Title
Resolved Outflow Kinematics in Lensed Galaxies at Cosmic Noon

Author
G.C., Keerthi Vasan

Publication Date
2024

Peer reviewed|Thesis/dissertation

Resolved Outflow Kinematics in Lensed Galaxies at Cosmic Noon

By

KEERTHI VASAN G.C.
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

PHYSICS

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Tucker Jones, Chair

Christopher Fassnacht

Andrew Wetzel

Committee in Charge

2024

Contents

Abstract

'Cosmic Noon' is an epoch spanning ∼2–6 billion years after the Big Bang when galaxies are undergoing rapid star formation, morphological evolution, gaining mass and ejecting significant amounts of gas. Strong gravitational lensing by massive foreground galaxies enables detailed studies of the kinematic and morphological properties of galaxies during their crucial formative epochs at cosmic noon. This thesis focuses on studying galactic outflows of gas using strong gravitationally lensed galaxies. Outflows play the crucial role of regulating the galaxy growth by modulating the amount of gas available in a galaxy's interstellar medium at any time. However, fundamental questions about galactic outflows remain unconstrained for galaxies during this epoch. Physical properties such as the velocity at which gas is launched in outflows, the amount of mass lost via these outflows, and the fate of the ejected gas are highly uncertain. Establishing these properties will markedly improve our current understanding of how feedback processes regulate galaxy formation and evolution. This thesis represents significant progress toward establishing these quantities in galaxies at cosmic noon.

The three projects presented in this thesis describe the methodology to identify and spectroscopically confirm the lensed nature of a large sample of newly discovered gravitational lens systems, and conduct spectrally and spatially resolved studies utilizing the magnification from lensed galaxies to measure their outflow properties. By employing state-of-the-art semi-supervised learning techniques within a deep learning architecture and utilizing a training dataset containing both simulated lenses and non-lensed survey images, I demonstrated that we can greatly reduce the human effort required to find lensed candidates from imaging surveys. These machine learning methods dedicated to identifying lensed candidates exhibit remarkable effectiveness, achieving a success rate of ∼90%. I studied a sample of 20 lensed galaxies with good spectral resolution to characterize the kinematic structure of outflows at cosmic noon, finding good agreement with predictions for momentum-driven outflows.

These observations confirm that galaxies at this epoch launch powerful, fast, metal-enriched outflows ubiquitously during their formative stages. I also performed a spatially resolved pilot study of the outflows in one lensed galaxy, which showed that the rates of mass loss due to these outflows are comparable to star formation rates. This suggests efficient coupling between stellar feedback and the driving of outflowing mass in galaxies at cosmic noon. Much of the ejected gas from these outflows, however, remains confined within the galaxy halo, indicating that outflows play a crucial role in shaping the circumgalactic medium and providing a reservoir of gas to sustain extended star formation at later times. Finally, I outline several strategies to study the impact of stellar feedback and its influence on galaxy evolution through future integral field spectroscopic observations of gravitationally lensed galaxies

Acknowledgments

I owe my deepest gratitude to numerous people for their role in getting me where I am today. Foremost among them are my parents (and family). I would not have made it through grad school and reached this far in my personal and professional life without your unconditional love and support.

I am deeply indebted to my PhD advisor, Tucker Jones, for his invaluable guidance and support throughout my academic journey. Tucker is an incredibly kind and down-to-earth mentor whose immense enthusiasm and passion for astronomy, research, and outdoor adventures is profoundly inspiring. Tucker, I aspire to one day become a researcher with your level of rigor, commitment, passion, and fun.

I am also grateful to the rest of the group members – Alex, Erin, Ivana, Jessie, Ryan, Sunny, and Yuguang for all their help and stimulating discussions. Ryan, Yuguang, and Ivana, thank you so much for answering all my questions regarding observations, reductions, and astronomy in general.

Thanks to all my wonderful collaborators, including Karl, Glenn, Vy, Richard, Dan, James, Tania, Nandini, Anowar, Stephen, Po, Themiya and many more. My thesis and qualifying exam committee members – Andrew Wetzel, Chris Fassnacht, David Wittman and James Sharpnack, thank you for all your invaluable insights. This thesis also wouldn't have been possible without the astronomers and staff at Keck. Mahalo! Special mention to Elinor Gates at Lick, under whose guidance I learned to write my first reduction pipeline.

Being part of the astronomy group at UCD has been a wonderful and rewarding experience. I've found great joy in our arXiv coffee conversations, pizza lunches, code review sessions, Astronomy on Tap events, and rooftop gatherings. A special recognition for all the Physics and Astronomy staff members – Nathan, Daniel, and Chris have sorted out all the IT issues I have encountered over the years; Angela, Lenna, Monda, William and Falicia for handling everything on the administrative side.

I am truly fortunate to have shared my PhD journey with an exceptional cohort of graduate students, housemates, and postdocs – Alex, Azalee, Brian, Debora, Erfan, Fei, Gabe, Guga, Imran, Isaiah, Karthik, Matt, Micheal, Patty, Pratik, Pritom, Ramesh, Rukku, Sachin, Sam, Stephen, Victoria, Yize and many more. You guys are the best! Thanks for everything. Special thanks to Alex, Guga, Nihal and Suresh for numerous chaotic intellectual musings. Guga, thanks for introducing me to skiing, serving as my ski mentor and being there with me in all weather conditions as I made it through the bunny slopes to the daring black runs. My skiing experience wouldn't have been the same without Gokul, Stephen, Marie, Alex and others. Here's to all the unforgettable memories we shared on the slopes and the adventures to follow!

CHAPTER 1

Introduction

As a human being on this fascinating planet, gazing upon the vast cosmos on a cloudless night is a profoundly humbling experience; one that transcends temporal and cultural boundaries. Galaxies, which make up much of the night sky and a key part of this experience, are quite special in that their existence and evolution is uniquely driven *only* by the laws of physics. Unlike things on our planet where we, as a species have asserted a sense of dominance by being able to control our environments, we cannot physically interact with and alter the thousands of galaxies in the night sky to our desires (yet). Our inability as a species to do so, perhaps, strangely, has also fueled our immense curiosity to investigate and understand these galaxies that make up the cosmos. However, the only information that we have at our disposal is the light that galaxies radiate and their location on the sky. Therefore, it is not surprising that throughout the history of modern physics and astronomy, we have built sensitive instruments on progressively larger telescopes to capture the light (via images and spectra) from galaxies and devoted our time analyzing their properties (e.g., morphology, kinematics). Doing this for a statistical sample of galaxies on the sky, we have now developed theoretical frameworks capable of describing how galaxies in the universe grow and evolve in time to render the cosmic canvas that we see today. I outline the seminal and widely accepted theoretical foundations of galaxy formation and evolution in Section 1.1.

However, there are still numerous fundamental gaps in our understanding of the physics used in these frameworks and this thesis will focus on establishing them in the most formative early phases of galaxy formation using resolved observations of strong gravitationally lensed galaxies. Specifically, the broad goal of this thesis is to focus on 'galactic ouflows' that are

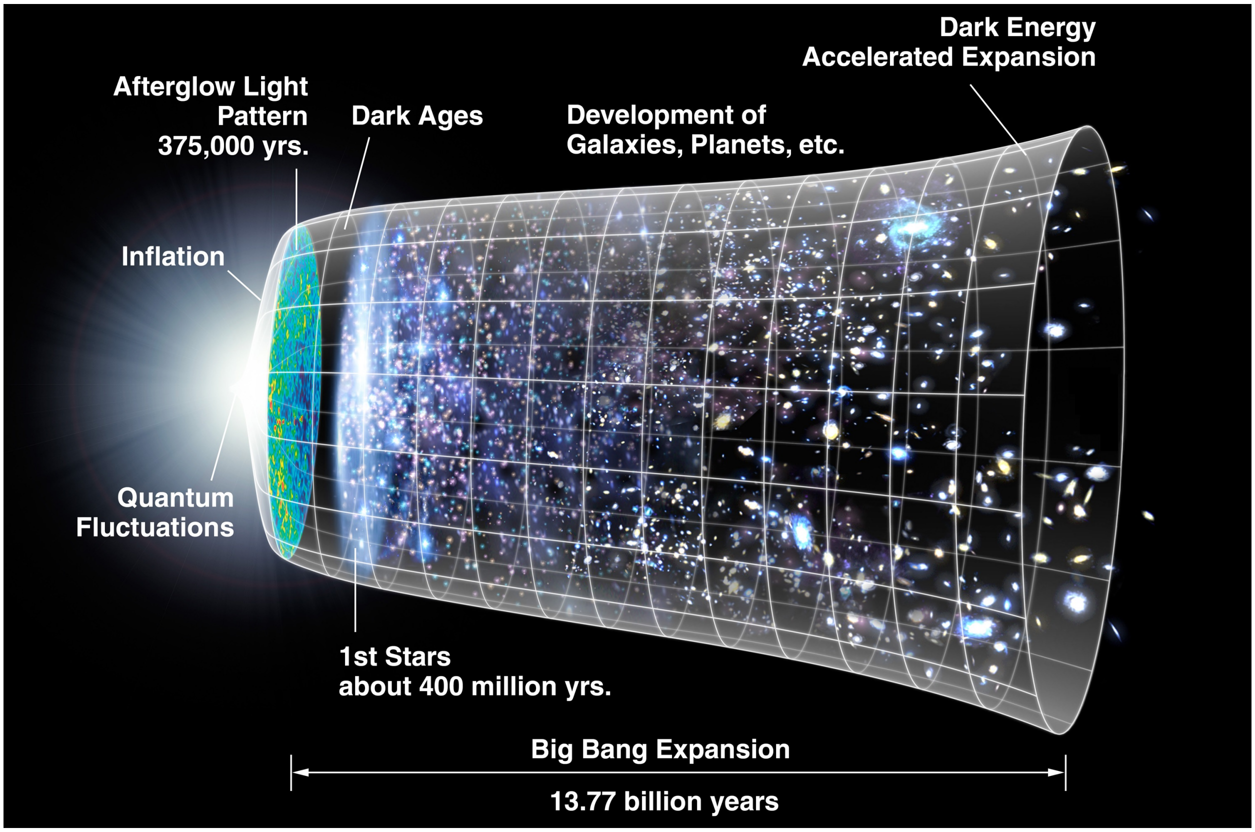

FIGURE 1.1. An illustrative history of the ~14 Billion years of the universe from the Big Bang to the present day (Image Credit: NASA).

powered by 'feedback' processes which play a vital role in shaping a galaxy's evolutionary trajectory.

1.1. Galaxies in the ΛCDM paradigm

Our current theoretical understanding of the universe is best described by the Big Bang and the Lambda Cold Dark Matter (ΛCDM) paradigms (illustrated in Figure 1.1). This model captures the universe's nearly 14 billion year history using a cosmological constant (Λ) associated with dark energy, Cold Dark Matter (CDM), radiation and baryonic matter. The theoretical foundations of these paradigms are described in great detail in many textbooks (e.g., Peebles, 1980; Kolb & Turner, 1990; Cimatti et al., 2019), and I summarize the key components essential to understand galaxy formation and evolution here.

According to this model, the universe began with a hot Big Bang, initially characterized by a hot and dense primordial state undergoing rapid expansion. During the early universe, matter and photons were tightly coupled (undergo frequent interactions) through Thomson scattering from free electrons and this radiation pressure prevented them from gravitating towards the denser regions and collapsing under the influence of gravity. However, this is not the case for cold dark matter particles, which are believed to not interact with the photons, and as a result, continue to grow into increasingly massive large-scale structures, seeding the growth of the cosmic web structures consisting of voids, walls, filaments, and halos that we see today. As the universe reaches around 3000 K ($\sim$380,000 years after the Big Bang), the temperature is finally cool enough for electrons and nuclei to combine into neutral atoms (almost entirely Hydrogen and Helium), thus decoupling from the radiation field. The photons that free stream across the universe from this epoch are our earliest observable probe of the universe and are detected today as a faint background radiation across the sky (known as the Cosmic Microwave Background Radiation; CMBR). Following this epoch is a period known as the 'dark ages' which is characterized mainly by gravitationally-driven growth of large scale structure. Eventually the density of baryonic matter at the centers of dark matter halos becomes sufficient for the first stars and galaxies to form, within $\sim$100-200 million years from the Big Bang. Star formation and hierarchical growth in the following $\sim$13 billion years leads to the galaxies we observe in the universe today.

Numerical simulations using this relatively simple ΛCDM model (e.g., Frenk & White, 2012; Somerville & Davé, 2015a) have been remarkably successful in explaining numerous observations at the large scales, notably the power spectrum and the properties of the CMB, the primordial nucleosynthesis yields and the large-scale structure of the universe. However, within this framework, the recipe to transform the baryonic material into the metal-rich galaxies with complex morphologies such as bars and spiral arms that we see today is not yet fully established.

1.2. The baryon cycle and the role of Outflows

The baryon cycle provides a comprehensive framework to study the formation and evolution of galaxies across cosmic time by disentangling the various physical processes involved. In this cycle, galaxies form and evolve through a combination of accretion, star formation, recycling, and outflows. Clouds of gas fall into cold dark matter halos, where they form stars from the gravitational collapse, converting primordial Hydrogen and Helium into heavier elements through stellar nucleosynthesis. Physical processes (collectively referred to as 'feedback') such as supernova explosions and accretion onto supermassive black holes impart large amounts of energy and momentum to their ambient Inter-Stellar Medium (ISM), driving outflows that transport the gas and heavy elements from nucleosynthesis throughout the galaxy's halo. Some of the gas launched through these outflows recycles back into the ISM at later times, providing the fuel for future star formation, and perpetuating this cycle across cosmic time. This can be summarized using the following simple mass balance equation

$$\dot{M}_{\rm gas} = \dot{M}_{in} - \dot{M}_{out} - \dot{M}_* + R\dot{M}_*$$

where $\dot{M}_{\rm gas}$ is the change in the gas mass, $\dot{M}_{in}$ is the gas accretion rate, $\dot{M}_{out}$ is the outflow rate, $\dot{M}_*$ is the rate at which gas in the ISM is converted to stars and R is the fraction of stellar mass that is returned to the ISM (e.g., Bruzual & Charlot, 2003). Outflows play a vital role in this cycle, acting as regulators of galaxy growth by modulating the gas available in the ISM at any given time. Since stars form from gas in the ISM, an increased mass loss rate via outflows would lead to a reduced gas mass and thus reduced star formation. Outflows, by transporting heavy elements from the ISM, enrich the circumgalactic medium (CGM) that surrounds galaxies. This process is crucial in shaping fundamental properties of galaxies such as their metallicity and the mass-metallicity relationship (see Tumlinson et al., 2017, and references therein).

Gaseous outflows are observed ubiquitously in galaxies where the star formation surface density exceeds $\gtrsim 0.1\mathrm{M}_{\odot}\mathrm{yr}^{-1}\mathrm{kpc}^{-2}$ (Heckman, 2002). Many studies have established the occurrence of outflows across all accessible redshifts (e.g., Heckman et al., 1990; Martin, 2005; Steidel et al., 2010; Rubin et al., 2014; Jones et al., 2018; Vasan G. C. et al., 2023). Star-forming galaxies at high redshifts typically exceed this star formation threshold, and outflows are accordingly quite common in the high redshift universe, contributing to the conclusion that outflows play a vital role in the formative early phases of galaxy formation. The role of outflows has been explored using simplified 'bathtub models' (e.g., Bouché et al., 2010; Lilly et al., 2013; Belfiore et al., 2019) which make reasonable approximations regarding galaxy properties (e.g., that higher star formation rates are associated with higher gas mass and higher outflow rates, and that total gas masses vary slowly on average). These models, along with more comprehensive cosmological simulations with careful treatment of outflows (e.g., Somerville & Davé, 2015a; Bullock & Boylan-Kolchin, 2017), have been largely successful in reproducing many observed scaling relations such as the evolution of the mass-metallicity relation and metallicity gradients, the typical specific star formation rates of galaxies, and the galaxy stellar mass function.

Based on our current understanding, outflows are driven by a combination of stellar and supermassive black hole feedback processes. Stellar feedback (e.g., supernova) is dominant in galaxies at the intermediate and low mass range (stellar mass $\log M_*/M_{\odot} \lesssim 11$) whereas AGN feedback plays a significant role at the high mass range ($\log M_*/M_{\odot} \gtrsim 11$). The galaxies studied in this thesis have intermediate masses, such that we expect stellar feedback processes to dominate. An example of a galactic scale outflow driven by stellar feedback in the nearby starburst galaxy M82 is shown in Figure 1.2. Below, I summarize the key processes contributing to stellar feedback in galaxies such as M82 as well as the rest of the galaxies presented in this work:

FIGURE 1.2. Outflows in the nearby starburst galaxy M82 driven primarily by stellar feedback processes (Image Credit: NASA, ESA, JPL-Caltech and the Hubble Heritage Team).

- Mass loss, radiation pressure and photo-ionization from young stars: Young stars dynamically interact with the ISM, contributing significantly to stellar feedback. These stars emit copious amount of photons, which can be absorbed and scattered by the ISM gas, injecting momentum into it. This interaction not only ionizes the gas but also increases its temperature. Furthermore, young stars expel winds from their surfaces at high velocities, reaching up to $\sim$ 1000 km/s. This shocks and provides a large amount of thermal energy to heat the surrounding gas.
- Supernovae: When stars explode as supernovae (core-collapse and Type Ia), they inject large amounts of energy, momentum and heavy elements into the ISM. Core-collapse supernovae occur on shorter timescales and more frequent compared to Type Ia (e.g., Tsujimoto et al., 1995). In a galaxy, several such supernovae events

can overlap to form 'Superbubbles' leading to large-scale outflows of gas from the ISM.

- Cosmic Rays: Approximately 10% of the $\sim 10^{51}$ ergs in kinetic energy released from a supernova explosion is thought to go into primary cosmic ray protons (and other nuclei). This could potentially play an important role in stellar feedback by applying an outward pressure force on the surrounding ISM. Several mechanisms have been proposed for their effect on the ISM and are still being investigated (e.g., Chan et al., 2019).

Theoretical analyses suggest that outflows are driven by a complex interplay between these various feedback processes (e.g., Murray et al., 2005, 2011; Hopkins et al., 2014; Recchia et al., 2016; Fielding et al., 2018). Since the resulting outflows are challenging to predict theoretically, careful observational measurements are important to determine the actual properties of outflows, and constrain the feedback models used in theory and simulations.

1.3. Galaxies at Cosmic Noon

'Cosmic Noon' is a dynamic and transformative period in the universe's history spanning roughly ∼2–6 Gyr after the Big Bang. This epoch is characterized by galaxies in various stages of development experiencing rapid star formation, morphological evolution, mass growth, and ejecting significant amounts of gas from their ISM via outflows (e.g., Figure 1.3). Examining galaxy formation during this epoch provides valuable insight into the complex evolutionary trajectories of galaxies in the universe. Stellar feedback during this epoch is predicted to be highly efficient with high mass loss rates and outflow velocities (e.g., Pandya et al., 2021; Nelson et al., 2019).

To reliably study outflows, we require resolved observations of galaxies. Spectrally resolved observations are essential to probe the outflow kinematics of the ISM gas such as maximum outflow velocity, and to establish global outflow scaling relations. Spatially resolved observations are vital to estimate the mass loss rates and dependence of outflow velocity on

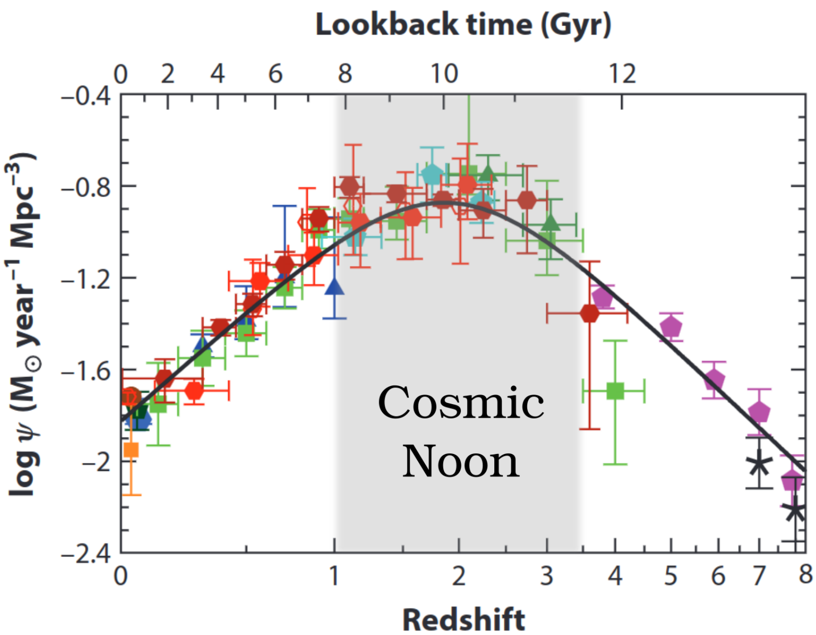

FIGURE 1.3. Plot of Star Formation Rate Density (star formation rate per cubic megaparsec) across cosmic time (modified from Madau & Dickinson, 2014) [1]. This thesis focuses on studying galaxies during the epoch of 'Cosmic Noon' (denoted by the gray region), a period roughly 2 – 6 billion years after the Big Bang during which the star formation rates peaked in galaxies before exponentially declining.

local global properties such as star formation surface densitiy. Such high-resolution observations have been carried out in great detail in nearby galaxies such as M82 (e.g., Lehnert et al., 1999; Xu et al., 2022b, Figure 1.2). However, galaxies at Cosmic Noon are faint and have small angular sizes on the sky making it extremely challenging to study them. Thus, fundamental quantities essential to understand feedback at Cosmic Noon still remain subject to large uncertainties.

This thesis makes use of strongly gravitationally lensed galaxies, which are brighter and magnified compared to unlensed galaxies, mitigating major observational challenges associated with studying galaxies at early times. Combined with strong lensing, the epoch of

[1]Republished with permission of Annual Reviews, Inc., from Annual Review of Astronomy and Astrophysics, 17 Aug 2014, Vol. 52, Issue 1, pages 415 - 486, Cosmic Star-Formation History (2014) by Piero Madau and Mark Dickinson; permission conveyed through Copyright Clearance Center, Inc.

Cosmic Noon is relatively accessible for outflow studies using rest-frame UV ISM spectral features, which are redshifted to optical wavelengths where they can be observed with sensitive instruments on 8-10m class telescopes. In particular, the advent of powerful integral field spectrographs (IFS) on large ground-based telescopes now allow us to spatially map lensed galaxies during their formative early phases. In recent years, numerous lensed galaxies at Cosmic Noon have been discovered in wide-area surveys (e.g. Tran et al., 2022a, discussed in Section 1.4), making this a prime time for their investigation.

1.4. Strong gravitationally lensed galaxies

Gravitational lensing (e.g., Schneider et al., 1992) is an observational phenomenon wherein a massive object deflects light from a background object, which can lead to large magnification of the total flux and angular size of the background object compared to its unlensed configuration. Strong gravitational lensing (e.g., Narayan & Bartelmann, 1996) is a special case of lensing when the background galaxy is brighter, magnified and spatially stretched out into multiple images on the sky in an arc or ring configuration. Figure 1.4 shows a montage of strong gravitationally lensed galaxies.

