Orbital Angular Momentum In High-Intensity Laser Interactions

Ramy Aboushelbaya
St Anne's College
University of Oxford

A Thesis submitted for the degree of
Doctor of Philosophy
Hilary 2020

"*He who seeks truth is not one who studies the writings of the ancients and, following his natural disposition, puts his trust in them, but rather the one who suspects his faith in them and questions what he gathers from them, the one who submits to argument and demonstration and not the sayings of human beings whose nature is fraught with all kinds of imperfection and deficiency. Thus the duty of the man who investigates the writings of scientists, if learning the truth is his goal, is to make himself an enemy of all that he reads, and, applying his mind to the core and margins of its content, attack it from every side. He should also suspect himself as he performs his critical examination of it, so that he may avoid falling into either prejudice or leniency.*"

---

al-Hasan Ibn al-Haytham

إن الباحث عن الحقيقة ليس هو من يدرس كتابات القدماء على حالتها ويضع ثقته فيها، بل هو من يُعلّق إيمانه بهم، ويتساءل ما الذي جناه منهم؟ هو الذي يبحث عن الحجة، ولا يعتمد على أقوال إنسان طبيعته يملؤها كل أنواع النقص والقصور؛ وبالتالي فإن من الواجب على من يحقق في كتابات العلماء، إذا كان البحث عن الحقيقة هدفه، أن يستنكر جميع ما يقرؤه، ويستخدم عقله حتى النخاع لبحث تلك الأفكار من كل جانب، وعليه أن يتشكك في نتائج دراسته أيضًا؛ حتى يتجنب الوقوع في أي تحيز أو تساهل.

الحسن ابن الهيثم

Acknowledgements

Out of all of the different parts of this thesis, this section was probably one of the hardest to write. The time I got to spend doing my DPhil at Oxford has been an incredible opportunity and an immense privilege. I am incredibly thankful to each and every person that either helped me on the way to starting my degree or during it, and it would be impossible for me to list them all now. So if you are one of those people and are not mentioned below, know that it is only for lack of pages not for lack of gratitude.

I am thankful to my mother Nariman who provides me with love and support each and every day of my life. Whatever success I may achieve, I owe it to the countless nights she spent teaching me how how to study and how to work. She not only gave me the tools to reach my goals but also instilled in me the values that define who I am.

I am thankful to my father Essam who always believes in me even when I don't believe in myself. He taught me the importance of science, and is one of the few people who supported my desire to pursue it. He always encourages me to follow my passions wherever they may lie, and without his support, be it moral or material, I would not be where I am today.

I am thankful to my brother Mohamed who always pushes me to be the best I can be. My years in Paris would not have been the same without him constantly having my back. I cannot remember a time where I did not look up to him and I feel incredibly lucky to have him as my role model.

I am thankful to my sister Safinaz who never tires of hosting me in her home. Only people who know me well, know how tiring that can truly be. Her patience, during my first years of adulthood, was one of the best gifts that I could have ever received.

I am thankful to Ahmed and Nada, on whom I know I can always count. Your home made this island a little less cold.

I am thankful to all my family: Intissar, Adel, Nabil, Idris, Nayla, Mai, Galal, Murad, Mama Nissa, Gedo Mohamed and many more. Each of you provided me with cherished memories that made it difficult to ever feel lonely.

I am thankful to my love, Salma. Although these past few years have been difficult and filled with ups and downs, I could always count on her to make everything easier. Even though we had to live apart, I never felt alone. She made every success more meaningful and every failure less hurtful.

I am thankful to the many wonderful teachers I have had over the years, whether in school or in university. I am particularly grateful to Karim Chikho who always encouraged us to question everything and never stop searching until we find an answer. I am also grateful to Rafik Marzouk who taught us that knowledge alone is meaningless without constant work and humility.

I am thankful to my supervisor Prof. Peter Norreys. He gave me an incredible opportunity to be a part of his group. His guidance all throughout my degree has been indispensable. He left me the space I needed to become an independent researcher while always offering support when I needed it the most.

I am thankful to all my colleagues and friends at Oxford. The memories I made here will forever stay with me. Like when Kevin, Alex, Luke, Muhammad and James welcomed me into the group during my first year and helped me through some of the most difficult days in the orange room. Or when Marko, Ben and I broke lead together in TA2. Or when Gabriel and I were looking for a shooting range in the middle of the swamps of Florida. Or the long discussions in the Rickty Press with Jimmy, Aimee, Picksley, Robin and Rob.

I am thankful to all the friends I had who shaped my life up till now. To Taimir, Naggar, Barakat, Shalakani, Baz, Gamal, Samaha, Habiba, Natalie, Menna, Menna, Antoine, Alia, Peryhane, Mahmoud, Yasser, Naguib and many more than I could possibly list, I am incredibly grateful to have had your presence in my life even if some of our paths have diverged.

I am thankful to the staff and members of St Anne's college for awarding me the Graduate Development Scholarship and providing me with an amazing level of support during my DPhil.

Finally, I would like to thank whoever is reading this thesis. This work is the product of a non-insignificant portion of my life, so if you are only reading a page, a chapter, or the whole thing, thank you for taking the time. I hope that you find within the upcoming pages something that is new or interesting or, hopefully, both.

Abstract

The orbital angular momentum (OAM) of light is one of the most intriguing properties of electromagnetic radiation. Although OAM is more commonly associated with the mechanical movement of massive particles, researchers have shown that, under certain conditions, laser beams can carry it. This is not merely a theoretical proposition, this idea was almost immediately experimentally proven by showing that OAM can be transferred between light and matter. This has, in turn, spurred an ever-increasing interest in leveraging the interesting proprieties of the OAM of light for various technological applications. This work focuses on the effect that OAM has on high-intensity laser interactions.

High-intensity lasers have been a boon to scientific investigations in their own right. They have allowed us to experimentally research astrophysical phenomena inside of laboratories, opened the possibility to tabletop particle accelerators and gotten us closer to useful fusion energy sources. More recently, we have been able to reach extreme intensities that allow us to probe the most fundamental interactions in the universe. Predictions that had been theorized decades ago by the pioneers of the quantum theory of matter are now close to being experimentally verifiable.

In the coming chapters, I look at the fundamental nature of the OAM of light and the many discussions it has spurred. I then show that it modifies an interaction known as vacuum photon-photon scattering where beams of light can interact with each other in the absence of any mediating matter violating the constraints established by the classical theory of electromagnetism. OAM provides an extra signal that makes this light-light interaction more identifiable in an experiment. On a more practical note, I continue to look specifically at high-intensity lasers and how they can be manipulated to produce high-intensity OAM-carrying beams and how said beams can be characterized.

ملخص

زخم الضوء الزاوي المداري (ز.ز.م) احدى خاصيات الإشعاعات الكهرومغناطيسية الأكثر إثارة للاهتمام. بالرغم من أن الز.ز.م خاصية مرتبطة أكثر بالحركة الميكانيكية للأجسام ذات الكتل، أثبت الباحثون أن أشعة الليزر يمكن أن تمتلك تلك الخاصية (أي خاصية الز.ز.م) في حالات محددة. هذه ليست مجرد افتراضية نظرية، إذ استطعنا أن نثبت عملياً أن الز.ز.م يُتبادل ما بين الضوء والمادة. هذان الاكتشافان حفزا محاولات لاستغلال الز.ز.م في العديد من الاختراعات العصرية. في هذه الرسالة، أناقش تأثير الز.ز.م على التفاعلات التي تقوم بها أشعة الليزر ذات الشدة العالية.

تلك أشعة الليزر كان لديها دوراً مهماً في عدة من الأبحاث العلمية الحديثة، فقد ساعدتنا في استكشاف بعض التفاعلات الفلكية في معامل أرضية، وخلقت فرص للحصول على مسرع للجسيمات صغير الحجم، كما أنها جعلت طاقة الانصهار النووي قرب المتناول. ومع الزيادة المهولة في شدة أشعة الليزر الحديثة، استطعنا أن نبدأ في استكشاف بعض التفاعلات المبدئية في كوننا من ناحية عملية، مما قربنا من إثبات العديد من التنبؤات النظرية التي قام بها مؤسسو نظرية الطبيعة الكمية.

في الفصول القادمة، أفحص طبيعة الز.ز.م الخاص بالضوء وأُثبِت كيف يُعدِل التفاعل المعروف بالتبعثر ما بين الفوتونات في الفراغ. هذا التفاعل ينافي أساسيات النظرية الكلاسيكية الكهرومغناطيسية حيث أنه يحدث ما بين إشعاعات ضوئية في الفراغ المنعدم من الجزيئات المادية. هذه الرسالة توضح كيف الز.ز.م يوفر علامة مميزة لهذا التفاعل تسمح لنا بأن نثبت وجوده في تجربة عملية. بالإضافة إلى الأوجه النظرية للز.ز.م، أتحدث أيضاً عن كيفية إنتاج أشعة ليزر ذات شدة عالية تحمل الز.ز.م وكيفية قياسها بشكل دقيق.

Contents

1 Introduction **1**

1.1 Overview and motivation 1
1.2 Thesis structure 3
1.3 Units and conventions 4
1.4 Author's contributions 5
1.5 Conferences and presentations 6

2 The angular momentum of light **9**

2.1 Context 9
2.2 Angular momentum in classical mechanics 11
2.3 Angular momentum in quantum mechanics 15
2.4 Electromagnetic waves 19
2.5 Angular momentum in field theory 29
2.6 Conclusion 41

3 Orbital angular momentum in photon-photon scattering **44**

3.1 Context 44
3.2 The classical electromagnetic field 46
3.3 Quantum electrodynamics 49
3.4 Effective field theory and the Euler-Heisenberg Lagrangian 61
3.5 The non-linear vacuum 63
3.6 Orbital angular momentum in photon-photon scattering 65
3.7 Proposed experimental setup 69
3.8 Conclusion 73

4 Generating high-intensity pulses carrying OAM **75**

4.1 Context 75
4.2 Chirped Pulse Amplification 76
4.3 The ASTRA GEMINI laser 79

4.4 Lasers carrying orbital angular momentum . . . 81
4.4.1 Conversion of Hermite-Gaussian modes . . . 82
4.4.2 Computer generated holograms . . . 84
4.5 Spiral phase optics . . . 90
4.6 Conclusion . . . 97

5 Characterising the OAM of high-intensity pulses 99
5.1 Context . . . 99
5.2 ASTRA experiment . . . 100
5.3 Cylindrical lens diagnostic . . . 103
5.4 Interferometric OAM diagnostic . . . 108
5.5 OAM projective measurement . . . 112
5.6 Conclusion . . . 116

6 Single-Shot Frequency Resolved Optical Gating 119
6.1 Context . . . 119
6.2 The FROG diagnostic . . . 121
6.3 Design . . . 125
6.4 Hybrid genetic algorithm . . . 130
6.4.1 Initialization . . . 133
6.4.2 Selection . . . 133
6.4.3 Crossover . . . 135
6.4.4 Mutation . . . 135
6.4.5 Local improvement . . . 135
6.4.6 Practical difficulties . . . 136
6.4.7 Comparing the hybrid algorithm to traditional algorithms . . . 136
6.5 Experimental results . . . 138
6.6 Conclusion . . . 141

7 Conclusions and future work 142
7.1 Summary . . . 142
7.2 Future work . . . 143
7.2.1 Experimental verification of photon-photon scattering . . . 143
7.2.2 Impact of OAM on laser-plasma interactions . . . 145
7.3 Afterword . . . 147

Bibliography 149

List of Figures

2.1 A classical point-like particle moving in uniform circular motion around a center O, which defines a standard right-handed Cartesian reference frame. 12

2.2 3D surface plot of the $|Y_2^0|^2$ spherical harmonic, clearly showing the azimuthal symmetry and the broken polar symmetry. 18

2.3 3D surface plot of the surfaces of constant phase for a standard Gaussian beam (upper) and for a Laguerre-Gaussian beam with $l = 1, p = 0$ (lower). This clearly shows that for the standard Gaussian the phase front is planar while it is helical for the Laguerre-Gaussian (LG) modes, illustrating why these modes have acquired the name "twisted light". 23

2.4 Simulated results for the LG envelope solution with $l = 1, p = 0$: (a) colour plot of the time-averaged transverse intensity at a particular transverse slice clearly showing the characteristic "doughnut" intensity profile that is due to the phase singularity at the origin; (b) colour plot of the instantaneous amplitude of $u_1^0(\mathbf{r})$ at a particular transverse slice. It shows the azimuthal variation of the amplitude from a positive maximum to a negative minimum. This profile rotates as the beam propagates and accomplish a full turn over a wavelength; (c) longitudinal slice showing the instantaneous amplitude of the envelope as the wave travels along the z-axis. It shows a similar focusing behaviour as the standard Gaussian mode, except of course for the singularity at the centre. 24

2.5 Diagram showing how the electric field vector of a circularly polarized light rotates drawing a helix as the wave propagates along the z-direction. It should be noted that, as opposed to the LG envelope, this does not provide any information on the shape of wavefront of the beam. Waves with planar wavefronts can still have circular polarization. 26

2.6 Energy diagram showing the quadrupole energy transitions $4S_{1/2} \rightarrow 3D_{5/2}$ in a calcium ion. The different states corresponding to the z projection of the total angular momentum m_j are split due the Zeeman effect caused by the constant external magnetic field $\mathbf{B}$. This breaks the symmetry of the system and makes the angular momentum projection states non-degenerate. The conservation of the total angular momentum in the stimulated absorption means that the sum of the incoming light's spin and orbital angular momentum must match the difference in the m_j number between the initial and final states. . . . 28

3.1 The two distinct momentum space Feynman diagrams representing the second-order expansion of the S-matrix for electron-positron interactions. Straight lines are Fermions and wavy lines represent photons. Left: This diagram represents electron-positron annihilation and recreation. Right: This diagram represents electron-positron scattering. . . 57

3.2 Feynman diagram representing one of the lowest-order contributions to photon-photon scattering. 58

3.3 Plot of the angular dependence of the field of the generated photons $|\mathbf{E}_4|^2$ to show the direction of propagation of the generated wave. θ is the polar angle defined, as it usually is in the spherical coordinate system, from the z-axis. 69

3.4 A sketch of the proposed geometry for the experiment showing the three main initial beams. The $2\omega_0$ (green) beams have a flat phase front whereas the ω_0 (red) beam has a helical front arising from the nonzero orbital angular momentum (OAM). The figure illustrates the generated $3\omega_0$ (blue) spiral photons. (a) and (b) are the frequency and OAM filters respectively, which will be used to remove background noise. 70

3.5 Plot showing the scaling the estimated number of generated photons as a function of initial laser power while fixing all other laser parameters (focal spot and pulse duration). 71

4.1 Schematic diagram describing the basic CPA procedure using a grating pair as its stretcher and compressor. 78

4.2 (a) The TA2 vacuum chamber where the high-intensity interactions take place (b) The compressor which takes the stretched and amplified pulse and compresses it before it is directed towards the interaction chamber. 80

4.3 Simulation of creating of a LG_0^1 modes using the two Hermite-Gaussian modes $HG_{1,0}$ and $HG_{0,1}$. (a)-(b) Show the field and intensity of $HG_{1,0}$, while (c)-(d) show those of $HG_{0,2}$ and (e)-(f) show those of $(1/\sqrt{2})(HG_{1,0} - iHG_{0,1})$. 83
4.4 Diagram of a typical holography setup: (a) information on the shape of the object is recorded onto the plate via the interference between the light scattered off of it and a reference laser. (b) When the holographic plate is illuminated with the same reference laser, a virtual image is then generated. 85
4.5 Simulated phase holograms for (a) $l = 1$ and (b) $l = 2$ with an 800 nm reference beam and $\theta = \pi/9$. 87
4.6 Simulated results comparing the action of a (a) binary phase hologram versus a (b) blazed hologram. Both holograms are generated to produce a state with $l = 1$ as can be seen from the form of the phase singularity at the centre. (c) and (d) show the far-field image produced by these holograms when a Gaussian laser, as in Eq.(4.4.16) is shown through the binary and blazed hologram respectively. 89
4.7 Model of a typical transmissive spiral phase plate optimized for converting an input Gaussian to the LG_0^l mode. 91
4.8 The focal spot of the ASTRA laser after passing through the transmissive continuous spiral phase plate. The beam has been apertured from its full size down to 20 mm to increase the size of the focal spot in order to better resolve the spatial structure imparted by the SPP and to avoid the step height inhomogeneity that existed on the plate away from its centre. As can be clearly seen from the image, the spot shows the distinctive singularity near the centre due to the phase discontinuity caused by the azimuthal dependence. 95

5.1 Diagram of the main parts of the ASTRA laser's target area showing the compressor chamber, the vacuum interaction chamber with its internal optical setup and the external diagnostics table with all the different OAM diagnostics. 102
5.2 (a) Picture of the inside of the ASTRA laser's interaction chamber (b) Picture of the diagnostics table outside of the interaction chamber showing the various OAM diagnostics. 104

5.3 (a-d) Simulated norm of a one-dimensional Fourier transform performed along the transverse x-axis of different Laguerre-Gaussian modes with $l = 0, +1, -1, +2$ respectively. The simulated norms clearly show the distinctive fringe pattern caused by the OAM. 105
5.4 (a) Optical setup for the diagnostic that was built at the ASTRA target area (b) Experimental result from a shot of the ASTRA laser after going through the SPP clearly showing the pattern similar to the $l = 1$ simulated pattern shown above (c) Experimental result from a continuous low-power 400 nm beam after it went through the same SPP placed in the opposite orientation. 106
5.5 (a) Shows the basic schematic of the optical setup for the OAM interferometer (b-c) Simulated results for the output of the interferometer in the case of normal Gaussian profile and a LG with $l = 1$ respectively. 109
5.6 Experimental results for the interferometric diagnostic on the ASTRA laser: (a) The result for the standard beam without the SPP, so $l = 0$, shows the typical fringes usually seen in interferometers; (b) When the SPP is placed in the beam path, the profile of the fringes change and a discontinuity similar to the one in Fig.(5.5.c) is seen indicating the presence of OAM. 111
5.7 Simulated results of the output of the interferometric diagnostic for (a) $l = 2$ and (b) $l = 3$. As can be seen from (a) as compared to Fig.(5.5), the LG_0^2 mode has a forked shaped but with two inner prongs instead of one. For LG_0^3, the pattern becomes extremely difficult to identify and exceeds the size of the window. 112
5.8 Schematic of an example of a projective diagnostic capable of projecting on $l = 0$ along path (a), $l = -1$ along path (b) and $l = 1$ along path (c). The remainder components of the beam can be inferred from the control measurement along path (d). 113
5.9 Simulated transverse profile of the the guides modes in a step-index fibre. The black circle represents the boundary between the core and the cladding. The two modes are (a) the LP_1^0 fundamental mode analogous to the free-space Gaussian and (b) the LP_1^1 mode analogous to the $l = 1$ mode. 115

5.10 Planned schematic for the photon-photon scattering detection experiment: (a) Narrow-band wavelength filters eliminates light at undesirable frequencies; (b) Mode converters only convert the appropriate OAM mode to Gaussian; (c) Lens to focus the radiation on the spatial mode filter; (d) Spatial mode filter leading to a photodetector. 118

6.1 Basic schematic of a generic SHG intensity autocorrelation diagnostic. 122
6.2 Simulated data for two pulses (A and B) with different intensity profiles: (a) The intensity autocorrelation signal for the two pulses; (b) The temporal intensity profile of the two pulses. This shows that although they have identical autocorrelation signals, they have vastly different intensity profiles. 124
6.3 Basic schematic of a generic SHG FROG diagnostic. 126
6.4 Diagram showing the two pulses mixing at angle θ on the second-harmonic generation crystal. The bigger the angle, the larger the maximum temporal window of the autocorrelator. 127
6.5 Schematic detailing the major optical components of the FROG. . . . 128
6.6 Picture showing the completed FROG device: (i) The second-harmonic generation autocorrelator used to generate the second harmonic signal for the FROG; (ii) The periscope and cylindrical lens used to flip the signal in the proper orientation and focus it onto the spectrometer's slit; (iii) The PI SP-300 spectrometer with the Andor NEO camera used to generate the FROG signal. 129
6.7 (a) Numerical FROG trace produced from a double pulse showing the distinct interference pattern that is usually used to calibrate FROG devices (b) 3D surface plot of the objective function across the two parameter scans of the guesses "around" the double pulse. 132
6.8 Flowchart of the hybrid algorithm. Purple rectangles represent steps of the genetic algorithm yellow represent steps of the PCGP algorithm and orange rectangles are the intermediate steps. 134
6.9 The graph shows the different progression for the two algorithms as they try to retrieve the original simulated double-pulse from its FROG trace. It should be noted that each iteration of the memetic algorithm contains multiple iterations of the local improvement using the PCGP. 137
6.10 (a) FROG trace produced by the femtosecond oscillator (b) Intensity of the retrieved pulse and its phase. 139

6.11 (a) FROG trace produced by the Vulcan short pulse beam (b) Intensity of the retrieved pulse and its phase. 140

7.1 Slices of the simulated intensity profile of various normalized LG modes at a particular y and z. The modes only differ in the l mode number while keep all other Gaussian parameters constant. As can be seen from the plot, for the same amount of energy, the peak intensity of an LG mode decreases with the increasing l number. 146

Abbreviations

AM angular momentum.

BEC Bose-Einstein condensate.

CBET cross-beam energy transfer.

CLF central laser facility.

CPA chirped pulse amplification.

CW continuous wave.

ELI extreme light infrastructure.

EPAC extreme photonics applications centre.

GVD group velocity dispersion.

HG Hermite-Gaussian.

ICF inertial confinement fusion.

LG Laguerre-Gaussian.

LHC large hadron collider.

LLE laboratory for laser energetics.

LPI laser-plasma interactions.

LWFA laser wakefield acceleration.

OAM orbital angular momentum.

OPA optical parametric amplification.

PMMA polymethyl methacrylate.

QCD quantum chromodynamics.

QED quantum electrodynamics.

QFT quantum field theory.

RABBITT reconstruction of attosecond beating by interference of two-photon transitions.

RAL Rutherford Appleton laboratory.

SAM spin angular momentum.

SBS stimulated Brillouin scattering.

SEL station of extreme light.

SPIDER spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses.

SPM spiral phase mirror.

SPO spiral phase optics.

SPP spiral phase plate.

SRS stimulated Raman scattering.

TA2 target area 2.

TA3 target area 3.

TiS Ti:Sapphire.

Symbols

A^μ four-potential.

$F^{\mu\nu}$ electromagnetic field tensor.

I moment of inertia.

J^μ four-current.

S action.

ϵ_0 vacuum permittivity.

$\eta_{\mu\nu}$ Minkowsky metric.

$\hat{\mathbf{J}}$ total angular momentum operator.

$\hat{\mathbf{L}}$ orbital angular momentum operator.

$\hat{\mathbf{S}}$ spin angular momentum operator.

$\mathbf{A}$ electromagnetic vector potential.

$\mathbf{B}$ magnetic field.

$\mathbf{E}$ electric field.

$\mathbf{\Psi}$ Dirac spinor.

$\mathbf{k}$ wavevector.

$\mathcal{H}$ Hamiltonian density.

$\mathcal{L}$ Lagrangian density.

LG_p^l Laguerre-Gaussian transverse envelope.

LP_m^l transverse envelope of linearly polarized mode.

μ_0 vacuum permeability.

ω angular frequency.

∂^μ four-gradient.

ϕ electric scalar potential.

k^μ four-wavevector.

l orbital angular momentum state.

Chapter 1

Introduction

The White Rabbit put on his spectacles.
"Where shall I begin, please your Majesty?" he asked.
"Begin at the beginning," the King said gravely,
"and go on till you come to the end: then stop."

Lewis Carroll, Alice in Wonderland

Clichés remind and reassure us that we're not alone,
that others have trod this ground long ago."

Miguel Syjuco

1.1 Overview and motivation

Light is one of the oldest studied natural phenomena. Throughout centuries of science and philosophy, our understanding of what light is and how it behaves has dramatically evolved. Thanks to the efforts and brilliance of many thinkers and scientists, we have come a long way since the days of the extramission theory of vision. Light or, more accurately, electromagnetic waves have come to form one half of the most accurate physical theory known to mankind, quantum electrodynamics (QED). The modern theoretical description of electromagnetic radiation, both in its "classical" and quantum mechanical form, has succeeded in explaining many observed properties and behaviours of light. QED has also made several predictions that contradicted

our intuitive understanding of how light behaves. Some of these predictions have been experimentally verified while others still remain unproven due to the difficulty of measuring their signals.

Despite this long history of extensive research, we are still learning new things about fundamental properties of light, which show us that our understanding of it may not be as complete as we think. One of the most recent examples of this was when the late Les Allen investigated the so-called Laguerre-Gaussian (LG) radiation modes in 1992 [1]. It was already well known that light carried an angular momentum associated with its spin [2]. What Allen showed was that certain types of electromagnetic radiation can carry another kind of angular momentum, which was termed orbital angular momentum (OAM).

Since its first discovery, OAM has attracted a lot of interest from a vast range of scientific fields ranging from telecommunication sciences [3] to quantum memories [4]. This is due to its many interesting qualities such as the fact that light carries a quantized amount of OAM, which has an infinite possible range [1]. This OAM also appeared to modify the way light interacts with matter in ways which can be beneficial for some use cases, such as modifying the selection rules of atomic transitions [5]. However, although many applications of OAM have been identified and even tested, this does not mean that we fully understand the fundamental physics behind it. There are still many questions surrounding the precise significance of OAM, its relationship to spin [6] and its quantum mechanical origin [7, 8].

Recently, researchers have even started to theoretically investigate the possible impact of OAM on high-power light-matter interactions. High-power laser science is a burgeoning field of experimental and theoretical physics. This is thanks to the ever increasing pulse powers that modern laser facilities are capable of producing such as the soon-to-be-completed ELI facility, which will generate a 10 PW pulse when fully operational [9] and the planned station of extreme light (SEL), which aims to produce a 100 PW one [10]. These incredible machines, and many others around the world, have allowed us to investigate phenomena which occur when matter transforms into a plasma - a "fourth" phase of matter that occurs when conditions are so extreme that electrons are ripped from their atoms, forming a quasi-neutral gas of ions and electrons [11]. Laser-plasma physics has offered new ways of achieving fusion energy [12], of creating small-scale particle accelerators [13] and of even studying astrophysical phenomena in terrestrial laboratories [14]. Researchers have shown that OAM affects laser-plasma interactions by modifying the plasma oscillations driven by the OAM-carrying lasers [15]. This has the potential of offering a new way of manipulating

the aforementioned laser plasma accelerators [16, 17], as well as controlling the laser-plasma instabilities that occur in laser-driven inertial confinement fusion.

Unfortunately, although the techniques to produce and characterise OAM modes had been honed to a great precision with small low-power lasers due to them being a subject of ongoing research for close to three decades now, this is not the case for high-power lasers. There are scarcely any experimental results on the production of high-power beams carrying OAM or their characterisation. As such, most of the previously mentioned findings concerning OAM in laser-plasma interactions remain theoretical predictions. This is, in part, because interest in exploring the effects of OAM on the interactions of high-power lasers is relatively new and the latter present some unique challenges due to the sheer scale of the optical intensities that would need to manipulated.

1.2 Thesis structure

In this thesis, I showcase the results of my investigations into the orbital angular momentum of light in the context of high-intensity laser interactions. My goal was to address parts of each of the open questions and problems outlined in the previous section. Although the work is by no means complete, I have arrived at some interesting results, which will be detailed and explored over the coming chapters. The thesis will be structured as follows:

1. Firstly, I outline the various types of angular momenta for point particles, as well as their origin both from a classical viewpoint and from the viewpoint of quantum mechanics. This will lead to the definition of the angular momentum of light and how it may be split into spin and orbital components. I showcase how both of these angular momenta can be derived for an arbitrary field simply from the symmetries obeyed by that field, ending with a discussion on the interpretation of OAM and the on-going controversy over its rigorous definition.

2. Then, an investigation into vacuum photon-photon scattering is detailed. By leveraging the effective field theory approach, a semi-classical framework is used to explore the effects of OAM on this interaction, deriving additional angular momentum conservation conditions for the scattering. This is followed by a proposed experimental setup that would leverage this additional conservation condition to make the signal produced by this incredibly weak interaction more detectable. Chapter 3 concludes with an approximate quantitative estimation

of the expected results of such an experiment as well as a discussion of some of the more practical aspects of its realisation.

3. Following the study of the interactions of high-intensity laser pulses carrying OAM, I turn to look at how these pulses can be generated. This chapter starts with a look at how standard high-intensity pulses are generated. Then, I go through a detailed rundown of the various methods that can be used for the generation of OAM-carrying lasers as well as their relative suitability for use on high-power lasers. Chapter 4 also contains experimental results of an OAM mode converter that was tested at the Central Laser Facility (CLF).

4. Of course, generating high-intensity pulses with OAM is only half the battle. For that reason, I turn my attention in Chapter 5 to the characterisation of the OAM of high-intensity pulses. There, I detail three very different measurement methods both from a theoretical point of view, and from a practical experimental point of view, using data gathered on the ASTRA laser at the CLF.

5. Measuring the OAM of high-power lasers is not the only important diagnostic for the successful analysis of experimental results. Most high-power lasers use a pulse length that is so short that it becomes impossible to characterise them using traditional electronic methods. As such, accurate ultra-fast pulse length measurements are one of the most important diagnostics in high-power laser facilities. Chapter 6 focuses on a novel implementation of the Frequency-Resolved Optical Gating (FROG) diagnostic, which is used to characterise both the pulse length and the phase evolution of laser pulses. It also outlined a new hybrid genetic algorithm for the analysis of the measurements of this diagnostic.

6. Finally, I conclude with a short summary of the main results as well as an overview of the future work that is planned based on them.

1.3 Units and conventions

The units and mathematical conventions that will be used throughout this thesis are as follows:

1. For most of the thesis, the standard international system of units (SI) will be used.

2. In some parts, specifically dealing with quantum field theory (QFT) models, the "natural" system of units will be used. There, all units will be normalized so that $\hbar = c = \epsilon_0 = 1$.

3. The convention used for the Minkowski metric is

$$\eta_{\mu\nu} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$$

4. Greek indices take values $\mu = 0, 1, 2, 3$ while Latin indices only take $i = 1, 2, 3$.

5. The four-gradient is defined as $\partial^\mu = (\partial_t, -\nabla)$

6. The totally anti-symmetric tensor $\epsilon^{\mu\nu\rho\sigma}$ is defined so that it is equal to 0 if any two of its indices are the same, $+1$ for any even number of permutations of $(0, 1, 2, 3)$ and -1 otherwise.

1.4 Author's contributions

- All the content of this thesis is my own writing, and the figures, plots and illustrations are my own creation.
- I have reproduced all the standard derivations found in the thesis.
- I derived the theoretical framework of including the treatment of OAM in the description of photon-photon scattering. I designed a potential experiment to detect it based on the predictions of the framework and calculated the estimated signal to be expected from such an experiment.
- I worked on the interpretation of OAM as an observable independent of spin angular momentum (SAM) from a quantum field theoretical point of view.
- I designed, tested and analysed the performance of the OAM mode converter for the ASTRA laser.
- I co-led the experimental campaign to field test various OAM diagnostics on the ASTRA laser. On that campaign, I designed, assembled, tested and analysed the results of the three different OAM diagnostics.

- I built the FROG diagnostic and tested it on the Vulcan laser. I also designed the hybrid genetic algorithm for the analysis of FROG measurements and coded its Python implementation.

Of course, I would be remiss if I did not acknowledge the support of my colleagues. I would have not been able to accomplish any of the above mentioned contributions without the constant assistance and ever present advice provided by them, especially by my supervisor Prof. Peter Norreys, and my collaborators Prof. Robert Bingham, Dr. Kevin Glize and Dr. Pedro Oliveira.

1.5 Conferences and presentations

1. Invited talk on the generation and measurement of OAM in high-intensity laser interactions and its effect on vacuum photon-photon scattering given at the 61st Annual Meeting of the American Physical Society's (APS) Division of Plasma Physics on 21-25 October 2019.

2. Contributed talk given at the Oxford Centre for High Energy Density Science (OxCHEDS) conference on 18-19 March 2019 on theoretical investigations of the effect of OAM on vacuum photon-photon scattering.

3. Contributed talk given at the European Physical Society conference on High Energy Physics (EPS-HEP) on 10-17 July 2019 on theoretical investigations of the effect of OAM on vacuum photon-photon scattering.

4. Poster presented at the International Conference on High Energy Density (ICHED) on 31 March5 April 2019 on theoretical investigations of the effect of OAM on vacuum photon-photon scattering.

5. Poster presented at the National Ignition Facility (NIF) user meeting on 4-7 February 2018 on improved algorithms for the analysis of the results of the Frequency-Resolved Optical Gating diagnostic.

6. Poster presented at the Central Laser Facility (CLF) Christmas meeting on 19-21 December 2016 on two-qubit geometric quantum gates for superconducting qubits.

Publications

[1] R. Aboushelbaya, K. Glize, A. F. Savin, M. Mayr, B. Spiers, R. Wang, N. Bourgeois, C. Spindloe, R. Bingham, and P. A. Norreys. Measuring the orbital angular momentum of high-power laser pulses. Submitted to Physics of Plasmas (February 2020).

[2] R. Aboushelbaya, K. Glize, A. F. Savin, M. Mayr, B. Spiers, R. Wang, J. Collier, M. Marklund, R. M. G. M. Trines, R. Bingham, and P. A. Norreys. Orbital angular momentum coupling in elastic photon-photon scattering. *Phys. Rev. Lett.*, 123:113604, Sep 2019.

[3] R. Aboushelbaya, A. F. Savin, L. Ceurvorst, J. Sadler, P. A. Norreys, A. S. Davies, D. H. Froula, A. Boyle, M. Galimberti, P. Oliveira, B. Parry, Y. Katzir, and K. Glize. Single-shot frequency-resolved optical gating for retrieving the pulse shape of high energy picosecond pulses. *Review of Scientific Instruments*, 89(10):103509, October 2018.

[4] M. W. Mayr, L. Ceurvorst, M. F. Kasim, J. D. Sadler, B. Spiers, K. Glize, A. F. Savin, N. Bourgeois, F. Keeble, A. J. Ross, D. R. Symes, R. Aboushelbaya, R. A. Fonseca, J. Holloway, N. Ratan, R. M. G. M. Trines, R. H. W. Wang, R. Bingham, L. O. Silva, P. N. Burrows, M. Wing, P. P. Rajeev, and P. A. Norreys. Wakefields in a cluster plasma. *Phys. Rev. Accel. Beams*, 22:113501, Nov 2019.

[5] J. D. Sadler, Y. Lu, B. Spiers, M. W. Mayr, A. Savin, R. H. W. Wang, R. Aboushelbaya, K. Glize, R. Bingham, H. Li, K. A. Flippo, and P. A. Norreys. Kinetic simulations of fusion ignition with hot-spot ablator mix. *Phys. Rev. E*, 100:033206, Sep 2019.

[6] J. D. Sadler, L. O. Silva, R. A. Fonseca, K. Glize, M. F. Kasim, A. Savin, R. Aboushelbaya, M. W. Mayr, B. Spiers, R. H. W. Wang, R. Bingham, R. M. G. M. Trines, and P. A. Norreys. Advantages to a diverging Raman amplifier. *Communications Physics*, 1(1):19, May 2018.

[7] A. F. Savin, A. J. Ross, R. Aboushelbaya, M. W. Mayr, B. Spiers, R. H.-W. Wang, and P. A. Norreys. Energy absorption in the laser-QED regime. *Scientific Reports*, 9(1):8956, June 2019.

Chapter 2

The angular momentum of light

> "*The mathematician plays a game in which he himself invents the rules while the physicist plays a game in which the rules are provided by nature, but as time goes on it becomes increasingly evident that the rules which the mathematician finds interesting are the same as those which nature has chosen*"
>
> Paul A. M. Dirac

2.1 Context

Angular momentum (AM) is one of the most important concepts in physics. It appears in practically all of its domains from classical mechanics to general relativity. From a fundamental point of view, AM and its conservation laws are related to the intrinsic symmetries of physical systems. It is thus a property that can be defined for all physical fields. However, its importance is in no way limited to abstract theoretical investigations; it has a wide range of applications. To name a few, AM is the cause behind the fine and hyperfine structures of atomic energy transitions, it is of vital importance to calculations of stellar dynamics and for the application of gyroscopic stabilisation technologies.

Usually, one thinks of angular momentum as a continuously-varying property that exists in massive particles that are performing some form of orbital motion, such as the earth around the sun or an electron orbiting a nucleus. The advent of quantum mechanics showed us that particles actually carry two types of *quantized* angular momenta. First, there is the regular one associated with the motion of the particle

which is called the OAM since it is the quantum analogue of the classical orbiting particle. OAM is quantized by an unbounded quantum number, i.e. a particle can be in any one of an infinite amount of discrete states. In addition to this, particles carry an intrinsic AM that is related to a purely quantum mechanical property called spin. This SAM is omnipresent no matter what kind of motion the particle is doing. Its possible values are finite and their range is determined by the type of particle under consideration.

However, AM is not only limited to massive particles. Quantum field theory (QFT) tells us that it can be defined for all physical fields including those which do not have mass, such as the electromagnetic field. Here again, it is split into a spin and an orbital component. In the case of the electromagnetic field, the spin of light had been a known and well studied phenomenon. Photons, the quantized excitations of electromagnetic fields, are considered spin-1 particles. This means that there are technically three possible spin states, but only two of them correspond to physically allowable photon states. The spin of light has been verified in experiments that studied the interaction between photons and matter. These experiments showed that photons with different spin states can impart different angular momenta onto matter, thereby confirming that this spin is indeed a type of angular momentum. Careful measurements have shown that the ratio between the energy of the photon and its SAM is σ/ω, where ω is the photon's frequency and σ is its spin quantum number that determines its spin state.

More recently, it has been shown that there are modes of electromagnetic radiation that can carry orbital angular momentum. These are called the Laguerre-Gaussian (LG) modes. Over the past three decades, they have been well characterised and there has been extensive research on how to generate these modes in all sorts of different regimes, from small on-chip solid state lasers to milliwatt tabletop lasers to the multi-terawatt beams available at high-power laser facilities. This is because the OAM of light has attracted a significant amount of research interest over a wide range of domains including, but not limited to, quantum information processing, quantum optics, classical optical communication, optical imaging, laser-plasma interactions and particle physics. In fact, the interest in OAM is no longer limited to academic research, with projects currently undergoing to use OAM multiplexing to enhance commercial telecommunication applications. In addition to methods of generating it, a significant amount of research has been undertaken into how to efficiently measure and characterise OAM-carrying light, with researchers now able to measure it with high precision at the single photon level.

Although research into the possible applications of OAM has greatly intensified and has shown some very promising results, this does not mean that its fundamental nature has been completely understood. In fact, there is a very active ongoing debate about the fundamental physics of the OAM of light and about how it can be defined from first principles. This mostly centres around whether or not it is possible to split the angular momentum of these fields in a consistent way into a spin and an orbital component or if that split is "artificial". This debate is quite relevant as it extends beyond just the definitions of angular momentum of photons and into current open questions in quantum chromodynamics (QCD) concerning the spin structure of nucleons. Experiments which have been performed over the past three decades dispelled earlier notions that the proton gains its spins almost exclusively from the spin of the constituent quarks [18, 19, 20]. This then led to an investigation of the SAM and OAM of gluons and quarks, the constituent particles of protons and the main fields of QCD, which revealed issues concerning the split of angular momentum of these fields similar to the ones present in the definition of the AM of the electromagnetic field [21, 7]. As such, investigating the OAM of light has the potential for broader applications to other interesting questions in fundamental physics.

This chapter will be organised as follows: first, I will discuss the various ways to define angular momentum in the classical and quantum descriptions of the dynamics of massive particles. Then, the angular momentum of classical electromagnetic waves will be defined. Following J. Allen's method, I will use the Laguerre-Gaussian solutions to the electromagnetic wave equation to show how the AM of light contains both an orbital and a spin component. Lastly, the symmetries of quantum fields will be used to show how two types of angular momenta can be defined from first principles for any generic quantum field, hopefully addressing some of the above-mentioned concerns about how to consistently define OAM for light.

2.2 Angular momentum in classical mechanics

Before I can address the AM of light, it is instructive to consider the classical definition of angular momentum in general, the most basic example of which can be seen in the case of an orbiting point-like particle. To this particle, one defines an angular momentum $\mathbf{J}$ that is related to its position and momentum via [22]

$$\mathbf{J} = \mathbf{r} \times \mathbf{p} \tag{2.2.1}$$

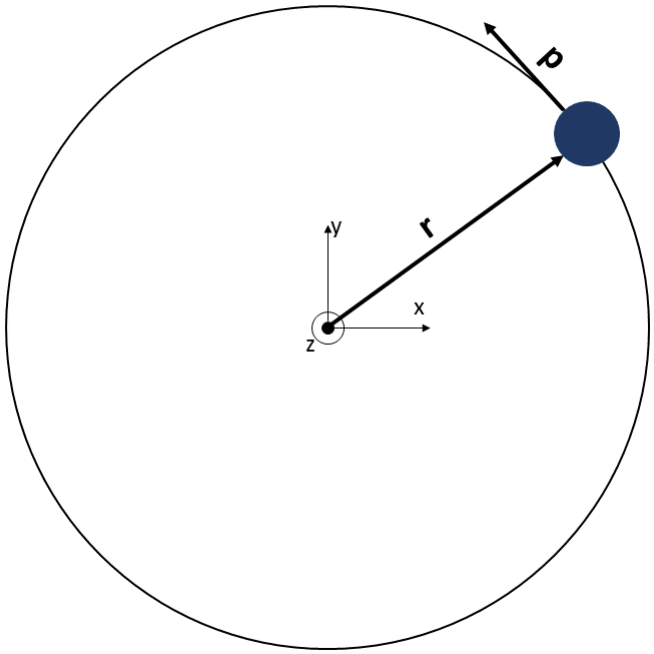

Figure 2.1: A classical point-like particle moving in uniform circular motion around a center O, which defines a standard right-handed Cartesian reference frame.

In the case of the simple setup in Fig.(2.1), the angular momentum takes on the standard form for uniform circular acceleration $\mathbf{J} = mR^2\omega\mathbf{e}_z = I\omega$, where R is the radius of the orbit, ω is the angular velocity of the particle and I is the moment of inertia of the system. The classical angular momentum of a system also obeys a law of motion analogous to that obeyed by the linear momentum:

$$\frac{\mathrm{d}\mathbf{J}}{\mathrm{dt}} = \mathbf{r} \times \mathbf{F} = \tau \tag{2.2.2}$$

where τ is the torque, or the moment of force, applied to the system. This means that, in the absence of torque, the total angular momentum of a system is conserved. A more fundamental origin of the definition of angular momentum and its conservation can be seen by considering the Lagrangian formulation of classical mechanics. There, a system can be completely described by its Lagrangian $\mathcal{L}(q_i, \dot{q}_i, t)$, where q_i and $\dot{q}_i$ are the generalised coordinates and their time derivatives respectively [22]. This Lagrangian defines the *action* of the system:

$$S = \int \mathcal{L}(q_i, \dot{q}_i, t)dt \tag{2.2.3}$$

By finding the extremum of this action over all possible paths between points in the phase space, the equations of motion that govern the dynamics of the system

can be found. This principle of least action can be summarised in the well-known Euler-Lagrange equations:

$$\frac{d}{dt}\left(\frac{\partial \mathcal{L}}{\partial \dot{q}_i}\right) = \frac{\partial \mathcal{L}}{\partial q_i} \tag{2.2.4}$$

The Lagrangian formulation offers an important insight into the quantities that are conserved during the evolution of the system through what is known as *Noether's theorem* [23]. Consider an infinitesimal transformation of the coordinates

$$q^i \to q^i + \epsilon f^i(\mathbf{q}) = q^i + \delta q^i \tag{2.2.5}$$

where ϵ is an small parameter and f is an arbitrary function of the coordinates, which models the action of the infinitesimal transformation. If one performs this transformation on the action, then it follows that the deviation in the action is written as

$$\delta S = \int_{t_1}^{t_2} \left(\frac{\partial \mathcal{L}}{\partial q_i} - \frac{\partial}{\partial t}\left(\frac{\mathcal{L}}{\partial \dot{q}_i}\right)\right)\delta q_i dt + \left[\frac{\partial \mathcal{L}}{\partial \dot{q}_i}\delta q_i\right]_{t_1}^{t_2} \tag{2.2.6}$$

where t_1 and t_2 are the initial and final times that parametrise the system's path in phase space. As seen in Eq.(2.2.4), classical systems that follow paths of least action obey the Euler-Lagrange equations and thus, for them, the first term in the integrand on the RHS of Eq.(2.2.6) cancels out. If the transformation defined in Eq.(2.2.5) is a symmetry of the system, i.e. leaves the action intact, then necessarily the boundary terms also needs to vanish for arbitrary t_1 and t_2. This means that one defines the following quantity, known as the Noether charge:

$$Q = \frac{\partial \mathcal{L}}{\partial \dot{q}_i} f_i(\mathbf{q}) \tag{2.2.7}$$

which is conserved along the path of the system. This signifies that for each symmetry transformation that leaves the action of the system unchanged, there exists an associated conserved quantity. For example, if the system is symmetric under simple spatial translations along a particular axis $\hat{\mathbf{n}}$, i.e. $f^i(\mathbf{q}) = n^i$, then the Noether charge is $Q = (\partial \mathcal{L}/\partial \dot{q}_i)n_i$. If the system is symmetric with respect to spatial translations along every axis, then all the $\partial \mathcal{L}/\partial \dot{q}_i$ are conserved. If one goes back to the previously defined system of an orbital particle, then it is obvious that the physics of the system remain unchanged under translations along the z-axis. Its Lagrangian is written as

$$\mathcal{L} = \frac{1}{2} m |\dot{\mathbf{r}}|^2 \tag{2.2.8}$$

It follows that $Q = \partial \mathcal{L}/\partial \dot{q}_3 = m\dot{z}$ is just the particle's momentum along the z-axis. Another way of seeing it is that the momentum in that direction is the quantity that

is conserved due to translational symmetry of the system along that axis. Consider instead an infinitesimal rotation around an axis defined by the unit vector $\mathbf{n}$, this rotational transformation can be written as

$$\mathbf{r} \to \mathbf{r} + \epsilon(\mathbf{n} \times \mathbf{r}) \tag{2.2.9}$$

If the system is left invariant by this transformation, then the associated Noether charged is written as

$$Q = \frac{\partial \mathcal{L}}{\partial \dot{\mathbf{q}}} \cdot (\mathbf{n} \times \mathbf{r}) = \mathbf{n} \cdot (\mathbf{r} \times \mathbf{p}) = J_n \tag{2.2.10}$$

where J_n is the projection of the angular momentum along the $\mathbf{n}$ direction. Again, if the system is symmetric under any rotation, the entire angular momentum is conserved. Hence, angular momentum can be thought of as the quantity that is conserved when the system is rotationally invariant. This confirms the conservation law for the system outlined in Fig.(2.1). It is clear that the system of the orbiting particle is rotationally invariant around the z-axis and, as such, the z-component of angular momentum is conserved.

Hamiltonian mechanics offers another way of calculating the dynamics of a system and its conserved quantities. It also lends itself more easily to the quantization of the dynamics. There, one defines the Hamiltonian of the system [22]:

$$H = \sum_i (q_i p_i - \mathcal{L}) \tag{2.2.11}$$

where $p_i = \partial \mathcal{L} / \partial \dot{q}_i$ is the generalised momentum associated with the generalised coordinate q_i. The dynamics of the system are then calculated from Hamilton's equations:

$$\frac{dp_i}{dt} = -\frac{\partial H}{\partial q_i} \qquad\qquad \frac{dq_i}{dt} = \frac{\partial H}{\partial p_i} \tag{2.2.12}$$

These equations allow one to also calculate the time evolution of any arbitrary function $f(q_i, p_i)$ defined over the system's conjugate positions and momenta

$$\frac{df}{dt} = \{f, H\} + \frac{\partial f}{\partial t} \tag{2.2.13}$$

where $\{\cdot, \cdot\}$ are the Poisson brackets defined for any two functions as

$$\{f, g\} = \sum_i \left(\frac{\partial f}{\partial q_i} \frac{\partial g}{\partial p_i} - \frac{\partial f}{\partial p_i} \frac{\partial g}{\partial q_i} \right) \tag{2.2.14}$$

If one inserts the definition of the angular momentum into the Poisson's bracket, then one finds that $\{J_i, H\} = 0 \; \forall i$ using the Hamiltonian of the orbiting particle. Since the angular momentum lacks any explicit time dependence, this means that the angular momentum of that system is conserved, i.e. $\mathrm{d}J_i/\mathrm{d}t = 0$.

2.3 Angular momentum in quantum mechanics

The classical angular momentum, as it was defined in the previous section, is a continuous quantity and it is still not clear how a part of it can be separated into distinct components. To see how this is possible, one must use the quantum mechanical description of physical systems. The simplest way to do so is to use the so-called canonical quantization procedure. This is based on the Hamiltonian formulation described above, the functions over the phase space $f(q_i, p_i)$ are "upgraded" to operators $\hat{\mathrm{F}}$ on the Hilbert space of the vectors describing the state of the system. Also, the Poisson brackets are "upgraded" to commutators $\{f, g\} \to (1/i\hbar)[\hat{\mathrm{F}}, \hat{\mathrm{G}}]$, where $[\hat{\mathrm{F}}, \hat{\mathrm{G}}] = \hat{\mathrm{F}}\hat{\mathrm{G}} - \hat{\mathrm{G}}\hat{\mathrm{F}}$ [24].

Consider the quantum mechanical model of the electron in the hydrogen atom. It can be described by the central potential Hamiltonian:

$$\hat{\mathrm{H}} = \frac{\hat{\mathrm{p}}^2}{2\mu} - \frac{e^2}{4\pi\epsilon_0 \hat{\mathrm{r}}} \tag{2.3.1}$$

where ϵ_0 is the vacuum permittivity, e is the electron charge, μ is the reduced mass of the electron-nucleus system, $\hat{\mathbf{r}}$ is the position operator and $\hat{\mathbf{p}}$ is the momentum operator. These two operators obey the canonical commutation relations [25]:

$$[\hat{\mathrm{x}}_i, \hat{\mathrm{x}}_j] = [\hat{\mathrm{p}}_i, \hat{\mathrm{p}}_j] = 0 \qquad\qquad [\hat{\mathrm{x}}_i, \hat{\mathrm{p}}_j] = i\hbar\delta_{i,j} \tag{2.3.2}$$

where δ_{ij} is the standard Kronecker delta. A convenient way to solve the dynamics of this system is to find the eigenvalues of the Hamiltonian in Eq.(2.3.1). This is equivalent to solving the time-independent Schrödinger equation and finding the stationary states of the system. To do so, it is convenient to work in the position representation, where the momentum operator can be expressed as $\hat{\mathbf{p}} = -i\hbar\boldsymbol{\nabla}$ while $\hat{\mathbf{r}} = \mathbf{r}$. Thanks to the spherical symmetry of the problem, the eigenvalue equation becomes

$$-\frac{\hbar^2}{2\mu}\left(\frac{1}{r^2}\frac{\partial}{\partial r}\Big(r^2\frac{\partial\psi}{\partial r}\Big) + \frac{1}{r^2\sin(\theta)}\frac{\partial}{\partial\theta}\Big(\sin(\theta)\frac{\partial\psi}{\partial\theta}\Big) + \frac{1}{r^2\sin^2(\theta)}\frac{\partial^2\psi}{\partial\phi^2}\right) - \frac{e^2}{4\pi\epsilon_0 r}\psi = E\psi \tag{2.3.3}$$

where $\psi(\mathbf{r})$ is the wavefunction of the system in position representation and θ and ϕ are the spherical polar and azimuthal angles respectively. It is clear that this equation can be solved by separation of variables, where the wavefunction is decomposed into a radial and angular part $\psi(\mathbf{r}) = R(r)T(\theta, \phi)$. By analogy to the classical case mentioned above, one defines the quantum angular momentum operator as

$$\hat{\mathbf{L}} = \hat{\mathbf{r}} \times \hat{\mathbf{p}} = -i\hbar(\mathbf{r} \times \nabla) \tag{2.3.4}$$

The reason for denoting this quantum angular momentum operator with the letter L instead of the J used before will be clear later on. Using the commutation relations for the momentum and position operators in Eq.(2.3.2), it can be easily shown that $\hat{\mathbf{L}}$ obeys what are known as the angular momentum commutation relations:

$$[\hat{\mathrm{L}}_i, \hat{\mathrm{L}}_j] = i\hbar\epsilon_{ijk}\hat{\mathrm{L}}_k \qquad\qquad [\hat{\mathrm{L}}^2, \hat{\mathrm{L}}_i] = 0 \tag{2.3.5}$$

where ϵ_{ijk} is the completely anti-symmetric tensor and $\hat{\mathrm{L}}^2 = \hat{\mathrm{L}}_1^2 + \hat{\mathrm{L}}_2^2 + \hat{\mathrm{L}}_3^2$ is the norm of the angular momentum vector operator. Unlike the position and momentum operators, the different projections of angular momentum operators do not commute. As such, it is not possible to simultaneously diagonalize two different projections of the angular momentum operator. This means that, while a state can have well defined linear momenta along different axes by being an eigenstate of two different projection operators, this is not possible for angular momenta operators so a particular state can only have a well defined angular momentum along one axis.

Via separation of variables, these operators allow one to write Eq.(2.3.3) in more compact eigenvalue equations:

$$\frac{\hbar^2}{2\mu}\frac{\partial}{\partial r}\Big(r^2\frac{\partial R}{\partial r}\Big) + \frac{e^2 r}{4\pi\epsilon_0}R + Er^2R = -\frac{\hbar^2 l(l+1)}{2\mu}R \tag{2.3.6}$$

$$\hat{\mathrm{L}}^2 T = \hbar^2 l(l+1)T \tag{2.3.7}$$

The radial equation is then solved using the *ansatz* $R = f(r)/r$, while the angular one is simply the eigenvalue equation for the L^2 operator which has the standard solutions using the spherical harmonics $Y_l^m(\theta, \phi)$, defined as [26]:

$$\hat{\mathrm{L}}^2 Y_l^m(\theta,\phi) = \hbar^2 l(l+1)Y_l^m(\theta,\phi) \qquad \hat{\mathrm{L}}_z Y_l^m(\theta,\phi) = \hbar m_l Y_l^m(\theta,\phi) \tag{2.3.8}$$

where l is the quantum number that represents the *absolute* angular momentum states, while m_l is the quantum number that represents the *projection* of angular momentum on the z-axis which obey the following constraints:

$$l \in \mathbb{Z} \qquad\qquad -l \leq m_l \leq l \tag{2.3.9}$$

This indicates that the stationary solutions to the hydrogen Hamiltonian are also eigenstates of L^2 and L_z and, as such, carry a definite absolute angular momentum and projection along the z-axis. As mentioned before, one cannot simultaneously diagonalize the various projections of angular momentum. As such, these solutions do not carry a definite projection along the x or y-axis.

However, it should be noted that the central potential Hamiltonian is completely spherically symmetric and, hence, the choice of the z-axis is completely arbitrary and has only been made here since it is the standard way to define the polar angle in spherical coordinates. In fact, since $[\hat{L}^2, H] = [\hat{L}_i, H] = 0$, Ehrenfest theorem specifies that the the expectation of absolute angular momentum as well as of all its projections are constants of motion for this system. This is similar to the case of the classical orbiting particle. Putting everything together, the stationary states of this system can then be written as

$$\psi_{nlm}(\mathbf{r}) = \mathcal{N} e^{-2r/na_0} \left(\frac{2r}{na_0}\right)^l L_{n-l-1}^{2l+1}\left(\frac{2r}{na_0}\right) Y_l^m(\theta, \phi) \tag{2.3.10}$$

where $\mathcal{N}$ is a normalization constant, a_0 is the Bohr radius and n is the quantum number representing the energy of the system. However, this is still not the full picture. As the Stern-Gerlach experiment has shown [27], electrons carry an intrinsic angular momentum that is quantized. This angular momentum is omnipresent and independent from any spatial motion. This is called the SAM, which for the electron is represented by the two-dimensional operator

$$\hat{\mathbf{S}} = \frac{\hbar}{2}\hat{\boldsymbol{\sigma}} \tag{2.3.11}$$

where $\boldsymbol{\sigma}$ is a vector operator containing the standard Pauli matrices

$$\hat{\sigma}_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad \hat{\sigma}_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \qquad \hat{\sigma}_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{2.3.12}$$

It is clear from the definition of the constituent Pauli matrices that the $\mathbf{S}$ operators obey the above-mentioned angular momentum commutation relations with similar eigenvalues for $\hat{\mathrm{S}}_i$ and $\hat{\mathrm{S}}^2$. Every physical electron state has a definite SAM with a definite projection along a particular axis with eigenvalues of $\hbar m_s = \pm\hbar/2$. This is because electrons belong to the Fermion class of particles, which all have half-integer spin [28]. Although this treatment of spin may seem ad-hoc as the spin component was merely added to the wavefunction to explain empirical observations, it can actually be derived from fundamental considerations via the relativistic treatment of electrons with the Dirac equation, showing the same results.

