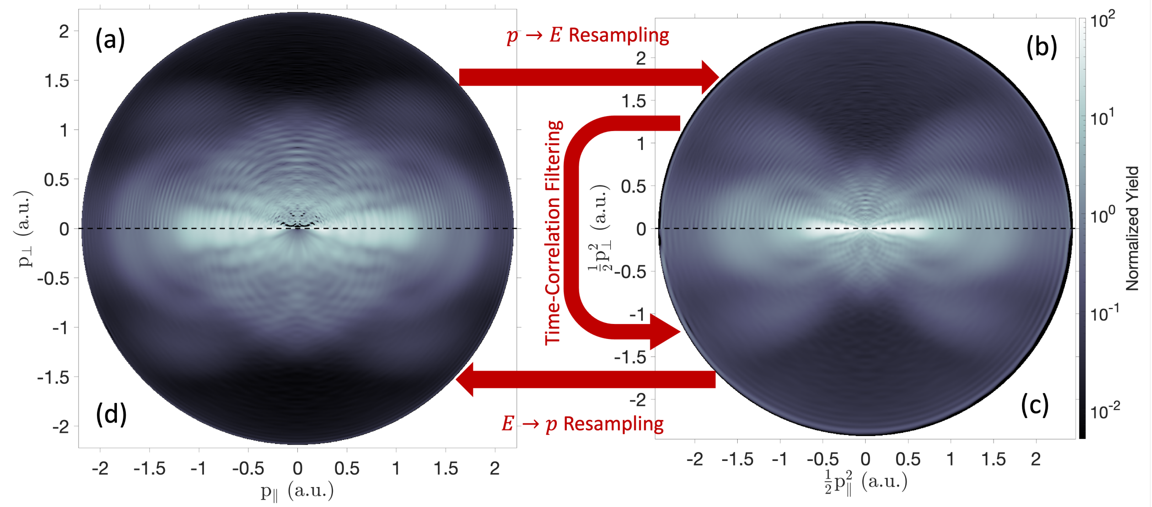

Figure 5.2: The full two-dimensional filtering process with a cutoff of 0.85 cycles. The steps are indicated by the annotated red arrows. (a) The top half of the panel shows $Y(p_\parallel, p_\perp)$ reproduced from Figure 4.3. (b) The top half of the panel shows the unfiltered energy-spaced distribution after having its anisotropy parameters interpolated into energy space. (c) The bottom half of the panel shows the filtered energy-spaced distribution after TCF^{85} is applied to the energy-spaced anisotropy parameters. (d) The bottom half of the panel shows $Y^{0.85}(p_\parallel, p_\perp)$ after interpolating $\beta_n^{0.85}(E)$ back to momentum space.

on top of the entire PMD. As mentioned towards the end of the previous chapter, because the ATI rings are an already well-studied and dominating holographic feature, they are of little interest to us. The time-correlation cutoff of 0.85 cycles in Figure 5.1 was selected strategically as the empirically optimal value to remove these rings without disrupting underlying holographic structure.

The full angle-resolved filtering process using TCF$^{0.85}$ is outlined in Figure 5.2. Panel 5.2(a) is constructed using Equation 4.2 and the original inverted anisotropy parameters $\beta_n(p_r)$. Each subsequent step (b) through (d) applies the transformation indicated beside the red arrow to the anisotropy parameters from the previous step, and then reconstructs the 2-dimensional representation by substituting the transformed anisotropy parameters into Equation 4.2. Note that in a typical TCF transformation, the two-dimensional spectra shown in panels (b) and (c) are not generated, but they are included here for illustrative purposes. Panels (a) and (d) directly compare $Y(p_\parallel, p_\perp)$ to $Y^{0.85}(p_\parallel, p_\perp)$. Looking at $Y^{0.85}(p_\parallel, p_\perp)$, we can clearly see that the ATI rings have been fully removed, and that the underlying structure somewhat visible in $Y(p_\parallel, p_\perp)$ has not only been preserved but has also been significantly uncovered.

The PMD $Y^{0.85}(p_\parallel, p_\perp)$ is discussed so frequently in this thesis that it is repeated in its own large figure, see Figure 5.3. One of the primary holographic structures revealed by TCF$^{0.85}$ is the Carpet, which is now clearly visible along the transverse axis in the approximate region of 0.7 a.u. $< p_\perp <$ 1.25 a.u. This characteristic checkerboard-like pattern of phase shifts which was

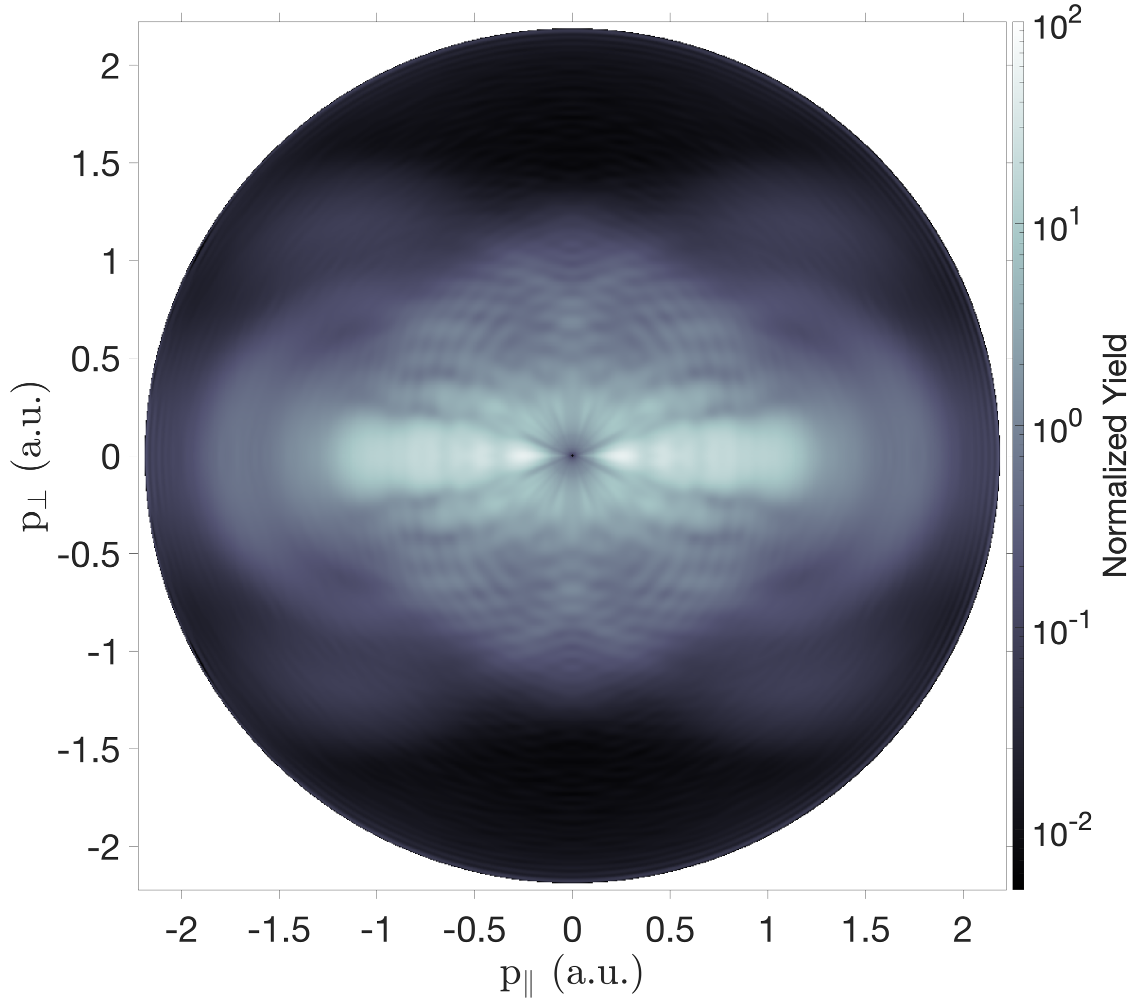

Figure 5.3: The PMD $Y^{0.85}(p_\parallel, p_\perp)$. Compared to the unfiltered PMD, $Y(p_\parallel, p_\perp)$, shown in Figure 4.3, the ATI rings have been completely removed without disrupting any underlying holographic structure.

completely obscured by the ATI rings is now visible in $Y^{0.85}(p_{\|}, p_{\perp})$. Another structure that becomes more prominent is a defined "wiggling" in the first and second order spider legs (though is present for all orders, see further discussion ahead). As previously discussed, the spider-leg structure is formed from two similar forward scattering orbits passing on opposite sides of the parent ion. This should yield regular horizontal fringes (See Figure 4.6(e)). The existence of the wiggling pattern on top of these spider legs is evidence of the contribution of a third orbit interfering with those of the spider legs. Going forward, I refer to these wiggles on top of the spider legs as "spider-leg modulations."

5.2.1 Intensity Calibrations Using the Time-Correlation Filter

In Chapter 4 we deferred the discussion of determining the peak intensity of the laser pulses. The Carpet and the spider-leg modulations revealed by time-correlation filtering are excellent references to the numerically accurate calculations of Qprop. The Carpet presents as clearly defined fringes along the transverse axis, whereas the spider-leg modulations form distinct nodes and antinodes along the first two spider legs. By generating multiple Qprop calculations using different laser peak intensities and directly comparing these structures, we can fairly accurately determine the peak intensity of the laser used in the experiment.

This comparison is shown in Figure 5.4. Again we are examining the experimental PMD $Y^{0.85}(p_{\|}, p_{\perp})$. In each PMD, the bottom left quadrant has been replaced by a Qprop calculation at the peak intensities (a) 100 TW/cm^2, (b) 125 TW/cm^2, and (c) 150 TW/cm^2. These Qprop calculations have also been filtered using TCF$^{0.85}$ to facilitate a direct comparison and remove the ATI rings. Panels 5.4(d) and 5.4(e) show the normalized electron yield along the parallel and transverse momentum axis respectively. On these plots, the experimental yield is shown in black, while the Qprop calculations for each selected intensity are colored according to the legends. We don't expect the calculated lines to fully overlap the experimental lines, as the experiment cannot replicate the contrast demonstrated in a Qprop calculation which is a result of having no background signal; rather, we want to see whether the local maxima and minima, which indicate the phase of the holograms, line up for each of them. Looking at the red line corresponding to a calculated laser intensity of 125 TW/cm^2, we see that for both the transverse and parallel axes there is better agreement between these maxima and minima than for the other two lines. Also by direct inspection of Panel 5.4(b), we see the calculated spectrum most closely matches the experimental spectrum. One especially notable feature that can be matched are the number of horizontal fringes corresponding to the spider-leg structure, which only aligns in Panel 5.4(b). Through these phase-matching comparisons to Qprop calculations at different intensities, we can conclude that the peak intensity of the experiment was approximately 125 ± 10 TW/cm^2, which has higher accuracy than what we are able to determine using simple classical cutoffs or estimations from the optical setup itself.

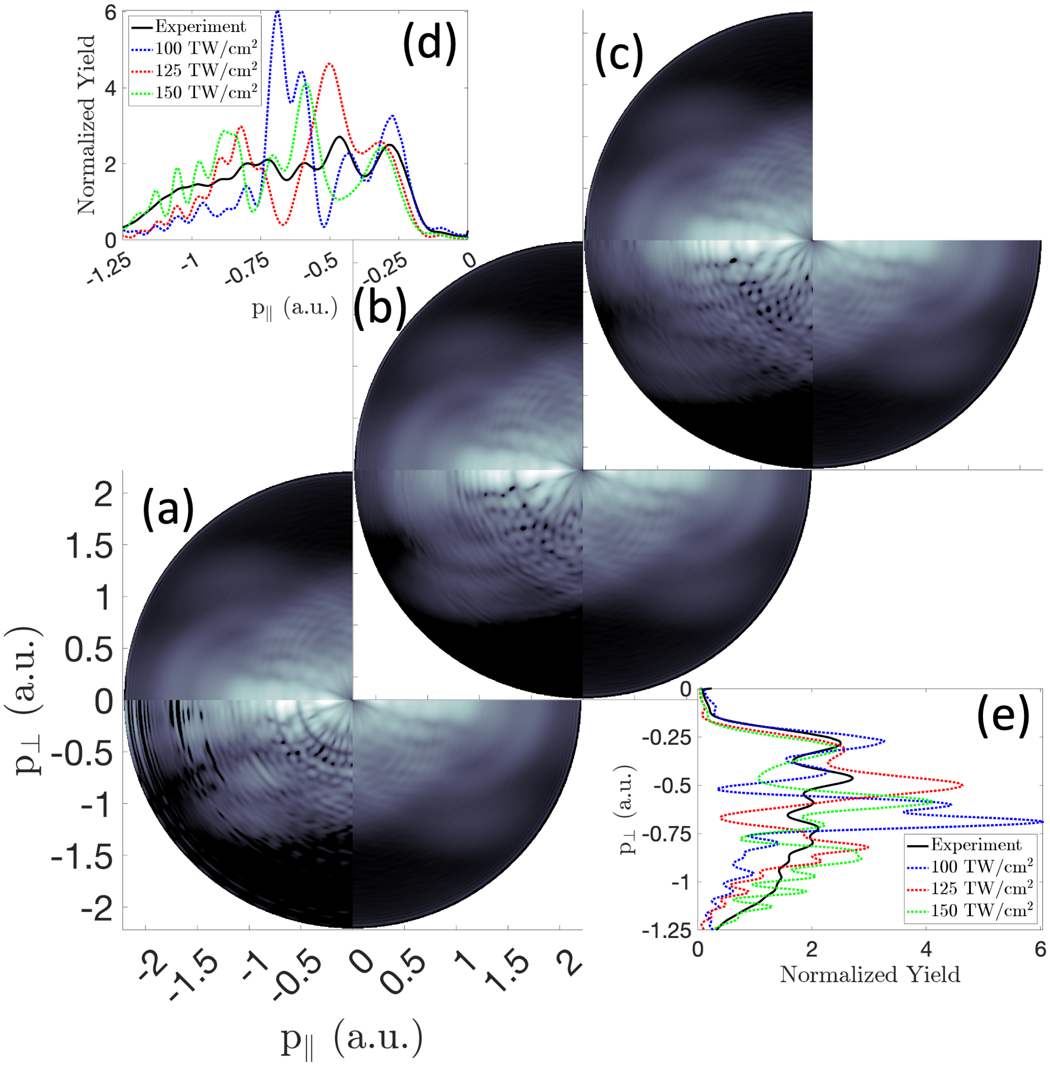

Figure 5.4: The experimental PMD $Y^{0.85}(p_\parallel, p_\perp)$ is presented in quadrants two and four of each of the panels and time-correlation filtered Qprop calculations at different laser peak intensities are presented in quadrant three of each panel. Panels (a), (b), and (c) are calculated using a peak intensity of 100 TW/cm^2, 125 TW/cm^2, and 150 TW/cm^2 respectively. Panels (d) and (e) show the normalized electron yield for the experimental PMD and each of the Qprop calculations along the parallel axis and transverse axis respectively. Panels (a), (b), and (c) are on the same colorscales and have the same momentum ranges as previous figures. Panels (b) and (c) have had the x- and y-axis labels removed for clarity.

5.3 Isolating Holograms with the Time-Correlation Filter

Direct measurements of strong-field ionization times are challenging for experimentalists, and often involve complicated shaped laser fields (see Chapter 6, as well as References [71] or [72] as examples). The TCF does not require any experimental modification and can be applied to all angularly resolve spectra to reveal temporal ionization information. The TCF that we have examined thus far is a useful tool to remove holographic structures generated by interfering pairs of trajectories whose ionization times are separated by more than some chosen time; however, by employing a pair of filtered spectra, we can also directly measure this time separation to within a quarter cycle of accuracy. Taking the normalized residual of two filtered spectra presents all the holographic structures that form in between the cutoffs of the component spectra. Explicitly we write:

$$N^{\alpha\rightarrow\beta}(p_{\|}, p_{\perp}) \equiv \frac{Y^{\beta}(p_{\|}, p_{\perp}) - Y^{\alpha}(p_{\|}, p_{\perp})}{Y^{\beta}(p_{\|}, p_{\perp}) + Y^{\alpha}(p_{\|}, p_{\perp})}. \tag{5.2}$$

Here I have introduced the quantity $N^{\alpha\rightarrow\beta}(p_{\|}, p_{\perp})$, which refers to the normalized residual of the filtered spectra with cutoffs α and β, with $\alpha < \beta$. This quantity reveals all holographic structures formed from trajectory pairs with time separations in the window $\alpha < t < \beta$, and thus I refer to $N^{\alpha\rightarrow\beta}(p_{\|}, p_{\perp})$ as a "time-correlation filter window." Using a normalized residual rather than a simple residual is useful for spectra with exponential dependence as it allows each pixel to be compared regardless of the overall magnitude of the subtracted pixels.

> Given that the TCF is constructed upon a low-pass filter to remove time separations above a given cutoff, it seems natural to use a band-pass filter to extract holographic structures within a given window of time separations. However, from empirical testing this method generates significantly more numerical artifacts than calculating the normalized residual of two low-pass filters. The numeric constraints on generating band-pass filters also result in a minimum pass-band width of approximately 0.4 cycles, which is of lower accuracy than the 0.25 cycle accuracy we can produce using normalized residuals.

5.3.1 Predictions of the Time-Correlation Filter

One way in which we might verify the validity of the TCF is by matching the features we can observe at specific time separations to specific predictions of our models of strong-field holograms. The three most prominent features to examine are the ATI rings, the spider-leg structure, and the Carpet. These features along with their launch-time separations are briefly summarized in Table 5.1. From these, we can make three concrete predictions for the behavior of the TCF. First, the

Structure	Time Separation (Cycles)	Qualitative Description
ATI Rings	1	Ring-like fringes centered on the origin with photon energy spacing
Spider	~ 0	Horizontal fringes offset from the parallel momentum axis
Carpet	~ 0.45–0.55	Checkerboard-like hologram near the transverse momentum axis

Table 5.1: A selection of three holographic structures whose predicted time separations can be used to verify the capability of the TCF.

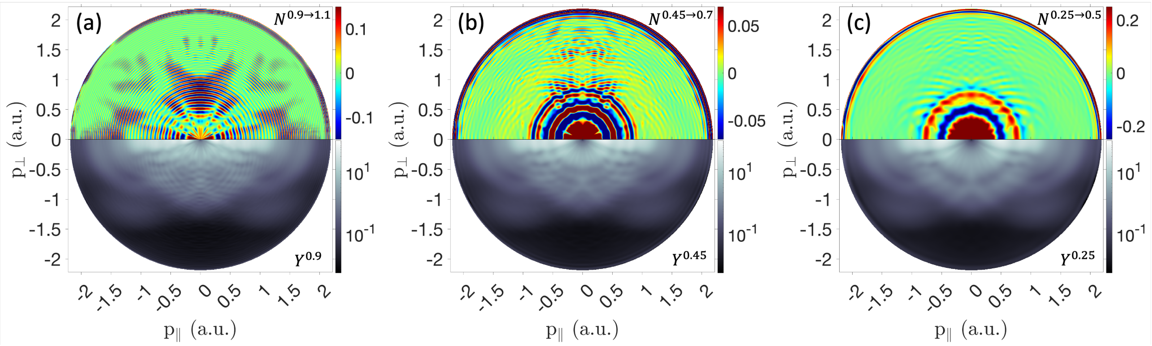

Figure 5.5: Three selected time-correlation filter windows with launch time-separation windows selected to emphasize notable holograms. The top half of each panel presents the time-correlation filter window indicated by the annotation at top right, and the bottom half shows the time-correlation filtered spectrum indicated by the annotation at bottom right for comparison. (a) The launch time-separation window 0.9 cycles to 1.1 cycles was selected to isolate the ATI interference. (b) The launch time-separation window 0.45 cycles to 0.7 cycles was selected to isolate the Carpet hologram. (c) The launch time-separation window 0.25 cycles to 0.5 cycles was selected to isolate holograms formed in this range which in particular include the spider-leg modulations (See Figure 5.7 below).