The main observational challenge with strong lensing is that the spatial alignment of a background galaxy with a suitably massive foreground galaxy (or cluster) is quite rare. For example, in the entire SDSS imaging survey, only of order 100 lensed galaxies have been discovered (e.g., Belokurov et al., 2009). This landscape has entirely changed recently, thanks to the rich availability of computing resources, progressive machine learning techniques and the advancements in imaging capabilities. These technological strides now enable efficient searches of wide-area imaging surveys to find thousands of lensed galaxy candidates (e.g., Sonnenfeld et al., 2018; Jacobs et al., 2019a; Petrillo et al., 2019; Huang et al., 2020; Rojas et al., 2022). For example, Jacobs et al. (2019a) used a convolutional neural network (CNN) approach on millions of images of galaxies in the entire DES/DECaLS fields, finding

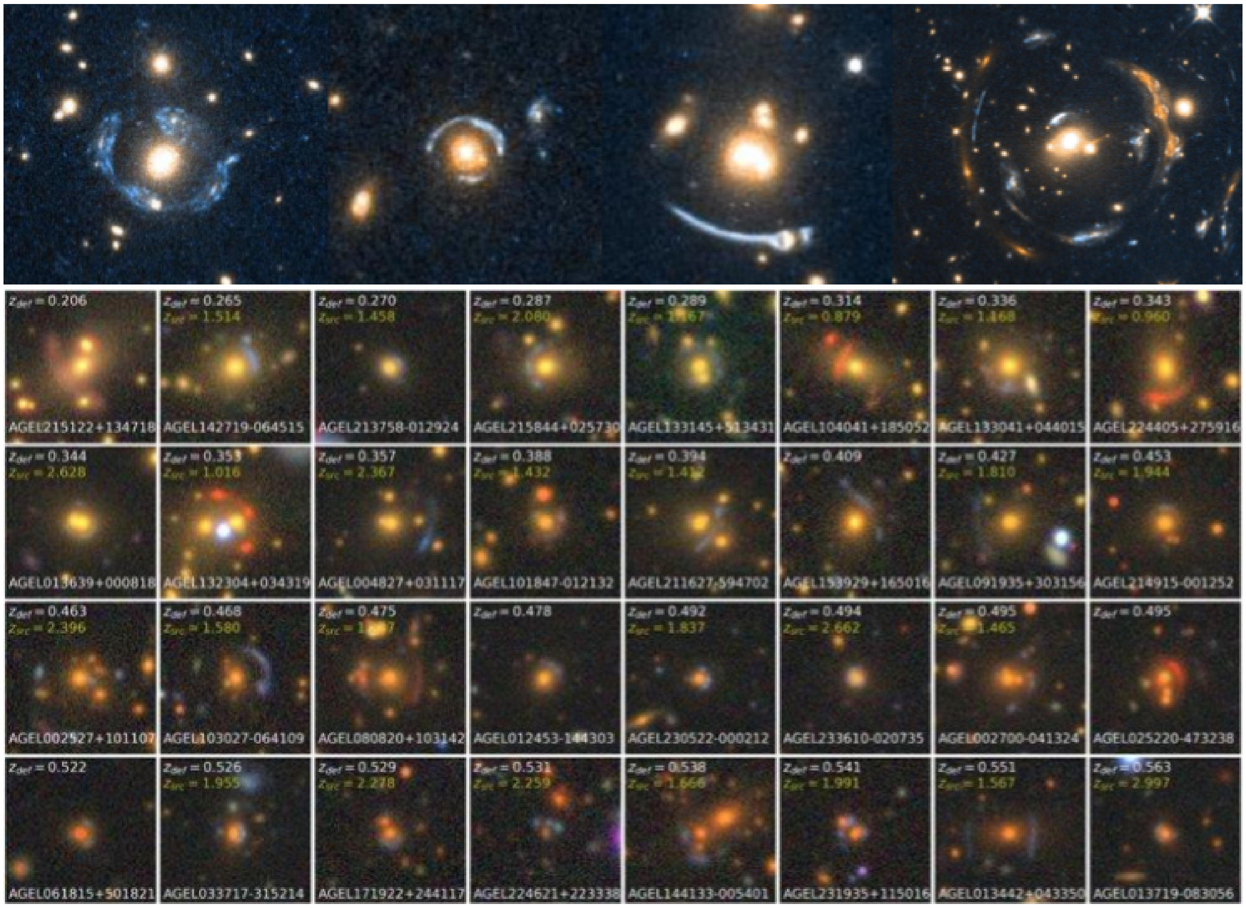

FIGURE 1.4. Color composite images of a few gravitationally lensed systems that were discovered using machine learning techniques (Jacobs et al., 2019b; Keerthi Vasan et al., 2023) in ground-based surveys, for which I led the spectroscopic confirmation campaign as part of the AGEL survey (Tran et al., 2022c). The lensed galaxies in these images are at high redshifts ($z \sim 1 - 4$), and are brighter and have enlarged apparent size compared to unlensed galaxies at this epoch thanks to typical magnification factors of ~ 10. These systems are therefore ideal for probing the nature of feedback in galaxies at high redshifts. Furthermore, thousands of such lensed systems will be uncovered by future space missions (e.g., Euclid, Roman) which will survey the sky at greater sensitivity and resolution compared to current data.

~ 500 lensed candidates. In Chapter 2, I demonstrate that pipelines employing state-of-the-art deep learning models along with data augmentation techniques tasked with finding strong lens systems can be highly efficient, minimizing the amount of human inspection required. Gravitational lens modeling techniques have seen similar improvements. Modelling the foreground galaxy mass distribution is essential to remove the lensing effect ('de-lens') and examine galaxy properties in the source plane, but this process can be highly time consuming and challenging, often requiring high resolution space-based or adaptive optics imaging. Several automated/semi-automated lens modelling pipelines have been developed in recent years (e.g., Ertl et al., 2023; Sahu et al., 2024) which have significantly reduced the time required by an investigator to model the foreground galaxy mass distribution, thus enabling rapid lens modelling of the hundreds of newly discovered lensed galaxy systems.

One of the notable advantages of the lenses selected morphologically using Machine Learning techniques is that the arcs are bright and well resolved in the ground-based imaging, i.e., the massive deflector galaxy and the arc are several arcseconds apart. This is contrasted with spectroscopic approaches where the sources cannot be resolved in ground-based follow-up programs (e.g., Bolton et al., 2006; Brownstein et al., 2012) where the lenses have smaller Einstein Radii ($\Theta_E \leq 1.5''$). As part of this dissertation, I have searched for and spectroscopically confirmed many such lens systems, including as part of the ASTRO 3D Galaxy Evolution with Lenses (AGEL) survey discussed below. A subset of confirmed systems from AGEL are shown in Figure 1.4, exhibiting dramatic strong lensing morphologies with typical image separations of several arcseconds

1.5. Thesis overview

I summarize the key projects that form the basis of my thesis below. This thesis is divided into two broad projects: (a) Identification and spectroscopic confirmation of lensed galaxies, and (b) Studying galactic outflows and stellar feedback in lensed galaxies in the early universe.

Chapter 2 (Keerthi Vasan et al., 2023) describes deep learning techniques to identify strong gravitational lenses on the sky. This project focuses on optimizing models and augmentations to boost the performance of the machine learning models to minimize the human effort of sorting through thousands (or even millions) of images. A companion paper focusing on the methodology used in this project is presented in Sheng et al. (2022, on which I am the second author). In addition to identifying lens candidates, I led a spectroscopic campaign to confirm the lensing nature of $\gtrsim 100$ such candidates found using machine learning techniques as part of the AGEL survey (described in detail in Tran et al., 2022c, and Barone et al., in prep). One of the cornerstones of this campaign was to establish that machine learning methods tasked with finding lensed candidates are highly effective, with a success rate (defined as the number of spectroscopically confirmed lenses among the total number of lensed candidates observed) of $\sim 90\%$. A subset of the data obtained from this survey are featured in the work presented in Chapters 2 and 3, and Chapter 5 outlines future projects that I will work on using this sample.

The rest of this thesis (Chapters 3 and 4) primarily focuses on lensed galaxies at Cosmic Noon. Key outflow properties are currently unconstrained for these galaxies, and the work presented in this thesis represents significant progress in characterizing typical outflow properties. Chapter 3 uses slit spectroscopy to probe the outflowing gas kinematics at good spectral resolution in a sample of 20 lensed galaxies, and Chapter 4 uses deep integral field spectroscopy (IFS) to spatially probe outflow properties in a single galaxy. In these chapters, I describe new measurements of resolved velocities and column densities of outflowing gas from individual galaxies, which constrain their outflow kinematics, mass loss rates and mass loading factors. This in turn provides a better understanding of the effects of feedback processes and their role in galaxy formation, particularly at the most formative epoch of Cosmic Noon.

CHAPTER 2

Optimizing machine learning methods to discover strong gravitational lenses in the Deep Lens Survey

Published as Keerthi Vasan et al. (2023) in Monthly Notices of the Royal Astronomical Society, Volume 524, Issue 4, October 2023.

2.1. Abstract

Machine learning models can greatly improve the search for strong gravitational lenses in imaging surveys by reducing the amount of human inspection required. In this work, we test the performance of supervised, semi-supervised, and unsupervised learning algorithms trained with the ResNetV2 neural network architecture on their ability to efficiently find strong gravitational lenses in the Deep Lens Survey (DLS). We use galaxy images from the survey, combined with simulated lensed sources, as labeled data in our training datasets. We find that models using semi-supervised learning along with data augmentations (transformations applied to an image during training, e.g., rotation) and Generative Adversarial Network (GAN) generated images yield the best performance. They offer 5–10 times better precision across all recall values compared to supervised algorithms. Applying the best performing models to the full 20 deg^2 DLS survey, we find 3 Grade-A lens candidates within the top 17 image predictions from the model. This increases to 9 Grade-A and 13 Grade-B candidates when 1% ($\sim$ 2500 images) of the model predictions are visually inspected. This is $\gtrsim 10\times$ the sky density of lens candidates compared to current shallower wide-area surveys (such as the Dark Energy Survey), indicating a trove of lenses awaiting discovery in upcoming deeper all-sky surveys. These results suggest that pipelines tasked with finding strong lens systems

can be highly efficient, minimizing human effort. We additionally report spectroscopic confirmation of the lensing nature of two Grade-A candidates identified by our model, further validating our methods.

2.2. Introduction

Under rare alignment configurations, the gravitational potential of a massive galaxy can cause light from a distant galaxy located behind it to take multiple paths around it. This results in the formation of several distinct images of the distant galaxy around the massive galaxy, a phenomenon known as strong gravitational lensing (e.g., Treu, 2010). These multiple images are magnified by factors that can reach >10 times, making them appear brighter and more spatially extended. Such magnification makes these systems ideal for studying the formation and evolution of galaxies across cosmic time (e.g., Wuyts et al., 2014a; Pettini et al., 2002a; Swinbank et al., 2009a; Koopmans et al., 2006; Leethochawalit et al., 2016a), while analysis of the lensing mass distribution enables insight into the nature of dark matter (e.g., Chiba, 2002; Bradač et al., 2002; Miranda & Macciò, 2007; Gilman et al., 2019; Shajib et al., 2022).

The main current challenge in working with strong lens systems is their scarcity on the sky. Therefore, methods which are able to efficiently identify lensed galaxies from wide-area sky surveys are extremely beneficial. Automated methods will be especially valuable for lens searches in upcoming wide-area sky surveys to be carried out by the Vera Rubin Observatory, Euclid, and Roman (e.g., LSST Science Collaboration et al., 2009; Laureijs et al., 2011; Spergel et al., 2015), whose improvements in sensitivity, angular resolution and sky coverage will enable detection of far more lens samples than are currently known.

Early approaches to finding strong lens systems included various algorithms searching for multiple lensed images or arc shapes, manual searches around massive galaxies, and citizen science projects (e.g., Moustakas et al., 2007; Paraficz et al., 2016; Seidel & Bartelmann, 2007; Gavazzi et al., 2014; Alard, 2006; Fassnacht et al., 2004; More et al., 2016; Belokurov et al., 2009; Diehl et al., 2009; Garvin et al., 2022). While successful, these methods are time-consuming and difficult to incorporate into an automated framework. Convolutional Neural Networks (CNNs; LeCun et al., 1989; Krizhevsky et al., 2012), which have been successfully developed into a standard tool in the field of computer vision in the past decade, are a promising approach to solving image recognition problems. Depending on the problem, there are various neural network architectures that can be optimized for the desired objectives. CNNs and machine learning techniques in general have indeed been used with success in the past few years to uncover gravitationally lensed candidates in wide-area imaging surveys (e.g., Jacobs et al., 2017, 2019a; Sonnenfeld et al., 2018; Pourrahmani et al., 2018; Huang et al., 2020; Li et al., 2020; Cañameras et al., 2020).

Most machine learning searches for lenses have relied primarily on supervised learning methods (i.e., using a data set consisting of labeled lensed and non-lensed galaxies to train a model). However, while non-lensed galaxies are plentiful, current surveys have very few known lenses to be used as positive labels. Instead, machine learning models are trained on simulated lenses, which can be generated in abundance (e.g., Jacobs et al., 2017). However, this presents a new problem, that the training data distribution (i.e., the simulated lenses) differs from the test data distribution (i.e., the real lenses) – a problem called distribution shift (Quinonero-Candela et al., 2008). To overcome distribution shift, machine learning researchers have repurposed semi-supervised learning methods, which use unlabeled data and data augmentation to adapt the trained model to the test data (Berthelot et al., 2021).

An advantage to the semi-supervised learning approach is that it can learn from the abundance of unlabeled images from the survey, which allows models to generalize better to unseen images. This is particularly useful to improve performance given millions of galaxy images that are detected in sky surveys but not included in the training data. The model performance is further improved through augmentations applied to images during training (e.g., translation and rotation). In addition to conventional transformations, a rich source of data augmentation can be derived by making use of unsupervised learning algorithms (e.g., Goodfellow et al., 2014; Kingma & Welling, 2014; Erhan et al., 2010). Given the range of methodologies available, we now seek to address the question of which combination of machine learning methods (supervised and semi-supervised) and augmentations are best suited for finding strong gravitational lenses.

We seek efficient models which minimize human effort by reducing the number of images that must be visually inspected to recover a given sample of lenses. In this work we apply CNN models to the Deep Lens Survey (DLS; Wittman et al. 2002), which has relatively good image quality and also remains relatively unexplored in terms of machine learning searches, thus serving as a good testbed for this study. Also, because of the small size of known lenses from the DLS survey, we reserve those for use only in our test dataset. Training and validation datasets will only contain simulated lenses. In our previous methodology paper (Sheng et al., 2022, hereafter S22), we discussed the CNN models and lens detection techniques used in this work. Herein, we describe our training data in detail and focus on evaluating the performance of the different models on the DLS dataset.

This paper is organized as follows. In Section 2.3 we give an overview of the Deep Lens Survey and our source selection used for this work. We summarize our machine learning architecture and learning methods in Section 2.4. Section 2.5 describes the method used to generate training, validation, and testing data from DLS images. Section 2.6 discusses our metric to evaluate the performance of the different CNN models. We discuss the results from our experiments in Section 2.7, including the sample of new lens candidates from DLS and spectroscopic confirmation of two systems. Finally, we summarize the main conclusions in Section 2.8. Throughout this paper we use the AB magnitude system and a ΛCDM cosmology with $\Omega_M = 0.3$, $\Omega_\Lambda = 0.7$ and $H_0 = 70$ km s^{-1} Mpc^{-1}.

2.3. Deep Lens Survey Data

Here we give a brief overview of imaging data from the Deep Lens Survey (DLS) which we use to test and optimize strong lens detection methods. The DLS consists of relatively deep imaging over 20 square degrees in five independent $2° \times 2°$ fields which are widely separated in the sky (Wittman et al., 2002). Each field was imaged in $BVRz$ photometric filters (Schmidt & Thorman, 2013) using the 4-meter Mayall telescope at Kitt Peak National Observatory or Blanco telescope at Cerro Tololo Inter-American Observatory, depending on declination. The survey was carried out over ∼120 nights. The survey was designed for weak gravitational lensing measurements, with stringent requirements on image quality and limiting magnitude, such that the data are naturally well suited for identifying strong lens systems. Typical 5σ point-source detection limits are 25.8, 26.3, and 26.9 AB magnitude in the B, V, and R filters respectively (Schmidt & Thorman, 2013). The R band limit is only ∼0.6 magnitudes shallower than the expected depth to be reached by Rubin observatory's 10-year survey (Ivezić et al., 2019). The seeing is by design best in the R band (FWHM$\lesssim$0″.9) and is typically $\gtrsim$0″.9 in the B, V, and z bands (Wittman et al., 2002). Images in the z band are shallowest and typically subject to worse seeing conditions. In this paper, we use only the BVR data.

2.3.1. Source selection and regions of interest. The DLS catalog includes ∼5 million detected galaxies across 20 square degrees. However, only those of moderate redshift and relatively high mass will act as detectable strong lenses (i.e., with Einstein radii $\Theta_E \gtrsim 1$ arcsecond). We applied a magnitude cut of $17.5 < R < 22$ (similar to that used by Jacobs et al. 2017) in order to remove objects which are unlikely to produce a detectable lensing effect. Additionally, we use SExtractor (Bertin & Arnouts, 1996) flags to eliminate saturated low-redshift galaxies, and exclusion masks to remove galaxies around bright stars or at the edge of the field. This results in 281,425 objects (hereafter referred to as the SurveyCatalog). We find that SExtractor flags and exclusion masks remove $\sim 5\%$ of the galaxies from the survey which reduces the effective sky area probed by our SurveyCatalog to $\sim$ 19 square degrees. We set aside 2277 ($\sim$0.8%) randomly sampled object images from this catalog to experiment and tune the HumVI scaling parameters (discussed in Section 2.5.1). All model training analysis in this paper pertains to the remaining set of 279,149 objects (hereafter referred to as the TrainCatalog).

For our analysis we extract image cutouts spanning $25''.7 \times 25''.7$ (100×100 pixels) centered on each object. This size is sufficient for galaxy- and group-scale lenses ($\Theta_E \lesssim 12$"); we do not focus on the most massive cluster lenses which are already well cataloged (Ascaso et al., 2014) and simpler to identify. We create color-composite images from the source BVR FITS files for all targets in the SurveyCatalog (Figure 2.4; discussed in detail in Section 2.5.1). These color composite images have smaller file sizes compared to original data, enabling us to keep the rest of the analysis computationally efficient. These images are still able to capture the detected low-suface brightness features, while not saturating the brightest objects of interest for this work.

Additionally, they are better suited for the machine learning architecture and methods used in this work (discussed in Section 2.4).

2.4. Deep Learning Architecture and learning methods used

The task at hand is to establish a machine learning (ML) algorithm that efficiently classifies the 281,425 color-composite images from the survey into lensed and non-lensed galaxies. Furthermore, by ranking the images from highest predicted probability of being a lens to lowest, we can order the images for human inspection. This requires the selection of an architecture (i.e., a function that takes images as input and gives prediction probabilities as output) and learning methods (i.e., a way for our function to learn from the data). The key components of our ML training pipeline are a supervised convolutional neural net (CNN), domain adaptation with semi-supervised learning, and augmenting training samples with

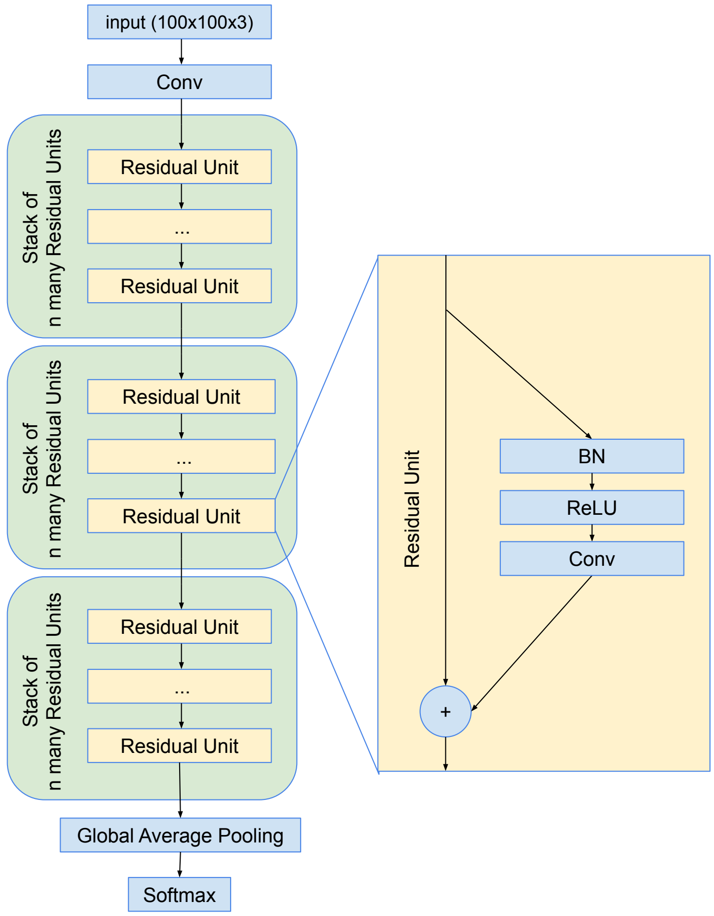

FIGURE 2.1. Schematic depiction of the ResNetV2 deep learning architecture used in this work. The input to the network is an RGB color-composite image generated from the BVR fits files (Section. 2.5.1), and the output is a value between 0 and 1 indicating the probability of the input image being a lens. The general network consists of three stacks, each containing $3n$ residual units. In this work, we use a stack size of $n = 1$ resulting in a total of three residual units. Each residual unit consists of two sets of Batch Normalization (BN), Rectified Linear Unit activation function (ReLU), and Conv units, where Conv denotes a convolutional layer with kernel size 3×3 and appropriate stride size. The network ends with global average pooling and a softmax layer.

generative adversarial nets (GAN). A more detailed account of our ML method can be found in S22.

2.4.1. Convolutional neural network architecture. CNNs have previously been used for classifying and identifying lens candidates (e.g., Jacobs et al., 2017). They are a specific form of neural network that learns translation invariant representations via trainable convolution kernels. This is particularly well suited to astronomical images where patterns are repeated throughout the sky. Deep CNNs are models where these learned non-linear representations of the image (called layers) are stacked on top of one another. Deep CNNs are trained using variations of stochastic gradient descent, where an objective function is evaluated on small subsets of the data, called mini-batches, and the parameters are updated by subtracting some fraction of the objective's gradient.

There are many choices of how precisely these layers are constructed and combined, such as selection of the convolutional kernel size, number of output channels for each convolution layer, the non-linear activation function, and the incorporation of other layers that improve performance such as Batch Normalization (Ioffe & Szegedy, 2015). All of these details together are called the model architecture.

We make use of the ResNet version-2 architecture (ResNetV2; He et al., 2016a,b) designed for the CIFAR10 dataset (Krizhevsky, 2009a), shown schematically in Figure 2.1. It is one of the widely used industry standard networks for image classification problems (e.g., Litjens et al., 2017; Gu et al., 2018; Madireddy et al., 2019). The ResNetV2 used in this work consists of three stacks (see Figure 2.1; green blocks) and each stack consists of n residual unit blocks, where n is a parameter to be chosen that controls the depth of the neural network. A deeper neural network has more learning capacity but requires more computational power and training samples. Each residual unit block consists of three convolution layers of kernel size 3×3 and one skip connection. To match the feature map dimensions (width, height) and the number of channels between stacks, a few extra convolution layers are included at the input and the beginning block of each stack. Therefore, $9n + 4$ convolution layers are present in the network in total. For all the models used in this work, we adopt $n = 1$. With strided convolutions, the feature map dimensions to each stack decrease by a factor of 1/2. The number of input and output channels to each stack are: $(16 \rightarrow 64)$, $(64 \rightarrow 128)$, $(128 \rightarrow 256)$.