In summary, the quantized dynamics of the electron in the hydrogen atom present two different types of angular momentum: the first is the spin angular momentum $\hat{\mathbf{S}}$ that is intrinsic to the particle itself. The second is $\hat{\mathbf{L}}$, which is related to the spatial dynamics of the electron, and is thus appropriately called the orbital angular

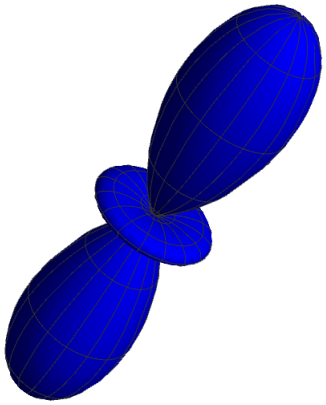

Figure 2.2: 3D surface plot of the $|Y_2^0|^2$ spherical harmonic, clearly showing the azimuthal symmetry and the broken polar symmetry.

momentum (OAM). In addition to these two operators, one can define the operator for *total* angular momentum

$$\hat{\mathbf{J}} = \hat{\mathbf{L}} + \hat{\mathbf{S}} \tag{2.3.13}$$

where $\hat{\mathbf{J}}$ operators also obey the angular momentum commutation relations. One can also similarly define $\hat{\mathrm{J}}^2$ operators with $[\hat{\mathrm{J}}^2, \hat{\mathrm{J}}_i] = 0$ and similar eigenvalue equations

$$\hat{\mathrm{J}}^2 \left|j, m_j\right\rangle = \hbar^2 j(j+1) \left|j, m_j\right\rangle \qquad \hat{\mathrm{J}}_z \left|j, m_j\right\rangle = \hbar m_j \left|j, m_j\right\rangle \tag{2.3.14}$$

Under certain conditions, it is more convenient to describe the states in terms of the total angular momentum, since that is the quantity that is conserved. For example, since the electron is a charged particle, each angular momentum is associated with a magnetic moment. This means that the magnetic moment due to the spin can couple to the one due to the OAM. This spin-orbit coupling is represented in the Hamiltonian via a term proportional to $\hat{\mathbf{L}} \cdot \hat{\mathbf{S}}$. In fact, if one wants to calculate the fine-structure of the energy levels of atoms, one needs to account for spin-orbit coupling among other relativistic corrections to the central potential Hamiltonian. Due to the spin-orbit coupling term, the projections of the two angular momenta no longer commute with the Hamiltonian and are no longer independently conserved, thus preventing one from writing the stationary states of the system as eigenstates of either of them. However, it is possible to index the stationary states using the $\hat{\mathrm{J}}^2$ and $\hat{\mathrm{J}}_z$ operators.

Although for bounds states, it is relatively straightforward to see the origin of the split of the angular momentum into spin and orbital components, it remains unclear how it would be possible to define this orbital momentum for freely moving particles, let alone for electromagnetic waves. To see how this is done, consider first the classical description of electromagnetic waves.

2.4 Electromagnetic waves

In order to show how an electromagnetic wave can carry angular momentum, the first step is to look at the joint classical dynamics of a charged particle in an generic electromagnetic field. These dynamics are governed by the well known Maxwell's equations for the electric $\mathbf{E}$ and magnetic $\mathbf{B}$ fields [29]

$$\boldsymbol{\nabla} \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \tag{2.4.1}$$

$$\boldsymbol{\nabla} \cdot \mathbf{B} = 0 \tag{2.4.2}$$

$$\boldsymbol{\nabla} \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{2.4.3}$$

$$\boldsymbol{\nabla} \times \mathbf{B} = \mu_0 \Big(\mathbf{j} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \Big) \tag{2.4.4}$$

as well as the Lorentz force equation

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \tag{2.4.5}$$

where μ_0 is the vacuum permeability, q is the charge of the particle and ρ and $\mathbf{j}$ are the charge and current densities respectively defined as

$$\rho = q\delta^{(3)}(\mathbf{r} - \mathbf{r}_0(t)) \qquad\qquad \mathbf{j} = q\mathbf{v}\delta^{(3)}(\mathbf{r} - \mathbf{r}_0(t)) \tag{2.4.6}$$

with $\delta^{(3)}(\mathbf{r})$ being the 3-dimensional Dirac delta function and r_0 the position of the particle at a specific time. Using Maxwell's equations, the charge continuity equation is then immediately derived, which relates the charge and current densities, and expresses the local conservation of charge as

$$\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot \mathbf{j} = 0 \tag{2.4.7}$$

The dynamics of the system are completely defined by the above-mentioned equations. The combined particle and field system presents the following different con-

stants of motion (one scalar and two vectors) [30]:

$$E = \frac{1}{2}mv^2 + \frac{\epsilon_0}{2}\iiint_V (E^2 + c^2B^2)d^3r \tag{2.4.8}$$

$$\mathbf{P} = m\mathbf{v} + \epsilon_0 \iiint_V (\mathbf{E} \times \mathbf{B})d^3r \tag{2.4.9}$$

$$\mathbf{J} = \mathbf{r} \times (m\mathbf{v}) + \epsilon_0 \iiint_V \mathbf{r} \times (\mathbf{E} \times \mathbf{B})d^3r \tag{2.4.10}$$

To show that these are, in fact, constants of motion, it suffices to take their time derivative and use Maxwell's equations, Lorentz equation and the definition for ρ and $\mathbf{j}$ to show that the right-hand side cancels out. The interpretation of these constants is quite straightforward: they are the total energy, linear and angular momentum of the system. It is clear that each constant is composed of two contributions: the first ones are just the classical constants of motion of a particle while the other ones can be attributed to the electromagnetic field itself.

This then allows one to define for the electromagnetic field an energy density as well as linear and angular momentum densities:

$$\mathcal{E} = \frac{\epsilon_0}{2}(E^2 + c^2B^2) \tag{2.4.11}$$

$$\boldsymbol{\mathcal{P}} = \epsilon_0(\mathbf{E} \times \mathbf{B}) \tag{2.4.12}$$

$$\boldsymbol{\mathcal{J}} = \epsilon_0\mathbf{r} \times (\mathbf{E} \times \mathbf{B}) \tag{2.4.13}$$

I now begin to address the question of how to define the angular momentum of the electromagnetic waves. To better isolate the problem, consider Maxwell's equations in the vacuum devoid of any sources of charge and current. This leads naturally to the vacuum electromagnetic wave equations

$$\left(\frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \nabla^2\right)\mathbf{E} = 0 \tag{2.4.14}$$

$$\left(\frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \nabla^2\right)\mathbf{B} = 0 \tag{2.4.15}$$

where c is the speed of light. The relationship between $\mathbf{E}$ and $\mathbf{B}$ is of course still governed by Maxwell's equations.

The simplest possible solution to these wave equations is the plane wave modes, which can be written in complex notation as

$$\mathbf{E} = \mathcal{E}e^{i\mathbf{k}\cdot\mathbf{r}-\omega t}\hat{\mathbf{n}} \tag{2.4.16}$$

$$\mathbf{B} = \frac{\hat{\mathbf{k}}}{c} \times \mathbf{E} \tag{2.4.17}$$

where $\hat{\mathbf{n}}$ is the polarization unit vector of the electric field, $\mathcal{E}$ is the amplitude, $\mathbf{k}$ is the wave-vector of the wave, $\hat{\mathbf{k}} = \mathbf{k}/|\mathbf{k}|$ is the unit vector for the direction of the wave-vector and ω is the frequency. The wave-vector and the frequency obey the following conditions

$$\mathbf{k} \cdot \hat{\mathbf{n}} = 0 \qquad\qquad \omega = c|\mathbf{k}| \tag{2.4.18}$$

If one was to take the simplest case where the coordinates are aligned such that the wave is propagating along the z-axis $\hat{\mathbf{k}} = \hat{\mathbf{e}}_z$ and the electric field is polarized along the x-axis $\hat{\mathbf{n}} = \hat{\mathbf{e}}_x$, the time-averaged linear momentum density of this wave is then readily calculated

$$\langle \boldsymbol{\mathcal{P}}_{\text{plane}} \rangle_t = \epsilon_0 \Re(\mathbf{E}^* \times \mathbf{B}) = \frac{\epsilon_0 |\mathcal{E}|^2}{c} \hat{\mathbf{e}}_z \tag{2.4.19}$$

This is intuitive since the plane wave is propagating solely in the z-direction and hence it would make sense for it to have its momentum aligned along the direction of its wave-vector. To calculate the time-averaged angular momentum density, it is more convenient to use the cylindrical coordinate system (ρ, ϕ, z)

$$\langle \boldsymbol{\mathcal{J}}_{\text{plane}} \rangle_t = (\rho \hat{\mathbf{e}}_\rho + z \hat{\mathbf{e}}_z) \times \langle \mathcal{P}_{\text{plane}} \rangle_t = -\frac{\epsilon_0 |\mathcal{E}|^2}{c} \hat{\mathbf{e}}_\phi \tag{2.4.20}$$

Once this is integrated over the transverse surface, it vanishes since its amplitude is azimuthaly symmetric and the unit polar $\hat{\mathbf{e}}_\phi$ vector would integrate to 0. Clearly, plane waves do not carry any net angular momentum. For that, one needs to look at more exotic beam modes. Going back to the wave equation in (2.4.14), instead of using the plane wave expansion, consider solutions of this form

$$\mathbf{E} = \Big(u(\mathbf{r}) \hat{\mathbf{e}}_x + \frac{i}{k} \frac{\partial u}{\partial x} \hat{\mathbf{e}}_z \Big) e^{i(kz - \omega t)} \tag{2.4.21}$$

$$\mathbf{B} = \frac{1}{c} \Big(u(\mathbf{r}) \hat{\mathbf{e}}_y + \frac{i}{k} \frac{\partial u}{\partial y} \hat{\mathbf{e}}_z \Big) e^{i(kz - \omega t)} \tag{2.4.22}$$

where $u(\mathbf{r})$ is a function which describes the envelope of the wave and $\omega = ck$ is the frequency of the carrier wave. This description is valid in the so-called paraxial approximation where the wave is propagating along a direction almost-parallel to the z-axis [31]. In this approximation, the products of the first derivatives are negligible and $\partial_z u \ll k u(\mathbf{r})$. This approximation is also called the slowly-varying envelope approximation because it considers that the envelope function has small spatial derivatives compared to the spatial oscillation of the carrier wave (represented by k). Using these approximations, the time-averaged momentum density for this wave is

$$\langle \boldsymbol{\mathcal{P}} \rangle_t = \frac{\epsilon_0}{c} \Big[k |u(\mathbf{r})|^2 \hat{\mathbf{e}}_z + \frac{i}{2k} (u \boldsymbol{\nabla}_\perp u^* - u^* \boldsymbol{\nabla}_\perp u) \Big] \tag{2.4.23}$$

where $\boldsymbol{\nabla}_\perp = \partial_x \hat{\mathbf{e}}_x + \partial_y \hat{\mathbf{e}}_y$ is the transverse gradient. This indicates that depending on the behaviour of $u(\mathbf{r})$, this class of solutions may carry a non-zero angular momentum. If the expression for the electric field in Eq.(2.4.33) is plugged into the vacuum wave equation, by again using the paraxial approximation, the envelope function will obey the following equation:

$$\nabla_\perp^2 u - 2ik\frac{\partial u}{\partial z} = 0 \tag{2.4.24}$$

The simplest solution to this equation is the Gaussian envelope, which leads one to the well-known Gaussian beam profile, which, in turn, is the most common mode output by lasers [31]. However, for the purposes of investigating the angular momentum that can be carried by light, it is more interesting to solve the equation in cylindrical coordinates where $\boldsymbol{\nabla} = \partial_\rho \hat{\mathbf{e}}_\rho + (1/\rho)\partial_\phi \hat{\mathbf{e}}_\phi + \partial_z \hat{\mathbf{e}}_z$. This leads to a class of solutions called the Laguerre-Gaussian (LG) modes

$$\begin{aligned} u_l^p(\mathbf{r},t) = \sqrt{\frac{2p!}{\pi(p+|l|)!}}\frac{w_0}{w(z)}\left(\frac{\sqrt{2}\rho}{w(z)}\right)^{|l|} L_p^{|l|}\left(\frac{2\rho^2}{w^2(z)}\right) \\ \times \exp\left(-\frac{\rho^2}{w^2(z)} + i\Big(kz - \omega t + l\phi + \frac{k\rho^2}{2R(z)} + \psi_p^l(z)\Big)\right) \end{aligned} \tag{2.4.25}$$

where (ρ, θ, z) are the standard cylindrical coordinates, l is the azimuthal mode number, p is the radial mode number, w_0 is the beam waist at focus. ω and k are the frequency and longitudinal wave-vector of the wave with $\omega_0 \approx ck$ since, in the paraxial approximation, the transverse wave-vector of the beam is assumed to be very small compared to the longitudinal one.

$L_p^{|l|}$ is the generalised Laguerre polynomial, which is the solution of the following differential equation

$$xf''(x) + (l+1-x)f'(x) + pf(x) = 0 \tag{2.4.26}$$

The rest of the functions which govern the behaviour of LG mode are the standard Gaussian parameters for the beam waist $w(z)$, radius of curvature $R(z)$ and Gouy phase shift $\psi_p^l(z)$:

$$w(z) = w_0\sqrt{1 + \Big(\frac{2z}{kw_0^2}\Big)^2} \tag{2.4.27}$$

$$R(z) = z\left(1 + \Big(\frac{kw_0^2}{2z}\Big)^2\right) \tag{2.4.28}$$

$$\psi_p^l(z) = (|l| + 2p + 1)\arctan\Big(\frac{2z}{kw_0^2}\Big) \tag{2.4.29}$$

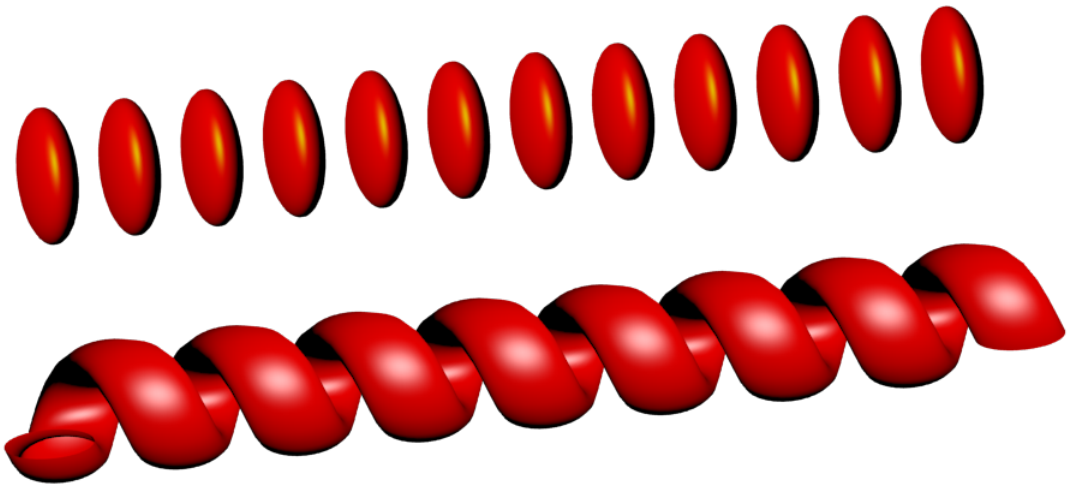

Figure 2.3: 3D surface plot of the surfaces of constant phase for a standard Gaussian beam (upper) and for a Laguerre-Gaussian beam with $l = 1, p = 0$ (lower). This clearly shows that for the standard Gaussian the phase front is planar while it is helical for the LG modes, illustrating why these modes have acquired the name "twisted light".

It is clear that when $l = 0$ and $p = 0$, $u_p^l(\mathbf{r})$ simplifies to the standard Gaussian envelope, which is technically the lowest-order LG mode.

Inserting the expression for the LG solution into Eq.(2.4.23), the time-averaged linear momentum density of these modes is then

$$\langle \boldsymbol{\mathcal{P}} \rangle_t = \frac{\epsilon_0}{c}\Big(\frac{\rho}{R(z)}\hat{\mathbf{e}}_\rho + \frac{l}{k\rho}\hat{\mathbf{e}}_\phi + \hat{\mathbf{e}}_z\Big)|u_p^l(\mathbf{r})|^2 \tag{2.4.30}$$

This clearly shows that, as well as the longitudinal component, the LG mode, has a transverse component in its linear momentum density. However, due to the azimuthal symmetry of the $\hat{\mathbf{e}}_\phi$ and $\hat{\mathbf{e}}_\rho$ components, they will integrate to 0 again due to the oscillation of the polar unit vectors. Therefore, the total linear momentum still integrates out to be directed along the z-axis. This can also be seen from the shape of the phase front shown in Fig.(2.3). The LG modes have a helical phase front indicating that their wave-vectors have a transverse component, which are continuously rotating around an average constant $\hat{\mathbf{e}}_z$ component.

As for the angular momentum density, it can again be calculated from the linear

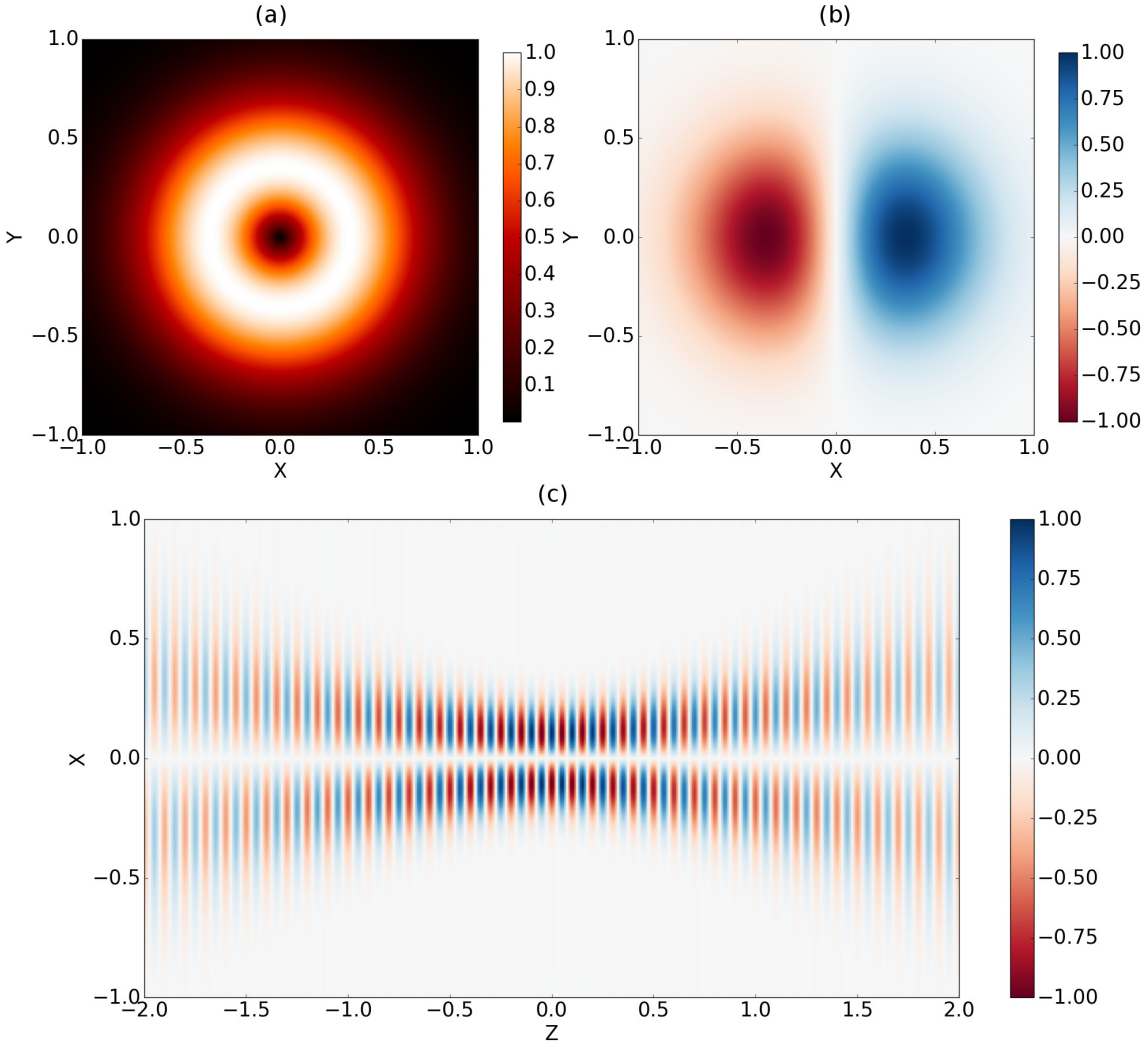

Figure 2.4: Simulated results for the LG envelope solution with $l = 1, p = 0$: (a) colour plot of the time-averaged transverse intensity at a particular transverse slice clearly showing the characteristic "doughnut" intensity profile that is due to the phase singularity at the origin; (b) colour plot of the instantaneous amplitude of $u_1^0(\mathbf{r})$ at a particular transverse slice. It shows the azimuthal variation of the amplitude from a positive maximum to a negative minimum. This profile rotates as the beam propagates and accomplish a full turn over a wavelength; (c) longitudinal slice showing the instantaneous amplitude of the envelope as the wave travels along the z-axis. It shows a similar focusing behaviour as the standard Gaussian mode, except of course for the singularity at the centre.

momentum using Eq.(2.4.20)

$$\langle \boldsymbol{\mathcal{J}} \rangle_t = \frac{1}{c}\Big(- \frac{zl}{k\rho}\hat{\mathbf{e}}_\rho + \rho\Big(\frac{z}{R(z)} - 1\Big)\hat{\mathbf{e}}_\phi + \frac{l}{k}\hat{\mathbf{e}}_z\Big)|u(\mathbf{r})|^2 \tag{2.4.31}$$

As before, the $\hat{\mathbf{e}}_\rho$ and $\hat{\mathbf{e}}_\phi$ components are azimuthally symmetric and thus would vanish when the total angular momentum density is integrated. However, this time, the angular momentum density has a $\hat{\mathbf{e}}_z$ component which does not vanish in the integration. This means that the LG mode carries a net angular momentum equal to

$$\mathbf{J}_{\mathrm{LG}} = \epsilon_0 \frac{l}{\omega}\hat{\mathbf{e}}_z \tag{2.4.32}$$

which is directed along the z-axis, which, in turn, is quantized by the azimuthal mode number l of the envelope.

Now this shows how modes of electromagnetic radiation can carry angular momentum; however, it is not evident how this can be split into separate spin and orbital components. To see this, instead of using linearly polarized LG modes as in Eq.(2.4.33), consider the circularly polarized modes, where the electric and magnetic fields can be written as

$$\mathbf{E}^\sigma_{l,p} = \frac{1}{\sqrt{2}}\Big(u^l_p(\mathbf{r})(\hat{\mathbf{e}}_x + i\sigma\hat{\mathbf{e}}_y) + \frac{i}{k}\Big(\frac{\partial u^l_p}{\partial x} + i\sigma\frac{\partial u^l_p}{\partial y}\Big)\hat{\mathbf{e}}_z\Big)e^{i(kz-\omega t)} \tag{2.4.33}$$

$$\mathbf{B}^\sigma_{l,p} = \frac{1}{c\sqrt{2}}\Big(u^l_p(\mathbf{r})(\hat{\mathbf{e}}_y - i\sigma\hat{\mathbf{e}}_x) + \frac{i}{k}\Big(\frac{\partial u^l_p}{\partial y} - i\sigma\frac{\partial u^l_p}{\partial x}\Big)\hat{\mathbf{e}}_z\Big)e^{i(kz-\omega t)} \tag{2.4.34}$$

where $\sigma = \pm 1, 0$ is a parameter determining whether the wave is left $(+1)$ or right (-1) circularly polarized or if it just linearly polarized.

Whereas in linearly polarized waves, the vector fields are oriented in a fixed direction and circularly polarized waves have vector fields which rotate while the wave is propagating as can be seen in Fig.(2.5). By plugging the expression for the circularly polarized LG fields in the equation for the angular momentum densities in Eq.(2.4.13), it can be shown that these modes carry a net angular momentum equal to [1]

$$\mathbf{J} = \frac{\epsilon_0(l+\sigma)}{\omega}\hat{\mathbf{e}}_z \tag{2.4.35}$$

It is now clear that the total angular momentum of these modes can be split into two contributions: the first is the contribution arising from the polarization of a mode represented by the σ parameter. This is solely due to the vector nature of electromagnetic waves and is thus identified by their intrinsic spin. The other is the contribution coming from the spatial distribution of the fields represented by

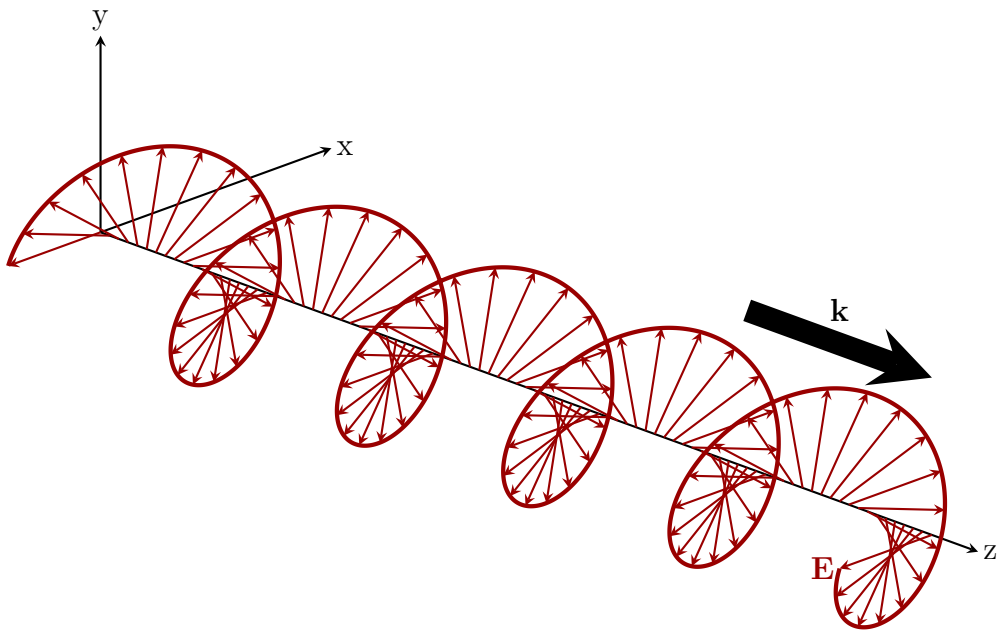

Figure 2.5: Diagram showing how the electric field vector of a circularly polarized light rotates drawing a helix as the wave propagates along the z-direction. It should be noted that, as opposed to the LG envelope, this does not provide any information on the shape of wavefront of the beam. Waves with planar wavefronts can still have circular polarization.

the azimuthal mode number l of the LG modes. This is then called the orbital angular momentum of the wave. Since the LG envelope function is normalized to that $\iiint d^3r|u(\mathbf{r})|^2 = 1$, the energy contained in the beam is simply $E = \epsilon_0$. The ratio between the angular momentum of the beam and its energy $(l + \sigma)/\omega$ takes an analogous form to the typical ratio of the photon's spin to its energy. This further confirms our interpretation of $\mathbf{J}$ as the angular momentum of light.

Calling the property caused by the LG envelope an orbital angular momentum is not just a mathematical formality. It does in fact act as a type of angular momentum, separate from that arising from the polarization of light. This can be seen in how it affects the interaction between light and matter. Very soon after LG modes were first characterised, researchers were able to directly observe the rotation of microscopic particles when they absorb LG radiation. This was done by monitoring the behaviour of micrometre sized particles suspended in fluid before and after they interact with lasers that have been converted to LG modes [32]. H. He and colleagues found that the particles start rotating in the same direction after they interact with the beam. They also confirmed that OAM is the origin of this induced rotation when they flipped the sign of the azimuthal phase dependence $l \to -l$; hence, the sign of the OAM, and

the particles' direction of rotation would be inverted accordingly.

That being said, it is very difficult to verify that the OAM of light is quantized using these types of experiments due to the strong thermal fluctuations inherent to any warm particle. Efforts were then put to try to use a superfluid or a Bose-Einstein condensate (BEC) to evidence this quantization. In 2006, this was achieved using a stimulated Raman scattering scheme on a BEC of sodium atoms [33]. The condensate coherently interacted with two counter-propagating beams, one LG_0^1 and the other a regular Gaussian. It absorbed the photons from the LG mode and emitted ones in the Gaussian mode via stimulated emission, thereby acquiring the difference in OAM between the two modes as well as the linear momentum. This was proportional to the number of photons involved in the process (1 or 2-photon Raman scattering). The experiments confirmed that the BEC acquired a quantized OAM since it gained either $\hbar$ or $2\hbar$ depending on the number of photons in the interaction. Since the BEC is a macroscopic object with quantum mechanical properties, it was actually possible to see a hole at the center of BEC that was due to the singularity caused by the OAM. The exact OAM state of the OAM was measured by interfering the BEC with one which had been interacting with two counter-propagating Gaussian beams and hence acquired no angular momentum.

These experiments confirm that the OAM of light is truly a form of angular momentum, and that it is distinct from the intrinsic spin. Since OAM is related to the spatial distribution of the electromagnetic wave, it is natural for it to interact with the spatial rotational degrees of freedom of matter. However, this is not the full story. For a long time, researchers in atomic spectroscopy had known that the polarization (spin) of the incoming light can affect the type of electronic energy transition that can happen in the absorbing atom. Electromagnetic waves can drive transitions between different energy states if the wavefunctions of both states overlap with the multipole expansion of the field of the wave and all the relevant conservation rules are obeyed (energy, momentum). The most common type of transitions are the dipole ones. There, if the magnetic moment due to the orbital angular momentum of the electron has a well-defined axis, for example due to an external constant magnetic field, polarized light will induce transitions between energy states where $\Delta m_j = \sigma$. Here, m_j refers to the quantum number specifying the projection of the electron's total angular momentum along the z-axis as it was defined in Eq.(2.3.14), while σ refers to the polarization of the incoming light.

Recently, experiments have proven that the OAM of light can also modify electronic transitions rules. Using precise measurements of the atomic transitions in

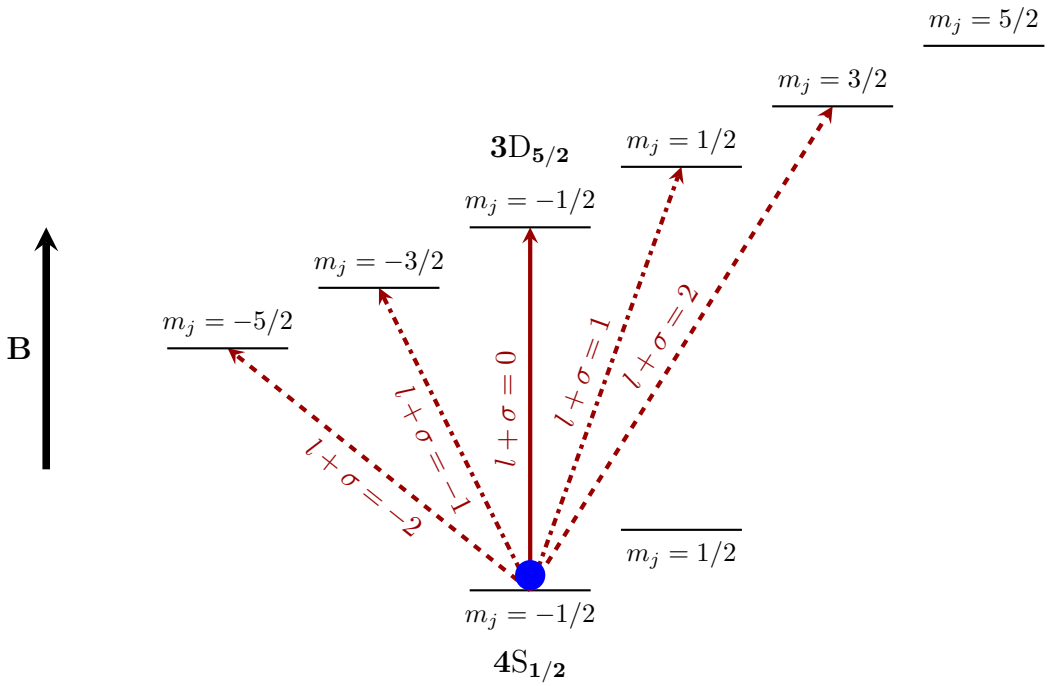

Figure 2.6: Energy diagram showing the quadrupole energy transitions $4S_{1/2} \rightarrow 3D_{5/2}$ in a calcium ion. The different states corresponding to the z projection of the total angular momentum m_j are split due the Zeeman effect caused by the constant external magnetic field **B**. This breaks the symmetry of the system and makes the angular momentum projection states non-degenerate. The conservation of the total angular momentum in the stimulated absorption means that the sum of the incoming light's spin and orbital angular momentum must match the difference in the m_j number between the initial and final states.

calcium ions [5], it was determined that the OAM of light has an additive effect on the change in the m_j quantum number where $\Delta m_j = l + \sigma$. Due to the exotic shape of the LG modes, which have a singularity at the centre with null intensity as seen in Fig.(2.4), the relevant transitions are not the dipole transitions but the quadrupole ones. These are dependent not on the oscillations of the field but on the oscillations of its gradient and the LG have strong field gradients at this singularity. This means that not only does the OAM of light affect the external rotation of an atom, it also couples to its internal degrees of freedom in the angular momentum of its electronic states.

It should be noted that the azimuthal mode number l should not be confused with the quantum number for the absolute orbital angular momentum from the previous section. The latter quantizes the eigenvalues of $\hat{L}^2$ while the former is actually more analogous to the m quantum number, which quantizes the projection of the angular

momentum operator on the z-axis $\hat{\mathbf{L}}_z$, as we have seen with the case of the atomic transitions.

I have used this notation because it is the standard one used throughout the literature on the OAM of light and the standard notation for the LG mode decomposition. As most of the physical models in this chapter and in the following ones use a classical or semi-classical framework, the distinction between the l and m is not important. Where a quantum mechanical approach is used, care must be taken to avoid any confusion as to which angular momentum operator is being referred to. Additionally, although the LG solutions are the best-known OAM-carrying modes, they are not the only ones. In fact, any mode with an azimuthal dependence in its phase can have non-zero OAM since that means that its wave-vector has an azimuthal component and, in turn, so will its linear momentum density, as can be seen from Eq.(2.4.23). The ubiquity of the LG modes in the literature for OAM is partly due to them being the subject of L. Allen's seminal 1992 paper on the angular momentum of light, as well as due to the fact that they are relatively easy to generate and manipulate, as will be seen in subsequent chapters.

So far, only the classical description of the electromagnetic field has been considered. The form of the AM of light arose simply from considerations of conserved quantities in the interaction between light and matter. Furthermore, its split into a spin and an orbital part has been made for a particular configuration of the electromagnetic field and it remains unclear if there is any fundamental reason for it, especially since they both seem to interact with matter in an identical way. Both spin and orbital angular momenta are transferred to macroscopic particles causing them to rotate and they both affect the internal electronic transitions in the same way by modifying Δm. Finally, a classical description does not show how this orbital angular momentum is quantized or whether it can be defined for a single photon. To that end, a more fundamental approach needs to be used to describe angular momentum.

2.5 Angular momentum in field theory

Much of the current debate concerning the angular momentum of light revolved around whether or not the split into a spin and orbital part has a fundamental reason and whether it is possible to do it in a consistent and a gauge invariant way [7]. Some questions have also been raised about the quantization of the spin and orbital angular momenta of light and whether the resultant quantum mechanical operators obey

the typical commutation relations for the angular momentum operators. To address this, angular momentum will be defined from first principles in the most generic way without making reference to a particular physical field. This definition follows from the typical way generators of symmetry are defined in textbooks on applied group theory in physics [34]. As much of this section will deal with relativistic quantum field theory, the l and m_l quantum numbers will refer to the absolute orbital angular momentum and its projection respectively. Also, to make the notation easier, natural units, as they are defined in the units and conventions section, will be used.

Let us move away from physical systems for a bit and consider mathematical groups of symmetry transformations and their representations. These are defined as

Definition 2.5.1. A group $(G, \cdot)$ is a finite or infinite set G of elements equipped with a binary operation $\cdot$ called the group law, which together satisfy four essential properties:

- **Closure**: If $g_1, g_2 \in G$ then $g_1 \cdot g_2 \in G$
- **Associativity**: For all $g_1, g_2, g_3 \in G$, $g_1 \cdot (g_2 \cdot g_3) = (g_1 \cdot g_2) \cdot g_3$
- **Identity element**: There exists a unique element $e \in G$ called the identity such that $\forall g \in G$, $e \cdot g = g \cdot e = g$
- **Inverse element**; For all $g \in G$ there exists a unique inverse element $g^{-1} \in G$ such that $g \cdot g^{-1} = g^{-1} \cdot g = e$.

Definition 2.5.2. A representation is a mapping D that maps each element in a group $g \in G$ to a linear operator $D(g) : V \to V$ acting on the vector space V such that:

- $D(e) = \mathbf{I}_V$, where $\mathbf{I}_V$ is the identity operator of the vector space
- $D(g_1)D(g_2) = D(g_1 \cdot g_2)$, so the mapping preserves the group's structure and maps the group law to the operator multiplication law of the vector space.

Care should be taken not to confuse the dimensions of the vector space on which the operators of a particular representation act and the size of the set of elements that form the group. In fact, the same group can have multiple different representations with different dimensions. For example, a group could have both a matrix representation acting on a vector space of finite dimensions and another functional representation acting on an infinitely dimensional space of functions.

As I am considering groups of symmetry transformations, I will now focus my attention onto Lie groups. A Lie group is a group where each element $g(\alpha^a)$ depends in a continuous and infinitely differentiable way on a set of continuous parameters $\alpha^a, a \in \{1, .., N\}$ [35]. It is convenient to choose the origin of these parameters such that $g(0) = e$. The fact that the elements vary continuously with the parameters means that the group's elements have a notion of "closeness" in the space spanned by the group via the closeness of their parameters. Of course, since the group itself is parameterized by α^a, then so will any of its representations. Thanks to the differentiability of the parameterization, the representation can be Taylor expanded around $\alpha^a = 0$ to give

$$D(\alpha^a) \approx \mathbf{I}_V + i\alpha_a X^a + O((\alpha^a)^2) \tag{2.5.1}$$

where

$$X^a = -i\frac{\partial D(\alpha^a)}{\partial \alpha_a}\Big|_{\alpha_a=0} \tag{2.5.2}$$

are called the generators of the group in the particular representation D. No matter the representation, there are of course as many generators as there are parameters for the group.

There exists a convenient parameterization for the elements of the group "far" from the identity. This is called the *exponential parameterization* [34] where no matter how far away from 0 the parameters are, one can write any element in the group as

$$D(\alpha^a) = \lim_{k\to\infty}\left(\mathbf{I}_V + i\frac{\alpha_a X^a}{k}\right)^k = e^{i\alpha_a X^a} \tag{2.5.3}$$

This is possible because, as the representation preserves the group's structure, when an element "close" to the identity, which can be written as Eq.(2.5.1), is successively multiplied it will generate another element of the group. It should be noted that, unlike in Eq.(2.5.1), the α_a are finite and can be arbitrarily large. The limit taking the denominator k to infinity will then arrange it so that the expression in the parentheses of Eq.(2.5.3) matches the Taylor expansion of group elements "near" the identity. This way of defining the group's elements is convenient because it allows one to express the elements of the group in terms of its generators and further allows one to define the Lie algebra of the group.

Consider a particular representation of two elements $D(g_1) = e^{i\alpha_a X^a}$ and $D(g_2) = e^{i\beta_a X^a}$. By definition, one must have $D(g_1)D(g_2) = D(g_1 \cdot g_2)$ so that the presentation would preserve the structure of the group. Hence it must be the case that

$$e^{i\alpha_a X^a} e^{i\beta_a X^a} = e^{i\delta_a X^a} \tag{2.5.4}$$

Since linear operators do not necessarily commute, it is not immediately obvious how this equation can be satisfied. To that end, one uses the Baker-Campbell-Hausdorff formula to find a solution to it [36]. Keeping only the terms up to the second order in the parameters, it can be shown that in order to satisfy the above mentioned condition one must have

$$i\delta_c X^c = i(\alpha_c + \beta_c)X^c - \frac{1}{2}[\alpha_a X^a, \beta_b X^b] + O(\alpha^2\beta, \alpha\beta^2) \tag{2.5.5}$$

$$\alpha_a\beta_b[X^a, X^b] = 2i(\alpha_c + \beta_c - \delta_c)X^c + O(\alpha^2\beta, \alpha\beta^2) \tag{2.5.6}$$

Since Eq.(2.5.6) must be true for all α_a and β_b, then it must be possible to factor it out of the equation. This then allows one to write the condition for Eq.(2.5.4) as what is called the Lie algebra of the group:

$$[X^a, X^b] = if_c^{ab}X^c \tag{2.5.7}$$

Where f_c^{ab} are called the structure constants. It should be noted that although the particular form of the generators of the group depends on the representation that is being used, the Lie algebra and, in particular, the structure constants are completely independent from it. This is because if the structure constants depended on the representation, then which element of the group $D(g_1)D(g_2)$ maps to, would also depend on the representation [37]. This violates the definition outlined above since, in this case, the representation fails to preserve the group law. This means that the Lie algebra is defined by the structure of the group itself regardless of the representation under consideration.

Now that the basic structure of groups has been laid out, consider one of the most important symmetry groups in physics: the Lorentz group. This group is of essential importance to relativistic classical and quantum mechanics. The laws that describe the electromagnetic strong and weak forces all exhibit Lorentz symmetry. Hence, understanding the structure of this group allows one to understand a lot of the properties of physical fields. The Lorentz group is defined as the group of spacetime transformations which leave the spacetime interval, defined as $\eta_{\mu\nu}x^\mu x^\nu = t^2 - x^2 - y^2 - z^2$, invariant thus leaving the speed of light constant during the transformation [38]. More generally, Lorentz transformations leave invariant the Minkowski norm of 4-vectors. To see how this works, consider the linear space of 4-vectors V^μ equipped with the Minkowski inner product where the norm of the 4-vector is defined by

$$V = \eta_{\mu\nu}V^\mu V^\nu = V_\mu V^\mu \tag{2.5.8}$$

where $\eta_{\mu\nu}$ is the Minkowski metric. A Lorentz scalar can then be made via the Minkowski inner product of any two 4-vectors. Consider then a generic element of the Lorentz group, the representation of this element on the above-mentioned 4-vector space is written as a 4×4 tensor $\Lambda^{\mu\nu}$. The Lorentz invariance condition tells us that the norm of the transformed 4-vector

$$V'^{\mu} = \Lambda^{\mu}{}_{\nu} V^{\nu} \tag{2.5.9}$$

should remain unchanged. This invariance condition can also be written as

$$\eta_{\rho\sigma} = \eta_{\mu\nu} \Lambda^{\mu}{}_{\rho} \Lambda^{\nu}{}_{\sigma} \tag{2.5.10}$$

which indicates that $\det(\mathbf{\Lambda}) = \pm 1$. Factoring out all discrete transformations such as the parity, spatial and temporal inversion transformations (transformations which flip the sign of an odd number of spatial components, or flip the sign of the time component), I need only consider the subgroup of Lorentz transformations called the $SO(3,1)$.

All the symmetry transformations that are left can be decomposed into the three spatial rotations around each spatial axis and the three Lorentz boosts along each one of them. Therefore, this representation, and by extension the group, depends on 6 continuous parameters: three angles and three velocities. This means that $SO(3,1)$ is a Lie group. It is convenient to instead use the rapidity ξ^i defined as $v^i = \tanh(\xi^i)$ to parameterise the boost. That way, a Lorentz boost can be thought of as hyperbolic rotation of an angle ξ^i in the (t, x^i) plane. This number of parameters is confirmed by looking at a Lorentz transformation very close to the identity $\Lambda^{\mu\nu} = \delta^{\mu\nu} + \omega^{\mu\nu}$. Then, the condition in Eq.(2.5.10) tells us that

$$\omega^{\mu\nu} = -\omega^{\nu\mu} \tag{2.5.11}$$

meaning that the $\omega^{\mu\nu}$ is anti-symmetric and an anti-symmetric 4×4 tensor can only depend on 6 parameters. This then allows one to write a generic element of the Lorentz group as a function of only six generators. Following the standard way or representing the Lorentz group found in quantum field theory textbooks [25, 38], it is convenient to maintain the notation with an anti-symmetric tensor of parameters $\omega^{\mu\nu}$. In that case, any $g \in SO(3,1)$ can be written as

$$D(g(\omega^{\mu\nu})) = \exp(-\frac{i}{2} \omega_{\mu\nu} S^{\mu\nu}) \tag{2.5.12}$$

where $(1/2)$ prevents double counting. Care should be taken to not confuse the indices (μ, ν) counting over the parameters and generators with the indices of whatever

representation is being used. Staying in the 4-vector representation, one rewrites an infinitesimal Lorentz transform as

$$V'^{\rho} = \Lambda^{\rho}{}_{\sigma} V^{\sigma} \approx (\delta^{\rho}_{\sigma} - \frac{i}{2}\omega_{\mu\nu}(S^{\mu\nu})^{\rho}{}_{\sigma})V^{\sigma} \tag{2.5.13}$$

To match the original definition of the Lorentz transform (Eq.(2.5.9)) and seeing that $\omega_{\mu\nu}$ is anti-symmetric, the generators in the 4-vector representation are then written explicitly as

$$(S^{\mu\nu})^{\rho}{}_{\sigma} = i(\eta^{\mu\rho}\delta^{\nu}_{\sigma} - \eta^{\nu\rho}\delta^{\mu}_{\sigma}) \tag{2.5.14}$$

leading one to the Lie algebra for $SO(3,1)$ [25]:

$$[S^{\mu\nu}, S^{\rho\sigma}] = i(\eta^{\nu\rho}S^{\mu\sigma} - \eta^{\mu\rho}S^{\nu\sigma} - \eta^{\nu\sigma}S^{\mu\rho} + \eta^{\mu\sigma}S^{\nu\rho}) \tag{2.5.15}$$

Although this algebra was derived using the explicit form of the generators in the 4-vector representation, it must hold for any representation of the $SO(3,1)$ group. As seen before, the transformations of this group are composed of three transformations affecting only the spatial components and three transformations mixing the time component with a spatial one. For that reason, it is convenient to redefine the generators into two types of operators $S^i = (1/2)\epsilon^{ijk}S^{jk}$ and $K^i = S^{i0}$ which act on separate subspaces. This leads to a more readable form for the Lie algebras

$$[S^i, S^j] = i\epsilon^{ijk}S^k \tag{2.5.16}$$

$$[S^i, K^j] = i\epsilon^{ijk}K^k \tag{2.5.17}$$

$$[K^i, K^k] = -i\epsilon^{ijk}S^k \tag{2.5.18}$$

It is clear that the S^i operators obey the typical angular momentum commutation rules as they were defined in Eq.(2.3.5). It is even possible to define an operator $\mathrm{S}^2 = (S^1)^2 + (S^2)^2 + (S^3)^2$ called the Casimir operator which commutes with all of the generators [25]. This shows the fundamental link between angular momentum operators and the spatial generators of Lorentz symmetry. Also, since these commutation relations are the Lie algebra of this group then they will be satisfied by the generators of symmetry no matter what system, i.e. representation, is under consideration. The S^i operators are called the generators of rotations. To see why, one rewrites a generic element of $SO(3,1)$ using the new S^i and K^i operators:

$$D(g(\omega^{\mu\nu})) = \exp(-i(\theta^i S^i - \eta^i K^i)) \tag{2.5.19}$$

where $\theta^i = (1/2)\epsilon^{ijk}\omega^{jk}$ and $\eta^i = \omega^{i0}$. In the 4-vector representation, the S^i are written as

$$S^1 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -i \\ 0 & 0 & i & 0 \end{pmatrix} \quad S^2 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & i \\ 0 & 0 & 0 & 0 \\ 0 & -i & 0 & 0 \end{pmatrix} \quad S^3 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \tag{2.5.20}$$

Consider then the action of a transformation solely generated by S^3, i.e. where $\eta^i = 0\ \forall i$ and $\theta^i = 0\ \forall i \neq 3$, on the 4-position vector x^μ

$$x'^\mu = \left[\exp(-i\theta^3 S^3)\right]^\mu_{\ \nu} x^\nu = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta^3) & -\sin(\theta^3) & 0 \\ 0 & \sin(\theta^3) & \cos(\theta^3) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} t \\ x \\ y \\ z \end{pmatrix} \tag{2.5.21}$$

This means that in the 4-vector representation, the S^3 generates the rotation transformation around the z-axis in (x, y) plane with the θ^3 parameter being the angle of rotation. It can be similarly shown that S^1, S^2 generate rotations in the (y, z) and (z, x) plane respectively. It is clear that all the transformations generated by the three S^i matrices form a subgroup since any composition of spatial rotations is itself a spatial rotation, any rotation can be inverted by rotating with the opposite angle and there is an identity element when all angles are null. This Lie group is called the $SO(3)$ and its generators obey the Lie algebra defined in Eq.(2.5.16). Staying in the 4-vector representation, the Casimir operator S^2 would be:

$$\mathrm{S}^2 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix} \tag{2.5.22}$$

It is clear that all three generators and the Casimir operator are block diagonal in this representation. This means that it is decomposed into two sub-spaces that do not mix. The Casimir operator in each of these sub-spaces is written as $s(s+1)\mathbf{I}_{2s+1}$ where $s \in \{0, 1\}$ and $\mathbf{I}_n$ is the n-dimensional identity matrix. This means that each subspace is labelled by a number s, which represents the eigenvalue of S^2 in it. Since all three generators commute with the Casimir operator but not with each other, one can only choose one of them to diagonalize simultaneously. Choosing S^3 gives us in each sub-space integer eigenvalues m_s that span $-s \leq m_s \leq s$. This means that when it comes to spatial rotations, a generic 4-vector representation is decomposed into two representations on two independent sub-spaces labelled $s = 0, 1$ each having $2s + 1$ different eigenstates for the S^3 operator. This is intuitive since temporal

components are unaffected by spatial rotations, all the generators are then simply 0 when restricted to this subspace.

While the 4-vector representation has allowed the elucidation of many of the properties of the Lorentz group, it is by no means the only physically interesting representation. All higher-order tensors, by definition, obey Lorentz symmetry. That being said, their representations are generally simple to find as their transformation laws are based on the transformation of the 4-vector where a generic tensor obeys

$$T'^{\alpha_1\alpha_2\cdots\alpha_n}{}_{\beta_1\beta 2\cdots\beta_n} = \Lambda^{\alpha_1}{}_{\delta_1}\Lambda_{\beta_1}{}^{\gamma_1}\cdots\Lambda^{\alpha_n}{}_{\delta_n}\Lambda_{\beta_n}{}^{\gamma_n}T^{\delta_1\delta_2\cdots\delta_n}{}_{\gamma_1\gamma 2\cdots\gamma_n} \tag{2.5.23}$$

allowing one to construct the higher order representations from the 4-vector one. This would also lead to eigenspaces that are labelled by even higher s numbers with a larger amount of eigenstates. For example, a generic rank-2 tensor representation would have an eigenspace labelled by $s = 2$ containing 5 states. It should also be noted that exists a type of representation called a spinor which has a different transformation under the Lorentz group and which has half-integer values for the s numbers [28]. This representation is the appropriate one to describe particles such as electrons.

However, what I really want to do is to see the symmetries obeyed by a particular field $\mathbf{\Phi}(x)$. The interesting bit here is that in addition to having its own tensor structure with its own transformation under the Lorentz group, a field depends on the 4-position which in turns also transforms under the Lorentz group. Since I am ultimately interested in applying this analysis to the electromagnetic field, consider the 4-vector field $A^\mu(x)$. At a particular point in spacetime, the previous discussion showed how the components of the field transform:

$$x'^\mu = \Lambda^\mu{}_\nu x^\nu \tag{2.5.24}$$

$$A'^\mu(x') = \Lambda^\mu{}_\nu A^\nu(x) \tag{2.5.25}$$

However, in order to figure out the representation of generators on the field, one needs to isolate its variation under the transformation from the variation in the coordinate x'^μ. One does this by considering an infinitesimal transformation and Taylor expanding $\delta_0 A^\lambda(x) = A'^\lambda(x) - A^\lambda(x)$ while keeping only terms linear in the transformation parameters [25]:

$$\begin{aligned}\delta_0 A^\lambda(x) &= A'^\lambda(x'^\rho + \frac{i}{2}\omega_{\mu\nu}(S^{\mu\nu})^\rho{}_\sigma x^\sigma) - A^\lambda(x) \\ &= A'^\lambda(x') + \frac{i}{2}\omega_{\mu\nu}(S^{\mu\nu})^\rho{}_\sigma x^\sigma\partial_\rho A^\lambda(x) - A^\lambda(x) \\ &= -\frac{i}{2}\omega_{\mu\nu}(S^{\mu\nu})^\lambda{}_\sigma A^\sigma(x) + \frac{i}{2}\omega_{\mu\nu}(S^{\mu\nu})^\rho{}_\sigma x^\sigma\partial_\rho A^\lambda(x)\end{aligned} \tag{2.5.26}$$

One then uses the explicit form of $S^{\mu\nu}$ in the 4-vector representation to define a new generator

$$L^{\mu\nu} = -(S^{\mu\nu})^{\rho}{}_{\sigma} x^{\sigma} \partial_{\rho} = i(x^{\mu}\partial^{\nu} - x^{\nu}\partial^{\mu}) \tag{2.5.27}$$

This generator clearly obeys the $SO(3,1)$ algebra. Also, it is possible to do the same decomposition that was done before by defining $L^i = (1/2)\epsilon^{ijk}L^{jk} = i\epsilon^{ijk}x^j\partial^k = -i(\mathbf{r}\times\boldsymbol{\nabla})^i$ thereby restricting the situation to purely spatial transformations. These operators obey the $SO(3)$ Lie algebra and have the same form as the orbital angular momentum operators defined for the hydrogen atom in Eq.(2.3.4). Unlike the S^i operators, these operators act on the functional aspect of the field. Although this operator is infinitely-dimensional, its eigenvalues are discrete. By analogy to the previously defined quantum angular momentum operators, a Casimir operator for the L^i operators can be constructed and diagonalized along with the L^3 operator on the basis of the spherical harmonics $Y_l^m(\theta,\phi)$ defined in the Eq.(2.3.8).

The $L^{\mu\nu}$ generator allows one to rewrite Eq.(2.5.26) as

$$\delta_0 A^{\lambda} = -\frac{i}{2}\omega_{\mu\nu}((S^{\mu\nu})^{\lambda}{}_{\sigma} + L^{\mu\nu})A^{\sigma} = -\frac{i}{2}\omega_{\mu\nu}(J^{\mu\nu})^{\lambda}{}_{\sigma}A^{\sigma} \tag{2.5.28}$$

with $J^{\mu\nu} = S^{\mu\nu} + L^{\mu\nu}$ being the representation of the generators of the Lorentz group in the linear space of the 4-vector fields. This corresponds to generators of $SO(3)$ on the 4-vector field that have the form $J^i = S^i + L^i$, indicating where the split between the two different kinds of angular momentum operators comes from. For fields, the total generators of rotation are composed of two contributions: one component is related to the rotation of the field itself. This is inherent to the type of field under consideration; further, the generators for these rotations have a form and eigenspace that depends on the particular representation to which the field belongs. For example, they are null on scalar fields since they are unchanged under any Lorentz transformation; on the 4-vector fields they are represented by the matrices in Eq.(2.5.20) and so on for higher order tensor fields. These generators are said to generate *internal* rotations and can thus be related to the spin angular momentum of the field and its operators.

The second component L^i is related to how the field transforms under the rotation of the spacetime on which it acts. The generators for this transformation have a general expression that is independent from the particular representation of the field under consideration and acts on the spatial dependence of the field rather than on the various components that make it up. These generators are always in the same infinitely dimensional representation. This is because they do not act on vectors or tensors, but on the infinitely dimensional linear space of functions of spacetime

coordinates. These generators are said to generate *spatial* rotations and can then be linked to the orbital angular momentum of the field.

In the previous sections, it was shown that the angular momentum of a system can be thought of as the quantity that is conserved due to its symmetry of rotations. To see how this then links it inextricably to the generators of said symmetry, consider the application of Noether's theorem to field dynamics which can be found in field theory textbooks [38]. Staying with the 4-vector field, the field Lagrangian density can be defined as $\mathcal{L}(A_\mu, \partial_\mu A^\nu)$, which in turns defines the following action

$$S = \int d^4x \mathcal{L}(A_\mu, \partial_\mu A_\nu) \tag{2.5.29}$$

Like in the case for classical particles, the evolution of the dynamics of the fields can be found from the solutions of the field Euler-Lagrange equations, defined as

$$\partial_\mu \frac{\partial \mathcal{L}}{\partial(\partial_\mu A_\nu)} = \frac{\partial \mathcal{L}}{\partial A_\nu} \tag{2.5.30}$$

A generic transformation of the field $\delta_0 A^\mu$ and spacetime coordinates δx^μ causes a variation in the action that is written, for fields which obey the Euler-Lagrange equations, as

$$\delta S = -\int d^4x \partial_\mu \Big(\frac{\partial \mathcal{L}}{\partial(\partial_\mu A_\nu)} \delta_0 A_\nu - \delta x^\mu \mathcal{L} \Big) \tag{2.5.31}$$

where the first component comes from the variation of the Lagrangian when the field is transformed but the coordinates kept constant; and the second from the variation of the volume element. If this transformation is a symmetry, then the action is unchanged $\delta S = 0$. This allows one to define a conserved *Noether current* $\partial_\mu j^\mu$ where

$$j^\mu = \frac{\partial \mathcal{L}}{\partial(\partial_\mu A_\nu)} \delta_0 A_\nu - \mathcal{L} \delta x^\mu \tag{2.5.32}$$

as well as a Noether charge

$$Q = \int d^3x j^0 \tag{2.5.33}$$

Using Stokes' theorem and the fact that, for physical systems, the fields vanish quickly enough at infinity, one can readily show that $\partial_0 Q = 0$. Thus, this charge represents a conserved quantity in the dynamics of the field.