ATI rings are removed for TCF$^{\alpha<1}$. Second, the spider-leg structure should not be removed by any TCF at all cutoffs. Third, the Carpet structure should appear in a time-correlation filter window if a time separation range is selected to include the range 0.45–0.55 cycles (See Figure 4.5(b) and 4.6(f)). Let us examine these predictions.

In Figure 5.5 I present the time-correlation filter windows containing interesting launch-time separations. The bottom half of each panel contains the time-correlation filtered spectrum with the lower cutoff in the filter window residual above it. The top half of Panel 5.5 presents the time-correlation filter window $N^{0.9\to1.1}(p_\parallel, p_\perp)$. In this window, we see very clearly the ATI rings which we expect to see in this time-separation range. Paired with the discussion earlier about selectively removing ATI rings with TCF$^{0.85}$, this serves as concrete evidence for our ATI-ring prediction. Looking solely at the bottom halves of each panel, the spider-leg structure remains intact at all

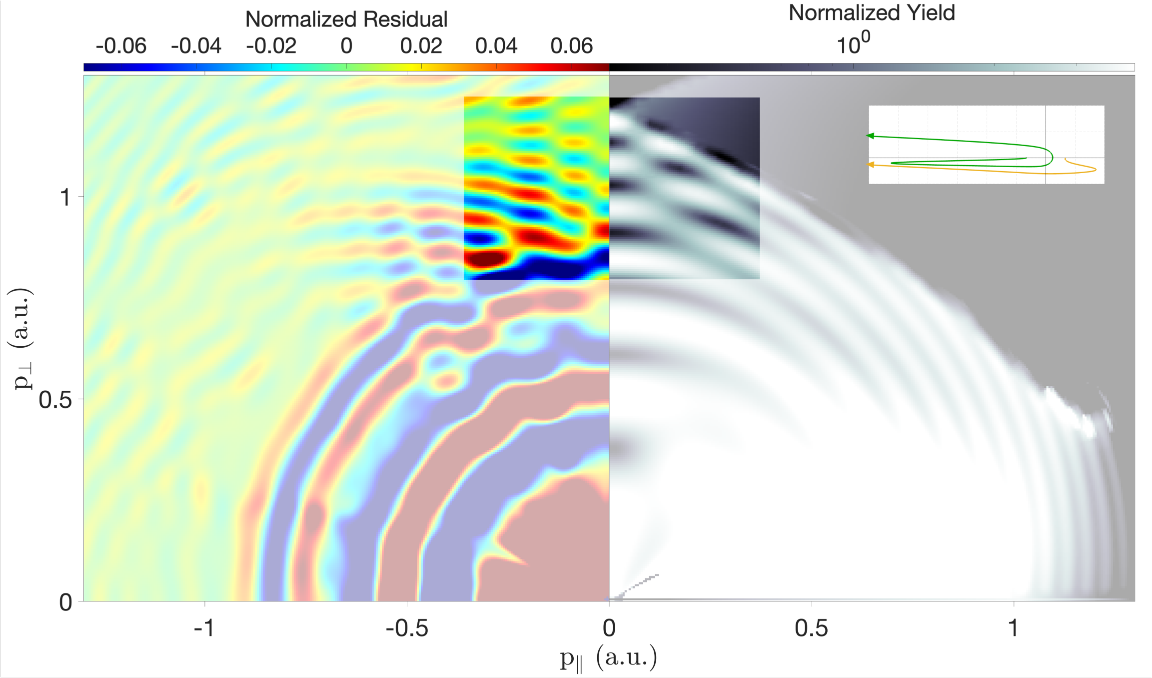

Figure 5.6: The experimental normalized residual $N^{0.45\to0.7}(p_\parallel, p_\perp)$ is presented in the left half of the plot, and the CQSFA calculation of the Carpet hologram formed from orbits #3 and #4 is presented in the right half. The highlighted rectangle in the center outlines the region of interest containing the Carpet hologram which can be unambiguously measured with the normalized residual. The inset at top right shows the form of the included CQSFA trajectories as in Figure 4.6. The colorscales and momentum limits are selected to emphasize the contrast in the region of interest.

cutoffs, even down to $Y^{0.25}(p_\parallel, p_\perp)$, presented in Panel 5.5(c). In the filter windows, there is also no evidence that the spider-leg structure has been removed; however, very interestingly Panel 5.5(c) shows that the spider-leg modulations have started being removed, which indicates that the third contributing trajectory is within the time-separation range of 0.25–0.5 cycles. This is explored in the following section to highlight how the TCF can make specific predictions leading to insight into which trajectories form holograms in strong-field spectra.

The final prediction says that the Carpet should form in the filter window presented in the top half of Panel 5.5(b). This is visible along the transverse momentum axis between approximately $0.8 < p_\perp < 1.2$ a.u.; however, it is difficult from just Panel 5.5(b) alone to determine whether we are truly seeing the Carpet hologram. In Figure 5.6 a zoomed-in version of $N^{0.45\to0.7}(p_\parallel, p_\perp)$ is presented alongside the CQSFA calculation of the interference of orbits #3 and #4 forming the Carpet. A region of interest is highlighted showing the portion of the momentum spectrum in which the Carpet forms. Outside of this region of interest, $N^{0.45\to0.7}(p_\parallel, p_\perp)$ reveals interference patterns which are not replicated in the CQSFA comparison, as they are not solely the result of the interference of CQSFA orbits #3 and #4. Within this region, however, there are no other predicted trajectories which may interfere and confound a direct comparison, which is why we restrict our attention to just this area. A similar checkerboard-like pattern appears for both the filter window and the CQSFA calculation, with similar fringe spacing and locations. This comparison allows for an unambiguous determination that the Carpet is observed as expected in the filter window $N^{0.45\to0.7}(p_\parallel, p_\perp)$ as predicted.

Interestingly, while it satisfied the prediction of the CQSFA, the measured result that the Carpet appears in the range 0.45 cycles to 0.7 cycles is non-trivial. In fact, the orbits used by the CQSFA to generate the Carpet are in conflict with the original explanation of the origin of the Carpet derived from the strong-field approximation (SFA), which attributed the interference pattern to the equivalent of CQSFA orbits #1 and #2 [40]. The developers of the CQSFA disagreed with this explanation, and claimed that the Carpet hologram observed in experiments and TDSE calculations [17] arises from the interference of orbits #3 and #4. This disagreement persisted because in their respective formalisms, any linearly-polarized, monochromatic field will generate interference patterns which have a 2ω fringe spacing near the transverse momentum axis [2], and so both sets of orbits could generate a carpet-like interference pattern. However, the measurement made here settles the debate, as due to the time-ordering of orbits #1 and #2, a carpet-like hologram forming from the interference of these orbits would occur from the combination of approximately 0.3 cycles of time separation and 0.7 cycles of time separation [2]. Because both of these values are outside the range of $N^{0.45\to0.7}(p_\parallel, p_\perp)$ and yet the Carpet is still observed, we can conclude that the CQSFA explanation is more correct. Furthermore, the momentum range at which the Carpet can be observed in the filter window provides further evidence that the SFA explanation is incorrect, because the maximum allowed energy of the direct electrons forming orbit #1 is less than the energy of the electrons observed to form the Carpet.

With the three predictions spanning a wide range of possible time separations satisfied, let us turn our attention to novel measurements using time-correlation filter windows to uncover new types of holographic interference structures.

5.3.2 Disentangling Multi-Path Interference

The CQSFA model used in this thesis calculates four distinct electron orbits. At the highest range of parallel momentum there is only the contribution from the backscattered orbit (#4). In the Carpet there is only the contribution from the backscattered orbit and the forward scattered orbit (#3). Everywhere else, all four orbits contribute simultaneously, which makes disentangling holograms (which always form from *pairs* of interfering orbits) exceptionally difficult. Ordinarily, experimental spectra cannot overcome this problem; however, time-correlation filter windows can circumvent the overlapping multi-path interference by isolating contributions from individual pairs of trajectories.

Consider as a representative example the spider-leg modulations. It is clear from their location that the spider-leg modulations are at least a three-path interference involving the orbits forming the Spider (#2 and #3) and one or more other orbits [36]. Since the spider-leg modulations are observed in the full CQSFA calculation with all four orbits, see Figure 4.6(g), we can surmise that either orbits #1 or #4 (or both) are contributing. If the hologram were formed from the pair of orbits #4 and #3 or the pair of orbits #4 and #2 then we would observe the spider modulations forming in the filter window $N^{0.45\to0.7}(p_\parallel, p_\perp)$, as the time separation between those pairs is approximately 0.5 cycles. Since we do not observe them, we can conclude orbit #4 does not contribute in a pair with orbits #2 or #3. If the pair included orbit #1, then the possible time separations would be either 0.3 cycles for orbits #2 or #3 preceding orbit #1 or 0.7 cycles for orbits #2 or #3 following orbit #1. Using the time-correlation filter window $N^{0.25\to0.5}(p_\parallel, p_\perp)$, we can determine not only the additional orbit in the multi-path hologram, but also the time-ordering of orbits giving rise to it.

In the top half of Panel (c) of Figure 5.5, a pattern of lobes can be observed at a radial momentum of 0.8 a.u. These lobes exactly trace the second-to-last spider-leg modulation which can be observed in Figure 5.3. Also in that panel are azimuthal distributions of other lobes corresponding to the other visible modulations; however, I focus on the pattern at 0.8 a.u. as it has the highest contrast in $N^{0.25\to0.5}(p_\parallel, p_\perp)$. In Figure 5.7, I convert the Cartesian filter window to rectangular polar coordinates and restrict the range of radial momenta p_r to focus on the pattern of lobes at 0.8 a.u. A dashed line is drawn at 0.8 a.u., and the amplitude along that line is plotted as a dashed line in Panel 5.7(c). A careful counting of the second-to-last spider-leg modulations at radial momentum $p_r = 0.8$ a.u. in Figure 5.3 reveals seven distinct lobes in a single quadrant, which matches with the seven peaks visible in the experimental line-out in Panel 5.7. Because this structure can be observed within the time-separation range of 0.25–0.5 cycles, we may further conclude that the modulations are the result of the pairwise interference of orbit #1 with both orbits #2 and #3 in the specific time-ordering where orbit #1 follows orbits #2 or #3.

For completeness, in Panel 5.7(b) a CQSFA calculation is presented for comparison. Recall from Equation 2.23 that we are able to modify the spectra calculated from CQSFA to remove interference contributions from orbits separated by more than a defined parameter τ cycles. In order to facilitate a one-to-one comparison with the filter window in Panel 5.7(a), the spectrum in Panel 5.7(b) is the normalized residual of the time-windowed CQSFA calculations at $\tau = 0.5$ and $\tau = 0.25$. This has the effect of restricting the observed holographic structure to only the combinations [#1, #2] and [#1, #3], with orbit #1 following the other in each case. The CQSFA residual also reveals a pattern of lobes with approximately the same position and spacing as the pattern uncovered by the experimental filter window. This agreement further corroborates the validity of the TCF as a tool to make quantitative measurements of ultrafast electron interference.

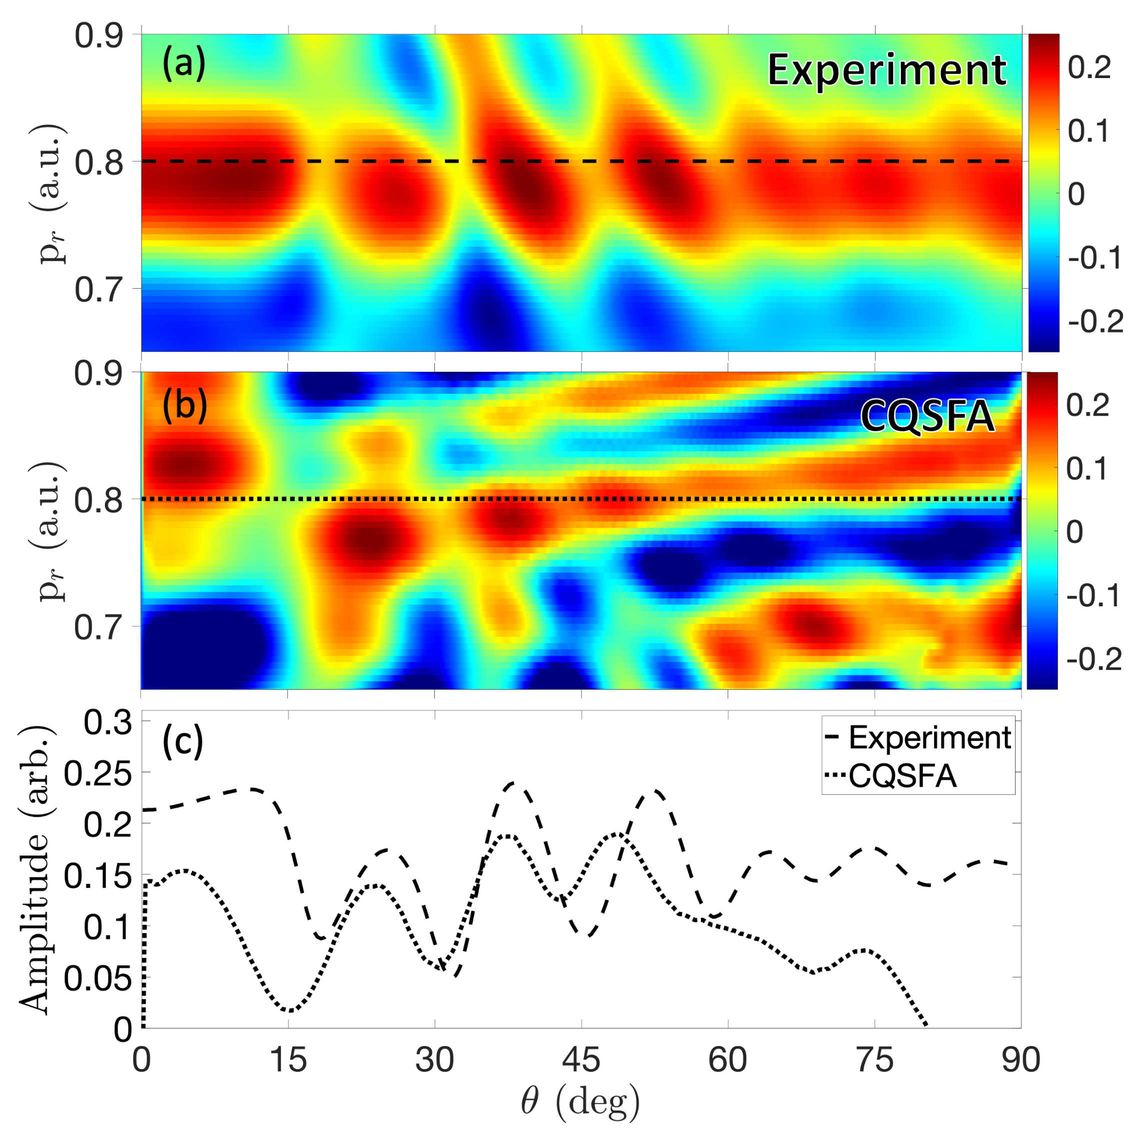

Figure 5.7: Three panels highlighting a hologram observable by restricting the range of the launch time separations of interfering trajectories to the window 0.25-0.5 cycles. Each are in rectangular polar coordinates, with θ as the azimuthal angle and p_r as the radial momentum. (a) This panel presents the experimental normalized residual $N^{0.25\rightarrow 0.5}(p_r, \theta)$. (b) This panel presents the polar normalized residual of the time-windowed CQSFA at $\tau = 0.5$ and $\tau = 0.25$. (c) This plot shows the amplitude of the above two panels at $p_r = 0.8$a.u. as indicated by the dashed and dotted lines.

5.4 Limitations of the Time-Correlation Filter

One of the clearest ways to evaluate where errors form as a result of the time-correlation filter is to directly compare the effect of the TCF and an analytical approach to time-windowing. Since we are able to analytically evaluate CQSFA calculations at defined time separations, we can simply compare time-correlation filtered CQSFA calculations to their equivalent analytically time-windowed counterparts. This is shown in Figure 5.8 for the orbits forming the Carpet. This figure examines three time-separation ranges: 0.75 cycles to 1 cycle, 0.5 cycles to 0.75 cycles, and 0.25 cycles to 0.5 cycles in the panels respectively from left to right. Because in the CQSFA the full two-orbit interference occurs in two separate ranges (0.3 cycles to 0.5 cycles for orbits #3 followed by #4 and 0.5 cycles to 0.7 cycles for orbits #4 followed by #3) the analytic residuals in the upper right quadrants of each plot exactly isolate these ranges, with Panel 5.8(b) identically zero as its time separation range is outside the range of allowed separations. Thus the most interesting question to answer would be whether or not the time-correlation filter can also separately extract these time-orderings. The time-correlation filter windows in Panels 5.8(e) and 5.8(i) nearly perfectly match their analytical counterparts along the transverse axis, indicating that the time-correlation filter does indeed have the time-resolution in this region to resolve time-orderings fairly closely spaced around 0.5 cycles.

However, Panel 5.8(a) shows some limitations. As a numerical approach, the time-correlation filter cannot effectively generate a result that is identically zero, and a small amount of the total signal is removed in the range 0.75 cycles to 1 cycle, which then is emphasized in the filter window. In these circumstances it is helpful to also directly examine the time-correlation filtered spectrum to assess the degree to which the spectrum has had holograms removed. In Panels 5.8(c) and 5.8(d) we have this comparison. We note that along the transverse axis, the Carpet hologram is fully intact, indicating that the TCF correctly excludes this region at this cutoff; however, along the parallel momentum axis the spectrum is rather significantly filtered, indicating that there is an azimuthal dependence on the error of the TCF.