The network ends with global average pooling, a fully-connected layer and softmax. The global average pooling constrains the output to be rotationally invariant. The softmax transforms the output to be a value between 0 and 1 which can be interpreted as a probability. Throughout this work, a value of 1 is designated for lensed candidates (referred to herein as Lenses) and 0 for nonlensed candidates (referred to as NonLenses).

2.4.2. Domain adaptation with semi-supervised learning. In supervised learning, our algorithm is trained via mini-batches of images X and corresponding labels y (1 for Lenses and 0 for NonLenses). The algorithm then tries to learn the neural network parameters, collectively referred to as Θ. The output of the neural network after the softmax activation produces a prediction $p_\Theta(X)$, which is our predicted probability of X being a lens. Our supervised learning objective function is the cross-entropy loss function, denoted ℓ_S, which is a measure of the quality of our predictions, $p_\Theta(X)$, when compared to the true labels, y. Merely using supervised learning does not perform well in the face of distributional shift, and we turn to semi-supervised learning (SSL) methods which make use of the unlabeled test data to adapt to this domain. There are many semi-supervised approaches to deep learning. The methods we explore are FixMatch[1] (Sohn et al., 2020), MixMatch (Berthelot et al., 2019), Virtual Adversarial Training (Miyato et al., 2019), Mean Teacher (Tarvainen & Valpola, 2017), Π-Model (Laine & Aila, 2017), and Pseudo-Labeling (Lee, 2013).

Most SSL algorithms follow the same template. We minimize an objective function consisting of a supervised component (i.e. ℓ_S losses), where the label is provided, plus an unsupervised component (i.e. ℓ_U losses). Both are optimized together over mini-batches, now consisting of labeled and unlabeled data, but without significant modification to the stochastic gradient descent algorithm. The main feature that distinguishes our setting from

[1]FixMatch was not part of the original lens search study since this technique had not been published at the time. We are including it in our results here to be thorough.

RGB-shuffle	Randomly perturb the order of the channels in the images
JPEG-quality	50-100%
Rot90	Randomly rotate the images by a multiple of 90 degrees
Translations	Randomly translate the images by at most 20 pixels in the up, down, left and right directions
Horizontal flips	Randomly flips the images across the x-axis
Color augmentation	Randomly perturb the brightness(-0.1-0.1), saturation(0.9-1.3) hue(0.96-1.00), and gamma(1.23-1.25) of the images

TABLE 2.1. Data augmentations used on images in the semi-supervised training pipeline.

typical SSL is that our training NonLenses and test set come from the same pool of data, while the simulated Lenses do not exist in the test data. This is in contrast to Jacobs et al. (2017) for example, in which they produce simulated NonLenses as well, but do not attempt domain adaptation.

In the Pseudo-Label algorithm (Lee, 2013), we assign pseudo-labels to unlabeled data by taking the model's predicted class as the label. We can then use the same loss as in the supervised task (i.e., $\ell_S = \ell_U$). The motivation is that we are implicitly enforcing *entropy minimization* by forcing the model to be confident on unlabeled samples. An alternative approach to SSL is *consistency regularization*, where two independently augmented samples of the same test image are encouraged to produce similar predictions. The Π-model algorithm (Laine & Aila, 2017) directly uses consistency regularization. The idea is to take two random augmentations of the same sample data point, X, and compute the squared difference of the model outputs for the augmented copies. We use $\mathrm{aug}, \widetilde{\mathrm{aug}}$ to denote two independent augmentations, which can be produced by selecting different randomization seeds. The unsupervised loss is then

$$\ell_U(X) = \left\| p_\Theta(\mathrm{aug}(X)) - p_\Theta(\widetilde{\mathrm{aug}}(X)) \right\|^2 . \tag{2.1}$$

The choice of stochastic augmentation function is up to the modeler and will often be domain specific.

The Mean Teacher algorithm (Tarvainen & Valpola, 2017) also uses consistency regularization, but replaces one of the augmentations in Equation 2.1 with the output of the model using an exponential moving average (the teacher model) of model parameters, Θ. FixMatch (Sohn et al., 2020) and MixMatch (Berthelot et al., 2019) employ both consistency regularization and entropy minimization. MixMatch was originally proposed as a heuristic approach, and FixMatch was later derived as a more principled simplification of MixMatch and other related SSL methods. Virtual adversarial training (VAT; Miyato et al., 2019) uses an adversarial, worst-case, augmentation. This adversarial augmentation pushes the image in the direction which will cause the greatest increase in loss. One downside to VAT is that the adversarial augmentations are not able to encode the domain specific prior information that random augmentations can provide (see Table 2.1).

2.4.3. Data augmentation and GANs. Data augmentation serves as a crucial regularizer in semi-supervised learning (SSL) algorithms. Several SSL algorithms, including those mentioned in this paper such as pi-model (Laine & Aila, 2017), MixMatch (Berthelot et al., 2019), and fixMatch (Sohn et al., 2020), utilize data augmentation techniques. The data augmentation techniques we employed in our study are provided in Table 2.1, and are particularly well-suited for DLS images.

RGB-shuffle randomizes the order of channels and Color augmentation perturbs the colors in the images. These have the effect of accounting for systematic bias in channel and color information introduced by the simulation pipeline. JPEG-quality augmentation accounts for varying levels of noise and image quality, and applies to any color composite image irrespective of the format that the image is saved in (e.g., in this case we use png format instead of jpeg). Rot90, Translations, and Horizontal flips induce translational and rotational invariance in the predictions. Examples of these augmentations are shown in Figure 2.2. We note that even though some augmentations (e.g., RGB-shuffle) result in unrealistic images, our empirical tests described in Section 2.7.1.1 indicate that these augmentations yield improved model performance. Domain adaptation problems employing semi-supervised algorithms (SSLs) have been shown to benefit greatly from data augmentations in general (e.g., Sohn et al., 2020), suggesting that this effect is not specific to our lens search.

A second tool that we use to augment our data is to generate new images that mimic the simulated lenses. In deep learning, the state-of-the-art method to produce generative models is by using Generative Adversarial Networks (GANs; Goodfellow et al., 2014; Arjovsky et al., 2017). GANs generate unseen samples that are distinct from the original images, but are

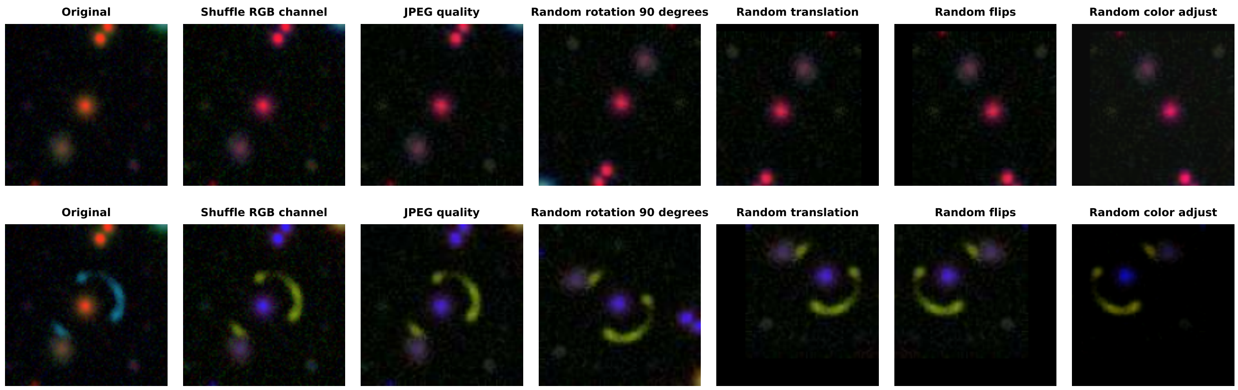

FIGURE 2.2. Example of augmentations used during training. From left to right, the original RGB color composite image undergoes the series of augmentations described in Table 2.1: RGB-shuffle, JPEG quality, Rot90, Translation, Flip, Color adjustment. The final image is then passed as input to the model.

distributionally quite similar. These generative models are trained along with an adversarial discriminator that is attempting to distinguish between the fake and real images.

We trained a WGAN-GP (Wasserstein GAN + Gradient Penalty; Gulrajani et al., 2017) on simulated lenses and add the generated images (see examples in Figures 2.3 and 2.4) to our training set as another form of data augmentation. The motivation is that GANs can provide a rich source of more exotic data augmentations.

Figure 2.3 gives a brief summary of the steps discussed thus far. The training, testing, and validation data along with the model checkpoints used in this paper are made available on our GitHub repository [2].

2.5. Training and Validation data

One of the challenges that we face in gravitational lens searches is trying to generate a training and testing dataset when having limited knowledge of the type of strong lenses that we might find in a survey. Prior to this work, Kubo & Dell'Antonio (2008) used a semi-automated method to search for lensed candidates in one of the DLS fields (F2) and uncovered two lens candidates. But in order to train a machine learning model to recognize

[2]https://github.com/sxsheng/SHLDN

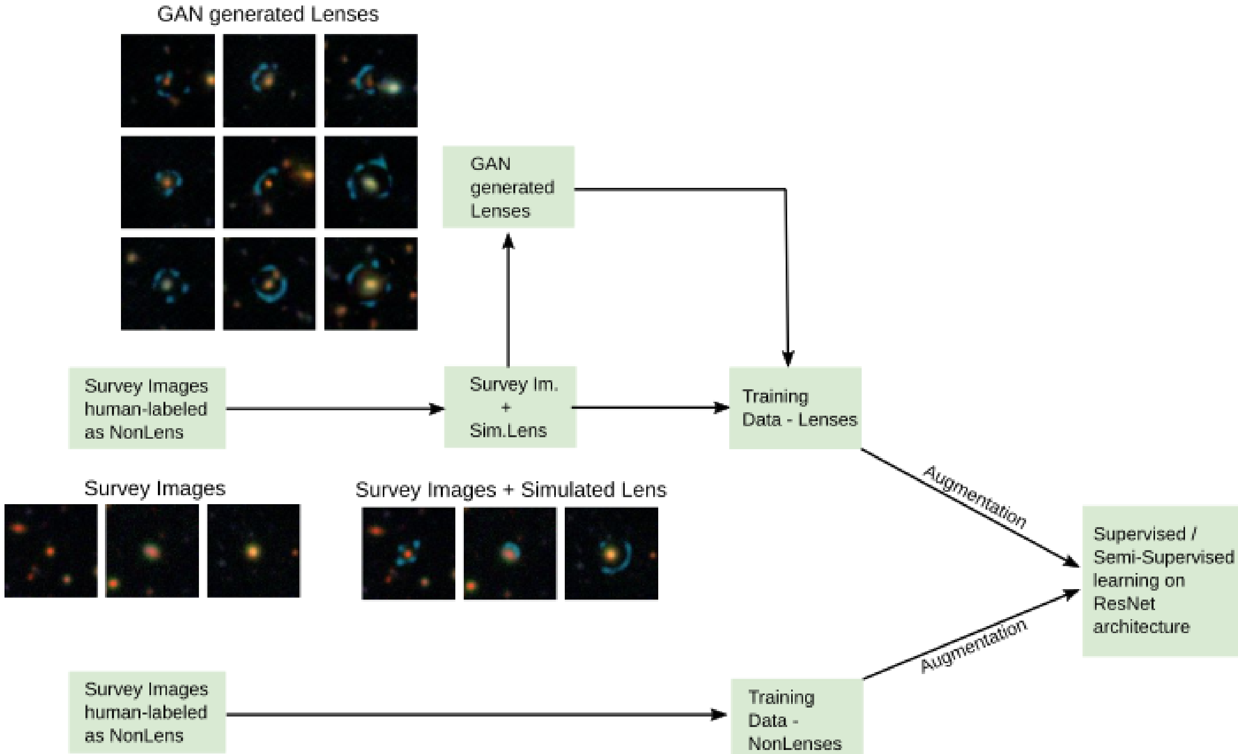

FIGURE 2.3. Schematic of the pipeline used in this work to test the performance of different learning methods described in Sections 2.4.2 and 2.4.3 (see text for details). The GAN generated lenses are only included in the training data for unsupervised learning methods (e.g., GAN+MixMatch).

lenses, we require Lens and NonLens image samples on the order of a few thousand. This is not a problem for NonLens galaxies, as they are abundant. But this is challenging for Lenses, as the known samples are extremely small compared to training requirements. We note that although the DLS area overlaps with other surveys used for strong lens searches (e.g., SDSS), no lens candidates have been published from these other surveys within the DLS footprint. This is likely due to the shallower depth of other surveys (see Section 2.7.4). We must therefore generate an artificial lens training set. We describe our process of generating the training and testing datasets in this section.

2.5.1. Generating the NonLenses dataset. Color png images centered on each object in the SurveyCatalog are constructed from BVR fits files using HumVI (Marshall et al., 2015). HumVI is based on the color composition algorithm described in Lupton et al. (2004) and offers several tunable parameters to control the output image (e.g., contrast). We randomly sample objects from the SurveyCatalog and visually inspect the effect of changing the HumVI parameters s and p which control the contrast and color balance respectively. Although there is a degeneracy in the choice of these values, we pick ones that reasonably represent both the bright and dim features in the data (i.e., spanning the range of detectable surface brightness). Table 2.2 lists our chosen HumVI parameters and Figure 2.4 (top panel) shows 4 randomly selected color-composite survey images generated using these values. The chosen HumVI parameters are kept constant and applied to all images in the survey. It is beyond the scope of this work to explore the effect of choosing different HumVI parameters on the performance of the models, but we note that color augmentations applied during training (Table 2.1; Section 2.7.1.1) have the effect of making our models invariant to small perturbations in color.

2.5.2. Generating the simulated Lenses dataset. As described above, the scarcity of known lensed galaxies requires us to generate simulated lens samples for training ML models. Our approach is to add simulated lensed galaxies onto survey images, as has been used successfully in prior work (e.g., Jacobs et al., 2017, 2019a). For this work, we adopt an agnostic procedure for simulating lensed arcs which does not rely on photometric measurements of the deflector galaxy. We consider all galaxies which satisfy the magnitude cut criteria described in Section 2.3.1 (regardless of their color) for simulating the lensed arcs. We note that $\sim 50\%$ of the galaxies in our SurveyCatalog have a BPZ best fit photometric template from Schmidt & Thorman (2013) indicating that they are massive early-type galaxies at intermediate redshifts, and are indeed likely to act as strong lenses. We discuss the actual color distribution for lens candidates in Section 2.7.2.3.

Given any object from the training dataset, we assume that the central galaxy ("deflector") is at a redshift $z_{\rm def} \in [0.3, 0.7]$ and is characterized by a Singular Isothermal Ellipsoid

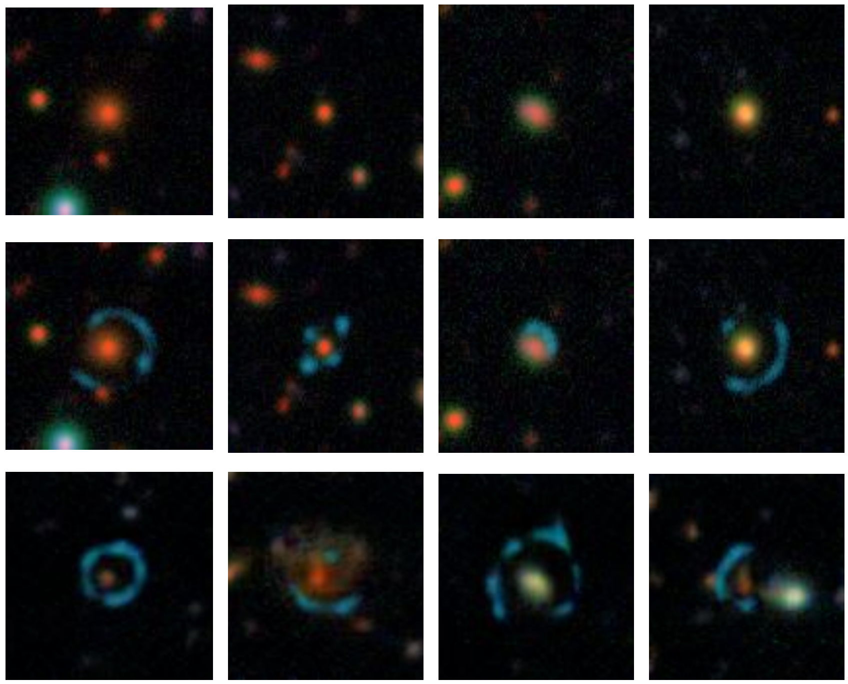

FIGURE 2.4. *Top row:* Four randomly selected color composite survey images generated by running HumVI on their respective BVR FITS files. These images are examples of NonLenses used for training the network. Each image spans 25.″7 × 25.″7 on the sky. Table 2.2 lists the HumVI parameters used to generate these images. *Middle row:* The same set of survey images as in the top row, but superimposed with simulated lens configurations generated with glafic. Section 2.5.2 discusses the steps involved in detail. These images are examples of Lenses used during training. *Bottom row:* GAN generated simulated lenses. These are added to the training data as Lenses for our unsupervised models (e.g., GAN+MixMatch; Section. 2.4.3).

(SIE) mass density (Kormann et al., 1994). The mass profile is dependent on the galaxy's position ($x_{\rm def}, y_{\rm def}$), ellipticity ($e_{\rm def}$), position angle ($\theta_{\rm def}$), velocity dispersion ($\sigma_{\rm def}$), and choice of $r_{\rm core,def}$. The values for these parameters are sampled from a uniform distribution spanning the ranges listed in Table 2.2. These values ensure that the resulting mass profile of the deflector is sufficient to produce a detectable lensing effect (i.e., $\Theta_E \gtrsim 1$ arcsecond). A background galaxy ("source") is assumed to lie at a redshift $z_{\rm src}$ with morphology given by a Sérsic profile parameterized by its position (x_{src}, y_{src}), central brightness (in units of counts/pix^2), ellipticity ($e_{\rm src}$), position angle ($\theta_{\rm src}$), and a Sérsic index of 1. The value for

	Parameter	Value
	glafic	
Position (arcseconds)	$x_{\rm def}, y_{\rm def}, x_{\rm src}, y_{\rm src}$	U(-0.5,0.5)
PA (degrees)	$\theta_{\rm def}, \theta_{\rm src}$	U(0,180)
Ellipticity	$e_{\rm def}, e_{\rm src}$	U(0.3,0.7)
Dispersion ($\mathrm{km\,s^{-1}}$)	$\sigma_{\rm def}$	U(250,450)
	$r_{\rm core,\ def}$	U(0,0.5)
Brightness (counts/pix^2)		U(200,600)
Redshift	$z_{\rm def}$	U(0.3,0.7)
Redshift	$z_{\rm src}$	U($z_{\rm def}$ + 0.5, $z_{\rm def}$ + 2.5)
	HUMVI	
	-s	0.2,0.7,1.3
	-p	2.5, 0.01
	-m	0.1

TABLE 2.2. Values for the glafic and HumVI parameters used to generate the simulated arcs and png color-composite images respectively. $U(x_{min}, x_{max})$ indicates that the value was sampled from a uniform distribution with x_{min} and x_{max} being the minimum and maximum values.

$z_{\rm src}$ is randomly chosen from a uniform distribution between $z_{\rm def} + 0.5$ and $z_{\rm def} + 2.5$. These values for the deflector and source redshifts are typical of spectroscopically measured values from previous strong lens surveys (e.g., Sonnenfeld et al., 2013; Bolton et al., 2008; Tran et al., 2022a).

The light from the background galaxy is traced using glafic (Oguri, 2010) to produce a simulated lensed arc in the image plane. The simulated lensed arcs are convolved with the point spread function (PSF) of the survey, scaled by a factor of (1,1.5,3) for the BVR filters, and then added to the BVR fits images of the galaxy. We model the PSF of the survey in all the three filters as a 2D Gaussian kernel with a FWHM of $\sim$1 arcsecond corresponding to the approximate average seeing conditions. In addition to smoothing, we add Poisson noise in order to produce more realistic simulated arc images. The fits images are converted to a color png image using HumVI (as described in Section 2.5.1). For this paper, we focus on generating moderately bright blue lensed arcs, and the parameter ranges that produce these configurations are listed in Table 2.2. Figure 2.4 illustrates common configurations of the arcs produced using this method. However, we note that the RGB-shuffle augmentation which is applied during training produces arcs of different colors (e.g., Figure 2.2). We find that such an approach, where the simulated arcs are not dependent on the photometric properties of the central deflector galaxy, likely serves as an additional form of augmentation. This approach prevents over-fitting of our deep learning models while allowing for rapid prototyping and testing.

2.5.3. Generating the training datasets: TrainingV1 and TrainingV2. Using the Lenses and NonLenses datasets, we construct two training sets: TrainingV1 and TrainingV2. The main difference between the two training sets is the number of labeled images used as Lenses and NonLenses. Prior work using CNNs (e.g., Jacobs et al., 2019a) have favored large training datasets (i.e., $\gtrsim$150,000 galaxies). Therefore, for TrainingV1 we use 266,301 images for non-lenses and 257,874 corresponding simulations as lenses (described in Section 2.5.2). Since semi-supervised training requires both labeled and unlabeled data, TrainingV1 cannot be used to test semi-supervised learning methods.

For TrainingV2, we choose the number of images for each class to be similar to those used in standard computer vision datasets such as Canadian Institute for Advanced Research-10 (CIFAR-10; Krizhevsky, 2009b) and Street View House Numbers (SVHN; Netzer et al., 2011) dataset. We use a set of 7,074 human-labeled objects as NonLenses and 6,929 corresponding simulations as Lenses. The human labeling was carried out on randomly chosen images from Field-1 (F1) of the DLS. We note that the choice of labeling the data only from F1 does not affect the results presented in the rest of the paper (see Appendix 2.8). The 259,248 NonLens images which are not part of TrainingV2 serve as unlabeled data for our semi-supervised learning methods (e.g., MixMatch; Section 2.4.2).

Counter-intuitively, we find that too much training data from simulated lenses and randomly selected NonLenses can hurt the performance of our algorithms. We refer readers to Section 2.7.1.2 and S22 for further discussion of sample size effects, which can also contribute to differences in performance between the training sets. We note that the TrainingV2 labeled datasets are comparable to the size where we find peak performance.

We performed a 90-10 split for both TrainingV1 and TrainingV2, where 90% of the data was allocated for training the ResNetV2 model and 10% was kept aside for validation. We chose the maximum number of epochs (passes through the training dataset) for each training combination as 100, since this was sufficient to observe a plateau in the validation metrics. For each of the training combination described in Section 2.4, we conducted four independent trials and selected the checkpoint with the best validation metrics for testing it on the survey data.

2.6. Metric to evaluate model performance

We have described several models which are each tuned to optimize validation accuracy, which is measured on the validation dataset (Section 2.5.3) consisting of simulated Lenses and survey NonLenses. In order to gauge the performance of the models on their ability to find real lenses from the survey, we require a testing dataset consisting of lenses from the survey, as well as a metric to evaluate them on.