Restricting oneself to infinitesimal Lorentz transformations, one can use Eq.(2.5.24) and Eq.(2.5.26) to rewrite δx^μ and $\delta_0 A^\mu$ respectively. Since Lorentz symmetry must be respected no matter the parameters for a particular transformation $\omega_{\mu\nu}$, one defines

multiple Noether charges associated with each generator, so that

$$\begin{aligned} Q^{(\mu\nu)} &= \int d^3x j^{(\mu\nu)0} \\ &= \int d^3x \Big(\frac{\partial \mathcal{L}}{\partial(\partial_0 A_\lambda)} (L^{\mu\nu} A_\lambda + [S^{\mu\nu}]_\lambda{}^\sigma A_\sigma) - \mathcal{L}[S^{\mu\nu}]^0{}_\sigma x^\sigma \Big) \end{aligned} \tag{2.5.34}$$

Restricting oneself even further to the spatial rotations in $SO(3)$ as before, the component coming from $\mathcal{L}[S^{\mu\nu}]^0{}_\sigma x^\sigma$ vanishes since the S^i generators have no temporal components in the 4-vector representation. This allows one to define the three charges related to the symmetries of spatial rotations of the field.

$$Q^{(i)} = \int d^3x \Big(\frac{\partial \mathcal{L}}{\partial(\partial_0 A_\lambda)} (L^i A_\lambda + [S^i]_\lambda{}^\sigma A_\sigma) \Big) \tag{2.5.35}$$

This therefore shows that it is possible to associate each generator of rotational symmetry with a conserved charge for any field that obeys the Euler-Lagrange equations. All of the conclusions that have been reached so far are valid for any generic 4-vector field. My goal now is to see how this applies to the electromagnetic field and how it can be linked to my earlier description of its angular momentum. To that end, I now need to define the appropriate Lagrangian and field operator.

Typically, the physically measurable quantities that describe a particular electromagnetic field are the electric $\mathbf{E}$ and magnetic $\mathbf{B}$ field vectors, which obey Maxwell's equations as they are defined in Eq.(2.4.4). However, these are not the actual field operators in field theory. It is convenient to describe the field using the electromagnetic potentials

$$\mathbf{E} = -\nabla\phi - \frac{\partial \mathbf{A}}{\partial t} \qquad\qquad \mathbf{B} = \nabla \times \mathbf{A} \tag{2.5.36}$$

with ϕ and $\mathbf{A}$ being the scalar and vector potentials. Although the electromagnetic interaction was the first physical theory to show Lorentz symmetry, the description of its dynamics using Maxwells equations does not make this immediately obvious. Fortunately, classical field theory offers a description that is both manifestly Lorentz covariant and further provides a clearer progression to the quantization that will be required when QED effects will be considered later on. Following the standard methods of classical electrodynamics, I start by defining some important tensors that are essential for the covariant formulation [29]. The electromagnetic field will now be completely described by the 4-potential A^μ:

$$A^\mu(x) = (\phi, \mathbf{A}) \tag{2.5.37}$$

This is a 4-vector field and thus transforms in the way defined in Eq.(2.5.25) under a Lorentz transformation. The physical fields are represented in the anti-symmetric electromagnetic field tensor $F^{\mu\nu}$ defined as:

$$F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu \tag{2.5.38}$$

Considering the definition of the physical electric and magnetic fields in Eq.(3.2.3), the electromagnetic field tensor is then written in terms of the traditional **E** and **B** fields [29] as

$$F^{\mu\nu} = \begin{pmatrix} 0 & -E_x & -E_y & -E_z \\ E_x & 0 & -B_z & B_y \\ E_y & B_z & 0 & -B_x \\ E_z & -B_y & B_x & 0 \end{pmatrix} \tag{2.5.39}$$

Using these objects, the classical electromagnetic field in the absence of sources is described, in Lagrangian formulation, in the following way [25]:

$$\mathcal{L}(A_\mu, \partial_\mu A_\nu) = -\frac{1}{4} F^{\mu\nu} F_{\mu\nu} \tag{2.5.40}$$

This means that

$$\frac{\partial \mathcal{L}}{\partial(\partial_0 A_0)} = 0 \tag{2.5.41}$$

$$\frac{\partial \mathcal{L}}{\partial(\partial_0 A_i)} = -(\partial^0 A^i - \partial^i A^0) = E^i \tag{2.5.42}$$

Applying this to Eq.(2.5.35) and using the fact that, when restricted to the spatial components, the generators of internal rotations can be written as $S^i = -i\epsilon^{ijk}$, the conserved Noether charges can be rewritten as

$$\begin{aligned} Q^{(i)} &= i \int d^3x (E^a(\epsilon^{ijk} x^j \partial^k) A^a - \epsilon^{iab} E^a A^b) \\ &= -i \int d^3x (E^a(\mathbf{r} \times \boldsymbol{\nabla})^i A^a + (\mathbf{E} \times \mathbf{A})^i) \end{aligned} \tag{2.5.43}$$

Since $\boldsymbol{\nabla} \cdot \mathbf{E} = 0$ in the absence of sources and $\mathbf{B} = \boldsymbol{\nabla} \times \mathbf{A}$, one can use the vector identity for the triple product and integrate by parts while assuming that the fields vanish at infinity to show that the above equation is equivalent to

$$\mathbf{Q} = \int d^3x (\mathbf{r} \times (\mathbf{E} \times \mathbf{B})) \tag{2.5.44}$$

This means that the conserved charges resulting from the symmetries generated via L^i and S^i are none other than the total angular momentum of the electromagnetic field as it was defined in Eq.(2.4.13).

2.6 Conclusion

The analysis that is outlined in this section leads to several conclusions. First, the quantum mechanical angular momentum operators have a fundamental origin in that they simply arise from the generators of rotations in the Lorentz group. For any physical field, there are two types of generators that can be defined, which are distinct in both their origin and the way they act on the field. One type of those is represented by generators of internal rotations whose representation and eigenspace depend on the kind of field under question. This generator acts on the tensor nature of the fields and is related to the way the tensor transforms under rotation. In the case of light, the generator takes the particularly simple form in Eq.(2.5.20). Technically, this should mean that there are four different states of light with definite (s, m_s) eigenvalues. However, only two of these states are physically allowable: $(1, 1)$ and $(1, -1)$. These correspond to the right and left circular polarizations of light respectively. The main reason why the other states are forbidden is seen from the transversality condition for electromagnetic waves $\mathbf{k} \cdot \mathbf{n} = 0$, meaning that there are only two independent components for the polarization of electromagnetic excitations. The other type corresponds to the generators of spatial rotations, which are independent from the representation and the field under consideration. They have the same form for the electromagnetic field as they do for other fields, such as electrons. Their eigenspace is discrete but infinite and they act on the spacetime dependence of the electromagnetic field. Together these generators are related, via Noether's theorem, to the same conserved quantity that was earlier identified, in Eq.(2.4.10), as being the total angular momentum of the electromagnetic field.

It should be noted, however, that this is not the only method of defining the angular momentum operators. Instead of identifying the operators with the generators of rotations, another approach defines the conserved Noether charge in Eq.(2.5.43) as being the total angular momentum operator of the electromagnetic field with the first and second terms in the integrand being the orbital and spin angular momentum densities respectively. This is generally called the "canonical" version. Other approaches include the "Belinfante" version which starts from the classical definition of angular momentum that is related to the Poynting vector of an electromagnetic field and upgrades that integral to an operator [39], as well as the version proposed to make the spin and orbital densities in the canonical version separately gauge-invariant by splitting the electromagnetic potential into transverse and longitudinal components [21, 30]. Each of these versions has its own weaknesses and strengths: the canonical

version has a fundamental link to symmetry transformations but lacks gauge invariance for each angular momentum operator separately, the Belinfante version can be mapped to the angular momentum that naturally arises in classical electromagnetism but has no fundamental separation between the spin and orbital components and the gauge invariant version does not obey the typical angular momentum commutation relations for its operators [8]. One way which has been suggested to adjudicate which of these versions is the best description, is to measure the angular momentum density of the electromagnetic field in an experiment where an electromagnetic wave interacts with matter [40].

Another interesting conclusion, that was reached in the previous section, is that the so-called angular momentum commutation relations in Eq.(2.3.5) have a fundamental origin in that they are the Lie algebra that all generators of rotations must obey. This is true for the two types of generators, corresponding to "internal" and spacetime rotations, that were defined above no matter which physical system is under consideration, be it the electromagnetic field or electrons. Some of the controversy surrounding the behaviour of the angular momentum operators for light and their commutation relations stems from the above-mentioned conflicting views on how they should be defined. If one were to use gauge-invariant definition of the angular momentum operators as being the total conserved charge due to rotations after the integration of the spacetime dependence, expanding the field operators in the integral in certain bases would cause the resultant integrated orbital and spin operators not to obey the Lie algebra, as was done in reference [8]. However, this method of expressing the angular momentum operators is flawed in that it involves the field operators in the definition of a physical observable. The conserved charges can be thought more as the expectation value of the angular momentum operators defined by the generators.

Lastly, there remains the question of whether the individual components of the angular momentum of light are conserved separately or whether only the total angular momentum is conserved. As in the case for the electron in the hydrogen atom, this also depends on which system is considered and how it is modelled. If spin-orbit coupling is included for the electron, the projections of its spin and orbital components are no longer separately conserved since they interact with each other. For light, a similar analysis is possible. For the case of a paraxial beam moving in free space, I have shown that the envelope function $u(\mathbf{r})$ that relates to the OAM is completely independent from the polarization of the beam defining its SAM. However, there have been spin-orbit effects observed for light in certain conditions such as whispering gallery mode emitters [41] and photonic crystal waveguides [42]. Additionally, if we go back to the

light-matter interaction considered previously, when an atom absorbs a photon it is the total angular momentum that is conserved during this interaction and not simply its separate components. In the next chapter, I will show how the conservation of the angular momentum of light can be leveraged to experimentally evidence one of the longest standing predictions of quantum electrodynamics: vacuum photon-photon scattering.

Chapter 3

Orbital angular momentum in photon-photon scattering

"Erwin with his psi can do
Calculations quite a few.
But one thing has not been seen:
Just what does psi really mean?"

Eric Huckel

3.1 Context

As mentioned in the introduction, light has been the basis for many scientific investigations due to how important it is in everyday life, from being indispensable for one of humanity's most important senses to being the current best carrier for high-speed information. Throughout the history of scientific research on light, there have been many models to describe its nature and behaviour ranging from geometric rays of particles to waves in an electromagnetic field. However, for the vast majority of that time, light beams were never considered to be able to interact with one another in the absence of matter. This was mostly in line with the empirical evidence that was available. Unlike colliding particles, colliding rays of light seem to be generally transparent to one another. The only times light seems to interact with itself are in the context of non-linear optics where effects, such as harmonic generation coupling various laser pulses, can occur. This happens because certain media are made of material that exhibits significant non-linear polarization in the presence of strong electromagnetic fields. This non-linear polarization can couple different modes of radiation together when intense light propagates in these non-linear media.

This all changed with the development of the theory of quantum electrodynamics (QED) in the early-to-mid 20$^{\text{th}}$ century. This new theory posited that even in the absolute vacuum, there exist fluctuations in the electromagnetic and Fermion fields which QED describes. These fluctuations are not simply mathematical curiosities, they lead to detectable effects such as the observed Lamb shift between two Hydrogen energy levels [43, 44] or the Casimir effect where a force can be detected between two conducting plates placed in the vacuum due to the difference between the electromagnetic fluctuations inside the cavity formed by the plates versus outside that cavity [45, 46]. One of the most fascinating consequences of these vacuum fluctuations is when virtual Fermion particle-antiparticle pairs are generated in the vacuum and annihilated almost instantly leading to a theorised *vacuum polarization*, which can act in the same way as the non-linear polarization of certain dielectric media, causing non-linear coupling between modes of electromagnetic radiation. This vacuum polarization makes it then possible for photons to scatter off of each other in the absence of matter. This vacuum photon-photon scattering is an effect that has attracted considerable interest over the past few decades [47, 48]. However, due to the weakness of vacuum polarization, experimentally observing this effect is no easy feat. Many different proposals have been made to detect such processes using microwave cavities [47], plasma channels [49] and high-power lasers [50, 51]. The main obstacle to these schemes has been the relatively weak predicted signal compared to the noise inherent to any experiment. As such, scattering between real photons has yet to be observed in the vacuum.

In the previous chapter, I showed how the OAM of light can be defined both from a classical point of view, in the Laguerre-Gaussian solutions to the electromagnetic wave equation, and from a quantum field theoretical point of view where it arises from symmetry considerations. In this chapter, I investigate the effect of OAM coupling in elastic photon-photon scattering and show that OAM contributes an additional unique signature in the photons generated via this elusive effect. This signature will help filter the aforementioned background noise and is thus predicted to make the detection of these weak signals significantly easier.

This chapter will be organised as follows: first, it will be shown why photon-photon scattering is purely a quantum mechanical effect not possible with the classical electromagnetic field. Then, the path integral formulation will be used to derive an effective field theory and a non-linear electromagnetic wave equation, allowing for a semi-classical treatment of light-light scattering. Afterwards, a harmonic analysis of the non-linear equation is performed to show some of its properties. Following that,

exact solutions to the linear electromagnetic wave equation which carry OAM will be discussed. These modes will then be used to investigate the effect of OAM on photon-photon interaction. Finally, a proposed experimental configuration will be presented that is expected to validate the accompanying theoretical work. As this chapter focuses on effects from quantum field theory, I will use the natural units throughout this chapter, except where is explicitly indicated otherwise.

3.2 The classical electromagnetic field

Before considering the modifications brought on by QED, I will first turn my attention to classical electromagnetism and how it models the dynamics of light. Recall the 4 Maxwell equations as well as the definition of the electric and magnetic fields in terms of the potentials:

$$\nabla \cdot \mathbf{E} = \rho \qquad \nabla \times \mathbf{B} = \mathbf{J} + \frac{\partial \mathbf{E}}{\partial t} \tag{3.2.1}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \qquad \nabla \cdot \mathbf{B} = 0 \tag{3.2.2}$$

$$\mathbf{E} = -\nabla\phi - \frac{\partial \mathbf{A}}{\partial t} \qquad \mathbf{B} = \nabla \times \mathbf{A} \tag{3.2.3}$$

with ϕ and $\mathbf{A}$ being the scalar and vector potentials. Since this chapter will, later on, focus on the effects arising from QED, I will be using the covariant formulation of classical electromagnetism outlined in the previous chapter. Whereas there I restricted myself to the Lagrangian of the vacuum electromagnetic field, here I want to introduce coupling to charges and currents. To that end, I define the four-potential A^μ and the four-current as

$$A^\mu = (\phi, \mathbf{A}) \qquad J^\mu = (\rho, \mathbf{J}) \tag{3.2.4}$$

where ρ is the charge density and $\mathbf{J}$ is the current density. Using these objects, along with the electromagnetic field tensor $F^{\mu\nu}$, the dynamics of the classical electromagnetic field in the presence of sources are described by the following Lagrangian density [25]:

$$\mathcal{L}(A_\mu, \partial_\mu A_\nu) = -\frac{1}{4} F^{\mu\nu} F_{\mu\nu} - A_\mu J^\mu \tag{3.2.5}$$

This Lagrangian is composed of two parts, the first represents the electromagnetic field in vacuum while the second part is the minimal coupling of the field to any electromagnetic sources of charge and current. The advantage of the Lagrangian

formulation is that the Lorentz covariance of the electromagnetic field is now explicit, as we have constructed the Lagrangian using only Lorentz scalars. Another useful symmetry is seen by recalling how the physically observable electric and magnetic fields are defined using the elements of the four potential in Eq.(3.2.3). Due to equations relating them with the scalar and vector potential, the electric and magnetic fields are invariant under the following transformation:

$$\mathbf{A} \to \mathbf{A} + \nabla\theta \qquad\qquad \phi \to \phi - \frac{\partial\theta}{\partial t} \tag{3.2.6}$$

with θ being an arbitrary scalar function of spacetime coordinates. Again, this is written more concisely in an explicitly Lorentz covariant way as:

$$A^\mu \to A^\mu - \partial^\mu\theta \tag{3.2.7}$$

This symmetry is what is called *gauge invariance*, where Eq.(3.2.7) is the associated gauge transformation, and further plays an essential role in the theoretical description of electromagnetic interactions. As an example for the important consequences of gauge invariance, consider the action of the electromagnetic field.

$$S = \int d^4x \mathcal{L}(A_\mu, \partial_\mu A_\nu) \tag{3.2.8}$$

Since the physical fields are gauge invariant, one would like for the action, which determines the equations of motion, to be so as well. Eq.(2.5.38) shows that the electromagnetic field tensor is gauge invariant, hence if I perform the gauge transformation outlined in Eq.(3.2.7) only the minimal coupling part leaves an additional term in the new action

$$\delta S = -\int d^4x J^\mu \partial_\mu\theta \tag{3.2.9}$$

By requiring the four-current to vanish at infinity, and by further requiring this difference in the action to vanish to respect gauge invariance, integration by parts leads to the familiar expression of charge conservation

$$\partial_\mu J^\mu = 0 \tag{3.2.10}$$

Now that all the relevant parts of the classical electromagnetic field and its symmetries have been defined, all that is left is to derive the equations of motion that govern its interactions. As usual for the Lagrangian formulation, this is done by extremising the action integral defined above in Eq.(3.2.8). This extremum is be found by using the Euler-Lagrange equations for fields

$$\partial_\mu \frac{\partial\mathcal{L}}{\partial(\partial_\mu A_\nu)} = \frac{\partial\mathcal{L}}{\partial A_\nu} \tag{3.2.11}$$

This equation, when applied to the Lagrangian defined in Eq.(3.2.5), leads to the following equations of motion:

$$\partial_\mu F^{\mu\nu} = \partial_\mu \partial^\mu A^\nu - \partial^\nu(\partial_\mu A^\mu) = J^\nu \tag{3.2.12}$$

Using Eq.(2.5.39) along with the definition of the four-current, it can be shown that this equation is equivalent to the first pair of Maxwell's of equations (Eq.(3.2.1)). The remaining Maxwell equations can be derived by considering the so-called *dual* electromagnetic tensor defined as

$$\tilde{F}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\alpha\beta}F_{\alpha\beta} \tag{3.2.13}$$

where $\epsilon^{\mu\nu\alpha\beta}$ is the totally anti-symmetric tensor. $\tilde{F}^{\mu\nu}$ is also gauge-invariant and, as such, provides a suitable alternative description for the physical electromagnetic fields. In fact, one can use it in the Lagrangian instead of the original tensor rewriting the former using the identity $\tilde{F}^{\mu\nu}\tilde{F}_{\mu\nu} = -F^{\mu\nu}F_{\mu\nu}$. From the definition of the dual tensor, it is clear that

$$\partial_\mu \tilde{F}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\alpha\beta}(\partial_\mu\partial_\alpha A_\beta - \partial_\mu\partial_\beta A_\alpha) = 0 \tag{3.2.14}$$

This is because $\epsilon^{\mu\nu\alpha\beta}$ is anti-symmetric in all of its indices while the first element in parentheses is symmetric under exchange of $\mu \leftrightarrow \alpha$ and the second element is symmetric under $\mu \leftrightarrow \beta$. This identity leads to the remaining Maxwell's equations. It should be noted that Eq.(3.2.14) was derived completely independently from the Lagrangian and the equations of motion: it is merely a consequence of the definition of the electromagnetic field tensor. This is reflected in the fact that there are no source terms involved no matter what medium or free sources are being considered. Let us now consider the equations of motion defined in the vacuum where there are no sources and $J^\mu = 0$. Taking advantage of the gauge invariance defined above, one defines the four-potential such that

$$A^0 = 0 \qquad\qquad \nabla \cdot \mathbf{A} = 0 \tag{3.2.15}$$

This is called the *radiation gauge* and it implies that $\partial_\mu A^\mu = 0$. This means that Eq.(3.2.12) can be written more simply as

$$(\frac{\partial^2}{\partial t^2} - \nabla^2)\mathbf{A} = 0 \tag{3.2.16}$$

with $\phi = 0$, as per the gauge condition. The easiest, most immediate solutions, to this electromagnetic wave equation are the so-called *plane wave* modes:

$$\mathbf{A}(x^\mu) = Ae^{-ik_\nu x^\nu}\mathbf{n}(\mathbf{k}) \tag{3.2.17}$$

Where A is a constant amplitude, $k^\mu = (\omega, \mathbf{k})$ is the four-wavevector and $\mathbf{n}(\mathbf{k})$ is the polarization of the vector potential. The wave equation implies that $k_\mu k^\mu = 0$, producing the familiar dispersion relation for electromagnetic waves in the vacuum, while the gauge condition implies that $\mathbf{n}(\mathbf{k}) \cdot \mathbf{k} = 0$. Choosing the coordinate system such that the plane wave is propagating along the z-direction, this means that the polarization vector lies in the transverse plane, giving the field two degrees of freedom for its polarization. The most general solution to Eq.(3.2.16) is then written as a plane wave decomposition analogous to the Fourier decomposition.

$$\mathbf{A}(x^\mu) = \int \frac{d^3k}{(2\pi)^3\sqrt{2\omega_k}} \sum_{\sigma=\pm} a_\sigma(\mathbf{k}) e^{-ik_\nu x^\nu} \mathbf{n}_\sigma(\mathbf{k}) + \text{c.c.} \tag{3.2.18}$$

Where $a_\sigma(\mathbf{k})$ is the amplitude of the specific plane wave component labelled by $(\sigma, \mathbf{k})$ and $\mathbf{n}_\pm(\mathbf{k}) = \mathbf{e}_x \pm i\mathbf{e}_y$ are the *circular* polarization vectors defined in the transverse (x, y) plane to the direction of propagation. This decomposition of the polarization is convenient as it expresses the field in terms of modes of definite helicity. The normalization of the integral is somewhat arbitrary but it was chosen to render the amplitude dimensionless, while the complex conjugate is added to ensure that the field is real. Since $\phi = 0$, the physical electric field is derived from the vector potential using $\mathbf{E} = -\partial\mathbf{A}/\partial t$. It should be noted that although the radiation gauge provides a convenient way of solving the equations of motion for the electromagnetic field and decomposing it in the familiar transverse plane waves basis, the treatment becomes no longer manifestly Lorentz covariant as can be seen in the gauge conditions in Eq.(3.2.15).

Looking at Eq.(3.2.16), it is clear that the electromagnetic wave equation is linear in the vacuum. This leads to what is called the *superposition principle*, meaning that any superposition of solutions to the equation will also be a solution to the equation. This is seen from the Lagrangian since in the absence of any sources of charge and current, it is merely a quadratic function of the fields, and as such the equations of motions are linear and there are no self-interaction terms. This seems to indicate that two electromagnetic waves cannot interact or scatter in the vacuum. However, as will be seen later on, once the Fermion and electromagnetic fields are quantized, the vacuum is not strictly speaking "empty".

3.3 Quantum electrodynamics

The dawn of the 20$^{\text{th}}$ century brought rapidly-moving transformative change to our understanding of the fundamental physics of the universe. The "two clouds" pointed

out by Lord Kelvin, in his 1900 lecture at the Royal Institution, quickly led to the development of the ground-breaking theories of Quantum Mechanics [52] and Special Relativity [53]. However, it was quickly realised that they did not appear to be compatible with one another as the fundamental equation in quantum mechanics, the Schrödinger equation [54], was not Lorentz covariant since it treated space and time differently. The quest for symmetry led to the development of QFT pioneered by Paul Dirac [55]. This new theory elegantly describes the elementary particles and their interactions as coupling between quantized excitations of certain fields. The very first interaction to successfully receive this treatment was the electromagnetic interaction in what came to be known as Quantum Electrodynamics. Currently, QED is the most rigorously tested physical theory and has provided the most precise predictions we have to date such as predicting the fine structure constant with an uncertainty better than one part in a *billion* [56].

Evidently, the first step to this description is to quantize all the relevant fields starting with the electromagnetic field described before. Since the primary goal of QED is to describe relativistic quantum mechanics, let us forgo the radiation gauge described in the previous section for the Lorentz invariant *Lorentz gauge*:

$$\partial_\mu A^\mu = 0 \tag{3.3.1}$$

Plugging this in the equations of motion in the vacuum leads to the wave equation for the entire four-potential $\partial_\mu \partial^\mu A^\nu = 0$. However, we can no longer decompose the field into helicity eigenstates since the condition $\mathbf{n}(\mathbf{k}) \cdot \mathbf{k} = 0$ is replaced with $k_\mu n^\mu = 0$.

The so-called canonical quantization in QFT is done in analogy to the one in non-relativistic quantum mechanics. There, a Hamiltonian $H(q^i, p^j)$ and its specific generalised coordinates q^i and conjugate momenta p^j are "upgraded" to quantum mechanical operators $\hat{\mathrm{H}}$, $\hat{\mathrm{q}}^i$, $\hat{\mathrm{p}}^j$ with the Poisson brackets $\{p^i, q^j\}_{\mathrm{PB}} = \delta^{ij}$ and $\{q^i, q^j\}_{\mathrm{PB}} = \{p^i, p^j\}_{\mathrm{PB}} = 0$ "upgraded" to the commutation relations $[\hat{\mathrm{q}}^i, \hat{\mathrm{p}}^j] = i\delta^{ij}$ and $[\hat{\mathrm{q}}^i, \hat{\mathrm{q}}^j] = [\hat{\mathrm{p}}^i, \hat{\mathrm{p}}^j] = 0$. The quantity of interest in QFT is not the specific coordinates or momenta of the particles but the fields that describe them $A^\mu(x^\mu)$. We can then, again by generalisation from classical Lagrangian mechanics, define the conjugate "momenta" [25]

$$\Pi^\mu(x^\nu) = \frac{\partial \mathcal{L}}{\partial(\partial_0 A_\mu(x^\nu))} \tag{3.3.2}$$

that can then be used to define the Hamiltonian density

$$\mathcal{H}(A^\mu, \Pi^\nu) = \Pi^\mu \partial_0 A_\mu - \mathcal{L} \tag{3.3.3}$$

Since the fields depend on time, when they are upgraded to operators, the commutation relations for them are only defined at *equal time*. The previously defined commutators generalise to

$$[\hat{\mathrm{A}}^{\mu}(\mathbf{x},t),\hat{\Pi}^{\nu}(\mathbf{y},t)] = i\eta^{\mu\nu}\delta^{(3)}(\mathbf{x}-\mathbf{y}) \tag{3.3.4}$$

$$[\hat{\mathrm{A}}^{\mu}(\mathbf{x},t),\hat{\mathrm{A}}^{\nu}(\mathbf{y},t)] = [\hat{\Pi}^{\mu}(\mathbf{x},t),\hat{\Pi}^{\nu}(\mathbf{y},t)] = 0 \tag{3.3.5}$$

To make the notation less cumbersome, the "hat" on the quantized field operators will be omitted, relying on the context to distinguish them from their classical counterpart. This quantization poses a problem when considering the vacuum Lagrangian for the electromagnetic field as it was previously defined. As can be seen in Eq.(3.2.5), the Lagrangian does not depend on $\partial_0 A_0$ meaning that $\Pi^0 = 0$. Respecting Eq.(3.3.4) is then impossible since the commutator of any field operator with Π^0 would be trivial and $A^0 \neq 0$ in the Lorentz gauge. The avoid this issue, while maintaining Lorentz invariance, consider a more general vacuum Lagrangian:

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - \frac{1}{2}(\partial_\mu A^\mu)^2 \tag{3.3.6}$$

At first glance, this Lagrangian appears not to be compatible with the electromagnetism of Maxwell's equations as the newly-added second term breaks the gauge invariance found in the physically observable fields. However, it does now offer a way to properly define all the necessary conjugate momenta. It also immediately leads to the wave equation $\partial^\nu\partial_\nu A^\mu = 0$ without the need for any gauge condition. The electromagnetic field operator can now be decomposed into plane waves

$$A^\mu = \int \frac{d^3k}{(2\pi)^3\sqrt{2\omega_k}} \sum_{\lambda=0}^{3} \left(n^\mu_{(k,\lambda)}\hat{\mathrm{a}}_{k,\lambda}e^{-ik_\nu x^\nu} + n^\mu_{(k,\lambda)}\hat{\mathrm{a}}^\dagger_{k,\lambda}e^{ik_\nu x^\nu}\right) \tag{3.3.7}$$

Where $k_\mu k^\mu = 0$ and n^μ is the completely unrestricted polarization four-vector with four independent degrees of freedom. For each four-wavevector, it useful to orient spacetime to have the spatial wavevector on the z-axis and define the four possible polarization along the axis of that spacetime. Since the field is now an operator, the field amplitudes in Eq.(3.2.18) also become operators. To maintain the commutation relations between the field operators defined in Eq.(3.3.4), these amplitude operators must obey the relations

$$[\hat{\mathrm{a}}_{k\lambda},\hat{\mathrm{a}}^\dagger_{k'\lambda'}] = -(2\pi)^3\delta^{(3)}(\mathbf{k}-\mathbf{k}')\eta_{\lambda\lambda'} \tag{3.3.8}$$

By analogy with the quantum harmonic oscillator, $\hat{\mathrm{a}}_{\lambda k}$, $\hat{\mathrm{a}}^\dagger_{\lambda k}$ are called the *annihilation* and *creation* operators respectively. This allows for the construction of a

Fock space which describes the state of the electromagnetic field in terms of excitations (photons) in particular modes (k) and polarizations (λ). The vacuum field $|0\rangle$, containing no photons, is defined as

$$\hat{\mathrm{a}}_{k\lambda}\,|0\rangle = 0 \tag{3.3.9}$$

while any other basis state in the Fock space can be constructed from the vacuum via

$$|k_1,\lambda_1;\cdots;k_n,\lambda_n\rangle = \sqrt{2\omega_{k_1}}\cdots\sqrt{2\omega_{k_n}}\hat{\mathrm{a}}^\dagger_{k_1\lambda_1}\cdots\hat{\mathrm{a}}^\dagger_{k_n\lambda_n}\,|0\rangle \tag{3.3.10}$$

Now that the completely covariant quantization of the electromagnetic field has been done, the issue of the lack of gauge invariance in the Lagrangian still remains to be addressed. Additionally, this method of quantization seems to allow the existence of unphysical longitudinally polarised and time polarized photons. These inconsistencies with physical electromagnetism can be resolved by what is known as the Gupta formalism [57]. Instead of considering the complete Fock space constructed with the creation and annihilation operators defined above, I will constrain myself to the subspace of physically allowable states ($\{|\mathbf{\Phi}_{\mathrm{Phys}}\rangle\}$ defined as

$$\langle\mathbf{\Phi}'_{\mathrm{Phys}}|\,\partial_\mu A^\mu\,|\mathbf{\Phi}_{\mathrm{Phys}}\rangle = 0 \tag{3.3.11}$$

By considering the expansion of the field operators in terms of Fock operators, this condition means that a generic physical single-photon state with wavevector k^μ is written as

$$|\mathbf{\Phi}_{\mathrm{Phys}}\rangle = |\mathbf{\Phi}_T\rangle + \alpha(\hat{\mathrm{a}}^\dagger_{k,0} - \hat{\mathrm{a}}^\dagger_{k,3})\,|0\rangle \tag{3.3.12}$$

with $|\mathbf{\Phi}_T\rangle$ being a generic transverse state created by the action of a superposition of $\hat{\mathrm{a}}^\dagger_{k,1}$ and $\hat{\mathrm{a}}^\dagger_{k,2}$ on the vacuum. The commutation relations for the Fock operators immediately lead to the conclusion that the second part of Eq.(3.3.12) has zero norm. By calculating the Hamiltonian density of the electromagnetic field, it can be shown that that second part of Eq.(3.3.12) does not contribute to the energy of the state. Hence, this physical state is equivalent to just $|\mathbf{\Phi}_T\rangle$ and the subspace of physical states is further constrained to only consider transverse photons, which matches with what is physically observable and what is predicted by Maxwell's equations.

The next step in describing the electromagnetic interaction is to define the fundamental particles that experience it. From experimental observations, the most common of these particles are massive spin 1/2 Fermions. The covariant field operators that describe the behaviour for these particles is the 4-dimensional Dirac spinor

field $\mathbf{\Psi}(\mathbf{x}, t)$. First, I will define some mathematical objects, which are useful when constructing the relevant Lagrangian, starting with the standard Pauli matrices.

$$\sigma^0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \qquad \sigma^1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad \sigma^2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \qquad \sigma^3 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \tag{3.3.13}$$

From those, the 4×4 γ-matrices can be defined, which in the so-called *chiral* representation are written as

$$\gamma^\mu = \begin{pmatrix} \mathbf{0} & \sigma^\mu \\ \bar{\sigma}^\mu & \mathbf{0} \end{pmatrix} \tag{3.3.14}$$

where $\bar{\sigma}^\mu = (1, -\sigma)$ and the γ-matrices satisfy the Dirac algebra

$$\{\gamma^\mu, \gamma^\nu\} = \gamma^\mu\gamma^\nu + \gamma^\nu\gamma^\mu = 2\eta^{\mu\nu}\mathbf{I}_4 \tag{3.3.15}$$

The Lagrangian density for the Dirac field can then be compactly written as [28]

$$\mathcal{L} = \bar{\mathbf{\Psi}}(i\gamma^\mu\partial_\mu - m)\mathbf{\Psi} \tag{3.3.16}$$

Here, one defines the Dirac adjoint $\bar{\mathbf{\Psi}} = \mathbf{\Psi}^\dagger\gamma^0$. As one would hope, this Lagrangian is a Lorentz scalar and thus invariant. Using the Euler-Lagrange equations, the famous Dirac equation for free Fermions can be derived

$$(i\partial\!\!\!/ - m)\mathbf{\Psi} = 0 \tag{3.3.17}$$

where $\partial\!\!\!/ = \gamma^\mu\partial_\mu$ uses Feynman's "slash" notation. Staying in the chiral representation, the Dirac spinor is written using the left and right 2-dimensional Weyl spinors $\psi_{L,R}$

$$\mathbf{\Psi} = \begin{pmatrix} \psi_L \\ \psi_R \end{pmatrix} \tag{3.3.18}$$

Eq.(3.3.17) then implies the simpler Klein-Gordon equation for each of these spinors

$$(\partial^\mu\partial_\mu + m^2)\psi_{L,R} = 0 \tag{3.3.19}$$

This, of course, admits simple plane wave solutions oscillating at $e^{-ip^\mu x_\mu}$ with $p^2 = E^2 - \mathbf{p}^2 = m^2$. From the previous dispersion relation, the solutions for these particles are thus split into positive and negative-energy solutions. Since the Dirac equation is a first-order PDE, there are two distinct classes of solutions for the Dirac spinor field [28]:

$$\mathbf{\Psi} = u^s(\mathbf{p})e^{-ip^\mu x_\mu} \quad \text{with} \quad u^s(\mathbf{p}) = \begin{pmatrix} \sqrt{\mathbf{p}\cdot\sigma}\xi^s \\ \sqrt{\mathbf{p}\cdot\bar{\sigma}}\xi^s \end{pmatrix} \quad \text{for} \quad p^0 > 0 \tag{3.3.20}$$

$$\mathbf{\Psi} = v^s(\mathbf{p})e^{ip^\mu x_\mu} \quad \text{with} \quad v^s(\mathbf{p}) = \begin{pmatrix} \sqrt{\mathbf{p}\cdot\sigma}\eta^s \\ -\sqrt{\mathbf{p}\cdot\bar{\sigma}}\eta^s \end{pmatrix} \quad \text{for} \quad p^0 < 0 \tag{3.3.21}$$

Where $s \in \{1,2\}$, and ξ^s and η^s are two-dimensional spinors that label the distinct spinor solutions $u^s(\mathbf{p})$ and $v^s(\mathbf{p})$ while satisfying the orthogonality conditions $\xi^{\dagger s}\xi^r = \eta^{\dagger s}\eta^r = \delta^{sr}$. Now that solutions to the field equations have been defined, the canonical quantization of the Dirac field is undertaken in the same way as for the electromagnetic field. After upgrading the fields to operators, the conjugate momenta are derived from Eq.(3.3.16):

$$\Pi = \frac{\partial \mathcal{L}}{\partial(\partial_0 \mathbf{\Psi})} = i\mathbf{\Psi}^\dagger \tag{3.3.22}$$

Since the Dirac field describes spin 1/2 Fermions, defining commutation relations for it is not appropriate. Instead, the spin-statistics theorem implies that field operators obey *anti-commmutation* relations instead:

$$\{\mathbf{\Psi}_a(\mathbf{x},t), \mathbf{\Psi}_b^\dagger(\mathbf{y},t)\} = \delta^{(3)}(\mathbf{x}-\mathbf{y})\delta_{ab} \tag{3.3.23}$$

where instead of using the full conjugate momenta defined in Eq.(3.3.22), I have divided all relevant relations by the imaginary number i. Similarly to what was done with the electromagnetic field operator, the Dirac field operator can now be expanded in the plane-wave basis

$$\mathbf{\Psi} = \int \frac{d^3p}{(2\pi)^3\sqrt{2E_{\mathbf{p}}}} \sum_{s=1,2} \left(\hat{\mathrm{a}}_{\mathbf{p},s} u^s(\mathbf{p}) e^{-ip^\mu x_\mu} + \hat{\mathrm{b}}^\dagger_{\mathbf{p},s} v^s(\mathbf{p}) e^{ip^\mu x_\mu}\right) \tag{3.3.24}$$

Unlike the electromagnetic field, there are no reality conditions on the Dirac field, allowing the $\hat{\mathrm{a}}_{\mathbf{p},s}$ and $\hat{\mathrm{b}}_{\mathbf{p},s}$ to be distinct. The anti-commutation relations for the fields impose similar ones on the momentum-space operators

$$\{\hat{\mathrm{a}}_{\mathbf{p},s}, \hat{\mathrm{a}}^\dagger_{\mathbf{q},r}\} = \{\hat{\mathrm{b}}_{\mathbf{p},s}, \hat{\mathrm{b}}^\dagger_{\mathbf{q},r}\} = (2\pi)^3\delta^{(3)}(\mathbf{p}-\mathbf{q})\delta^{sr} \tag{3.3.25}$$

allowing one to consider these operators as creation and annihilation ones. Eq.(3.3.24) then shows that the positive energy solutions can be interpreted as particles created by $\hat{\mathrm{a}}^\dagger_{\mathbf{p},s}$ while the negative energy solutions are anti-particles created by $\hat{\mathrm{b}}^\dagger_{\mathbf{p},s}$. The main free fields involved in the electromagnetic interaction have now been defined and quantized with their total Lagrangian density being

$$\mathcal{L} = \bar{\mathbf{\Psi}}(i\gamma^\mu\partial_\mu - m)\mathbf{\Psi} - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} - \frac{1}{2}(\partial_\mu A^\mu)^2 \tag{3.3.26}$$

So far, we have only treated the free non-interacting fields, what is missing to complete the QED description is the coupling between them.

To that end, symmetry is again relied upon. Along with the typical Lorentz invariance, the Lagrangian in Eq.(3.3.26) exhibits another interesting symmetry with respect to the $U(1)$ transformation

$$\mathbf{\Psi} \to e^{iq\theta}\mathbf{\Psi} \tag{3.3.27}$$

It is clear that if θ is constant, then the Lagrangian remains invariant. This type of symmetry, which depends on a constant parameter, is called *global*. However with a slight modification to the Lagrangian, it can be promoted to a *local* symmetry, i.e. one where $\theta = \theta(x^\mu)$. Consider the following Lagrangian:

$$\hat{\mathcal{L}} = \bar{\mathbf{\Psi}}(i\gamma^\mu\partial_\mu - m)\mathbf{\Psi} - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} - \frac{1}{2}(\partial_\mu A^\mu)^2 - qA_\mu\bar{\mathbf{\Psi}}\gamma^\mu\mathbf{\Psi} \tag{3.3.28}$$

It is invariant over the previously defined physical subspace when the two following transformation are done concurrently

$$\mathbf{\Psi} \to e^{iq\theta}\mathbf{\Psi} \qquad\qquad A^\mu \to A^\mu - \partial^\mu\theta \tag{3.3.29}$$

with θ being an arbitrary scalar function of spacetime. Eq.(3.3.28) is the full QED Lagrangian and it shows why QED is called a *gauge theory*. The gauge invariance of the physical electromagnetic field was used to promote the global $U(1)$ symmetry to a local one, with the consequence being the coupling of the electromagnetic field to the Dirac field.

However, this coupling had an unfortunate effect. It is no longer possible to easily solve the equations of motions for the fields or decompose them into the convenient plane waves modes. Exactly solving any equations derived from the full Lagrangian is no longer feasible. In order to gain some insight into the interactions between the fields, one must turn to perturbation theory. Eq.(3.3.28) can be split into three parts

$$\mathcal{L} = \mathcal{L}_{\text{EM}} + \mathcal{L}_{\text{Dirac}} + \mathcal{L}_{\text{Int}} \tag{3.3.30}$$

where $\mathcal{L}_{\text{EM}}$ and $\mathcal{L}_{\text{Dirac}}$ are the free electromagnetic and Dirac Lagrangian, and $\mathcal{L}_{\text{Int}} = -qA_\mu\bar{\mathbf{\Psi}}\gamma^\mu\mathbf{\Psi}$ is the interaction Lagrangian responsible for the coupling between them.

Consider an input state consisting of an electron and a positron $|\mathbf{p}_e\mathbf{k}_p\rangle_{\text{in}}$ that are initially, at $T_i \to -\infty$, far enough apart that they do not interact. As they evolve, their state is modified by the coupling. One of the most important applications of QED is to determine what is the probability that after their interaction, the electron and the positron state end up, at $T_f \to \infty$, in the state $|\mathbf{p}'_e\mathbf{k}'_p\rangle_{\text{out}}$ where they are again considered to be far enough apart to no longer be interacting. One of the ways to

solve this is to derive what is called the *scattering amplitude* of the interaction, which is quantified by the S-matrix element ${}_{\text{out}}\langle \mathbf{p}'_e \mathbf{k}'_p | S \, | \mathbf{p}_e \mathbf{k}_p \rangle_{\text{in}}$. The S-matrix contains all the information about the evolution of the input state via the interaction. It is useful to decompose it into $S = \mathbf{1} + iT$ where iT contains all the non-trivial interactions that do not leave the input state unchanged. It can be then shown that [28]

$$ {}_{\text{out}}\langle \mathbf{p}'_e \mathbf{k}'_p | iT \, | \mathbf{p}_e \mathbf{k}_p \rangle_{\text{in}} = {}_0\langle \mathbf{p}'_e \mathbf{k}'_p | \mathcal{T}(\exp(-i \int d^4x \mathcal{H}_I)) \, | \mathbf{p}_e \mathbf{k}_p \rangle_0 \tag{3.3.31} $$

where $\mathcal{T}$ is the time-ordering operator, $\mathcal{H}_I$ is the Hamiltonian density in the interaction picture and the states with subscript 0 are the free unperturbed states. In the case of the electromagnetic interaction, since $\mathcal{L}_{\text{int}}$ has no dependence on derivatives of the field, the interaction Hamiltonian is simply $\mathcal{H}_I = qA_\mu \bar{\mathbf{\Psi}} \gamma^\mu \mathbf{\Psi}$, where all the fields are also in the interaction pictures and can thus be expanded as plane waves, in the same way as for the free fields.

Even with all the prior simplifications, Eq.(3.3.31) is practically impossible to solve exactly. However, the exponential can be expanded into powers of the coupling terms with each higher order either revealing new possible interactions or adding successive corrections to pre-existing ones. The field operators in the coupling terms can then be expanded into creation and annihilation operators via the canonical quantization procedure introduced in Eq.(3.3.7) and Eq.(3.3.24). Therefore, the contribution of each expansion can be computed. The zeroth order term is trivial as it represents no interaction between the input and output states. For the particular electron-positron interaction under consideration, the first interesting contribution to the scattering amplitude comes from the second-order term. Due to the complex nature of these calculations, a set of rules were devised to keep track of all the different terms that could arise from Eq.(3.3.31) along with a graphical representation of the calculations. This framework is called the Feynman diagrams. For the electron-positron interaction, there are two distinct diagrams that arise from the second-order expansion, these are shown in Fig.(3.1).

Along with providing a useful way of keeping track of S-matrix calculations, Feynman diagrams provide an intuitive interpretation of the various terms arising from the coupling. For example, as shown in the right part of Fig.(3.1), the electromagnetic interaction that leads to electron-positron scattering is no longer just an abstract invisible force between them but is caused by the exchange of a "virtual" photon, coming from the electromagnetic field that couples them, which mediates the momentum transfer. These two diagrams are, of course, not the only contributors to the electron-positron interaction. Higher-order corrections can be derived corresponding

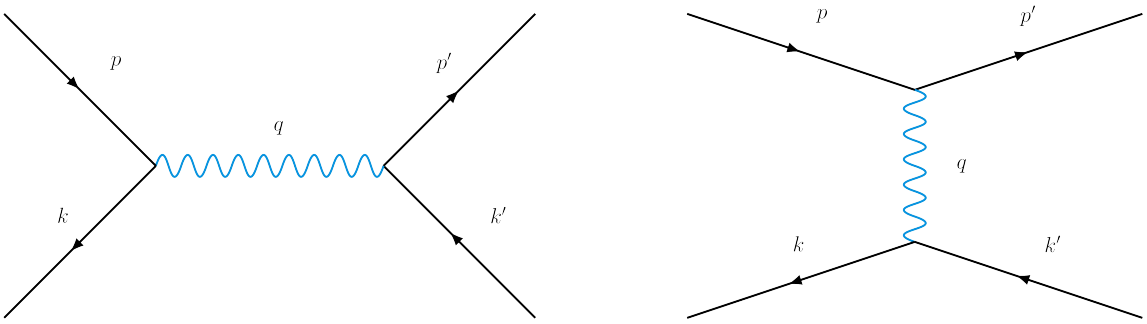

Figure 3.1: The two distinct momentum space Feynman diagrams representing the second-order expansion of the S-matrix for electron-positron interactions. Straight lines are Fermions and wavy lines represent photons. Left: This diagram represents electron-positron annihilation and recreation. Right: This diagram represents electron-positron scattering.

to ever more complex Feynman diagrams that make it possible to calculate more accurate predictions. This is what has allowed QED to be the most accurately tested physical theory we have to date. In fact, the previously mentioned result on the fine-structure constant was achieved by calculating 8^{th} order corrections corresponding to 891 Feynman diagrams.

All the elements are now in place to start tackling the problem of photon-photon scattering in vacuum. Unlike in classical electromagnetism, the effect of the coupling between the Dirac and the electromagnetic fields in QED is clearly non-linear, and all the various interactions can now be interpreted as various Feynman diagrams showing exchange of momenta between the particles of the fields. To see how this can make photon-photon scattering possible, recall some of the rules of Feynman diagrams for QED [38]:

1. Every successive expansion has the same number of interaction vertices in its diagram as its order;

2. Each vertex has to connect one photon line, one incoming Fermion line and one outgoing Fermion line;

3. The initial and final states are represented by external lines while internal lines represent "virtual" particles which mediate the interaction;

4. Charge is conserved in the interaction represented at every vertex.

These rules show that it is not possible to have scattering between just two photons in vacuum as there would not be enough interaction vertices to couple them.

However, the same is not true when considering four photons. Expanding Eq.(3.3.31) to fourth order gives the required four interaction vertices and thus provides the lowest order contribution to photon-photon scattering in the vacuum. One of the resultant coupling terms can be represented by the Feynman diagram in Fig.(3.2).

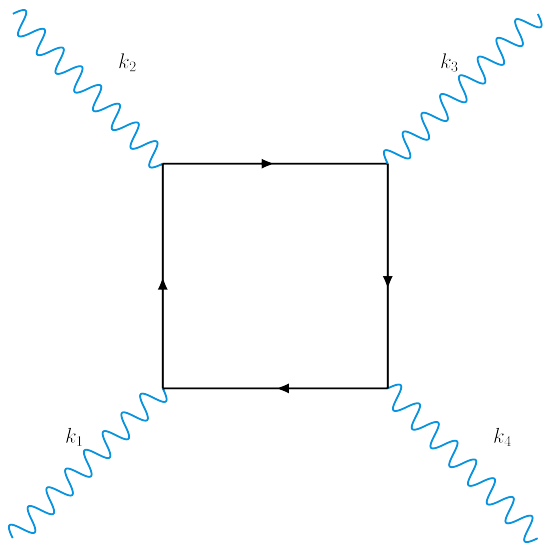

Figure 3.2: Feynman diagram representing one of the lowest-order contributions to photon-photon scattering.

This diagram now shows that scattering between photons in vacuum is possible via mediating virtual electron-positron pairs that form what is called a loop diagram. There are two different ways to interpret these virtual electron-positron pairs which can arise even from low-energy photons. The first is the so-called "old-fashioned perturbation theory" where the Heisenberg uncertainty relations predict that energy conservation can be violated over a very short time interval leading to the possibility of the production of these virtual pairs since they get annihilated almost immediately. The other way to interpret this is that there is always a finite, if small, probability of a photon fluctuating into a virtual pair while still conserving energy because the virtual particles do not obey the typical mass-energy dispersion relations ($m^2 = E^2 - \mathbf{p}^2$) and as such have properties different than their real counterparts which have to obey these relations.

This effect, where the virtual electron-positron pairs generated by a photon modify the electromagnetic proprieties of the vacuum, is generally called vacuum polarization [28]. Photon-photon scattering is only one of the many predicted consequences of vacuum polarization. Ever since the latter was first theoretically modelled [58, 59, 60], many of its predicted effects have been experimentally verified. These include

Delbrück scattering, where a photon scatters off the Coulomb field generated by a nucleus, which was first observed to occur with Lead atoms [61] as well as photon splitting, which is when a photon interacting with an electromagnetic field splits into two photons, which was observed in atomic fields [62]. Other vacuum polarization effects remain without experimental verification so-far but there have been many experimental proposals and campaigns attempting to evidence them [63]. One such effect is vacuum birefringence where the modifications brought on by the Fermion loop induced by an electromagnetic field can make the vacuum act like a birefringent material affecting the polarization of other electromagnetic waves. There have been many ongoing initiatives to experimentally verify this effect either by using magnetic fields interacting with electromagnetic waves in a Fabry-Përot cavity such as in the BMV and PVLAS campaigns [64, 65, 66], or by using a high-power optical laser interacting with an x-ray free electron laser (XFEL) such as the experiment that is proposed for the HIBEF facility [67], among others. Beyond vacuum birefringence, there are many other unmeasured consequences of vacuum polarization which have driven various interesting experimental proposals such as vacuum photon acceleration [68] and vacuum harmonic generation [69].

Among the effects that remain unverified, I have chosen to focus on photon-photon scattering. Although vacuum polarization has already been confirmed in previous experiments as detailed above, researching the photon-photon interaction is still interesting and has the potential to lead to impactful results. Not only is this because its experimental verification can add further confirmation to the validity of QED, but also because any deviations in the detected signal can point to new and interesting physics beyond the framework of the standard model. This is because vacuum polarization is not only predicted to occur due to non-linear QED effects, there are also theories which hypothesize that it can occur due to coupling mediated by axions [70], milli-charged particles [71] or due to other models of non-linear electrodynamics such as the Born-Infeld model [72, 73]. As such, the results of an investigation of photon-photon scattering can confirm or put bounds on the parameters of these non-standard models.

It should be noted that scattering due to vacuum polarization does not only occur between real photons. It can also occur between the virtual photons which mediate the electromagnetic interaction. In fact, an experiment performed using the ATLAS detector at the large hadron collider (LHC) was the first to report photon-photon scattering between virtual photons [74]. This occurred because strong electromagnetic fields, generated during the ultra-peripheral collisions between heavy relativistic ions,

can be modelled as photons with very low virtuality. As mentioned above, virtual particles can be interpreted to violate the standard mass-energy dispersion relation in order to preserve energy conservation. As such, The virtuality of a particle Q^2 can be defined as the deviation of its 4-momentum p^μ from the standard dispersion relation:

$$Q^2 = p^2 - m^2 \tag{3.3.32}$$

where m is the mass of the particle and $p^2 = p^\mu p_\mu$ making it so that $Q^2 = 0$ for real particles. The electromagnetic fields caused by the aforementioned relativistic ions can be treated as photons with virtuality as low as $Q^2 < 10^{-3}\,\mathrm{GeV}^2$. Due to the relativistic nature of these collisions, the electromagnetic fields can be strong enough for the higher-order QED corrections, such as the ones due to the one-loop photon-photon interaction diagram shown above, to be relevant and have a detectable impact. Although this detection is an important step in experimentally proving photon-photon scattering, there is still no experimental verification for this interaction between real photons in the vacuum.

One of the ideal ways to show this would be simply to study the interaction between high-power lasers in the vacuum. This approach has been steadily attracting more interest over the past two decades thanks in part to the ever increasing power of new laser systems that will be capable of delivering multi-petawatt pulses [9, 75]. Many proposals have been put forth to leverage these high-power lasers in order to investigate the photon-photon scattering which is predicted to occur when they collide [51, 76, 77]. Although these proposals predict a theoretically measurable signal using modern laser systems, they have a low signal-to-noise ratio due to the presence of other competing interactions that occur due to imperfect experimental conditions. An example of such sources of noise is the Compton scattered photons produced when the intense laser pulses couple with the particles that still remain in the interaction region in imperfect vacuum chambers.

Over the next sections, I will show how leveraging the orbital angular momentum of light can help increase the expected signal-to-noise of the photon-photon scattering signal. To that end, a quantitative model must first be made for this interaction when using laser radiation so that the effects of laser spatial modes can be included. Although the diagram in Fig.(3.2) gives a very intuitive and clear picture of how this interaction can happen, calculating the full scattering amplitude using the S-matrix formulation is a daunting task, even at the lowest order contributions, and requires careful consideration to avoid any divergences that may arise. To that end, it is more convenient to treat this problem from an effective field theory point of view.

3.4 Effective field theory and the Euler-Heisenberg Lagrangian

The canonical quantization procedure introduced is not the only method for calculating scattering amplitudes, and in fact it is not always the most optimal way of doing so. An alternative procedure exists called the path integral formulation of quantum mechanics in which the scattering amplitude between initial and final states is [38]

$$\langle i, T_i | f, T_f \rangle = \int \mathcal{D}\mathbf{\Psi}\mathcal{D}\bar{\mathbf{\Psi}}\mathcal{D}A^\mu \exp(iS_{\text{QED}}) \tag{3.4.1}$$

where the initial and final state ($|i, T_i\rangle$ and $|f, T_f\rangle$) are in the time independent Heisenberg picture, the integration in the functional integral runs over all possible physical field configurations between the initial and final states and $S_{\text{QED}} = \int d^4x \mathcal{L}_{\text{QED}}$ is the QED action. Eq.(3.4.1) shows that the scattering amplitude is the sum of all possible paths in configuration space between the initial and final field configurations, each weighted by the complex exponential of their action. Applying the stationary phase approximation to this integral leads to the classical principle of least action.

To derive an effective field theory for the vacuum electromagnetic field, one simply integrates out all the contributions of the Dirac field and then approximates by only keeping terms up to the one loop contributions to the QED vacuum shown in the previous section. To that end, it is possible to expand the action in Eq.(3.4.1), group together all the terms that depend on the Dirac fields and then integrate them out to define an effective Lagrangian dependent only on the electromagnetic field:

$$\exp\left(i \int d^4x \mathcal{L}_{\text{eff}}(A_\mu)\right) = \int \mathcal{D}\mathbf{\Psi}\mathcal{D}\bar{\mathbf{\Psi}} \exp\left(i \int d^4x (\bar{\mathbf{\Psi}}(i\gamma^\mu \partial_\mu - m)\mathbf{\Psi} - qA_\mu \bar{\mathbf{\Psi}}\gamma^\mu \mathbf{\Psi})\right) \tag{3.4.2}$$

This integral is quite challenging to calculate, even if it is just a perturbative expansion. However, it satisfies a few properties that can restrict what kind of terms the effective Lagrangian might contain. The integral on the RHS of Eq.(3.4.2) is manifestly Lorentz invariant since everything inside of it is a Lorentz scalar. It is also, by construction, gauge invariant under the previously defined QED gauge transformation of the Dirac and electromagnetic fields. Finally, since one is integrating all the Dirac fields, it can only depend on the electromagnetic field. These restrictions imply that the effective Lagrangian is only made of powers of $F_{\mu\nu}F^{\mu\nu}$ and $F_{\mu\nu}\tilde{F}^{\mu\nu}$, bearing in mind that $\tilde{F}_{\mu\nu}\tilde{F}^{\mu\nu} \propto F_{\mu\nu}F^{\mu\nu}$. By expanding the integral and keeping only terms up to ones corresponding to the one loop interaction terms in the Feynman

diagrams ($\mathrm{O}(m_e^{-4})$), the so-called Euler-Heisenberg effective Lagrangian is found [58]:

$$\mathcal{L}_{\mathrm{E-H}} = \mathcal{L}_{\mathrm{EM}} + \mathcal{L}_{\mathrm{eff}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \frac{\xi}{8}\Big((F_{\mu\nu}F^{\mu\nu})^2 + \frac{7}{4}(F_{\mu\nu}\tilde{F}^{\mu\nu})^2\Big) \tag{3.4.3}$$

where $\xi = \dfrac{4\alpha^2}{45m_e^4}$ with $\alpha = \dfrac{e^2}{4\pi}$ being the fine-structure constant. It should be noted that the modification that was made to allow for the covariant quantization of the electromagnetic field (the added $(\partial_\mu A^\mu)^2$ term) has been omitted, as I am only concerned with physical fields and am no longer going to quantize them. Also, since this merely an expansion keeping the lowest order contributing terms, this effective Lagrangian has a limit on its validity. Specifically, it is only adequate for fields varying slowly with respect to the electron Compton frequency ($\lambda >> 10\,\mathrm{pm}$) and whose strengths are small relative to the Schwinger limit ($E_c = 1.32 \times 10^{18}\,\mathrm{V\,m^{-1}}$).