The TCF implicitly relies on the observation that holographic fringe spacings in strong-field PMDs are dominated by a term proportional to $E(t_2 - t_1)$, where $(t_2 - t_1)$ is the ionization launch-time separation of the contributing pairs of electron orbits [2]. In regions of the PMD in which this observation holds, the TCF accurately extracts holograms within narrow ranges of time separations. When evaluating the difference in the semi-classical action (Equation 2.14) between two orbits, the CQSFA predicts additional terms which are not proportional to the launch time separation [1]. These additional terms drop to zero around the transverse momentum axis and are maximal along the parallel momentum axis. When numerically evaluated in the maximal case along the parallel momentum axis, we arrive at a maximum disagreement of approximately 0.25 cycles of launch time separation [2]. This suggests that the launch-time separations measured by the TCF have the least error along the transverse axis, with error smoothly increasing to a maximum of 0.25 cycles as we

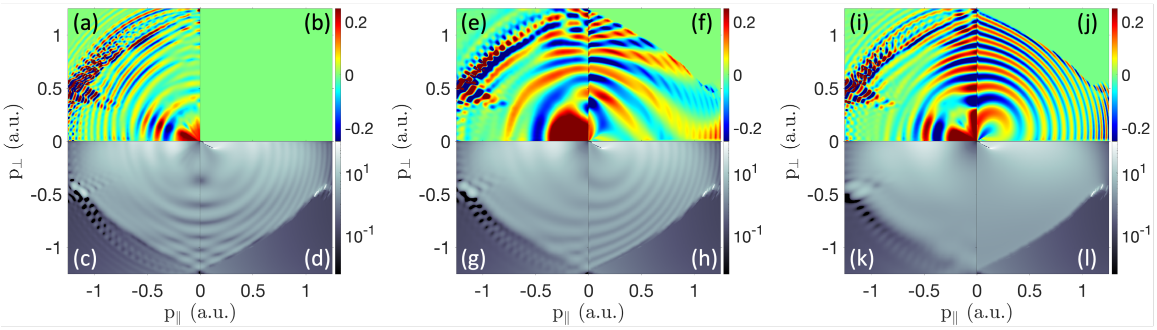

Figure 5.8: Comparison between the TCF acting on the CQSFA calculation of the Carpet and the equivalent analytically calculated time-windowed CQSFA spectra. The left half of each panel are results from the TCF, while the right halves are the results of analytic time-windowing. The upper halves show the normalized residuals of time-correlation filtered calculations (left) and analytic time-windowed CQSFA calculations (right). The bottom halves show the time-correlation filtered calculation at the lower cutoff in the residual above (left) and the analytic time-windowed CQSFA calculation at the equivalent value of τ. (a, b) Residuals evaluated in the time-separation range of 0.75 cycles to 1 cycle. (c, d) Calculations evaluated at 0.75 cycles. (e, f) Residuals evaluated in the time-separation range of 0.5 cycles to 0.75 cycles. (g, h) Calculations evaluated at 0.5 cycles. (i, j) Residuals evaluated in the time-separation range of 0.25 cycles to 0.5 cycles. (k, l) Calculations evaluated at 0.25 cycles.

progress azimuthally towards the parallel momentum axis. Even in its worst case, this error still allows us to make definitive conclusions regarding whether an observed hologram is formed from orbits both launched from the same half-cycle, or whether it is the result of electron orbits launched from opposing half-cycles. Along the transverse momentum axis the TCF's accuracy is even greater, allowing for unambiguous determination of sub-quarter-cycle temporal resolution in launch-time separations.

Chapter 6

Probing Ionization Phases with Shaped Laser Fields

The previous chapter focused on an analysis technique which revealed ionization time separations between pairs of trajectories in any angle-resolved strong-field experiment. In this chapter I discuss how shaping the strong laser field with two or more laser pulses can reveal ionization phases for individual electron trajectories. I introduce two new experiments and associated data collections, one involving two overlapped laser pulses and one involving three. In particular, I dive into the experimental technique of "phase-of-the-phase," which involves overlapping the fundamental strong laser field with its relative-phase-controlled second harmonic perturbation and measuring the phase of the pixel-by-pixel oscillations in the resulting spectrum. After discussing the strengths and weaknesses of the technique, I introduce a three-color replacement technique to succeed the phase-of-the-phase technique that accurately retrieves electron trajectory ionization phases from multiple trajectories.

6.1 Shaping Strong Fields with Harmonic Perturbations

When two laser pulses overlap in the interaction region of the VMI, the resulting strong field is the sum of the electric fields of the contributing pulses. This can be expressed rather simply for any number of pulses:

$$E(t) = \sum_n E_n(t). \tag{6.1}$$

Let's continue to assume the form of $E_n(t) = E_n \cos(\omega_n t - \phi_n)$, which omits the carrier envelope of the pulse and is a valid approximation for examining sub-cycle phenomena emerging from laser pulses exceeding approximately 10 cycles in duration. If the summed pulses have frequencies which are some integer multiple of the fundamental ionizing frequency, then the summed field will maintain the periodicity of the original ionizing pulse. This is a crucial fact that allows us to design an experiment which modulates the probabilities of electron trajectories in a controlled manner across each contributing cycle of the ionizing laser.

The first half of this chapter focuses on the simplest such case: modulating the fundamental (ω) pulse with a small second harmonic (2ω) perturbation. In this case we can treat the summed field as

$$E(t) = E_\omega \cos(\omega t - \phi_\omega) + E_{2\omega} \cos(2\omega t - \phi_{2\omega}). \tag{6.2}$$

Since the frame of reference of our experimental measurements is locked with respect to the phase of the ionizing pulse, we define the origin of our coordinate system such that the phase $\phi_\omega = 0$, and thus $\phi_{2\omega}$ describes the relative phase between the fundamental pulse and the second harmonic pulse. Although it can be determined by Equation 6.2, it can be difficult to keep track of the harmonic relative phases and how they relate to the fundamental. For clarity, note that advancing $\phi_{2\omega}$ by 360 degrees is equivalent to advancing $\phi_{2\omega}$ by *half* a laser cycle of the fundamental. This can be observed in Figure 6.2(a)-(e). Throughout this chapter, I refer to quantities of phase in units of degrees and restrict values of phase to the range [0,360], except when units of laser cycles are needed to facilitate a discussion.

6.2 Two-Color Data

The results discussed in the upcoming sections use data collected from the following experiment. We again are looking at the PMDs of argon gas photoionized by two overlapped collinear laser pulses, one at 800 nm and its second harmonic at 400 nm. The 800 nm pulse was focused to an approximate intensity of 90 TW/cm^2, corresponding to $\gamma \approx 1.211$, while the 400 nm pulse was focused to an approximate intensity of 0.9 TW/cm^2. The intensity of the fundamental was evaluated by comparing to Qprop calculations filtered with TCF$^{0.85}$ as detailed in Figure 5.4, and the intensity of the second harmonic was determined from the ratio of the measured pulse powers on the optical table prior to the experiment. The relative phase of the 800 nm and 400 nm pulses is controlled using glass wedges as outlined in Section 3.4.3. A PMD was collected at 151 evenly spaced glass wedge positions, spanning a relative phase shift of approximately 12 full cycles of the second harmonic. This collection was tuned such that each interaction event between the laser and the pulsed argon gas generated approximately 200 photoelectrons. In total, this collection ran for 18 hours, and 2.5 billion electron counts were recorded, resulting in an approximate distribution of 17 million electron counts per PMD. Other than these parameters, all data collection parameters are the same as described in Section 4.1.

A representation of the data is presented in Figure 6.1. In panel 6.1(a) a stacked subset of the first five raw PMDs is shown. A region of interest is highlighted by a dotted rectangle corresponding to a small region of backscattered electrons. Panel 6.1(b) shows the summed electron counts within

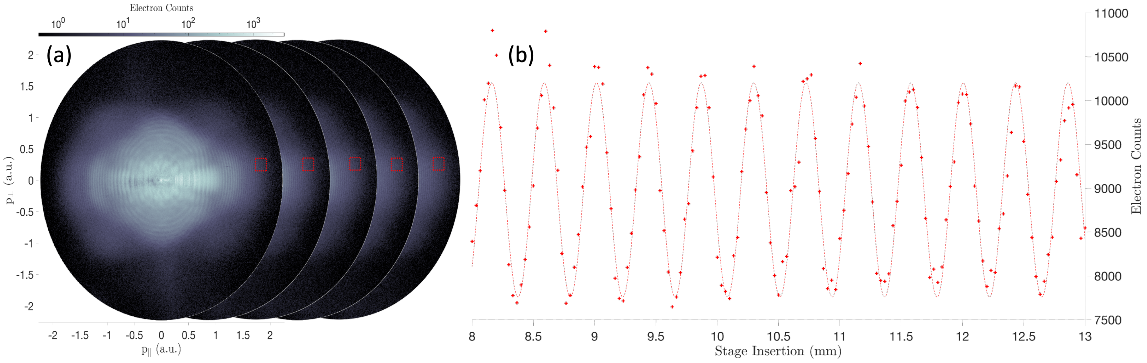

Figure 6.1: Representation of the two-color PMDs. (a) Five PMDs corresponding to the first five stage insertion positions are shown. A region of interest is highlighted within a dotted red rectangle. (b) The summed electron counts within the highlighted region of interest for all stage insertion positions. The dashed red sinusoid is the cosine fit to the data.

that region of interest with respect to glass wedge stage insertion. The glass wedge stage insertion maps linearly to the relative phase of the second harmonic, and can be calibrated by observing the amount of insertion insertion needed for one full oscillation of electron counts, which corresponds to delaying the 400 nm perturbation by a full cycle. We may also note that the amplitude of these oscillations drops slightly with increased stage insertion, which results from shifting the center of the envelope of the two pulses off from each other. This effect is fairly small over nearly 12 cycles of 400 nm shift, indicating that the pulse duration is sufficiently long that the envelope shift from the difference of the group velocities of the pulses in glass can mostly be ignored.

6.3 Effects of the Second Harmonic Perturbation

From Figure 6.1 we see that the electron counts of the backscattered electrons vary in accordance with the relative phase of the second harmonic. This can be explained quite effectively within the electron trajectory model of strong-field ionization. First let's assume that only a single electron trajectory can result in the final momentum of a given pixel on the detector. In our semi-classical approximation we can define electron trajectories by their ionization phase, ωt_0, and their initial transverse momentum, v_0, before tracing their motion in the electric field. Then we can express the predicted amplitude of the pixel that trajectory arrives at as

$$Y(p_{\|}, p_{\perp}) \propto W_0(t_0) W_{\perp}(t_0, v_0), \tag{6.3}$$

where $W_0(t_0)$ is the ADK tunnel ionization rate of electrons at the ionization phase ωt_0 given by Equation 2.7 and $W_{\perp}(t_0, v_0)$ is the value of the transverse momentum distribution given by Equation 2.9 for the specified initial velocity v_0.

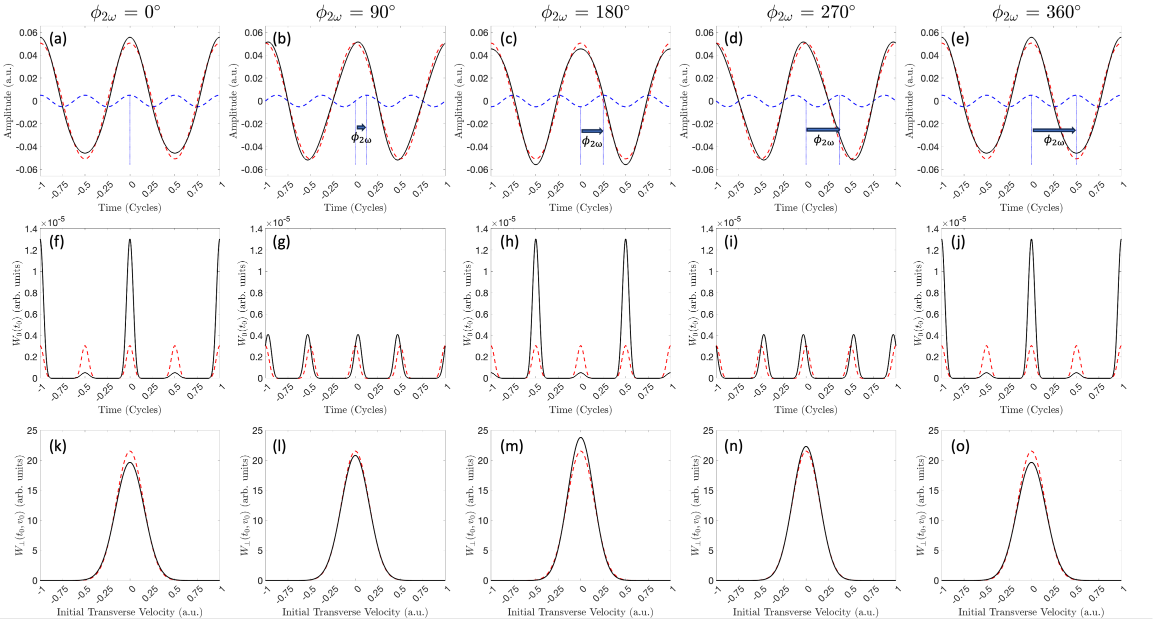

Figure 6.2: The result of summing the fundamental electric field with a small second harmonic perturbation. Each column is calculated from the value of $\phi_{2\omega}$ shown at the top. Panels (a) through (e) show the amplitude of the summed electric field (solid black) compared to the the fundamental field (red dashed) and the second harmonic perturbation (blue dashed). Panels (f) through (j) show the ionization rate of the summed field (solid black) compared to the ionization rate of the fundamental field alone (red dashed) as calculated from Equation 2.7. Panels (k) through (o) show the initial velocity distribution of the summed fields (solid black) compared to the initial velocity distribution (red dashed) as calculated from Equation 2.9 at the value $\omega t_0 = 10°$.

What happens when the second harmonic perturbation is added in? First, because the amplitude of the electric field is now dependent on the value of $\phi_{2\omega}$ as in Equation 6.2, both the tunnel ionization rate and the transverse momentum distribution become dependent on this phase as well. In Figure 6.2, I show the effect this phase has on the quantities $W_0(t_0)$ and $W_\perp(t_0, v_0)$ at five values of $\phi_{2\omega}$ indicated at the top of each column. The summed electric field is shown in the upper row, the ionization rate $W_0(t_0)$ is shown in the middle row, and the velocity distribution $W_\perp(t_0, v_0)$ with the selected value $\omega t_0 = 10°$ is shown in the bottom row. In this figure, the ratio $E_\omega/E_{2\omega} = 10$ in accordance with the experimental two-color data collected, and I continue to use these values for the rest of this discussion. Despite the relatively small difference in the shape and amplitude of the summed electric field shown in the first row, we can see that the second harmonic causes an enhancement of $W_0(t_0)$ when the two colors constructively interfere, and a suppression $W_0(t_0)$ when the two colors destructively interfere. This indicates that the magnitude of $W_0(t_0)$ is maximal for a specific ionization phase ωt_0 when the appropriate positive or negative peak of the 2ω perturbation exactly aligns with ωt_0 to maximize the value of $|E(t_0, \phi_{2\omega})|$. On the other hand, the value of $W_\perp(t_0, v_0)$ changes relatively little as $\phi_{2\omega}$ changes, indicating that the amplitude of a given trajectory in the presence of a second harmonic perturbation depends primarily on the change in amplitude of the ionization rate, $W_0(t_0)$.

Note that the final momentum and phase of a given trajectory is set by the path the trajectory follows within the laser field, which is also affected by the changing phase of the second harmonic perturbation. However, the force acting on the electron is linear with respect to the laser field, and as such a small second harmonic perturbation such as what we are examining here has only a minor effect on electron trajectories. From calculations using these field parameters, this deviation is on the order of a single pixel of final momentum deflection, which is mostly negligible. For the upcoming sections, I operate under the approximation that in any given pixel on the detector the specific trajectories arriving at that pixel do not change regardless of the value of $\phi_{2\omega}$, and thus the only changes are the amplitudes of the trajectories as determined primarily by the changing value of $W_0(t_0)$.

6.4 Two-Color Phase-of-the-Phase

Based on the previous observations, we can analyze each pixel of the PMD individually to reveal information about the ionization phases of the trajectories arriving there. Specifically we model each pixel's oscillation as [35]

$$Y(p_\parallel, p_\perp, \phi_{2\omega}) = P(p_\parallel, p_\perp)\cos(\phi_{2\omega} - \phi_{opt}(p_\parallel, p_\perp)) + C(p_\parallel, p_\perp), \tag{6.4}$$

where $(p_\parallel, p_\perp)$ are the coordinates of the examined pixel, $P(p_\parallel, p_\perp)$ is the amplitude of the oscillation which I refer to as the “phase contrast,” $C(p_\parallel, p_\perp)$ is some constant offset, and $\phi_{opt}(p_\parallel, p_\perp)$ is the phase at which the modeled signal is maximal.

Of these quantities, ϕ_{opt} is of particular interest for directly measuring ionization phases of electron trajectories. As stated earlier, the amplitude of a trajectory arriving at a pixel is maximal when the second harmonic aligns with that trajectory's ionization phase, which may be recorded by ϕ_{opt}. The analysis technique in which the sinusoidal model from Equation 6.4 is fit using linear least squares to each pixel in two-color spectra to extract the parameter ϕ_{opt} is called "phase-of-the-phase" (PP) [35]. With this terminology, I also describe Equation 6.4 as the "phase-of-the-phase fit," or simply "PP fit."

A PP analysis of the experimental data is presented in Figure 6.3. On the top half of the figure is the value of ϕ_{opt} extracted from each pixel. On the bottom half is the phase contrast, $P(p_{\parallel}, p_{\perp})$. The uncertainty in these parameters at each pixel is retrieved from the least-squares fit by calculating half the range of each parameter's 95% confidence interval. Note that the PP analysis is performed on the raw PMDs before any Abel inversion is performed. This is because the number of electrons collected in each individual raw PMD is insufficient to perform a high-fidelity Abel inversion. Using the raw PMDs slightly complicates the interpretation of the results near the center of the spectra, but the effect is relatively small. Before proceeding with the results of the PP analysis, some effort must be taken to understand how to properly interpret the quantities $\phi_{opt}(p_{\parallel}, p_{\perp})$ and $P(p_{\parallel}, p_{\perp})$.