2.6.1. Generating the Testing dataset. Curating testing data in our case is a challenging task. As discussed earlier, only two strong lenses in the entire survey were known prior to this work, which is insufficient for meaningful evaluation. Therefore, we use an ensemble of 5 ResNet models trained on simulated lenses but using polar transformed images as input to the network. The exclusive task of this model is to find real lens candidates to add to our test dataset. We emphasize that this model is independent of the rest of the models discussed so far in this paper, and does not influence their performance in any way. Details of its implementation are discussed in S22. It is beyond the scope of this paper to quantify the performance of ensemble models or the effect of polar transformation during training, but it is an interesting avenue for future work.

We find 52 likely lens candidates from this model, of which 27 are deemed to be good candidates upon visual inspection. Therefore, we create two testing datasets: TestV1 and TestV2. TestV1 contains all the 52 lens candidates found using our ensemble model approach, while TestV2 contains the 27 best visual candidates. NonLenses for both TestV1 and TestV2 were formed by randomly selecting 874 of our 8734 human-labeled non-lenses (Section 2.5).

2.6.2. Precision and Recall. A standard metric widely used in machine learning to evaluate the performance of test data on a trained model is the Precision-Recall curve (PR curve), where precision and recall are defined as follows:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad \text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}. \tag{2.2}$$

Here TP, FP, and FN are the number of True Positive, False Positive, and False Negative images respectively. These values are computed by passing a labeled test dataset (TestV1 and TestV2 in this case) through a trained model (e.g., GAN+Mixmatch) and setting different prediction thresholds.

Since the primary goal of this work is to find models which minimize the number of nonlensed images that an investigator encounters while maximizing the number of lensed images found (i.e., less FP and FN values), we seek models which have high precision at high recall. We present the results from our PR curve analysis in the next section.

2.7. Results and Discussion

2.7.1. Semi-supervised algorithms with GANs and Augmentations have superior performance. We consider 17 variations on the learning approaches described in Section 2.4: 4

supervised, 6 semi-supervised, and 7 semi-supervised with GANs. SupervisedV1 and SupervisedV2 are our baseline models. They were trained using a supervised learning approach with no data augmentation on TrainingV1 ($\sim$250,000 Lenses and NonLenses) and TrainingV2 ($\sim$7000 Lenses and NonLenses) respectively. On the other hand, SupervisedV1+DA and SupervisedV2+DA were trained using supervised learning with data augmentation (DA). The rest of the models were trained on TrainingV2 using semi-supervised learning methods with DA or with DA + GANs. In this subsection, we summarize the performance of these different models. We primarily use the PR curve (Section 2.6.2) evaluated on our TestV1 and TestV2 sets to gauge which models perform best. We note that our methodology paper S22 includes an additional discussion of these results.

We plot the PR curve obtained for our best-performing baseline models (SupervisedV1, SupervisedV2) along with a subset of semi-supervised and GAN+semi-supervised models in Figure 2.5 (see Tables 3 and 4 of S22 for additional model results). Table 2.3 lists the precision value obtained for a subset of models at 100% recall. We find that our models tend to generalize poorly when trained without any augmentations. Our baseline models, trained without any data augmentation, performed worst out of all models at every recall level. For example, at 100% recall, the baseline SupervisedV1 and SupervisedV2 have a precision of $\sim 3\%$ on our TestV2 set, whereas the GAN+Π-model has a precision of $\sim 22\%$. The poor precision values of our supervised models may reflect challenges in simulating the characteristics of lenses from a survey given limited priors. Fortunately, we find that data augmentation methods are able to address this problem. We find a factor $\sim$5-10$\times$ improved precision across almost all recall levels when applying the full set of augmentations (Table 2.1) to our supervised models.

The improvement of semi-supervised over supervised algorithms suggests that valuable features can in fact be extracted from the mostly unlabeled NonLenses, providing benefits in the classification of real lenses. Adding GAN images to our training pipelines had a seemingly profound impact at all recall levels, especially at higher recalls where more difficult-to-classify images come into play. This suggests that GAN-generated images contain subtle variations which, while not necessarily significant to the naked eye, do in fact produce a strong regularizing effect when used in training.

2.7.1.1. *Ablation study on data augmentations.* We investigated the impact of each of the data augmentations we used by doing an ablation study using TrainingV2. The results from this study are tabulated in Table 2.7. We find that removing GAN images from the training sets causes a noticeable decrease in model performance at all recall levels, which agrees with our earlier conclusion. It also appears that color augmentations and JPEG quality play a very significant role in model performance. Including these three augmentations in our training pipelines is apparently what allows our model to generalize so well, despite relying on simulated lenses for training. A curious result from this ablation study is that multiples of 90-degree rotations actually had a negative effect on model performance. The difference in performance is relatively small compared to that seen for other augmentations (e.g., GANs), but persists at all recall rates. A possible reason for this could be our small validation and test sets. Because the validation set is small, model selection may be biased towards certain orientations of the image. Likewise, an equally small test set may have preferred orientations that the model does not generalize to, resulting in degraded performance.

2.7.1.2. *Larger non-lens training samples can degrade the classifier's performance.* To understand why our larger training set (TrainingV1) led to poorer generalization, we also performed a test where we fixed the number of simulated Lenses and varied the number of NonLenses in the dataset (see S22, Table 5). As we gradually increased the number of NonLenses in the training data from 0 to 256,000, we saw that precision gradually increased and peaked at around 8000-16000 NonLenses, then started to significantly decrease to around $\sim$6% precision for nearly all recall levels. One possible explanation for this effect is that as we increase the number of NonLenses in training, we also increase the number of NonLens

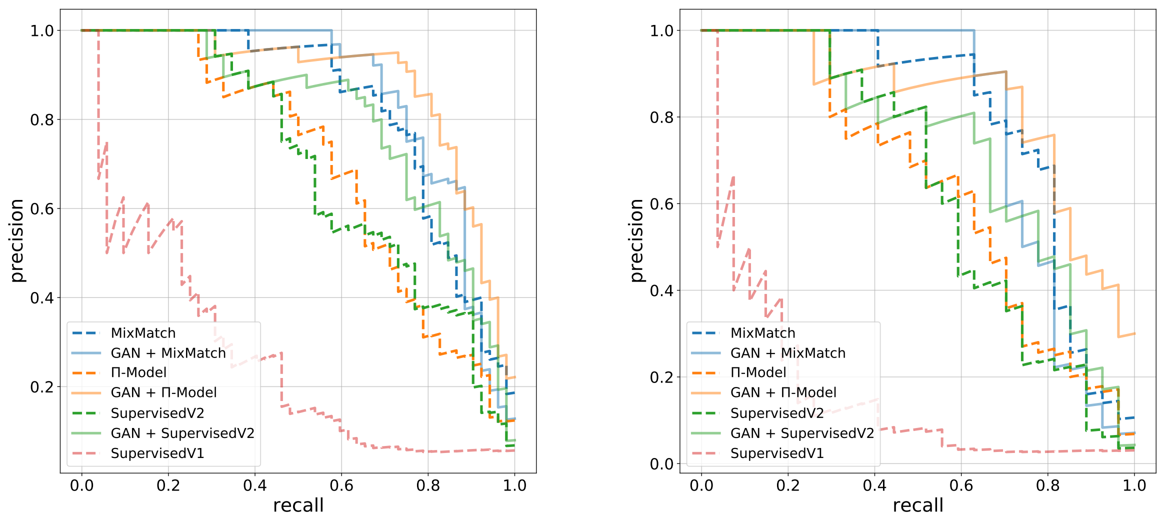

Figure 2.5. Precision-Recall curves (PR curves) for a subset of the models described in Section 2.4 obtained using TestV1 (*Left*) and TestV2 (*Right*). TestV1 contains 52 lens candidates found using our ensemble model approach, while TestV2 contains the 27 best visual candidates (Section 2.6.1). SupervisedV1 and SupervisedV2 are our baseline models. They were trained using a supervised learning approach with no data augmentation on TrainingV1 (∼250,000 Lenses and NonLenses) and TrainingV2 (∼7000 Lenses and NonLenses) respectively. The rest of the models were trained on TrainingV2 with augmentations. MixMatch and Π-Model are semi-supervised learning approaches, whereas GAN+MixMatch and GAN+Π-Model use GAN generated images along with semi-supervised learning (see Figure 2.3 and Section 2.4 for details). GAN+SupervisedV2 uses supervised learning with GAN generated images. Models which use semi-supervised learning along with GANs clearly outperform our baseline supervised learning models at all recall values, with GAN+Π-model having the highest precision at 100% recall (see results in Table 2.3; we note that Table 2.3 reports the average of our four runs while this figure shows the runs with the best precision).

false positives which appear similar to real lenses in the survey data (and perhaps even more similar to real lenses than the simulations we use). As a result, the decision boundary for non-lenses overlaps more with the regions occupied by real lenses, leading to higher levels of misclassification. Therefore, care must be taken in constructing training data based

Model	Training data used	TestV1 Precision(%)	TestV2 Precision(%)
SupervisedV1	TrainingV1 w/ no augmentation	5.62±0.01	3.01±0.02
SupervisedV2	TrainingV2 w/ no augmentation	5.65±0.02	3.06±0.04
MixMatch	TrainingV2 w/ augmentation	12.28±5.09	6.84±3.00
Π-Model	TrainingV2 w/ augmentation	13.41±2.33	8.68±1.49
GAN + Supervised	TrainingV2 w/ augmentation	8.25±2.85	6.05±2.69
GAN + MixMatch	TrainingV2 w/ augmentation	14.13±6.53	7.97±3.93
GAN + Π-Model	TrainingV2 w/ augmentation	**15.2±6.21**	**22.27±7.71**

TABLE 2.3. Average precision values were obtained for a subset of the models tested at 100% recall. We note that a table with the performance of all the models at various recall values is presented in S22. Here the average is computed from the performance of four independent runs on the test sets. The uncertainties are 1σ standard deviations from the mean.

FIGURE 2.6. Grade-A lenses found in the DLS along with the rank (Section 2.7.2) assigned to them by GAN+MixMatch(MM) and GAN+Π-model(PI) models. All Grade-A lenses have a clear arc morphology and are located near a moderately massive galaxy or group, making them convincing lens candidates. Among these candidates, 212072337 and 432021600 have been spectroscopically confirmed to be true strong lens systems (Section 2.7.2.2).

on simulations. Arbitrarily increasing the size of the training data can evidently lead to significantly worse performance than using a smaller well-curated training set.

To summarize, we find that models trained with a semi-supervised learning approach using TrainingV2 and GAN-generated images along with all of our proposed list of data augmentations have high precision values at all recall values. In particular, among the

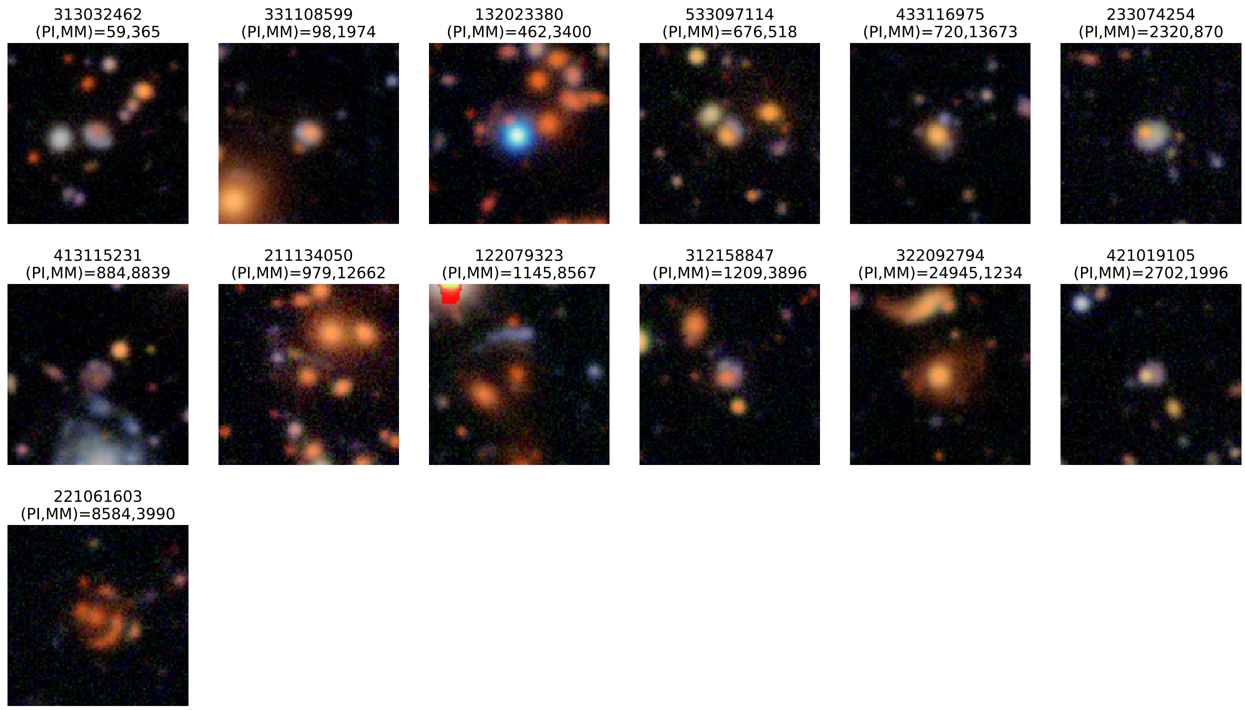

FIGURE 2.7. Grade-B lenses found along with the rank (Section 2.7.2) assigned to them by GAN+MixMatch(MM) and GAN+PiModel(PI) models. Targets in this category have either a tentative nebulous arc-like feature surrounding a massive galaxy, or have approximately linear extended morphology near an apparent galaxy group or cluster. It is hard to discern if these features correspond to lensed arcs or are caused by blending of multiple sources, hence the uncertain Grade-B classification.

models tested, the top two performing models are GAN+MixMatch and GAN+Π-model. In the following subsection, we turn to apply these models to the full set of DLS survey images (i.e., SurveyCatalog in Section 2.3.1)

2.7.2. Catalog of Lens candidates found. Having established which of our trained models perform best on our test set in terms of PR curves, we now turn to the key question of how many lenses are identified in the DLS and importantly, how much human inspection effort is required to find them.

We obtain a ∼97% and ∼86% precision at 50% recall (i.e., to find 50% lenses from our test set) for the GAN + MixMatch and GAN + Π-model respectively. On the other hand,

Rank threshold	Number of unique lenses investigated (G+MM,G+PI)	Number of Grade-A lenses (G+MM)	Number of Grade-A lenses (G+PI)	Total lenses Grade-A (both models)	Number of Grade-A lenses (SupervisedV2)	Number of Grade-A lenses (SupervisedV2+DA)	Number of Grade-A lenses (SupervisedV2+DA+GAN)
12	9, 9	1	1	1	0	0	0
25	19, 16	1	3	3	0	0	0
100	67, 56	2	3	3	0	1	2
800	513, 430	4	3	4	1	2	3
2800	1735, 1459	6	5	8	-	-	-
4000	2459, 2076	7	5	9	-	-	-

TABLE 2.4. Comparison of the number of Grade-A lenses found by different models tested. The predictions from the models are ranked such that the most likely predicted lens has rank=1. The rank threshold value sets the number of lenses that an investigator has to visually inspect. The left two columns show the chosen rank threshold and the number of unique lenses that it corresponds to (removing duplicates as described in Section 2.7.2). Our best performing models GAN+MixMatch (G+MM) and GAN+PiModel (G+PI) find 4 and 3 lensed candidates each among the top $\sim$500 unique images (top 800 ranks), and 7 lensed candidates each when the top $\sim$ 2300 images are investigated. Combining the results from both the models, we find 9 Grade-A candidates (shown in Figure 2.6). The right three columns show the number of lenses found from the SupervisedV2, SupervisedV2+Data Augmentation(DA) and SupervisedV2+DA+GAN. Although they find fewer ($\lesssim$ 50%) lens candidates than our best performing models, we can see that DA and GANs are able to boost the number of lenses found from 1 to 3 at a rank threshold of 800.

Rank threshold	Number of unique lenses investigated (G+MM,G+PI)	Number of Grade-B lenses (G+MM)	Number of Grade-B lenses (G+PI)	Total lenses Grade-B lenses (both models)	Total lenses Grade-A+B lenses (both models)
12	9, 9	0	0	0	1
25	19, 16	0	0	0	3
100	67, 56	0	2	2	5
800	513, 430	2	5	5	9
2800	1735, 1459	6	11	12	20
4000	2459, 2076	9	11	13	22

TABLE 2.5. Similar to Table 2.4 but for Grade-B lenses.

We note that the relative ranks which we use in this study will be unaffected under such scaling transformations.

One substantial caveat when looking at the top n predictions is that, due to the density of galaxies in the sky and our image selection method, the top predictions are not necessarily unique. For example, the top 25 predictions from the GAN+Π-Model contain 17 unique sources and 8 duplicates centered on different nearby objects (shown in Figure 2.13 in the Appendix). For the top 2800 predictions, the number of unique candidates is $\sim$ 1600 on average (i.e., $\sim$ 40% are repeated). Since this is a significant portion of the number of images and would increase human effort during labeling, we remove such repetitions based on their sky coordinates. Given our image size, we remove duplicates within a radius of 26 arcseconds of each object in the top n predictions.

The remaining images are then replaced with a larger field of view, ensuring that a given region of the sky needs to be visually inspected only once. We note that removing duplicates is strictly a post-processing step. Two of us (KVGC and TJ) visually inspected the lens candidates and classified them into confidence categories: Grade-A, Grade-B, Grade-C, and non-lenses. Grade-A indicates a high likelihood of being a strong lens system, on the basis of a clear arc morphology and/or coincidence with a moderately massive group of galaxies. Grade-B lenses generally have a nebulous arc-like feature surrounding a massive galaxy and/or have approximated linear extended arc morphology near a group or cluster of galaxies. It is uncertain if these features are from the lens or the effect of blending multiple sources. Grade-C lenses (not discussed in this paper) are the lowest-confidence candidates which typically show blended arc-like features likely arising from spiral arms, tidal features, or asymmetric diffuse light from the onset of mergers.

Figures 2.6 and 2.7 show the color composite images for the 9 Grade-A and 13 Grade-B lenses found from the survey upon visually inspecting $\sim$ 2500 unique candidates (the top 4000 by rank). Their sky coordinates are listed in Table 2.8 in the Appendix. Several of the Grade-A lenses appear to be compound lenses or part of a moderately massive group or cluster of galaxies. This is interesting since our training data consists of only galaxy-galaxy lenses. This is likely due to the addition of GAN-generated images to our training data, as the GAN-generated images (Figure 2.3) include irregularly shaped arcs.

2.7.2.1. *Human inspection effort.* We now examine how much human effort is required to find the 22 Grade-A and B lens candidates. To quantify the effort we consider the number of lenses found at different ranks, listed in Table 2.4. The rank threshold determines the number of unique images which must be visually inspected. Looking at the top 800 predictions from the GAN+MixMatch and GAN+Π-model (corresponding to 513 and 430 unique lens candidates respectively), we find 4 and 3 Grade-A lenses, and 2 and 5 Grade-B lenses respectively. This is several times ($\gtrsim 3\times$) higher sky density than has been found from the shallower ground-based DES survey, and smaller than the density found in COSMOS with HST, as expected. The number of lens candidates found increases to 9 Grade-A and 13 Grade-B candidates when the top 4000 candidates ($\sim$2500 unique images) are considered. This corresponds to $\sim$1 lens per deg^2 searched, which is $\gtrsim 10\times$ higher sky density of lenses compared to previous shallower ground-based surveys (as we discuss in Section 2.7.4).

In comparison, our supervised models (e.g., SupervisedV1, SupervisedV2) find $\lesssim$ 50% of these top lens candidates. They also have lower precision values (Table 2.3), with no compelling lenses found within the top 17 candidates inspected (whereas G+PI finds 3 within this threshold range). This again highlights the value of adding data augmentation and GAN images. The SupervisedV2+DA+GAN model finds 3 times more lenses than SupervisedV2 within the same threshold range. These results demonstrate the efficiency with which the models explored in this work can find strong lenses.

2.7.2.2. *Spectroscopic confirmation of two Grade-A lenses.* While image morphology can provide compelling evidence for strong gravitational lensing, spectroscopic redshifts are the standard to unambiguously establish the lensing nature of a system. We have obtained

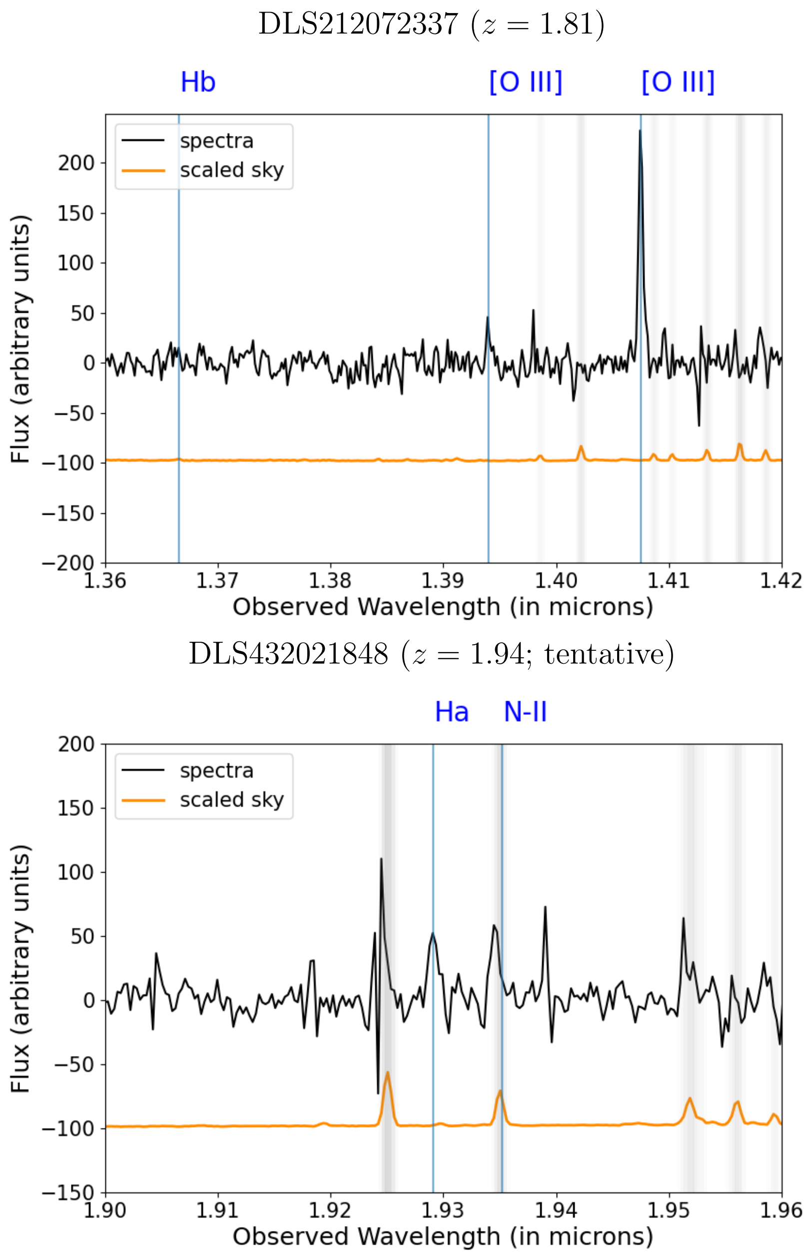

FIGURE 2.8. (Top): NIRES spectra of Grade-A lens DLS212072337 at a redshift of $z = 1.81$ with prominent [O III] emission lines marked in blue. (Bottom): NIRES spectra of DLS432021848 showing the single emission line detected at $1.93\mu m$ which we tentatively identify as Hα at $z = 1.94$. In both panels the scaled sky spectrum is shown in orange (offset by -100), with gray shading denoting regions affected by strong sky lines.