The Euler-Heisenberg Lagrangian allows for a semi-classical treatment of the photon-photon scattering. To see how it implies the non-linearity of the electromagnetic vacuum, I will derive the new effective equations of motions. By applying the Euler-Lagrange equations to Eq.(3.4.3), they can be written as

$$\partial_\mu F^{\mu\nu} = \xi\partial_\mu\Big((F_{\alpha\beta}F^{\alpha\beta})F^{\mu\nu} + \frac{7}{4}(F_{\alpha\beta}\tilde{F}^{\alpha\beta})\tilde{F}^{\mu\nu}\Big) \tag{3.4.4}$$

It is more convenient to rewrite this equation in terms of the traditional electric and magnetic fields. Remembering the definition of the electromagnetic field tensor from Eq.(2.5.39), it can be readily shown that $(F_{\alpha\beta}F^{\alpha\beta}) = -2(\mathbf{E}^2 - \mathbf{B}^2)$ and $F_{\alpha\beta}\tilde{F}^{\alpha\beta} = -4(\mathbf{E}\cdot\mathbf{B})$. Thus, by projecting Eq.(3.4.4) along the various rows of $F^{\mu\nu}$, a modified version of the vacuum Maxwell equations is found:

$$\nabla\cdot\mathbf{E} = -\xi\nabla\cdot\Big(2(\mathbf{E}^2 - \mathbf{B}^2)\mathbf{E} + 7(\mathbf{E}\cdot\mathbf{B})\mathbf{B}\Big) \tag{3.4.5}$$

$$\begin{aligned}\nabla\times\mathbf{B} - \frac{\partial\mathbf{E}}{\partial t} = \xi\Bigg[&\nabla\times\Big(7(\mathbf{E}\cdot\mathbf{B})\mathbf{E} - 2(\mathbf{E}^2 - \mathbf{B}^2)\mathbf{B}\Big)\\ &+\frac{\partial}{\partial t}\Big(2(\mathbf{E}^2 - \mathbf{B}^2)\mathbf{E} + 7(\mathbf{E}\cdot\mathbf{B})\mathbf{B}\Big)\Bigg]\end{aligned} \tag{3.4.6}$$

$$\nabla\times\mathbf{E} = -\frac{\partial\mathbf{B}}{\partial t} \qquad\qquad \nabla\cdot\mathbf{B} = 0 \tag{3.4.7}$$

The last two equations are unchanged from their original classical form since they are a consequence of the structure of the electromagnetic field tensor itself ($\partial_\mu\tilde{F}^{\mu\nu} = 0$) no matter the Lagrangian.

It is instructive to compare these equations to Maxwell's equations in non-linear media, which are usually written, in the absence of free sources of charge and current, as

$$\nabla \cdot \mathbf{E} = -\nabla \cdot \mathbf{P} \qquad \nabla \times \mathbf{B} = (\nabla \times \mathbf{M} + \frac{\partial \mathbf{P}}{\partial t} + \frac{\partial \mathbf{E}}{\partial t}) \tag{3.4.8}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \qquad \nabla \cdot \mathbf{B} = 0 \tag{3.4.9}$$

where $\mathbf{P}$ and $\mathbf{M}$ are the medium's polarization and magnetization respectively. From Eq.(3.4.6) and Eq.(3.4.8), it is clear that we can define a *vacuum polarization* $\mathbf{P}_{\text{vac}}$ and a *vacuum magnetization* $\mathbf{M}_{\text{vac}}$:

$$\mathbf{P}_{\text{vac}} = \xi\Big(2(\mathbf{E}^2 - \mathbf{B}^2)\mathbf{E} + 7(\mathbf{E}\cdot\mathbf{B})\mathbf{B}\Big) \tag{3.4.10}$$

$$\mathbf{M}_{\text{vac}} = \xi\Big(7(\mathbf{E}\cdot\mathbf{B})\mathbf{E} - 2(\mathbf{E}^2 - \mathbf{B}^2)\mathbf{B}\Big) \tag{3.4.11}$$

This means that the electromagnetic vacuum, as described by the Euler-Heisenberg Lagrangian, is non-linear implying that the superposition principle no longer applies and that electromagnetic waves are no longer transparent to one another. To see this more explicitly, the non-linear electromagnetic wave equation can be derived by taking the curl of the left equation in Eq.(3.4.9) and replacing the resultant curl of the magnetic field with its value in Eq.(3.4.8) leaving us with

$$\partial^\mu \partial_\mu \mathbf{E} = (\nabla(\nabla\cdot \mathbf{P}_{\text{vac}}) - \partial_t(\partial_t \mathbf{P}_{\text{vac}} + \nabla \times \mathbf{M}_{\text{vac}})) \tag{3.4.12}$$

This equation indicates that one of the consequences of the modifications to the electromagnetic Lagrangian, brought on by the effective field theory, is to add a field-dependent source term to the electromagnetic wave equation.

3.5 The non-linear vacuum

To investigate the effect of the source term in Eq.(3.4.12) and how it allows for light-light scattering, consider the case with three initial electromagnetic plane waves so that the total field is $\mathbf{E}(\mathbf{r},t) = \sum_{i=1}^{3} \mathbf{E}_i(\mathbf{r},t)$ where each wave can be written in complex notation as

$$\mathbf{E}_i(\mathbf{r},t) = E_j e^{i(\mathbf{k}_i\cdot\mathbf{r}-\omega_i t)}\mathbf{n}_i(\mathbf{k}_i) \tag{3.5.1}$$

with $\mathbf{k}_i^2 = \omega_i^2$, $\mathbf{n}_i(k)$ being the polarization of the field necessarily satisfying the orthogonality condition $\mathbf{n}_i \cdot \mathbf{k}_i = 0$ and $i \in \{1,2,3\}$ labelling each individual wave. It should be noted that since the complex notation is being used for the fields, the

squared fields are calculated from their real part (e.g. $\mathbf{E}^2 = \Re(E)^2$). From Eq.(3.4.10) and Eq.(3.4.11), it is clear that the non-linearity is cubic, so the inhomogeneous wave equation, Eq.(3.4.12), can be rewritten to first order as

$$\begin{aligned}\partial^\mu\partial_\mu\mathbf{E} = \sum \mathbf{A}_{ijk}E_iE_jE_ke^{i((\mathbf{k}_i+\mathbf{k}_j+\mathbf{k}_k)\cdot\mathbf{r}-(\omega_i+\omega_j+\omega_k)t)} \\ +\mathbf{B}_{ijk}E_iE_jE_k^*e^{i((\mathbf{k}_i+\mathbf{k}_j-\mathbf{k}_k)\cdot\mathbf{r}-(\omega_i+\omega_j-\omega_k)t)} \\ +\mathbf{C}_{ijk}E_iE_j^*E_k^*e^{i((\mathbf{k}_i-\mathbf{k}_j-\mathbf{k}_k)\cdot\mathbf{r}-(\omega_i+\omega_j-\omega_k)t)}\end{aligned} \tag{3.5.2}$$

with $i, j, k \in 1, 2, 3$. This equation shows that, due to the non-linearity of the vacuum, a multitude of different harmonics are generated via the different photon-photon interactions between the three initial waves. It is possible to investigate this equation via harmonic analysis in order to simplify it to a more solvable form. In particular, one of the interactions mixes all three initial waves generating a fourth mode $(\mathbf{k}_4, \omega_4)$, satisfying the following energy and momentum matching conditions:

$$\mathbf{k}_1 + \mathbf{k}_2 = \mathbf{k}_3 + \mathbf{k}_4 \qquad \omega_1 + \omega_2 = \omega_3 + \omega_4 \tag{3.5.3}$$

There are, of course, many other interaction terms that involve different combinations of the three initial waves. However, in the rest of this chapter, only terms that correspond to the matching conditions in Eq.(3.5.3) will be considered. This is because that interaction produces the most convenient generated wave to measure experimentally and allows for flexibility in the geometrical setup of the incoming beams as will be seen in the following sections. This does not necessarily mean that the other interactions terms are particularly weaker. However, they are also not specifically stronger, and waves generated through them will have a different frequency and direction of propagation (and different OAM state as will be seen later on) to the interaction in Eq.(3.5.3), meaning that they can safely be neglected when estimating the number of detected photons in the experimental geometry that will be considered later on. This means that the generated field of interest can be written as

$$\mathbf{E}_4(\mathbf{r}, t) = E_4(\mathbf{r}, t)e^{i(\mathbf{k}_4\cdot r-\omega_4 t)}\mathbf{n}_4(\mathbf{k}_4) \tag{3.5.4}$$

with $\mathbf{k}_4 = \mathbf{k}_1 + \mathbf{k}_2 - \mathbf{k}_3$, $\omega_4 = |\mathbf{k}_4| = \omega_1 + \omega_2 - \omega_3$ and $E_4(\mathbf{r}, t)$ is a slowly varying function of space and time. For this generated field, the left-hand side of Eq.(3.5.2) can be expanded to first order as

$$\partial^\mu\partial_\mu\mathbf{E}_4(\mathbf{r}, t) \approx -2i(\omega_4\frac{\partial}{\partial t} + \mathbf{k}\cdot\nabla)E_4(\mathbf{r}, t)e^{i(\mathbf{k}_4\cdot r-\omega_4 t)}\mathbf{n}_4(\mathbf{k}_4) \tag{3.5.5}$$

Multiplying both sides of Eq.(3.5.2) by $e^{i(\mathbf{k}_4\cdot r-\omega_4 t)}$, it can then be Fourier transformed in space leaving only the time derivative of $E(\mathbf{r}, t)$. On the RHS, only the resonant

source term will have its oscillatory phase cancelled out while all other sources will remain oscillating at $\Delta\omega = \omega_{ijk} - \omega_4$. When integrating over time, all the non-resonant terms oscillating at $\Delta\omega \neq 0$ will vanish and the only relevant non-linear coupling term will be the resonant $E_1 E_2 E_3^*$ term. Hence Eq.(3.5.2) can be simplified to

$$\partial^\mu \partial_\mu \mathbf{E}_4(\mathbf{r}, t) \approx \mathbf{B}_{123} E_1 E_2 E_3^* e^{i((\mathbf{k}_1+\mathbf{k}_2-\mathbf{k}_3)\cdot\mathbf{r}-(\omega_1+\omega_2-\omega_3)t)} \tag{3.5.6}$$

Therefore, one is left with a somewhat familiar-looking electromagnetic wave equation with a source term which describes the growth of the electric field generated via photon-photon scattering. Of course, this equation is only valid if $|\mathbf{E}_4|$ remains relatively small, which as we will see is fortunately (and unfortunately) the case. The most straightforward way to derive an analytical form for the evolution of the generated field is to use the Green's function method. Recalling the Green's function for the wave equation operator [29], the general solution to Eq.(3.5.6) is

$$\mathbf{E}_4(\mathbf{r}, t) = \mathbf{A}_4 E_1 E_2 E_3^* \int \frac{e^{i(\mathbf{k}_4\cdot\mathbf{r}' - \omega_4 t_R(\mathbf{r}'))}}{|\mathbf{r} - \mathbf{r}'|} d^3x' \tag{3.5.7}$$

where $\mathbf{A}_4$ is a constant containing all factors arising from fundamental or geometrical considerations and $t_R(\mathbf{r}') = t - |\mathbf{r}' - \mathbf{r}|$ is the retarded time. Any further analysis of this equation requires that the interaction geometry be specified, which will be done in the following section.

3.6 Orbital angular momentum in photon-photon scattering

Now that the plane wave solution to the non-linear wave equation has been detailed, I want to investigate the effect of OAM on the photon-photon interaction. To do so, I need to solve Eq.(3.4.12) using OAM-carrying modes. This can potentially be done using the Laguerre-Gaussian modes described in the previous chapter. However, considering the mathematical description of these modes

$$\begin{aligned} E_l(\mathbf{r}, t) = \sqrt{\frac{2p!}{\pi(p+|l|)!}} \frac{w_0}{w(z)} \left(\frac{\sqrt{2}\rho}{w(z)}\right)^{|l|} L_p^{|l|}\left(\frac{2\rho^2}{w^2(z)}\right) \\ \times \exp\left(-\frac{\rho^2}{w^2(z)} + i\Big(kz - \omega_0 t + l\phi - \frac{k\rho^2}{2R(z)} + \psi(z)\Big)\right) \end{aligned} \tag{3.6.1}$$

It is clear that they are not the most convenient ones to treat the problem of photon-photon scattering even though they are considered the quintessential OAM-carrying

modes. To that end, my goal is to find a way to express OAM-carrying modes in a way that makes it simpler to solve the non-linear equation.

Going back to the linear electromagnetic wave equation, the simplest solution to it, as was seen in the previous sections, is the plane wave decomposition, where a plane wave is described by:

$$\mathbf{E}(\mathbf{r},t) = E_0 e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)}\hat{\mathbf{n}} \tag{3.6.2}$$

where $\mathbf{k}$ is the wavevector, ω is the angular frequency and $\hat{\mathbf{n}}$ is the polarization unit vector of the plane wave. We then consider an infinite superposition of plane waves whose wave vectors lie on the surface of a cone with half-angle α around the z-axis:

$$\mathbf{E}_l(\mathbf{r},t) = \frac{E_0}{2\pi}\int_0^{2\pi}(i)^l e^{i(\mathbf{k}(\phi_k)\cdot\mathbf{r}-\omega_0 t+l\phi_k)}\hat{\mathbf{n}}(\phi_k)d\phi_k \tag{3.6.3}$$

$$\mathbf{k}(\phi_k) = k\cos(\alpha)\hat{\mathbf{z}} - k\sin(\alpha)(\cos(\phi_k)\hat{\mathbf{x}} + \sin(\phi_k)\hat{\mathbf{y}}) \tag{3.6.4}$$

$$\hat{\mathbf{n}}(\phi_k) = \sin(\phi_k)\hat{\mathbf{x}} - \cos(\phi_k)\hat{\mathbf{y}} \tag{3.6.5}$$

where $l \in \mathbb{Z}$ is the azimuthal mode number. $\mathbf{E}_l$ is obviously then an exact solution to the electromagnetic wave equation as it is a linear superposition of exact solutions. Its integral representation can be simplified by decomposing the trigonometric functions into their exponential form and using the fact that $(1/2\pi)\int_0^{2\pi} e^{i(l\phi - x\sin(\phi))}d\phi = J_l(x)$ [78], where $J_l(x)$ is the l-th order Bessel function of the first kind. Thus, the electric field in Eq.(3.6.3) can be rewritten in cylindrical coordinates (z,ρ,ϕ) as:

$$\mathbf{E}_l(z,\rho,\phi,t) = \mathcal{E}_0 e^{i(\kappa z-\omega t+l\phi)}\left[-\frac{lJ_l(\beta\rho)}{\beta\rho}\hat{\mathbf{e}}_\rho + i\frac{J_{l+1}(\beta\rho)-J_{l-1}(\beta\rho)}{2}\hat{\mathbf{e}}_\phi\right] \tag{3.6.6}$$

Where $\kappa = k\cos(\alpha)$ and $\beta = k\sin(\alpha)$. The corresponding magnetic field can then be calculated from Eq.(5) using Maxwell's equations ($\nabla\times\mathbf{E}_l = -\partial\mathbf{B}_l/\partial t$). This solution can be seen as an OAM-carrying extension of the electromagnetic Bessel modes. One of the interesting properties of Bessel modes is that they propagate without diffraction as can be seen by noting that the radial profile of the field is independent of the longitudinal position along the direction of propagation. True Bessel beams are unphysical as they carry infinite energy. However, Bessel-like modes, which retain some of the aforementioned interesting properties over a known coherence length, have been generated using axicon lenses [79]. When considering the case of $l \neq 0$, the field has an azimuthal phase dependence, and therefore as shown in [1], it carries nonzero

OAM. This can be confirmed by calculating the time-averaged angular momentum density of the field $\langle \mathfrak{L} \rangle = (1/2)\mathbf{r} \times \Re(\mathbf{E}^* \times \mathbf{B})$.

$$\langle \mathfrak{L} \rangle = E_0^2 \left(\frac{l}{k_0} J_l^2(\beta\rho)\hat{\mathbf{z}} - \frac{lz\sin(\alpha)}{\beta\rho} J_l^2(\beta\rho)\hat{\rho} - \rho\cos(\alpha)\frac{J_{l+1}^2(\beta\rho) - J_{l-1}^2(\beta\rho)}{2}\hat{\phi} \right) \quad (3.6.7)$$

If $\langle \mathfrak{L} \rangle$ is integrated over a finite circular transverse profile of the field, only the z-component of the angular momentum density contributes as the polar unit vectors integrate to 0. This results in a total angular momentum $\langle \mathbf{L} \rangle \propto l\hat{\mathbf{z}}$.

It should be noted that this method of defining light modes that carry OAM obviously differs from the standard approach of considering Laguerre-Gaussian (LG) modes shown before. However, it offers two advantages, the first being that the field defined in Eq.(3.6.6) is, by construction, divergence-free and hence an exact solution of Maxwell's equations. Secondly, it offers a simple representation of OAM beams as a superposition of plane waves. This proves highly advantageous when calculating the nonlinear interaction terms compared to the LG modes, which contain a complicated envelope that makes the calculation intractable.

Now that OAM-carrying modes have been well defined, we can go about introducing them in the photon-photon interaction. Let the three initial plane waves be defined by

$$\mathbf{E}_1 = E_1 e^{i(k_0 x - \omega_0 t)}\hat{\mathbf{y}} \qquad \mathbf{E}_2 = E_2 e^{i(-k_0 x - \omega_0 t)}\hat{\mathbf{z}} \quad (3.6.8)$$

$$\mathbf{E}_3 = (i)^l E_3 e^{i(\mathbf{k}_3(\phi_k)\cdot\mathbf{r} - \omega_0 t)} e^{il\phi_k}\hat{\mathbf{n}}(\phi_k) \quad (3.6.9)$$

where $\mathbf{k}_3(\phi_k) = k_0\cos(\alpha)\hat{\mathbf{z}} - k_0\sin(\alpha)(\cos(\phi_k)\hat{\mathbf{x}} + \sin(\phi_k)\hat{\mathbf{y}})$ and $\hat{\mathbf{n}}(\phi_k) = \sin(\phi_k)\hat{\mathbf{x}} - \cos(\phi_k)\hat{\mathbf{y}}$. It is clear that $\mathbf{E}_3$ is a plane wave that lies on the surface of a cone with half-angle α similar to the aforementioned case. Further, α is considered to be very small, which is a reasonable assumption with Bessel beams as they are generally generated with axicon lenses which have a very shallow angle. Keeping only the resonant interaction terms, Eq.(3.5.6) for the generated fourth wave $\mathbf{E}_4$ can then be written as:

$$\partial^\mu \partial_\mu \mathbf{E}_4 \approx (-i)^l \omega_0^2 \xi \hat{\mathbf{v}}(\phi_k) E_1 E_2 E_3^* e^{i(\mathbf{k}_4(\phi_k)\cdot\mathbf{r} - \omega_0 t - l\phi_k)} \quad (3.6.10)$$

where $\hat{\mathbf{v}}(\phi_k) = \cos(\alpha)\cos(\phi_k)\hat{\mathbf{x}} + \cos(\alpha)\sin(\phi_k)\hat{\mathbf{y}} + \sin(\alpha)\hat{\mathbf{z}}$ and $\mathbf{k}_4 = -\mathbf{k}_3$.

Consider now the case where $\mathbf{E}_3$ is not a plane wave, but an l-th order OAM mode whose average wavevector is directed along the positive z-axis. As seen above, the OAM beam can be decomposed into an appropriate superposition of plane waves. It is only then necessary to perform an integration on Eq.(3.6.10), similar to the one

presented above in Eq.(3.6.3), and find that the source term, in this case, can be rewritten as Bessel functions with an overall azimuthal phase dependence.

$$\partial^{\mu}\partial_{\mu}\mathbf{E}_4 \approx \omega_0^2 \xi E_1 E_2 E_3^* \mathbf{J}_l(\rho,\phi) e^{i(-\kappa z - \omega_0 t - l\phi)} \tag{3.6.11}$$

where $\beta = k_0\cos(\alpha)$, $\kappa = k_0\sin(\alpha)$ and $\mathbf{J}$ is a vector containing the transverse radial dependence of the source.

$$\begin{aligned}\mathbf{J}_l(\rho,\phi) =& \sin(\alpha)J_l(\beta\rho)\hat{\mathbf{z}} + i\cos(\alpha)\frac{J_{l+1}(\beta\rho) - J_{l-1}(\beta\rho)}{2}\hat{\rho} \\ &- \frac{l\cos(\alpha)J_l(\beta\rho)}{\beta\rho}\hat{\phi}\end{aligned} \tag{3.6.12}$$

Eq.(3.6.11) can be solved using the standard Green's function method, and the generated field, far from the source, can be written as:

$$\mathbf{E}_4 = \frac{(i)^l}{4\pi r}\omega_0^2 \xi E_1 E_2 E_3^* \mathbf{\Lambda}_l(\theta) e^{i(k_0 r - \omega_0 t - l\phi)} \tag{3.6.13}$$

where:

$$\begin{aligned}\mathbf{\Lambda}_l(\theta) =& 2\pi(\sin(\alpha)A_l\hat{\mathbf{z}} - \cos(\alpha)\frac{A_{l+1} + A_{l-1}}{2}\hat{\rho} \\ &- i\cos(\alpha)\frac{A_{l+1} - A_{l-1}}{2}\hat{\phi})\end{aligned} \tag{3.6.14}$$

$$A_l = \int\int_{V'} e^{-i(\kappa + k_0\cos(\theta))} J_{-l}(k_0\sin(\alpha)\rho') J_l(\beta\rho')\rho' d\rho' dz' \tag{3.6.15}$$

and $r \in \mathbb{R}^+$ is the spherical radial distance from the origin of the finite interaction volume V', whose extent will depend on the size of the focal spots of the three beams being used in an experimental setup. The interaction volume is artificially limited in this calculation because the beams used in the analytical derivation have infinite extent and are non-integrable over all of space, while physical beams will always have an integrable intensity envelope.

The direction of propagation of this generated wave can be seen in the angular dependence of $\mathbf{\Lambda}_l(\theta)$ in Fig.(3.3). From it one concludes that the generated wave propagates along the same axis but in the opposite direction as the initial OAM mode $\mathbf{E}_3$, which is intuitive as $\mathbf{E}_1$ and $\mathbf{E}_2$ are anti-parallel, and linear momentum is conserved. It can also be clearly seen in Eq.(3.6.13) that the generated field has the opposite azimuthal phase dependence of $\mathbf{E}_3$. This means that the OAM of the generated photons will be opposite in sign, which is to be expected as a consequence of conservation of angular momentum since the two other initial waves $\mathbf{E}_1$ and $\mathbf{E}_2$ were plane waves. One then uses Eq.(3.6.13) to calculate the generated intensity $I_4^{\mathrm{SI}} = (1/2)c\epsilon_0|E_4^{\mathrm{SI}}|^2$ which can be integrated over the surface of the interaction volume to calculate the number of generated photons for specific initial laser parameters.

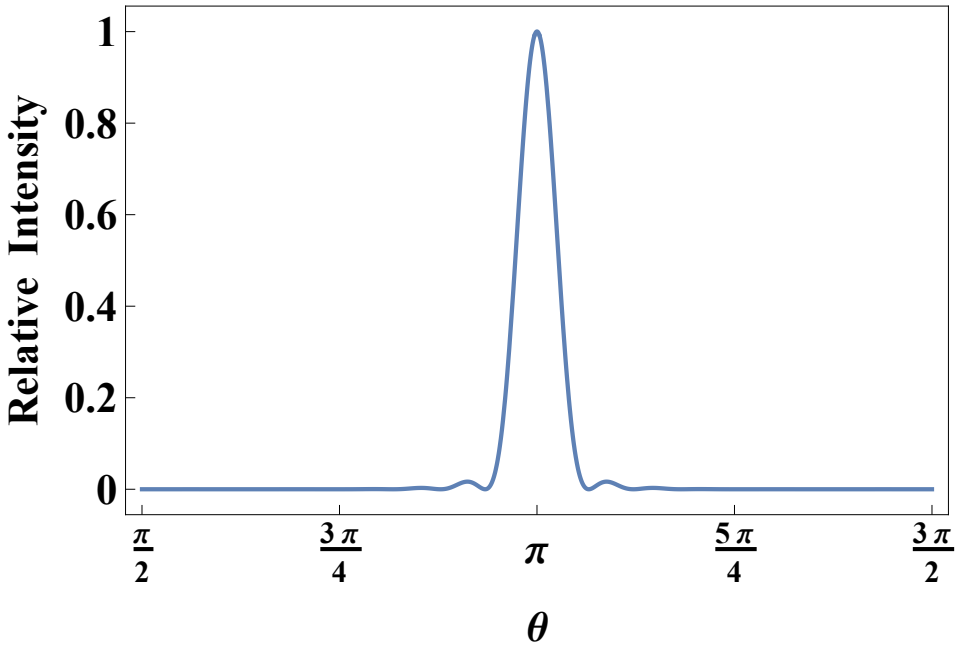

Figure 3.3: Plot of the angular dependence of the field of the generated photons $|\mathbf{E}_4|^2$ to show the direction of propagation of the generated wave. θ is the polar angle defined, as it usually is in the spherical coordinate system, from the z-axis.

3.7 Proposed experimental setup

In the previous section, Eq.(3.6.13) shows that for a specific set of three-beam geometries, if one of the lasers is carrying some OAM (l), then the photons generated due to photon-photon scattering will carry the opposite value $l_4 = -l$ as can be seen from the azimuthal dependence in the phase of the generated field. One can use this additional signature to design a high-power laser experiment to verify this hitherto undetected phenomenon. In order to make the generated photons even easier to distinguish, the geometry will be modified slightly from the one considered previously. Consider instead three initial waves, where $(\mathbf{k}_1 = -2(\omega_0)(1/2\hat{\mathbf{z}} + \sqrt{3}/2\hat{\mathbf{x}})$, $\hat{\mathbf{n}}_1 = \hat{\mathbf{y}})$ and $(\mathbf{k}_2 = -2(\omega_0)(1/2\hat{\mathbf{z}} - \sqrt{3}/2\hat{\mathbf{x}})$, $\hat{\mathbf{n}}_2 = \hat{\mathbf{y}})$ are plane waves, and $\mathbf{E}_3$ is an OAM beam with $l_3 = 1$ and $\omega_3 = \omega_0$ that is propagating along the positive z-axis. According to Eq.(3.5.3) and Eq.(3.6.13), one then expects to detect photons generated at $\omega_4 = \omega_1 + \omega_2 - \omega_3 = 3\omega_0$ with $l_4 = -1$, propagating along the negative z direction similar to the one considered previously. These generated photons are then distinct, from any of the initial radiation in the interaction, in frequency, direction of travel and OAM state. The expected number of photons generated by photon-photon scattering

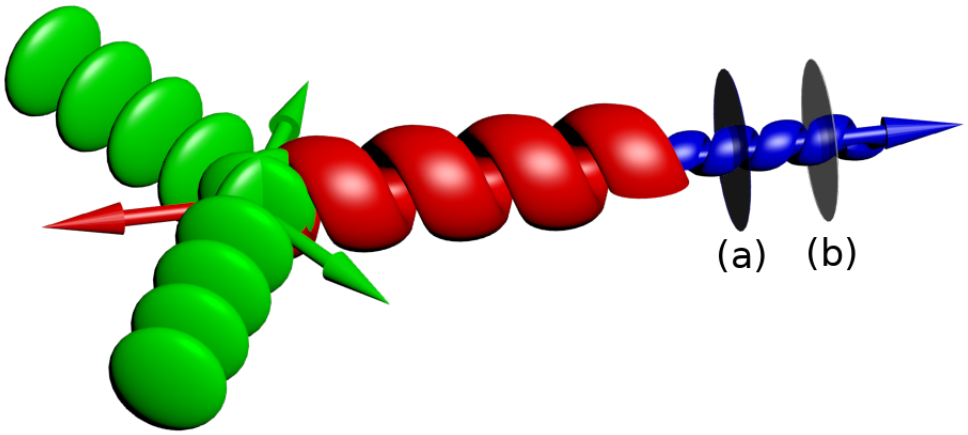

Figure 3.4: A sketch of the proposed geometry for the experiment showing the three main initial beams. The $2\omega_0$ (green) beams have a flat phase front whereas the ω_0 (red) beam has a helical front arising from the nonzero OAM. The figure illustrates the generated $3\omega_0$ (blue) spiral photons. (a) and (b) are the frequency and OAM filters respectively, which will be used to remove background noise.

N_γ can be calculated from Eq.(3.6.13) using:

$$N_\gamma^{\mathrm{SI}} = \frac{\epsilon_0 c\tau}{2\hbar\omega_4}\int_0^{\pi}\int_0^{2\pi} |E_4^{\mathrm{SI}}(r,\theta,\phi)|^2 r^2 \sin(\theta) d\theta d\phi \tag{3.7.1}$$

where τ is the duration of the interaction. In the case of a high-power laser of 10 PW at a wavelength of 800 nm, where the laser pulse duration is $\tau = 30$ fs, the beam must be split into three, where two of the component parts are frequency doubled. Considering a second harmonic conversion efficiency of 30% and a spot size of 5 µm, Eq.(3.7.1) predicts that N_γ is expected to be on the order of 100 photons per shot. It should be noted that since unphysical light modes were used for this derivation, it can only provide an initial estimate of the number of generated photons through photon-photon scattering. This is, however, not a major concern as any additional realistic envelope (such as a Gaussian envelope) will vary slowly compared to the oscillations of the field and as such will have negligible contributions to the photon-photon scattering interaction.

Although the number of photons may seem small, especially compared to the other 3 high-power beams (the main incoming beam has around 10^{20} photons), its multiple different signatures will allow for efficient filtering of any other background radiation. There have been many successful experiments, in quantum optics and computation,

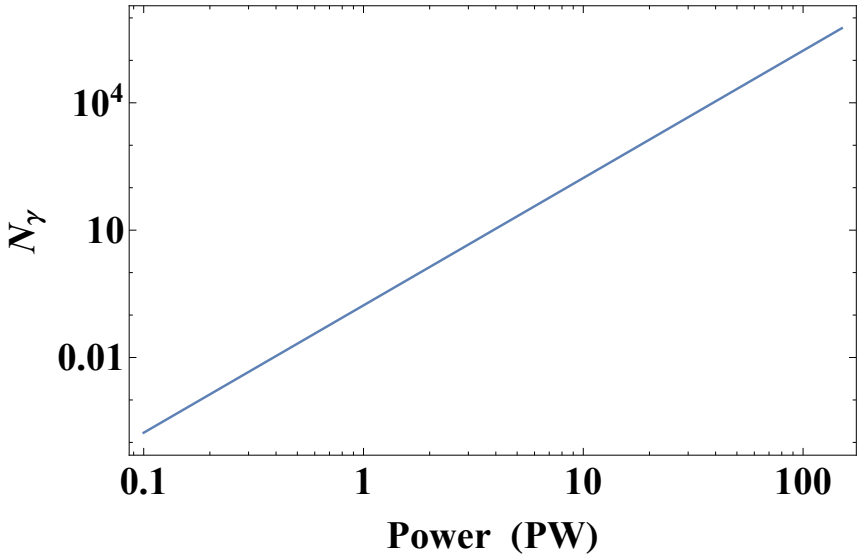

Figure 3.5: Plot showing the scaling the estimated number of generated photons as a function of initial laser power while fixing all other laser parameters (focal spot and pulse duration).

where even single photons have been filtered by their OAM state and detected using highly efficient single-photon detectors [4, 80]. Along with OAM filtering, frequency filtering with ultra narrow-band bandpass filters would block any photons whose frequency is sufficiently different from what is expected from the previously defined three-beam interaction. The single-photon detectors can also be gated to only be sensitive to scattered photons that are generated during the interaction time in order to further reduce the noise. Furthermore, performing this experiment on a high-power and high-repetition rate laser facility, such as the ELI facility [81, 82] or the Apollon facility [75], would allow for gathering the statistics essential for verification of the scattering event. The scattering signal would also be boosted by orders of magnitude by performing the experiment on the upcoming 100 PW Station of Extreme Light (SEL) laser facility as may be seen from the scaling of the predicted number of generated photons with the incoming laser power (while keeping all the other parameters constant) shown in Fig.(3.5).

The question now remains as to why one should use OAM-carrying beams in this interaction. After all, Eq.(3.5.6) can be much more easily solved with the regular plane waves and generating OAM in the lab is far from easy or ideal. The main advantage to using beams with OAM in photon-photon scattering is the additional signature that will be carried by the generated photons distinguishing them from the background radiation. This is predicted to improve the signal-to-noise ratio of

interaction making it much more detectable.

To see this, first consider the signal of the generated photons. The non-linear equation for an incoming beam carrying OAM was derived by first considering a superposition of incident plane waves with an angle α with respect to z-axis and an azimuthal phase shift and then integrating them with respect to the phase. The strength of the coupling with no OAM is derived by solving Eq.(3.6.10) for $\alpha = 0$ and no phase variance. By comparing the predicted number of generated photons through photon-photon scattering for each case (with and without OAM), the number of generated photons is practically unaffected when parameters of the incoming beams (power, duration and focal spot) are held constant. This is intuitive as the effect of OAM would only be to modify the OAM state of the generated photons and not the coupling strength of the photon-photon interaction.

Of course, high-power lasers typically do not carry OAM, which must thus be imparted onto them using mode converters such as so-called spiral phase plates (SPP) [83]. Even if the coupling strength is not affected, the conversion efficiency to the desired OAM mode is critical as a poorly designed converter would weaken the incoming beam, thus potentially undermining the positive effect of the additional signature. Efficiently generating OAM-carrying beams has been a very active topic of research due to their potential applications in quantum computation and communication, and high-density coding of information [4, 84]. In the coming chapter, I will discuss the different methods of efficient conversion of laser pulses into OAM modes specifically in the context of high-power lasers. I will show that spiral phase optics have been proven to be highly efficiency, having a theoretical maximum efficiency of up to 93%, and will further discuss experimental results of using spiral phase optics on the high-power ASTRA laser at the Central Laser Facility (CLF). This means that, by using a well-adapted mode converter, the conversion of one of the input beams into an OAM-carrying one will have a negligible effect on the number of photons generated through photon-photon scattering.

Considering now the noise, other than the background radiation that will come from any scattering of the incoming beams off the surfaces inside the experimental chamber, there is another important source of radiation noise. It is practically impossible to achieve a perfect vacuum even in a laboratory setting. This means that even within a vacuum chamber, there are still some particles that can interact with the lasers and generate some noise. The main interaction contributing to this concern is the aforementioned Compton scattering, which occurs when photons scatter off charged particles. This could be a problem because it could potentially generate

photons of a similar frequency as the ones being generated due to photon-photon scattering and in the same direction. An analysis done by Lundstrom and colleagues in [50] estimated the number of Compton-scattered photons that would be generated in the conditions of an experiment involving high-power laser facilities. They found that if the vacuum is of sufficiently high-quality, the noise due to the Compton scattering is relatively small compared to the signal of the photon-photon interaction. However, in realistic conditions, this will still cause problems as vacuum chambers in high-power systems tend to be very large and only capable of moderately high vacuums. However, as proven in [85], when photons carrying OAM Compton scatter off charged matter, the OAM state of these scattered photons will be the opposite of the OAM carried by those generated due to photon-photon scattering. In that case, the additional unique OAM signature carried by the photons, emanating from the photon-photon interaction, allows for the use of OAM filtering techniques developed in quantum optics research [86] to distinguish them from those generated due to Compton scattering. This will make the experiment to detect photon-photon scattering between high-power laser pulses more feasible. In the coming chapters, I will discuss these techniques and how they can be implemented on high power laser facilities.

3.8 Conclusion

In this chapter, I have considered the photon-photon interaction caused by the non-linear source term that is added to the vacuum electromagnetic wave equation. As I have shown, this is caused by virtual electron-positron pairs being exchanged between real photons in the vacuum. Using the Euler-Heisenberg formalism, I have shown how one can get from QED considerations to a classical non-linear wave equation with a cubic source term. Before introducing OAM to this interaction, I first solved this non-linear equation modelling the interaction between three standard plane waves to show how it predicts the generation of a fourth wave obeying momentum and energy conservation conditions.

Turning my attention back to OAM, to avoid having to use the complex functional form of the Laguerre-Gaussian modes, I have shown how another class of OAM-carrying electromagnetic modes can be constructed as a relatively simple superposition of plane wave modes. This allows for an analytical solution of the non-linear wave equation modelling the interaction between two plane waves and one OAM-carrying wave. As expected, the solution shows that for a simple planar interaction geometry,

such as the one in Fig.(3.4), angular momentum is also conserved in addition to linear momentum and energy. Thus, the introduction of OAM to photon-photon scattering provides a unique signature for the interaction that allows one to distinguish the photons generated due to this interaction from background noise, specifically that coming from Compton scattering off any residual matter in the vacuum chamber. This will hopefully allow us to finally have experimental evidence of photon-photon scattering in vacuum, confirming one of the most intriguing and long-standing predictions made by quantum electrodynamics.

Chapter 4

Generating high-intensity pulses carrying OAM

"A theory is something nobody believes, except the person who made it.
An experiment is something everyone believes, except the person who made it."

Albert Einstein

4.1 Context

Starting with the discovery of the chirped pulse amplification (CPA) by Donna Strickland and Gérard Mourou [87], laser powers have been steadily increasing with currently under-construction experimental facilities planning to generate 10 PW ultrashort laser pulses [9]. This will offer an incredible opportunity to explore non-linear and QED effects driven by high-intensity lasers in plasma and in the vacuum.

The vast majority of laser sources, including all high-power sources, generate modes with a relatively planar wavefront. As such, in order to be able to experimentally verify the effects of OAM in high-intensity laser interactions, one has to be able to generate OAM-carrying high-power laser modes or convert the standard ones to OAM-carrying modes. Generally, high-power lasers have a large diameter in order to limit their local intensity. This has the dual purpose of making it possible to manipulate them with optical components without risking damage to the optical surfaces and avoiding any unwanted non-linear effects which could damage the wavefront of the beam and its spectrum. However, this large diameter makes it challenging to design and fabricate OAM mode converters for these beams.

Additionally, it is vitally important to be able to completely characterise and measure the OAM carried by both intense light and single photons in order to be able to perform the vacuum photon-photon scattering experiment described in the previous chapter. This can prove challenging since it is not immediately clear how one would go about measuring the wavefront of a photon.

This chapter will showcase multiple methods for generating lasers carrying OAM while explaining the theoretical basis of each method and highlighting their advantages and disadvantages as well as their suitability for different regimes. It will also show experimental results for the conversion method that is most relevant for high-intensity laser-matter interactions - data which was collected during an experimental campaign on the ASTRA laser at the Central Laser Facility (CLF).

4.2 Chirped Pulse Amplification

Since the first laser was built by Theodore Maiman [88], laser technology has developed to become an integral part of many scientific investigations and applications ranging across multiple diverse domains from astronomy to medical sciences. The general operating principles for lasers were theoretically outlined by multiple research groups working in parallel in the 1950s [89, 90]. They are based on the concept of stimulated emission of electromagnetic radiation, first outlined by Albert Einstein [91]. The vast majority of lasers consist of a gain medium in which radiation is coherently amplified via stimulated emission, a pumping mechanism to provide the energy for amplification, and a cavity that allows the radiation to perform multiple pass-throughs in the gain medium for further amplification.

Scientists have been constantly seeking to achieve higher and higher intensities of laser radiation in order to investigate multiple interesting phenomena. This includes non-linear effects, such as self-focusing or high-harmonic generation, which occur when light polarizes the material with which it is interacting, i.e. when the electric field of the laser radiation becomes comparable to the atomic binding field [31]. This occurs when the intensity of the laser becomes around $I_{\text{laser}} \sim 10^9\,\text{W/cm}^2$. Beyond that, lasers can start ionizing the material, i.e. stripping the atoms of electrons, forming a plasma and allowing for the study of laser-plasma phenomena. With even higher intensities ($I_{\text{laser}} \sim 10^{18}\,\text{W/cm}^2$), the speed of electron motion driven by laser light in the aforementioned plasma gets closer to the speed of light allowing for the investigation of relativistic effects. Finally, ultra-intense lasers near the Schwinger limit (corresponding to $I_{\text{laser}} \sim 10^{29}\,\text{W/cm}^2$) can start driving non-linear effects in

the vacuum and generating electron-positron pairs [92], allowing for the investigation of strong field QED effects.

After its initial inception, laser technology progressed gradually with techniques such as Q-switching and mode locking, allowing for the generation of intense short-pulses [93]. This allowed for the investigation of non-linear effects. Unfortunately, and ironically, it is these very effects that also made it difficult to progress further and generate lasers with even higher-intensities using the standard methods. Once laser intensities start getting near the threshold for dielectric non-linear effects, it becomes extremely difficult to amplify them further in the same gain media since, by definition, they can now trigger these damaging effects in those media. To compound this problem, non-linear self-focusing could occur, effectively transforming the material into a lens and increasing the peak intensity of the laser even more and potentially triggering ionization effects and destroying the gain medium. One way to circumvent these issues was to increase the diameter of the beam and all the relevant optical components in order to reduce the peak intensity. Once sufficiently amplified, the beam could then be refocused to get a high-intensity focal spot. However, these obstacles had effectively put a limit on achievable laser intensities since they meant that laser diameters needed to be extremely large, making the required gain media prohibitively expensive and impractical. They also made the development of high-intensity conventional lasers impossible. This led to a stagnation in the peak achievable intensity for decades.

A breakthrough came when Donna Strickland and Gerard Mourou invented a technique that was termed chirped pulse amplification [87]. The idea behind CPA was deceptively simply. Instead of stretching the pulse in space and increasing the size of the system, the pulse would be stretched in time. This is done by chirping the pulse, i.e. making it so that the higher frequency components of the pulse lag behind the lower ones or vice versa. The stretching lowers the peak intensity of the pulse allowing it to be amplified in a standard amplification medium. The pulse is then compressed back to its original duration. This process allows for amplification by orders of magnitude more than what can be achieved without it.

The first experimental setup using CPA exploited the group velocity dispersion inherent to fibre optic cables to chirp and stretch the pulse, which was then subsequently compressed using a grating compressor. This achieved a 1.5 ps pulse with a power of 0.5 GW. The early attempts at CPA were constantly plagued with the issue of matching the stretching and compressing of the pulse, which were driven by the

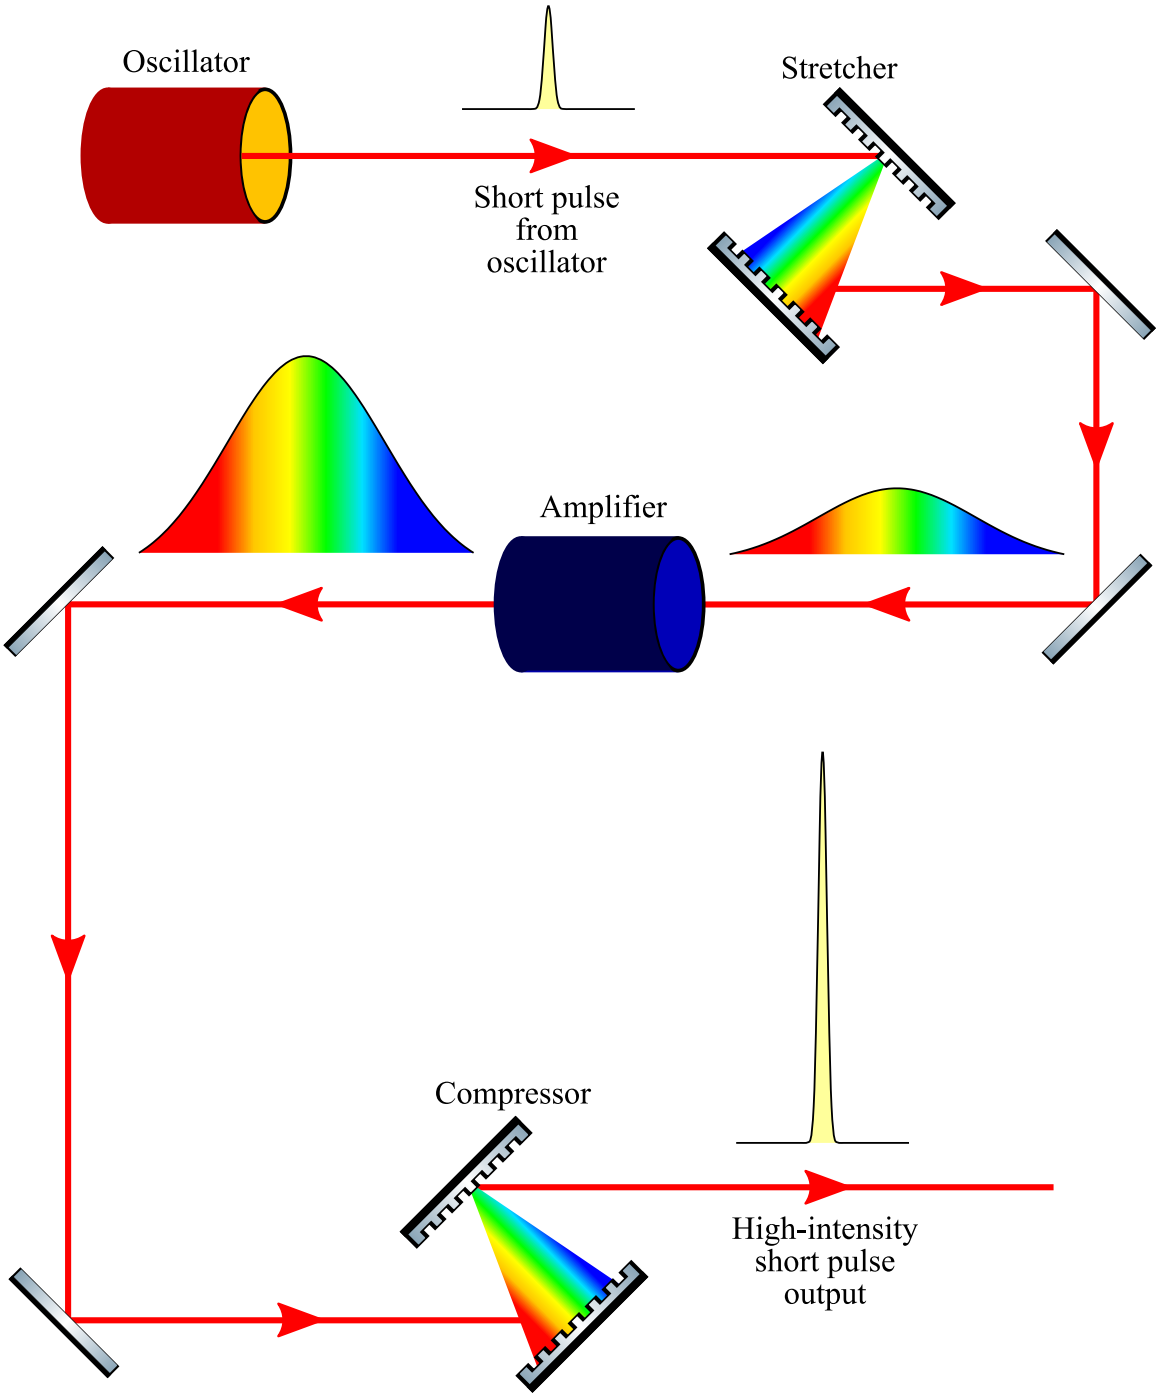

Figure 4.1: Schematic diagram describing the basic CPA procedure using a grating pair as its stretcher and compressor.

fibre cable and grating compressor respectively. This was solved when Oscar Martinez discovered that a telescope system placed within a double grating setup would modify the phase shift experienced by the different wavelengths in a way that would invert the chirp caused by the system [94]. This was leveraged to create a grating *stretcher*, the dispersion of which can be much more easily matched with a compressor for optimal CPA.

In the years since the invention of CPA, technological advancements have helped optimise the CPA procedure to produce extremely intense pulses. This included better seed pulse generation, higher compression ratios and the use of optical parametric amplification (OPA) instead of traditional gain media, which has allowed for the generation of shorter pulses at a higher repetition rate and with more energy [95]. As of today, even more than 30 years after its first use, CPA remains the most used method of generating high-power pulses by far. CPA powers facilities such as the central laser facility (CLF) at the Rutherford Appleton laboratory (RAL), the OMEGA EP laser at the laboratory for laser energetics (LLE) and even the soon-to-be-completed extreme light infrastructure (ELI) project which is aiming to achieve a 10 PW. Even more impressively, the basic scheme for CPA has not been significantly altered over all these years. As can be seen in Fig.(4.1), the CPA procedure starts with an oscillator producing a short pulse, which then goes through a grating stretcher to be chirped and subsequently amplified. The amplified pulse is then re-compressed, producing an extremely high-intensity short pulse which is ready to be used.

4.3 The ASTRA GEMINI laser

The ASTRA GEMINI facility is one of the flagship high-power laser facilities in the United Kingdom [96]. Located at the Central Laser Facility in the Rutherford Appleton Laboratory, it is a Ti:Sapphire (TiS) short pulse high-intensity laser operating at a 800 nm wavelength. The full laser is capable of delivering two beams each carrying 15 J of energy achieved using CPA. Although this is a relatively low energy beam, it can be compressed to an ultra-short pulse length of 30 fs, producing a 0.5 PW pulse. At an optimal, diffraction-limited focal spot, this corresponds to a theoretical peak intensity of around $3 \times 10^{21}\,\mathrm{W/cm^2}$. This kind of short ultra-intense pulse is ideal for many scientific investigations including, but not limited to, laser particle acceleration, higher-harmonic generation and X-ray generation.

The facility is divided into two main target areas. There is target area 3 (TA3) where the two pulses, as they were described above, are available. The other is target

(a)

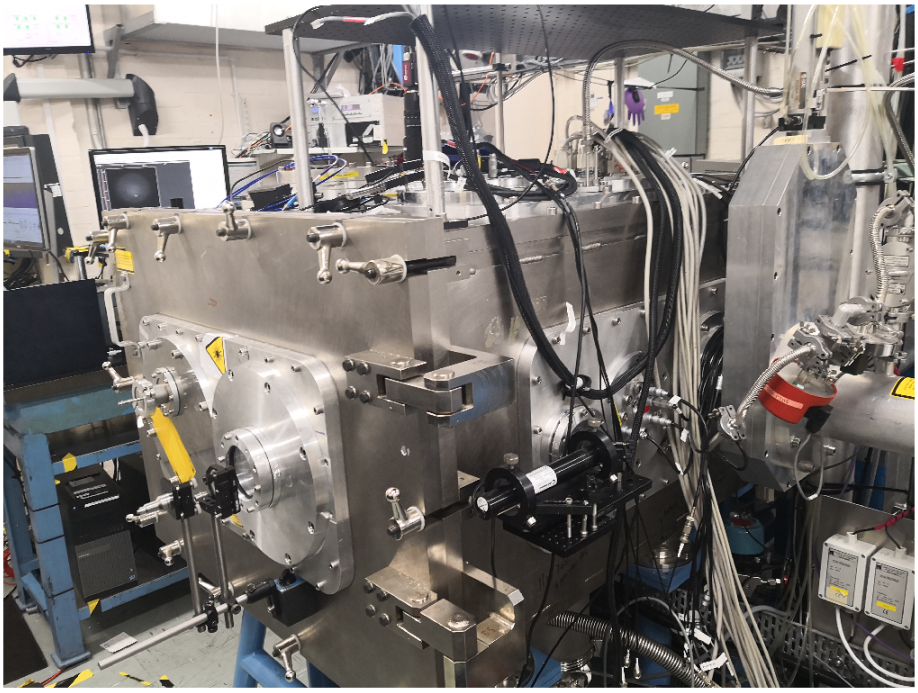

(b)

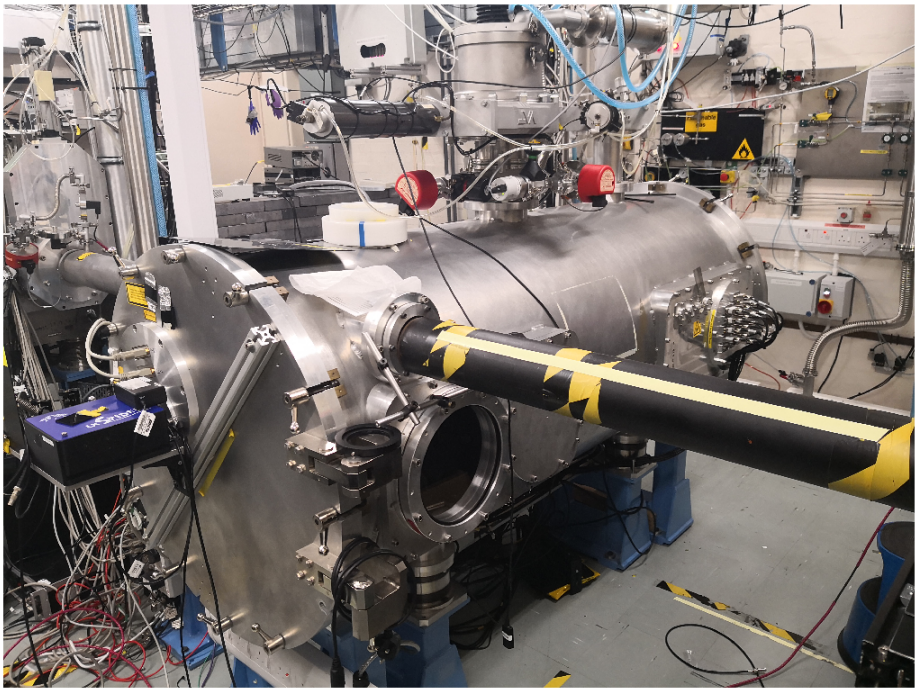

Figure 4.2: (a) The TA2 vacuum chamber where the high-intensity interactions take place (b) The compressor which takes the stretched and amplified pulse and compresses it before it is directed towards the interaction chamber.

area 2 (TA2) where only one, not fully amplified pulse is available (TA1 is not currently active). While this pulse has all the pulse length and wavelength characteristics stated earlier, it only carries 600 mJ. Fig.(4.2) shows the experimental vacuum chamber and the compressor of TA2. At the exit of this compressor, the laser pulse has a diameter of 60 mm. It is in TA2 where we carried out an experimental campaign to investigate the optimal methods for generating and characterising the OAM of high-power pulses. All experimental results showcased in this chapter and the following one were gathered from that campaign. For our purposes, the pulse length was set at 50 fs at the exit of the compressor. To maintain the terminology preferred by the CLF, the TA2 pulse shall be referred to as the ASTRA beam.

4.4 Lasers carrying orbital angular momentum

Generating beams which carry angular momentum is not a straightforward task. Although the previously defined LG modes form a complete set of solutions to the electromagnetic wave equation, creating a laser which naturally operates at a higher LG mode is quite difficult. One of the main reasons behind is that, as seen before, different OAM modes will modify the way light interacts with matter, thereby changing how the light will be amplified through a particular gain medium. Although there has been some progress towards creating a laser system which naturally produces OAM light, especially with solid state lasers [97], it remains an impractical option especially if one wants to control the amount of OAM generated.

As such, methods for converting other laser radiation into OAM modes have been under development ever since the initial characterisation of LG beams. The simplest way to evaluate the conversion efficiency of a particular method is to define an *inner product* on the space V of functions $u \in V : \mathbb{R}^2 \to \mathbb{R}$, describing the transverse profile of a beam. Consider a method that is described by an operator $\mathcal{T} : V \to V$, the inner product describing the conversion efficiency into a particular Laguerre-Gaussian mode at a particular plane $z = z_0$ can be written as

$$\langle \mathcal{T}(u), \mathrm{LG}_p^l \rangle = \iint \mathcal{T}(u)(x,y)\mathrm{LG}_p^l(x,y)^* dxdy \tag{4.4.1}$$

where $\mathrm{LG}_p^l(x,y)$ are the functions describing the Laguerre-Gaussian modes, as previously defined in Eq.(2.4.25), evaluated at the aforementioned specific plane:

$$\mathrm{LG}_p^l(x,y) = E_p^l(\mathbf{r},0)\Big|_{z=z_0} \tag{4.4.2}$$

Depending on the form of $\mathcal{T}(u)$, it may be easier to switch to circular polar coordinates and in that case, the area element will transform to $\rho d\rho d\phi$.