Because ϕ_{opt} is a phase related to the shifting of the second harmonic perturbation, it must be translated to a phase of the fundamental ionizing laser in order to make useful conclusions about the ionization phases of contributing electron trajectories. Unfortunately, ϕ_{opt} does not uniquely map onto the values of ionization phase, ωt_0. If we look back at the first row of Figure 6.2, it can be observed that if we select a value of ωt_0 in the range $[0, 90°]$ where the fundamental field is positive, then the value of ϕ_{opt} must be set in the range $[0, 180°]$ such that the positive peak of the second harmonic may constructively interfere with the positive fundamental field. However, if ωt_0 is selected in the range $[90°, 270°]$ where the fundamental field is negative, then the value of ϕ_{opt} must be set in the range $[0, 360°]$ to overlap the *negative* peak of the second harmonic. Thus the measured value of ϕ_{opt} always yields two distinct candidate ionization phases, which can be described by

$$\phi_{C,1} = \frac{[(\phi_{opt} + 180°) \mod 360°] - 180°}{2} \mod 360° \tag{6.5}$$

where $\mod 360°$ is the modulo operator with a modulus of 360° and $\phi_{C,1} \in [0, 90°]$ or $[270°, 360°]$, and

$$\phi_{C,2} = \frac{\phi_{opt} + 180}{2} \tag{6.6}$$

where $\phi_{C,2} \in [90°, 270°]$. In some circumstances, we are able to reduce the number of ionization phase candidates to a single unique ionization phase either by incorporating predictions from theory or by selecting the candidate whose value is closer to 0 or 180° as they have the highest rate of ionization. However, as is discussed in Section 6.4.1, there often isn't a rigorous way to proceed to

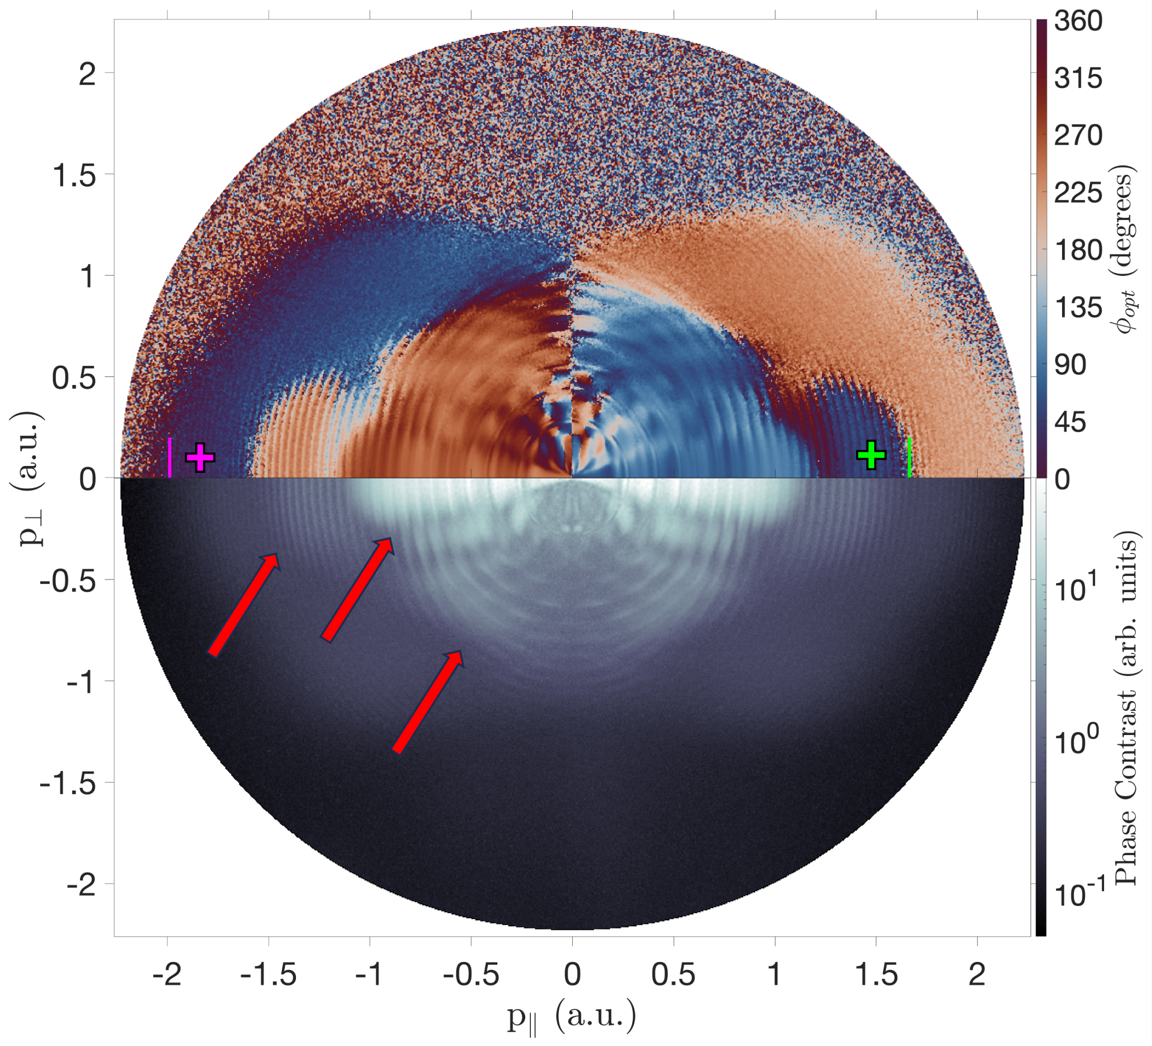

Figure 6.3: The result of a PP analysis on the experimental data showcased in Figure 6.1. The top half shows $\phi_{opt}(p_\parallel, p_\perp)$ and the bottom half shows $P(p_\parallel, p_\perp)$ calculated by performing the fit shown in Equation 6.4 on each pixel of the experimental spectrum. On the $\phi_{opt}(p_\parallel, p_\perp)$ spectrum, a magenta line on the left indicates the $10U_p$ cutoff for first-return backscattered electrons ionized when the fundamental field was positive. The green line on the right indicates the $7U_p$ cutoff for second-return backscattered electrons ionized on the same half-cycle. The magenta cross indicates the region at which $\phi_{2\omega}$ is calibrated to correctly enhance first-return backscattered electrons. The green cross indicates a lobe which is ionized a half-cycle apart from the neighboring backscattered electrons (see text). The red arrows on the $P(p_\parallel, p_\perp)$ spectrum point to minima which indicate transitions in the subset of trajectories arriving at regions on the detector.

determine a proper ionization phase due to the number of trajectories contributing to a single ϕ_{opt} measurement. Importantly, this limitation only exists in cases where we are interested in converting ϕ_{opt} to an ionization phase. It is a trivial matter to determine the ϕ_{opt} that should be measured given an ionization phase by selecting the value of ϕ_{opt} that satisfies either Equation 6.5 or Equation 6.6 depending on the phase range of the given ionization phase.

The phase contrast is a measure of the amount of enhancement caused by the relative phase of the second harmonic. Primarily, the phase contrast is used to determine the boundaries between areas on the PMD with different sets of contributing trajectories. As is explained in more detail in Section 6.4.1, when two trajectories contribute to the amplitude of a pixel in the spectrum and have values of ϕ_{opt} offset by half a cycle with respect to each other, their total contribution destructively interferes. When crossing from a region dominated by one trajectory into a region dominated by another, as the amplitudes of each trajectory start becoming equal, the phase contrast approaches a minimum.

When performing the fits described by Equation 6.4 on real data, the proper values of $\phi_{2\omega}$ to use are not clearly calibrated during the data collection. There is clearly a “correct” phase ruler to use, which would define $\phi_{2\omega} = 0$ as a point in which the second harmonic perturbation perfectly constructively interferes with the fundamental at its maximum positive value, which can be seen in Figure 6.2(a). To calibrate these values, I return to the discussion of the SMM in Chapter 4. During that discussion, we calculated that the maximum momentum which can be achieved by an electron in the negative direction occurs from first-return backscattered electrons, which ionize at an ionization phase of 15.48 degrees. Thus if we first perform the fit at the location of first-return backscattered electrons in the PMD, we can use Equation 6.5 to calibrate the values of $\phi_{2\omega}$ to yield $\phi_{opt} = 31°$, which would correspond to an enhancement at $\omega t_0 = 15.5°$ on the fundamental laser. In Figure 6.3, the $10U_p$ cutoff is indicated by a magenta line on the left side of the ϕ_{opt} spectrum. Slightly to the right, a magenta cross indicates the region on the detector used to calibrate $\phi_{2\omega}$ according to the predicted phase of first-return backscattered electrons.

On the right side of the ϕ_{opt} spectrum, the $7U_p$ cutoff is indicated by a green line. This cutoff denotes the maximum extent of the second-return rescattered electrons predicted by the SMM. Observing the region indicated by the green cross to the left of the $7U_p$ cutoff, we see that the value of ϕ_{opt} in this area is 180° out of phase with the region of first-return backscattered electrons immediately to the right of the cutoff. This indicates that the electrons within this region are ionized on the opposite half-cycle compared to the neighboring first-return backscattered electrons, which matches the predictions of SMM for second-return rescattered electrons. In the high-resolution experimental dataset Figure 4.4(b), this feature corresponds to the ring of enhanced signal bordered by the blue ring indicating the maximum extent of second-return rescattered electrons. The prominent phase shift visible in the ϕ_{opt} spectrum around this region is an unambiguous measurement of a single-trajectory feature formed from second-return rescattered electrons.

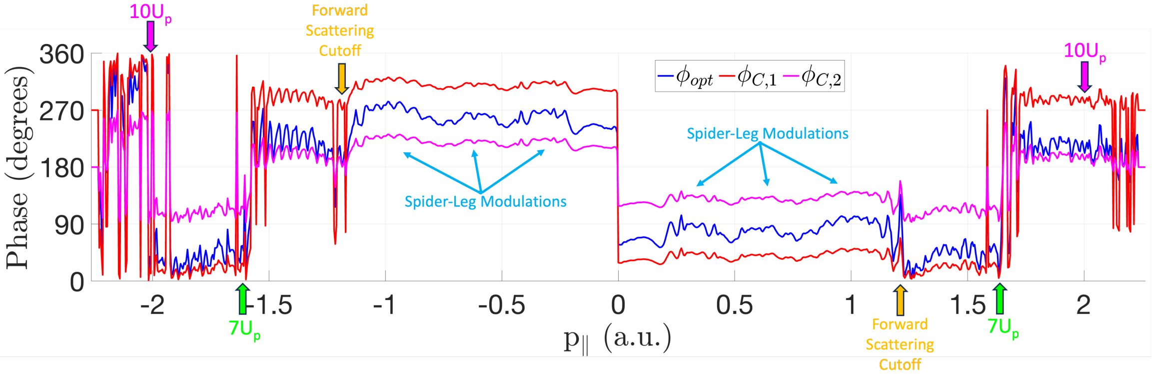

Figure 6.4: The value of ϕ_{opt} (in blue) from the PP analysis in Figure 6.3 plotted along the parallel momentum axis. In addition, the two candidate ionization phases $\phi_{C,1}$ and $\phi_{C,2}$ calculated from Equations 6.5 and 6.6 are plotted. Arrows with annotations highlight features of interest.

In Figure 6.4 I have plotted the values of ϕ_{opt} and its associated ionization phase candidates along the parallel momentum axis. The most obvious feature we can observe is a 180° phase shift in ϕ_{opt} occurring at $p_{\|} = 0$, which is also visible in the full $\phi_{opt}(p_{\|}, p_{\perp})$ spectrum in Figure 6.3. The 180° phase shift is also visible in the ionization phase candidates except that $\phi_{C,1} \leftrightarrow \phi_{C,2}$ after the shift. There are three distinct phase transitions that can be observed on each side of the momentum spectrum, occurring at approximately $p_{\|} = \pm 1.2$ a.u., ± 1.65a.u., and ± 2 a.u. The final two of these correspond to the $7U_p$ and $10U_p$ cutoffs respectively denoting the rescattering transitions as described above. The shift at $p_{\|} = \pm 1.2$ a.u. corresponds to the maximum extent along the polarization axis of forward scattered electrons ionized on the same half-cycle as the second-return rescattering electrons. In the region of the ϕ_{opt} line-out from 0 to this cutoff, cyan arrows point to small observable ridges. These correspond to the primary spider-leg modulations along the first-order leg. From the previous chapter, we determined that these modulations are caused by the pairwise interference of the direct electrons with each of the forward scattered orbits. These direct electrons have ionization phases on the opposite half cycle as the forward scattered orbits, which implies that there are contributing trajectories in this region with very different ionization phases from each other. The fact that the observable ridges in the ϕ_{opt} line-out are fairly small in magnitude indicate that the amplitude of the direct electron trajectories are much smaller than the amplitude of the forward scattered trajectories, and so have a lesser contribution to ϕ_{opt} (see Section 6.4.1).

On the phase contrast spectrum in the bottom half of Figure 6.3, I have placed three red arrows pointing to specific trajectory transitions which can be identified in two-color PP analysis. From left to right these transitions are the first-return backscattered trajectories to second-return rescattered trajectories, second-return rescattered trajectories to direct trajectories, and first-return backscattered trajectories to first-return forward scattered trajectories. These transitions can be unambiguously identified because each of them involves a transition from a region with exactly one

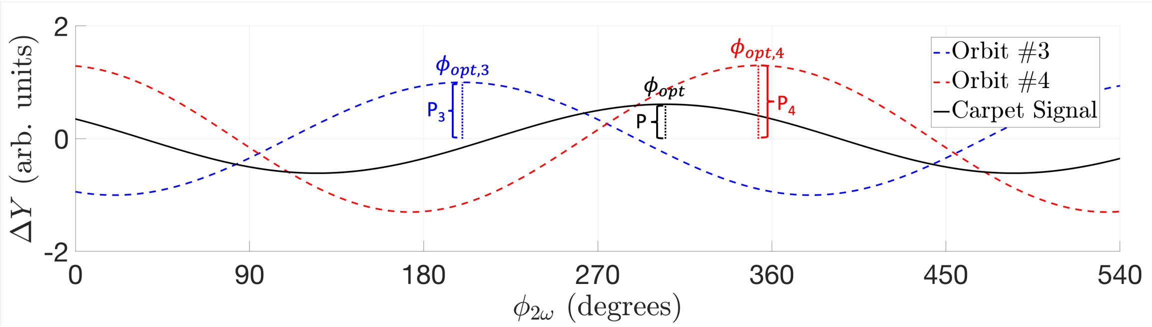

Figure 6.5: A simulated single pixel PP fit in the center of the Carpet hologram at $(p_{\parallel}, p_{\perp}) = (-0.1$ a.u., 0.8 a.u.$)$. Exactly two trajectories contribute at this pixel, orbit #3 (blue dashed) and orbit #4 (red dashed). The measurable PP fit is the sum of the contributions from each orbit, resulting in the black curve. The horizontal positions of each dotted line indicate the optimal phase of that PP fit, while the vertical length is the phase contrast, as annotated in the plot.

contributing trajectory. As I demonstrate in the upcoming section, this is a critical limitation, and indeed it is very difficult to use two-color PP spectra to measure properties of electron trajectories in regions with more than one trajectory arriving at a given pixel.

6.4.1 Limitations of Two-Color Phase of the Phase

Two-color PP analysis can reveal the PMD regions in which different types of electron rescattering occur; however, the only reason accurate ionization phases can be retrieved in this area is because only one type of trajectory contributes to the PP fit.

When two or more trajectories arrive at a single pixel, each of them generates a signal which is modeled by the PP fit. The measurable signal then is the sum of the contributions from each trajectory. The sum of the two sinusoidal functions of equal frequency but different phases and amplitudes is another sinusoidal function of the same frequency, but whose phase is determined by the ratio of the amplitudes of the two component sinusoids and whose amplitude is determined by the coherent sum of the amplitudes of the two component sinusoids. In general, this sum can never be decomposed into the individual components as it is indistinguishable from the PP fit of a single trajectory.

In Figure 6.5 I present a simulated example of the PP fit from a pixel in the center of a fringe in the Carpet hologram with momenta $(p_{\parallel}, p_{\perp}) = (-0.1$ a.u., 0.8 a.u.$)$. Using the CQSFA, the PP fits for the component trajectories, orbits #3 (blue dashed) and #4 (red dashed), are generated from their ionization phases and relative amplitudes. The calculated ionization phases are $\omega t_{0,3} = 190°$ for orbit #3 yielding $\phi_{opt,3} = 200°$, and $\omega t_{0,4} = 356.5°$ for orbit #4 yielding $\phi_{opt,4} = 353°$. The calculated relative amplitudes of the two trajectories provide their relative phase contrasts, yielding $P_3 = 1$ and $P_4 = 1.3$. The measurable PP fit (solid black) parameters for these values ends up being $\phi_{opt} = 305.02°$ and $P = 0.611$. Neither of these two measurable values yield any useful information regarding the fundamental parameters of either of the two contributing trajectories. The two ionization phase candidates for this value of ϕ_{opt} are $\phi_{C,1} = 332.5°$ and $\phi_{C,2} = 242.51°$, neither of which are particularly close to the ionization phases of the contributing trajectories.

Measuring the value of $\phi_{opt}(-0.1, 0.8)$ directly from Figure 6.3 yields $\phi_{opt}(-0.1, 0.8) = 299° \pm 11°$, which is quite similar to the simulated value. This similarity indicates that the $\phi_{opt}(p_{\|}, p_{\perp})$ spectrum in the upper half of Figure 6.3 is measuring summed PP fits as predicted when there are two or more trajectories arriving at a single pixel. Because these fits cannot be decomposed into the component trajectory's PP fits, all pixels in the PMD containing the contribution of two or more trajectories cannot be effectively analyzed with the PP analysis method. This effectively voids the value of phase-of-the-phase for analyzing any holograms, as by definition they are composed of pairwise interference of more than one trajectory.

6.5 The Third Harmonic Correction

So how do we fix this? The main problem causing PP analysis to fail for multiple trajectories is that there is no way to decompose a PP fit into the contributions from more than one trajectory. Thus we must design an experiment which can effectively distinguish between the contributions from multiple trajectories, and allow us to separately extract an ionization phase from each of them. With only one harmonic perturbation, our single pixel analysis is trapped in one dimension which doesn't have enough free parameters to resolve multiple trajectories. With two harmonic perturbations, however, each pixel may be represented by a two-dimensional manifold, which can resolve any number of trajectories, as long as there is sufficient resolution and phase stability in the collected data.

Introducing the third harmonic perturbation to our combined laser field results in the following electric field:

$$E(t, \phi_{2\omega}, \phi_{3\omega}) = E_{\omega} \cos(\omega t) + E_{2\omega} \cos(2\omega t - \phi_{2\omega}) + E_{3\omega} \cos(3\omega t - \phi_{3\omega}), \tag{6.7}$$

where $\phi_{2\omega}$ is the relative phase between the peak of the fundamental and the second harmonic as defined previously, and $\phi_{3\omega}$ is the equivalently defined relative phase between the fundamental and the third harmonic.

Given this new field we are free to simulate the effect varying the relative phases of both $\phi_{2\omega}$ and $\phi_{3\omega}$ would have on a single pixel in the final PMD. For parity with the previous discussion, in Figure 6.6 I present a three-color ionization phase extraction procedure to disentangle contributions from multiple trajectories acting on a simulation of the same pixel presented in Figure 6.5 at the momentum coordinates $(p_{\|}, p_{\perp}) = (-0.1 \text{ a.u.}, 0.8 \text{ a.u.})$. Once again, there are the contributions of two trajectories, #3 and #4, which have ionization phases $\omega t_{0,3} = 190°$ and $\omega t_{0,4} = 356.5°$. This phase extraction procedure is quite involved and requires significant explanation.

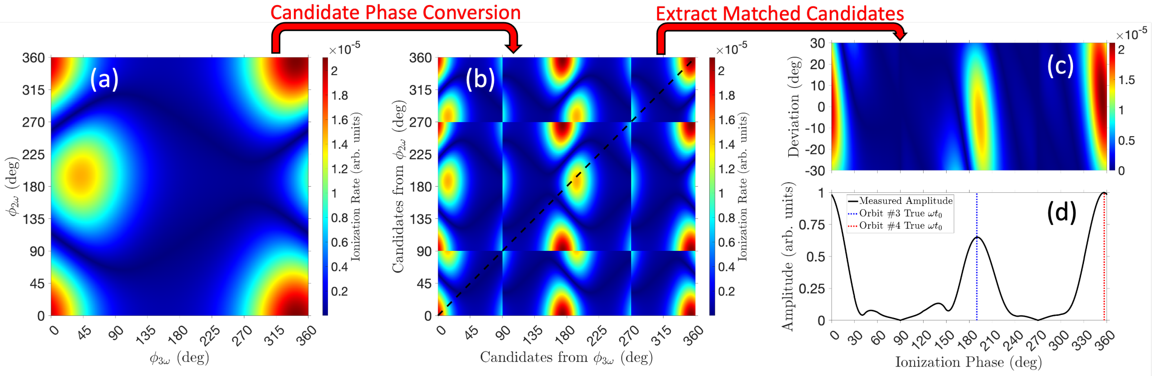

Figure 6.6: A sample three-color ionization phase extraction for a simulated pixel in the Carpet hologram at $(p_{\|}, p_{\perp}) = (-0.1 \text{ a.u.}, 0.8 \text{ a.u.})$ with exactly two contributing trajectories. (a) The simulated amplitude of the pixel plotted as a manifold against the relative phases $\phi_{2\omega}$ and $\phi_{3\omega}$. (b) The two relative phases of the single pixel amplitude are translated into their ionization phase candidates to yield the candidate phase distribution. A dashed black line drawn along the diagonal indicates the portion of the candidate phase distribution where the candidate ionization phases from each of the two relative phases match. (c) The portion of the candidate distribution along the dashed black line is plotted, along with up to 30° of deviation above and below the line. Around the discontinuities at 90° and 270° a windowing function (see Equation 6.13) is applied to suppress the amplitude. (d) The weighted sum across the phase deviation in Panel (c) is plotted as the solid black line. The initial ionization phases used to generate the single pixel amplitude in Panel (a) are shown as the blue dotted line for Orbit #3 and the red dotted line for Orbit #4.