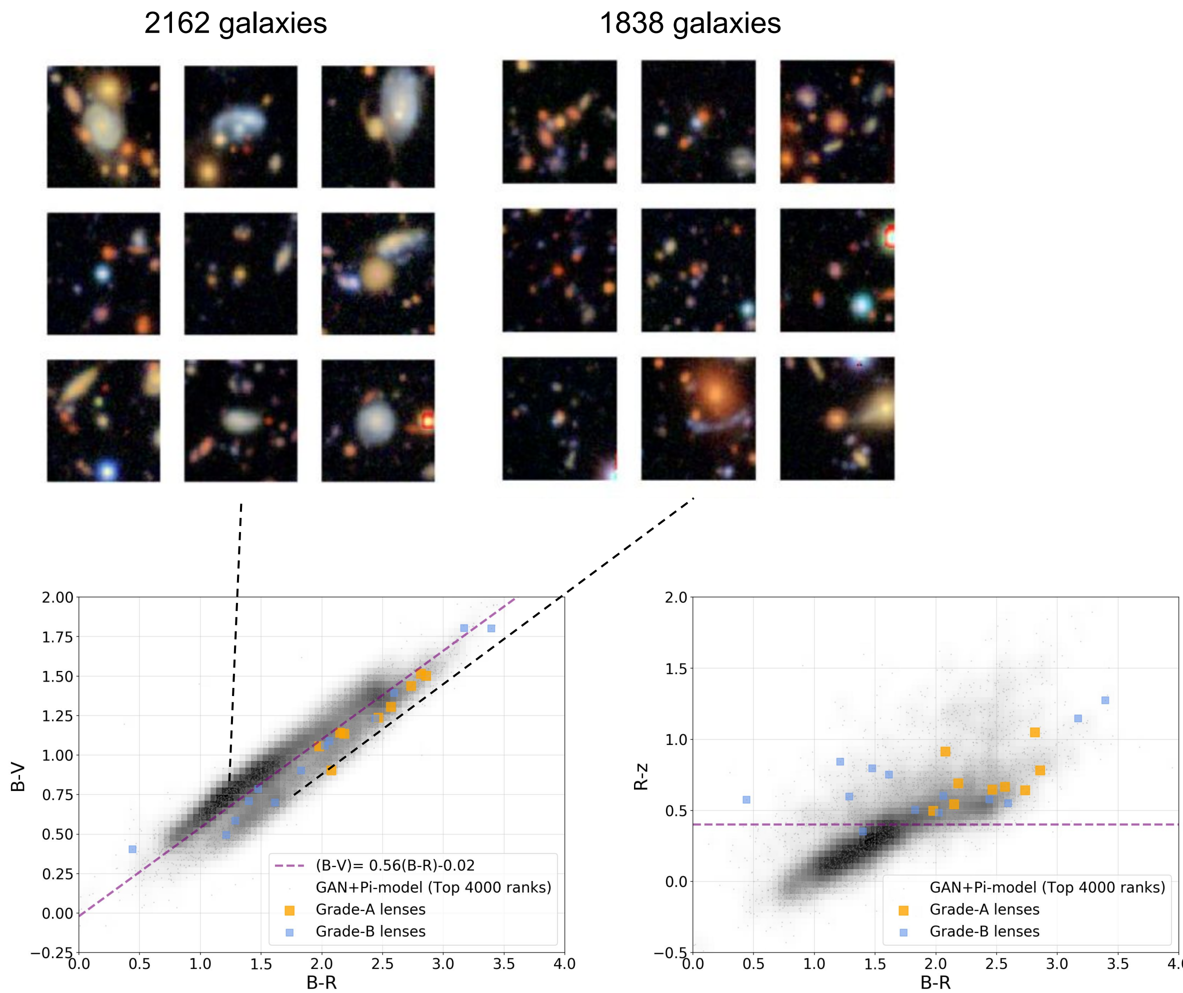

FIGURE 2.9. Distribution of the top 4000 lenses found by GAN+Π-model in color-color space. The left panel shows $B-R$ vs $B-V$ and the right panel shows $B-R$ vs $R-z$. The images above show examples of galaxies found in the two regions of the left panel separated by the purple line. Low-z galaxy candidates are clustered in the region above the trend line whereas all of the Grade-A lens candidates are below it. The right panel additionally shows that lens candidates are typically redder in $R-z$ colors ($\gtrsim 0.5$). A color selection based on the purple lines in each panel would yield higher precision in our lens candidate samples while retaining nearly all of the most probable lenses.

spectroscopy with Keck Observatory to confirm the lensing nature of two Grade-A systems presented herein: DLS212072337 and DLS432021848 (Figure 2.8). Observations of the arcs were conducted with NIRES (Wilson et al., 2004) on the Keck II telescope. Full details of the observations and data reduction are described in Tran et al. (2022a), along with spectroscopic redshifts for DLS212072337 (reported as AGEL091935+303156). We find a secure redshift of $z_{\rm arc} = 1.81$ for DLS212072337 from detection of Hα λ6564 and [O III] $\lambda\lambda$4960,5008 emission lines. The deflector galaxy is at a redshift of $z_{\rm def} = 0.43$, based on stellar absorption features from optical SDSS/BOSS spectra.

We observed DLS432021848 with NIRES on 12 January 2022 using the same methodology. We obtained 6 exposures of 300 seconds each. We detect a single emission line at $\lambda = 1.93\mu$m which we tentatively identify as either Hα at $z_{\rm arc} = 1.94$ or [O III] λ5008 at $z_{\rm arc} = 2.85$. However, we are unable to confirm the redshift with other strong lines, which fall in regions of poor atmospheric transmission at both potential redshifts. We find further support for the lensing nature of DLS432021848 from its morphology in follow-up HST imaging (discussed in Section 2.7.4), which shows clear kurtosis and evidence of multiple lensed images. Thus we are reasonably confident that this is indeed a strong lensing system on the basis of high-resolution imaging, despite the limited spectroscopic information. Together with DLS212072337, these results give additional confidence in the sample of lens candidates presented in this paper and demonstrate that our methods are successful.

We note that redshifts are known for two additional Grade-A candidate deflectors (DLS212148326, DLS421095124) from archival data. DLS212148326 is at $z_{\rm def} = 0.424$ from SDSS/BOSS spectra, while DLS421095124 is part of a massive galaxy cluster spectroscopically confirmed at $z_{\rm def} = 0.680$ (Wittman et al., 2003, 2006, reported as DLSCL J1055.2-0503). These redshifts are promising, as the distances and approximate masses are consistent with the deflection angles implied by the strong lensing interpretation of these images.

2.7.2.3. *Distribution of lensed candidates in color-color space.* The analysis and model performance described thus far in the paper is based on a source selection using an intentionally simple R band magnitude cut and SExtractor flags (Section 2.3.1). We have demonstrated in the above sections that such cuts are sufficient to search for lensed candidates in the DLS. However, more sophisticated selections can increase the efficiency of lens searches. Here we briefly consider how color selection can provide higher-purity samples.

In Figure 2.9 we show the distribution of Grade-A and B lenses from Section 2.7.2 in various color-color spaces, along with the top 4000 ranked images from the GAN-II-model as an example. These colors generally correspond to the central (candidate deflector) galaxy. The top lens candidates are not distributed uniformly, and we demonstrate two color-color selections where the top candidates are clustered: $(B - V) < 0.56(B - R) - 0.02$ (purple line in left panel), and $R - z \gtrsim 0.4$ (right panel). Such simple color cuts can retain all Grade-A lenses while removing the majority of false positives, thereby reducing the required human inspection effort. Physically, these colors are indicative of 4000 Å breaks at redshifts $z \gtrsim 0.25$ (i.e. in the V or R band) whereas lower-z galaxies are less likely to act as strong lenses.

The distribution of lens candidates in color space suggests that the precision of our models can be further improved by adopting color criteria as a pre- or post-processing step, with minimal loss of the best candidates. Using photometric redshift and mass estimates is a similar and potentially even more promising method (Schmidt & Thorman, 2013) although it is beyond the scope of this paper. Alternatively, a state-of-the-art automated means to address this would be by using self-similarity based approaches (e.g., Stein et al., 2021), wherein a CNN further classifies the lens probabilities based on their similarity with each other.

2.7.3. Lensing signatures identified by the models. We now examine which features of the lens candidate images are most relevant for the model predictions. Deep neural networks (such as ResNetV2 used in this work) are often considered as "black boxes" with all input information collapsed to a simple prediction for the user to interpret. Having only a single output, it is impossible to discern which distinguishing features of a gravitational

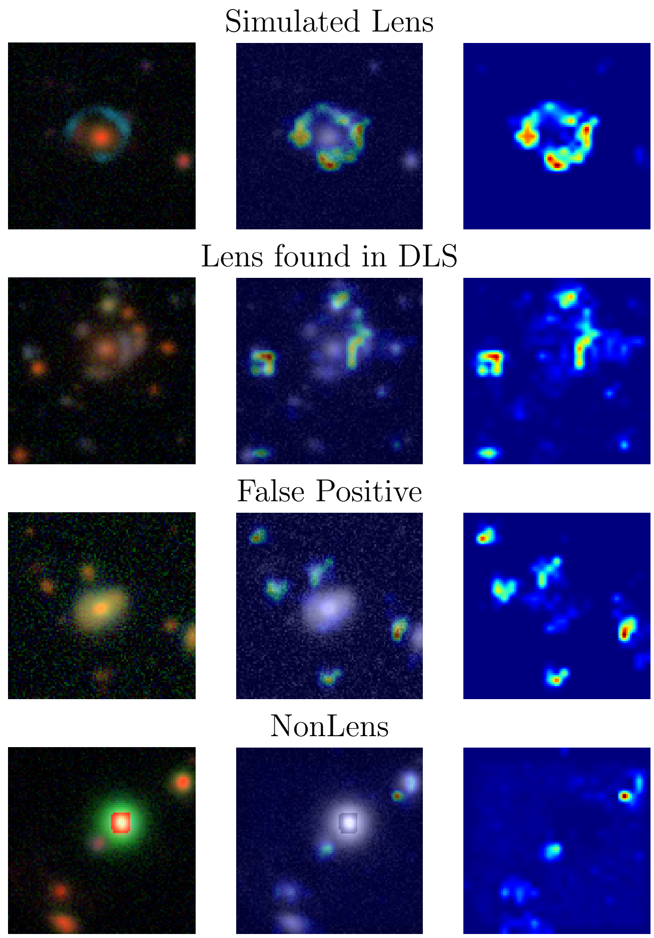

FIGURE 2.10. Grad-CAM++ heatmaps for an example simulated lens, two Grade-A lenses, two false positive lenses, and a NonLens. The left column shows the color composite image obtained from HumVI and passed to the model. The right column shows the Gradcam++ heatmaps. The red and green shading indicates regions of high and moderate importance to the model, respectively, whereas blue represents low importance. The middle column shows the heatmaps superimposed on input images for visualization purposes. For the simulated lens, we can clearly see that the entire lensed arc region is taken into consideration. For the Grade-A lens candidates found in DLS, we also find that the lensed arc features are considered important by the model, despite a range of lensing morphologies and colors. This suggests that models have indeed successfully generalized to the survey data. Notably, the massive deflector (i.e., the luminous red galaxy) causing the lensing effect is not highlighted in the simulated or candidate lens systems. Additional objects in the field are also highlighted in heatmaps for the Grade-A lenses, which is also apparent in the False Positive and NonLens examples. In the case of the False Positives, the highlighted object distributions resemble an “Einstein cross” lens configuration. Heatmaps for all the Grade-A lenses are provided in Figure 2.15 in the appendix.

lens are actually being identified and considered by the models. Fortunately, in the past few years, there have been a variety of methods proposed to alleviate this such as occlusion methods, Guided Backprop (Springenberg et al., 2015), CAM (Zhou et al., 2016), Grad-CAM (Selvaraju et al., 2017), Grad-CAM++ (Chattopadhay et al., 2018), and DeepSHAP (Fernando et al., 2019).

Gradient-based interpretation methods (e.g., Grad-CAM++) effectively compute gradients on intermediate feature maps of the network to determine the importance of a feature. These gradient maps can then be overlaid on top of the original input image, in order to assess which image regions are contributing most to the predicted output from the classifier. These methods are not without drawbacks (e.g., Adebayo et al., 2018) but can provide valuable insight. Here we use Grad-CAM++ to analyze some of our trained models.

Figure 2.10 shows Grad-CAM++ heatmaps obtained for a few illustrative examples. We consider a simulated lens from the training data, real Grade-A lenses from the survey, false positive images (i.e., images which are classified as lenses but show no visual evidence of lensing), and a non-lens. In the case of the simulated lens, it is clear that the model is indeed making its prediction based on the lensed arc features. For the Grade-A lenses, the model does indeed discern the lensed arcs, but there are additional unrelated regions within the images that also influence its decision. Curiously, the central massive deflector galaxy is not highlighted in these cases. In the case of the false positives, the model encouragingly is not misled by the extended central galaxies, but rather the heatmap highlights multiple sources of similar color which surround the central galaxy. For example in the spiral galaxy false-positive image, it is clear that the model picks up on the three nearby red objects. The location and color of these nearby objects is indeed similar to plausible multiple-image lensing configurations. It thus appears that the model has successfully learned to identify the astrophysical signatures of strong lensing.

2.7.3.1. *Finding red arcs.* As discussed in Section 2.5.2, our Lens dataset used for training only consists of lensed arcs with blue optical colors. However, it is encouraging that the models have also identified red arcs such as the system DLS212148326 (Figure 2.6). The network may be learning to identify red arcs through color augmentations (Figure 2.2). Although red-lensed arcs are known to exist, presumably a training dataset consisting of only blue arcs is not ideal to robustly search for and quantify them. It could be the case that adding more augmentations or fine-tuning existing ones might suffice to search for arcs of various colors. Alternatively, a broader range of arc colors could be used in the simulated training set, or a separate classifier could be constructed from a training set of red arcs. Given our adopted training set, we consider the number of red-lensed arcs found from this work to be a lower limit (relative to the blue arcs). Additionally, there are likely many fainter blue or red arcs which our training set does not represent, although the detection of fainter objects is naturally more challenging.

2.7.4. Implications for future large-area sky surveys: sensitivity and angular resolution. The next generation of wide-area sky surveys is expected to uncover $\gtrsim 10^5$ strong lens systems (e.g., Oguri & Marshall, 2010; Collett, 2015). Here we consider the gain in lens detection with survey depth and angular resolution based on our DLS sample from Section 2.7.2. We compare the sky density of detected lens candidates with two other illustrative examples of CNN-based searches in Table 2.6. In our DLS search, we find $\sim$0.5 Grade-A lenses per square degree (or $\sim$1 Grade-A+B lenses per square degree). This is considerably larger than found in shallower surveys such as SDSS and DES, which have uncovered $\sim$0.1 lenses per square degree (in regions far from the galactic plane). While these surveys have a comparable seeing-limited resolution, sharper image quality enables more lenses to be found. An example is the search of COSMOS HST imaging by Pourrahmani et al. (2018) using a CNN approach, which found 13 Grade-A candidates and 70 Grade-A+B candidates in the 2 square degree field (i.e., $\sim$35 per square degree). Therefore, we see that the sky density of detectable

Survey	Lenses found per sq.deg	5σ point source detection (r/R/F814W-band magnitude)	FWHM	References
DES/DECaLS	~ 0.1	23.6 (r)	$0\farcs98$	J19
DLS	1	26.7 (R)	$0\farcs9$	This work
COSMOS	~ 35	27.2 (F814W)	$0\farcs07$	P18,K07

TABLE 2.6. Number of lenses found using machine learning methods per square degree of sky in different surveys, along with the 5σ point source detection depth and median angular resolution (given as the FWHM: full-width at half maximum). We note that CNN and grading methods employed to find lenses in each survey are different; the density of lenses should thus be treated as an approximate comparison. References are as follows. J19: Jacobs et al. (2019a), P18: Pourrahmani et al. (2018), K07: Koekemoer et al. (2007).

strong lens systems increases by $\sim$10 times when going from shallower ground-based surveys (e.g., SDSS) to the DLS, and by another factor of $\gtrsim$10 when the angular resolution is improved by an order of magnitude with space-based HST imaging at modest depth. These results generally support the predictions of large lens samples which will become detectable with near-future surveys planned with the Rubin (LSST Science Collaboration et al., 2009), Roman (Spergel et al., 2015), and Euclid (Laureijs et al., 2011) observatories.

To visually illustrate the detection of lenses at different depths and angular resolutions, Figure 2.11 compares DECaLS, DLS, and HST imaging[3] for the Grade-A lens candidate DLS432021848 found in this work. A blue arc is clearly visible in the DLS image and shows typical lensing morphology in the high-resolution HST image. However, the arc is only marginally visible in shallower DECaLS imaging. Indeed, most (if not all) of the Grade-A lens candidates found from this work would be difficult to detect in shallower imaging surveys (e.g., DECaLS; hence for example they are not included in the catalog of Huang et al. 2020).

[3]The HST image was secured as part of program HST-GO-16773 targeting lens candidates identified primarily in DES and DECaLS imaging (Tran et al., 2022a). In brief, the HST image in Figure 2.11 was taken with WFC3-IR in the F140W filter with $\sim$30 minutes of exposure time (<1 orbit), and reduced using standard procedures. Details of the HST program will be described in a forthcoming paper.

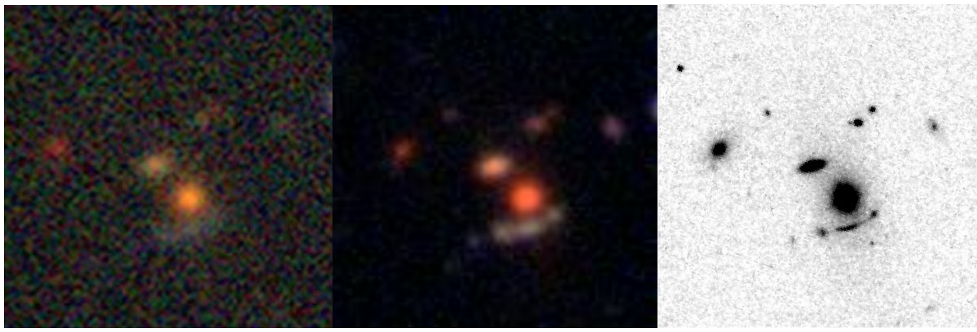

Figure 2.11. Comparison of the image quality from different observations of the lens system DLS432021848, which shows a prominent blue arc in DLS imaging (below center of images; all panels show the same field of view). *Left*: The arc is apparent but not well detected in DECaLS imaging, which has modest sensitivity. This image would likely be flagged in a low-confidence category and indeed was not identified in previous lens searches (e.g., Huang et al., 2020). *Middle*: DLS image of the target showing a prominent blue arc-like feature below the red deflector galaxy, characteristic of a gravitational lens system. The increased sensitivity of DLS compared to DECaLS imaging (Table 2.6) enables clear arc detection. *Right*: Near-infrared image of the same target observed with HST, with a diffraction-limited angular resolution approximately 6 times sharper than DLS or DECaLS images. The HST image reveals the lensed arc morphology at a high signal-to-noise ratio. This demonstrates the capabilities of a ground-based telescope at good depth (e.g., DLS), and a diffraction-limited space-based telescope with moderate exposure time (e.g., HST).

Given the detectability of many lens systems with upcoming surveys, it is clear that machine learning approaches (such as those we have explored here) will be vitally important for the efficient selection of large samples. We have also demonstrated the feasibility of spectroscopically following up on these moderately faint arc systems (Section 2.7.2.2), which will be vital for confirmation and subsequent analyses.

2.8. Conclusions

In this paper, we have evaluated the performance of different CNN learning approaches and data augmentations on their ability to efficiently find gravitational lens candidates in the Deep Lens Survey. We make use of the deep learning architecture ResNet for our experiments, along with a training dataset consisting of simulated Lenses and survey image NonLenses. We demonstrate that by using these state-of-the-art semi-supervised learning approaches, we can greatly reduce the human effort required to find lensed candidates from a survey. We summarize our key results below.

(1) Among 17 variants of learning approaches tested in this work, we find that our best performing models (i.e., those which have high precision and minimize false positives during human inspection) are GAN+MixMatch and GAN+$\Pi-$model. They have a precision of $\sim 86\%$ and $\sim 97\%$ at 50% recall and, $\sim 22\%$ and $\sim 8\%$ at 100% recall respectively. In comparison, our supervised models have a precision of $\sim 3\%$ at 100% recall. This increase in the performance of the best models can be attributed largely to three factors. (1) They leverage data augmentation (Table 2.1) during training, which helps them to generalize better. (2) The datasets used to train these models to contain simulated Lenses as well as GAN-generated images (Section 2.4), which serves as an additional form of data augmentation. (3) Both of these top models employ a semi-supervised learning approach (MixMatch, Π-model) which enables our methods to adapt to distributional shift (Section 2.4.2). These results indicate that data augmentation, GANs, and semi-supervised learning are highly effective approaches for building an efficient lens classifier.

(2) We investigated the Grad-CAM++ feature maps (Section 2.7.3) used by our best performing models to make their predictions, finding that they indeed are influenced mostly by lensed arc regions and are generally not misled by other galaxies/artifacts (e.g., diffraction spikes) in the images. This supplements our results presented above that salient information regarding the arcs needed for classification has been successfully learned by the models through our methods. This is encouraging for future

lens searches, since simulated Lenses used in this work are generated without relying on photometric data of the deflector galaxy (Section 2.5), making it simpler to automate the task of generating a training dataset.

(3) Applying the GAN+MixMatch and GAN+Π-model to the entire DLS survey, and visually inspecting the top ~ 2500 lens candidates, we find 9 Grade-A and 13 Grade-B lensed candidates (22 in total). 3 out of the 9 Grade-A candidates are found within the top 17 ranked images. The number of lenses found in the DLS corresponds to $\sim 10\times$ higher sky density of lenses per deg^2 compared to the shallower DES/DECaLS survey imaging and supports predictions that vast numbers of lens systems ($\gtrsim 10^5$) will be detectable in the upcoming generation of sky surveys. We further confirmed the lensed nature of 2 Grade-A candidates with spectroscopy and high-resolution imaging, demonstrating that our methods are successful.

We have generally explored methods intended to find as many lenses as possible while minimizing human inspection effort. While there are likely additional detectable lenses beyond those we have identified, it is encouraging that our models have been able to identify lenses that are not represented in the training set. In particular, our training set focused on blue lensed arcs, while our models also find red arc candidates such as DLS212072337 (Section 2.7.3.1), although at a lower rank compared to the bluer lenses. Additional augmentation methods and/or training datasets may be able to provide further improvement for diverse lens system properties. Another straightforward improvement to our lens search efficiency is to include simple cuts in color-color space as demonstrated in Section 2.7.2.3. Such cuts can help increase the model precision by excluding sources that are not likely to act as strong lenses based on their color and magnitude (which is physically related to their mass and distance). Since our sample is agnostic to color information, our results are well-suited for assessing the color space distribution of the best lens candidates.