4.4.1 Conversion of Hermite-Gaussian modes

One of the earliest methods for generating LG modes consisted of converting the so-called Hermite-Gaussian (HG) modes. HG are also a complete orthogonal family of solutions to the paraxial wave equation in Cartesian coordinates that can expressed as [31]

$$\begin{aligned} \mathrm{E}_{n,m}(\mathbf{r},t) = \frac{w_0}{w(z)} H_n\left(\frac{\sqrt{2}x}{w(z)}\right) H_m\left(\frac{\sqrt{2}y}{w(z)}\right) \exp\left(-\frac{\rho^2}{w^2(z)}\right) \\ \times \exp\left(i\Big(kz - \omega t + \frac{k\rho^2}{2R(z)} + \psi_{n,m}(z)\Big)\right) \end{aligned} \tag{4.4.3}$$

where $H_n(x)$ is the "physicist's form" of the Hermite polynomial of order n. The radius of curvature $R(z)$ and the beam width $w(z)$ are the same as in Eq.(2.4.25) and the Gouy phase shift is only slightly modified to be

$$\psi_{n,m} = (n+m+1)\arctan\left(\frac{2z}{kw_0^2}\right) \tag{4.4.4}$$

As was done above with $\mathrm{LG}_p^l(x,y)$, we can define at a particular plane

$$\mathrm{HG}_{n,m}(x,y) = E_{n,m}(\mathbf{r},0)\Big|_{z=z_0} \tag{4.4.5}$$

For both the Hermite- and Laguerre-Gaussian modes, it is useful to define the N-number which can be expressed using the mode numbers as $N = n + m = |l| + 2p$. These HG transverse field distributions can be used to describe the resonant modes of an optical cavity such as the ones used in lasers. As such, it is possible to describe the output of a laser in terms of these modes and, depending on how the cavity is constructed, it is possible to have lasers which produce higher-order HG modes.

Since the HG modes are a complete and orthogonal set of solutions, any $\mathrm{LG}_p^l(x,y)$ profile can be expressed a sum of various $\mathrm{HG}_{n,m}(x,y)$ ones:

$$\mathrm{LG}_p^l(x,y) = \sum_{n,m} \langle \mathrm{HG}_{n,m}, \mathrm{LG}_p^l \rangle \, \mathrm{HG}_{n,m}(x,y) \tag{4.4.6}$$

A closed form solution for this change of basis can be readily found in the literature of basis modes and so Eq.(4.4.6) can be rewritten as [98]

$$\langle \mathrm{HG}_{n,m}, \mathrm{LG}_p^l \rangle = \begin{cases} i^m f\left(\frac{N+l}{2}, \frac{N-l}{2}, m\right), & \text{if } n+m = |l|+2p \\ 0, & \text{else} \end{cases} \tag{4.4.7}$$

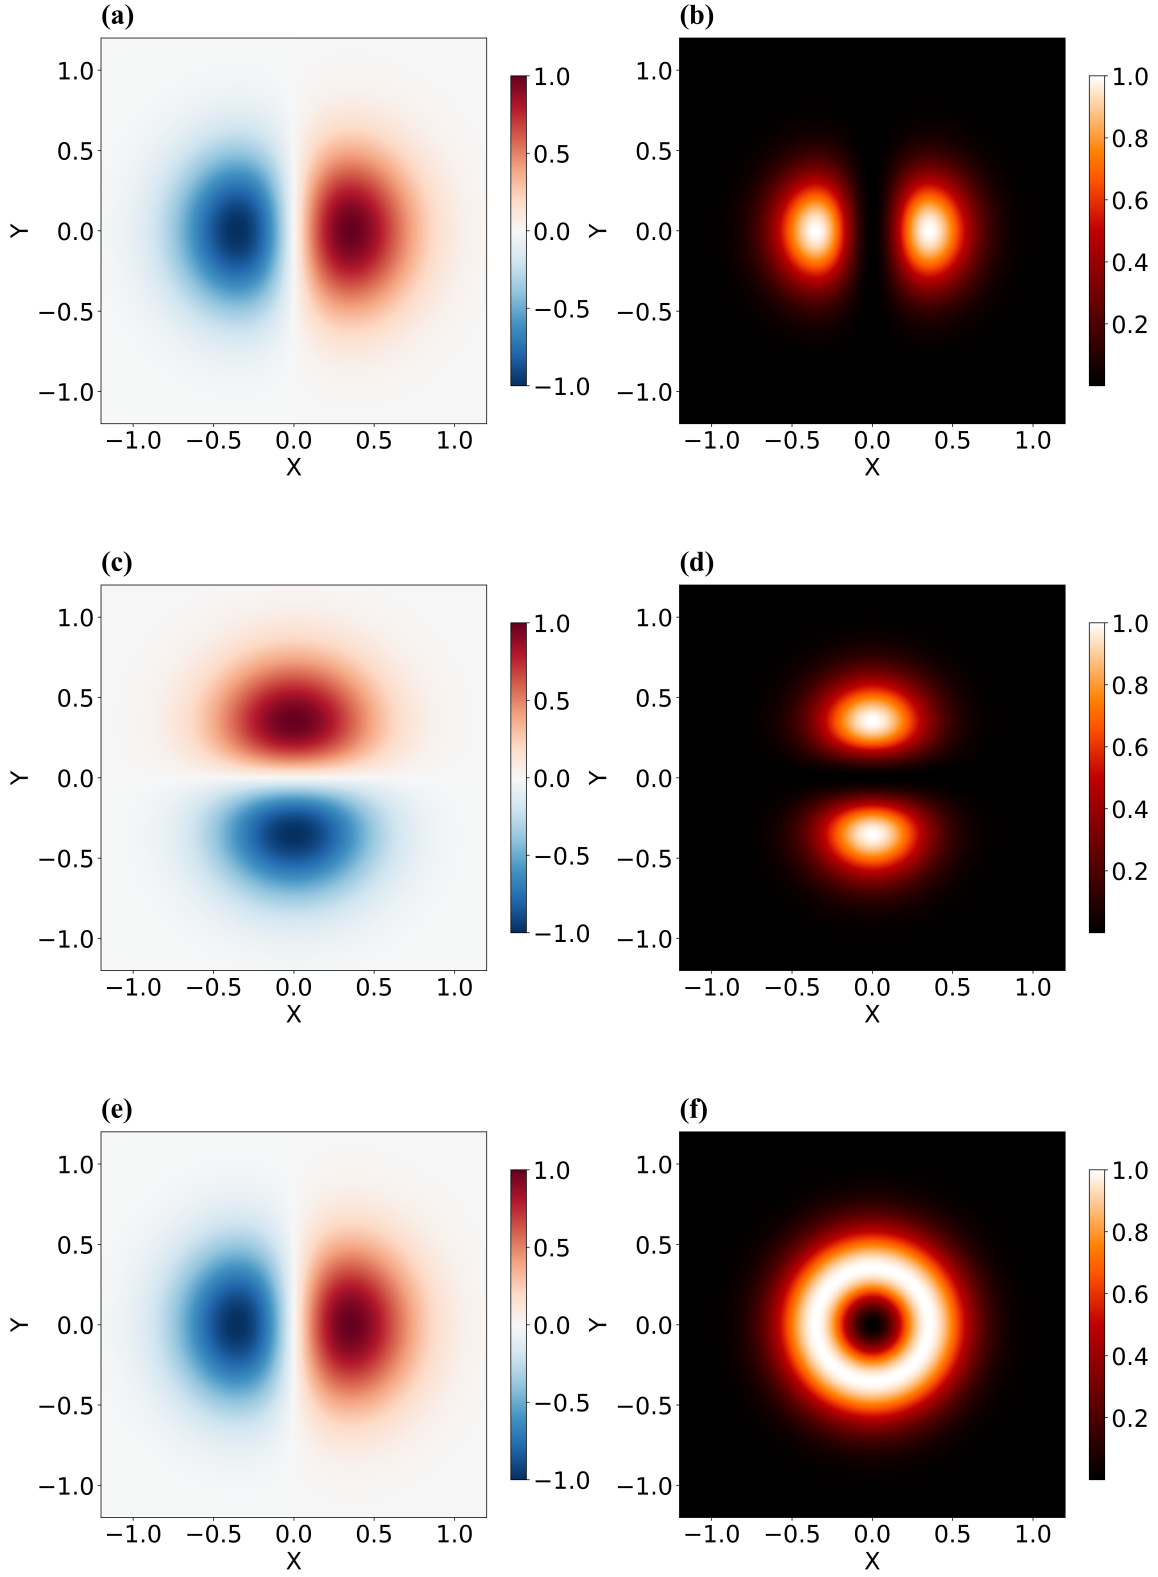

Figure 4.3: Simulation of creating of a LG_0^1 modes using the two Hermite-Gaussian modes $HG_{1,0}$ and $HG_{0,1}$. (a)-(b) Show the field and intensity of $HG_{1,0}$, while (c)-(d) show those of $HG_{0,2}$ and (e)-(f) show those of $(1/\sqrt{2})(HG_{1,0} - iHG_{0,1})$.

where N is the previously defined number shared by the HG and LG modes in the case where the inner product is finite and

$$f(n,l,m) = \sqrt{\frac{(n+l-m)!m!}{2^{n+l}n!l!}} \frac{1}{m!} \frac{d^l}{dt^l} \Big[(1-t)^n (1+t)^l\Big]\Big|_{t=0} \tag{4.4.8}$$

As can be seen in Fig.(4.3), the lowest order OAM-carrying mode LG_0^1 can be written in terms of the first two Hermite modes $\mathrm{HG}_{1,0}$ and $\mathrm{HG}_{0,1}$:

$$\mathrm{LG}_0^1 = \frac{1}{\sqrt{2}}(\mathrm{HG}_{1,0} - i\mathrm{HG}_{0,1}) \tag{4.4.9}$$

This then allows for a relatively easy generation of OAM modes as long as one has access to HG beams. In fact, it was shown that just by using cylindrical lenses, one can readily convert HG modes into LG ones [99]. If the lens was of high quality, then this conversion would be highly efficient, since it relied on the mathematical relationship between the two mode families. However, this method has severe limitations. Firstly, it is still quite difficult to tune the desired OAM mode since this requires having access to the appropriate HG laser. Additionally, if one wants to generate higher-order LG modes, many HG ones (on the order of $O(l)$) would be needed. Finally, and much more critically, modern lasers are made in such a way so as to avoid producing HG radiation, thus making this method unviable, particularly for high-intensity lasers. What I am looking for is a method that is somewhat tunable and that can take as an input the types of modes that are currently being generated by high-power sources.

4.4.2 Computer generated holograms

A hologram is an optical element which leverages the diffractive properties of light to recreate a particular image [100]. Unlike a regular photographic image which only contains information about the intensity of light, a hologram also holds information about its phase. As can be seen in Fig.(4.4), the basic principle of holographic imaging is using a reference laser beam and interfering it with light containing information about a particular image, encoding that image onto a holographic plate. Later, when a laser similar to the encoding one is shone on the holographic plate at the same angle, light scatters off, generating a virtual image similar to the one that was encoded. The fact that there is an angular dependence to the reconstruction of the image has been exploited to create a stereoscopic 3D effect using holograms and is indeed the driver behind most current major commercial uses of holograms.

In order to gain a bit more insight on how holographic techniques could be used to generate LG modes, let us consider the transverse profile of a generic electromagnetic

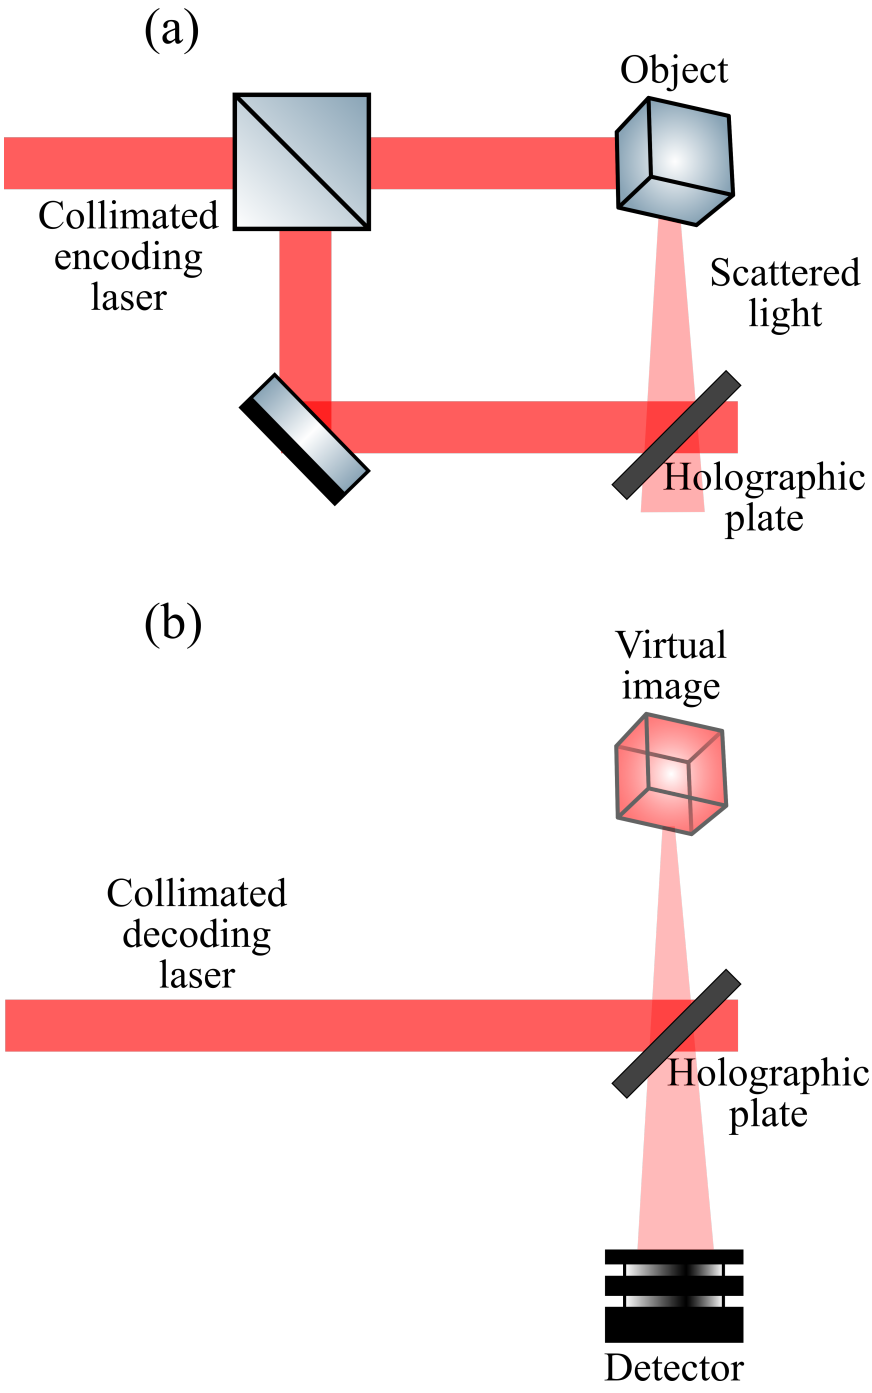

Figure 4.4: Diagram of a typical holography setup: (a) information on the shape of the object is recorded onto the plate via the interference between the light scattered off of it and a reference laser. (b) When the holographic plate is illuminated with the same reference laser, a virtual image is then generated.

wave $E_i(x, y)$, as well as a particular reference plane wave $E_0(x, y)$, which intersects with the former with an angle θ. At the plane (x, y), which is transverse to the direction of propagation of E_i and placed where the two waves intersect, the profile of each beam can be written as

$$E_i(x, y) = \sqrt{I_i(x, y)} \exp(i\Phi(x, y)) \qquad E_0 = \sqrt{I_0} \exp(ik \sin(\theta)x) \tag{4.4.10}$$

where $k = 2\pi/\lambda_0$ is the wavenumber of the reference plane wave. The superposition of the two beams creates an interference pattern at this intersection plane where the intensity can be written as

$$|E_i(x, y) + E_0(x, y)|^2 = I_i(x, y) + I_0(x, y) + 2\sqrt{I_i I_0} \cos(k \sin(\theta)x - \Phi(x, y)) \tag{4.4.11}$$

This interference pattern is then used to construct a hologram by recording it in an optical plate. This plate acts as a converter on any input radiation $E(x, y)$ with a conversion operator that can be expressed as multiplication by a particular transmission function $\mathcal{T}(E) = T_H(x, y)E(x, y)$. Broadly speaking, there are two major sub-types of holographic plates: ones where T_H modifies the amplitude of a beam going through it (amplitude modulation), and ones where T_H only modifies the phase (phase modulation). When the reference laser E_0 is again shone on this plate, the diffraction of light off the plane of the plate reconstructs the recorded beam E_i in the far-field similar to what can be seen in Fig.(4.4-b).

For our particular purposes, we are only interested in recording the phase information contained in the beam we wish to reconstruct, since this is what contains information about the OAM state. Hence, we can discard the intensity information I_i, creating a so-called *phase-only* hologram [101]. In a phase modulation setup, discarding constant phase terms, the transmission function of such a hologram is written as

$$T_H(x, y) = \exp(i\gamma \cos(k \sin(\theta)x - \Phi(x, y))) \tag{4.4.12}$$

where γ is the amplitude of the phase modulations. To keep the terminology concise, I shall henceforth refer to the aforementioned holograms simply as phase holograms.

To use this technique to generate OAM beams, it is of course not possible to try interfering a reference mode with the desired LG mode and recording the resultant pattern, like in Fig.(4.4), since this would require having access to said LG mode. Instead, one must create a *computer-generated* hologram where this process is simulated and a pattern is generated based on assumptions on the reference and desired

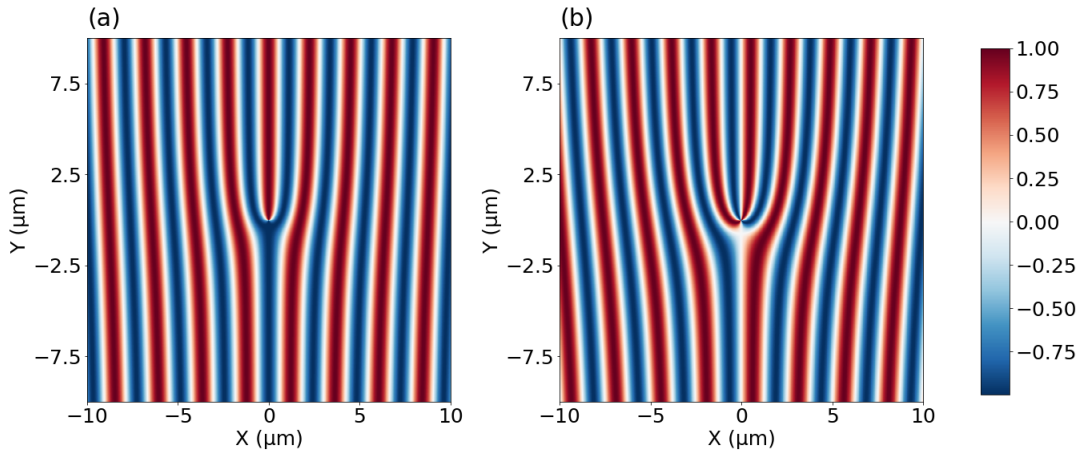

Figure 4.5: Simulated phase holograms for (a) $l = 1$ and (b) $l = 2$ with an 800 nm reference beam and $\theta = \pi/9$.

mode. Since the goal is to make a phase hologram, only the phase information of an LG_0^l mode is required, which can be expressed in polar coordinates as $\exp(il\phi)$. This then allows us to rewrite Eq.(4.4.12) as

$$T_H(x, y) = \exp(i\gamma \cos(k \sin(\theta)x - l\phi)) \tag{4.4.13}$$

Fig.(4.5) shows a simulated phase hologram designed for a 800 nm laser. The hologram shows the distinctive "fork" in the centre which is caused by the phase singularity. Away from the fork, the hologram looks like a transmissive diffraction grating. In fact, as mentioned earlier, its action is quite similar to diffraction gratings. When 800 nm light passes, directed along the normal, through the hologram several diffraction orders are produced. The zeroth order contains the unmodified wavefront, while the first order, produced at the angle θ, contains the LG_0^l wavefront. In fact, a diffraction grating can be thought of as a phase hologram simply reconstructing the same reference beam but with the propagation along an angle θ, with different wavelengths subtending different angles.

Unfortunately, it is quite difficult to fabricate phase holograms where the phase acquired varies sinusoidally such as those in Fig.(4.4.13). In practice, they are replaced by binary phase holograms where the variance in phase is step-wise. For these, the phase function $\cos(k \sin(\theta)x - l\phi)$ is replaced with

$$T_H(x, y) = \exp(i\gamma \mathrm{square}(k \sin(\theta)x - l\phi)) \tag{4.4.14}$$

where square(x) = sgn$(cos(x))$ is the square wave function. Another problem with this type of hologram is that it is quite inefficient. Similar to a diffraction grating, most of the energy remains in the zeroth order with only very little of it going into the desired first order.

Thankfully, this problem can be remedied by taking a cue from diffraction gratings and making the hologram *blazed* to the first order. This means that it is optimised to maximise the energy deposited in the chosen order. The transmission function of such a hologram can be written as [102]

$$T_H(x,y) = \exp(i\gamma \text{Mod}(k\sin(\theta)x - l\phi, 2\pi)) \tag{4.4.15}$$

where $\text{Mod}(x,y) = x - y\lfloor(x/y)\rfloor$ is the modulus function where $\lfloor x \rfloor$ is the floor function which outputs the largest integer smaller than its input. While these blazed holograms are more difficult to fabricate than the binary holograms, especially due to the fork at the centre, they have a theoretical 100% energy deposition in the first order.

To model these phase holograms, their action on input modes can be readily modelled using diffractive optics simulations. Even more simply, since we are interested in the far field limit, the effect of phase holograms can be modelled using Fraunhofer diffraction. This allows us to write the intensity in the far field as the magnitude of the Fourier transform of the field in the plane of the hologram $T_H(x,y)E_0(x,y)$. Since the most common transverse mode output for lasers is the Gaussian mode, this is generally what is taken to be the input which can be then written, in polar coordinates, as

$$E_0(\rho,\phi) = \sqrt{\frac{2}{w_0^2\pi}}\exp(-\frac{\rho^2}{w_0^2}) \tag{4.4.16}$$

where w_0 is the waist of the beam in the focal plane and the pre-factor is added to normalise the mode.

As can be seen in Fig.(4.6), by simulating the action of these holograms using the above-mentioned method, the difference between the binary and blazed hologram becomes clear. In the far field of the binary hologram, the various diffraction orders can be seen with the doughnut shape corresponding to the LG_0^1 mode visible in the ± 1 orders. However, as mentioned previously, most of the energy of the incoming laser is deposited into the unchanged zeroth order. On the other hand, the far field of the blazed hologram shows that the entirety of the energy has been deposited into the desired first order of diffraction.

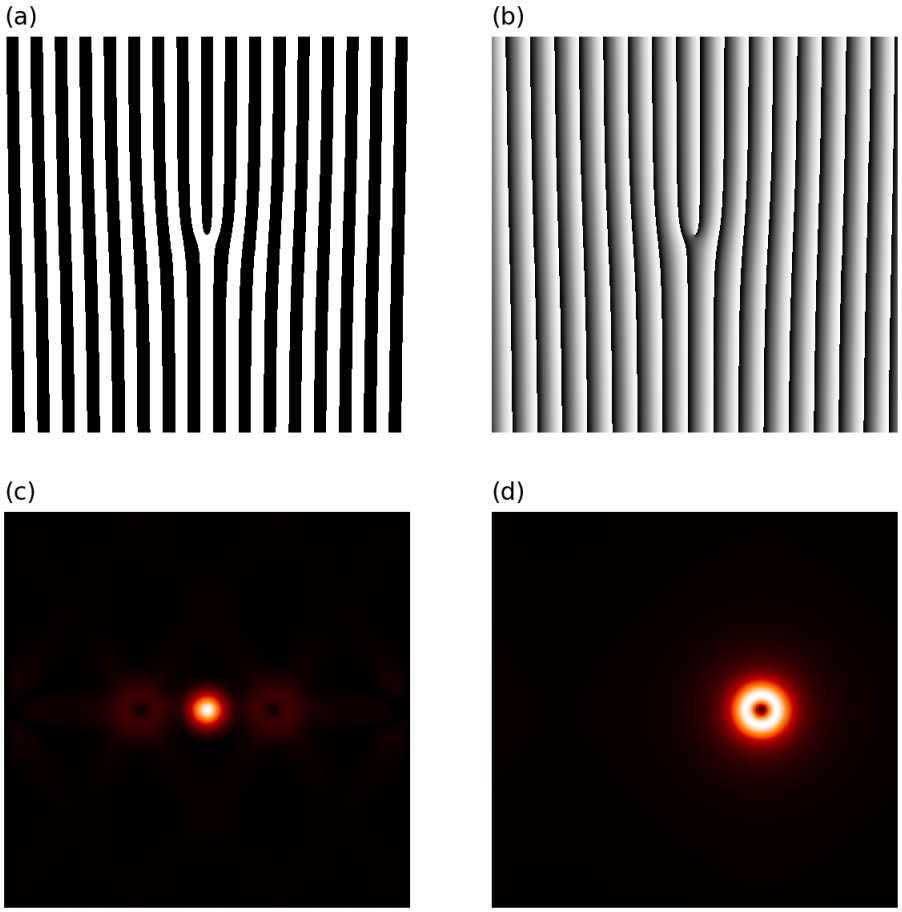

Figure 4.6: Simulated results comparing the action of a (a) binary phase hologram versus a (b) blazed hologram. Both holograms are generated to produce a state with $l = 1$ as can be seen from the form of the phase singularity at the centre. (c) and (d) show the far-field image produced by these holograms when a Gaussian laser, as in Eq.(4.4.16) is shown through the binary and blazed hologram respectively.

Phase holograms were some of the earliest methods used to produce OAM-carrying modes. They have the advantage of relying on the same input beam no matter what OAM state one wants to achieve. This input beam can be a simple Gaussian mode, which is already the standard output of a vast majority of lasers. They are also quite tunable. By using the same technique and simply modifying the hologram, one can generate any arbitrary OAM state, including LG modes with a non-zero radial mode number $p \neq 0$ [103]. These advantages made them the go-to method for many use cases. Over the decades since the first introduction of phase holograms, better fabrication methods meant that higher efficiencies were possible, with modern phase holograms capable of over 80% conversion efficiency into the desired mode [86].

All of this, however, does not mean that holograms are suited for all purposes. Although they can achieve very high conversion efficiencies, blazed phase holograms are very difficult to fabricate. As can be seen in Fig.(4.5), for an 800 nm, the detail in the structure is extremely small and requires a very high precision fabrication process. This may not have been a large impediment for designing these holograms for low-power lasers, such as the ones used in quantum optics, where the beam diameter is generally less than a 5 mm. Unfortunately, this does become an issue when trying to make a hologram suited for use in high-power laser systems. This is because, as mentioned previously, they generally have a large beam diameter $\geq$50 mm to reduce their peak intensity and avoid any non-linear effects. This means that any hologram that would be suited for these purposes would also need to be at least as big as the beam itself, making its fabrication prohibitively difficult. Furthermore, even if it were possible to make a hologram of the appropriate size, it would still not be suited for PW facilities. This is because, even with the large beam diameters, PW pulses are intense enough to be distorted when passing through any transmissive optical component and are generally only handled in vacuum and with reflective optics to preserve their spectrum and pulse length. All of this means that even though they can be highly efficient, phase holograms do not provide a reliable practical method for producing high-power beams carrying OAM.

4.5 Spiral phase optics

A relatively simpler alternative is to turn to what are called spiral phase optics (SPO). These are optical components that impart an azimuthally varying phase on the wavefront of any incoming beam either through the use of a transmissive spiral phase plate (SPP) or a reflective spiral phase mirror (SPM). Consider for example the

Figure 4.7: Model of a typical transmissive spiral phase plate optimized for converting an input Gaussian to the LG_0^l mode.

model shown in Fig.(4.7) of a typical continuous transmissive SPP. When a beam with a planar phasefront passes through the SPP, it acquires different relative phase shifts depending on the azimuthal angle due to the azimuthally-varying thickness of the plate. Since we want the beam to acquire a relative phase factor $\propto e^{il\phi}$ after passing through the SPP, the thickness of the plate as a function of the azimuthal angle can be mathematically modelled as:

$$h(\phi) = \frac{Q\lambda\phi}{2\pi(n_{\text{plate}} - n_{\text{ext}})} + h_0 \tag{4.5.1}$$

where λ is the wavelength of the light in question, Q is the topological "charge" of the phase discontinuity of the plate, h_0 is the thickness of the base of the SPP before the helical section, n_{plate} is the refractive index of the plate and n_{ext} is the refractive index of the external medium in which the plate is located. Generally, these plates are operated in air, as such the latter is taken to be $n_{\text{ext}} \approx 1$. For a spiral phase mirror, although their end goal is the same, the operating principle is slightly different. Instead of a varying thickness of some dielectric medium causing a varying phase shift for propagating light, spiral phase mirrors are designed to have an azimuthally-varying height from a particular base. That way, when light beams reflects off these mirrors, different parts of the beam travel paths of varying lengths to reach the reflective surface, thus acquiring different phases. In all cases, the most

significant parameter, which defines a particular spiral optic, is the height of the step discontinuity of the plate, which, in the case of SPPs is simply

$$h_{\text{step}} = h(2\pi) - h_0 = \frac{Q\lambda}{n_{\text{plate}} - n_{\text{ext}}} \tag{4.5.2}$$

Due to the relative ease of their fabrication and use, interest in spiral phase optics has steadily increased over the past two decades [104]. In particular, their conversion efficiency has been the subject of much research, both theoretical and experimental. The action of all spiral phase optics, transmissive or reflective, can be simply modelled as a multiplication by a phase factor. For a standard continuous SPP, this can be written as

$$\mathcal{T}_{\text{SPO}}(E_0(x,y)) = \exp(-iQ\phi)E_0(x,y) \tag{4.5.3}$$

As such, evaluating the inner product which quantifies the conversion efficiency to a particular LG_p^l mode of these plates gives us

$$\left| C \iint \exp(-iQ\phi)E_0(\rho,\phi)\left(\frac{\sqrt{2}\rho}{w_{\text{LG}}}\right)^{|l|} L_p^{|l|}\left(\frac{2\rho^2}{w_{\text{LG}}^2}\right)\exp\left(-\frac{\rho^2}{w_{\text{LG}}^2}\right)\exp(il\phi)\rho d\rho d\phi \right|^2 \tag{4.5.4}$$

where C is a normalisation constant, w_{LG} is the waist of the desired LG mode and the integral is evaluated in cylindrical coordinates as they are more convenient for describing LG modes.

As mentioned previously, we are only interested in modes with a null radial mode $p = 0$. Since $L_0^{|l|}(x) = 1$, this greatly simplifies Eq.(4.4.1). When a simple Gaussian input, such as the one in Eq.(4.4.16), is considered, it is possible to find a closed form analytical expression for the conversion efficiency with [105]

$$|\left\langle \mathcal{T}_{\text{SPO}}(\text{LG}_0^0), \text{LG}_0^l \right\rangle|^2 = \frac{2^{|l|+2}\left(\frac{w_{\text{LG}}}{w_0}\right)^2 \Gamma^2(|l|/2+1)}{|l|!\left(1+\left(\frac{w_{\text{LG}}}{w_0}\right)^2\right)^{|l|+2}}\delta_l^Q \tag{4.5.5}$$

where LG_0^0 is simply another way of expressing a Gaussian profile, w_0 is the waist of the input Gaussian and $\Gamma(x)$ is the Gamma function. The Kronecker delta shows that the conversion is only non-zero when the charge of the SPP is matched with the OAM state of the desired. It is possible to extend the calculations to non-integer charges but, in that case, the efficiency would not have a simple analytical solution. The remaining parts of the equation show that the conversion efficiency is only dependent on the ratio between the wait of the desired mode and that of the input Gaussian. If it assumed that the waist of the beam is unchanged as it goes through the SPP, then

the maximum theoretical conversion efficiency is $\Gamma^2(|l|/2+1)/|l|!$, which evaluates to $\approx 77\%$ for $l = 1$.

However, there is no physical reason for this assumption to be true. In fact, there have been reports of conversion efficiencies of over 90%, exceeding the theoretical maximum for when the waists are matched [106]. A higher-order LG mode has a very different intensity profile from that of a standard Gaussian and, as such, matching the waist would not satisfy any "continuity" requirement. The conversion efficiency is measuring the decomposition of the converted beam into particular modes of radiation. Thus, there is no reason to *a priori* exclude any parameter from consideration [105]. If the waist ratio is left as a variable, then the optimal waist for the LG component of the converted beam can be found by finding the maximum of the function

$$\frac{x^2}{(1+x^2)^{|l|+2}} \tag{4.5.6}$$

This can be easily done to show that the optimal waist is ratio is

$$\frac{w_{\mathrm{LG}}}{w_0} = (|l|+1)^{-1/2} \tag{4.5.7}$$

Plugging this into Eq.(4.5.5) gives us a theoretical maximum conversion efficiency of $\approx 93\%$. Since we have shown that it is impossible to have contributions from modes $l \neq Q$, the remaining energy of the converted beam is to be found in the various other radial modes where $p \neq 0$.

All of this indicates that SPOs seem to provide a reliable and efficient method of converting lasers from standard output modes to OAM-carrying modes. The construction techniques of these plates have been undergoing constant improvements with modern plates capable of reaching very high conversion efficiencies. Since there are no complex micrometric patterns, scaling the plates to larger diameters to accommodate high-power beams is not too challenging. It is even possible to use a segmented plate instead of a continuous one. These plates are formed by placing circular sector layers of varying thicknesses forming a helical ladder-like structure. The transmission function of an N-step segmented plate can be modelled as

$$T(\rho,\phi) = \exp\left(iQ\sum_{j=1}^{N-1} \mathbb{1}_{A_j}(\phi)\frac{2\pi j}{N-1}\right) \tag{4.5.8}$$

where $\mathbb{1}_{A_j}(\phi)$ is the indicator function which is equal to 1 if $\phi \in A_j$ and zero otherwise, and A_j is a set of the partition of the domain of ϕ

$$A_j = \left[\frac{2\pi j}{N}, \frac{2\pi(j+1)}{N}\right] \tag{4.5.9}$$

In this case the quality of the conversion will depend on the number of segments. It has been shown that with just 16 steps, there is minimal efficiency loss for $l = 1$ [107]. Segmented plates can be easier to fabricate especially if the fabrication relies on deposing layers of material instead of etching.

For our purposes, we have built a continuous SPP made of polymethyl methacrylate (PMMA) with a diameter of 100 mm. This particular plate was designed for the aforementioned ASTRA laser. Considering the refractive index of PMMA and the wavelength of the ASTRA laser, the step of this SPP needs to be $h(2\pi) = 1.65\,\mu\text{m}$ in order for it to have a charge of $Q = 1$. The plate was fabricated in the SciTech fabrication facilities at the RAL

In order to fabricate the plate, it was first turned out of a stock material. The rear surface of the plate is skimmed to ensure that it is parallel to the machine axis and is flat with no waviness as any variation would cause the step in the SPP to be lost. The active face is then skimmed flat with respect to that zero plane. The machine then runs the cutting program to shave the form into the surface. The plate is not held with a vacuum chuck as that would distort the surfaces when machining. The spacing of the cuts, the tool speed, start diameter and helix depth are all optimised to ensure the best cutting performance; further, the depth of the step is limited by the swarf pick on the tool. The plate is measured using a white light interferometer to determine the step height of the SPP across its radius, which is then fed back into the algorithm to optimise the programming. Fig.(4.8) shows the focal spot of the ASTRA laser after it had gone through the SPP, where the beam was apertured to 20 mm in order to be able to see the distinctive "doughnut" shot typical of an OAM beam.

There are various difficulties and challenges that arise when using transmissive spiral phase plates to generate OAM-carrying beams. Firstly, for the plate to work as expected, the incoming beam must have a relatively flat phase front in order to produce the best quality OAM beams at the output. This generally can be arranged as most high-power laser facilities, which have a way to correct any wavefront aberrations using adaptive optics and wavefront sensors. These adaptive optics can similarly be used to optimise the quality of the converted beam [108].

Additionally, there are mechanical concerns when it comes to fabricating these plates. Since they are transmissive, we would like to have them as thin as possible in order to minimise any B-integral or non-linear effects driven by the high-powered lasers that will be propagating through them. Although the thickness of the step in any given SPP is quite small (on the order of the wavelength of the light for which

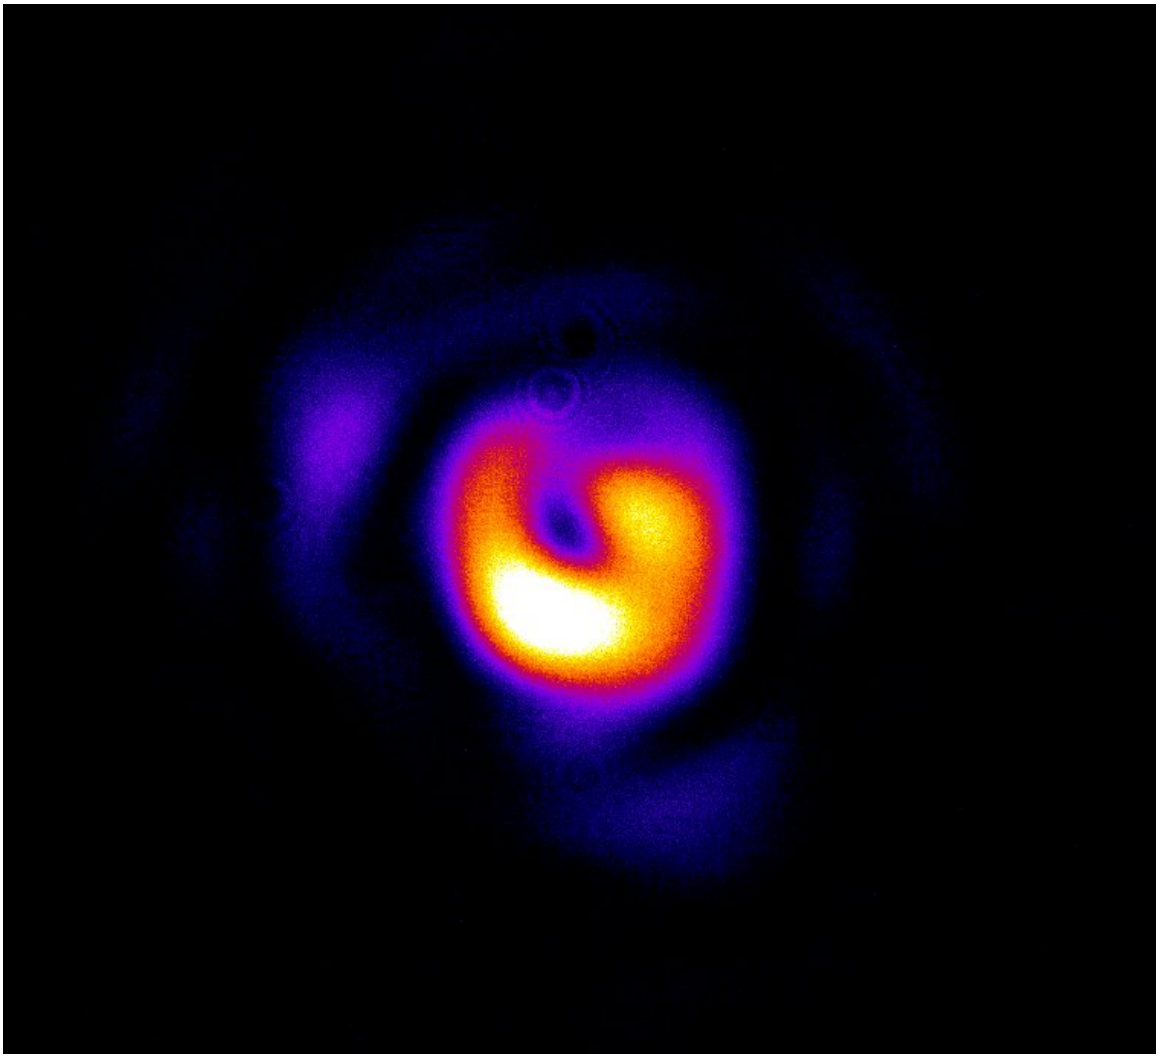

Figure 4.8: The focal spot of the ASTRA laser after passing through the transmissive continuous spiral phase plate. The beam has been apertured from its full size down to 20 mm to increase the size of the focal spot in order to better resolve the spatial structure imparted by the SPP and to avoid the step height inhomogeneity that existed on the plate away from its centre. As can be clearly seen from the image, the spot shows the distinctive singularity near the centre due to the phase discontinuity caused by the azimuthal dependence.

it is designed), this only contributes to a small fraction of the overall thickness of the device. The base of the plate, which corresponds to the h_0 in Eq.(4.5.1), needs to be thick enough to prevent any warping effects and to keep the SPP structurally stable and flat. This is because even a slight warping in the plate could potentially wash out the step and damage the quality of the wavefront that it produces. This becomes even more crucial as the size of the plate increases to accommodate larger laser diameters.

There is then a trade-off between maintaining a thick enough plate for it to be structurally sound and avoiding non-linear effects. This is compounded by the fact that with higher-power lasers the diameter needs to be even larger, making it necessary to design even thicker plates to maintain structure stability. All of this effectively puts a power-threshold on the usability of these transmissive plates, especially if one is very sensitive to any distortions in the pulse length or spectrum that can occur when using transmissive optics with high power beams.

For our particular plate, there were many difficulties with maintaining a consistent step height along an entire radial line. This was due to the way the plate was fabricated, its continuous nature and its large diameter. This inhomogeneity was severe enough that it made the plate incapable of adequate conversion for the full-sized 60 mm ASTRA beam. As such, in order to perform our diagnostic tests, we were forced to only use it with the beam apertured to 20 mm as the height of the SPP was fairly consistent over that length. This reduces the energy in the pulse to about 10% of the full energy.

A more promising method of OAM conversion is found in reflective spiral plates. As mentioned previously, they operate in a similar way to SPP while avoiding the transmission of beams in dielectric media. As such, they circumvent any issues associated with non-linear dielectric effects caused by high-power beams. That way, they can be as thick as needed to make them stable and can be further made from materials that withstand even the most powerful lasers in operation. This means that, in principle, they can withstand any practically realisable beam power in the same way that other optical components are built for these beams. They can even be fabricated in such a way so as to operate off-axis [109], reflecting the beam at an angle and avoiding any geometrical issues with their inclusion in experimental setups that might have been caused by a normal-incidence-only mirror.

4.6 Conclusion

In this chapter, I have looked at how high-power laser pulses are generated in general as well as the specifications and parameters of the ASTRA GEMINI laser facility in particular. I then looked at various methods of generating lasers which carry OAM and the suitability of each method for use on high-power lasers. The three main methods that were explored for producing OAM-carrying lasers are

1. Using Hermite-Gaussian modes to create Laguerre-Gaussian ones either via superposition of various modes or via the action of appropriate optical components. Although this method has a theoretical 100% conversion efficiency if the input HG modes are ideal (due to the fact that any LG mode can be represented as a sum of HG modes), it presents many other disadvantages that make it impractical. Chief among those are the fact that higher-order LG modes will be more complicated to produce and, more importantly, the fact that modern lasers, especially high-power ones, are optimised to avoid producing HG modes.

2. Using holographic phase plates to convert standard laser modes into OAM-carrying ones. This method presents the advantage that making higher-order modes using it is not particularly more difficult than lower-order ones and that it uses standard laser outputs as its input modes. This method can also achieve very high efficiencies when the hologram is blazed. However, due to the incredibly complex process of fabricating these plates for the large beam diameters required for high-power lasers, they are not a practical solution for generating high-intensity lasers with OAM even if they are ideal for certain low-power use cases due to the increasing ease of creating small-scale computer-generated holograms.

3. Using spiral phase optics to convert standard laser modes. This method has a very high theoretical conversion efficiency and there have been multiple reports of efficient conversions using it. It also has the advantage of a relatively straightforward fabrication process that can be scaled to accommodate large beam sizes. Another important advantage offered by this method is the fact that it can use reflective optics instead of transmissive ones by using spiral phase optics. This circumvents any issues with non-linear effects distorting the pulse length or spectrum of very high-power pulses when they are transmitted in dielectric media. It also effectively means that the only limitation to its use with high-power optics is the damage threshold of the material used in the coating of

the mirrors, which generally also limits all other reflective optical components used in high-power laser systems.

Now that there is a reliable way of generating high-intensity pulses which carry OAM, I will turn my attention to the best methods of characterising these pulses and the amount of OAM they carry. Not only is this important for experimentally gauging, in an accurate manner, the quality of conversion methods, it is also important for diagnosing the results of any experiment which investigates OAM effects, including the one proposed in the previous chapter.

Chapter 5

Characterising the OAM of high-intensity pulses

"I often say that when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of science, whatever the matter may be."

William Thomson, First Baron Kelvin.

5.1 Context

Although OAM-carrying beams have a very distinctive transverse spatial structure with the singularity present at the centre of the intensity profile, that alone does not suffice to characterise these modes. The singularity could potentially be due to other effects and even if it is simply due to the presence of OAM, it would not be possible to accurately gauge the amount of OAM present nor its orientation. As such, more sophisticated methods are required in order to be able to analyse the results of experiments involving the interaction of high-power OAM-carrying modes. Of course, the most comprehensive method would be to use a wavefront sensor to measure the phase information contained within the beam profile. However, this is

highly impractical as most wavefront sensors lack the resolution required to completely acquire all the phase information over the relatively large diameter of these beams. Additionally, these sensors are highly sensitive to alignment and any other aberrations that are bound to exist in physical situations.

As with the generation of OAM-carrying beams, much of the research into characterising the OAM of light has been focused on low-power beams and even single-photons [86]. This is due to the fact that the vast majority of experiments involving OAM have been in that regime, while the application of OAM to high-power interactions has mostly remained in theoretical and numerical research so far. For a characterisation method to be suited for measure the OAM of a high-power laser, it must meet some important criteria: firstly, it needs to be able to measure collective electromagnetic effects (not single-photons). Depending on the type of high-power facility, additional constraints must be met. For low-energy high-intensity facilities, the characterisation method also needs to be able to measure short to ultra-short pulses (ps down to fs-scale). For high-energy facilities with ns pulses, this method further needs to be able to perform the measurement in a single-shot, since the repetition rate on high-energy laser facilities can be extremely low and there are often issues with shot-to-shot reproducibility.

This chapter explores three different methods for characterising the OAM and their applicability to experiments involving high-power laser pulses. Firstly, the experiment on the ASTRA laser, where the different diagnostic methods were tested and the experimental results on them were collected, is described. This is followed by detailing the three different methods: the cylindrical lens diagnostic, the interferometric diagnostic and the projective diagnostic. The theoretical basis of each method is discussed and their experimental results are showcased. The chapter concludes with a brief discussion on the relative advantages and disadvantages of each method and how they fit in the context of the planned photon-photon scattering experiment.

5.2 ASTRA experiment

The ASTRA laser, which was described in the previous chapter, provided the optimal platform for testing the conversion and characterisation of the OAM of high-power laser pulses. Although it is by no means the most powerful laser available, the ASTRA laser and the associated experimental facility are ideal for field testing new diagnostics. This is due to the fact that they provide a reasonably powerful laser (in the range of 10s of TW) for long experimental runs with a very similar setup to what would

be found on the most powerful laser systems: a large vacuum interaction chamber, a large beam size, diagnostic optical tables inside and outside the chamber, and adaptive optics control of the beam's wavefront.

The particular experimental campaign which produced the results that are outlined in this chapter was a twelve-week long campaign aimed at identifying the best ways of measuring the OAM of high-power pulses. The experiment was also specifically focused on the context of high-intensity laser-plasma interactions. This was in preparation for an experimental campaign to investigate the effects of OAM on laser-plasma instabilities. As such, the experimental setup was arranged to emulate what a laser-plasma experiment looks like.

Fig.(5.1) shows an overview of the main components of the setup. The ASTRA beam enters the compressor where it is compressed to $\approx 50\,\mathrm{fs}$. At this point, the compressed pulse is carrying a power of 12 TW and as such is only propagated in vacuum. The compressor is connected to the vacuum interaction chamber through a gate valve which transports the pulse between the two chambers and allows for maintaining separate vacuums if the interaction chamber needs to be opened. Inside of the interaction chamber, the ASTRA pulse passes through the SPP described in the previous chapter where it acquires an OAM of $l = 1$.

The converted beam then passes by a 99 : 1 beamsplitter, which backreflects a portion of the laser pulses with an angle. This portion is diverted out of the interaction chamber towards the diagnostic table. This is meant to emulate the radiation produced during laser-plasma interactions, which would generally be a fraction of the input laser's energy. The main beam continues to be focused using an off-axis parabola and then another portion of it is diverted out of the interaction chamber to a focal spot diagnosis. This focal spot diagnosis is what produced the "doughnut" image in Fig.(4.8),which is used to adjust the adaptive optics and correct any aberrations in the wavefront of the beam. Before placing the SPP, a Haso wavefront sensor is placed in the focal spot diagnostic where it measures the wavefront of the beam. This measurement is used to adjust the adaptive optics and produce the flattest wavefront possible in order for optimal OAM conversion.

The portion of the beam which was diverted before the focusing optic is directed onto the external diagnostic table. There, the beam is split into the three different diagnostics allowing for simultaneous measurements on each shot. The beamsplitters can also be replaced with mirrors to focus on each diagnostic independently. Each diagnostic was initially assembled and pre-aligned independently using a small continuous wave (CW) laser. It should also be noted that, although the diagnostics in this

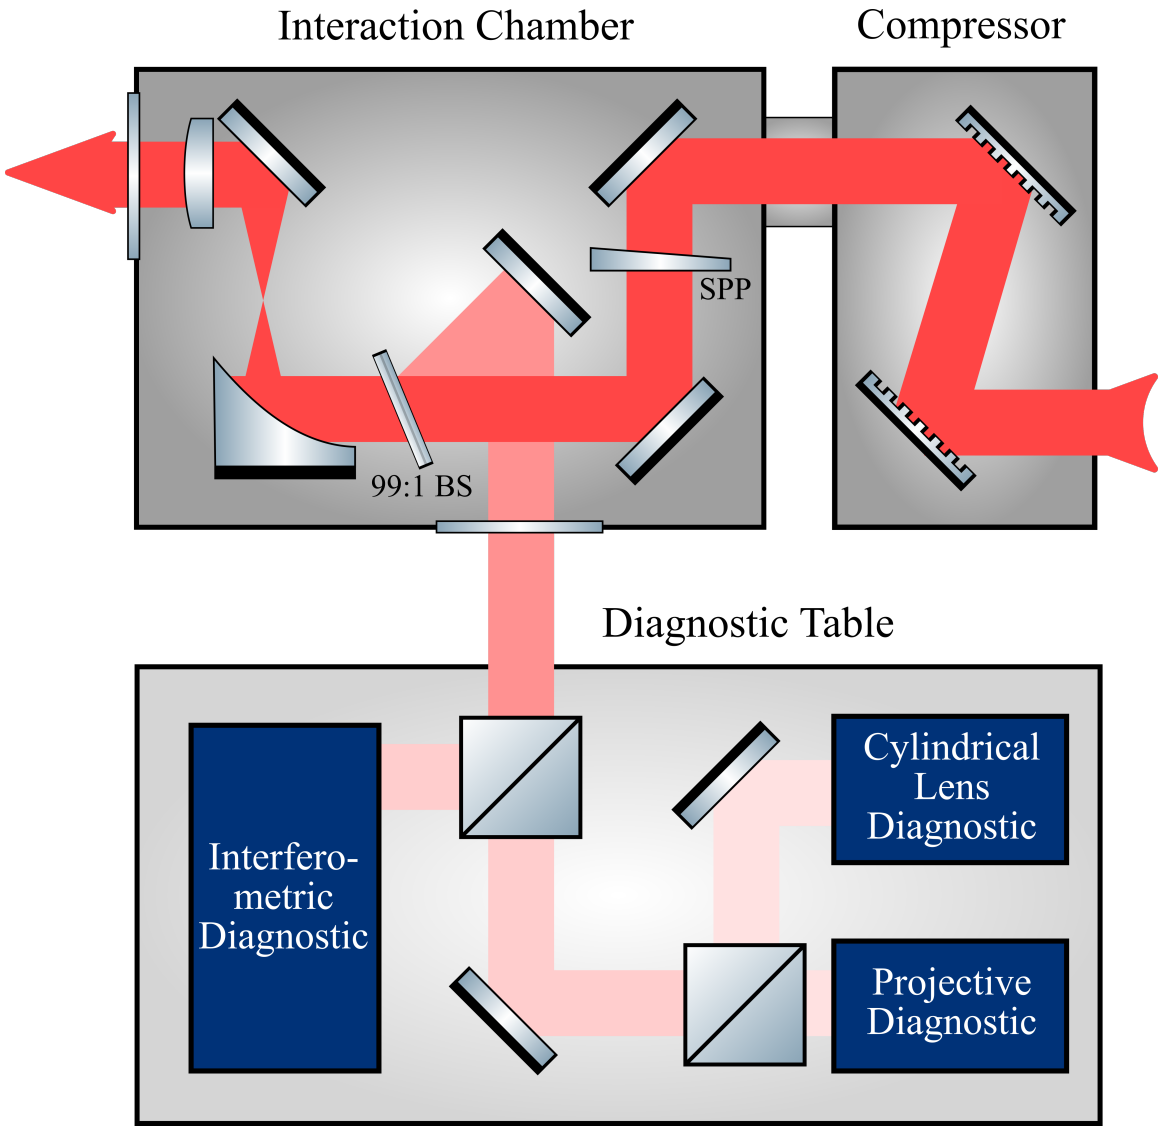

Figure 5.1: Diagram of the main parts of the ASTRA laser's target area showing the compressor chamber, the vacuum interaction chamber with its internal optical setup and the external diagnostics table with all the different OAM diagnostics.

experiment were placed in air outside the interaction chamber due to size limitations, this is not necessary and they would work similarly under vacuum. This could be needed if the beam to be measured is very high-intensity and cannot be manipulated in air. However, if that is the case, some of the diagnostics may not be suitable due to the use of transmissive components. Once the setup of each diagnostic is complete, reference shots were taken without the presence of the SPP, which were then followed by the data shots, which, in turn, were taken after placing the spiral phase plate in the path of the beam.

5.3 Cylindrical lens diagnostic

It is well known that lenses, arranged in particular configurations, can be used to immediately compute the spatial Fourier transform of transverse spatial profile of the input light field. In particular, a one-lens imaging system would produce, at the focal plane, a field whose magnitude is equivalent to the magnitude of the spatial Fourier transform of the input field [110]. However, as seen in the previous section, the information contained in the focal plane of a spherical lens would not be sufficient to characterise the OAM on the input field. However, it has been shown that cylindrical lenses, which only focus along one particular spatial axis leaving the other intact, produce very distinctive shapes at their focal planes when a collimated OAM-carrying mode is incoming on them [111].