In Panel 6.6(a) the amplitude of a single pixel is simulated using the ADK tunnel ionization rates:

$$Y(p_{\parallel}, p_{\perp}, \phi_{2\omega}, \phi_{3\omega}) = \Bigg| P_3 W_0(t_{0,3}, \phi_{2\omega}, \phi_{3\omega}) + P_4 W_0(t_{0,4}, \phi_{2\omega}, \phi_{3\omega}) \exp(i\Delta S) \Bigg|. \tag{6.8}$$

Here, the tunnel ionization rates W_0 are calculated at each value of the relative phases $\phi_{2\omega}$ and $\phi_{3\omega}$ at the ionization phase of each of the two included trajectories, and $\Delta S = \pi$ is the value of the phase difference between the two trajectories, which is calculated from the CQSFA in the absence of the two perturbations. Note that in this formulation I have omitted the contributions of the transverse momentum distributions which vary only slightly with the perturbations. As before, the scaling factors $P_3 = 1$ and $P_4 = 1.3$ are included to match the ratio of the amplitudes of the two trajectories in the CQSFA. The single pixel amplitude is calculated across the full range of relative phases, from 0 to 360°, and as such the full manifold is cyclic across the left and right boundary and also separately across the upper and lower boundary.

Immediately we see that there are two distinct lobes in the single pixel amplitude, corresponding to each of the two trajectories. In the summed two-dimensional manifold, these lobes resolve separately and do not generate a confounding signal unrelated to either trajectory. Reminiscent of performing a PP fit, we may examine the optimal phases of each of these lobes in both relative phase directions, extracting a pair of values for each. As an example, the lobe appearing on the left half of the amplitude spectrum has an apparent set of optimal phases $(\phi_{2\omega,opt},\ \phi_{3\omega,opt}) = (200°,\ 30°)$. From Equations 6.5 and 6.6 we can determine the ionization phase candidates for the second harmonic perturbation, which come out to be $\phi_{2\omega,C,1} = 280°$ and $\phi_{2\omega,C,2} = 190°$. In our simulation, we already know the correct value is 190°, but it would be useful to determine this algorithmically, and so we must also incorporate the candidate ionization phases for the third harmonic perturbation.

The ionization phase candidates calculated from the third harmonic perturbation are different than those from the second harmonic perturbation, and must be defined over specific ranges of $\phi_{3\omega,opt}$:

$$\phi_{3\omega,C,1} = \frac{\phi_{3\omega,opt}}{3} \quad : \quad \phi_{3\omega,opt} \in [0, 270°], \tag{6.9}$$

$$\phi_{3\omega,C,2} = \frac{\phi_{3\omega,opt} - 90°}{3} + 90° \quad : \quad \phi_{3\omega,opt} \in [90°, 360°], \tag{6.10}$$

$$\phi_{3\omega,C,3} = \frac{\phi_{3\omega,opt}}{3} + 180° \quad : \quad \phi_{3\omega,opt} \in [0, 270°], \tag{6.11}$$

$$\phi_{3\omega,C,4} = \frac{\phi_{3\omega,opt} - 90°}{3} + 270° \quad : \quad \phi_{3\omega,opt} \in [90°, 360°]. \tag{6.12}$$

Here, each of the ionization phase candidates occupy a unique 90° region of phase space: $\phi_{3\omega,C,1} \in [0, 90°]$, $\phi_{3\omega,C,2} \in [90°, 180°]$, $\phi_{3\omega,C,3} \in [180°, 270°]$, and $\phi_{3\omega,C,4} \in [270°, 360°]$.

From these equations, we can calculate the two possible ionization phase candidates for the selected lobe as $\phi_{3\omega,C,1} = 10°$ and $\phi_{3\omega,C,3} = 190°$. As we have a match between a candidate phase from the second harmonic perturbation and a candidate phase from the third harmonic perturbation, we may conclude that the true ionization phase of the trajectory forming the lobe is 190° as expected.

The previous procedure was done manually based on an observed local maximum on the single pixel amplitude manifold. This would not work for an experiment with noisy data and lower phase resolution. Instead, let us define the "Candidate Phase Distribution" (CPD) as the amplitude of all ionization phase candidates calculated from the amplitude of every ordered pair of relative phases $(\phi_{2\omega},\ \phi_{3\omega})$. This translates the amplitude at $(\phi_{2\omega},\ \phi_{3\omega})$ into either four or six points in the CPD depending on the value of $\phi_{3\omega}$. The CPD calculated from the simulated single pixel amplitude is shown in Panel 6.6(b). The two axes of this distribution are both in units of degrees on the fundamental ionizing laser pulse. Thus, the only points in this distribution which are physically realizable phases lie along the line $\phi_{2\omega,C} = \phi_{3\omega,C}$ along the central diagonal.

There are a few reasons why these matched candidates might not perfectly align with the proper ionization phases of contributing trajectories in an experiment. First, the initial set of calibrations which generate the values of $\phi_{2\omega}$ and $\phi_{3\omega}$ may have errors which propagate into the CPD. Second, the correction factor calculated in Equation 3.10 may be imperfect, leading to the value of $\phi_{3\omega}$ being improperly coupled to the value of $\phi_{2\omega}$, causing systemic aberrations. Third, statistical variations in the collected electrons may coincidentally result in lower signal along the matched candidates compared to slightly off the main diagonal.

In order to somewhat account for these sources of error, a phase deviation of 30° is defined, and the total amplitude along the main diagonal as well as 30° above and below that line is collected. I'll refer to the amplitude of this region along the diagonal as the "Matched Candidates Strip" (MCS). This is shown in Panel 6.6(c). Additionally, in this step, the amplitudes are windowed by the function

$$\begin{aligned}\mathrm{CPD}(\phi_{2\omega,C},\phi_{3\omega,C}) \times \big(&H(|\cos(\phi_{2\omega,C})| - 0.5) + H(0.5 - |\cos(\phi_{2\omega,C})|)|\cos(\phi_{2\omega,C})|\big)\\ &\times\big(H(|\cos(\phi_{3\omega,C})| - 0.5) + H(0.5 - |\cos(\phi_{3\omega,C})|)|\cos(\phi_{3\omega,C})|\big),\end{aligned} \tag{6.13}$$

where $\mathrm{CPD}(\phi_{2\omega,C},\phi_{3\omega,C})$ is the value of the candidate phase distribution at the ordered pair of phase candidates $(\phi_{2\omega,C},\phi_{3\omega,C})$ and $H(x)$ is the Heaviside step function. This complicated looking window simply leaves the amplitude of points in the CPD the same if they're farther than 30° from either 90° or 270° in either phase direction, and causes them to decay according to $\cos(\theta)$ as they approach 90° or 270°. This is important to reduce the effect of nonphysical discontinuities at these phase transitions which interfere with accurate retrieval of ionization phases from the matched candidates. Note that at this point the MCS is no longer a manifold of candidate ionization phases along either axis, but instead is plotted against true ionization phase and phase deviation.

Once the MCS is collected, we can perform a weighted sum along the phase deviation to extract a one-dimensional amplitude. This weighted sum is given by

$$\Omega(\phi_0) = \sum_{\Delta\phi} \mathrm{MCS}(\phi_0, \Delta\phi) \cos(3\Delta\phi), \tag{6.14}$$

where ϕ_0 are the range of possible ionization phases, $\Omega(\phi_0)$ is the weighted sum, and $\Delta\phi$ is the range of phase deviations in the MCS. This weighted sum assigns the most weight to points exactly along the diagonal, and decreases the weight of points the farther they are from the matched candidate diagonal.

The result of this weighted sum in the simulation is shown in Panel 6.6(d). The maxima in this weighted sum correspond to the predicted values of the ionization phases of the contributing trajectories forming the single pixel amplitude in Panel 6.6(a). On this plot I have overlaid the true values of the ionization phases used in the simulation. We see that the maxima of $\Omega(\phi_0)$ exactly match the correct ionization phases, and thus we may use this technique to successfully extract the ionization phases from pixels with multiple contributing trajectories. Additionally, the relative heights of the two peaks are closely related to the ratio of the amplitudes, P_3 and P_4, of the two trajectories that were used to generate the single pixel distribution according to Equation 6.8.

> Given the form of Equation 6.4 and its importance to two-color PP analysis, you may be wondering why there is no similar fitting process outlined in the procedure for ionization phase extraction with three colors. Originally when I approached this problem, I did in fact design a fitting algorithm which looked for single pixel signal taking the form
>
> $$\begin{aligned} Y(p_\parallel, p_\perp, \phi_{2\omega}, \phi_{3\omega}) = \sum_i \Big(& P_{2\omega,i}(p_\parallel, p_\perp) \cos\left(\phi_{2\omega} - \phi_{2\omega,i,opt}(p_\parallel, p_\perp)\right) \\ & + P_{3\omega,i}(p_\parallel, p_\perp) \cos\left(\phi_{3\omega} - \phi_{3\omega,i,opt}(p_\parallel, p_\perp)\right) \Big) + C(p_\parallel, p_\perp). \end{aligned} \tag{6.15}$$
>
> Here, the sum is over an indeterminate number of trajectories indexed by i. In principle, such a fit does exist for each pixel, and if it could be accurately determined it would most likely yield correct values for the optimal phases for each contributing trajectory, which could then be translated into the proper ionization phases for each trajectory. However, in practice I found this to be an insurmountable challenge. Because there's no *a posteriori* method to know the number of trajectories arriving in a given pixel, there's no way that I discovered to set the fitting parameters up correctly for each pixel without manual intervention. Additionally, even when the fit is set up correctly, such as when using two trajectories for the simulated pixel we examined throughout this section, the number of free parameters ends up

overwhelming common fitting algorithms and yields a poor fit.

In contrast, the procedure outlined in the rest of this section is computationally inexpensive and relatively robust. Since the relative phases of two colors can uniquely determine a set of matched ionization phase candidates, we are able to bypass the need to do complex manifold geometry fitting and work with approximately one-dimensional results that sidestep the ambiguities plaguing two-color PP analysis.

6.6 Three-Color Data

The three-color experimental data analyzed in the rest of this chapter is described here. Argon gas was photoionized by three overlapped collinear laser pulses. As before, the fundamental 800 nm ionizing pulse was focused to an approximate intensity of 90 ± 10 TW/cm^2, corresponding to $\gamma \approx 1.2$; the second harmonic 400 nm pulse was focused to an approximate intensity of 0.45 TW/cm^2; and the third harmonic 266 nm pulse was focused to an approximate intensity of 0.9 TW/cm^2. In following the notation used in Equation 6.7, $E_\omega = 5.1 \times 10^{-2}$ a.u., $E_{2\omega} = 3.6 \times 10^{-3}$ a.u., and $E_{3\omega} = 5.1 \times 10^{-3}$ a.u. The intensity of the fundamental was determined with Qprop as previously described in Figure 5.4, while the intensities of the other two pulses were determined from the measured pulse powers as in the two-color case. The two relative phases are controlled via glass wedges mounted on two linear translation stages, as described in Section 3.4.5. A PMD was collected at 17 evenly spaced positions of the second harmonic glass wedges, and at positions of the glass wedge controlling the third harmonic such that 17 evenly spaced values of relative phase were collected. Each of the sets of relative phase were calibrated to span approximately one full cycle of second or third harmonic relative phase. Approximately 9.7 billion electron counts were recorded over the span of approximately 15 hours, resulting in a distribution of about 33.6 million electron counts per PMD. All other parameters are the same as described in Section 4.1.

When collecting a strong-field dataset which includes laser light around 266 nm, a significant number of spurious electrons are generated that are unrelated to strong-field ionization processes. This is because the work function of stainless steel is around 4.4 eV and the photon energy from 266 nm photons is approximately 4.66 eV. As a result, scattered photons which impact the stainless steel plates liberate electrons due to the photoelectric effect, which are accelerated by the VMI plates and impact the detector. This spurious signal causes streaks across the detector that interfere with Abel inversion techniques. In the three-color dataset examined in this chapter, the laser field was carefully focused so that the streaks were predominantly located at high transverse momenta. While these streaks are independent of the relative phase of the third harmonic and as such should not significantly disrupt the three-color phase extraction procedure previously outlined, they do introduce statistical error. As a result, the experimental spectrum has been truncated at approximately $p_\perp = \pm 1$ a.u., which is just outside the Carpet hologram.

As in the two-color case, the relative phase of the third harmonic is calibrated by examining the yield of electrons in the backscattered regime just shy of $10U_p$ on the left half of the raw PMD. Summing over the $\phi_{2\omega}$ coordinate provides enhanced statistical resolution. The phases $\phi_{3\omega}$ are then set by finding the optimal phase in this region, and setting $\phi_{3\omega,C,1} = 15^\circ$ for these electrons.

6.7 Three-Color Ionization Phase Extraction

In actual experimental data, the weighted sum $\Omega(\phi_0)$ is not nearly as smooth as the curve shown in Panel 6.6(d). In order to successfully extract ionization phases, I need to add one more layer of processing. For each point on the $\Omega(\phi_0)$ curve, whether or not there is a peak at that point, I calculate the centroid (or center of mass) of a proposed peak according to

$$\phi_{cm} = \frac{\sum_n \phi_{0,n}\Omega(\phi_{0,n})}{\sum_n \Omega(\phi_{0,n})}, \tag{6.16}$$

where n indexes the range of points over which to calculate the centroid. For each centroid, the moving average is calculated as

$$\Omega_{cm}(\phi_{cm}) = \frac{1}{N}\sum_n \Omega(\phi_{0,n}), \tag{6.17}$$

where N is the total number of points in the defined range. For the analysis shown here, $N = 5$, and the window is centered on the examined point. Points at the end of the range are treated cyclically for the purpose of centroiding. This moving average is then normalized to have a maximum of 1, in order to facilitate direct comparisons between the relative amplitudes of multiple trajectories. The centroids corresponding to the two peaks of greatest amplitude are determined, and recorded as the extracted ionization phases of the two most probable contributing trajectories. In principle, with high enough resolution more than two peaks could be calculated, but the dataset I collected did not have sufficient phase stability between the laser fields to adequately resolve a third trajectory.

The result of applying this phase extraction to the single pixel located at $(p_\parallel, p_\perp) = (-0.1$ a.u., 0.8 a.u.$)$ is shown in Figure 6.7. The signal extracted from this pixel is the direct analog to the simulation shown in Figure 6.6. From the single pixel distribution Panel 6.7(a), we can roughly make out the two distinct lobes corresponding to the two contributing trajectories, which were very prominent in the much higher resolution simulation. Proceeding through the full phase extraction procedure, we arrive at the centroided moving average in Panel 6.7(b). Also shown in dashed cyan is the noisy $\Omega(\phi_0)$. The centroids corresponding to the top two peaks are indicated respectively by

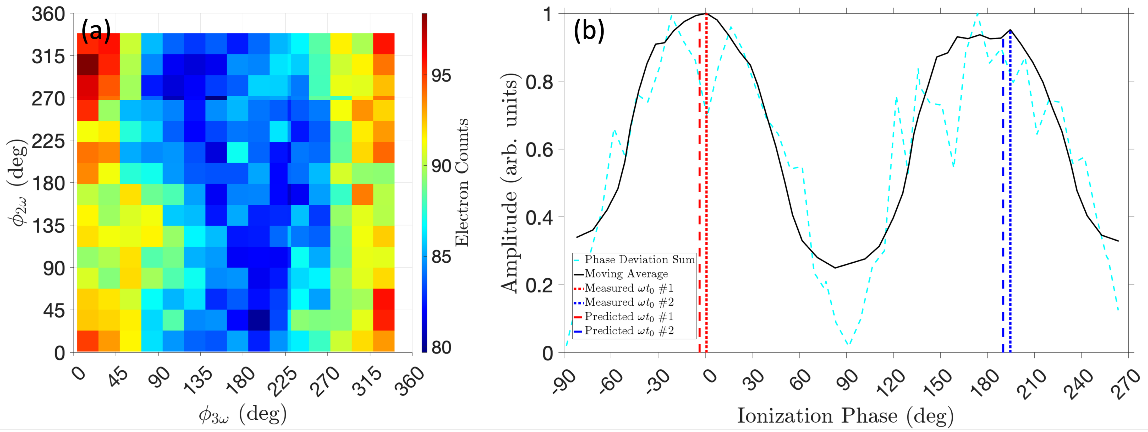

Figure 6.7: The result of three-color phase extraction on the pixel located at $(p_{\parallel}, p_{\perp}) = (-0.1$ a.u., 0.8 a.u.$)$. (a) This panel shows the raw electron counts of this pixel expressed against the relative phases $\phi_{3\omega}$ and $\phi_{2\omega}$. (b) This panel plots the centroided moving average (solid black, see text) of $\Omega(\phi_0)$ extracted from the single pixel in (a). $\Omega(\phi_0)$ is also plotted (dashed cyan) for reference. The two highest peaks corresponding to the measured ionization phases are indicated by dotted red and blue lines. The ionization phases predicted from CQSFA are plotted as dashed red and blue lines for orbits #4 and #3 respectively.

a red dotted line and a blue dotted line. These are the measured ionization phases we are able to extract from the single pixel. The blue and red dashed lines are the predicted ionization phases corresponding to CQSFA orbit #3 and #4 respectively, and are replicated from the analysis in Figure 6.6. For clarity, on this panel the ionization phases range from $-90°$ to $270°$ as the predicted and measured ionization phase corresponding to orbits #4 are just barely on opposite sides of 0. The measured ionization phases come out to be $0.8°$ and $194.6°$ compared to the predicted values of $356.5°$ and $190°$. This is a deviation of about $5°$ of phase for the ionization phases of both of the contributing trajectories predicted by CQSFA to arrive at this pixel. Expressed in units of time, this represents an agreement with theory to within 37 attoseconds within a given laser cycle.