The scope of our models is currently limited to the DLS. However, our methodology can be adapted for other data sets, and we note that the DLS fields overlap with wide-area surveys such as DECaLS and SDSS. Exploring ways to translate these models across surveys would be greatly beneficial. Finally, confirming the lensing nature of new candidates either through spectroscopy (Section 2.7.2.2) or via arc morphology (Section 2.7.4) is essential for a variety of investigations, including probes of galaxy evolution and cosmology. We have demonstrated the feasibility of confirming moderately faint arcs in our sample. Accomplishing confirmation for the thousands of lenses that will be discovered in forthcoming surveys (such as with Rubin/LSST, Roman, and Euclid) will aid in our understanding of the formation and evolution of galaxies and the contents of the Universe.

Appendix: Model performance and final lens sample

In this appendix, we provide some additional details of the model performance and the top lens candidates identified in this work.

Figure 2.12 shows the distribution of model scores across the different DLS fields (F1 to F5), demonstrating similar performance in each field. This is generally expected given the similar image quality across the DLS survey. Importantly it shows that our use of labeled training data from only F1 does not substantially affect the model performance in the other fields.

Table 2.7 lists the results from our ablation study discussed in Section 2.7.1.1. We show the precision across recall rates from 50-100%. The performance differences are generally similar across all recall rates. Color augmentation, JPEG quality, and GAN images appear to most prominently improving the model performance (i.e., the models perform significantly worse when these augmentations are removed).

We show the top 25 predicted lens candidates from the GAN+Π-model and GAN+MixMatch models in Figures 2.13 and 2.14, respectively. These include several of our top lens candidates based on human inspection (see Figures 2.6) and 2.7), but many do not show obvious signs

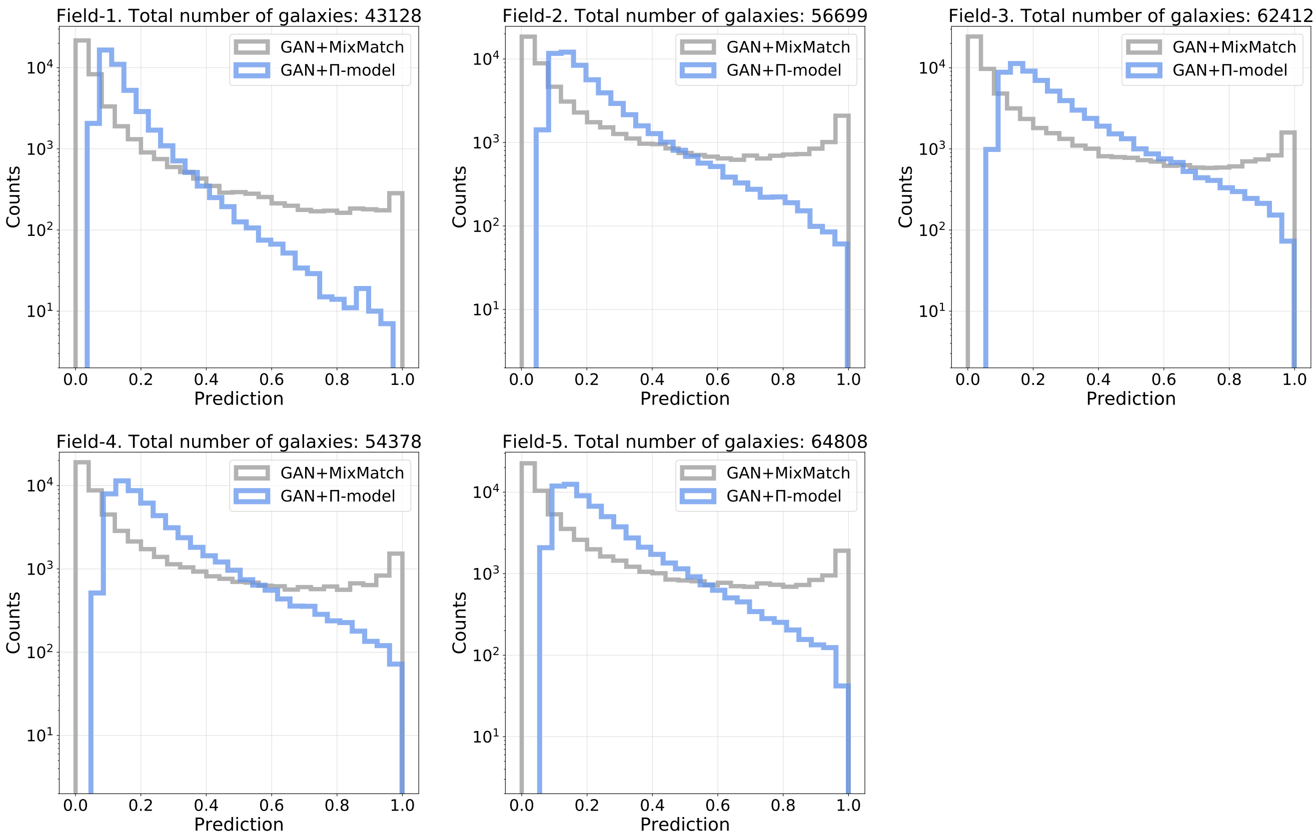

FIGURE 2.12. Histogram of the scores obtained by the GAN+MixMatch and GAN+Π-model in the five independent DLS Fields F1 through F5. As discussed in Section 2.5, the training set used to train our models (TrainingV2) contains human labeled NonLenses which were randomly sampled only from Field F1. But as we clearly see, the distribution of scores (and performance of the models as a result) is independent of the field chosen.

of strong lensing. There are several duplicate images at slightly different sky positions as discussed in the main text. In Figure 2.15 we include the GradCAM++ heatmaps obtained for all the Grade-A candidates (analogous to the example subsets shown in Figure 2.10). These heatmaps were generated using our best performing models: GAN+MixMatch or GAN+Π-model (discussed in Section 2.7.3). Finally, we list the sky coordinates of all Grade-A and Grade-B lenses in Table 2.8.

TABLE 2.7. Ablation performance for 50–100% recall rates (in steps of 10%) for the GAN+Supervised model using TrainingV2. The first row of each recall rate shows the baseline precision value obtained from the model on the test sets (TestV1, TestV2) when none of the augmentations are removed. In the subsequent rows, we report the precision obtained when the model was trained without the specified augmentation. For example, the baseline model at 50% recall has a precision of 80.19% for TestV2 and decreases to 55.77% when GAN images are removed during training. The difference in the obtained precision values are quoted in the last two columns. Augmentations which improve model performance (i.e., improve precision when included and decrease decrease precision when removed) are shown in red, while those which decrease model performance are shown in blue. Overall, the models perform worse when color augmentations, JPEG quality and GANs are not included, indicating that these augmentations are important for optimal performance. The errors quoted here are 1σ.

Augmentation removed	TestV1 Precision(%)	TestV2 Precision(%)	TestV1 baseline difference(%)	TestV2 baseline difference(%)
		Performance at 50% recall rate		
None	84.95 ± 8.70	80.19 ± 17.08	-	-
GAN	65.46 ± 15.00	55.77 ± 17.43	-19.49	-24.42
RGB shuffle	44.8 ± 24.32	35.45 ± 21.75	-40.15	-44.74
JPEG quality	65.30 ± 13.84	56.84 ± 14.06	-19.65	-23.35
Rot90	91.81 ± 4.41	89.96 ± 6.33	+6.86	+9.77
Translations	89.47 ± 10.34	84.59 ± 13.04	+4.52	+4.4
Horizontal flips	84.63 ± 10.85	75.74 ± 13.96	-0.32	-4.45
Color augmentation	71.88 ± 17.35	68.23 ± 15.2	-13.07	-11.96
		Performance at 60% recall rate		
None	79.88 ± 6.30	78.03 ± 12.45	-	-
GAN	55.54 ± 13.13	44.09 ± 12.94	- 24.34	-33.94
RGB shuffle	34.25 ± 24.49	26.14 ± 15.9	- 45.63	-51.89
JPEG quality	52.75 ± 24.68	51.69 ± 17.91	- 27.13	-26.34
Rot90	84.94 ± 7.04	78.99 ± 5.23	+5.06	+ 0.96

Continued on next page

Table 2.7

Augmentation removed	TestV1 Precision(%)	TestV2 Precision(%)	TestV1 baseline difference(%)	TestV2 baseline difference(%)
Translations	74.95 ± 16.11	72.26 ± 24.65	-4.93	-5.77
Horizontal flips	71.15 ± 11.46	60.64 ± 12.27	-8.73	-17.39
Color augmentation	63.24 ± 15.67	50.91 ± 15.53	-16.64	-27.12
Performance at 70% recall rate				
None	69.42 ± 4.60	68.05 ± 12.96	-	-
GAN	46.45 ± 14.49	37.24 ± 12.85	-22.97	-30.81
RGB shuffle	26.68 ± 17.28	18.75 ± 11.78	-42.74	-49.3
JPEG quality	40.38 ± 18.64	40.52 ± 21.74	- 29.04	-27.53
Rot90	79.97 ± 12.15	73.69 ± 7.11	+10.55	+5.64
Translations	67.69 ± 12.67	59.31 ± 22.34	-1.73	-8.74
Horizontal flips	67.36 ± 11.95	55.94 ± 14.71	-2.06	-12.11
Color augmentation	52.79 ± 14.92	45.93 ± 13.97	-16.63	-22.12
Performance at 80% recall rate				
None	54.35 ± 4.57	40.98 ± 17.12	-	-
GAN	33.02 ± 10.26	24.37 ± 10.50	-21.33	-16.61
RGB shuffle	15.34 ± 7.59	11.75 ± 4.90	-39.01	-29.23
JPEG quality	21.33 ± 8.29	15.25 ± 3.41	-33.02	-25.73
Rot90	62.63 ± 12.43	52.2 ± 13.65	+8.28	+11.22
Translations	52.17 ± 12.92	35.09 ± 16.64	-2.18	-5.89
Horizontal flips	60.54 ± 9.78	36.46 ± 9.03	+6.19	-4.52
Color augmentation	39.47 ± 13.24	33.84 ± 10.92	-14.88	-7.14
Performance at 90% recall rate				
None	34.12 ± 7.47	16.33 ± 8.661	-	-
GAN	22.60 ± 4.78	10.99 ± 4.25	-11.52	-5.34

Continued on next page

Table 2.7

Augmentation removed	TestV1 Precision(%)	TestV2 Precision(%)	TestV1 baseline difference(%)	TestV2 baseline difference(%)
RGB shuffle	10.37 ± 4.68	5.88 ± 2.61	-23.75	-10.45
JPEG quality	15.15 ± 6.37	6.18 ± 2.41	-18.97	-10.15
Rot90	40.72 ± 12.74	21.79 ± 3.08	+6.6	+5.46
Translations	32.88 ± 3.35	16.14 ± 6.72	-1.24	-0.19
Horizontal flips	30.81 ± 20.44	14.72 ± 6.59	-3.31	-1.61
Color augmentation	23.86 ± 8.13	14.25 ± 8.45	-10.26	-2.08
Performance at 100% recall rate				
None	8.25 ± 2.85	6.05 ± 2.69	-	-
GAN	8.13 ± 2.49	5.79 ± 3.93	-0.12	-0.26
RGB shuffle	6.43 ± 0.97	3.84 ± 1.02	-1.82	-2.21
JPEG quality	5.88 ± 0.21	4.14 ± 2.09	-2.37	-1.91
Rot90	14.76 ± 8.42	13.00 ± 5.16	+6.51	+6.95
Translations	10.83 ± 2.46	7.37 ± 3.66	+2.58	+1.32
Horizontal flips	15.74 ± 3.06	8.86 ± 1.84	+7.49	+2.81
Color augmentation	11.42 ± 7.25	6.84 ± 4.12	+3.17	+0.79

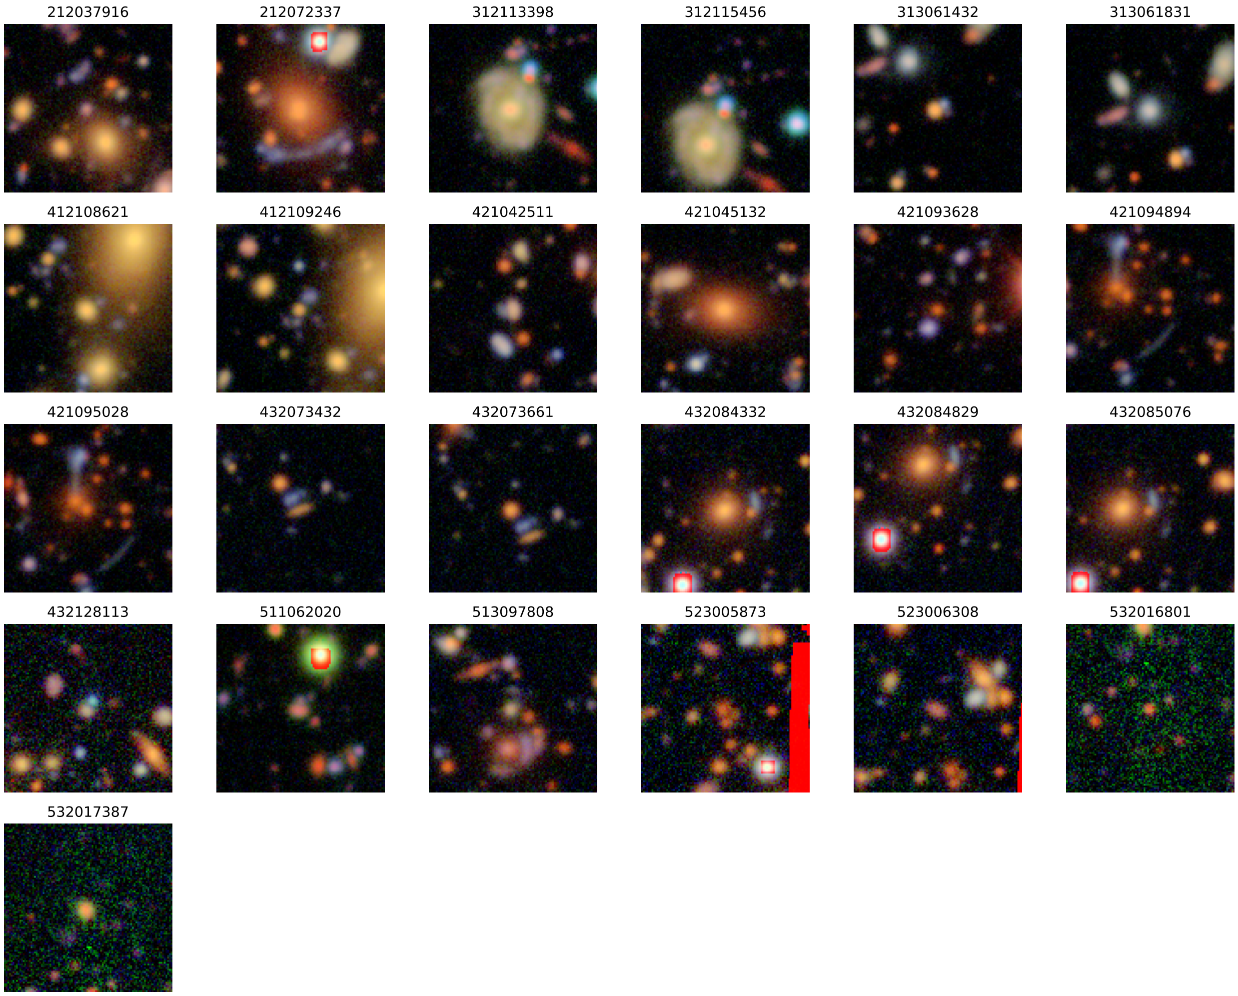

FIGURE 2.13. DLS images of the top 25 predictions from GAN+Π-model. Several show clear evidence of strong lensing, while other images appear to be false positives. We note that many images are duplicates (at overlapping regions of the sky), which we remove before visual inspection.

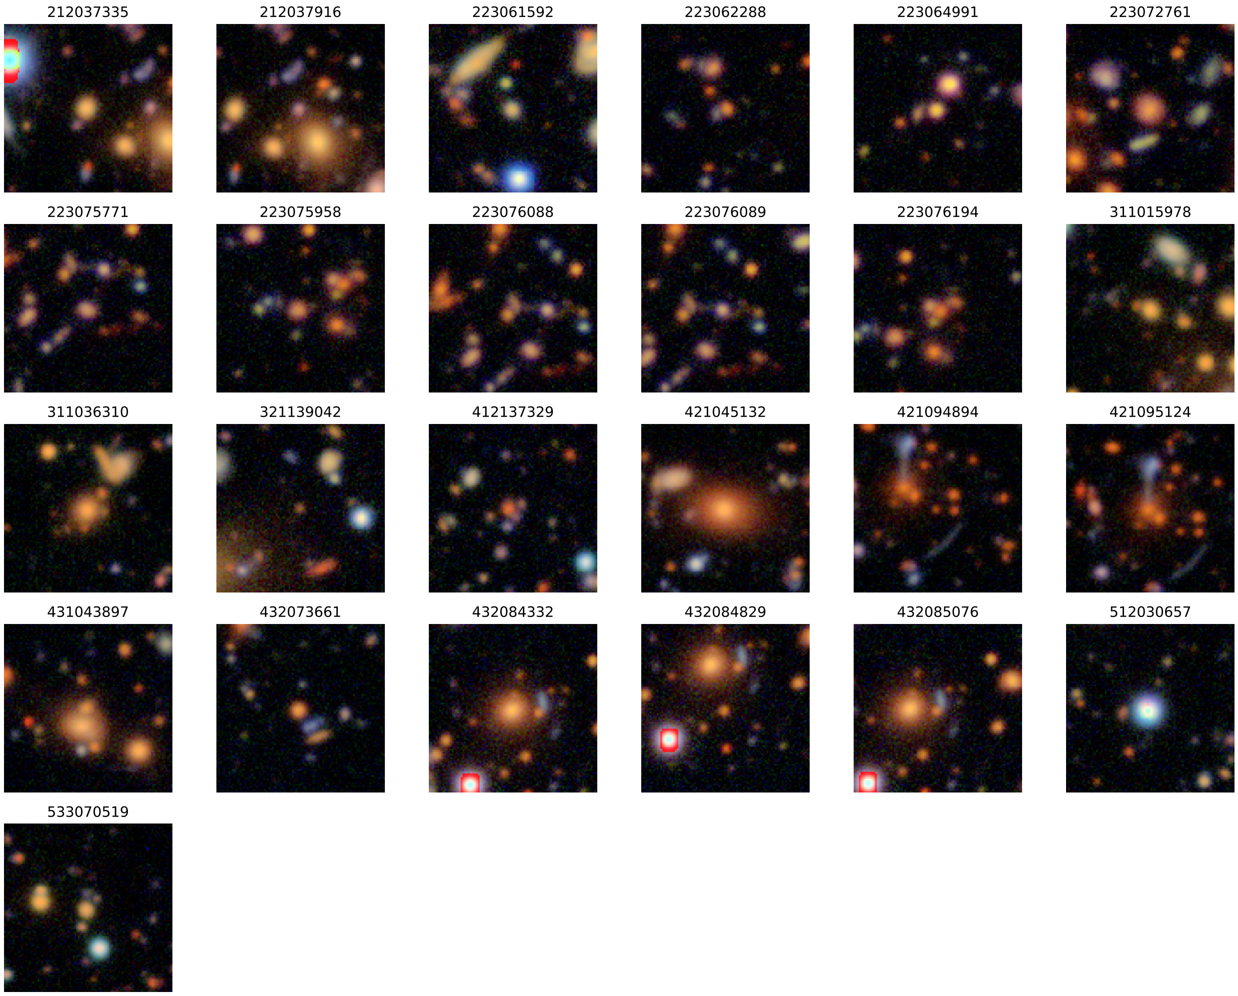

FIGURE 2.14. Equivalent to Figure 2.13, showing the top 25 predictions from GAN+Mimatch.

object ID	RA	DEC	Field	Rank (GAN+MixMatch)	Rank (GAN+Π-model)
			Grade-A candidates		
421095124	163.792076	-5.070373	F4	2	12
513097468	209.340092	-10.244328	F5	38	13
212072337	139.896040	30.532355	F2	181	21
322054393	79.839914	-48.949647	F3	733	8326
432021600	162.750073	-5.941902	F4	1262	12424
431010921	163.364259	-5.789092	F4	1279	19826
512037933	209.677055	-10.687652	F5	7461	2068
421117552	163.897903	-5.054885	F4	4799	2768
212148326	139.512033	30.953524	F2	3579	23223
			Grade-B candidates		
313032462	78.742878	-48.149829	F3	365	59
331108599	81.300608	-49.432676	F3	1974	98
132023380	13.551513	11.794606	F1	3400	462
533097114	209.328083	-11.993324	F5	518	676
433116975	162.551767	-5.697394	F4	13673	720
233074254	139.046712	29.298535	F2	870	2320
413115231	162.585545	-4.498548	F4	8839	884
211134050	140.304878	30.471131	F2	12662	979
122079323	13.182525	12.323637	F1	8567	1145
312158847	80.455801	-48.489660	F3	3896	1209
322092794	80.115321	-49.246309	F3	1234	24945
421019105	163.411890	-4.870280	F4	1996	2702
221061603	140.662872	29.846367	F2	3990	8584

TABLE 2.8. Grade-A and Grade-B Lens candidates found from this work with their object ID, RA and DEC coordinates, DLS field (F1 through F5), and their corresponding ranks from GAN+MixMatch and GAN+Π-models. The rank is obtained by passing all the survey images (281,425 objects in total; Section 2.3) through the models and sorting them based on their prediction scores. High-confidence Lens candidates have lower ranks and high prediction scores. For example, the Grade-A lens candidate DLS212072337 whose lensing nature has been spectroscopically confirmed (Section 2.7.2.2) has a rank of 21 from the GAN+Π-model and a prediction score of $\simeq 1$. The ranks quoted here represent an upper bound on the number of images an investigator has to look at to find the lens candidate, as they do not account for duplicated sky regions which we remove before visual inspection (as discussed in Section 2.7.2), reducing the number of unique lens candidates investigated.