In fact, the intensity $I(x_f, y_f)$ at the focal plane (x_f, y_f) of a simple one-cylindrical lens optical system can be modelled as the magnitude of a one-dimensional Fourier transform of the two-dimensional function describing the transverse profile of the input LG mode $\mathrm{LG}_p^l(x_i, y_i)$ at the input plane (x_i, y_i):

$$I_p^l(x_f, y_f) \propto \left| \int \mathrm{LG}_p^l(x_i, y_f) e^{i2\pi x_f x_i/\lambda f} dx_i \right|^2 \tag{5.3.1}$$

where λ is the wavelength of the incoming light and f is the focal length of the cylindrical lens. There are two subtleties here that should be noted. This form assumes that the cylindrical lens has been perfectly aligned so as to act along some well-defined spatial x-axis. However, the choice of the x-axis is of course completely arbitrary and, in the case of LG modes, irrelevant since rotations in the transverse plane would only amount to adding a constant phase factor which is physically meaningless. In experimental conditions, the rotation of the cylindrical lens around the axis of propagation of the incoming beam would be adjusted for the convenience of

(a)

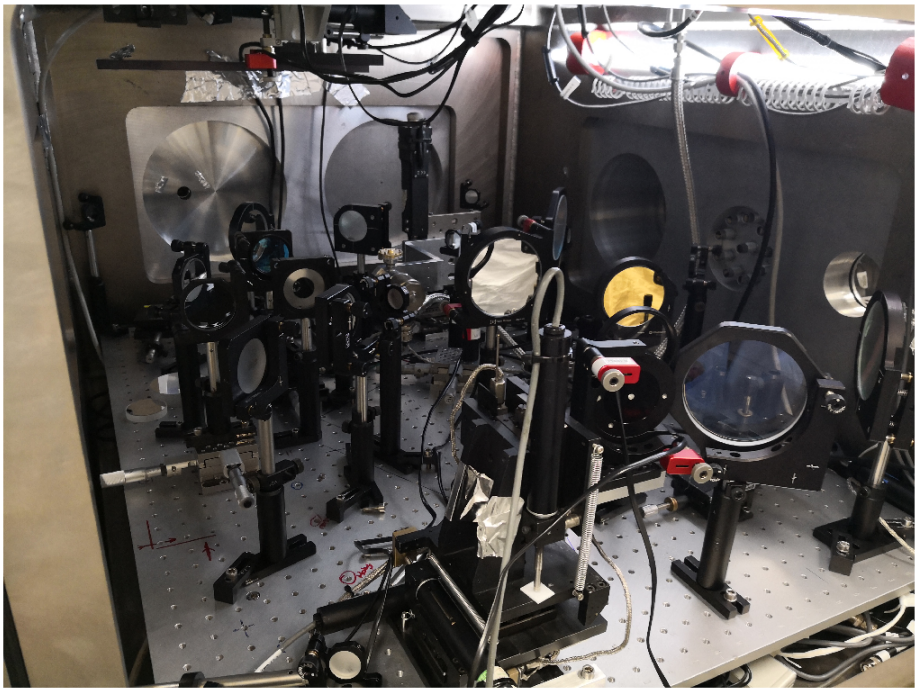

(b)

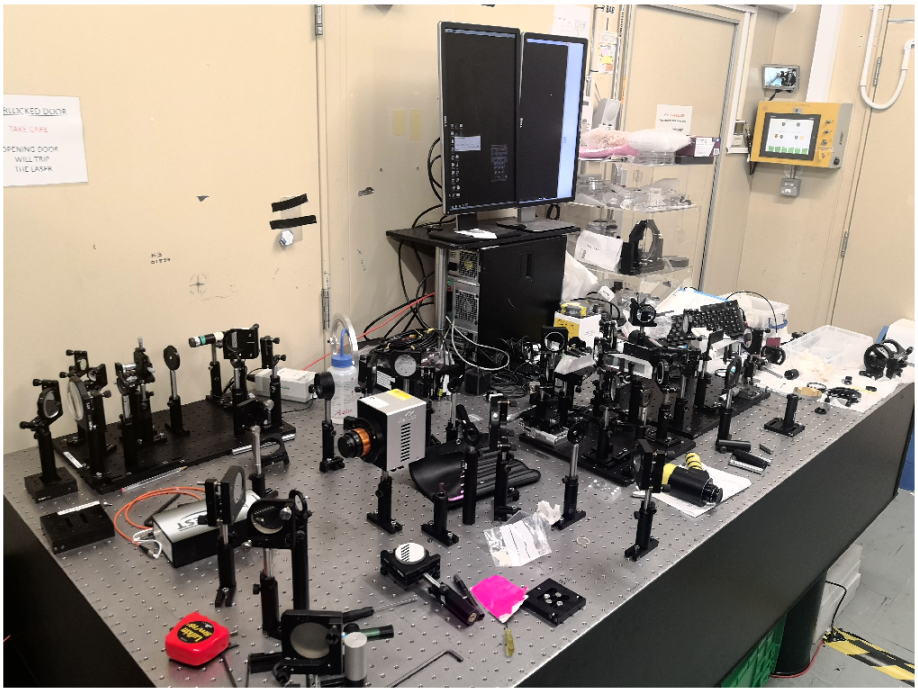

Figure 5.2: (a) Picture of the inside of the ASTRA laser's interaction chamber (b) Picture of the diagnostics table outside of the interaction chamber showing the various OAM diagnostics.

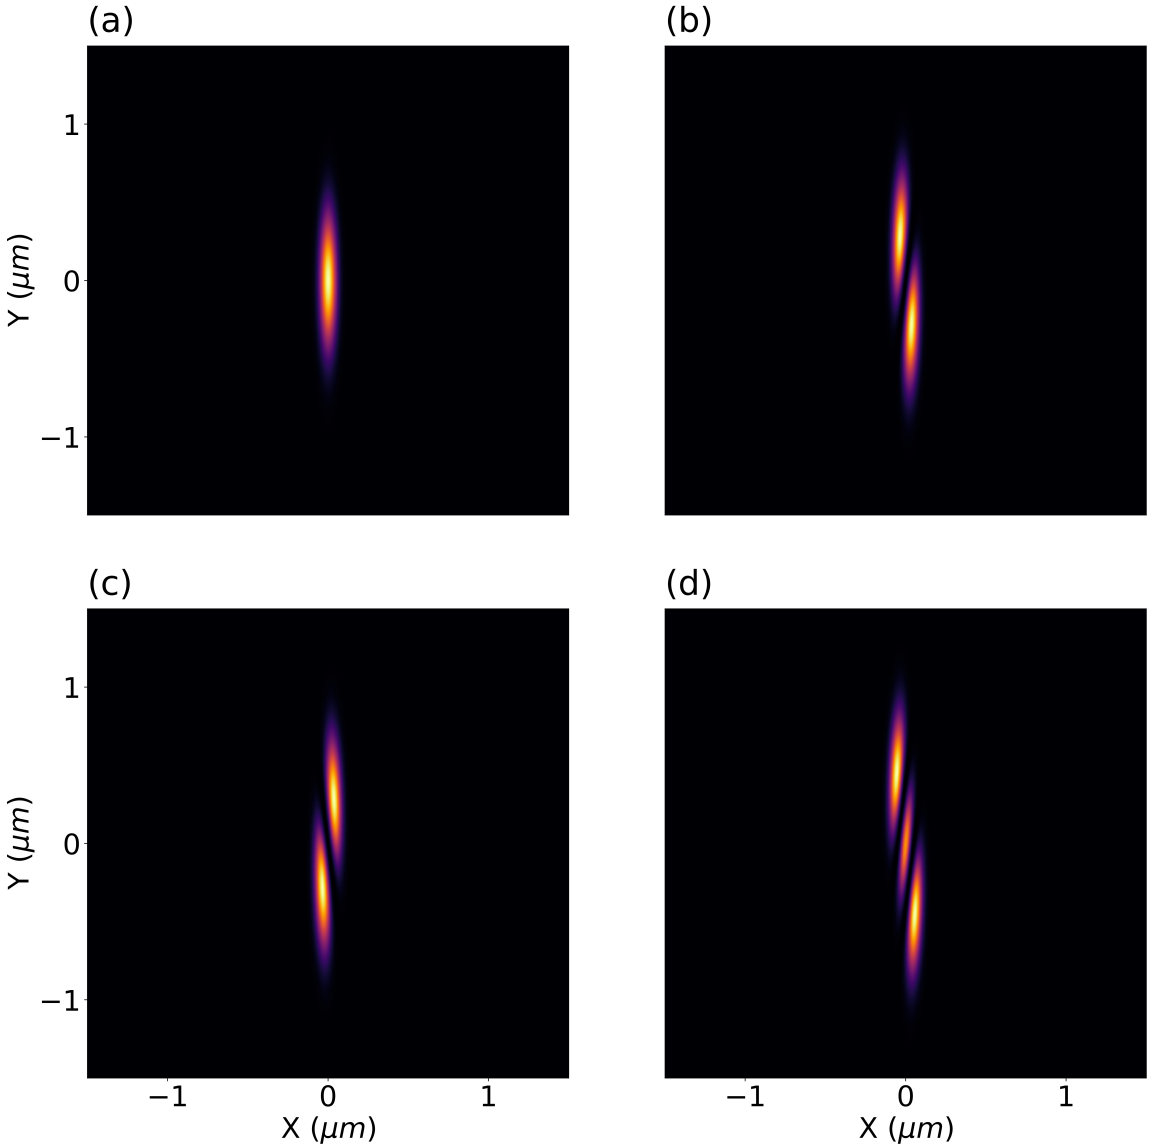

Figure 5.3: (a-d) Simulated norm of a one-dimensional Fourier transform performed along the transverse x-axis of different Laguerre-Gaussian modes with $l = 0, +1, -1, +2$ respectively. The simulated norms clearly show the distinctive fringe pattern caused by the OAM.

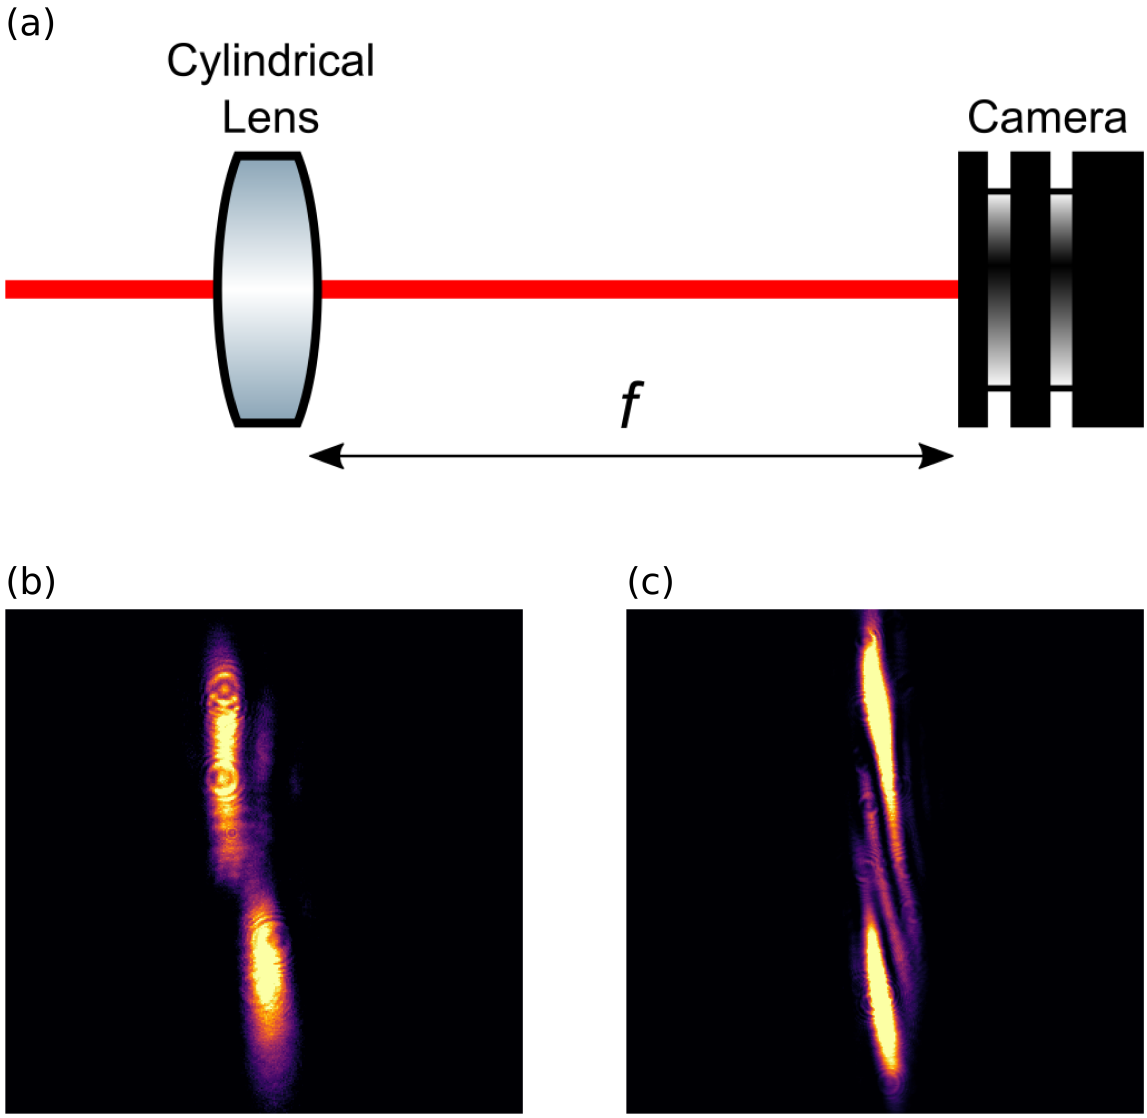

Figure 5.4: (a) Optical setup for the diagnostic that was built at the ASTRA target area (b) Experimental result from a shot of the ASTRA laser after going through the SPP clearly showing the pattern similar to the $l = 1$ simulated pattern shown above (c) Experimental result from a continuous low-power 400 nm beam after it went through the same SPP placed in the opposite orientation.

the setup and subsequent data analysis but would not impact the information contained within. The second point is that, although the cylindrical lens only acts along the x-axis, technically there are diffraction effects on the y-axis from the plane of the lens y_i to the focal plane y_f. However, since I am considering a collimated input beam with a Rayleigh length that is larger than any of the relevant physical dimensions of the system, these effects can be ignored.

Fig.(5.4.a)-(5.4.d) show simulated results of the action of a cylindrical lens on various LG modes $\mathrm{LG}_0^l(x_i, y_i)$. These simulations clearly show that for LG modes with $l \neq 0$, there are fringes which appear. The number of dark fringes correspond to the absolute value of the different l mode numbers. The simulations further show that the tilt of the fringes changes with the sign of l. This means that in one shot, this diagnostic is capable of providing information concerning both the magnitude and the orientation of the OAM carried by the mode. It is important to note that there is no "absolute" sign for the OAM of light and it is only meaningful to speak of the relative sign of one with respect to another. In fact, Eq.(5.3.1) can be solved for particular values of l so that the experimental data may be compared to the ideal solution. For example, solving the equation for the first LG modes with unit waist gives:

$$I_0^{\pm 1}(x_f, y_f) \quad \propto \exp\Big[-\Big(\frac{k^2}{4} + y^2\Big)\Big](k \pm 2y) \tag{5.3.2}$$

$$I_0^{\pm 2}(x_f, y_f) \quad \propto \exp\Big[-\Big(\frac{k^2}{4} + y^2\Big)\Big](-2 + (k \pm 2y)^2) \tag{5.3.3}$$

Fig.(5.4.f)-(5.4.g) show some experimental results using this diagnostic both on an ASTRA main beam shot and using a much weaker mW CW laser diode emitting light at $2\omega_{\mathrm{ASTRA}}$ after they had gone through the SPP. The output of the diagnostic fits well with what is to be expected from these beams. For the laser diode, the SPP was put in the opposite direction to showcase the ability of this diagnostic to distinguish the sign of the OAM, which is clear from the image seen at the focal plane. Additionally, since the laser diode is emitting at twice the frequency of the ASTRA beam, it will acquire twice the azimuthal charge going through the SPP, and subsequently twice the OAM. As expected, the result of the diagnostic on the converted laser diode shows two dark fringes characteristic of $l = 2$.

This diagnostic has many important advantages that make it well suited for use in high-power laser systems. Indeed, it is based on a single optical component and a single sensor component and, as such, requires very little alignment. Additionally, reflective cylindrical optics are available that would allow for the characterisation of even

more high-powered beams where the intensity would damage traditional transmissive cylindrical lenses. This method also gives us information both on the magnitude and the orientation of the OAM in one easy-to-interpret output. Space-wise, since this diagnostic requires so few components, it can be made arbitrarily compact by using a cylindrical optic of appropriate focal length. It should be noted, however, that if the F-number (f/D with f being the focal length and D the diameter of the beam) of the system is too small, this will result in a very small spatial extent in the direction of the focusing. This might make interpreting the output difficult, or even impossible, depending on the spatial resolution of the sensor being used. However, this problem can be circumvented by aperturing the input beam reducing the F-number of the system. A shortcoming that cannot be avoided, unfortunately, is that if the input beam to the diagnostic is composed of a superposition of various OAM modes, then it will be relatively difficult to accurately decompose them, since the only data that would be in the output is the magnitude of the sum of their various Fourier transforms.

5.4 Interferometric OAM diagnostic

Another common way of characterising structured spatial phase variations is by interfering the beam with another reference, such as one carrying a well-defined OAM (e.g., a Gaussian $l = 0$ beam). This allows for extracting certain information about the phase of the original beam. However, using a reference beam is impractical when it comes to high-power laser experiments, as this requires precise alignment (in space and time) of the pulse with another reference pulse that is similar enough in amplitude and wavelength. A way around these restrictive conditions is to interfere the pulse with a copy of itself. However, for this to give any useful information on the azimuthal charge, there must be a way to modify the transverse profile of one copy relative to another. This can be done using optical components known as prisms [112]. In the case of the setup that will be discussed in this section, I used what is called a dove prism. A dove prism is a transparent hexahedron where the four lateral sides are regular trapezoids, and the top and base are different sized rectangles. It can be used to rotate the transverse profile of an input beam or to create a mirror image of it. Here, I am interested in its ability to rotate an input OAM beam, which, as mentioned previously, amounts to adding a relative phase shift.

My setup is based on the well-known Mach-Zehnder (M-Z) interferometer. As shown in Fig.(5.5.a), two dove prisms are introduced into each of the M-Z interferometer's arms. The prisms are rotated by $\alpha = \pi/2$ with respect to each other. This

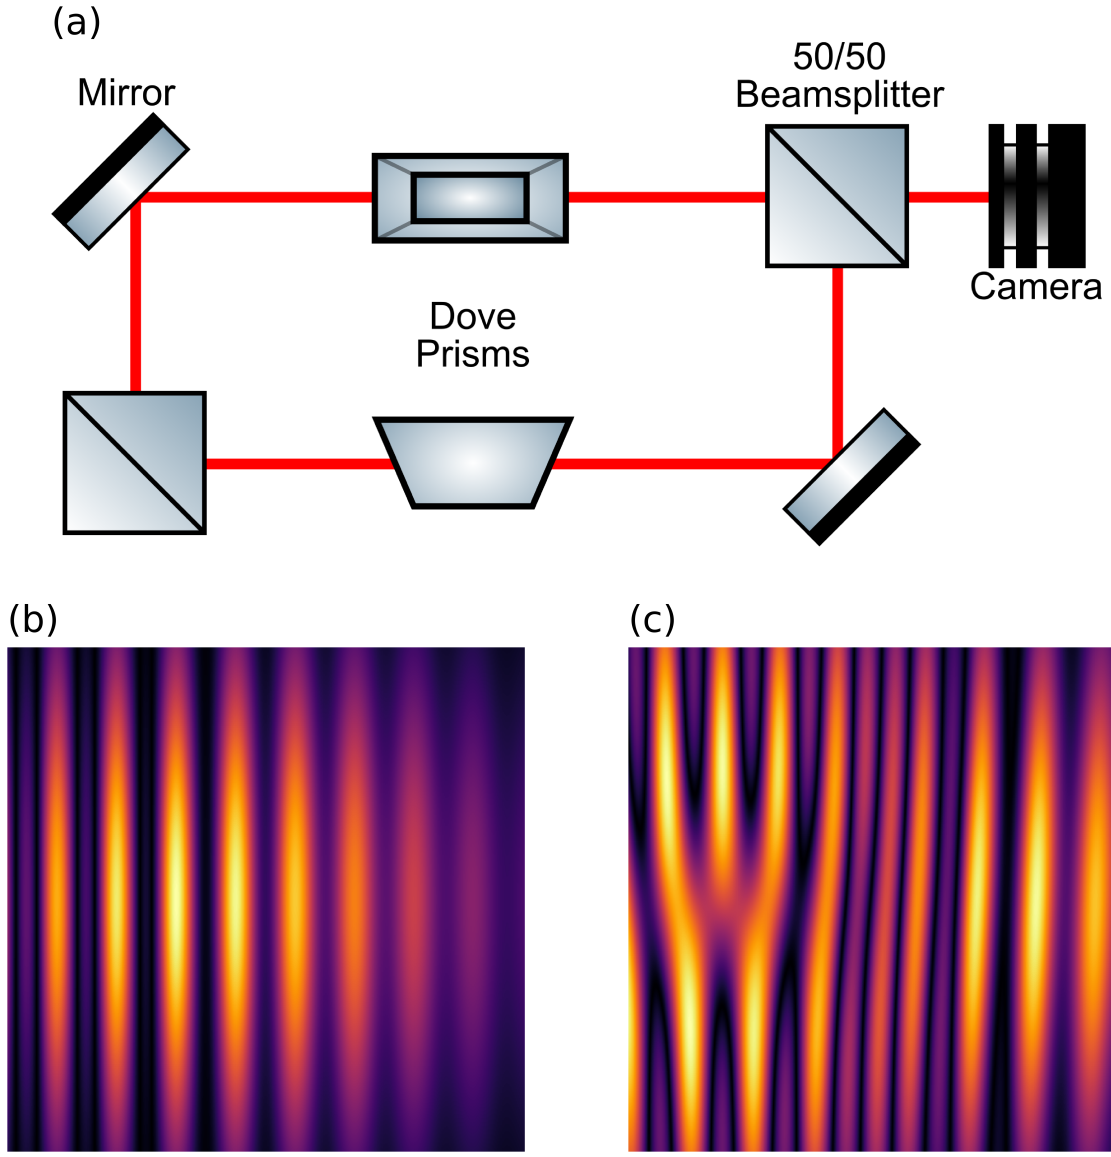

Figure 5.5: (a) Shows the basic schematic of the optical setup for the OAM interferometer (b-c) Simulated results for the output of the interferometer in the case of normal Gaussian profile and a LG with $l = 1$ respectively.

causes one copy of the input beam to be rotated by $2\alpha = \pi$ relative to the other. In the case of an LG mode, this would mean that the transverse spatial variation of the phase of one copy will be 180° out of phase with the other. This signifies that when the M-Z inteferometer is perfectly aligned, even OAM modes would favour one path after the final beamsplitter while odd modes would favour another. This can be used to create an "OAM-splitter" for single photons [113].

However, this splitting is not very relevant for my purposes as I am currently interested in characterising macroscopic fields and not single photons. Another consequence of this setup can be seen when the interferometer is not perfectly aligned so as to create interference fringes at the outputs of the final beamsplitter. As can be seen in Fig.(5.5.b-5.5.c), the interference between a beam and a rotated copy of itself creates distinctive patterns depending on the l mode number. For a regular Gaussian $l = 0$, the pattern corresponds to the standard fringes normally seen in such interferometers. For $l = 1$, the fringes present a distinctive "fork" shaped discontinuity similar to the one present in the phase holograms mentioned in the previous chapter. These varying patterns make it possible to diagnose the different OAM states of the beam. Again, they not only hold information on the magnitude of the beam but also on its relative orientation. For example, the pattern for $l = -1$ would be inverted relative to the one for $l = 1$.

The setup shown in Fig.(5.5.a) was built and aligned at the CLF and tested on the ASTRA laser. Due to the relative complexity of the setup, it was done on multiple stages. First, a standard Mach-Zehnder interferometer was built and aligned using a continuous source to first align the direction of the two copies in order to see the standard interference fringes in one of the two outputs of the final beamsplitter. Although the arms of the interferometer were adjusted so as to have two approximately equal path lengths, some final alignment is always needed using a delay stage on which one of the mirrors is fixed. This allowed me to fine-tune the path length to optimise the interference when switching to pulsed beams. This became especially crucial when dealing with fs pulses such as those from the ASTRA laser. Once the alignment and optimisation of the Mach-Zehnder was complete, Dove prisms with the desired relative rotations were inserted into each path. The orientation and position of the Dove prisms were then adjusted to first retrieve the interference fringes with the continuous beam. Although the Dove prisms were chosen to be of the same size, uncertainties in the manufacturing can be large enough to cause temporal misalignment for fs pulses. This was compensated for by using the aforementioned delay stage. As for the experimental performance of this diagnostic, this can be seen by looking at the

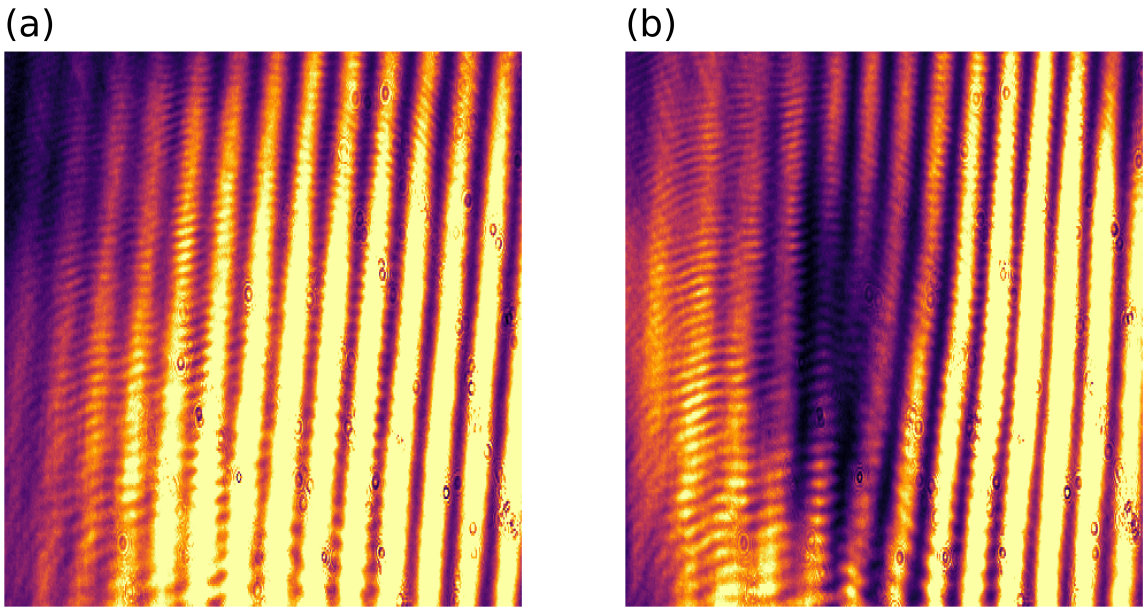

Figure 5.6: Experimental results for the interferometric diagnostic on the ASTRA laser: (a) The result for the standard beam without the SPP, so $l = 0$, shows the typical fringes usually seen in interferometers; (b) When the SPP is placed in the beam path, the profile of the fringes change and a discontinuity similar to the one in Fig.(5.5.c) is seen indicating the presence of OAM.

results shown in Fig.(5.6), which compares the output of the interferometer between the unconverted ASTRA pulse (a) and the converted one (b). The behaviour of the beam is similar to the one shown in the previously discussed simulated results. The interferometric diagnostic allows us to distinguish between regular Gaussian beam and a $l = 1$ beam.

Although this diagnostic offers a way to quantify the OAM state (both magnitude and relative orientation) of a pulsed beam, it suffers from many shortcomings. The setup is complex and the alignment is quite involved and time-consuming. Further, any slight drift in the direction of propagation of the input pulse might cause the diagnostic to fail. This alignment problem is only gets exacerbated as the pulse length decreases because the system then becomes even more sensitive to temporal misalignment. Hence comes the need for the delay stage in the Mach-Zehnder interferometer. For example, the ASTRA pulse, at 50 fs, would get temporally misaligned if the two paths in the interferometer drifted by 15 µm. Since many high-power lasers tend to use even shorter pulses, this makes this drawback particularly critical. Besides the practical implementation aspects of this diagnostic, it has other more fundamental disadvantages. Although it is relatively easy to diagnose the $l = 0, \pm 1$ OAM states

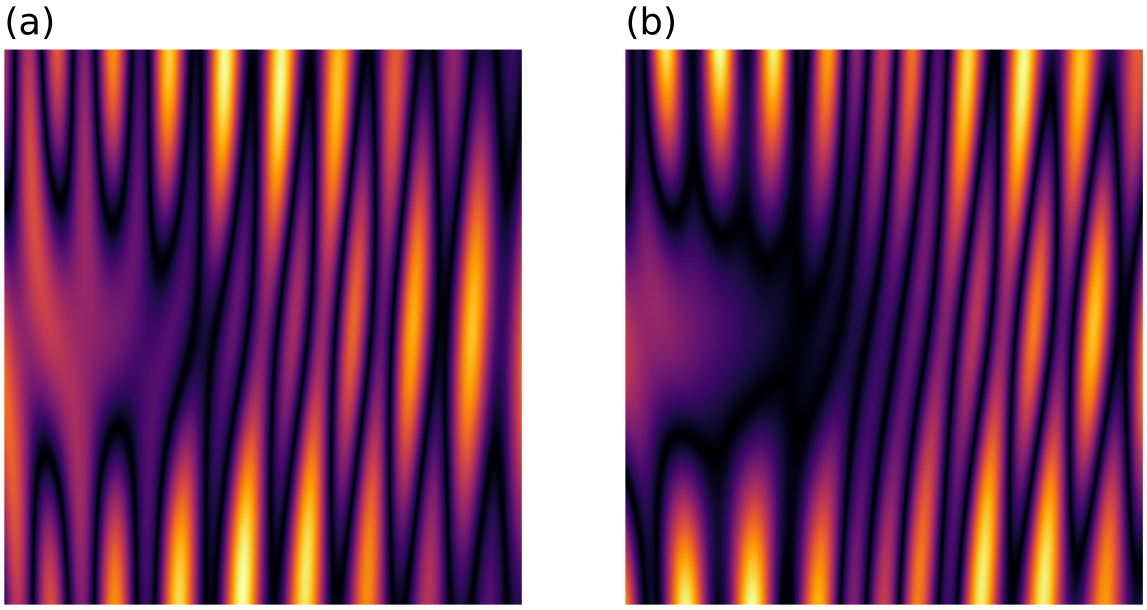

Figure 5.7: Simulated results of the output of the interferometric diagnostic for (a) $l = 2$ and (b) $l = 3$. As can be seen from (a) as compared to Fig.(5.5), the LG_0^2 mode has a forked shaped but with two inner prongs instead of one. For LG_0^3, the pattern becomes extremely difficult to identify and exceeds the size of the window.

from the pattern in the output image, this is not at all the case for higher order modes. Fig.(5.7) showcases the complex interference patterns that emerge from pure higher-order OAM states. In fact, higher-order modes would generate patterns that would require ever increasing spatial windows to properly diagnose. Additionally, in an actual experiment, the modes being diagnosed are rarely pure OAM modes. As such, the emerging patterns would be even more difficult to interpret and fit to simulated data.

5.5 OAM projective measurement

So far, the diagnostic methods that have been discussed have shown some weaknesses when faced with non-pure OAM states. The last one that will be discussed offers a way around this problem. This is done by "projecting" the input beam onto various OAM basis states. The basic principle around this projective method is the fact that any beam can be filtered from its non-Gaussian components by using a spatial mode filter, such as a pinhole. Consider a setup such as the one in Fig.(5.8). If one splits the input beam with l_{in} into multiple paths and inserts a known OAM mode converter onto each path (such as phase holograms or spatial light modulators (SLM)). Each

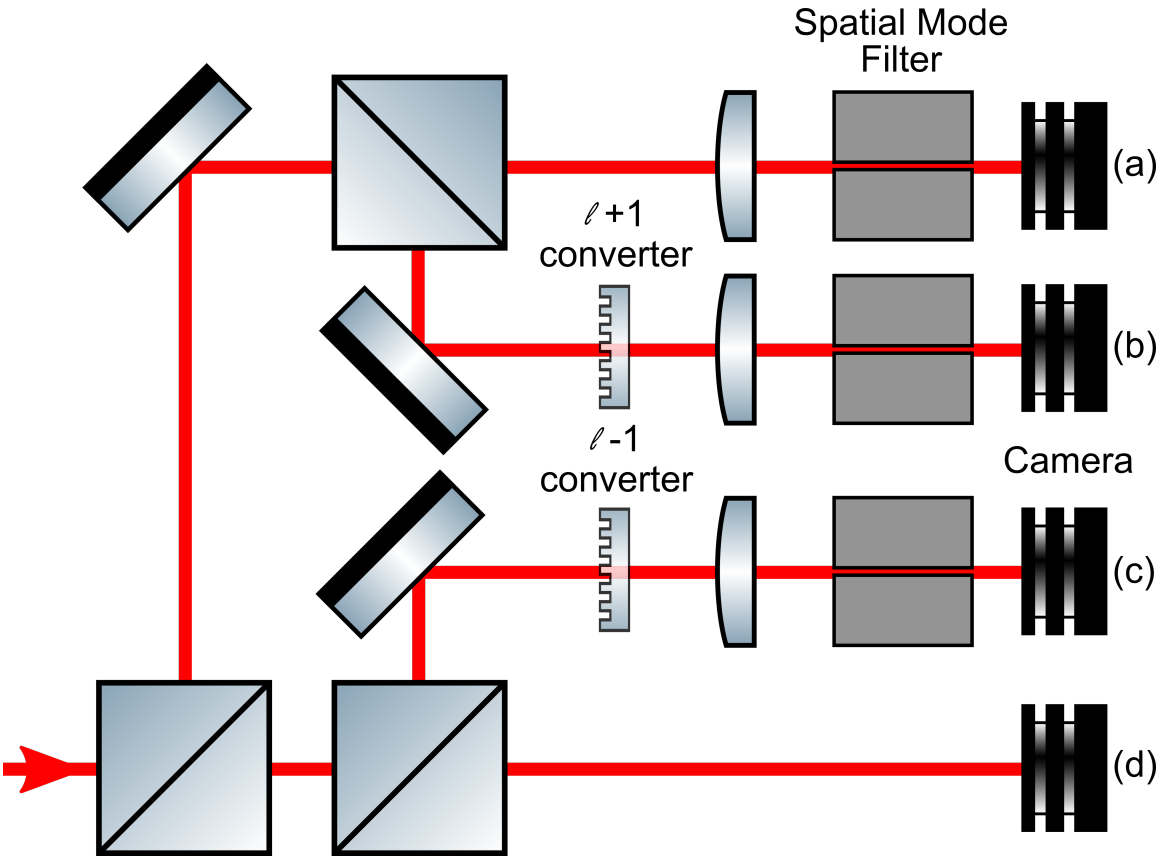

Figure 5.8: Schematic of an example of a projective diagnostic capable of projecting on $l = 0$ along path (a), $l = -1$ along path (b) and $l = 1$ along path (c). The remainder components of the beam can be inferred from the control measurement along path (d).

beam on each path j thus has an OAM state of $l_{\text{in}} + l_j$. Consider a beam that is a superposition of various OAM modes; on each path j, the component of the beam that has OAM of $-l_j$ will be converted into a Gaussian mode ($l = 0$), while other components will be converted to various other modes. This means that if the beam is focused on a spatial mode filter placed after the mode converter followed by a detector, only the $-l_j$ component will pass through the filter and be detected. An additional path can be added which has no filters or mode converters. This allows for the detection of the power that is in the rest of modes not being projected by the other paths.

To test the operation of this diagnostic method, I opted to project on just one state in order to calculate the Gaussian component that is left in the beam after the OAM conversion. This was done because fabricating spatial mode converters for beams of the size typically used in high-power laser experiments, is quite a complex and

difficult process. The spatial mode filter that was used was a single mode step-index optical fibre cable optimised for use for 800 nm light.

A step-index fibre cable is a cable formed of two distinct dielectric zones: the internal core and the surrounding cladding. The refractive index is uniform inside the core and decreases rapidly, going from the core to the cladding where it remains uniform again. When an external beam is coupled into a fibre-optical cable, it has the potential to excite a variety of what are called waveguide modes. For the vast majority of use cases, the waveguide modes of interest in step-index fibre optic cables are the so-called weakly-coupled linearly polarized modes LP_m^l [114], where l and m serve a similar purpose as in the free-space LG_m^l modes. These LP modes have a similar appearance to the free-space LG because they both arise from solving the wave equation in cylindrical coordinates. Fig.(5.9) shows the first two LP modes. It should be noted that, as opposed to the free space modes, the radial mode number for the guides linearly polarised modes starts at $m = 1$ not $m = 0$.

A single mode step-index fibre is one where the V number is less than 2.4 for its wavelength of operation λ, where V is defined as

$$V = \frac{2\pi a}{\lambda}\sqrt{n_1^2 - n_2^2} \tag{5.5.1}$$

with a being the diameter of the core, n_1 the refractive index of the core and n_2 the refractive index of the cladding. When the aforementioned condition is satisfied, only the fundamental LP_1^0 (analogous to the free-space Gaussian) is guided in the fibre. In such a cable, any higher-order mode that gets excited decays quickly in an exponential manner $I \propto e^{-\alpha L}$, where α is the attenuation coefficient and L is the length of the cable.

It should be noted that a fibre optic cable is only single-mode above a particular wavelength, termed the "cut-off" wavelength. Below this threshold, high-order modes can be guided. In practical use, the cut-off wavelength is generally defined as that where the attenuation of the second mode LP_1^1 is $20\,\mathrm{dB\,m^{-1}}$ more than the fundamental mode [115]. The higher the wavelength relative to the cut-off, the more high-order modes are attenuated. Experiments have shown that the attenuation can be as high as $25\,\mathrm{dB\,m^{-1}}$ for single photons [86]. In my setup, I used a one-meter long single-mode optical fibre where the cut-off was at 730 nm, i.e., well below the wavelength of the Astra beam. Detecting the power of the beam after the filter, and adjusting for the action of the filter as well as the transmission of the SPP, shows that 27% of the power remains in the Gaussian beam. This signifies that the transmissive SPP has a conversion efficiency of around 73% into non-Gaussian modes.

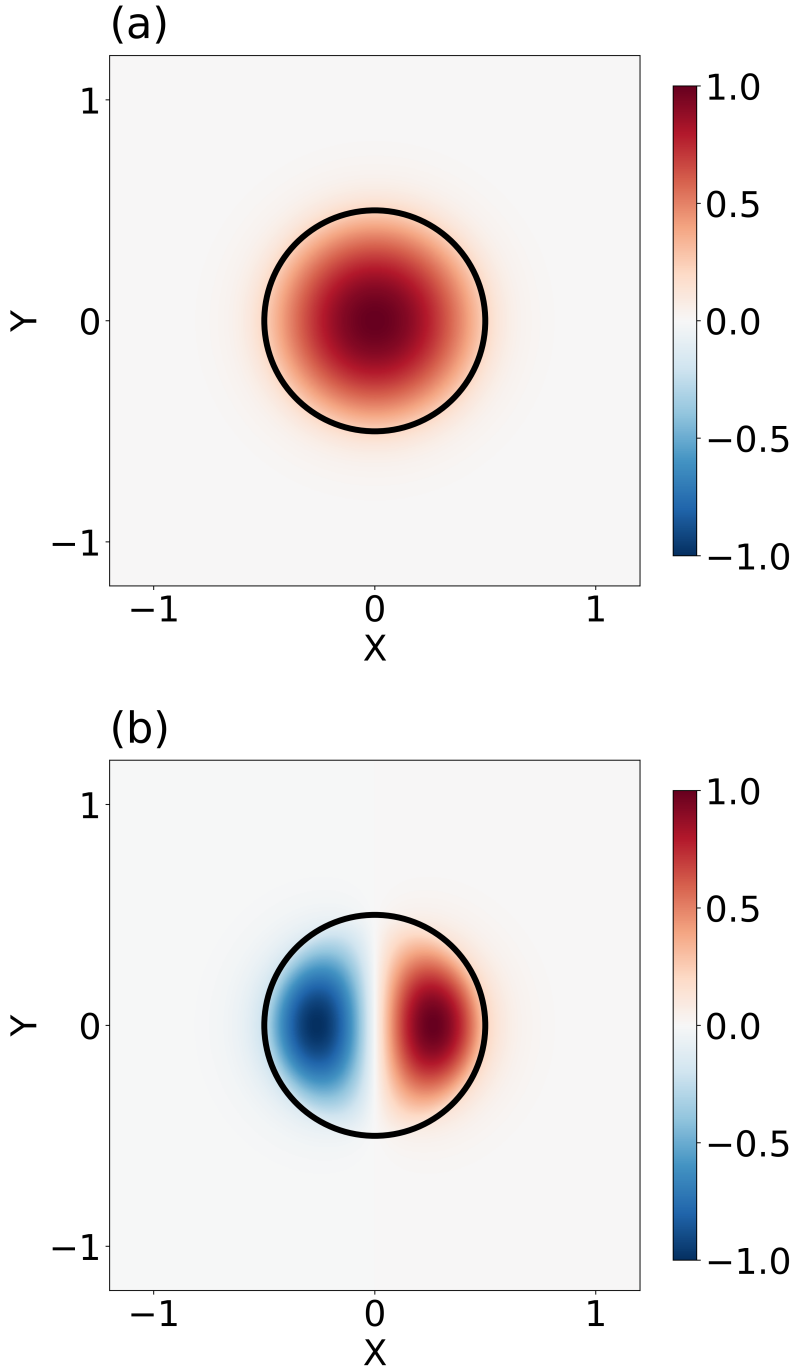

Figure 5.9: Simulated transverse profile of the the guides modes in a step-index fibre. The black circle represents the boundary between the core and the cladding. The two modes are (a) the LP_1^0 fundamental mode analogous to the free-space Gaussian and (b) the LP_1^1 mode analogous to the $l = 1$ mode.

Since this characterisation method is capable of identifying and quantifying an arbitrary number of LG_0^l components present in a beam with very little post-processing, it is technically the most complete way to characterise the OAM of any beam. However, this method does have several important drawbacks that limit its practicality. First, the alignment procedure can be very involved depending on the number of modes one wishes to project on. Each additional path requires at least four components aligned: a mode converter, a lens, a mode filter and a camera/photodetector. This makes the device poorly scalable. This is also aggravated by the fact that any misalignment in one particular path relative to the others would make it more difficult to interpret the results of the diagnostic. Additionally, mode converters can be quite difficult to make, especially if they need to accommodate the beam size or if they need to avoid dielectric transmission to accommodate the intensity of some high-power laser pulses. This brings us to the final limitation of the device when it comes to its use with high-power beams. Although there are fibre cables designed for high-power use, almost all high-power beams would require heavy attenuation if fibre cables are to be used to directly diagnose their OAM. An alternative, and more suitable, solution would be to use pinhole spatial filters, which can made be from materials that can easily withstand much higher intensities. There have even been pinhole filters capable of withstanding up to $3 \times 10^{12}\,\mathrm{W/cm^2}$ [116].

5.6 Conclusion

In this chapter, I have considered the various ways to characterise OAM both from a theoretical point of view and from a practical implementation standpoint in the context of high-power laser pulses. Looking at the results of the ASTRA experimental campaign, I summarise the various advantages and disadvantages of each method as follows:

- The cylindrical lens diagnostic is by far the easiest diagnostic to setup and requires the fewest components. By leveraging the Fourier transform properties of cylindrical focusing optics, it is possible to assess the magnitude and orientation of the OAM of an incoming beam in one shot. Thanks to the availability of cylindrical mirrors, this diagnostic can even be used with high-intensity beams, although there will remain the issue of finding an appropriate detector. These advantages make it ideal for quick and simple characterisation of very short pulses. Further, this method will prove especially useful in ensuring that OAM mode converters are well aligned and are working properly. The main weakness

of the cylindrical lens diagnostic is that it is unsuitable for when there is a superposition of various OAM modes.

- Although the interferometric diagnosis seemingly offers a single-shot way to characterise both the magnitude and the orientation the OAM of laser pulses, it has distinct disadvantages compared to the cylindrical one. Indeed, the interferometric diagnosis is incredibly difficult to align since it requires aligning a complete Mach-Zehnder interferometer and then adding the dove prisms to each arm without losing the alignment: a task whose difficulty is compounded by the fact that many high-power pulses are extremely short in duration. Another important shortcoming is that the pattern produced becomes increasingly complex to analyse with higher-order OAM modes and requires a larger spatial window. Finally, this method also shares the problem of the cylindrical lens diagnostic when measuring a superposition of modes.

- The projective diagnostic has the powerful advantage of being able to completely characterise the OAM content of both pure and mixed modes. However, this comes with the caveat that it can only characterise as many or as few modes as its setup allows. If one wishes to increase the number of projection states, an additional path needs to be setup and aligned. The precise relative alignment of each path is also crucially important since the way the results are analysed is by looking at the relative measurement on the photodetector placed on each path. The difficulty of aligning a single arm varies depending on whether simple holographic phase plates and single mode fibre cables can be used, or whether the intensity of the beam requires reflective mode converters and high-damage threshold pinhole filters are needed.

As such, it appears that the cylindrical lens diagnostic and the projective diagnostic are the two most suitable ones. Whether one is more suitable than the other depends upon the particular experiment and its requirements. In the context of the photon-photon scattering detection experiment proposed in Chapter 3, there are two important steps in that experiment which will require precise OAM characterisation.

The first is to gauge the efficiency of the conversion of one of the driver beams into an OAM-carrying mode in order to maximise it as much as possible. This is essential as it will determine the amount of photons that we should expect in the predicted OAM mode. It is also important to fix a specific orientation for the OAM of the driving laser pulse since the generated photons are expected to have the opposite

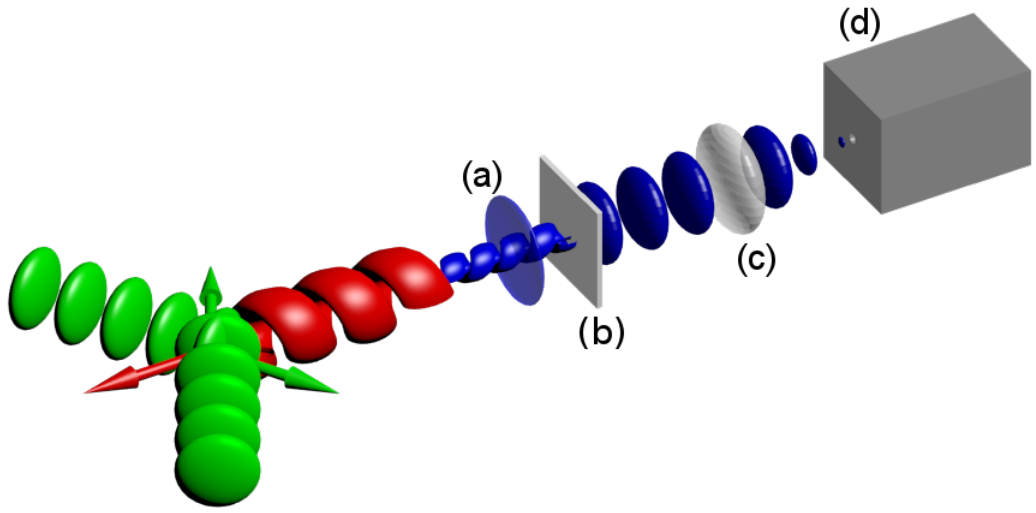

Figure 5.10: Planned schematic for the photon-photon scattering detection experiment: (a) Narrow-band wavelength filters eliminates light at undesirable frequencies; (b) Mode converters only convert the appropriate OAM mode to Gaussian; (c) Lens to focus the radiation on the spatial mode filter; (d) Spatial mode filter leading to a photodetector.

orientation. To accomplish this task, it will actually be easier to use two diagnostics. First, one should use the cylindrical lens diagnostic to obtain the rough alignment of the mode converter, followed by the projective diagnostic with only one projection state required: the desired OAM state.

Secondly, one needs to filter out light based on its OAM state while detecting a small number of photons. Again, here, the idea is to use the projective diagnostic method with only one projection state as an OAM filter. Since I am interested in measuring a few hundred photons, it suffices to use a blazed holographic plate mode converter and a single-mode fibre cable. As can be seen from Fig.(5.10), this filter needs to be placed after a wavelength filter along the expected line of sight of the interaction. Also, from Reference [86], one expects a 25 dB reduction of the background noise coming from other sources thanks to this filter. This drastically improves the signal-to-noise ratio, thus rendering this experiment more feasible.

Chapter 6

Single-Shot Frequency Resolved Optical Gating

"We have to remember that what we observe is not nature herself, but nature exposed to our method of questioning."

Werner Heisenberg

6.1 Context

As discussed in Chapter 4, in order for lasers to reach the extreme intensities that are needed for experiments such as the photon-photon scattering experiment, it is generally necessary to compress them into ultra-short fs and ps pulses. As CPA became more advanced, the use of high-power ultrashort pulses became widespread in scientific experiments in physics, chemistry and biology. As such, the need arose to properly measure and fully characterise these pulses.

Unfortunately, the response time of electronic circuits is generally significantly slower than the time-scale of these pulses. For sufficiently long pulses (> 30ps), high temporal resolution optical streak cameras used in conjunction with optical spectrometers are able to characterise the pulses spectrally and temporally. However, for shorter pulses, other methods must be used to measure the temporal evolution of the pulse such as the spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses (SPIDER) method used for femtosecond-scale pulses [117] and the reconstruction of attosecond beating by interference of two-photon transitions (RABBITT) method used for attosecond-scale pulses [118]. Another notable example

of a diagnostic, that is used to characterize femtosecond-scale pulses, is frequency-resolved optical gating (FROG) - originally developed by Rick Trebino [119]. Over the decades since its invention, this diagnostic has become so ubiquitous that there are multiple manufacturers supplying commercial versions of it. However, with most current commercial single-shot FROG devices, the maximum pulse length that can be measured is limited to $\sim$ 5ps [120], thus leaving a sizeable interval of pulse lengths with no readily available method for characterisation.

This chapter will discuss a novel single-shot FROG device based on a second-harmonic generation that has been developed for the short pulse of the Vulcan laser at the Central Laser Facility (CLF) [121]. This instrument can measure pulses up to a theoretical limit of 25 ps. This device was commissioned specifically because commercially available FROGs could not be used on the Vulcan short pulse. As previously discussed, high power laser facilities, such as the CLF, are home to many important high energy density (HED) scientific experiments. To be able to properly interpret the results of experiments conducted at these facilities [121], accurate characterisation of the laser pulses is required.

The raw data output from FROGs do not immediately give us the required information about the pulses they are measuring. A specific algorithm is required to extract the information on the temporal evolution of the intensity of the pulse as well as its phase. This process is called phase retrieval. Traditionally, FROG retrieval algorithms have been based on an iterative Fourier transform loop. However, these algorithms tend to stagnate in the myriad of local minima that generally plague the solution space of possible laser pulses. To that end, a novel hybrid genetic algorithm has been developed and tested which circumvents the stagnation problem by iterating between randomly searching the solution space and optimising the current solution in a particular region of it.

This chapter is structured as follows: first, the general principles behind FROG diagnostics are discussed. Then, the design of the single-shot FROG diagnostic device is presented. Following that, the new genetic algorithm is described, as are its advantages over the traditional iterative Fourier transform algorithms. Finally, some experimental results are presented for two distinctly different pulse measurements and the limitations of the device in range and resolution are discussed.

6.2 The FROG diagnostic

Ignoring the separate issue of polarisation, the problem of fully characterising an ultra-short electric pulse can be reduced to retrieving its scalar electric field $E(t)$. The latter can be written in complex notation as

$$E(t) = \sqrt{I(t)}\exp(i\phi(t)) \tag{6.2.1}$$

where $I(t)$ is the function which describes the temporal evolution of its intensity, while $\phi(t)$ describes the evolution of its phase. Or equivalently, since they are related by a simple Fourier transform, it would suffice to know scalar field in the spectral domain $\tilde{E}(\omega) = \mathcal{FT}(E(t))$

$$\tilde{E}(\omega) = \sqrt{S(\omega)}\exp(i\phi(\omega)) \tag{6.2.2}$$

where $S(\omega)$ is the spectrum of the pulse and $\phi(\omega)$ is its spectral phase.

Although it might seem that the spectral field is relatively simple to characterise since it no longer has the ultra-short time dependence, this is not necessarily the case. When a pulse is spectrally measured, such as in a spectrometer, one generally only acquires the spectrum $S(\omega)$ with no access to the spectral phase. This is because when any beam of light is measured, one can only measure its intensity since that is what electronic sensors respond to. This means that in order to use this measurement, one needs to be able to retrieve $E(t)$ from just the magnitude of its Fourier transform $|\tilde{E}(\omega)|$. Mathematically, this problem is known as phase retrieval.

For a one-dimensional function like the electric field, it can be shown that this has no unique solution in the vast majority of cases. There are a few ambiguities that are immediately obvious, such as shifting $E(t)$ by a global phase $e^{i\phi}$, displacing it in time and inverting it in time. However, even up to these ambiguities, a given solution is not unique. To see why, consider the case of a discretised electric field E_t with $t \in \{0, 1, \cdots, N-1\}$. Its discrete Fourier transform (DFT) can be written as

$$\tilde{E}_k = \sum_{j=0}^{N-1} E_j \exp(-i\frac{2\pi}{N}jk) = \sum_{j=0}^{N-1} E_j z^j \tag{6.2.3}$$

where $z = \exp(-i(2\pi/N)k)$ provides a concise way of expressing the DFT as a sum of complex polynomials. The fundamental theorem of algebra tells us that any on-dimensional complex polynomial can be factorised in such a way that it can be expressed as:

$$\tilde{E}_k = \prod_{j=0}^{N-1}(z - z_j) \tag{6.2.4}$$

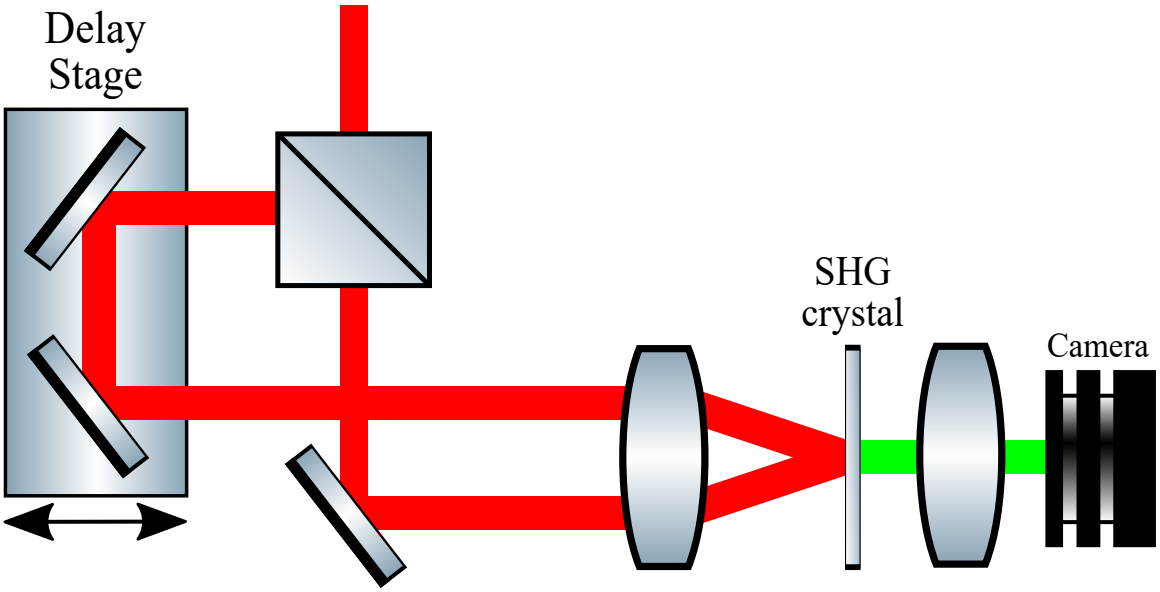

Figure 6.1: Basic schematic of a generic SHG intensity autocorrelation diagnostic.

Since one only has access to $|\tilde{E}_k|^2 = \tilde{E}_k\tilde{E}_k^*$, replacing one of the factors in the decomposition of $\tilde{E}_k$ with its complex conjugate leaves the magnitude of the observed signal unchanged. In fact, it has been shown that there are up to 2^{N-2} vectors of length N which share the same magnitude of their DFT.

Even if one where to abandon trying to retrieve the phase evolution $\theta(t)$ and focus just on $I(t)$, this task remains quite difficult. A common way of circumventing the issue of the relative slowness of electronic response is to try to measure the light pulse with a copy of itself that is shifted in time. This is called an intensity autocorrelation measurement. One of the most common methods of generating these signals is to mix the signal and its time-shifted copy on a second-harmonic generating (SHG) crystal. The output of the crystal can then be modelled as $E(t)E(t-\tau)$, where τ is the shift in time between the two copies. By imaging this signal on a camera, one obtains the intensity autocorrelation signal $A^{(2)}(\tau)$ which can be mathematically expressed as

$$A^{(2)}(\tau) = \int_{-\infty}^{\infty} I(t)I(t-\tau)dt \tag{6.2.5}$$

A basic schematic for such a system can be seen in Fig.(6.1). There, a pulse is split into two parts using a 50/50 beamsplitter and one half is propagated through a variable delay stage. The two halves are then recombined on a SHG crystal and the frequency-doubled signal is measured using a camera or photodetector. By varying the delay, one then constructs a discretised autocorrelation signal. If the mathematical shape of the pulse is known, then the autocorrelation can give information about

some of its parameters. For example, if the input pulse is known to have a Gaussian intensity envelope

$$I(t) = \exp\left[-\left(\frac{2\sqrt{\ln 2}t}{\Delta\tau_I^{\mathrm{FWHM}}}\right)^2\right] \tag{6.2.6}$$

then its autocorrelation should also have a Gaussian shape

$$A^{(2)}(\tau) = \exp\left[-\left(\frac{2\sqrt{\ln 2}\tau}{1.41\Delta\tau_I^{\mathrm{FWHM}}}\right)^2\right] \tag{6.2.7}$$

then its temporal full-width at half-maximum τ_I^{FWHM} can be estimated by fitting the measured autocorrelation signal. The same is true for other common pulse shapes such as those following a sech^2 or Lorentzian distribution. Since laser pulses usually have a theoretical model which describes their output, the intensity autocorrelation is frequently used to characterise the pulse length of standardised laser systems.