It is difficult to quantify the error on these measurements and the technique overall as there are several possible sources, but I lay out my reasoning here. Firstly, each pixel has statistical noise spread over its relative phase positions resulting from imperfect distributions of electrons at every pixel and PMD. Since we are less interested in the absolute values of the electron counts at each phase position and rather are focused on the positions of the centroided local maxima across the candidate phases, this source of error should be relatively small. Secondly, an imperfect calibration translating the glass wedge insertions into relative phases of the two harmonics would result in a systematic non-linear ionization phase offset resulting from the mapping of each pixel into its CPD. The inclusion of the integration across the MCS step in the technique should help account for this. Thirdly, an imprecise calibration between the two sets of wedges during the experiment resulting in the position of the 2ω wedges being coupled to the phase of third harmonic would disrupt the

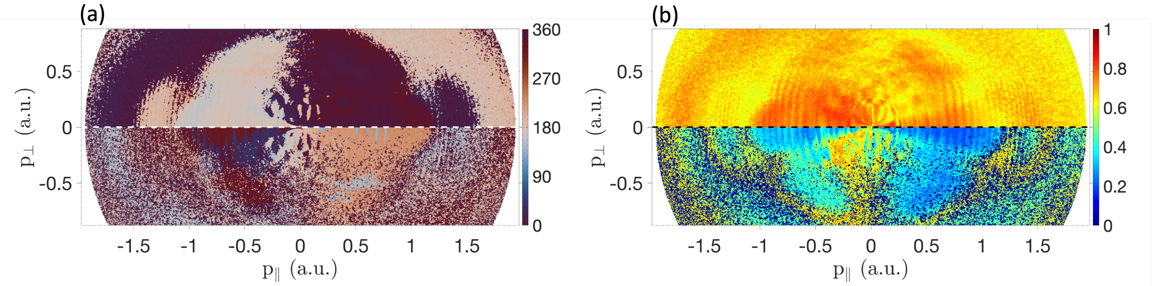

Figure 6.8: The result of the full three-color ionization phase extraction. (a) This panel presents the measured ionization phases of the two contributing trajectories with the highest amplitudes, with the top half showing the most probable trajectory and the bottom half showing the second-most probable trajectory. (b) This panel presents the prominences of the peaks of the $\Omega_{cm}(\phi_{0,cm})$ corresponding to the ionization phases in panel (a).

accurate generation of the CPD. This has the potential to be the largest contribution to the error. It is possible that slight mismatches could be corrected after-the-fact in the analysis by observing the response of the backscattering region on both halves of the spectrum. These two regions should be symmetric with respect to 3ω and anti-symmetric with respect to 2ω, so a correction function could be applied to enforce this relation. With these corrections in place, the largest source of error would then be the relative-phase resolution of the experiment itself. The experiment shown here was collected at 17 positions of each of the two sets of wedges. When converted to their candidate ionization phases and integrated along the MCS to yield $\Omega(\phi_0)$, this translates into an ionization phase axis with approximately 8° of separation between neighboring values. This is an appropriate value to assign to the error of the ionization phase of the primary extracted trajectory. Additional trajectories are a bit more complicated to assess and require a bit of additional reasoning, which I outline roughly below.

The full three-color phase extraction for the entire PMD is shown in Figure 6.8. Panel 6.8(a) shows the measured ionization phases of the highest amplitude extracted trajectory (top) and the second highest amplitude extracted trajectory (bottom). Panel 6.8(b) presents the prominences of the corresponding peak in the moving average $\Omega_{cm}(\phi_{0,cm})$. The prominence of a peak is defined as the vertical distance between the maximum value of the peak and the lowest point in the interval between the peak and either the endpoints or a signal with higher amplitude than the peak. The relative prominence of two trajectories is a useful metric which describes both the relative amplitude of the trajectories found in the phase extraction analysis and also the confidence in the measured ionization phase. In order to distinguish whether a low secondary prominence is due to the relative amplitudes or due to the inaccuracy of the measurement, some prior knowledge from calculation may be required.

Both the measured ionization phases in Panel 6.8(a) and the prominences in Panel 6.8(b) should be used in concert to reach accurate conclusions about the trajectories they describe. For instance, consider the region of first-return backscattered electrons on the left half of the PMD. We already know that there is only one type of trajectory arriving in this region. The measured ionization phases for the second most probable trajectory appear to suggest that there is a second trajectory with an ionization phase around 135° in this region; however, from the prominence map we see that this peak has particularly low prominence, indicating that it is likely spurious. Since the phase extraction procedure always generates at least two ionization phases for each pixel, we must rely on the prominence map to discard the second trajectory if none should exist.

In contrast, consider the region around $7U_p$ corresponding to second-return rescattering electrons. The ionization phase of the primary trajectory accurately reports the ionization phase of these trajectories; however, the measured ionization phase for the secondary trajectory indicates an ionization phase 180° apart. Looking at the prominence of this secondary trajectory, we observe a clear boundary around $7U_p$ where the prominence is higher than in the inner area below $7U_p$. This indicates a transition from the area where first-return rescattered electrons dominate into the area where second-return rescattered electrons dominate. Inside the area dominated by second-return rescattered electrons, the low prominence is an indication of the relatively higher amplitude of the second-return electrons compared to the first-return in this lower momentum regime.

One area where the secondary prominence is relatively high is around the Carpet hologram. This makes sense because the single pixel signal forms maximally distinct lobes for trajectories separated by 180° of phase, as is the case for the Carpet. It's also the main reason why the measured ionization phases for pixels in the Carpet region have such a stellar agreement with calculation. Based on observation, while the ionization phase of the most probable electron trajectory matches closely with expectations set by the CQSFA, the secondary ionization phase can acquire more error.

One of the most noticeable features in the prominence map is the very low prominence of the spider-leg structures. These low prominences are an expected result for this hologram based on the discussions we've had about the spider legs thus far. The Spider is formed from the interference of two trajectories whose ionization phases are nearly identical, and thus there is no way to distinguish the two via this phase extraction procedure. However, the phase extraction does appear to pick out the ionization phases of the direct trajectories on the opposite half-cycle in this region as the second-most probable trajectory. We know from the discussion around the time-correlation filter that the direct trajectories are the additional trajectories which give rise to the spider-leg modulations we have previously observed. The particularly low prominence we can measure here indicates that these trajectories have a much lower amplitude than the two forward scattering orbits which form the bulk of the spider-leg structures. While this trajectory is certainly not the second-most probable trajectory in this region, it is interesting that the technique is able to resolve and describe this low-amplitude trajectory.

6.8 Conclusion

Overall, the three-color ionization phase extraction technique is a powerful tool for accurately measuring the ionization phases of the trajectories contributing to a hologram in strong-field spectra. It resolves the major shortcoming of two-color phase-of-the-phase type approaches and yields quantitative results that have significant agreement with CQSFA calculations. Except for the ATI rings which have no specific ionization phases to extract, each of the major strong-field features we have examined in this thesis can be effectively described by this technique, demonstrating its value and versatility.

The three-color ionization phase extraction technique suffers from one major shortcoming: it is difficult to perform the experiment in a timely manner and with sufficient phase-stability and resolution. Overlapping the three laser pulses in the interaction region and accurately determining the relative phase calibration required to calibrate the correction factor needed to decouple the motion of the two stages is quite difficult. If this can be achieved routinely, however, then the use of two relative-phase-controlled harmonic perturbations of the ionizing laser field is a powerful probe of ultrafast electron dynamics.

Ultimately, the goal of this technique is to develop a new kind of probe pulse in studies of strong-field processes of atoms and molecules. The three-color field employed here is the minimally complex laser system which provides sufficient degrees of freedom to separately resolve the contributions of multiple trajectories in each hologram. Incorporating a three-color probe pulse in future strong-field studies could lead to new and exciting ways to image the electronic structures of atoms and molecules undergoing complex structural dynamics by revealing how electron trajectories change in response to those dynamics.

Chapter 7

Speculations on Future Strong-Field Discovery

The previous chapters demonstrated techniques to disentangle electron trajectories in strong-field spectra to make quantitative measurements of ultrafast electron dynamics. These techniques are possible due to the wealth of information encoded in strong-field PMDs. In various forms, strong-field ionization has been studied for the better part of the past fifty years, and new discoveries are being made each year. New types of strong-field experiments and innovative analysis techniques are continuously revealing intriguing features which appear in the data and indicate entirely new avenues for exploration. In this chapter I briefly touch on two of the thought-provoking features I have encountered in my studies which do not have sufficient explanation. I do not attempt to provide rigorous analyses of these phenomena, but I do speculate on their sources and discuss possible directions for further study.

7.1 Tunnel Barrier Exits in Position-Correlation Distributions

In my development of the time-correlation filter, I had many opportunities to study the results of Fourier transforming strong-field spectra. Eventually this work culminated in the discovery that the Fourier transforms of the anisotropy parameters in energy space yielded an angularly-resolved time-correlation distribution of contributing electron trajectories. Alongside that discovery was the realization that directly Fourier transforming the PMD itself yields a "position-correlation distribution" (PCD), which relates on two axes the position correlation along the direction of laser polarization and the position correlation transverse to the laser polarization. This immediately begs the question, "a position correlation of *what*?" While I still do not have a clear answer to this question, constructing the PCDs from one experiment has provided a clue. This experiment examines strong-field PMDs of argon gas at a broad range of laser intensities, which causes the tunnel barrier exit (TBE) of electron trajectories at the peak of the laser field to shift accordingly. Directly examining the PCDs of this dataset reveals a feature which appears to align with the predictions of the adiabatic TBE calculated by Equation 2.11.

7.1.1 Intensity Sweep Data

The intensity sweep experiment followed exactly the parameters and processing techniques outlined in Chapter 4, except at an evenly spaced range of different 800 nm pulse powers. The pulse powers were controlled by programmatically rotating the half-wave plate in the pulse-power throttle described in Section 3.4.1. After careful calibration with Qprop, it was determined that the intensities range from 56.5 TW/cm^2 to 131.5 TW/cm^2 at even steps across 51 PMDs. This corresponds to $\gamma \approx 1.527$ to $\gamma \approx 1.001$. Because of the differing laser intensities, the number of collected electrons differs significantly between the PMDs, ranging from 1.26 million electrons for the lowest intensity PMD to 197.4 million electrons for the highest intensity PMD.

7.1.2 Constructing the Position-Correlation Distribution

We begin from the fully inverted stack of PMDs, $Y(p_{\|}, p_{\perp}, I)$, where I indexes the intensity of the laser field. The PCD is then calculated from the 2-dimensional discrete Fourier transform acting on the momentum bins of each PMD:

$$\tilde{Y}[u,v] = \frac{1}{N^2} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} Y[j,k] \exp\left(-2\pi i \Big(\frac{uj+vk}{N}\Big)\right) \tag{7.1}$$

where N is the number of pixels along one dimension, j and k index the momentum bins $p_{\|}$ and $p_{\perp}$ respectively, u and v index the position correlation bins $z_{\|}$ and $z_{\perp}$, $\tilde{Y}(z_{\|}, z_{\perp})$ is the PCD, and the square brackets surrounding the arguments of Y and $\tilde{Y}$ indicate to use the momentum or position correlation bin referenced by the indices. Because the PMDs are square matrices by construction, we may use N for each dimension. Using this transform, the step size of the resulting range of position correlations is given by

$$\delta z = \frac{2\pi}{\delta p N}, \tag{7.2}$$

where δp is the width of each momentum bin. Then the position correlations cover the range

$$(-z_{\max}, z_{\max}) = \left(\frac{-N\delta z}{2}, \frac{(N-1)\delta z}{2}\right), \tag{7.3}$$

where $z_{\max}$ is the maximum extent of the position correlations in either dimension. For the data presented here, $\delta p = 0.0049$ a.u. and $N = 921$, which means the resolution of the PCD is $\delta z = 1.39$

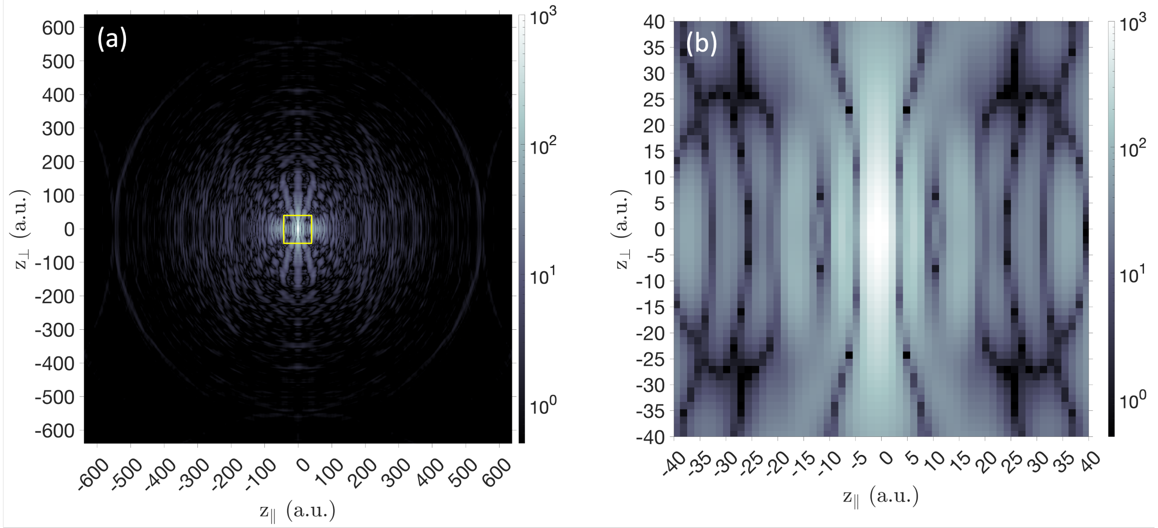

Figure 7.1: The PCD generated from the two-dimensional discrete Fourier transform of the lowest intensity PMD in the dataset: $Y(p_{\parallel}, p_{\perp}, 56.5\ \text{TW/cm}^2)$. (a) This panel plots the full PCD using the maximum extent of the position correlations in each dimension. (b) This panel shows just the zoomed-in portion of the PCD bounded by the yellow square in panel (a).

a.u. and the extent of the PCD is $z_{\text{max}} = 638.1$ a.u.

In Figure 7.1 the full PCD is plotted at the lowest intensity: 56.5 TW/cm^2. The full range of position correlations, shown in Panel 7.1(a), contains mostly features with little to no structure, which is not particularly interesting. In Panel 7.1(b) I zoom into the region bounded by the yellow square in Panel 7.1(a). In this region, there are many structural features, but there is no obvious origin for any of them. When we scan over the intensities, however, something interesting occurs. The vertical feature located at approximately $z_{\parallel} = 30$ a.u. begins to shift to the left. This is an intriguing result, as it is at a similar position to twice the adiabatic tunnel barrier exit distance calculated from Equation 2.11.

7.1.3 Revealing the Tunnel Barrier Exit Distances

To probe this, I integrate the PCD for each intensity along the $z_{\perp}$ axis from -15 a.u. to 15 a.u. before normalizing by the mean value of the curve. Plotting these sums for each intensity on a single two-dimensional grid yields Figure 7.2. We immediately see that there are at least two features which drift to the left, one originating at approximately $z_{\parallel} = 15$ a.u. and one originating at approximately $z_{\parallel} = 30$ a.u. The most prominent of these features is the one on the right at $z_{\parallel} = 30$ a.u. On top of these figures I have plotted two quantities. The red points are the adiabatic TBEs given by $z_{\text{TBE}} = \frac{I_p}{E_0}$. The blue points are double this quantity: $2z_{\text{TBE}} = \frac{2I_p}{E_0}$. These curves appear to match the observed features well, especially at higher intensities where the adiabatic TBEs are predicted

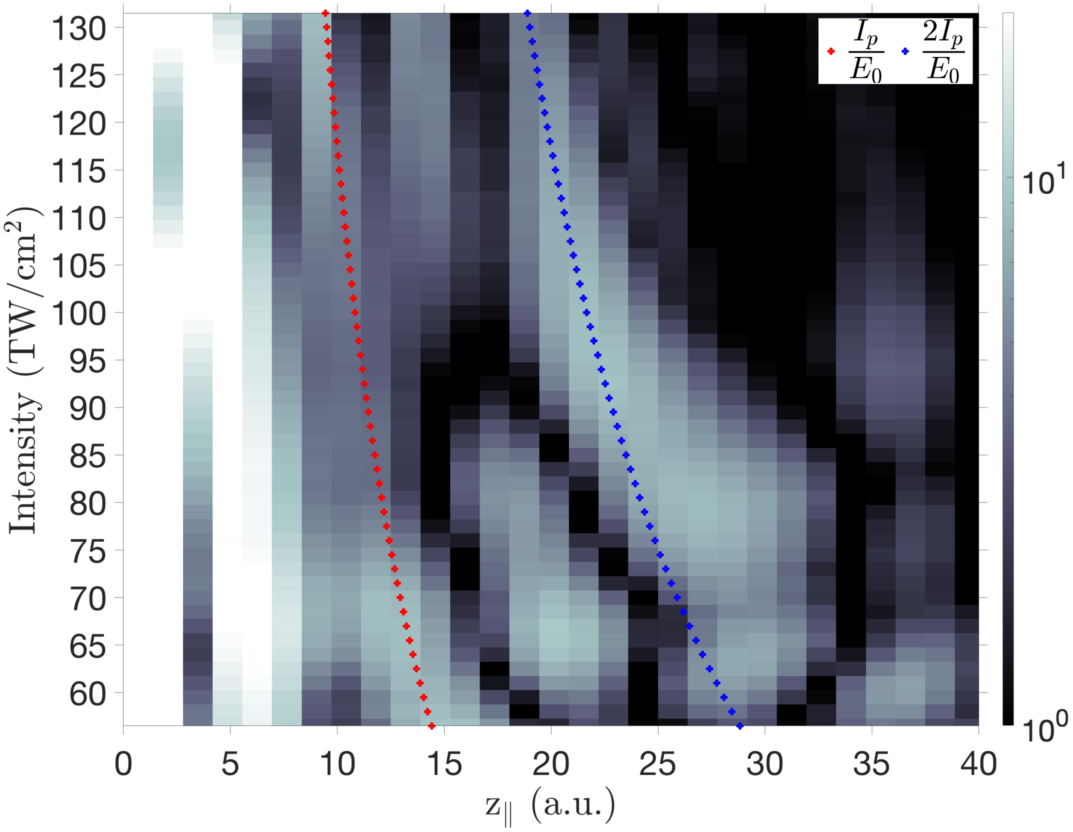

Figure 7.2: The region of the PCD shown in Panel 7.1(b) is summed along $z_\perp$ from -15 a.u. to 15 a.u. and normalized for each intensity. Overlaid on top of this plot are the TBEs calculated from Equation 2.11 for each intensity (red) and double this quantity (blue).

to be more accurate [39].

The physical significance of z_{TBE} is reasonably clear. If these features indicate position correlations within the electron wave packets which form the PMD, then it would be reasonable to suppose that the position of the portion of the electron wave packet at the moment of tunnel ionization would be correlated in some way to the residual electron wave packet remaining bound to the ion. This could manifest in a signal such as the one we observe highlighted by the red curve in Figure 7.2. Why then is the larger of the two signals occurring at $2z_{\text{TBE}}$? In many ways this quantity makes much more sense than z_{TBE}, as it could relate to the distance between the TBEs on each side of the ion. This connects cleanly to the electron trajectory model applied throughout this thesis. Electron trajectories ionized on opposite sides of the ion result in the electron wave packet self-interfering and forming some of the patterns we observe in the PMDs. The only position with any significance to these trajectories is their TBE, and so a major position correlation between these two trajectories must occur at the distance between their TBEs.

Ultimately, this explanation is conjecture at this point, but the apparent agreement between the adiabatic TBEs and the observed position-correlation features invites further study. A reasonable way to approach these PCDs is to perform experiments on samples which have more complicated spatial geometries than atomic argon can have. One example of an experiment which may simultaneously probe the properties of the PCD and examine these TBE features would be to use a pump-probe laser system to impulsively align a diatomic gas, and interrogate the system at a sufficient number of delay steps so as to capture the full alignment revival. As a diatomic molecule aligns with the laser polarization, its rate of ionization increases, suggesting there may also be a shift in the tunnel barrier as well. Additionally, the spatial geometry of the molecules would result in the TBE being farther from the center-of-mass in the case where the laser polarization aligns with the molecular axis compared to the antialigned case, resulting in larger distances between the TBEs on each side of the molecule.

7.2 Second-Return Rescattering Resonances

We've had several opportunities in this thesis to discuss how second-return rescattering electrons can be characterized in strong-field PMDs. In brief, second-return rescattering electrons attain momentum along the parallel axis just higher than the extent of forward scattered electrons, and have a classical cutoff of 7U_p. In Chapter 4, when we examined the SMM trajectories corresponding to the maximum extent of second-return rescattering electrons, we noted that they fell on a ring centered approximately at A_0 and extending out to 7U_p along the parallel axis. We saw that for the data presented in that chapter, this ring closely outlined a prominent circular feature in the experimental data. This feature is commonly referred to as a classical rescattering cutoff or caustic, and is generally assumed to always appear along the rescattering ring extending to 7U_p [35, 73, 74].