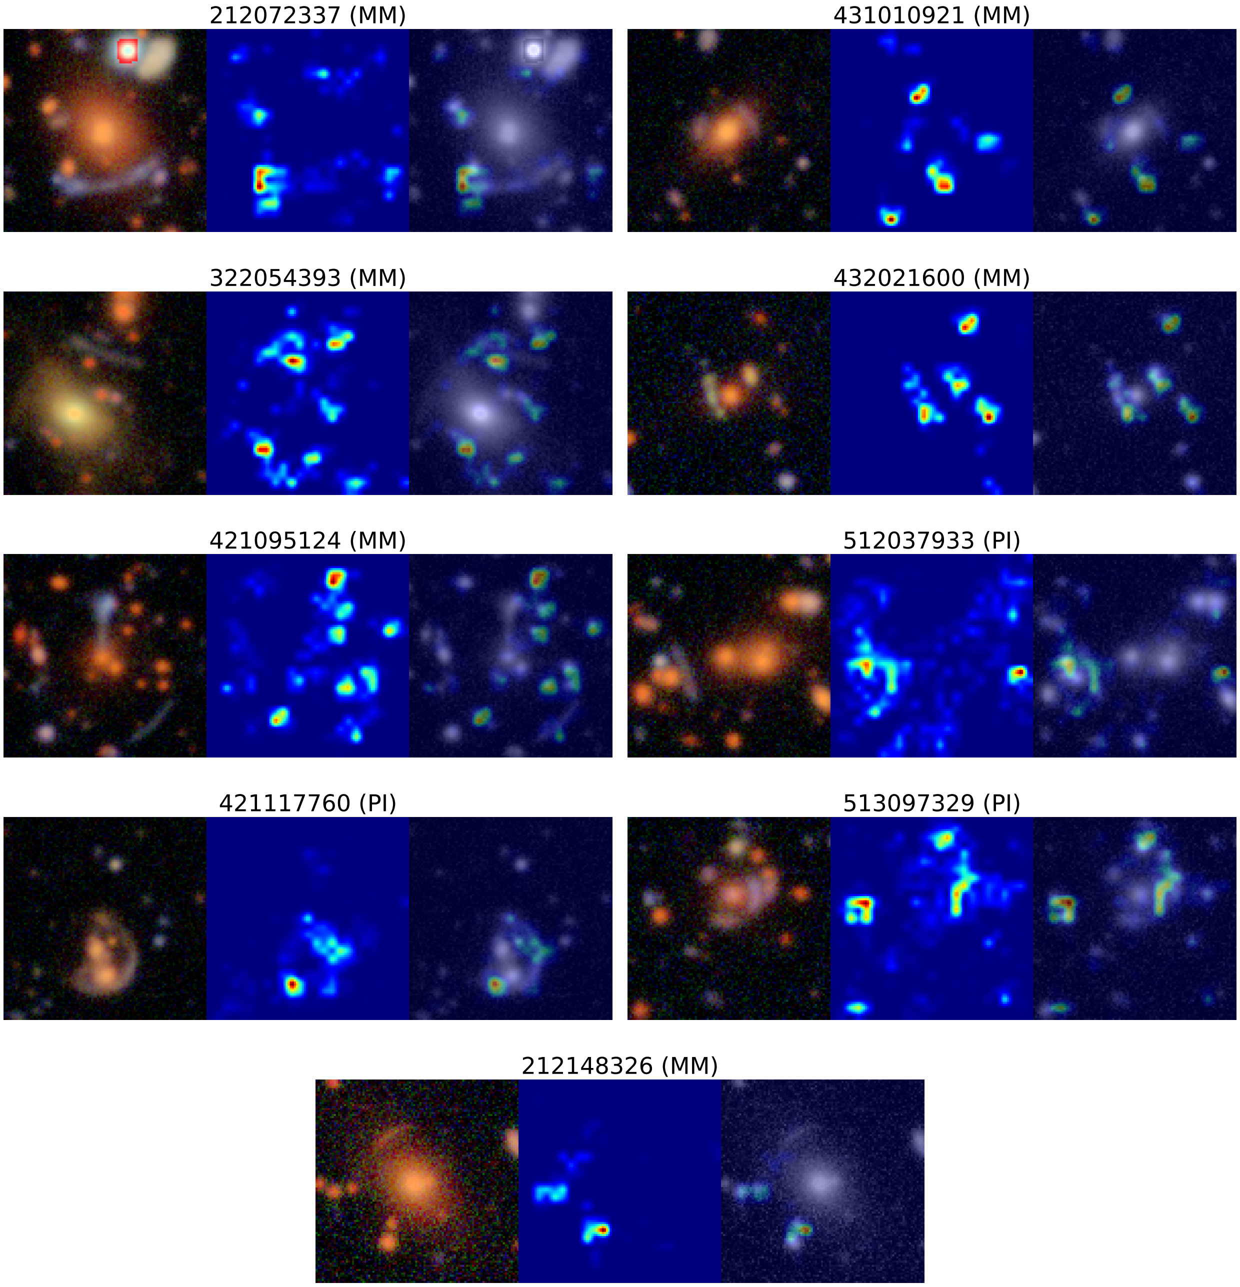

FIGURE 2.15. GradCAM++ heatmaps for all Grade-A lenses, equivalent to Figure 2.10. Each image is labeled with its object ID, and the model corresponding to the heatmaps (MM = GAN+MixMatch, PI = GAN+Π-model).

CHAPTER 3

Resolved velocity profiles of galactic winds at Cosmic Noon

Published as Vasan G. C. et al. (2023) in The Astrophysical Journal, Volume 959, Number 2, 2023.

3.1. Abstract

We study the kinematics of the interstellar medium (ISM) viewed "down the barrel" in 20 gravitationally lensed galaxies during Cosmic Noon ($z = 1.5 - 3.5$). We use moderate-resolution spectra ($R \sim 4000$) from Keck/ESI and Magellan/MagE to spectrally resolve the ISM absorption in these galaxies into ∼10 independent elements and use double Gaussian fits to quantify the velocity structure of the gas. We find that the bulk motion of gas in this galaxy sample is outflowing, with average velocity centroid $\langle v_{cent} \rangle = -141\ \mathrm{km\,s^{-1}}$ ($\pm 111\ \mathrm{km\,s^{-1}}$ scatter) measured with respect to the systemic redshift. 16 out of the 20 galaxies exhibit a clear positive skewness, with a blueshifted tail extending to $\sim -500\ \mathrm{km\,s^{-1}}$. We examine scaling relations in outflow velocities with galaxy stellar mass and star formation rate (SFR), finding correlations consistent with a momentum-driven wind scenario. Our measured outflow velocities are also comparable to those reported for FIRE-2 and TNG50 cosmological simulations at similar redshift and galaxy properties. We also consider implications for interpreting results from lower-resolution spectra. We demonstrate that while velocity centroids are accurately recovered, the skewness, velocity width, and probes of high velocity gas (e.g., v_{95}) are subject to large scatter and biases at lower resolution. We find that $R \gtrsim 1700$ is required for accurate results for the gas kinematics of our sample. This work represents the largest available sample of well-resolved outflow velocity structure at $z > 2$, and highlights the need for good spectral resolution to recover accurate properties.

3.2. Introduction

The formation and evolution of galaxies is regulated by feedback from star formation and supermassive black hole growth (e.g., King, 2003; Veilleux et al., 2005a; Di Matteo et al., 2005; Fabian, 2012; Somerville & Davé, 2015b). The energy released by high star formation or black hole accretion rates can drive powerful galactic-scale outflows of gas and dust, limiting future star formation (e.g., Zhang, 2018; Naab & Ostriker, 2017; Hayward & Hopkins, 2017). At redshifts $z \simeq 2-3$, corresponding to the peak period of cosmic star formation activity ("Cosmic Noon"; e.g., Madau & Dickinson, 2014), virtually *all* star-forming galaxies exhibit outflows (e.g., Frye et al., 2002; Shapley et al., 2003; Sugahara et al., 2019). This is indeed expected based on their high star formation rate (SFR) surface densities (Heckman, 2002; Cicone et al., 2016).

Outflows in high-redshift galaxies are typically identified by interstellar medium (ISM) features in the rest-frame ultraviolet spectrum. Outflowing gas produces blueshifted absorption, and redshifted emission in Lyα and other resonant lines. This signature is observed ubiquitously in $z > 2$ star-forming galaxies (Weiner et al., 2009; Shapley et al., 2003; Vanzella et al., 2009; Steidel et al., 2010; Jones et al., 2012; Du et al., 2018). However, while large samples are available, the spectral resolution R is typically too low to resolve the outflow velocity structure. At $R \sim 600$ the full width at half-maximum (FWHM) resolution is $\sim 500\ \mathrm{km\,s^{-1}}$, which is comparable to the maximum observed velocities, whereas in this work we will focus on $R \gtrsim 4000$ corresponding to FWHM $\lesssim 75\ \mathrm{km\,s^{-1}}$. Furthermore, many studies rely on stacking analyses which preclude characterizing individual systems. Our current knowledge is thus largely limited to the average velocity centroid, which encompasses both outflows and ambient interstellar material. This leaves key questions unanswered, such as the proportion of gas which is able to escape the galaxy halo (as opposed to low-velocity gas which will remain in the circumgalactic medium (CGM) or recycle back to the galaxy), and the covering fraction of low-ionization gas which regulates the escape of ionizing photons (e.g., Du et al., 2018). Low-resolution data are likewise unable to disentangle outflows from the non-outflowing ISM component.

A promising way forward is to observe bright gravitationally lensed galaxies, which can be magnified by factors of $\sim 10\times$. Such bright sources enable moderate resolution spectroscopy with good sensitivity on 8–10m telescopes. Early studies of a few individual systems at $z \simeq 2-3$ revealed the velocity structure of ISM and outflowing gas spanning $\sim$1000 $\mathrm{km\,s^{-1}}$ (Pettini et al., 2002b; Quider et al., 2009, 2010; Dessauges-Zavadsky et al., 2011). Similarly, Jones et al. (2013a) and Leethochawalit et al. (2016d) used deep spectroscopy of seven strongly lensed $z > 4$ galaxies to measure their covering fraction profiles, revealing a considerable diversity among the star-forming population.

The number of well-characterized strongly lensed systems has grown tremendously over the last decade thanks to all-sky surveys and dedicated lens searches (e.g., Sonnenfeld et al., 2018; Jacobs et al., 2019a; Huang et al., 2020). Previously, Jones et al. (2018) conducted a study of 9 bright lensed galaxies from the `CASSOWARY` survey (Belokurov et al., 2009; Stark et al., 2013), quantifying their bulk outflow velocities and chemical compositions. This work aims to compile a larger sample of 20 targets observed at moderate spectral resolution ($R \sim 2530 - 6300$) with the main goal of quantifying the ISM outflow velocity structure in a statistical sense. With these results we seek to aid and improve upon the interpretation of larger samples at lower spectral resolution, by comparing trends in outflow velocities between low and moderate resolution data. Finally, we seek to compare the measured outflow velocities with those obtained in simulations with different feedback prescriptions, and provide a benchmark data sample for future comparison with cosmological simulations.

This paper is organized as follows: Section 3.3 describes the lensed galaxy sample and moderate resolution spectroscopy. In Section 3.4 we derive velocity profiles of the interstellar and outflowing gas, while Section 3.5 discusses the kinematic features of the ISM. Section 3.6 compares the observations of outflow velocities with scaling relations from previous work and simulations. We summarize the main conclusions of this work in Section 3.7. Throughout this paper, we use the AB magnitude system and a ΛCDM cosmology with $\Omega_M = 0.3$, $\Omega_\Lambda = 0.7$ and $\mathrm{H}_0 = 70\ \mathrm{km\,s^{-1}\ Mpc^{-1}}$.

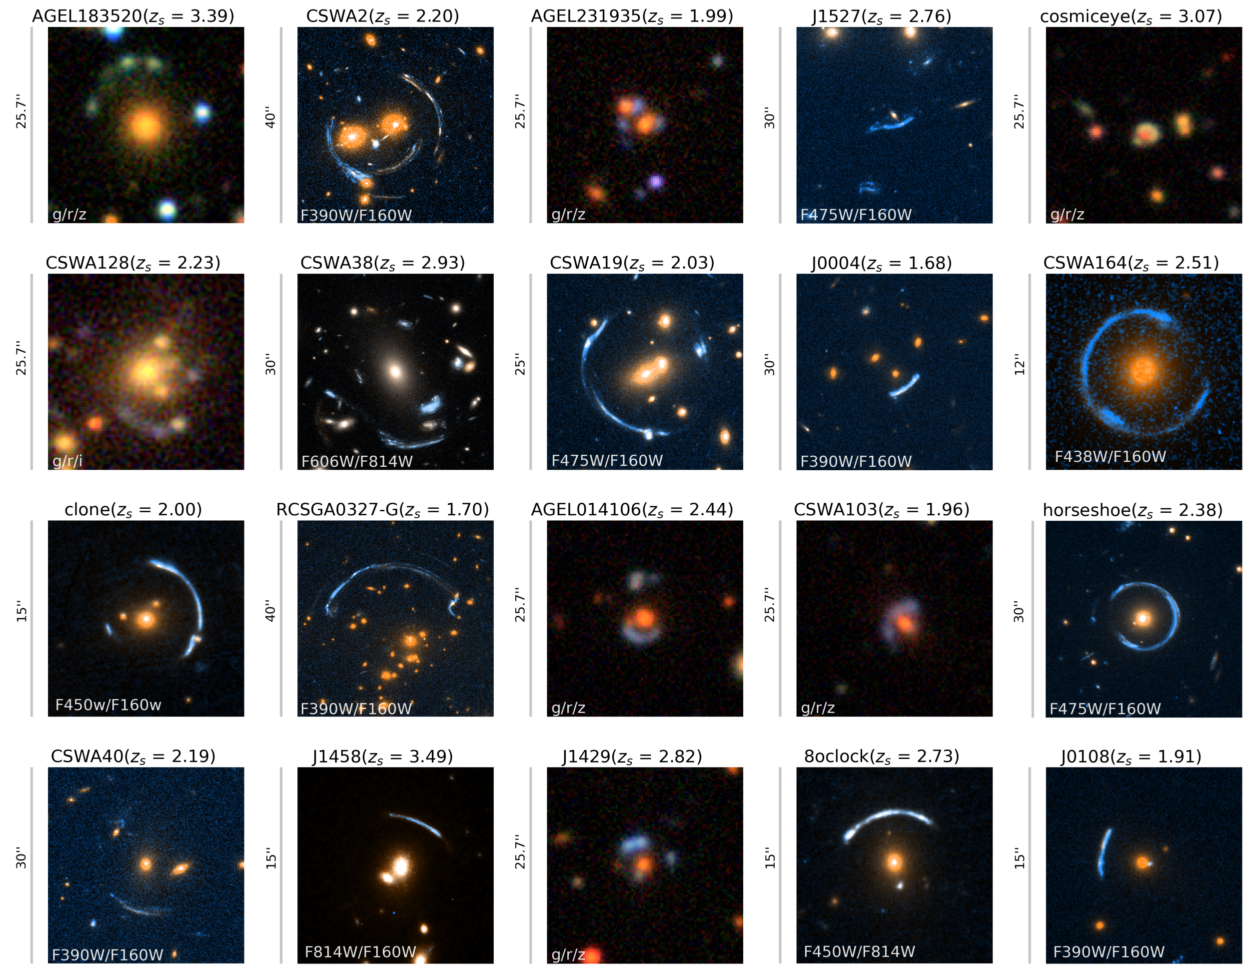

FIGURE 3.1. Color composite images of the gravitationally lensed galaxies used in this paper obtained either from DECaLS, SDSS, or the Hubble Space Telescope archive. The galaxies are arranged in order of increasing outflow velocity parameter ($v_{75,V2}$), from top to bottom. Each image is centered on the deflector(s) contributing to the lensing potential. These lensed galaxies are bright and appear highly magnified on the sky with a mean magnification value of $\mu = 9$. The images are oriented North-up, East-left with the image sizes labeled in arcseconds. The RA, Dec slit position and the position angle (PA) used for observations can be found in Table 3.1 and the references provided therein.

TABLE 3.1. Table of Galaxy properties. References for each survey are as follows. MEGASAURA: Rigby et al. (2018), CASSOWARY: Jones et al. (2018), AGEL: Tran et al. (2022b), KOA: Keck Observatory Archive. z_s is the source galaxy systemic redshift, and R is the spectral resolution. RA, Dec and PA (position angle) correspond to the slit locations used to observe the galaxies. The spectral features used to determine z_s are listed under Notes.

Object ID	RA (slit)	Dec (slit)	PA (slit)	z_s	Notes	Survey	R
J0004	00:04:51.685	−01:03:20.86	parallactic	1.6812	stellar absorption	MEGASAURA	2750
RCSGA0327-G	03:27:26.626	−13:26:15.30	parallactic	1.70385	nebular emission	MEGASAURA	2830
J0108	01:08:42.206	+06:24:44.41	parallactic	1.9099	stellar absorption	MEGASAURA	4380
CSWA103	01:45:04.38	−04:55:50.8	115	1.95978	C III]	CASSOWARY	6300
AGEL231935+115016	23:19:34.66	+11:50:18.1	-40	1.99256	ISM absorption	AGEL	4700
Clone	12:06:10.65	+51:44:44.1	40	2.0026	stellar absorption	KOA	4700
CSWA19	09:00:02.80	+22:34:07.1	86	2.03237	C III]	CASSOWARY	6300
CSWA40	09:52:40.29	+34:34:39.2	70	2.18938	stellar absorption	CASSOWARY	6300
CSWA2	10:38:41.88	+48:49:22.4	17	2.19677	C III]	CASSOWARY	6300
CSWA128	19:58:35.44	+59:50:52.2	60	2.22505	O III]	CASSOWARY	6300
HorseShoe	11:48:33.264	+19:29:59.11	parallactic	2.3814	stellar absorption	MEGASAURA	3980
AGEL014106-171324	1:41:06.1273	−17:13:23.545	320	2.43716	ISM absorption	AGEL	4700
CSWA164	02:32:49.93	−03:23:25.8	158	2.51172	stellar absorption	CASSOWARY	6300
8oclock	00:22:40.36	+14:31:27.6	-276	2.735	stellar absorption	KOA	4700
J1527	15:27:45.116	+06:52:19.57	parallactic	2.76238	stellar absorption	MEGASAURA	2740
J1429	14:29:54.857	+12:02:38.68	parallactic	2.8241	stellar absorption	MEGASAURA	3500
CSWA38	12:26:51.48	+21:52:17.9	130	2.92556	stellar absorption	CASSOWARY	6300
CosmicEye	21:35:12.7	−01:01:42.9	parallactic	3.0734	stellar absorption	MEGASAURA	2530
AGEL183520+460627	18:35:20.55	+46:06:35.4	16	3.38845	nebular emission	AGEL	4700
J1458	14:58:36.143	−00:23:58.17	parallactic	3.487	stellar absorption	MEGASAURA	4000

3.3. Sample and spectroscopic data

The goals of this work require moderate resolution spectroscopy ($R \gtrsim 4000$) in order to sample the ISM absorption profiles with ~10 independent spectral resolution elements. We have compiled a sample from our previous work, other archival data, and new observations from an ongoing survey of bright lensed galaxies discovered in wide area imaging surveys. The full sample used in this work is listed in Table 3.1, and color images of each source are shown in Figure 3.1. Below we describe the spectroscopic data sets.

(1) `CASSOWARY`: The Cambridge And Sloan Survey Of Wide ARcs in the skY (CASSOWARY, abbreviated CSWA) consists of bright lensed galaxies discovered in Sloan Digital Sky Survey (SDSS) imaging (Belokurov et al., 2009; Stark et al., 2013). Followup echellete spectra were taken with ESI (Sheinis et al., 2002) at Keck Observatory using an $0\farcs75$ slit width, resulting in $R = 6300$ resolution (FWHM = 48 $\mathrm{km\,s^{-1}}$) covering a wavelength range of 3900–11000 Å. These data are described in Jones et al. (2018) including an analysis of the ISM chemical composition. 7 targets from this sample (CSWA2, CSWA19, CSWA38, CSWA40, CSWA103, CSWA128, and CSWA164) have sufficient data quality and coverage of the ISM lines needed for this work.

(2) `MEGASAURA`: The Magellan Evolution of Galaxies Spectroscopic and Ultraviolet Reference Atlas (MEGaSaURA) consists of spectra of lensed galaxies taken with the MagE spectrograph on the Magellan telescopes, extracted over the wavelength range 3200–8280 Å (Rigby et al., 2018). 8 targets from this sample (J0004, J0108, J1429, J1458, J1527, CosmicEye, HorseShoe, and RCSGA0327-G) are used in this paper. A range of MagE slit widths were used resulting in spectral resolution ranging from 2530-4400, with an average $R = 3300$.

(3) `AGEL`: As part of the ASTRO3D Galaxy Evolution with Lenses (AGEL) project, we have obtained Keck/ESI spectra of bright lensed galaxies discovered from a

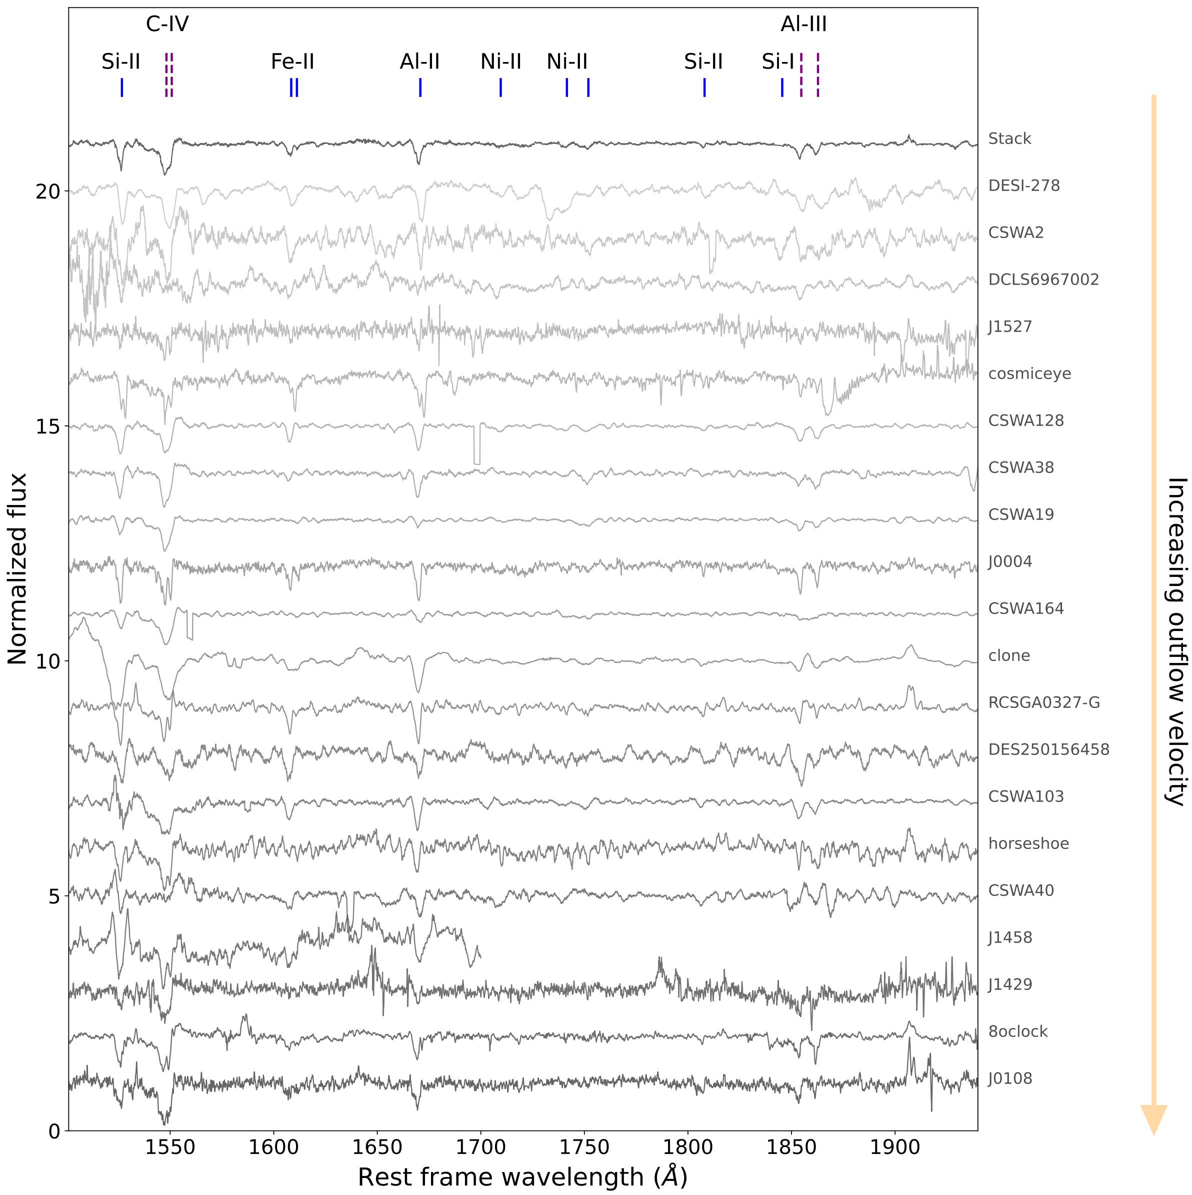

FIGURE 3.2. Spectra of the 20 lensed galaxies used in this work showing the rest-frame wavelength range 1500–1950 Å. The spectra are sorted from top to bottom in the increasing value of $v_{75,V2}$ (75% outflow velocity measured considering only gas with $v < 0$; see Table 3.2), and offset for clarity. Low ionization (e.g., Si II, Al II) and high ionization (e.g., C IV) lines are marked in blue and purple respectively. The interstellar transitions probe a range of optical depth, from the strong lines such as Si II $\lambda1526$ and Al II $\lambda1670$ to the weak (optically thin) Ni II features. A median-stacked spectrum is shown in the top row to demonstrate the various ISM absorption features with higher signal-to-noise.

machine learning search in wide area imaging. The search methodology and a subset of targets are described in Jacobs et al. (2019a). Spectra were taken with a $1''.0$ slit providing $R = 4700$ resolution (FWHM = 64 $\mathrm{km\,s^{-1}}$) covering a wavelength range of 3900-11000 Å. The observations are described in Tran et al. (2022b). 3 targets (AGEL231935+115016, AGEL122651+215218, AGEL183520+460627) from the AGEL sample are used in this paper.