However, if the assumption about the shape of the laser pulse is incorrect, then the estimated width can be wildly incorrect. Without some knowledge about the shape of the pulse, it becomes impossible to characterise it from its intensity autocorrelation. Even though one is using a sufficiently fast signal to sample the intensity of the pulse, namely a copy of itself, this signal does not provide sufficient information about the pulse in order to retrieve even just $I(t)$. To see why, consider the Fourier transform of Eq.(6.2.5):

$$\tilde{A}^{(2)}(\omega) = \mathcal{FT}(I(t) * I(t)) = |\tilde{I}(\omega)|^2 \tag{6.2.8}$$

Since knowing the signal is equivalent to knowing its Fourier transform, one is left again with the phase retrieval problem to obtain $I(t)$ from its autocorrelation. This means that an autocorrelation does not map to a unique pulse, even if one accounts for the above-mentioned trivial ambiguities. This can be seen from Fig.(6.2), where two simulated pulses with different intensity profiles have exactly the same autocorrelation. Another common problem with autocorrelation signals is that the more complex a pulse shape becomes, the more its autocorrelation "washes" out the high-frequency features. This makes this measurement unusable when one wants to characterise the laser pulse after an interaction.

In order to resolve all these issues, Rick Trebino and Daniel Kane invented the so-called Frequency-Resolved Optical Gating (FROG) method [119]. The setup of the FROG diagnostic is relatively straightforward. It starts, like an intensity autocorrelator, by mixing a pulse with a time-shifted copy of itself in a non-linear process. However, instead of directly imaging the output, it is first Fourier transformed by measuring its spectrum, which is then imaged on a camera. The basic setup of a

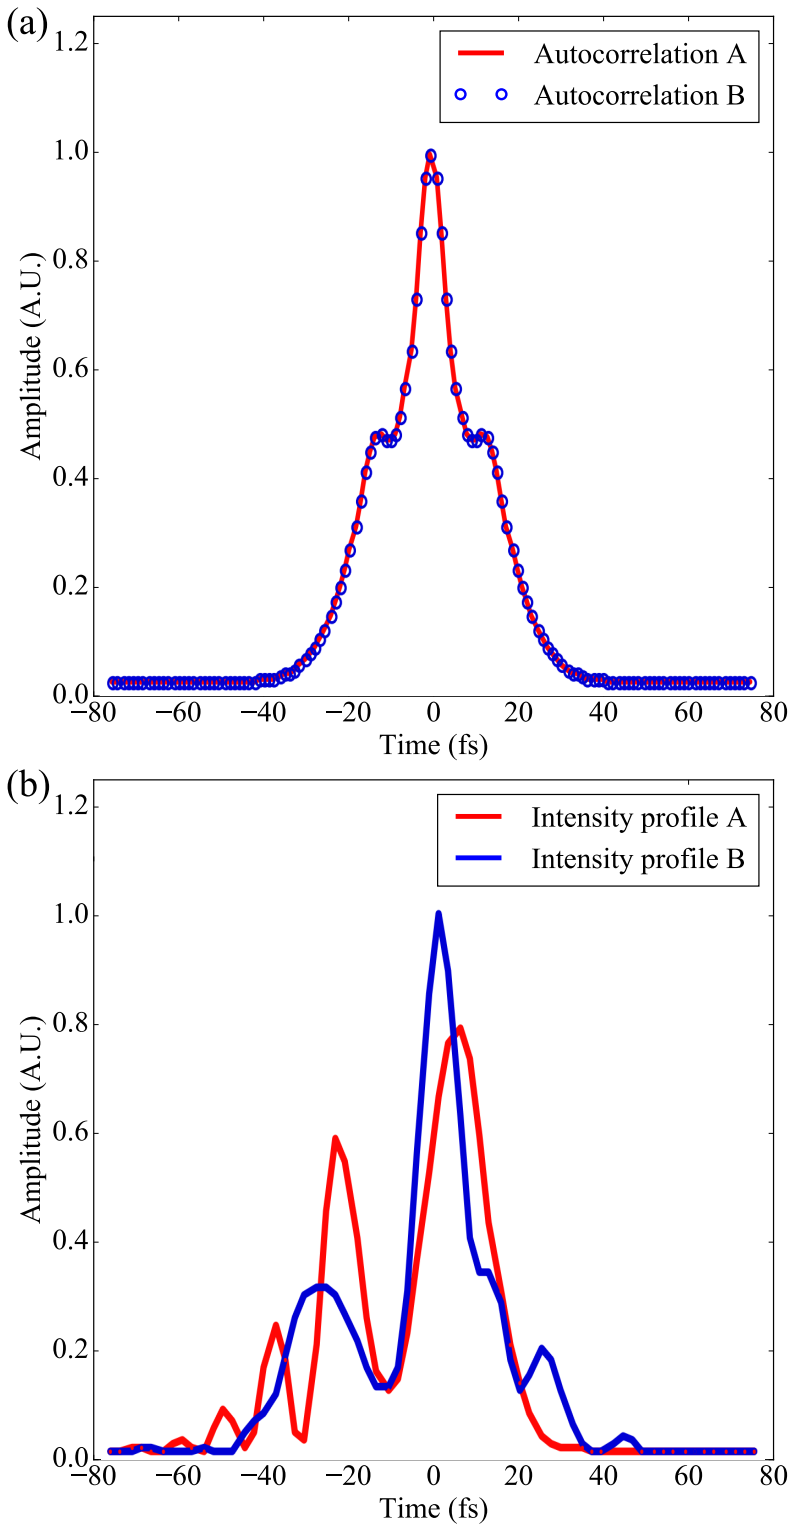

Figure 6.2: Simulated data for two pulses (A and B) with different intensity profiles: (a) The intensity autocorrelation signal for the two pulses; (b) The temporal intensity profile of the two pulses. This shows that although they have identical autocorrelation signals, they have vastly different intensity profiles.

SHG-FROG can be seen in Fig.(6.3), which is basically the same as the SHG intensity autocorrelation setup but with the camera replaced by spectrometer.

As the shift in time of the copy is varied, the 2-dimensional spectrogram of the pulse $S(\omega, \tau)$ can thus be formed. The actual measured signal is the squared magnitude of this spectrogram, which is conventionally called the FROG trace $I(\omega, \tau)$ and can be mathematically expressed as:

$$I(\omega, \tau) = |S(\omega, \tau)|^2 = \left| \int_{-\infty}^{\infty} E(t)E(t-\tau)e^{-i\omega t}dt \right|^2 \tag{6.2.9}$$

It may seem that we are back in the same old spot, having to retrieve a signal from the magnitude of its Fourier transform. However, counter-intuitively, this 2-dimensional phase retrieval is an easier problem to solve than the previously discussed 1-dimensional case. One way to see why is that, although the fundamental theorem of algebra specifies that one can expand the 1-D DFT into factors such as in Eq.(6.2.4), there is no equivalent theorem in 2-D. Thus, it can be shown that, given a measured FROG trace, 2-D phase retrieval gives, in most cases, a unique solution for the spectrogram $S(\omega, \tau)$. This is, of course, up to the trivial ambiguities that were outlined before. However, these ambiguities do not affect the physically relevant fields. This means that, given the right retrieval algorithm, it is possible to use the FROG diagnostic to not only retrieve the intensity of the pulse $I(t)$, but also to retrieve its phase $\theta(t)$ and characterise it completely. This has made the FROG diagnostic one of the most used measurements in ultra-fast laser experiments.

6.3 Design

One disadvantage of traditional FROG systems, such as the one described in Fig.(6.3), is that to generate the full trace, multiple shots are needed to capture the signal at different delays. Therefore, the temporal resolution is limited by the size of the steps of the variable delay stage. This is especially problematic when one wants to use this diagnostic to characterise high-power laser pulses where the repetition rate of the laser is often less than 1 Hz and where shot-to-shot reproducibility is rarely guaranteed. For that reason, the FROG device that was built for the Vulcan laser was based on a single-shot design capable of acquiring the full FROG trace in only one shot.

My design is based on second-harmonic generation, where the pulse halves are recombined on a barium borate (BBO) crystal to generate the field autocorrelation in the second-harmonic pulse that is produced by the crystal. However, instead of

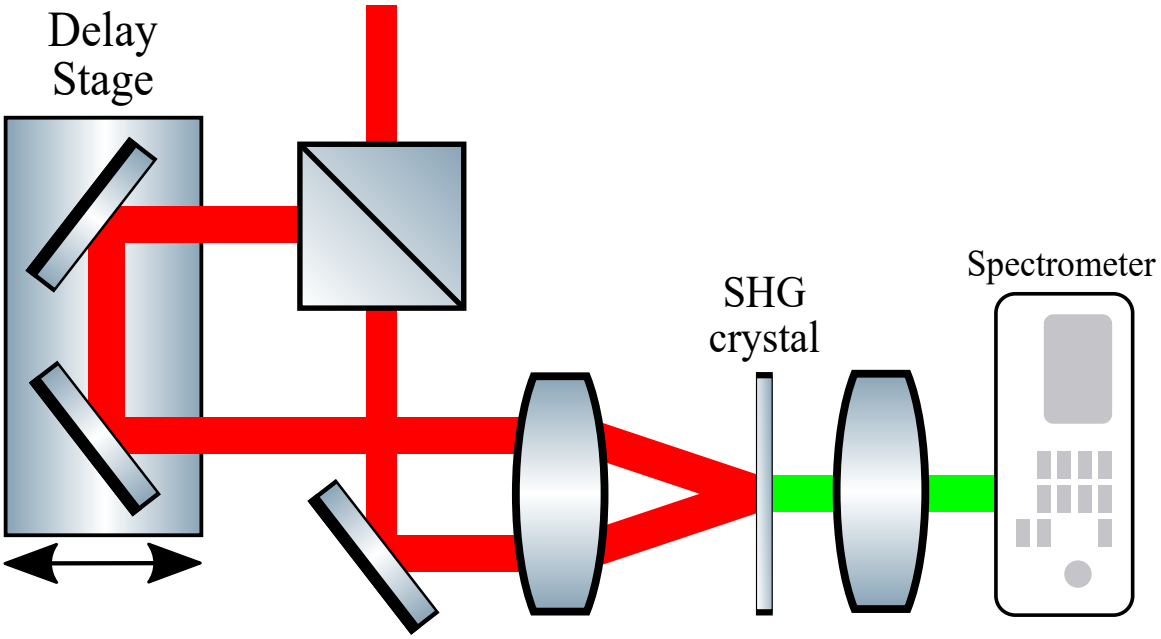

Figure 6.3: Basic schematic of a generic SHG FROG diagnostic.

using a variable delay stage, the pulses are recombined on the crystal at an angle. As shown in Fig. 6.4, this causes different points on the crystal's axis to have different delays between the pulse halves. That way, the time delay axis of the autocorrelation is mapped onto the transverse spatial axis of the crystal, thus getting a full autocorrelation in a single shot. The temporal window of the FROG can be expressed as:

$$\tau(x) = 2(x/c)\sin(\theta/2) \tag{6.3.1}$$

The maximum temporal window of the FROG can then be shown to be $\tau_{max} = (2d/c)\tan(\theta/2)$ [122], where d is the diameter of the beam and θ is the recombination angle between the two pulse halves as shown in Fig.(6.4). This technique allows for a full FROG trace with each shot. Both the traditional scanning FROG devices and their single-shot counterparts have their advantages and disadvantages. The former is agnostic to the quality of the transverse profile of the beam and does not require the use of large imaging spectrometers. However, as the latter nets the full FROG trace in a single-shot, it is better suited for use with high power laser systems as the previously mentioned low repetition rate and the significant variation between shots would make the use of the scanning FROG highly impractical.

The single-shot FROG has been used extensively in many experiments across a varied number of scientific disciplines [123, 124, 125]. However, with most current commercially available single-shot FROG devices, their temporal window is limited to

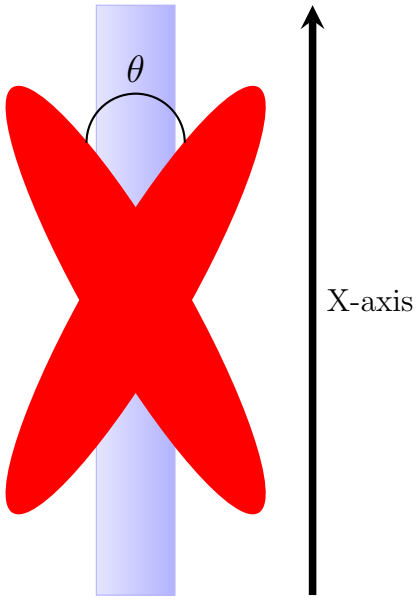

Figure 6.4: Diagram showing the two pulses mixing at angle θ on the second-harmonic generation crystal. The bigger the angle, the larger the maximum temporal window of the autocorrelator.

around ~ 5 ps [120]. In order to be able to measure longer pulses, a diffraction grating has been used to add pulse-front tilt to the laser pulse. By doing so, the recombination angle increases, increasing the maximum temporal window up to a theoretical limit of 25 ps. It should be noted that there is a trade-off between the size of the temporal window and the temporal resolution of the device. This can be adjusted by changing the diffraction order used by the autocorrelator. With the zeroth order, there is no pulse-front tilt and thus the temporal window is at its minimum. Further, by using higher orders, the size of the window is increased.

Fig.(6.5) shows a schematic of our FROG design, clearly indicating that the autocorrelation is then split using a 50/50 beamsplitter, where one half of the signal is imaged onto a charged coupled device (CCD) camera. This camera helps with the alignment of the pulse in the autocorrelator so that second-harmonic generation can first be confirmed before aligning the rest of the device. The second half is then focused using a cylindrical lens onto the slit of a Princeton Instruments Spectra Pro SP300 imaging spectrometer to measure the Fourier transform of the autocorrelation along one axis. This is because for a single-shot measurement, one axis contains the spectral information and the other contains all the time delay information. Due to the design of the slit of this spectrometer, the second harmonic had to be inverted using a periscope in order to obtain the proper axis on the grating, as shown in Fig.(6.6).

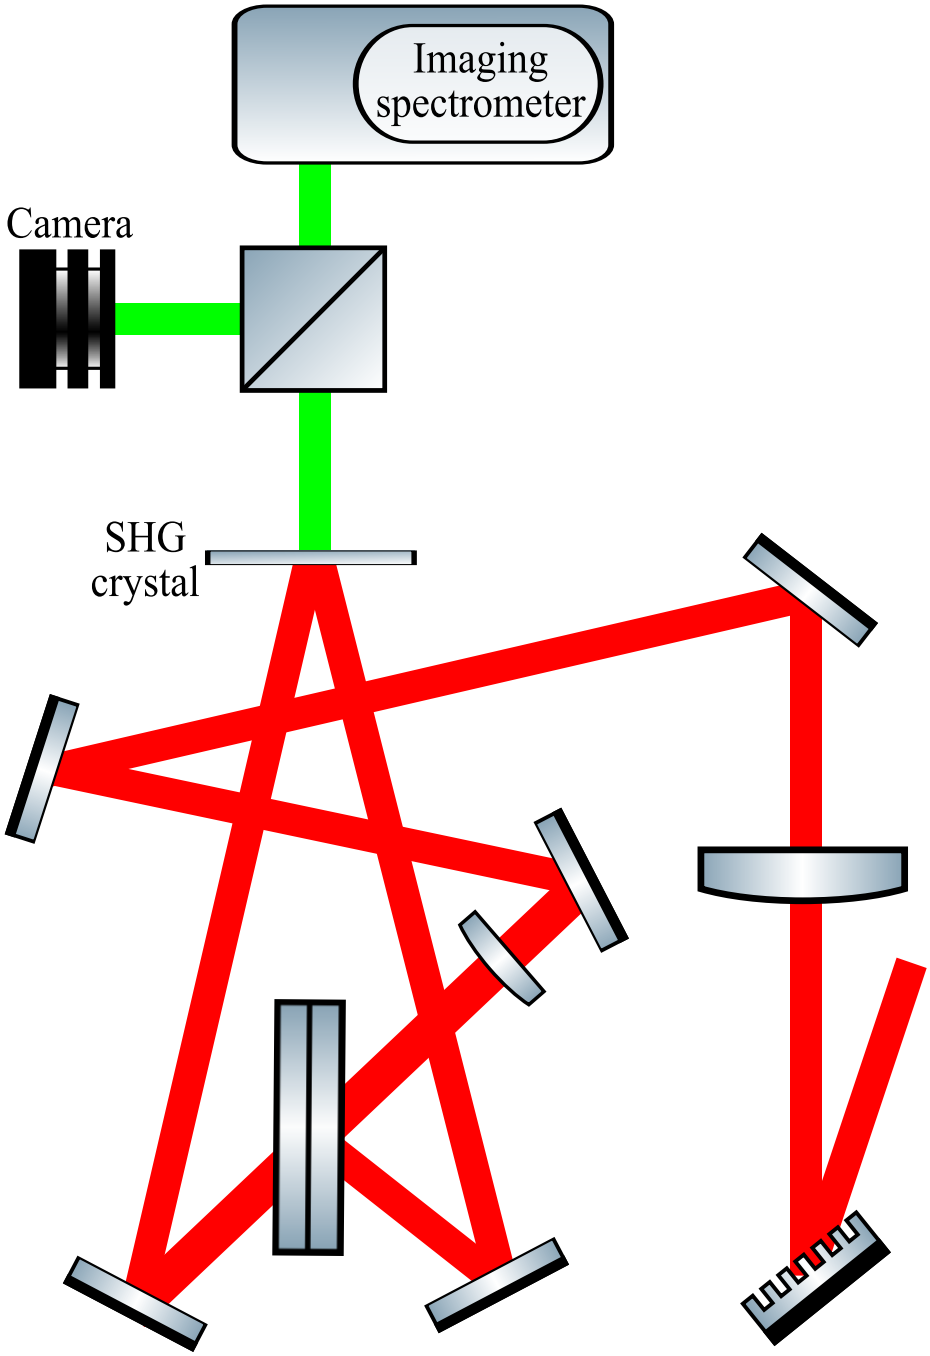

Figure 6.5: Schematic detailing the major optical components of the FROG.

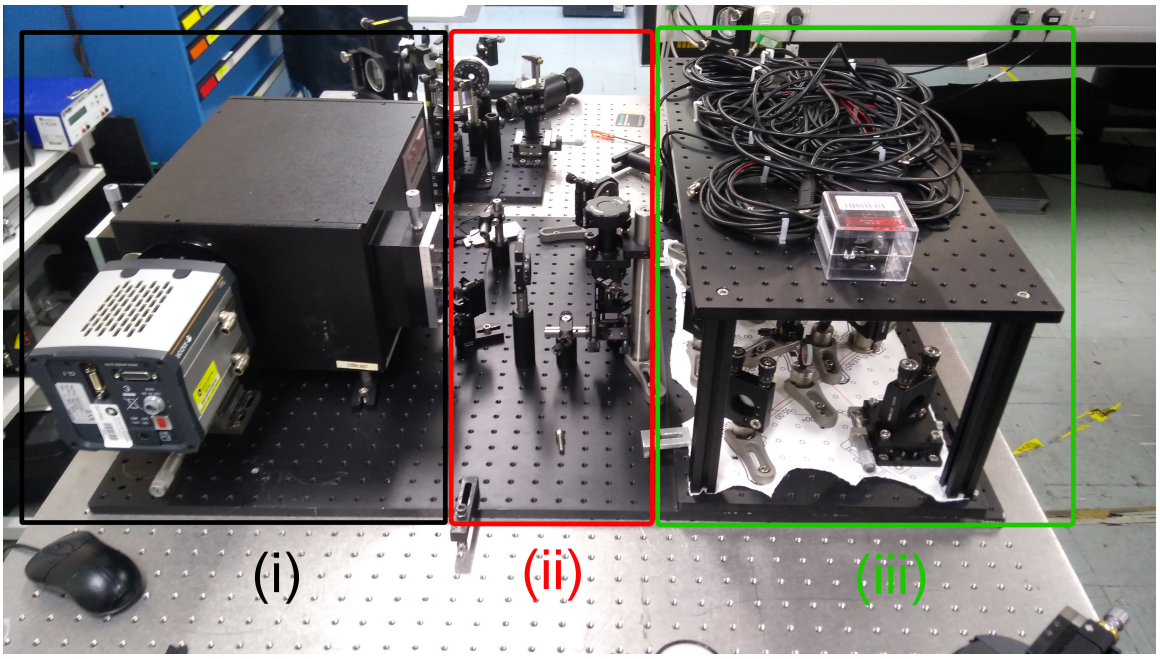

Figure 6.6: Picture showing the completed FROG device: (i) The second-harmonic generation autocorrelator used to generate the second harmonic signal for the FROG; (ii) The periscope and cylindrical lens used to flip the signal in the proper orientation and focus it onto the spectrometer's slit; (iii) The PI SP-300 spectrometer with the Andor NEO camera used to generate the FROG signal.

The device was aligned using a high-repetition oscillator that produces short pulses that are only visible using an infrared (IR) viewer. The challenge with aligning the second-harmonic into the spectrometer lies within the fact that the former is very weak. To combat this, a green diode laser is fed into the device and aligned with the second-harmonic. This is accomplished by putting two iris and using the autocorrelation camera to ensure that the green laser and the second-harmonic are co-propagating. Using this method, I have managed to transport the device between different locations on the site and align it successfully at these different locations using low-power high repetition rate pulsed sources.

6.4 Hybrid genetic algorithm

As discussed earlier, it is possible to recover the temporal evolution of both the intensity and the phase of an ultra-fast pulse from its measured FROG trace up to the typical ambiguities. In the case of the second-harmonic generation-FROG, there is an additional ambiguity as the direction of time cannot be determined due to the time reversal symmetry of the second-harmonic generation process.

Although it can be shown that the 2-D phase retrieval required to analyse FROG traces usually produces a unique solution, this does not mean that finding this solution is straightforward. Much research has been done into designing efficient and effective retrieval algorithms. This is because 2-D phase retrieval is not only important for FROG characterisation, it is also indispensable for important experimental techniques such as x-ray crystallography and coherent diffractive imaging. Traditional FROG algorithms are usually based on an iterative Fourier transform/projection process that seeks to minimise an objective function, which is measuring the error between a numerically simulated FROG trace (generated from an initial guess using Eq.(6.2.9)) and the experimentally measured trace. The most commonly used function for this is the mean-squared error which is defined at the i$^{\text{th}}$ iteration to be [126]:

$$F(E^{(i)}) = \sqrt{\frac{1}{N^2}\sum_{j,k=0}^{N} |I_{FROG}^{meas}(\tau_j,\omega_k) - I_{FROG}^{(i)}(\tau_j,\omega_k)|^2} \tag{6.4.1}$$

Where $I_{FROG}^{meas}(\tau_j,\omega_k)$ is the value of the measured FROG trace and $I_{FROG}^{(i)}(\tau_j,\omega_k)|$ is the value of the simulated FROG trace generated from the i-th iteration of the estimation of the pulse's electric field $E^{(i)}$. Since these functions are discrete, the temporal and spectral axes are labelled by j and k respectively. The iterative algorithm then seeks to find the electric field $E(t)$ that generates a FROG trace resembling most

the trace that was measured. This method is known to generally converge well if the objective function is convex over the space of possible solutions [126]. Unfortunately, this is often not the case for FROG measurements.

To illustrate this, the FROG trace of a simulated unchirped double pulse, shown in Fig.(6.7(a)), was numerically generated. Many guesses "close" to the true signal were then generated by varying two parameters: the duration of the double pulse and its chirp. The objective function was computed for all of these guesses. In general, the actual solution space has many more than two dimensions; pulses can vary not only in terms of the parameters that govern their mathematical functional expressions but also in the functions themselves. Here, I have restricted myself to these two parameters for purely illustrative reasons. The results for this parameter scan are shown in Fig.(6.7(b)).

It is clear from Fig.(6.7(b)) that, in addition to the global minimum at the correct parameters, there are multiple local minima at other parameter combinations. Within these minima, the traditional algorithm can stagnate and never reach the best possible estimation for the pulse. To combat this, a hybrid genetic algorithm that is based on a routine used in phase retrieval in x-ray diffraction [127] (a mathematically similar problem that faces many of the same obstacles) was used. Instead of starting with a single guess, this algorithm initialises many guesses simultaneously, which means it explores multiple areas of the solution space at the same time. After performing a few iterations of the traditional iterative algorithm on each guess, it selects a sample of them quasi-randomly (with a bias for the "closest" guesses) and combines them in order to generate the population for the next iteration. There is also a probability of a random "mutation" happening that alters some components in the guesses. A beneficial consequence of these steps is that it is very difficult for the algorithm to be stuck in a local minimum since it is constantly stochastically sampling different areas of the solution space, thus driving it out of any local minima.

This stochastic exploration can be quite slow, especially when very large high-dimensional solution spaces are considered, such as the ones for laser pulses. However, combined with an efficient local improvement step that exploits each guess, the algorithm runs quite efficiently [128]. As such, I developed a hybrid algorithm that combines the stochastic element of genetic evolutionary algorithms with the traditional iterative FROG algorithm as the local improvement step. Another advantage of this hybrid algorithm is that it can be fine-tuned so that it focuses more on the stochastic element or on the local improvement step. The flowchart in Fig.(6.8)

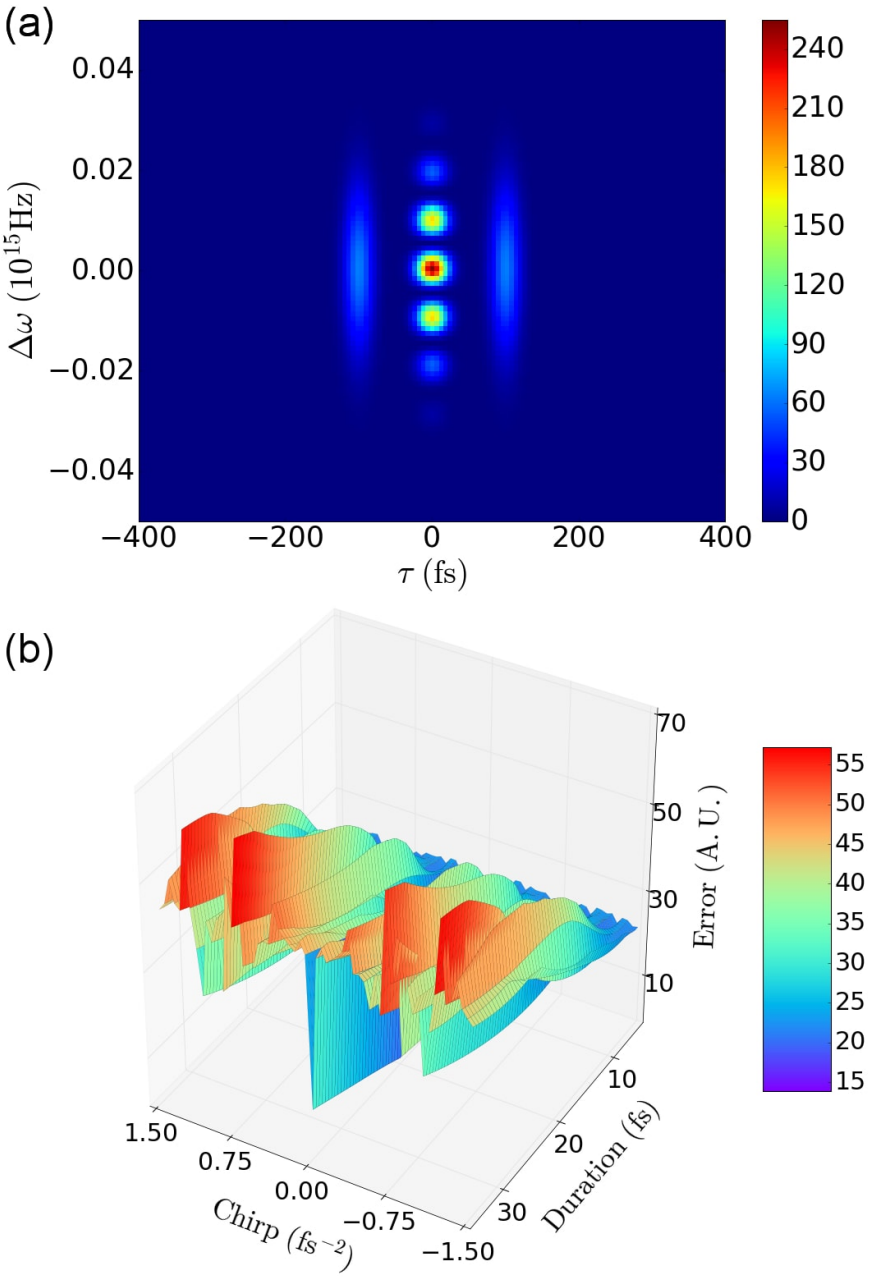

Figure 6.7: (a) Numerical FROG trace produced from a double pulse showing the distinct interference pattern that is usually used to calibrate FROG devices (b) 3D surface plot of the objective function across the two parameter scans of the guesses "around" the double pulse.

provides the basic overview of one iteration of the retrieval routine, which is then expanded upon in subsections 1-4 below.

6.4.1 Initialization

The first step in the algorithm is to generate an initial population of guesses. To do so, I start with some initial guesses for the pulse, $E_{init}(t)$. The initial population $\{E_i(t)\}_{i=1,..,N}$ is then generated by randomly changing the phase at each timestep of each pulse.

$$\begin{aligned} E_i(t_j) &= |E_{init}(t_j)| \exp(\mathrm{i}\phi_j) \\ \phi_j &= \arg(E_{init}(t_j)) + R \cdot \mathrm{rand}[-\pi, \pi] \end{aligned} \tag{6.4.2}$$

The parameters that govern this step are $N \in \mathbb{N}$, the size of the population used in the algorithm, and $R \in [0, 1]$, the degree of randomness in generating the initial population from the initial guess. N can greatly affect the performance of the algorithm. By increasing N, there are more guesses that explore the solution space. However, this comes at the cost of considerably slowing down the algorithm as there are many elements that have to go through the local improvement.

6.4.2 Selection

This is a crucial step in the algorithm, since it is the process by which elements of the population are chosen to become "parents" for the next generation. First, the current population is ordered with respect to the objective function (in our case the mean squared error) from best to worst. The algorithm then draws from a biased distribution to obtain the ranking of the element that will be selected. A bias that heavily favours the best elements will speed up the convergence. However, it increases the risk of the algorithm becoming stuck in a local minimum as it does not sufficiently explore the solution space. On the other hand, while choosing elements completely at random increases the exploration, it runs the risk of taking too much computing time. The ranking is drawn using the following distribution:

$$n = \left\lfloor (\mathrm{rand}(0, 1))^{\beta}) \cdot N \right\rfloor + 1 \tag{6.4.3}$$

Where $\mathrm{rand}(a, b)$ is the uniform distribution over the interval (a, b). The parameter that controls this step is $\beta \geq 1$, which adjusts the bias of the distribution.

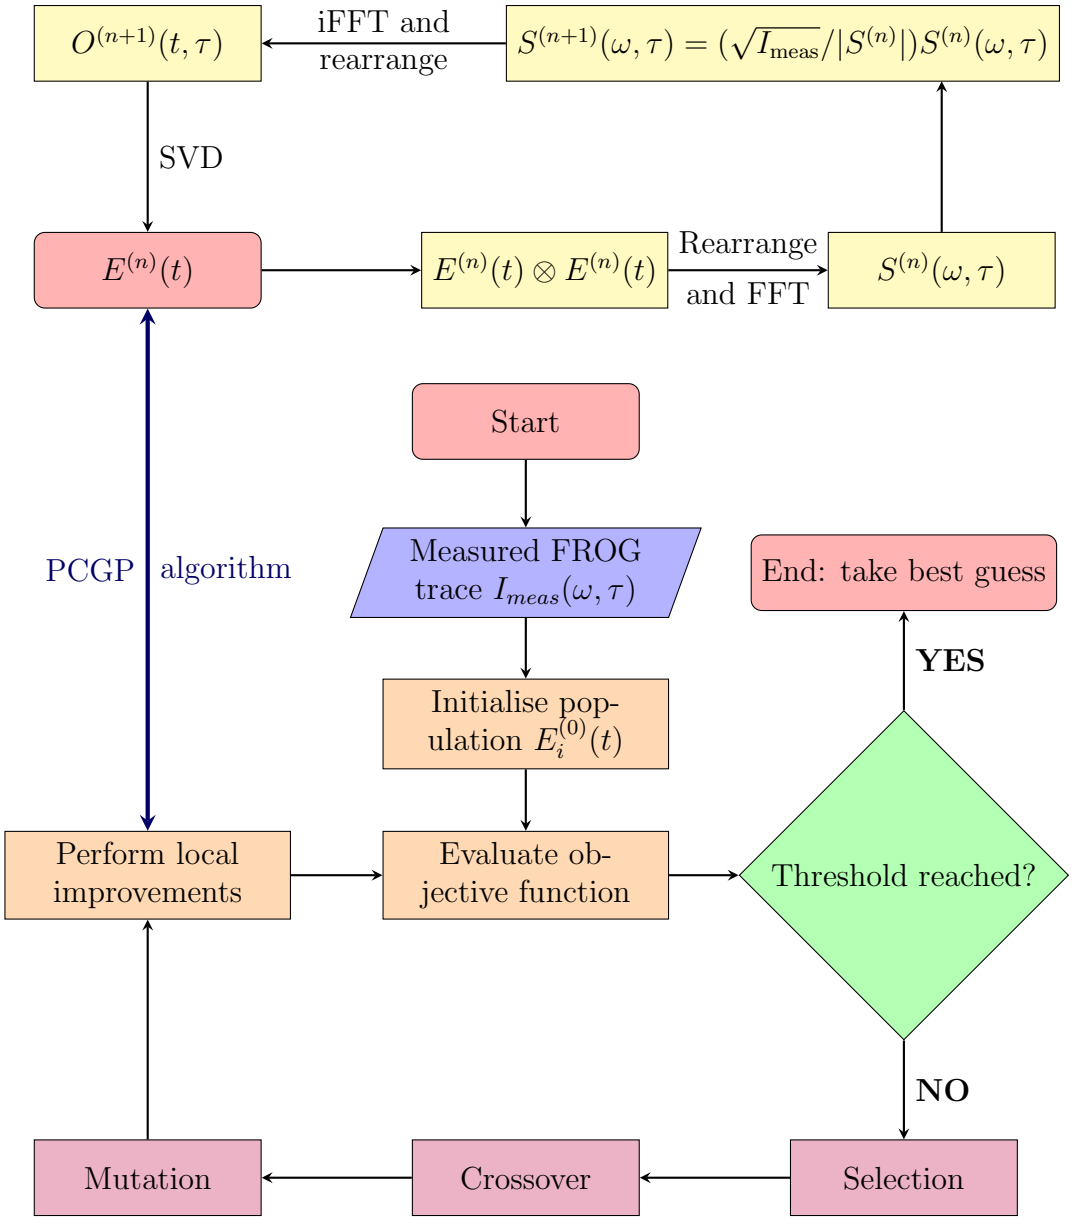

Figure 6.8: Flowchart of the hybrid algorithm. Purple rectangles represent steps of the genetic algorithm yellow represent steps of the PCGP algorithm and orange rectangles are the intermediate steps.

6.4.3 Crossover

In this step, the next generation of guesses is produced. Inspired by the genetic crossover that is the driver of natural evolution[129], differential crossover randomly combines different components from the parents. Each time I want to generate an offspring, four parents are chosen using the selection procedure outlined above. The offspring are then created following:

$$E_{new}(t_j) = \begin{cases} E_1(t_j) & \text{if rand}[0,1) \geq C \\ E_2(t_j) + D \cdot (E_3(t_j) - E_4(t_j)) & \text{else} \end{cases} \tag{6.4.4}$$

Where $C \in [0,1]$, called the balancing coefficient, controls how much crossover happens at each iteration and D, called the differential coefficient, controls the contribution of the third and fourth parent.

The number of offspring N_o can in fact be smaller than the current population size, such that $N_o = G \cdot N$, where $G \in [0,1]$ is usually called the genetic function [127]. In those cases, the next generation is supplemented with elements from the current one, which are also chosen using the biased distribution outlined in the previous section, so that all the generations have the same size as the initial one.

6.4.4 Mutation

In this step, for each element, the phase of each time step is modified by adding a random phase.

$$E_i(t_j) = \begin{cases} E_i(t_j) & \text{if rand}[0,1) \leq M \\ E_i(t_j)\exp(\mathrm{i}R_M \cdot \text{rand}[-\pi,\pi]) & \text{else} \end{cases} \tag{6.4.5}$$

Where $M \in [0,1]$ is the mutation coefficient, which controls the probability of a mutation occurring at any given time step. $R_M \in [0,1]$ controls the range of the added phase, hence controlling the 'degree' of the mutation. This step increases the exploration of the solution space to regions that might not have been reached in the initialisation step by randomly pushing the members of the population towards those regions.

6.4.5 Local improvement

The last step in an iteration of the hybrid algorithm is subjecting each element of the current generation to several iterations of a traditional retrieval algorithm. In our case, the iterative Principal Component Generalised Projections (PCGP) algorithm, which is well suited for the second-harmonic generation FROG [130], was used. For

each guess vector $\mathbf{E}$, an autocorrelation matrix is created by computing the outer product $\mathbf{E}\mathbf{E}^{\mathrm{T}}$. By rearranging the matrix and taking the squared magnitude of its Fourier transform, a numerical FROG trace can be created, which can then be compared to the measured one. If the guess is not "close" enough, the magnitude of the rearranged matrix is replaced with the square root of the measured FROG. By performing a singular value decomposition (SVD) on the resulting matrix, it can then be written as a product of two orthogonal matrices and a diagonal matrix. By taking the columns corresponding to the largest singular value in the diagonal matrix, the "improved" guess is attained. While this process can be computationally costly, since larger matrices are used, it can be sped up by implementing what is called the power method [130].

The parameter that controls this step is N_i, which determines the number of iterations of the PCGP algorithm that are run on each generation. At the end of the iteration, the objective function is again computed for all the members of the current population and the "best" guess is kept in memory as the routine goes into the next iteration.

6.4.6 Practical difficulties

A number of difficulties usually arise with this algorithm, the first being the choice of the parameters. Due to the large number of parameters that significantly impact the result of the algorithm, trial and error must be employed to determine the optimum set of parameters for each class of "similar" pulses. Also, as with most FROG retrieval algorithms, since there is an iterative Fourier transform step, the size of the temporal window and the frequency resolution are intrinsically related, as are the extent of the frequency window and the time resolution. More explicitly, if a large temporal window is required, the frequency resolution must be small enough to enable its calculation and vice versa, or else preforming the numerical Fourier transform would be impossible.

6.4.7 Comparing the hybrid algorithm to traditional algorithms

To compare the performance of the hybrid algorithm relative to the traditional PCGP algorithm, both of them were used on the FROG trace generated by the double pulse used above to showcase the solution space in Fig.(6.7(a)). The parameters used for

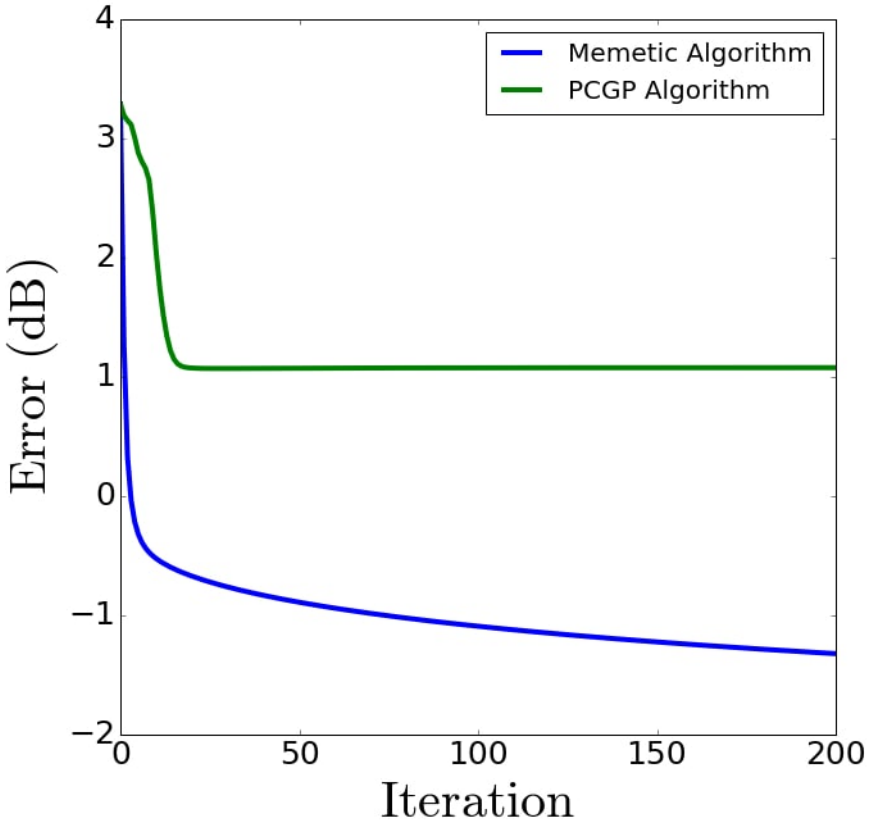

Figure 6.9: The graph shows the different progression for the two algorithms as they try to retrieve the original simulated double-pulse from its FROG trace. It should be noted that each iteration of the memetic algorithm contains multiple iterations of the local improvement using the PCGP.

the memetic algorithm were: $N = 40, N_i = 10, R = 0.5, M = 1, C = 0.5, D = 0, G = 1, \beta = 0.15$.

As can clearly be seen in Fig. 6.9, the traditional algorithm encounters a local minimum and is unable to improve beyond it, no matter how many iterations are performed. Conversely, the memetic algorithm does not become stuck and is able to continue and reduce the error by more than two orders of magnitude over ∼100 iterations.

It should be noted that the hybrid algorithm runs in ∼ 30 minutes on a windows laptop with an Intel Core i7 processor clocked at 2.6 GHz. The main parameters that affect its runtime are N, N_i, the number of iterations of the full algorithm and the size of the FROG trace, the effect of the latter of which can be mitigated by using the power method instead of the SVD.

6.5 Experimental results

The device was tested in two very different conditions, with the first measuring the output of a femtosecond oscillator that produces a train of pulses at a 50 Hz repetition rate with a pulse duration on the order of hundreds of femtoseconds. This oscillator is the source that was used to first calibrate the FROG. The calibration was done by introducing a prepulser between the oscillator and the device. The prepulser splits each pulse into two with equal energies, separated in time by a controllable delay. When two pulses enter the FROG, the trace has a very distinct shape, similar to the one shown in Fig.(6.7(a)), which has three separate 2D peaks separated in time by the delay between the two pulses [131]. The central peak also shows a distinct spectral interference pattern where the peaks are separated with the reciprocal of the delay. By varying the delay and measuring the effect it has on the FROG traces, I calibrated the spectral and temporal axis of the FROG trace. I found that the device had a pixel temporal resolution of 65 fs and a pixel spectral resolution of 0.01 nm. After calibrating the device, the FROG trace produced by the oscillator, without the prepulser, is then run through the hybrid retrieval algorithm. Fig.(6.10(b)) shows the results of this process. The retrieved pulse fits well with a Gaussian intensity profile with a FWHM of 390.01 fs, while the phase shows a slight chirp. The direction of this chirp is ambiguous since, as noted in the previous section, both positive and negative chirps have the same effect on the trace of second-harmonic generation-FROG. There is, however, a way to break this symmetry by introducing a small prepulse [131]. This breaks the time reversal symmetry of the pulse, which thus lifts the ambiguity in the retrieved pulse.

The device was then used to measure the pulse duration and relative phase shift of a high-power short pulse from the Vulcan laser system [121]. Fig.(6.11) shows the results of a measurement of the 10 J short pulse. It should be noted that the beam profile was smoothed using a random phase plate [132], both for the purposes of the experiment and because the single-shot FROG is quite sensitive to spatial non-uniformity. The FWHM of the retrieved intensity is 4.7 ps, which is in the expected range, with the phase showing an ambiguous chirp. I have also confirmed this measurement by taking the spectrum of the retrieved pulse and comparing it to the independently measured spectrum of the laser. It should be noted that, in all cases, the phase points of the retrieved signal that are beyond the temporal width of the intensity envelope were dropped because they convey no physical meaning. That is because these phase points represent points where the actual measured signal was

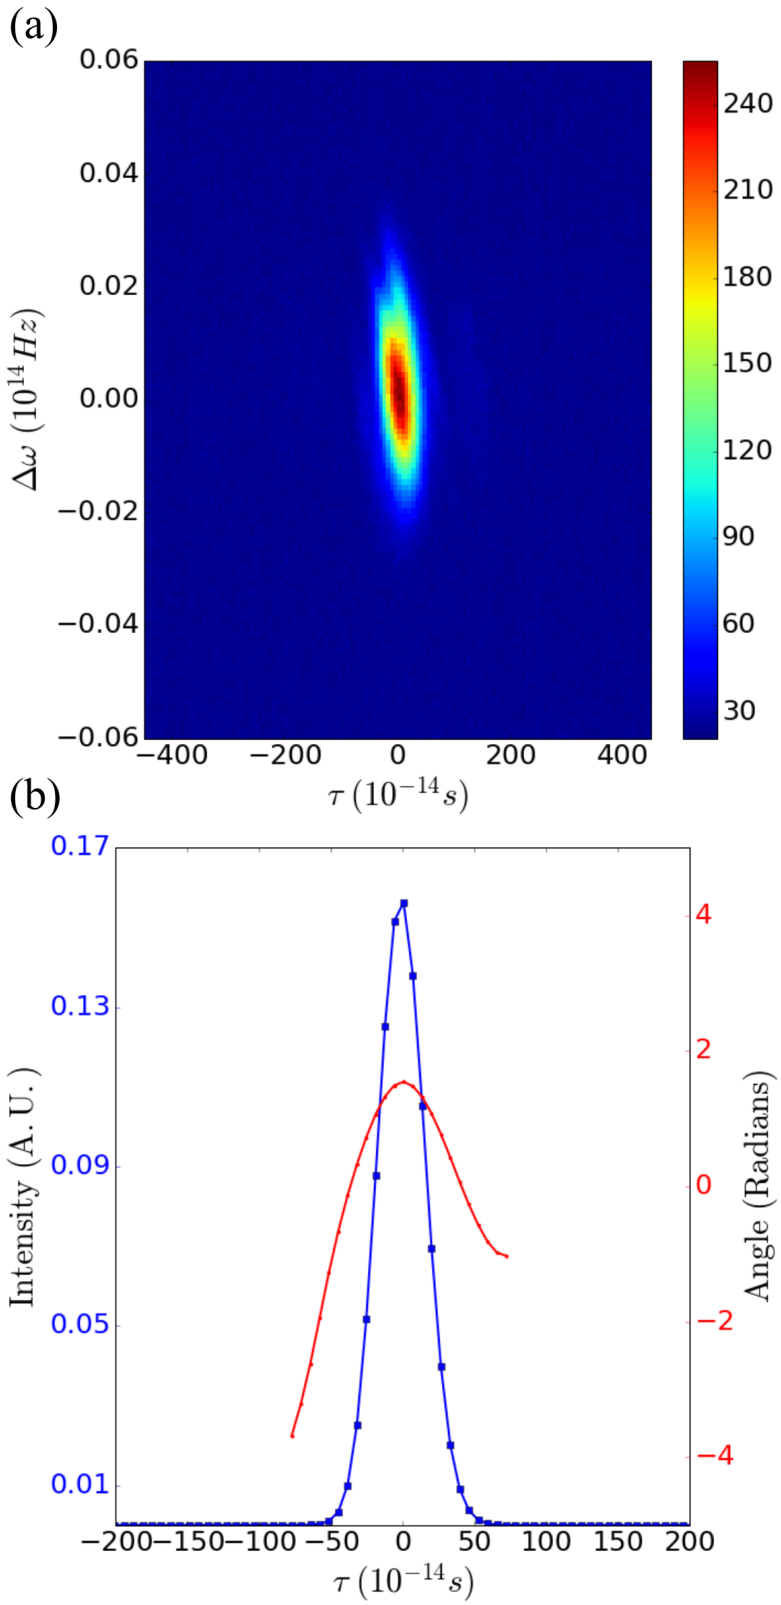

Figure 6.10: (a) FROG trace produced by the femtosecond oscillator (b) Intensity of the retrieved pulse and its phase.

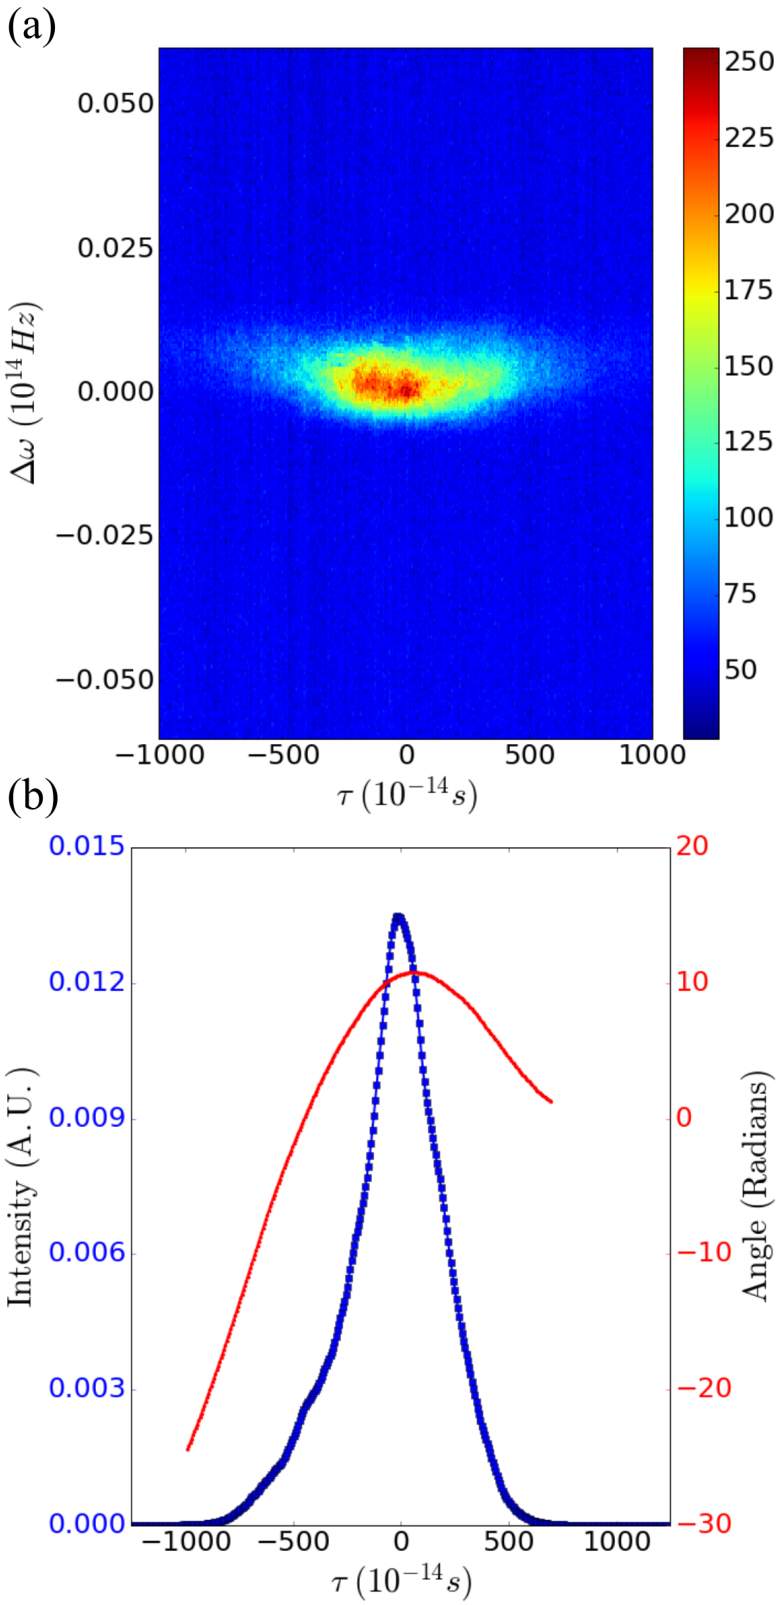

Figure 6.11: (a) FROG trace produced by the Vulcan short pulse beam (b) Intensity of the retrieved pulse and its phase.

below the intensity threshold required to register as a constituent part of the laser pulse.

6.6 Conclusion

In conclusion, a new single-shot FROG device based on second-harmonic generation was developed, which can measure the duration and relative phase of high intensity pulses up to $\sim 25\,\mathrm{ps}$ with high temporal resolution. This was achievable primarily due to the use of dispersive optical elements in order to introduce pulse front tilt to a laser pulse so as to increase the maximally measurable temporal window. This is very relevant to characterise picosecond laser pulses used in many different high energy density physics experiments. The device was developed and tested on the Vulcan laser at the Central Laser Facility, Rutherford Appleton Laboratory. The new, relatively low-cost instrument, is suitable for implementation on any short-pulse laser facilities, such as the OMEGA laser [133] and the National Ignition Facility [134].

A novel hybrid genetic algorithm was developed for the retrieval of pulses from the FROG trace. This algorithm avoids the stagnation problem that can occur with the traditional algorithm when it encounters a local minimum. By randomly exploring the solution space, the evolutionary algorithm seems to be resilient to the stagnation problem and continues to improve the guess beyond the traditional algorithm.

Chapter 7

Conclusions and future work

"*Everything has to come to an end, sometime.*"

L. Frank Baum

7.1 Summary

Over the course of this thesis, I have detailed the calculations and results supporting the following key conclusions:

- Using the standard formulations of quantum field theory, I have reproduced the derivation of the fundamental origins of the Orbital Angular Momentum (OAM) of light. I have also argued how, on the level of a single photon, OAM can be modelled as a true quantum mechanical observable, separate from the one associated with the spin of photons.

- Using the effective field theory developed by Euler and Heisenberg, I have shown how OAM modifies the vacuum photon-photon scattering interaction and creates an additional constraint to respect the conservation of angular momentum.

- Leveraging the additional conservation relation, I have proposed a novel design for an experiment to detect photon-photon scattering using high-power lasers carrying OAM. I also estimate the number of detectable scattered photons to be around 100 photons for a 10 PW input beam while reducing the impact of background radiation by an estimated 25 dB.

- I have described the different methods of generating OAM-carrying laser pulses and have shown experimental results of using continuous transmissive spiral phase plates on the ASTRA laser.

- I have shown how different methods of characterising OAM can be used on high-power laser facilities as well as discuss their differing effectiveness. I have outlined experimental results comparing each method on the ASTRA laser, showing that, while the cylindrical lens method is ideal for quick setup and analysis, it is not optimised for handling a superposition of different modes, which is a task better suited for the more complex, but more complete, projective diagnostic.

- I have outlined the design, assembly and experimental performance results of a novel Frequency-Resolved Optical Gating (FROG) diagnostic designed to characterise the pulse shape and phase evolution of ultra-short laser pulses up to a pulse length of 25 ps with a 65 fs pixel resolution.

- I have detailed a novel hybrid genetic algorithm which addresses the local minimum stagnation problem common with two-phase retrieval tasks, such as analysing FROG outputs to retrieve pulse diagnostic data. I have demonstrated that the hybrid algorithm improves upon the standard FROG retrieval algorithms and reduces the retrieval error by at least 2 dB, while never stagnating.

7.2 Future work

Based on the above, I have planned to perform several studies to leverage and expand upon the results of this thesis. The envisioned work will continue to focus on OAM in the context of high-intensity lasers and can be mostly split into two broad categories: the first is continuing the work on the impact of OAM on the vacuum photon-photon interaction and planning an experimental campaign to verify it. The second part is furthering theoretical work on the impact of OAM on laser-plasma interactions as well as planning experimental campaigns to verify the various predictions that have been made on this topic. These campaigns will leverage the findings concerning the effectiveness of various methods of producing and characterising the OAM of high-power lasers.

7.2.1 Experimental verification of photon-photon scattering

There are two main aspects in which the theoretical work on OAM in photon-photon scattering, detailed in this thesis, can be extended. Firstly, the variety of possible geometric setups for the three high-power beams that are driving the photon-photon interaction can be greatly expanded. The configuration that was chosen in Fig.(3.4)

was selected because it greatly simplified the calculation of the field that was generated via the interaction between the three beams. This was because all three beams were on the same plane, only one was carrying OAM and the two others were symmetric around it. This necessarily caused the generated photons (for a particular interaction term) to be anti-parallel with the OAM-carrying beam. This is due to the conservation of linear momentum, thus rendering the calculation of the conservation of angular momentum easier since it was restricted to a single one-dimensional axis. However, this arrangement is not necessarily the easiest geometric set-up from an experimental point of view since the generated photons are counter-propagating with a high-power beam, complicating the placement of the measurement and OAM filtration system. Allowing the beams to occupy different planes would simplify the experimental setup. Thus, further work needs to be done to investigate how the OAM is transferred with an off-axis interaction.

Another important extension of the work is to consider physically realistic laser spatial modes in the calculation of the expected scattered photons. To make it possible to reach an analytical solution to the Euler-Heisenberg Lagrangian, the input modes were considered to be either plane wave modes or the OAM extension to plane waves. However, both of these solutions technically hold infinite energy. A more realistic approach would instead consider Gaussian and Laguerre-Gaussian modes, or other modes more commonly found in high-power laser facilities, such as the supergaussian modes. This has the advantage of more accurately modelling the interaction, as it would happen in real world facilities and not requiring any artificially-imposed limits on the integration volume. Unfortunately, there is no such thing as a free lunch. Indeed, using realistic modes would make it impossible to find a generally analytical solution to the problem and would thus require a computational solution. Thankfully, since the Euler-Heisenberg wave equation, Eq.(3.4.12), is simply a non-linear wave equation with a cubic source term, it can be numerically solved for arbitrary input fields using similar numerical approaches to those used to solve other equations found in non-linear optics.