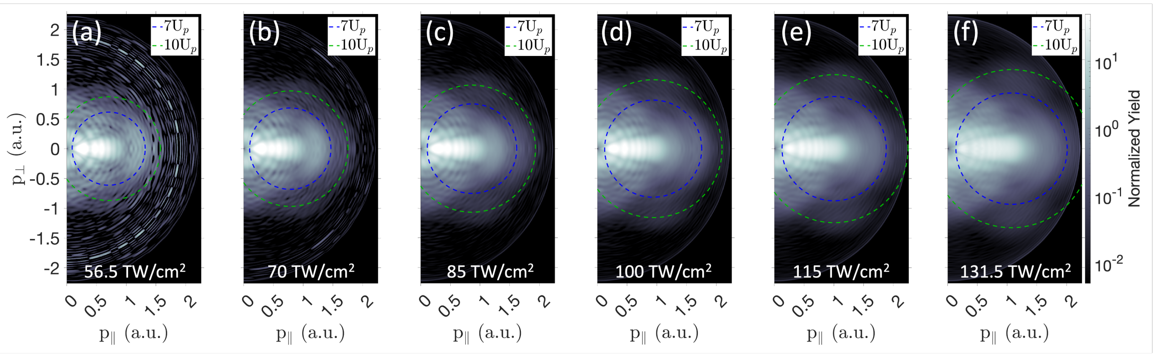

Figure 7.3: A selection of the filtered PMDs $Y^{0.85}(p_\parallel, p_\perp, I)$ at six selected intensities indicated by the annotation in the bottom of each panel. Overlaid on the PMDs are rings corresponding to the maximum extent of electron scattering on the first (green) or second return (blue). The maximum extent of these rings on the $p_\parallel$ axis reach 10U$_p$ and 7U$_p$ respectively.

Using the same intensity sweep dataset used to analyze the TBEs, I observe that this feature does *not* scale with intensity the way that simple explanation would suggest. Instead, as the intensity increases, I observe this feature remains static at exactly one momentum, and first appears as the rescattering ring at 7U$_p$ passes through its position on the detector. This effect is shown in Figure 7.3. Panels 7.3(a)-(f) plot the PMDs $Y^{0.85}(p_\parallel, p_\perp, I)$ at six selected intensities: 56.5 TW/cm^2, 70 TW/cm^2, 85 TW/cm^2, 100 TW/cm^2, 115 TW/cm^2, and 131.5 TW/cm^2, respectively. Each PMD has been filtered with TCF$^{0.85}$ to remove the ATI rings and expose the underlying structure. These PMDs have all been normalized by the average value of their pixels, so as to be plotted on the same color scale. Two rescattering rings have been overlaid on the data: in blue is the rescattering ring corresponding to the maximum extent of second-return rescattered electrons and in green is the rescattering ring corresponding to the maximum extent of first-return rescattered electrons. Because the feature does not form until Panel 7.3(d), we can see that the feature located at approximately $p_\parallel = 1.75$ a.u. does not appear until the intensity of the laser is somewhere in the range 85 TW/cm^2 to 100 TW/cm^2. In Panels 7.3(d)-(f) the feature does not maintain its distance with respect to the blue rescattering ring, as it must if it was a simple rescattering cutoff.

In Figure 7.4 I examine this effect more closely. Here I've integrated the PMDs over a 15° sector extending from the origin to the maximum momentum of the detector. These integrated sectors are then independently normalized by their mean value, and stacked for each value of the laser intensity. Two curves corresponding to 7U$_p$ and 10U$_p$ are overlaid on the data stack. Here we can clearly observe that the feature first appears at a laser intensity of approximately 95 TW/cm^2 exactly as the 7U$_p$ curve passes through 1.75 a.u. The vertical orientation of this feature indicates that it indeed does not scale alongside 7U$_p$.

The behavior of this structure is reminiscent of resonance. Once the laser intensity reaches beyond a certain threshold, all lower values of laser intensity may still be experienced by electrons

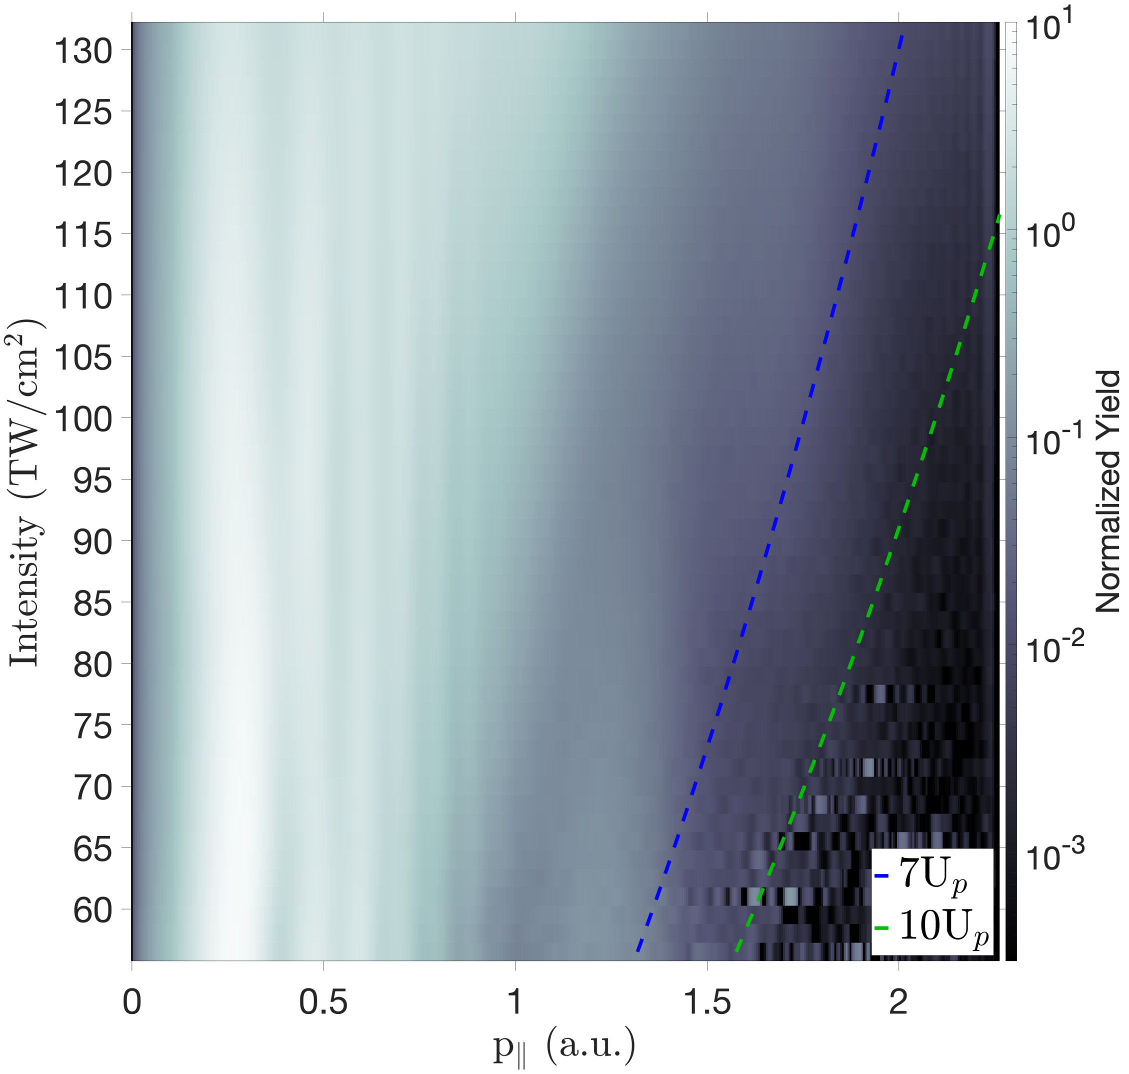

Figure 7.4: The azimuthal integration of all of the filtered PMDs in the sector bounded by azimuthal angle $\theta \in [0, 15°]$. Each integrated sector is plotted as the amplitude of each horizontal strip for each intensity. Overlaid on these plots are the values of 7U$_p$ (blue) and 10U$_p$ (green) for each intensity.

simply by moving away from the center of the laser focus. If, for whatever reason, one particular value of laser intensity causes the second-return rescattering electrons to primarily attain one set of final momentum along the $7U_p$ rescattering ring, and higher intensities may cause a broader range of final momenta, then an effect such as the one we observe here would occur. Furthermore, as the intensity of the laser extended much beyond the resonant value, the feature would gradually fade away as the overall amplitude of electron trajectories from the higher intensity portion of the laser would dominate. Given its shape, I believe this effect is also occurring for the vertical feature appearing at approximately $p_{\|} = 1.25$ a.u. The lowest value of the laser intensity used to generate a PMD in this dataset still results in a $7U_p$ curve which is above this feature, but I suspect if sufficient strong-field ionization of argon could be recorded at an intensity range of approximately 35 TW/cm^2 to 75 TW/cm^2 corresponding to $7U_p \approx 1$ to $7U_p \approx 1.5$, we would observe this feature "turn on" the same way the one at $p_{\|} = 1.75$ a.u. does. Unfortunately, this is a particularly difficult set of intensities to collect strong-field spectra at, as the rate of ionization is so low that in order to collect sufficient electron counts a very long data collection would be necessary.

While it is unclear what about the system, if anything, could form resonant features such as the ones we observe here, the mere fact that commonly recognized features of strong-field spectra do not follow their predicted intensity scaling is worthy of further study. It may also be discovered that the threshold intensity of these features aligning with the $7U_p$ momentum cutoff is purely a coincidence, and the second-return rescattered electrons simply underlie these prominent resonant features. In any case, an accurate characterization of these features will unlock further insight into the dynamics of strong-field ionization, and will surely be explored in the future.

Chapter 8

Conclusion

Throughout this dissertation, I took on the role of a cinematographer, carefully examining new techniques to disentangle and visualize the electron dynamics resulting in the intricate web of interference patterns which compose strong-field ionization spectra. The ultimate goal was to directly measure the fundamental properties of the electron trajectories whose pairwise interference results in overlapping layers of holograms. Developing tools that extract information on the timing of these electrons, particularly in terms of their phase of ionization, opens the door to more detailed studies of ultrafast electron dynamics and quantum control. To that end, I introduced two powerful analytical and experimental techniques: the time-correlation filter and three-color ionization phase extraction.

The time-correlation filter relied on the key insight that the Fourier transform of an angle-resolved power spectrum yields a temporal autocorrelation function relating the sum of the products of all electron trajectories whose ionization times are separated by a specific time separation. This allows me to directly manipulate the anisotropy parameters generated in the Abel inversion processing step to isolate specific combinations of electron trajectories. In particular, I looked at re-interpolating the anisotropy parameters to yield power spectra before low-pass filtering the results at defined time separations. When these filtered anisotropy parameters are used to generate the filtered PMD, every hologram formed from the interference of two trajectories separated by more than the selected time separation is removed, exposing the underlying structure. This is particularly useful for removing ATI interference rings, which encode no useful information of subcycle electron dynamics and obscure other, more interesting electron holograms. Simply removing the ATI rings immediately revealed a new type of hologram, the spider-leg modulations, which had not previously been identified [36]. These spider-leg modulations in tandem with the interference fringes of the Carpet provide unambiguous references for comparisons to TDSE calculations in order to determine the intensity of the ionizing laser.

In order to more accurately make quantitative measurements of electron trajectories, the normalized residual of a pair of time-correlation filtered PMDs can be calculated, which reveal all of the holograms whose contributing pairs of trajectories are separated within the window of the time separations of the time-correlation filters acting on the PMD. I used this windowing technique to directly extract the type of trajectory which contributes to the spider-leg modulations, concluding that the modulations are the result of the pairwise contribution of the direct trajectory interfering with each of the two types of forward scattered trajectories. I also isolated the Carpet as forming within the time separation range of 0.45 cycles to 0.7 cycles, which is a direct confirmation of the predictions of the CQSFA and settles a point of some controversy among the holography community [2].

I confirmed the validity of the time-correlation filter by testing it against well-known holographic parameters and by directly applying it to the CQSFA and comparing to analytical results. The time-correlation filter accurately revealed the ATI rings and the Carpet holograms in their predicted time-separation ranges, and preserved the spider-leg structure at all cutoffs. When applied to the CQSFA calculation of the Carpet, I saw reasonable agreement, but noted some limitations to the technique. Most notably, along the $p_\perp$ axis error is minimized, but increases as we progress azimuthally towards the $p_\parallel$ axis, where it attains a maximum value of approximately 0.25 cycles. For the applications presented in this thesis, that error is mostly subsumed into the width of the time-correlation filter windows we constructed, but the azimuthal dependence is an important factor to consider when assessing the accuracy of quantitative measurements of the time-correlation filter.

Ultimately, the time-correlation filter is significant to the study of strong-field spectra primarily because of how easily it can be applied to existing angularly- and energy-resolved data. As a pure analysis technique, the time-correlation filter can be applied across many measurements, even (or especially) retroactively, without requiring any modification to existing experimental designs. The technique unlocks a new observable in strong-field ionization studies which directly examines the properties of electron trajectories at the moment of their ionization. This opens the door to exciting new measurements of strong-field dynamics, and may serve to revitalize previously analyzed datasets to derive new insight from well-studied experiments.

Whereas the time-correlation filter is purely an analytical technique, alternative methods of measuring electron trajectory ionization phases can be achieved by shaping the probing field itself. I examined how to retrieve these phases by modifying the ionizing laser field with small perturbations of its second and third harmonic whose relative phase to the fundamental field can be finely controlled. I began by examining the two-color case first, and explored the strengths and weaknesses of the phase-of-the-phase technique. This technique relies on generating a single sinusoidal fit for each pixel in the PMD which relates the pixel's amplitude against the second harmonic phase. By examining the optimal second harmonic phase that results in the maximum enhancement of the sinusoidal fit, the two-color phase-of-the-phase technique attempts to retrieve the ionization phase of the trajectory arriving at that pixel. For pixels in which only one electron trajectory contribute to the final amplitude, this technique produces quantitative measurements of those trajectories' ionization phase; however, within regions of the PMD which contain multiple contributing trajectories, the method breaks down and cannot produce useful measurements of the ionization phase of any trajectory. Instead, two-color phase-of-the-phase can be used in these regions to determine the boundaries between different types of electron holograms at which the amplitudes of the contributing trajectories for the two holograms become roughly equal. Additionally, small modulations of the phase-of-the-phase signal can be traced back to subtle holographic features, though these modulations can do little more than provoke speculation as to their origin.

As my focus in this thesis has been on probing electron holograms, I must examine regions of the PMD which include more than one contributing trajectory, and thus cannot rely on two-color phase-of-the-phase measurements. I proposed the addition of a second phase-controlled perturbation of the ionizing field, such that the shaped ionizing pulse contains three colors. Using a calculation involving the ADK ionization rate of the two trajectories that appear in the Carpet hologram, I generated a single-pixel simulation that showed how the amplitude of a highly resolved three-color experiment would respond to the two dimensions of relative-phase control. This single-pixel simulation was used as the model to demonstrate the analysis technique of three-color ionization phase extraction, which aims to directly measure the ionization phases of multiple contributing trajectories to each individual pixel in strong-field PMDs.

The method of three-color ionization phase extraction first converts the pixel into its corresponding candidate phase distribution, which relates the possible ionization phases predicted by the second harmonic perturbation to the possible ionization phases predicted by the third harmonic perturbation. As the only physical values of this distribution lie along the central diagonal where these ionization phases agree for each color, a strip along this diagonal is collected and integrated to yield an amplitude versus ionization phase curve. Centroiding this curve and finding the ionization phases corresponding to the peaks of the result yield accurate measurements of ionization phases of each trajectory. The prominences of these peaks can then be used either as a measure of the confidence of the measurement or the relative amplitudes of the trajectories themselves. It is unfortunately a downside of the method that these two metrics are both coupled to the prominence of the measured peaks, and calculations from theory can be used to help distinguish the two.

I presented the technique of three-color ionization phase extraction on a three-color dataset. Directly examining the single pixel within the Carpet hologram at the same momenta as the simulation showed the ionization phases corresponding to the two contributing trajectories could be measured to within approximately 5° of phase. This is a stellar result, and represents a quantitative measure of an ultrafast process to within 37 attoseconds of its calculated value. Applying the phase extraction technique on the full three-color PMD accurately produced the ionization phases of the most probable electron trajectories in all regions, and a second ionization phase could be extracted for applicable regions by examining the prominence ratio of the peaks.

The methodology of overlapping three laser colors with two degrees of phase control is a powerful approach to probe electron trajectories while they are still local to the parent ion. The success of three-color ionization phase extraction allows experimentalists to modify their ionizing probe to measure the properties of electron trajectories in many contexts. By employing a probe pulse shaped with two harmonic perturbation in pump-probe experiments, experimentalists can make molecular movies tracking how the fundamental parameters of electron trajectories change in response to excitation or dissociation events. In electron-ion coincidence or covariance measurements, experimentalists can further correlate the types of electron trajectories coinciding with specific charged fragment combinations.

There is so much left to discover in the world of strong-field ionization. Features which defy simple explanation regularly appear in routine measurements. By designing new methods that can observe previously inaccessible properties of strong-field ionization, the electron cinematographer expands the set of tools which help explain these anomalous features. For my own part, I have high hopes that the techniques of time-correlation filtering and three-color ionization phase extraction will be employed and expanded upon to make exciting new discoveries in more complex systems.

Appendix A

How to Clean Contaminated MCPs

If you are reading this appendix because it is directly relevant to your current situation, I am very sorry for your misfortune. I also, unfortunately, had the type of accident that required me to invest significant time in puzzling out how to clean something as costly and delicate as a pair of MCPs. Given that there is very little available documentation on this topic, I decided it was worthy of a small passage in the appendix of this thesis.

While working to reassemble our detector stack after replacing old MCPS with new ones, a ceramic spacer separating the back MCP from the phosphor screen buckled, and the back MCP scraped across the phosphor. This coated the back MCP in a fine layer of the phosphor, and completely ruined the phosphor screen. Not wanting to discard an entire detector stack, I looked up the common ways to clean an MCP. The results were not encouraging. There were two main suggestions: use a small hand-held puffer to gently blow off loose particulates or use a single-haired brush to very gentle flick individual grains of dust off the surface of the MCP. Neither of these were effective solutions for how contaminated these MCPs were, so I had to get creative.

One (unsuccessful) avenue I pursued was using a polymer optic cleaning kit. These advertise as effective tools for removing heavy particulate contamination, while being gentle enough that the underlying optic remains undamaged. They work by applying a thin layer of liquid polymer over the contaminated surface, waiting for the polymer to dry, and then using a removal tool to peel the dried polymer off the optic along with any contamination. I tested this polymer cleaner on a previously discarded MCP that was too damaged to ever go back in service, and the results were disastrous. After applying a polymer layer to the MCP, the liquid polymer began bleeding through the plate's micro-channels, filling them and coating the backside of the plate as well. When it was time to remove the polymer, it was very clear that removal would be impossible to do without damaging the MCP. When I attempted to remove it anyway to see the extent of the damage, the MCP split in half before any portion of the polymer had been removed. I deemed this technique to be a poor option.