Finally, beyond the theoretical aspects, we are currently organising an experimental campaign to verify vacuum photon-photon scattering using OAM as an identifying marker for the scattering signal. There are multiple candidate facilities for the main experiment such as the ELI project's 10 PW beamline, the extreme photonics applications centre (EPAC) at the CLF and 100 PW SEL. These facilities are currently being commissioned and it will take a few years before they are fully operational. In the meantime, the plan is to work on the diagnostics aspect of the experiment with

a view to optimising the single-photon detection and OAM filtration setup so that it can efficiently operate at the expected wavelength and in the presence of various sources of background noise. This will be done over two stages, where first the detection setup will be assembled and tested using low-power sources to optimise the optical alignment of the system, which will then will be tested on currently operational PW-scale facilities such as the ASTRA-GEMINI facility in order to optimise the background noise filtration. Because this project has so many different aspects to it, the plan is to gather a team with a variety of expertise, ranging from ultra-fast quantum optics and strong field low-energy QED, to experimental high-power laser science and statistical methods from high energy physics.

7.2.2 Impact of OAM on laser-plasma interactions

The study and applications of laser-plasma interactions (LPI) have been one of the biggest drivers behind the development of high-power laser facilities around the world. As mentioned in the introduction, understanding laser-plasma interactions is essential for many scientific investigations. These include laser-driven inertial confinement fusion (ICF), where a capsule of fusion fuel is ionised and compressed via the direct, or indirect, action of an intense laser beam until it reaches the necessary temperature and pressure conditions required for nuclear fusion. Necessarily, since the compression process is driven by the laser, the interactions between the laser and the plasma that forms become crucial to a successful implosion. Another important application of LPI can be found in laser wakefield acceleration (LWFA), where an intense laser pulse drives a particular structure in the plasma, creating a strong electric field which can be used to accelerate particles [13]. The electric fields generated via this method can be a 1000 times the strength of those found in conventional accelerators, thus allowing for high acceleration over very short distances.

Laser-driven parametric instabilities are a particularly important type of LPI. These occur when a laser couples to a plasma wave, i.e., an oscillation of the constituents of the plasma, amplifying and generating additional electromagnetic or plasma oscillations. The two most important instabilities are known as stimulated Raman scattering (SRS) and stimulated Brillouin scattering (SBS), where a laser couples with an electron plasma wave and an ion acoustic wave respectively [135, 136]. Controlling these instabilities is essential for the aforementioned ICF and LWFA. The large accelerating electric field in LWFA is caused by a laser driven electron plasma wave. As for ICF, unwanted SRS, caused by the laser driver, can generate so-called hot electrons in the plasma, making the compression of the fuel, and by extension the

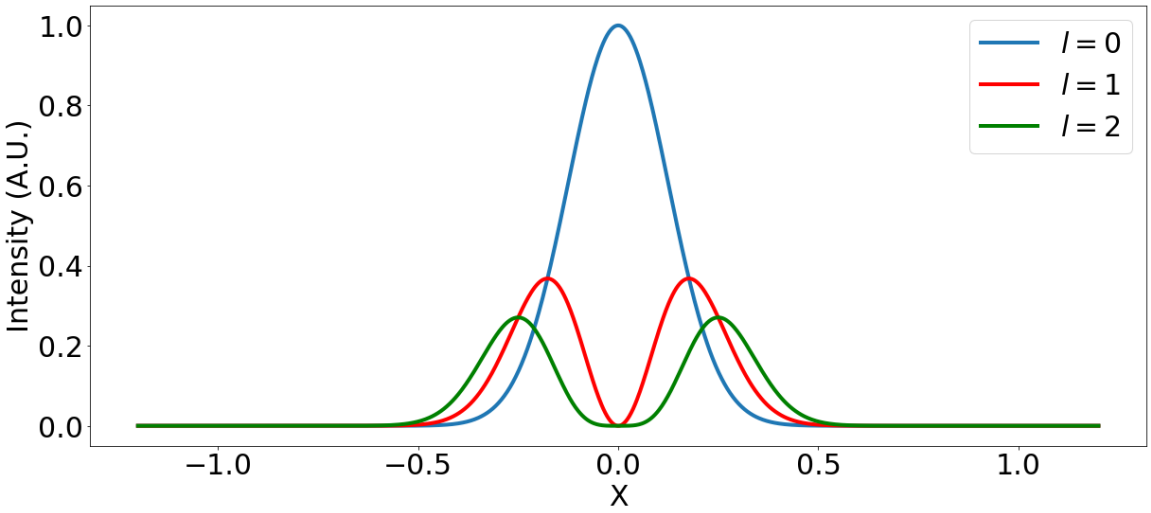

Figure 7.1: Slices of the simulated intensity profile of various normalized LG modes at a particular y and z. The modes only differ in the l mode number while keep all other Gaussian parameters constant. As can be seen from the plot, for the same amount of energy, the peak intensity of an LG mode decreases with the increasing l number.

achievement of fusion conditions, more difficult [137]. Additionally, SBS can cause cross-beam energy transfer (CBET), where the energy from one laser beam can be coupled away from the plasma into another backpropagating beam through a mediating ion acoustic wave, reducing the amount of laser energy that can be used in the compression, thereby reducing the efficiency of the process [138].

A lot of theoretical work has been dedicated to exploring the effect of OAM on certain LPI, producing many interesting predictions. For instance, it has been theorised that OAM can be transferred from lasers into different plasma waves [15], that OAM modifies the landau damping experienced by electron plasma waves [139] and that OAM can be used to control the topology of wakefields and further affects LWFA [16]. However, there are very few experimental results on OAM in LPI. This is because it was only recently that we have been able to generate high-intensity lasers carrying OAM.

The ASTRA experimental campaign, which served to field test the OAM conversion and characterisation methods on high-power facilities, was envisioned as a starting point for multiple planned experiments with the goal of investigating LPI with OAM. One particular focus of these experiments would be to examine whether OAM can mitigate the growth of laser-plasma instabilities. This is important because it can be detrimental to the use of OAM in LWFA, since it depends on laser-driven

plasma oscillations. This mitigation, if verified, will also be beneficial in the context of ICF since it will reduce the harmful impact of plasma instabilities.

There are multiple reasons to suspect why OAM may affect the growth of laser-plasma instabilities. The most straightforward reason is simply the difference in the spatial energy distribution between the fundamental Gaussian mode and higher-order LG modes. As can be seen in Fig.(7.1), for the same amount of energy, the peak intensity decreases as the l mode number increases. In fact, an analytical expression quantifying how much the peak intensity decreases can be easily derived. Considering LG modes with radial modes $p = 0$, the ratio between the peak intensity of the mode to the peak intensity of a Gaussian, both normalised so that $\iint I(x, y)d\Omega = 1$, can be written as

$$\frac{I^l_{\text{peak}}}{I^0_{\text{peak}}} = \frac{l^l}{l!}e^{-l} \xrightarrow[l\to\infty]{} \frac{1}{\sqrt{2\pi l}} \tag{7.2.1}$$

where I^l_{peak} is the peak intensity of an LG^l_0 profile, remembering that $l = 0$ is the standard Gaussian. Since the growth rate of laser-driven plasma instabilities is determined by the intensity of the driving laser, this decrease in intensity could mean that there needs to be more energy in a higher-order LG mode as compared to a Gaussian one in order to drive the same level of plasma waves.

Another more fundamental reason is the different damping rates that are experienced by OAM-carrying plasma waves. Depending on the type of lasers driving the plasma oscillation, the latter can carry different levels of OAM. It has already been shown that OAM can affect the collisionless Landau damping experienced by plasma waves. An open question which remains to be explored is whether OAM also modifies the kinetic collisional damping experienced by plasma waves. Identifying the effect of OAM on the damping rate will be crucial, since damping works against the growth of plasma instabilities and can slow them down considerably [140].

As such, one important focus of the planned future work would be to continue investigating the impact of OAM on LPI. This will be done by looking at its effect on collisional damping by using kinetic theory models, and by running multiple experimental campaigns to evidence the various predictions that were already made about LPI with OAM.

7.3 Afterword

High-intensity lasers have had a tremendous impact on scientific discovery allowing us to both verify decades-old predictions and discover new physics. Their impact will

only increase as technology allows us to reach intensities once-thought unattainable. The orbital angular momentum (OAM) of light is an incredibly powerful tool, which is both interesting in its own right from a fundamental physics point of view as well as in the ways in which it can be used in light-matter interactions. This work aims to combine these two tools and show how they can be used to shed light on some of the most fundamental interactions in the universe. The potential benefits that OAM can have in the context of high-intensity laser science are gathering a rapidly-growing amount of interest from the scientific community, producing discoveries which have the potential for wide-reaching impact on various physical phenomena, ranging from laser-QED effects to laser-plasma interactions.

Bibliography

[1] L. Allen, M. W. Beijersbergen, R. J. C. Spreeuw, and J. P. Woerdman. Orbital angular momentum of light and the transformation of Laguerre-Gaussian laser modes. *Physical Review A*, 45(11):8185–8189, June 1992.

[2] Richard A. Beth. Direct Detection of the Angular Momentum of Light. *Physical Review*, 48(5):471–471, September 1935.

[3] Juan P. Torres. Multiplexing twisted light. *Nature Photonics*, 6(7):420–422, July 2012.

[4] A. Nicolas, L. Veissier, L. Giner, E. Giacobino, D. Maxein, and J. Laurat. A quantum memory for orbital angular momentum photonic qubits. *Nature Photonics*, 8(3):234–238, March 2014.

[5] Christian T. Schmiegelow, Jonas Schulz, Henning Kaufmann, Thomas Ruster, Ulrich G. Poschinger, and Ferdinand Schmidt-Kaler. Transfer of optical orbital angular momentum to a bound electron. *Nature Communications*, 7:12998, October 2016.

[6] K. Y. Bliokh, F. J. Rodrguez-Fortuo, F. Nori, and A. V. Zayats. Spinorbit interactions of light. *Nature Photonics*, 9(12):796–808, December 2015.

[7] Elliot Leader and Cédric Lorcé. The angular momentum controversy: Whats it all about and does it matter? *Physics Reports*, 541(3):163–248, August 2014.

[8] S. J. Van Enk and G. Nienhuis. Commutation Rules and Eigenvalues of Spin and Orbital Angular Momentum of Radiation Fields. *Journal of Modern Optics*, 41(5):963–977, May 1994.

[9] B. Rus, P. Bakule, D. Kramer, G. Korn, J. T. Green, J. Nóvak, M. Fibrich, F. Batysta, J. Thoma, J. Naylon, T. Mazanec, M. Vtek, R. Barros, E. Koutris, J. Hřebíček, J. Polan, R. Baše, P. Homer, M. Košelja, T. Havlíček, A. Honsa,

M. Novák, C. Zervos, P. Korous, M. Laub, and J. Houžvička. ELI-Beamlines laser systems: status and design options. In *High-Power, High-Energy, and High-Intensity Laser Technology; and Research Using Extreme Light: Entering New Frontiers with Petawatt-Class Lasers*, volume 8780, page 87801T. International Society for Optics and Photonics, May 2013.

[10] Baifei Shen, Zhigang Bu, Jiancai Xu, Tongjun Xu, Liangliang Ji, Ruxin Li, and Zhizhan Xu. Exploring vacuum birefringence based on a 100 PW laser and an x-ray free electron laser beam. *Plasma Physics and Controlled Fusion*, 60(4):044002, February 2018.

[11] Francis F. Chen. *Introduction to Plasma Physics and Controlled Fusion*. Springer International Publishing, 3 edition, 2016.

[12] R. S. Craxton, K. S. Anderson, T. R. Boehly, V. N. Goncharov, D. R. Harding, J. P. Knauer, R. L. McCrory, P. W. McKenty, D. D. Meyerhofer, J. F. Myatt, A. J. Schmitt, J. D. Sethian, R. W. Short, S. Skupsky, W. Theobald, W. L. Kruer, K. Tanaka, R. Betti, T. J. B. Collins, J. A. Delettrez, S. X. Hu, J. A. Marozas, A. V. Maximov, D. T. Michel, P. B. Radha, S. P. Regan, T. C. Sangster, W. Seka, A. A. Solodov, J. M. Soures, C. Stoeckl, and J. D. Zuegel. Direct-drive inertial confinement fusion: A review. *Physics of Plasmas*, 22(11):110501, November 2015.

[13] S. M. Hooker. Developments in laser-driven plasma accelerators. *Nature Photonics*, 7(10):775–782, October 2013.

[14] S. V. Bulanov, T. Zh. Esirkepov, M. Kando, J. Koga, K. Kondo, and G. Korn. On the problems of relativistic laboratory astrophysics and fundamental physics with super powerful lasers. *Plasma Physics Reports*, 41(1):1–51, January 2015.

[15] J. T. Mendonça, B. Thidé, and H. Then. Stimulated Raman and Brillouin Backscattering of Collimated Beams Carrying Orbital Angular Momentum. *Physical Review Letters*, 102(18):185005, May 2009.

[16] J. Vieira, J.T. Mendonça, and F. Quéré. Optical Control of the Topology of Laser-Plasma Accelerators. *Physical Review Letters*, 121(5):054801, July 2018.

[17] J. Vieira and J.T. Mendonça. Nonlinear Laser Driven Donut Wakefields for Positron and Electron Acceleration. *Physical Review Letters*, 112(21):215001, May 2014.

[18] J. Ashman, B. Badelek, G. Baum, J. Beaufays, C. P. Bee, C. Benchouk, I. G. Bird, S. C. Brown, M. C. Caputo, H. W. K. Cheung, J. Chima, J. Ciborowski, R. W. Clifft, G. Coignet, F. Combley, G. Court, G. D'Agostini, J. Drees, M. Düren, N. Dyce, A. W. Edwards, M. Edwards, T. Ernst, M. I. Ferrero, D. Francis, E. Gabathuler, J. Gajewski, R. Gamet, V. Gibson, J. Gillies, P. Graftström, K. Hamacher, D. Von Harrach, P. Hayman, J. R. Holt, V. W. Hughes, A. Jacholkowska, T. Jones, E. M. Kabuss, B. Korzen, U. Krüner, S. Kullander, U. Landgraf, D. Lanske, F. Lettenström, T. Lindqvist, J. Loken, M. Matthews, Y. Mizuno, K. Mönig, F. Montanet, J. Nassalski, T. Niinikoski, P. R. Norton, G. Oakham, R. F. Oppenheim, A. M. Osborne, V. Papavassiliou, N. Pavel, C. Peroni, H. Peschel, R. Piegaia, B. Pietrzyk, U. Pietrzyk, B. Povh, P. Renton, J. M. Rieubland, A. Rijllart, K. Rith, E. Rondio, L. Ropelewski, D. Salmon, A. Sandacz, T. Schröder, K. P. Schüler, K. Schultze, T. A. Shibata, T. Sloan, A. Staiano, H. Stier, J. Stock, G. N. Taylor, J. C. Thompson, T. Walcher, S. Wheeler, W. S. C. Williams, S. J. Wimpenny, R. Windmolders, W. J. Womersley, and K. Ziemons. A measurement of the spin asymmetry and determination of the structure function g1 in deep inelastic muon-proton scattering. *Physics Letters B*, 206(2):364–370, May 1988.

[19] V. Yu. Alexakhin, Yu. Alexandrov, G. D. Alexeev, M. Alexeev, A. Amoroso, B. Badeek, F. Balestra, J. Ball, J. Barth, G. Baum, M. Becker, Y. Bedfer, C. Bernet, R. Bertini, M. Bettinelli, R. Birsa, J. Bisplinghoff, P. Bordalo, F. Bradamante, A. Bressan, G. Brona, E. Burtin, M. P. Bussa, V. N. Bytchkov, A. Chapiro, A. Cicuttin, M. Colantoni, A. A. Colavita, S. Costa, M. L. Crespo, N. d'Hose, S. Dalla Torre, S. Das, S. S. Dasgupta, R. De Masi, N. Dedek, D. Demchenko, O. Yu. Denisov, L. Dhara, V. Diaz, A. M. Dinkelbach, S. V. Donskov, V. A. Dorofeev, N. Doshita, V. Duic, W. Dünnweber, A. Efremov, P. D. Eversheim, W. Eyrich, M. Faessler, P. Fauland, A. Ferrero, L. Ferrero, M. Finger, M. Finger, H. Fischer, J. Franz, J. M. Friedrich, V. Frolov, R. Garfagnini, F. Gautheron, O. P. Gavrichtchouk, S. Gerassimov, R. Geyer, M. Giorgi, B. Gobbo, S. Goertz, A. M. Gorin, O. A. Grajek, A. Grasso, B. Grube, A. Guskov, F. Haas, J. Hannappel, D. von Harrach, T. Hasegawa, S. Hedicke, F. H. Heinsius, R. Hermann, C. He, F. Hinterberger, M. von Hodenberg, N. Horikawa, S. Horikawa, I. Horn, C. Ilgner, A. I. Ioukaev, I. Ivanchin, O. Ivanov, T. Iwata, R. Jahn, A. Janata, R. Joosten, N. I. Jouravlev, E. Kabu, D. Kang, B. Ketzer, G. V. Khaustov, Yu. A. Khokhlov, Yu. Kisselev, F. Klein,

K. Klimaszewski, S. Koblitz, J. H. Koivuniemi, V. N. Kolosov, E. V. Komissarov, K. Kondo, K. Königsmann, I. Konorov, V. F. Konstantinov, A. S. Korentchenko, A. Korzenev, A. M. Kotzinian, N. A. Koutchinski, O. Kouznetsov, K. Kowalik, D. Kramer, N. P. Kravchuk, G. V. Krivokhizhin, Z. V. Kroumchtein, J. Kubart, R. Kuhn, V. Kukhtin, F. Kunne, K. Kurek, M. E. Ladygin, M. Lamanna, J. M. Le Goff, M. Leberig, A. A. Lednev, A. Lehmann, J. Lichtenstadt, T. Liska, I. Ludwig, A. Maggiora, M. Maggiora, A. Magnon, G. K. Mallot, C. Marchand, J. Marroncle, A. Martin, J. Marzec, L. Masek, F. Massmann, T. Matsuda, D. Matthiä, A. N. Maximov, W. Meyer, A. Mielech, Yu. V. Mikhailov, M. A. Moinester, T. Nagel, O. Nähle, J. Nassalski, S. Neliba, D. P. Neyret, V. I. Nikolaenko, K. Nikolaev, A. A. Nozdrin, V. F. Obraztsov, A. G. Olshevsky, M. Ostrick, A. Padee, P. Pagano, S. Panebianco, D. Panzieri, S. Paul, D. V. Peshekhonov, V. D. Peshekhonov, G. Piragino, S. Platchkov, J. Pochodzalla, J. Polak, V. A. Polyakov, G. Pontecorvo, A. A. Popov, J. Pretz, S. Procureur, C. Quintans, S. Ramos, G. Reicherz, E. Rondio, A. M. Rozhdestvensky, D. Ryabchikov, V. D. Samoylenko, A. Sandacz, H. Santos, M. G. Sapozhnikov, I. A. Savin, P. Schiavon, C. Schill, L. Schmitt, W. Schroeder, D. Seeharsch, M. Seimetz, D. Setter, O. Yu. Shevchenko, H. W. Siebert, L. Silva, L. Sinha, A. N. Sissakian, M. Slunecka, G. I. Smirnov, F. Sozzi, A. Srnka, F. Stinzing, M. Stolarski, V. P. Sugonyaev, M. Sulc, R. Sulej, V. V. Tchalishev, S. Tessaro, F. Tessarotto, A. Teufel, L. G. Tkatchev, S. Trippel, G. Venugopal, M. Virius, N. V. Vlassov, R. Webb, E. Weise, Q. Weitzel, R. Windmolders, W. Wilicki, K. Zaremba, M. Zavertyaev, E. Zemlyanichkina, J. Zhao, R. Ziegler, and A. Zvyagin. The deuteron spin-dependent structure function g1d and its first moment. *Physics Letters B*, 647(1):8–17, March 2007.

[20] HERMES Collaboration, A. Airapetian, N. Akopov, Z. Akopov, A. Andrus, E. C. Aschenauer, W. Augustyniak, R. Avakian, A. Avetissian, E. Avetissian, S. Belostotski, N. Bianchi, H. P. Blok, H. Böttcher, A. Borissov, A. Borysenko, A. Brüll, V. Bryzgalov, M. Capiluppi, G. P. Capitani, G. Ciullo, M. Contalbrigo, P. F. Dalpiaz, W. Deconinck, R. De Leo, M. Demey, L. De Nardo, E. De Sanctis, E. Devitsin, M. Diefenthaler, P. Di Nezza, J. Dreschler, M. Düren, M. Ehrenfried, A. Elalaoui-Moulay, G. Elbakian, F. Ellinghaus, U. Elschenbroich, R. Fabbri, A. Fantoni, L. Felawka, S. Frullani, A. Funel, D. Gabbert, Y. Gärber, G. Gapienko, V. Gapienko, F. Garibaldi, K. Garrow, G. Gavrilov, V. Gharibyan, F. Giordano, O. Grebeniouk, I. M. Gregor, H. Guler, A. Gute,

C. Hadjidakis, M. Hartig, D. Hasch, T. Hasegawa, W. H. A. Hesselink, A. Hillenbrand, M. Hoek, Y. Holler, B. Hommez, I. Hristova, G. Iarygin, A. Ivanilov, A. Izotov, H. E. Jackson, A. Jgoun, R. Kaiser, T. Keri, E. Kinney, A. Kisselev, T. Kobayashi, M. Kopytin, V. Korotkov, V. Kozlov, B. Krauss, P. Kravchenko, V. G. Krivokhijine, L. Lagamba, L. Lapiks, P. Lenisa, P. Liebing, L. A. Linden-Levy, W. Lorenzon, J. Lu, S. Lu, B.-Q. Ma, B. Maiheu, N. C. R. Makins, Y. Mao, B. Marianski, H. Marukyan, F. Masoli, V. Mexner, N. Meyners, T. Michler, O. Mikloukho, C. A. Miller, Y. Miyachi, V. Muccifora, M. Murray, A. Nagaitsev, E. Nappi, Y. Naryshkin, M. Negodaev, W.-D. Nowak, H. Ohsuga, A. Osborne, R. Perez-Benito, N. Pickert, M. Raithel, D. Reggiani, P. E. Reimer, A. Reischl, A. R. Reolon, C. Riedl, K. Rith, G. Rosner, A. Rostomyan, L. Rubacek, J. Rubin, D. Ryckbosch, Y. Salomatin, I. Sanjiev, I. Savin, A. Schäfer, G. Schnell, K. P. Schüler, J. Seele, B. Seitz, C. Shearer, T.-A. Shibata, V. Shutov, K. Sinram, M. Stancari, M. Statera, E. Steffens, J. J. M. Steijger, H. Stenzel, J. Stewart, F. Stinzing, U. Stösslein, J. Streit, P. Tait, H. Tanaka, S. Taroian, B. Tchuiko, A. Terkulov, A. Trzcinski, M. Tytgat, A. Vandenbroucke, P. B. van der Nat, G. van der Steenhoven, Y. van Haarlem, D. Veretennikov, V. Vikhrov, C. Vogel, S. Wang, C. Weiskopf, Y. Ye, Z. Ye, S. Yen, B. Zihlmann, and P. Zupranski. Precise determination of the spin structure function ${g}_{1}$ of the proton, deuteron, and neutron. *Physical Review D*, 75(1):012007, January 2007. Publisher: American Physical Society.

[21] Xiang-Song Chen, Xiao-Fu Lü, Wei-Min Sun, Fan Wang, and T. Goldman. Spin and Orbital Angular Momentum in Gauge Theories: Nucleon Spin Structure and Multipole Radiation Revisited. *Physical Review Letters*, 100(23):232002, June 2008. Publisher: American Physical Society.

[22] Herbert Goldstein, Charles P. Poole, and John L. Safko. *Classical Mechanics*. Pearson, August 2013. Google-Books-ID: bhVnngEACAAJ.

[23] E. Noether. Invariante Variationsprobleme. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1918:235–257, 1918.

[24] P. A. M. Dirac. *The Principles Of Quantum Mechanics*. Oxford University Press, U.S.A., Oxford, 4 edition edition, January 1988.

[25] Michele Maggiore. *A Modern Introduction to Quantum Field Theory*. Oxford University Press, 2005. Google-Books-ID: yykTDAAAQBAJ.

[26] R. Courant & D. Hilbert. *Methods of Mathematical Physics volume 1.*

[27] Walther Gerlach and Otto Stern. Der experimentelle Nachweis der Richtungsquantelung im Magnetfeld. *Zeitschrift für Physik*, 9(1):349–352, December 1922.

[28] Michael E. Peskin and Daniel V. Schroeder. *An Introduction To Quantum Field Theory.* Avalon Publishing, October 1995. Google-Books-ID: EVeNNcslvX0C.

[29] John David Jackson. *Classical Electrodynamics, 3rd Edition.* July 1998.

[30] Claude Cohen-Tannoudji, Jacques Dupont-Roc, and Gilbert Grynberg. *Photons and Atoms: Introduction to Quantum Electrodynamics.* Wiley, March 1997. Google-Books-ID: Ow3FiqmKrocC.

[31] Bahaa E. A. Saleh and Malvin Carl Teich. *Fundamentals of Photonics.* Wiley, February 2013. Google-Books-ID: Qfeosgu08u8C.

[32] H. He, M. E. J. Friese, N. R. Heckenberg, and H. Rubinsztein-Dunlop. Direct Observation of Transfer of Angular Momentum to Absorptive Particles from a Laser Beam with a Phase Singularity. *Physical Review Letters*, 75(5):826–829, July 1995.

[33] M. F. Andersen, C. Ryu, Pierre Cladé, Vasant Natarajan, A. Vaziri, K. Helmerson, and W. D. Phillips. Quantized rotation of atoms from photons with orbital angular momentum. *Physical Review Letters*, 97(17):170406, October 2006.

[34] Howard Georgi. *Lie Algebras In Particle Physics: from Isospin To Unified Theories.* Hachette UK, October 1999. Google-Books-ID: dhY5DgAAQBAJ.

[35] Sophus Lie and Friedrich Engel. *Theorie der Transformationsgruppen.* American Mathematical Soc., 1970. Google-Books-ID: H07V8WKLgnAC.

[36] J. E. Campbell. On a Law of Combination of Operators (Second Paper). *Proceedings of the London Mathematical Society*, s1-29(1):14–32, 1897.

[37] Brian C. Hall. *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction.* Graduate Texts in Mathematics. Springer International Publishing, 2 edition, 2015.

[38] Steven Weinberg and Weinberg S. *The Quantum Theory of Fields.* Cambridge University Press, June 1995. Google-Books-ID: doeDB3_WLvwC.

[39] Elliot Leader. The photon angular momentum controversy: Resolution of a conflict between laser optics and particle physics. *Physics Letters B*, 756:303–308, May 2016.

[40] Elliot Leader. A proposed measurement of optical orbital and spin angular momentum and its implications for photon angular momentum. *Physics Letters B*, 779:385–387, April 2018.

[41] A. T. Rosenberger, Elijah B. Dale, Khoa V. Bui, Erik K. Gonzales, D. Ganta, Limu Ke, and Sreekul Raj Rajagopal. Cross-polarization coupling of whispering-gallery modes due to the spin-orbit interaction of light. *Optics Letters*, 44(17):4163–4166, September 2019.

[42] Moïse Sotto, Kapil Debnath, Isao Tomita, and Shinichi Saito. Spin-Orbit Coupling of Light in Photonic Crystal Waveguides. *Physical Review A*, 99(5):053845, May 2019.

[43] Willis E. Lamb and Robert C. Retherford. Fine Structure of the Hydrogen Atom by a Microwave Method. *Physical Review*, 72(3):241–243, August 1947. Publisher: American Physical Society.

[44] H. A. Bethe. The Electromagnetic Shift of Energy Levels. *Physical Review*, 72(4):339–341, August 1947. Publisher: American Physical Society.

[45] H.B.G. Casimir. On the Attraction Between Two Perfectly Conducting Plates. *Indag. Math.*, 10:261–263, 1948.

[46] S. K. Lamoreaux. Demonstration of the Casimir Force in the 0.6 to {\SI{6}{\mu\metre}} Range. *Physical Review Letters*, 78(1):5–8, January 1997. Publisher: American Physical Society.

[47] Gert Brodin, Mattias Marklund, and Lennart Stenflo. Proposal for Detection of QED Vacuum Nonlinearities in Maxwell's Equations by the Use of Waveguides. *Physical Review Letters*, 87(17):171801, October 2001.

[48] A. E. Kaplan and Y. J. Ding. Field-gradient-induced second-harmonic generation in magnetized vacuum. *Physical Review A*, 62(4):043805, September 2000.

[49] Baifei Shen, M. Y. Yu, and Xiaofang Wang. Photonphoton scattering in a plasma channel. *Physics of Plasmas*, 10(11):4570–4571, October 2003.

[50] E. Lundström, G. Brodin, J. Lundin, M. Marklund, R. Bingham, J. Collier, J. T. Mendonça, and P. Norreys. Using High-Power Lasers for Detection of Elastic Photon-Photon Scattering. *Physical Review Letters*, 96(8):083602, March 2006.

[51] B. King, H. Hu, and B. Shen. Three-pulse photon-photon scattering. *Physical Review A*, 98(2):023817, August 2018.

[52] Max Planck. On the Law of Distribution of Energy in the Normal Spectrum. *Annalen Phys.*, 4:553, 1901.

[53] Albert Einstein. Concerning an heuristic point of view toward the emission and transformation of light. *Annalen Phys.*, 17:132–148, 1905.

[54] E. Schrödinger. An undulatory theory of the mechanics of atoms and molecules. *Phys. Rev.*, 28:1049–1070, Dec 1926.

[55] Dirac Paul Adrien Maurice. Relativistic wave equations. *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences*, 155(886):447–459, July 1936.

[56] G. Gabrielse, D. Hanneke, T. Kinoshita, M. Nio, and B. Odom. New determination of the fine structure constant from the electron g value and qed. *Phys. Rev. Lett.*, 97:030802, Jul 2006.

[57] Suraj N. Gupta. Theory of Longitudinal Photons in Quantum Electrodynamics. *Proceedings of the Physical Society. Section A*, 63(7):681–691, July 1950.

[58] W. Heisenberg and H. Euler. Folgerungen aus der Diracschen Theorie des Positrons. *Zeitschrift für Physik*, 98(11):714–732, November 1936.

[59] E. A. Uehling. Polarization effects in the positron theory. *Phys. Rev.*, 48:55–63, Jul 1935.

[60] Robert Serber. Linear modifications in the maxwell field equations. *Phys. Rev.*, 48:49–54, Jul 1935.

[61] M. Schumacher, I. Borchert, F. Smend, and P. Rullhusen. Delbr\"uck scattering of \SI{2.75}{\megaelectronvolt} photons by lead. *Physics Letters B*, 59(2):134–136, October 1975.

[62] Sh. Zh. Akhmadaliev, G. Ya. Kezerashvili, S. G. Klimenko, R. N. Lee, V. M. Malyshev, A. L. Maslennikov, A. M. Milov, A. I. Milstein, N. Yu. Muchnoi, A. I. Naumenkov, V. S. Panin, S. V. Peleganchuk, G. E. Pospelov, I. Ya. Protopopov, L. V. Romanov, A. G. Shamov, D. N. Shatilov, E. A. Simonov, V. M. Strakhovenko, and Yu. A. Tikhonov. Experimental Investigation of High-Energy Photon Splitting in Atomic Fields. *Physical Review Letters*, 89(6):061802, July 2002. Publisher: American Physical Society.

[63] B. King and T. Heinzl. Measuring vacuum polarization with high-power lasers. *High Power Laser Science and Engineering*, 4:e5, 2016.

[64] Federico Della Valle, Aldo Ejlli, Ugo Gastaldi, Giuseppe Messineo, Edoardo Milotti, Ruggero Pengo, Giuseppe Ruoso, and Guido Zavattini. The PVLAS experiment: measuring vacuum magnetic birefringence and dichroism with a birefringent FabryPerot cavity. *The European Physical Journal C*, 76(1):24, January 2016.

[65] R. Battesti, B. Pinto Da Souza, S. Batut, C. Robilliard, G. Bailly, C. Michel, M. Nardone, L. Pinard, O. Portugall, G. Trénec, J.-M. Mackowski, G. L.J.A. Rikken, J. Vigué, and C. Rizzo. The BMV experiment: a novel apparatus to study the propagation of light in a transverse magnetic field. *The European Physical Journal D*, 46(2):323–333, February 2008.

[66] M. T. Hartman, R. Battesti, and C. Rizzo. Status of the Vacuum Birefringence Search at BMV. In *2018 Conference on Precision Electromagnetic Measurements (CPEM 2018)*, pages 1–2, July 2018. ISSN: 2160-0171.

[67] Hans-Peter Schlenvoigt, Tom Heinzl, Ulrich Schramm, Thomas E. Cowan, and Roland Sauerbrey. Detecting vacuum birefringence with x-ray free electron lasers and high-power optical lasers: a feasibility study. *Physica Scripta*, 91(2):023010, January 2016. Publisher: IOP Publishing.

[68] J. T. Mendonça, M. Marklund, P. K. Shukla, and G. Brodin. Photon acceleration in vacuum. *Physics Letters A*, 359(6):700–704, December 2006.

[69] A. Di Piazza, K. Z. Hatsagortsyan, and C. H. Keitel. Harmonic generation from laser-driven vacuum. *Physical Review D*, 72(8):085005, October 2005. Publisher: American Physical Society.

[70] Paola Arias, Jorge Gamboa, Horacio Falomir, and Fernando Méndez. Effects of heavy bosonic excitations on qed vacuum. *Modern Physics Letters A*, 24(16):1289–1299, May 2009. Publisher: World Scientific Publishing Co.

[71] Holger Gies, Joerg Jaeckel, and Andreas Ringwald. Polarized Light Propagating in a Magnetic Field as a Probe for Millicharged Fermions. *Physical Review Letters*, 97(14):140402, October 2006. Publisher: American Physical Society.

[72] Max Born, L. Infeld, and Ralph Howard Fowler. Foundations of the new field theory. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 144(852):425–451, March 1934. Publisher: Royal Society.

[73] José Manuel Dvila, Christian Schubert, and Mara Anabel Trejo. Photonic processes in BornInfeld theory. *International Journal of Modern Physics A*, 29(30):1450174, December 2014. Publisher: World Scientific Publishing Co.

[74] M. Aaboud, G. Aad, B. Abbott, J. Abdallah, O. Abdinov, B. Abeloos, S. H. Abidi, O. S. AbouZeid, N. L. Abraham, H. Abramowicz, H. Abreu, R. Abreu, Y. Abulaiti, B. S. Acharya, S. Adachi, L. Adamczyk, J. Adelman, M. Adersberger, T. Adye, A. A. Affolder, T. Agatonovic-Jovin, C. Agheorghiesei, J. A. Aguilar-Saavedra, S. P. Ahlen, F. Ahmadov, G. Aielli, S. Akatsuka, H. Akerstedt, T. P. A. kesson, A. V. Akimov, G. L. Alberghi, J. Albert, M. J. Alconada Verzini, M. Aleksa, I. N. Aleksandrov, C. Alexa, G. Alexander, T. Alexopoulos, M. Alhroob, B. Ali, M. Aliev, G. Alimonti, J. Alison, S. P. Alkire, B. M. M. Allbrooke, B. W. Allen, P. P. Allport, A. Aloisio, A. Alonso, F. Alonso, C. Alpigiani, A. A. Alshehri, M. Alstaty, B. Alvarez Gonzalez, D. lvarez Piqueras, M. G. Alviggi, B. T. Amadio, Y. Amaral Coutinho, C. Amelung, D. Amidei, S. P. Amor Dos Santos, A. Amorim, S. Amoroso, G. Amundsen, C. Anastopoulos, L. S. Ancu, N. Andari, T. Andeen, C. F. Anders, J. K. Anders, K. J. Anderson, A. Andreazza, V. Andrei, S. Angelidakis, I. Angelozzi, A. Angerami, F. Anghinolfi, A. V. Anisenkov, N. Anjos, A. Annovi, C. Antel, M. Antonelli, A. Antonov, D. J. Antrim, F. Anulli, M. Aoki, L. Aperio Bella, G. Arabidze, Y. Arai, J. P. Araque, V. Araujo Ferraz, A. T. H. Arce, R. E. Ardell, F. A. Arduh, J.-F. Arguin, S. Argyropoulos, M. Arik, A. J. Armbruster, L. J. Armitage, O. Arnaez, H. Arnold, M. Arratia, O. Arslan, A. Artamonov, G. Artoni, S. Artz, S. Asai, N. Asbah, A. Ashkenazi,

L. Asquith, K. Assamagan, R. Astalos, M. Atkinson, N. B. Atlay, K. Augsten, G. Avolio, B. Axen, M. K. Ayoub, G. Azuelos, A. E. Baas, M. J. Baca, H. Bachacou, K. Bachas, M. Backes, M. Backhaus, P. Bagiacchi, P. Bagnaia, J. T. Baines, M. Bajic, O. K. Baker, E. M. Baldin, P. Balek, T. Balestri, F. Balli, W. K. Balunas, E. Banas, Sw. Banerjee, A. A. E. Bannoura, L. Barak, E. L. Barberio, D. Barberis, M. Barbero, T. Barillari, M.-S. Barisits, T. Barklow, N. Barlow, S. L. Barnes, B. M. Barnett, R. M. Barnett, Z. Barnovska-Blenessy, A. Baroncelli, G. Barone, A. J. Barr, L. Barranco Navarro, F. Barreiro, J. Barreiro Guimares da Costa, R. Bartoldus, A. E. Barton, P. Bartos, A. Basalaev, A. Bassalat, R. L. Bates, S. J. Batista, J. R. Batley, M. Battaglia, M. Bauce, F. Bauer, H. S. Bawa, J. B. Beacham, M. D. Beattie, T. Beau, P. H. Beauchemin, P. Bechtle, H. P. Beck, K. Becker, M. Becker, M. Beckingham, C. Becot, A. J. Beddall, A. Beddall, V. A. Bednyakov, M. Bedognetti, C. P. Bee, T. A. Beermann, M. Begalli, M. Begel, J. K. Behr, A. S. Bell, G. Bella, L. Bellagamba, A. Bellerive, M. Bellomo, K. Belotskiy, O. Beltramello, N. L. Belyaev, O. Benary, D. Benchekroun, M. Bender, K. Bendtz, N. Benekos, Y. Benhammou, E. Benhar Noccioli, J. Benitez, D. P. Benjamin, M. Benoit, J. R. Bensinger, S. Bentvelsen, L. Beresford, M. Beretta, D. Berge, E. Bergeaas Kuutmann, N. Berger, J. Beringer, S. Berlendis, N. R. Bernard, G. Bernardi, C. Bernius, F. U. Bernlochner, T. Berry, P. Berta, C. Bertella, G. Bertoli, F. Bertolucci, I. A. Bertram, C. Bertsche, D. Bertsche, G. J. Besjes, O. Bessidskaia Bylund, M. Bessner, N. Besson, C. Betancourt, A. Bethani, S. Bethke, A. J. Bevan, R. M. Bianchi, M. Bianco, O. Biebel, D. Biedermann, R. Bielski, N. V. Biesuz, M. Biglietti, J. Bilbao De Mendizabal, T. R. V. Billoud, H. Bilokon, M. Bindi, A. Bingul, C. Bini, S. Biondi, T. Bisanz, C. Bittrich, D. M. Bjergaard, C. W. Black, J. E. Black, K. M. Black, D. Blackburn, R. E. Blair, T. Blazek, I. Bloch, C. Blocker, A. Blue, W. Blum, U. Blumenschein, S. Blunier, G. J. Bobbink, V. S. Bobrovnikov, S. S. Bocchetta, A. Bocci, C. Bock, M. Boehler, D. Boerner, D. Bogavac, A. G. Bogdanchikov, C. Bohm, V. Boisvert, P. Bokan, T. Bold, A. S. Boldyrev, M. Bomben, M. Bona, M. Boonekamp, A. Borisov, G. Borissov, J. Bortfeldt, D. Bortoletto, V. Bortolotto, K. Bos, D. Boscherini, M. Bosman, J. D. Bossio Sola, J. Boudreau, J. Bouffard, E. V. Bouhova-Thacker, D. Boumediene, C. Bourdarios, S. K. Boutle, A. Boveia, J. Boyd, I. R. Boyko, J. Bracinik, and ATLAS Collaboration. Evidence for light-by-light scattering in heavy-ion collisions with the ATLAS detector at the LHC. *Nature Physics*, 13(9):852–858, September 2017. Number: 9 Publisher: Nature Publishing Group.

[75] D. N. Papadopoulos, J. P. Zou, C. Le Blanc, G. Chériaux, P. Georges, F. Druon, G. Mennerat, P. Ramirez, L. Martin, A. Fréneaux, A. Beluze, N. Lebas, P. Monot, F. Mathieu, and P. Audebert. The Apollon 10 PW laser: experimental and theoretical investigation of the temporal characteristics. *High Power Laser Science and Engineering*, 4, 2016.

[76] B. King and C. H. Keitel. Photonphoton scattering in collisions of intense laser pulses. *New Journal of Physics*, 14(10):103002, October 2012. Publisher: IOP Publishing.

[77] J. Lundin, M. Marklund, E. Lundström, G. Brodin, J. Collier, R. Bingham, J. T. Mendonça, and P. Norreys. Analysis of four-wave mixing of high-power lasers for the detection of elastic photon-photon scattering. *Phys. Rev. A*, 74:043821, Oct 2006.

[78] Milton Abramowitz. *Handbook of Mathematical Functions, With Formulas, Graphs, and Mathematical Tables,*. Dover Publications, Inc., New York, NY, USA, 1974.

[79] D McGloin and K Dholakia. Bessel beams: Diffraction in a new light. *Contemporary Physics*, 46(1):15–28, 2005.

[80] Jonathan Leach, Miles J. Padgett, Stephen M. Barnett, Sonja Franke-Arnold, and Johannes Courtial. Measuring the Orbital Angular Momentum of a Single Photon. *Physical Review Letters*, 88(25):257901, June 2002.

[81] G. Cheriaux J. Collier R. Dabu P. Dombi A. M. Dunne K. Ertel P. Georges J. Hebling J. Hein C. Hernandez-Gomez C. Hooker S. Karsch G. Korn F. Krausz C. Le Blanc Zs. Major F. Mathieu T. Metzger G. Mourou P. Nickles K. Osvay B. Rus W. Sandner G. Szab D. Ursescu K. Varj J.-P. Chambaret, O. Chekhlov. Extreme light infrastructure: laser architecture and major challenges, 2010.

[82] Gérard Mourou and Toshiki Tajima. The Extreme Light Infrastructure: Optics Next Horizon. *Optics and Photonics News*, 22(7):47–51, July 2011.

[83] M. Caño-García, X. Quintana, J. M. Otón, and M. A. Geday. Dynamic multilevel spiral phase plate generator. *Scientific Reports*, 8(1):15804, October 2018.

[84] Djenan Ganic, Xiaosong Gan, Min Gu, Mathias Hain, Somakanthan Somalingam, Svetomir Stankovic, and Theo Tschudi. Generation of doughnut laser

beams by use of a liquid-crystal cell with a conversion efficiency near 100%. *Optics Letters*, 27(15):1351–1353, August 2002.

[85] U. D. Jentschura and V. G. Serbo. Generation of High-Energy Photons with Large Orbital Angular Momentum by Compton Backscattering. *Physical Review Letters*, 106(1):013001, January 2011.

[86] Adrien Nicolas, Lucile Veissier, Elisabeth Giacobino, Dominik Maxein, and Julien Laurat. Quantum state tomography of orbital angular momentum photonic qubits via a projection-based technique. *New Journal of Physics*, 17(3):033037, March 2015.

[87] Donna Strickland and Gérard Mourou. Compression of amplified chirped optical pulses. *Optics Communications*, 56(3):219–221, December 1985.

[88] T. H. Maiman. Stimulated Optical Radiation in Ruby. *Nature*, 187(4736):493–494, August 1960.

[89] A. L. Schawlow and C. H. Townes. Infrared and Optical Masers. *Physical Review*, 112(6):1940–1949, December 1958.

[90] R Gordon Gould et al. The laser, light amplification by stimulated emission of radiation. In *The Ann Arbor conference on optical pumping, the University of Michigan*, volume 15, page 92, 1959.

[91] Albert Einstein. Zur Quantentheorie der Strahlung. *Phys.Z.*, 18:121–128, 1917.

[92] Stepan S. Bulanov, Timur Zh. Esirkepov, Alexander G. R. Thomas, James K. Koga, and Sergei V. Bulanov. Schwinger Limit Attainability with Extreme Power Lasers. *Physical Review Letters*, 105(22):220407, November 2010.

[93] L. E. Hargrove, R. L. Fork, and M. A. Pollack. Locking of hene laser modes induced by synchronous intracavity modulation. *Applied Physics Letters*, 5(1):4–5, 1964.

[94] O. Martinez. 3000 times grating compressor with positive group velocity dispersion: Application to fiber compensation in 1.3-1.6 m region. *IEEE Journal of Quantum Electronics*, 23(1):59–64, January 1987.

[95] V. L. Bochenko, I. I. Zasavitski, Yu V. Kosichkin, Aleksandr P. Tarasevich, V. G. Tunkin, and A. P. Shotov. Picosecond parametric oscillator amplifying radiation from a tunable semiconductor laser. *Soviet Journal of Quantum Electronics*, 14(1):141, 1984.

[96] C. J. Hooker, S. Blake, O. Chekhlov, R. J. Clarke, J. L. Collier, E. J. Divall, K. Ertel, P. S. Foster, S. J. Hawkes, P. Holligan, B. Landowski, W. J. Lester, D. Neely, B. Parry, R. Pattathil, M. Streeter, and B. E. Wyborn. Commissioning the Astra Gemini Petawatt Ti:Sapphire Laser System. In *Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies (2008), paper JThB2*, page JThB2. Optical Society of America, May 2008.

[97] Albertas ukauskas, Mangirdas Malinauskas, and Etienne Brasselet. Monolithic generators of pseudo-nondiffracting optical vortex beams at the microscale. *Applied Physics Letters*, 103(18):181122, October 2013.

[98] Anna T. O'Neil and Johannes Courtial. Mode transformations in terms of the constituent HermiteGaussian or LaguerreGaussian modes and the variable-phase mode converter. *Optics Communications*, 181(1):35–45, July 2000.

[99] M. W. Beijersbergen, L. Allen, H. E. L. O. van der Veen, and J. P. Woerdman. Astigmatic laser mode converters and transfer of orbital angular momentum. *Optics Communications*, 96(1):123–132, February 1993.

[100] Dennis Gabor and William Lawrence Bragg. Microscopy by reconstructed wavefronts. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 197(1051):454–487, July 1949.

[101] W. T. Cathey. Phase Holograms, Phase-Only Holograms, and Kinoforms. *Applied Optics*, 9(6):1478–1479, June 1970.

[102] H. He, N. R. Heckenberg, and H. Rubinsztein-Dunlop. Optical Particle Trapping with Higher-order Doughnut Beams Produced Using High Efficiency Computer Generated Holograms. *Journal of Modern Optics*, 42(1):217–223, January 1995.

[103] J. Arlt, K. Dholakia, L. Allen, and M. J. Padgett. The production of multiringed LaguerreGaussian modes by computer-generated holograms. *Journal of Modern Optics*, 45(6):1231–1237, June 1998.

[104] M. W. Beijersbergen, R. P. C. Coerwinkel, M. Kristensen, and J. P. Woerdman. Helical-wavefront laser beams produced with a spiral phaseplate. *Optics Communications*, 112(5):321–327, December 1994.

[105] Andrew Longman and Robert Fedosejevs. Mode conversion efficiency to Laguerre-Gaussian OAM modes using spiral phase optics. *Optics Express*, 25(15):17382–17392, July 2017.

[106] G. Ruffato, M. Massari, and F. Romanato. Generation of high-order Laguerre-Gaussian modes by means of spiral phase plates. *Optics Letters*, 39(17):5094–5097, September 2014.

[107] K. Sueda, G. Miyaji, N. Miyanaga, and M. Nakatsuka. Laguerre-Gaussian beam generated with a multilevel spiral phase plate for high intensity laser pulses. *Optics Express*, 12(15):3548–3553, July 2004.

[108] Guodong Xie, Yongxiong Ren, Hao Huang, Martin P. J. Lavery, Nisar Ahmed, Yan Yan, Changjing Bao, Long Li, Zhe Zhao, Yinwen Cao, Moshe Willner, Moshe Tur, Samuel J. Dolinar, Robert W. Boyd, Jeffrey H. Shapiro, and Alan E. Willner. Phase correction for a distorted orbital angular momentum beam using a Zernike polynomials-based stochastic-parallel-gradient-descent algorithm. *Optics Letters*, 40(7):1197–1200, April 2015.

[109] Andrew Longman, Carlos Salgado, Ghassan Zeraouli, Jon I. Apinaniz, Jose Antonio Perez-Hernandez, M. Khairy Eltahlawy, Luca Volpe, and Robert Fedosejevs. Off-Axis Spiral Phase Mirrors for Generating High Intensity Optical Vortices. *arXiv:2001.07813 [physics]*, January 2020. arXiv: 2001.07813.

[110] Joseph W. Goodman. *Introduction to Fourier Optics*. Roberts and Company Publishers, 2005. Google-Books-ID: ow5xs_Rtt9AC.

[111] Samuel N. Alperin, Robert D. Niederriter, Juliet T. Gopinath, and Mark E. Siemens. Quantitative measurement of the orbital angular momentum of light with a single, stationary lens. *Optics Letters*, 41(21):5019–5022, November 2016.

[112] M. Harris, C. A. Hill, P. R. Tapster, and J. M. Vaughan. Laser modes with helical wave fronts. *Physical Review A*, 49(4):3119–3122, April 1994.

[113] Jonathan Leach, Miles J. Padgett, Stephen M. Barnett, Sonja Franke-Arnold, and Johannes Courtial. Measuring the orbital angular momentum of a single photon. *Phys. Rev. Lett.*, 88:257901, Jun 2002.

[114] Ajoy Ghatak and K. Thyagarajan. *An Introduction to Fiber Optics*. Cambridge University Press, 1998.

[115] K. A. H. van Leeuwen and H. T. Nijnuis. Measurement of higher-order mode attenuation in single-mode fibers: effective cutoff wavelength. *Optics Letters*, 9(6):252–254, June 1984.

[116] Peter M. Celliers, Kent G. Estabrook, Russell J. Wallace, James E. Murray, Luiz B. Da Silva, Brian J. MacGowan, Bruno M. Van Wonterghem, and Kenneth R. Manes. Spatial filter pinhole for high-energy pulsed lasers. *Applied Optics*, 37(12):2371–2378, April 1998.

[117] C. Iaconis and I. A. Walmsley. Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses. *Opt. Lett.*, 23(10):792–794, May 1998.

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

[119] Rick Trebino and Daniel J. Kane. Using phase retrieval to measure the intensity and phase of ultrashort pulses: frequency-resolved optical gating. *JOSA A*, 10(5):1101–1111, May 1993.

[120] Dongjoo Lee, Selcuk Akturk, Pablo Gabolde, and Rick Trebino. Experimentally simple, extremely broadband transient-grating frequency-resolved-optical-gating arrangement. *Opt. Express*, 15(2):760–766, Jan 2007.

[121] C. N. Danson, P. A. Brummitt, R. J. Clarke, J. L. Collier, B. Fell, A. J. Frackiewicz, S. Hancock, S. Hawkes, C. Hernandez-Gomez, P. Holligan, M. H. R. Hutchinson, A. Kidd, W. J. Lester, I. O. Musgrave, D. Neely, D.R. Neville, P. A. Norreys, D. A. Pepler, C. J. Reason, W. Shaikh, T. B. Winstone, R. W. W. Wyatt, and B. E. Wyborn. Vulcan Petawattan ultra-high-intensity interaction facility. *Nuclear Fusion*, 44(12):S239, 2004.

[122] Rick Trebino. Geometrical Issues: Single-shot FROG. In *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses*, pages 141–156. Springer, Boston, MA, 2000.

[123] S. Palaniyappan, R. C. Shah, R. Johnson, T. Shimada, D. C. Gautier, S. Letzring, D. Jung, R. Hörlein, D. T. Offermann, J. C. Fernández, and B. M.

Hegelich. Pulse shape measurements using single shot-frequency resolved optical gating for high energy (80 j) short pulse (600 fs) laser. *Review of Scientific Instruments*, 81(10):10E103, 2010.

[124] Sasi Palaniyappan, B. Manuel Hegelich, Hui-Chun Wu, Daniel Jung, Donald C. Gautier, Lin Yin, Brian J. Albright, Randall P. Johnson, Tsutomu Shimada, Samuel Letzring, Dustin T. Offermann, Jun Ren, Chengkun Huang, Rainer Hörlein, Brendan Dromey, Juan C. Fernandez, and Rahul C. Shah. Dynamics of relativistic transparency and optical shuttering in expanding overdense plasmas. *Nature Physics*, 8:763 EP –, Aug 2012. Article.

[125] W. Theobald, R. Häßner, R. Kingham, T. Feurer, H. Schillinger, G. Schäfer, and R. Sauerbrey. Probing black hole physics in the laboratory using high intensity femtosecond lasers. In *Ultrafast Phenomena XI*, pages 410–414, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

[126] Marco A. Krumbuegel and Rick Trebino. The FROG Algorithm. In *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses*, pages 157–178. Springer US, Boston, MA, 2000.

[127] Alessandro Colombo, Davide Emilio Galli, Liberato De Caro, Francesco Scattarella, and Elvio Carlino. Facing the phase problem in Coherent Diffractive Imaging via Memetic Algorithms. *Scientific Reports*, 7:42236, February 2017.

[128] T. El-Mihoub, Adrian A. Hopgood, L. Nolle, and A. Battersby. Hybrid genetic algorithms - a review. 2006.

[129] Harriet B. Creighton and Barbara McClintock. A Correlation of Cytological and Genetical Crossing-Over in Zea Mays. *Proceedings of the National Academy of Sciences of the United States of America*, 17(8):492–497, August 1931.

[130] Daniel J. Kane. Principal Component Generalized Projections FROG Algorithm. In *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses*, pages 357–366. Springer, Boston, MA, 2000.

[131] E. Zeek, A. P. Shreenath, P. O'Shea, M. Kimmel, and R. Trebino. Simultaneous automatic calibration and direction-of-time removal in frequency-resolved optical gating. *Applied Physics B*, 74(1):s265–s271, June 2002.

[132] S. N. Dixit, I. M. Thomas, B. W. Woods, A. J. Morgan, M. A. Henesian, P. J. Wegner, and H. T. Powell. Random phase plates for beam smoothing on the nova laser. *Appl. Opt.*, 32(14):2543–2554, May 1993.

[133] T. R. Boehly, R. S. Craxton, T. H. Hinterman, J. H. Kelly, T. J. Kessler, S. A. Kumpan, S. A. Letzring, R. L. McCrory, S. F. B. Morse, W. Seka, S. Skupsky, J. M. Soures, and C. P. Verdon. The upgrade to the OMEGA laser system. *Review of Scientific Instruments*, 66(1):508–510, January 1995.

[134] C. W. Siders, J. K. Crane, M. C. Rushford, L. C. Haefner, J. E. Hernandez, J. W. Dawson, R. J. Beach, W. J. Clark, D. J. Trummer, G. L. Tietbohl, and C. J. Barty. 10-kJ Status and 100-kJ Future for NIF PetaWatt Technology, July 2007.

[135] M. S. Sodha, R. P. Sharma, and S. C. Kaushik. Interaction of intense laser beams with plasma waves: Stimulated Raman scattering. *Journal of Applied Physics*, 47(8):3518–3523, August 1976.

[136] A. A. Offenberger, M. R. Cervenan, A. M. Yam, and A. W. Pasternak. Stimulated Brillouin scattering of CO2 laser radiation from underdense plasma. *Journal of Applied Physics*, 47(4):1451–1458, April 1976.

[137] P. Michel, L. Divol, E.L. Dewald, J.L. Milovich, M. Hohenberger, O.S. Jones, L. Berzak Hopkins, R.L. Berger, W.L. Kruer, and J.D. Moody. Multibeam Stimulated Raman Scattering in Inertial Confinement Fusion Conditions. *Physical Review Letters*, 115(5):055003, July 2015.

[138] J. F. Myatt, R. K. Follett, J. G. Shaw, D. H. Edgell, D. H. Froula, I. V. Igumenshchev, and V. N. Goncharov. A wave-based model for cross-beam energy transfer in direct-drive inertial confinement fusion. *Physics of Plasmas*, 24(5):056308, April 2017.

[139] J. T. Mendonça. Kinetic description of electron plasma waves with orbital angular momentum. *Physics of Plasmas*, 19(11):112113, November 2012.

[140] R. E. Turner, Kent Estabrook, R. L. Kauffman, D. R. Bach, R. P. Drake, D. W. Phillion, B. F. Lasinski, E. M. Campbell, W. L. Kruer, and E. A. Williams. Evidence for Collisional Damping in High-Energy Raman-Scattering Experiments at 0.26 µm. *Physical Review Letters*, 54(3):189–192, January 1985.