After pitching and discarding a few other candidate procedures, I settled on what appeared at first glance to be one of the more potentially damaging techniques: sonication and baking. MCPs are exceptionally brittle and easily damaged, and so have to be handled delicately. Sonication is not a delicate procedure; however, sonicating components in a bath of acetone and then sonicating them in a bath of isopropanol is a standard procedure for cleaning materials for use in high-vacuum operations. I began by filling a beaker with a sufficient amount of acetone that I could fully submerge the MCP in, and held the MCP by hand and swirled it in the liquid. Then I gently placed the MCP flat on the bottom of the beaker and placed it into an Ultra Sonicator. I let it sonicate on the lowest setting for twenty minutes, then I transferred the MCP to a beaker filled with isopropanol and repeated the same process. I did this for both MCPs. After sonicating, I baked the MCPs in a sterilizing furnace. I initially set the temperature to fifty degrees Celsius, and increased the temperature in ten degree steps to a final value of 200 degrees over the course of about an hour. I let the MCPs bake at this temperature for three hours before removing them. At this point, the MCPs appeared pristine, but it was unclear if there was still contamination or physical damage within the micro-channels themselves. I tested them by reassembling the detector stack, inserting the assembly back into our VMI chamber, and collecting some test data. Once everything was ready, we followed the standard MCP initialization procedure, and saw that everything was normal. The detector assembly was collecting electrons exactly the way new MCPs should!

Bibliography

[1] K. Amini, J. Biegert, F. Calegari, A. Chacón, M. F. Ciappina, A. Dauphin, D. K. Efimov, C. Figueira de Morisson Faria, K. Giergiel, P. Gniewek, A. S. Landsman, M. Lesiuk, M. Mandrysz, A. S. Maxwell, R. Moszyński, L. Ortmann, J. Antonio Pérez-Hernández, A. Picón, E. Pisanty, J. Prauzner-Bechcicki, K. Sacha, N. Suárez, A. Zaïr, J. Zakrzewski, and M. Lewenstein. Symphony on strong field approximation. *Reports on Progress in Physics*, 82(11):116001, November 2019.

[2] N. Werby, A. S. Maxwell, R. Forbes, C. Figueira de Morisson Faria, and P. H. Bucksbaum. Probing two-path electron quantum interference in strong-field ionization with time-correlation filtering. *Phys. Rev. A*, 106:033118, September 2022.

[3] C. Figueira de Morisson Faria and A. S. Maxwell. It is all about phases: ultrafast holographic photoelectron imaging. *Reports on Progress in Physics*, 83(3):034401, January 2020.

[4] F. Krausz and M. Ivanov. Attosecond physics. *Rev. Mod. Phys.*, 81:163–234, February 2009.

[5] P. B. Corkum and F. Krausz. Attosecond science. *Nat. Phys.*, 3(6):381–387, 2007.

[6] C. D. Lin, A. Le, C. Jin, and H. Wei. *Attosecond and Strong-Field Physics: Principles and Applications*. Cambridge University Press, Cambridge, 2018.

[7] S. Ghimire, G. Ndabashimiye, A. D. DiChiara, E. Sistrunk, M. I. Stockman, P. Agostini, L. F. DiMauro, and D. A. Reis. Strong-field and attosecond physics in solids. *J. Phys. B: At. Mol. Opt. Phys.*, 47(20):204030, October 2014.

[8] X. Bian and A. D. Bandrauk. Attosecond time-resolved imaging of molecular structure by photoelectron holography. *Physical Review Letters*, 108(26):263003, June 2012.

[9] G. Porat, G. Alon, S. Rozen, O. Pedatzur, M. Krüger, D. Azoury, A. Natan, G. Orenstein, B. D. Bruner, M. J. J. Vrakking, and N. Dudovich. Attosecond time-resolved photoelectron holography. *Nature Communications*, 9(1):2805, December 2018.

[10] M. Hentschel, R. Kienberger, C. Spielmann, G. A. Reider, N. Milosevic, T. Brabec, P. B. Corkum, U. Heinzmann, M. Drescher, and F. Krausz. Attosecond metrology. *Nature*, 414(6863):509–513, 2001.

[11] I. Gabalski, M. Sere, K. Acheson, F. Allum, S. Boutet, G. Dixit, R. Forbes, J. M. Glownia, N. Goff, K. Hegazy, A. J. Howard, M. Liang, M. P. Minitti, R. S. Minns, A. Natan, N. Peard, W. O. Rasmus, R. J. Sension, M. R. Ware, P. M. Weber, N. Werby, T. J. A. Wolf, A. Kirrander, and P. H. Bucksbaum. Transient vibration and product formation of photoexcited CS2 measured by time-resolved x-ray scattering. *The Journal of Chemical Physics*, 157(16):164305, October 2022.

[12] T. J. A. Wolf, D. M. Sanchez, J. Yang, R. M. Parrish, J. P. F. Nunes, M. Centurion, R. Coffee, J. P. Cryan, M. Gühr, K. Hegazy, A. Kirrander, R. K. Li, J. Ruddock, X. Shen, T. Vecchione, S. P. Weathersby, P. M. Weber, K. Wilkin, H. Yong, Q. Zheng, X. J. Wang, M. P. Minitti, and T. J. Martínez. The photochemical ring-opening of 1,3-cyclohexadiene imaged by ultrafast electron diffraction. *Nature Chemistry*, 11(6):504–509, June 2019.

[13] C. Cheng, R. Forbes, A. J. Howard, M. Spanner, P. H. Bucksbaum, and T. Weinacht. Momentum-resolved above-threshold ionization of deuterated water. *Phys. Rev. A*, 102:052813, November 2020.

[14] D. Strickland and G. Mourou. Compression of amplified chirped optical pulses. *Optics Communications*, 56(3):219–221, 1985.

[15] F. Crameri. Scientific colour maps, October 2023.

[16] P. B. Corkum. Plasma perspective on strong field multiphoton ionization. *Physical Review Letters*, 71(13):1994–1997, September 1993.

[17] A. S. Maxwell, C. Figueira de Morisson Faria, X. Lai, R. Sun, and X. Liu. Spiral-like holographic structures: Unwinding interference carpets of coulomb-distorted orbits in strong-field ionization. *Phys. Rev. A*, 102:033111, September 2020.

[18] A. S. Maxwell, A. Al-Jawahiry, T. Das, and C. Figueira de Morisson Faria. Coulomb-corrected quantum interference in above-threshold ionization: Working towards multi-trajectory electron holography. *Phys. Rev. A*, 96(2):023420, 2017.

[19] A. S. Maxwell, A. Al-Jawahiry, X. Y. Lai, and C. Figueira de Morisson Faria. Analytic quantum-interference conditions in coulomb corrected photoelectron holography. *J. Phys. B: At. Mol. Opt. Phys.*, 51:044004, 2018.

[20] A. S. Maxwell and C. Figueira de Morisson Faria. Coulomb-free and coulomb-distorted recolliding quantum orbits in photoelectron holography. *J. Phys. B: At. Mol. Phys.*, 51:124001, 2018.

[21] N. Werby, A. Natan, R. Forbes, and P. H. Bucksbaum. Disentangling the subcycle electron momentum spectrum in strong-field ionization. *Physical Review Research*, 3(2):023065, April 2021.

[22] X. Y. Lai, C. Poli, H. Schomerus, and C. Figueira de Morisson Faria. Influence of the coulomb potential on above-threshold ionization: A quantum-orbit analysis beyond the strong-field approximation. *Phys. Rev. A*, 92:043407, October 2015.

[23] W. Koch and D. J. Tannor. A three-step model of high harmonic generation using complex classical trajectories. *Annals of Physics*, 427:168288, 2021.

[24] S. V. Popruzhenko, G. G. Paulus, and D. Bauer. Coulomb-corrected quantum trajectories in strong-field ionization. *Phys. Rev. A*, 77:053409, May 2008.

[25] A. Rouzée, S. Guérin, O. Faucher, and B. Lavorel. Field-free molecular alignment of asymmetric top molecules using elliptically polarized laser pulses. *Physical Review A*, 77(4):043412, April 2008.

[26] S. Pabst, P. J. Ho, and R. Santra. Computational studies of x-ray scattering from three-dimensionally-aligned asymmetric-top molecules. *Physical Review A*, 81(4):043425, April 2010.

[27] M. Li, M. Liu, J. Geng, M. Han, X. Sun, Y. Shao, Y. Deng, C. Wu, L. Peng, Q. Gong, and Y. Liu. Experimental verification of the nonadiabatic effect in strong-field ionization with elliptical polarization. *Physical Review A*, 95(5):053425, May 2017.

[28] X. Ren and J. Zhang. Broken interference structure during ionization process of H atom in strong laser fields. *Optik (Stuttg).*, 181:287–292, 2019.

[29] C. M. Maharjan, A. S. Alnaser, I. Litvinyuk, P. Ranitovic, and C. L. Cocke. Wavelength dependence of momentum-space images of low-energy electrons generated by short intense laser pulses at high intensities. *J. Phys. B: At. Mol. Opt. Phys.*, 39(8):1955, 2006.

[30] T. Weber, H. Giessen, M. Weckenbrock, G. Urbasch, A. Staudte, L. Spielberger, O. Jagutzki, V. Mergel, M. Vollmer, and R. Dörner. Correlated electron emission in multiphoton double ionization. *Nature*, 405:658, 2000.

[31] T. Zuo, A. D. Bandrauk, and P. B. Corkum. Laser-induced electron diffraction: A new tool for probing ultrafast molecular dynamics. *Chemical Physics Letters*, 259(3):313–320, 1996.

[32] M. V. Ammosov, N. B. Delone, and V. P. Krainov. Tunnel ionization of complex atoms and atomic ions in electromagnetic field. In John A. Alcock, editor, *High Intensity Laser Processes*, volume 0664, pages 138–141. International Society for Optics and Photonics, SPIE, 1986.

[33] T. Yan, S. V. Popruzhenko, M. J. J. Vrakking, and D. Bauer. Low-energy structures in strong field ionization revealed by quantum orbits. *Phys. Rev. Lett.*, 105(25):253002, December 2010.

[34] C. D. Lin, A. Le, C. Jin, and H. Wei. Elements of the quantitative rescattering theory. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 51(10):104001, April 2018.

[35] S. Skruszewicz, J. Tiggesbäumker, K. H. Meiwes-Broer, M. Arbeiter, T. Fennel, and D. Bauer. Two-color strong-field photoelectronspectroscopy and the phase of the phase. *Physical Review Letters*, 115(4):043001, July 2015.

[36] N. Werby, A. S. Maxwell, R. Forbes, P. H. Bucksbaum, and C. Figueira de Morisson Faria. Dissecting subcycle interference in photoelectron holography. *Physical Review A*, 104(1):013109, July 2021.

[37] P. B. Corkum. Plasma perspective on strong field multiphoton ionization. *Phys. Rev. Lett.*, 71:1994–1997, September 1993.

[38] N. B. Delone and V. P. Krainov. Energy and angular electron spectra for the tunnel ionization of atoms by strong low-frequency radiation. *JOSA B*, 8(6):1207–1211, June 1991.

[39] A. S. Landsman and U. Keller. Attosecond science and the tunnelling time problem. *Physics Reports*, 547:1–24, 2015.

[40] P. A. Korneev, S. V. Popruzhenko, S. P. Goreslavski, T. M. Yan, D. Bauer, W. Becker, M. Kübel, M. F. Kling, C. Rödel, M. Wünsche, and G. G. Paulus. Interference carpets in above-threshold ionization: From the coulomb-free to the coulomb-dominated regime. *Phys. Rev. Lett.*, 108:223601, May 2012.

[41] D. Bauer and P. Koval. Qprop: A schrödinger-solver for intense laser–atom interaction. *Computer Physics Communications*, 174(5):396–421, March 2006.

[42] H. Kang, A. S. Maxwell, D. Trabert, X. Lai, S. Eckart, M. Kunitski, M. Schöffler, T. Jahnke, X. Bian, R. Dörner, and C. Figueira De Morisson Faria. Holographic detection of parity in atomic and molecular orbitals. *Physical Review A*, 102(1):013109, 2020.

[43] X. M. Tong and C. D. Lin. Empirical formula for static field ionization rates of atoms and molecules by lasers in the barrier-suppression regime. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 38(15):2593–2600, August 2005.

[44] H. Kleinert. *Path Integrals in Quantum Mechanics, Statistic, Polymer Physics and Financial Markets*. World Scientific, fifth edition, 2009.

[45] W. Becker, F. Grasbon, R. Kopold, D. Milošević, G. Paulus, and H. Walther. Above-threshold ionization: From classical features to quantum effects. *Advances in Atomic, Molecular, and Optical Physics*, 48:35–98, December 2002.

[46] S. Brennecke, N. Eicke, and M. Lein. Gouy's phase anomaly in electron waves produced by strong-field ionization. *Physical Review Letters*, 124(15):153202, 2020.

[47] X. Y. Lai, S. G. Yu, Y. Y. Huang, L. Q. Hua, C. Gong, W. Quan, C. Figueira de Morisson Faria, and X. J. Liu. Near-threshold photoelectron holography beyond the strong-field approximation. *Phys. Rev. A*, 96:013414, July 2017.

[48] A. S. Maxwell, S. V. Popruzhenko, and C. Figueira de Morisson Faria. Treating branch cuts in quantum trajectory models for photoelectron holography. *Phys. Rev. A*, 98:063423, 2018.

[49] 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(9):3477–3484, September 1997.

[50] C. S. Lehmann, N. B. Ram, and M. H. M. Janssen. Velocity map photoelectron-photoion coincidence imaging on a single detector. *Review of Scientific Instruments*, 83(9):093103, 2012.

[51] M. F. Kling, J. Rauschenberger, A. J. Verhoef, E. Hasović, T. Uphues, D. B. Milošević, H. G. Muller, and M. J. J. Vrakking. Imaging of carrier-envelope phase effects in above-threshold ionization with intense few-cycle laser fields. *New Journal of Physics*, 10(2):025024, February 2008.

[52] 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. *The Journal of Chemical Physics*, 141(22):221101, December 2014.

[53] G. M. Roberts, J. L. Nixon, J. Lecointre, E. Wrede, and J. R. R. Verlet. Toward real-time charged-particle image reconstruction using polar onion-peeling. *Review of Scientific Instruments*, 80(5):053104, May 2009.

[54] T. Szołdra, M. F. Ciappina, N. Werby, P. H. Bucksbaum, M. Lewenstein, J. Zakrzewski, and A. S. Maxwell. Femtosecond pulse parameter estimation from photoelectron momenta using machine learning. *New Journal of Physics*, 25(8):083039, August 2023.

[55] F. Allum, N. Anders, M. Brouard, P. Bucksbaum, M. Burt, B. Downes-Ward, S. Grundmann, J. Harries, Y. Ishimura, H. Iwayama, L. Kaiser, E. Kukk, J. Lee, X. Liu, R. S. Minns, K. Nagaya, A. Niozu, J. Niskanen, J. O'Neal, S. Owada, J. Pickering, D. Rolles, A. Rudenko, S. Saito, K. Ueda, C. Vallance, N. Werby, J. Woodhouse, D. You, F. Ziaee, T. Driver, and R. Forbes.

Multi-channel photodissociation and XUV-induced charge transfer dynamics in strong-field-ionized methyl iodide studied with time-resolved recoil-frame covariance imaging. *Faraday Discussions*, 228(0):571–596, May 2021.

[56] S. Li, Z. Guo, R. N. Coffee, K. Hegazy, Z. Huang, A. Natan, T. Osipov, D. Ray, A. Marinelli, and J. P. Cryan. Characterizing isolated attosecond pulses with angular streaking. *Opt. Express*, 26(4):4531–4547, February 2018.

[57] U. Even. The Even-Lavie valve as a source for high intensity supersonic beam. *EPJ Techniques and Instrumentation*, 2(1):17, December 2015.

[58] D. A. Dahl. Simion 3D version 6.0 user's manual. Technical Report INEL-95/0403-Rev.4, EG and G Idaho, Inc., Idaho Falls, ID (United States), November 1995.

[59] P. Hoess and K. Fleder. Response of very-fast-decay phosphors in image intensifier tubes for CCD readout devices. In Kazuyoshi Takayama, Tsutomo Saito, Harald Kleine, and Eugene V. Timofeev, editors, *24th International Congress on High-Speed Photography and Photonics*, volume 4183, pages 127–132. International Society for Optics and Photonics, SPIE, 2001.

[60] L. J. Zipp. *Probing ultrafast electron dynamics in atoms and molecules.* Stanford University, 2017.

[61] 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(11):4989–4996, November 2004.

[62] W. Becker, S. P. Goreslavski, D. B. Milošević, and G. G. Paulus. The plateau in above-threshold ionization: The keystone of rescattering physics. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 51(16):162002, August 2018.

[63] Y. Liu, J. Tan, M. He, H. Xie, Y. Qin, Y. Zhao, M. Li, Y. Zhou, and P. Lu. Photoelectron holographic interferences from multiple returning in strong-field tunneling ionization. *Optical and Quantum Electronics*, 51(5):145, April 2019.

[64] Z. Chen, A. Le, T. Morishita, and C. D. Lin. Quantitative rescattering theory for laser-induced high-energy plateau photoelectron spectra. *Physical Review A*, 79(3):033409, March 2009.

[65] H. G. Muller. Reconstruction of attosecond harmonic beating by interference of two-photon transitions. *Applied Physics B*, 74(1):s17–s21, June 2002.

[66] R. Jafari and R. Trebino. Highly reliable frequency-resolved optical gating pulse-retrieval algorithmic approach. *IEEE Journal of Quantum Electronics*, 55(4):1–7, 2019.

[67] D. Lee, S. Akturk, P. Gabolde, and R. Trebino. Experimentally simple, extremely broadband transient-grating frequency-resolved-optical-gating arrangement. *Opt. Express*, 15(2):760–766, Jan 2007.

[68] M. E. Anderson, A. Monmayrant, S. P. Gorza, P. Wasylczyk, and I. A. Walmsley. SPIDER: A decade of measuring ultrashort pulses. *Laser Physics Letters*, 5(4):259, January 2008.

[69] N. Wiener. Generalized harmonic analysis. *Acta Mathematica*, 55:117–258, 1930.

[70] J. O. Smith. *Introduction to Digital Filters with Audio Applications.* W3K Publishing, http://www.w3k.org/books/, 2007.

[71] D. Würzler, S. Skruszewicz, A. M. Sayler, D. Zille, M. Möller, P. Wustelt, Y. Zhang, J. Tiggesbäumker, and G. G. Paulus. Accurate retrieval of ionization times by means of the phase-of-the-phase spectroscopy, and its limits. *Physical Review A*, 101(3):033416, March 2020.

[72] J. Tan, Y. Zhou, M. He, Y. Chen, Q. Ke, J. Liang, X. Zhu, M. Li, and P. Lu. Determination of the ionization time using attosecond photoelectron interferometry. *Phys. Rev. Lett.*, 121:253203, December 2018.

[73] L. Cruz Rodriguez, T. Rook, B. B. Augstein, A. S. Maxwell, and C. Figueira de Morisson Faria. Forward and hybrid path-integral methods in photoelectron holography: Sub-barrier corrections, initial sampling, and momentum mapping. *Phys. Rev. A*, 108:033114, September 2023.

[74] T. Rook, L. Cruz Rodriguez, and C. Figueira de Morisson Faria. Influence of catastrophes and hidden dynamical symmetries on ultrafast backscattered photoelectrons. ArXiv:2403.02264, 2024.