(4) KOA: Data for 2 additional bright lensed galaxies (Clone, 8oclock) were obtained from the Keck Observatory Archive (KOA) and reduced using MAKEE written by Tom Barlow[1] for inclusion in this analysis. These observations were taken with the same setting as the AGEL sample, using the $1''.0$ slit. The reduction was performed following the same methods used for the AGEL data (Tran et al., 2022b), with default settings prescribed for ESI. A manual extraction region covering the entire galaxy light was taken to be the continuum, with the rest of the slit considered as the sky to generate the error spectra. Extracted 1D spectra are binned to a common dispersion of 11.5 $\mathrm{km\,s^{-1}}$ per pixel.

Our sample is comprised of moderately massive, star-forming main-sequence galaxies (Section 3.6), which show no evidence of AGN in the available spectra. Those with resolved spectroscopic observations exhibit a wide range of kinematic structure (e.g., Stark et al., 2008; Jones et al., 2013b; Wuyts et al., 2014b; Bordoloi et al., 2016; Leethochawalit et al., 2016b; Chisholm et al., 2018; James et al., 2018a; Shaban et al., 2022) with HorseShoe, Clone and CosmicEye being rotationally supported whereas CSWA2, CSWA19, CSWA38, CSWA128, and RCSGA0327 appear to be mergers/interacting systems. Figure 3.2 plots spectra of the full sample between 1500-1950 Å, with prominent absorption features labelled.

3.3.1. Systemic Redshifts. Systemic redshifts are needed to characterize ISM kinematics with respect to the stars. Table 3.1 lists the redshifts for galaxies in our sample, which span

[1] `https://www2.keck.hawaii.edu/inst/esi/makee.html`

$z = 1.6$–3.5, along with the type of features used for these measurements. In most cases, the systemic redshift is based on stellar photospheric absorption lines with a typical uncertainty of $\leq 30\,\mathrm{km\,s^{-1}}$. In some cases where suitable stellar features are not reliably measured (e.g., CSWA19), we use nebular C III] or O III] emission lines to establish the redshift. Photospheric absorption or nebular emission lines are available for 18 of the targets in our sample. For 2 targets (AGEL231935+115016, AGEL014106-171324) where none of these features are securely measured, we estimate the systemic redshift from the ISM absorption lines themselves in the following way: we find the velocity corresponding to the maximum covering fraction and then apply an offset of $+172 \pm 19\ \mathrm{km\,s^{-1}}$ (i.e. $z_{sys} = z_{Cf,max} + 172\,\mathrm{km\,s^{-1}}/c \pm 19\,\mathrm{km\,s^{-1}}/c$). This offset value is derived as the median difference and sample standard deviation between the systemic and $z_{Cf,max}$ from the 18 galaxies in the sample with robust systemic redshifts. The offset between systemic and ISM absorption velocities in our sample is comparable to measurements from non-lensed galaxies at similar redshifts (e.g., Steidel et al., 2010; Jones et al., 2013a).

3.3.2. Continuum normalization. In this work, we are interested in the strength of interstellar absorption relative to the stellar continuum. In order to achieve a constant continuum level around the ISM lines, the spectra from all targets are initially normalized by a running median of 2001 pixels ($\sim$20,000 $\mathrm{km\,s^{-1}}$) which removes any large scale structures in the spectra arising from effects such as dust attenuation, flux calibration, and flat fielding uncertainties (e.g., Figure 3.3 - top). To remove any local scale structures, we consider a region spanning -2000 to $2000\ \mathrm{km\,s^{-1}}$ around the ISM line of interest and divide it by the median value of the local region. We then average the ISM lines (Section 3.4.1) and divide by a third order polynomial fit to the continuum around the absorption profile, to account for any residual structure. This achieves a continuum level close to 1 for the mean absorption profile in all target galaxies (e.g., Figure 3.3). Any absorption can then be interpreted as gas present along the line-of-sight in front of this continuum starlight.

We note that using a third order polynomial normalization increases the Δv_{90} line widths by 48 ± 36 km s^{-1} on average compared to using only a median normalization, although there is no effect on the centroid (-3 ± 7 km s^{-1}). Additionally, we estimate the typical uncertainty in continuum level using the third order polynomial normalization to be approximately $\sim$1%, which propagates to a $\sim$3% average change in the width of ISM absorption as parameterized by Δv_{90} (Section 3.5) or similar quantities, while velocity centroids remain consistent within the statistical uncertainties. The effect is such that an underestimated continuum implies an underestimated Δv_{90} and absorption equivalent width from best-fit profiles. This uncertainty does not significantly affect the main results and conclusions presented herein.

The normalization procedure described here is relatively insensitive to the ISM line itself. In some cases the lensed galaxy spectra are subject to blending with the deflector light due to the nature of the observations, especially with `AGEL` and `MEGASAURA` data. This can affect the relative depth and equivalent width of ISM absorption profiles. However, the kinematic measurements used in this work are robust to blending with other sources, provided they have smooth continuum spectra. Any strong spectral features which interfere with the ISM lines of interest are masked out and not used in our analysis. In some cases, there are strong intervening absorption systems at lower redshift, which are likewise masked and not used in this analysis.

3.4. Velocity structure of ISM gas

Ultraviolet ISM absorption lines probe the velocity structure of gas seen along the line of sight toward ("in front of") the young stars in a galaxy. Spectrally resolving the absorption velocity profile is a practical and powerful way to probe the baryon cycle, as illustrated schematically in Figure 3.4. Interstellar gas within the galaxy will absorb at the systemic redshift (i.e., $v = 0$) with a velocity range set by the galactic rotation curve and velocity dispersion. Inflowing gas gives rise to redshifted absorption (at $v > 0$), while outflows result in blueshifted absorption ($v < 0$) which may even exceed the escape velocity. Recycling gas

– which transitions from outflowing to inflowing at moderately low velocity – would result in absorption near $v \approx 0$.

In this section we describe our methodology to determine spectrally resolved ISM absorption profiles, in order to characterize the gas kinematics and geometric covering fractions in our sample. The observed intensity I for an interstellar absorption line is

$$\frac{I(v)}{I_0} = 1 - \Psi(v), \tag{3.1}$$

where I_0 is the intensity of stellar continuum and $\Psi(v)$ describes the absorption depth as a function of velocity v. It is dependent on the covering fraction of gas $C_f(v)$ and the optical depth τ in the following way:

$$\Psi(v) = C_f(v)(1 - e^{-\tau}) \approx C_f(v) \tag{3.2}$$

where the latter approximation is valid for the case of optically thick absorption $\tau \gg 1$. In this work, we are interested in studying the ISM gas kinematics by measuring the covering fraction as a function of velocity $C_f(v)$ from galaxy-integrated slit spectra. We describe the ISM absorption profiles and velocity structure for lines of different optical depth in Section 3.5.3.1. The profiles are consistent among the stronger transitions indicating $\tau \gtrsim 1$ around the line center, suggesting that they largely trace the covering fraction. For our analysis we use these strongest lines with $\tau \gtrsim 1$, such that Equation 3.2 is a reasonable approximation. We note that if the gas is not optically thick (e.g., as may be the case at higher velocities), then these represent a lower limit on the covering fraction.

3.4.1. Kinematics of the low-ionization gas. The rest-frame UV spectra used in this work include interstellar absorption from both low- and high-ionization species, as well as stellar features, Lyα in absorption and/or emission, and other features such as nebular and fine structure emission (see Figure 3.3 for an example). We focus on ISM kinematics of the low-ionization phase, from which there are numerous prominent transitions of Si II, O I, C II,

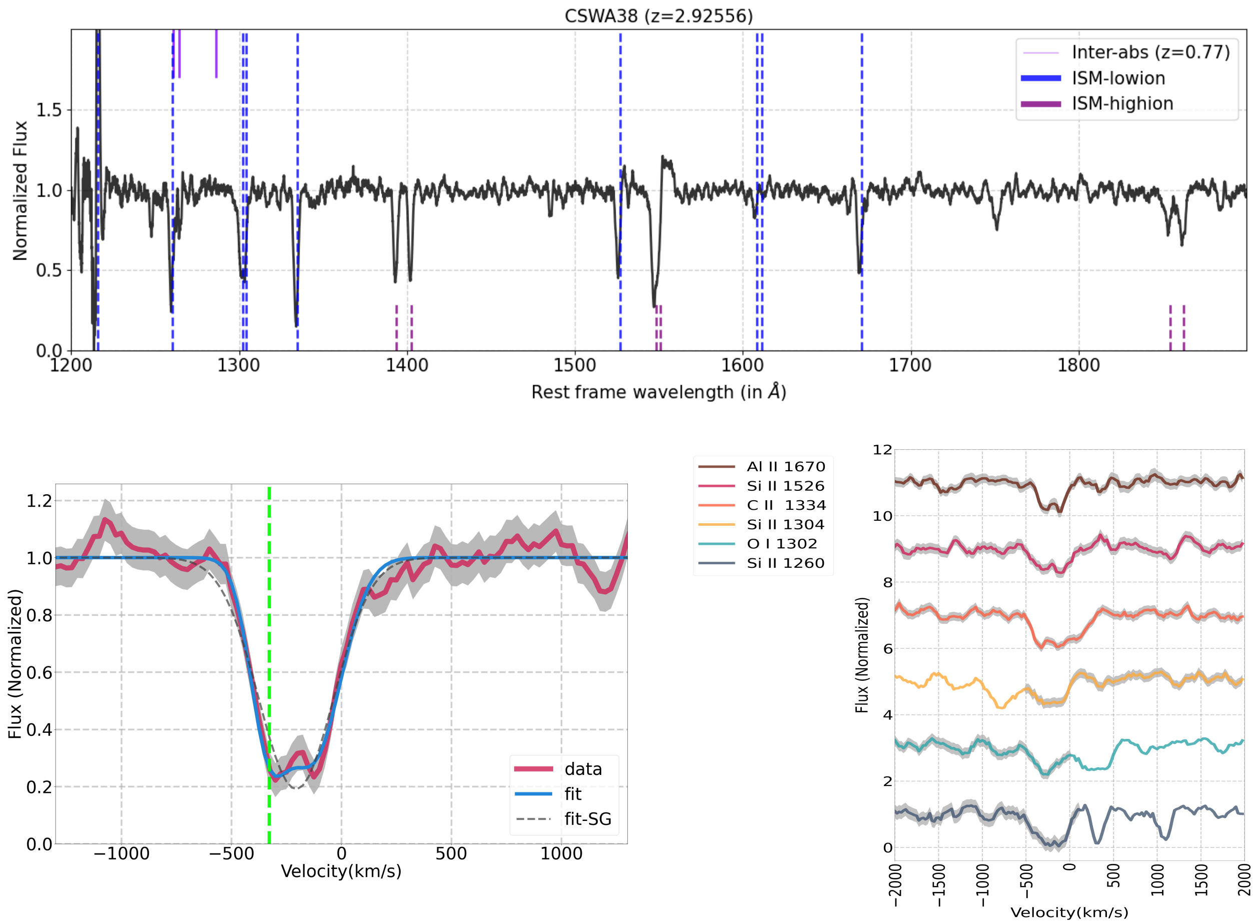

FIGURE 3.3. *Top:* Example normalized spectrum of one of the lensed galaxies in our sample, CSWA38. Prominent low ionization and high ionization ISM lines from the source galaxy at $z = 2.92$ are marked in blue and purple dashed lines. Mg II absorption from an intervening absorber galaxy at $z = 0.77$ (Mortensen et al., 2021a) is marked in violet. *Bottom Left:* Mean absorption profile of gas as a function of velocity obtained from a weighted average of the low-ionization Si II $\lambda1260$, O I $\lambda1302$, Si II $\lambda1304$, C II $\lambda1334$, Si II $\lambda1526$, and Al II $\lambda1670$ lines. The mean absorption profile is related to the covering fraction as 1-$C_f(v)$ (Equation 3.2). The gray shaded regions represent the error spectrum. The blue line is a double Gaussian fit to the data, showing good agreement, while the dashed line is a best-fit single Gaussian (SG) profile. The green vertical line indicates the outflow velocity parameter $v_{75,V2}$ defined as the velocity (in $\mathrm{km\,s^{-1}}$) at 75% absorption considering only absorption with $v < 0$ (see Figure 3.6). *Bottom Right:* Spectra of low-ionization lines used to obtain the covering fraction profile. Regions which have no error bars (gray shading) are not included for the weighted average. In these cases the regions are excluded due to absorption from an intervening galaxy (e.g., $> 200\,\mathrm{km\,s^{-1}}$ in Si II $\lambda1260$) or the blended nature of the O I $\lambda1302$, Si II $\lambda1304$ lines.

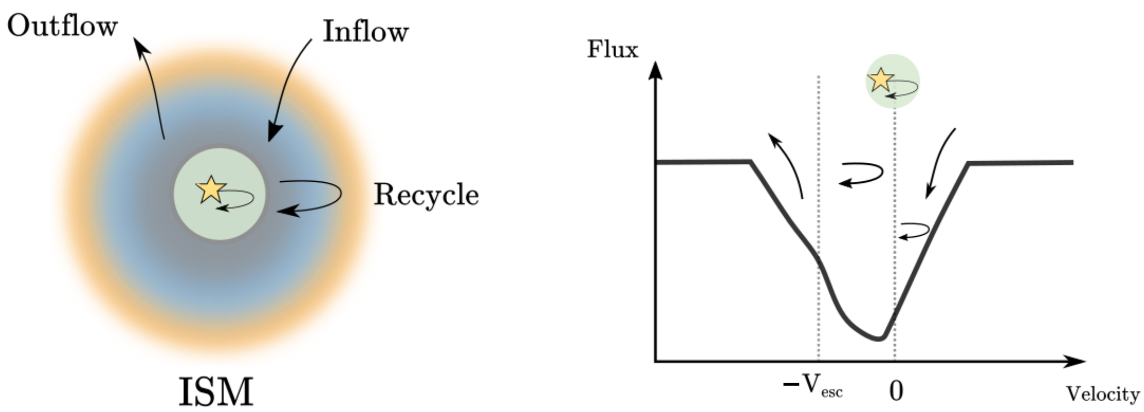

FIGURE 3.4. A guide to interpreting ISM velocity profiles in terms of the baryon cycle, shown as a schematic on the left with corresponding ISM absorption signatures on the right. Outflowing gas has blueshifted absorption (i.e. $v < 0$) whereas inflowing gas has redshifted absorption ($v > 0$). Recycling gas which arises from the outflowing gas transitioning to inflowing gas has velocities $v \approx 0$. The systemic velocity of the stars and their dispersion is centered at $v = 0$ by definition. The region $v < -v_{esc}$ corresponds to gas with velocities greater than the escape velocity of the ISM.

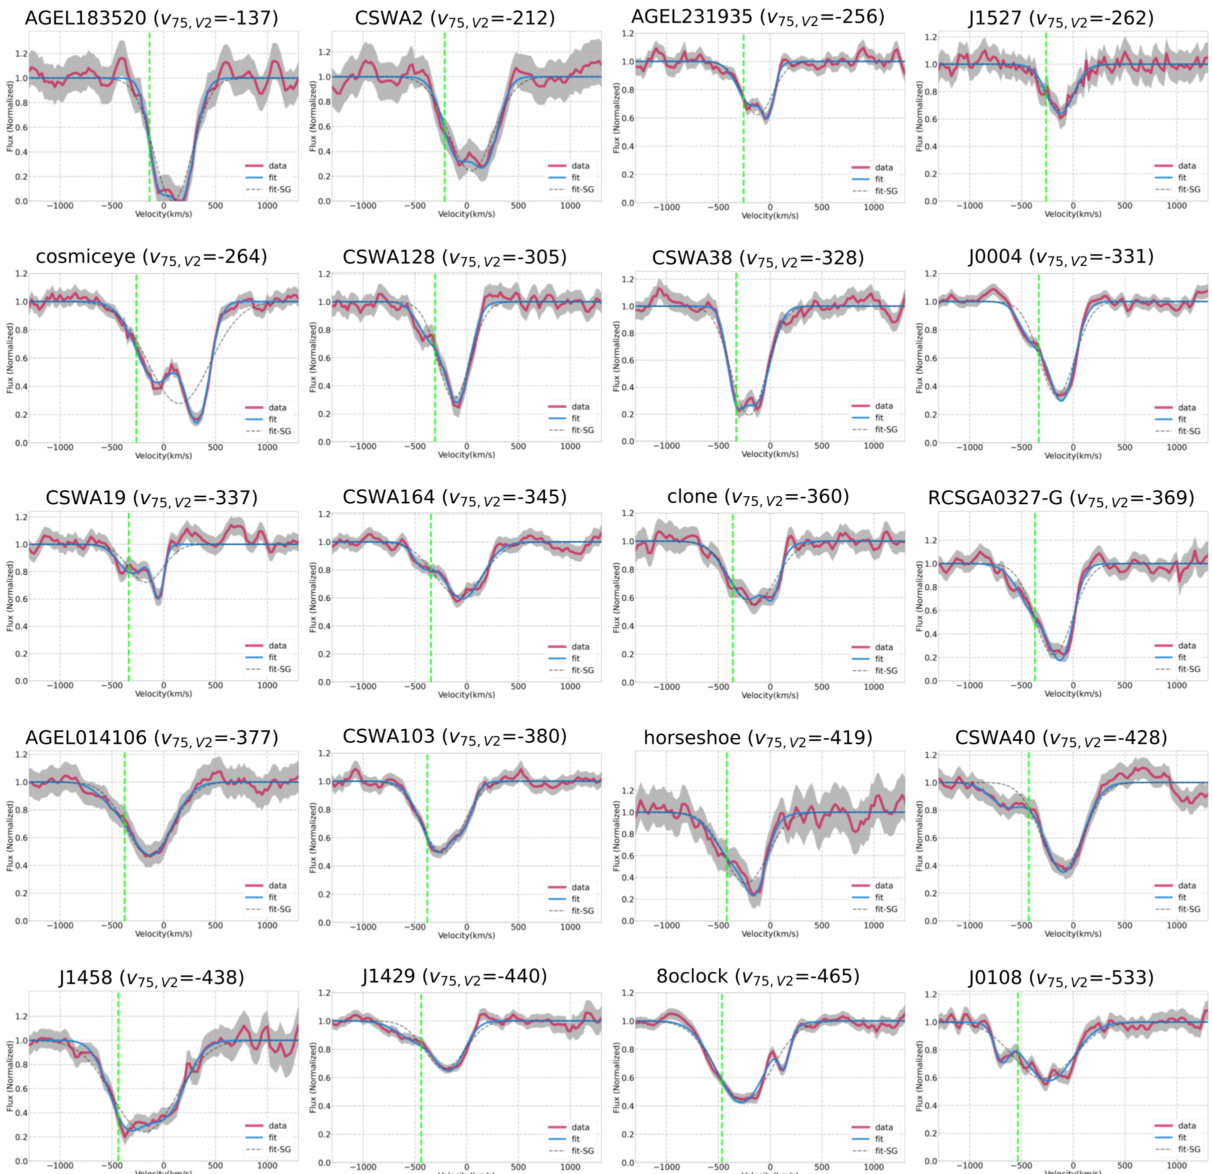

FIGURE 3.5. Plots of ISM absorption profiles for the full sample, sorted by increasing values of the outflow velocity parameter $v_{75,V2}$ (given in $\mathrm{km\,s^{-1}}$). The normalized flux profiles are related to covering fraction as $1 - C_f(v)$, and $v_{75,V2}$ is the 75% percentile of absorption measured by considering gas only at $v < 0$, where $v = 0$ is the systemic velocity (see Table. 3.2 and Section. 3.5 for more details). Denoted in red is the observed velocity profile, gray regions show the 2σ confidence interval, blue lines are double Gaussian (DG) fits to the data, and dotted lines are single Gaussian (SG) fits. The green vertical lines indicate the measured $v_{75,V2}$ in each case with the value (in $\mathrm{km\,s^{-1}}$) given above each plot.

Al II and Fe II. These metal ion transitions are often optically thick, approximately tracing the gas covering fraction as a function of velocity (Equation 3.2). This is in contrast to the H I Lyα profile which is complicated by resonant emission and damping wings.

For each spectrum we measure the ISM absorption profile I/I_0 from an average of the best available strong low-ion metal lines. We select those with good continuum sensitivity which appear to be saturated (based on multiple lines showing similar absorption profiles). Each absorption line is interpolated to a common velocity grid of 25 $\mathrm{km\,s^{-1}}$ and we take an inverse-variance weighted mean of the median-continuum-normalized flux at each velocity. Those which are affected by features such as strong sky emission, telluric absorption, bad pixels, or intervening absorption systems are excluded from this analysis. For blended transitions (e.g., O I λ1302 and Si II λ1304), only regions of interest corresponding to the transition are taken into account. Specifically, we use typical velocity ranges $v \lesssim 0$ $\mathrm{km\,s^{-1}}$ and $v \gtrsim -500$ $\mathrm{km\,s^{-1}}$ for the λ1302 and λ1304 transitions, respectively, similar to the approach of Jones et al. (2018). Other ISM lines are affected to a lesser extent by blending with weak features such as [S II] λ1259 (blended with Si II λ1260), C II* λ1335 (affecting C II λ1334), and stellar photospheric features near the O I λ1302 line. These features and their effects on derived ISM absorption profiles are typically not detected in individual galaxy spectra. We therefore do not mask these regions, effectively treating them as part of the stellar continuum (which is generally full of lines with low equivalent width). From analysis of the high-SNR stacked spectrum, we find that these blends can cause an increase in the measured Δv_{90} by up to $\lesssim 70$ $\mathrm{km\,s^{-1}}$ depending on the lines used. This represents a source of systematic uncertainty in the absorption profiles, with magnitude comparable to uncertainty arising from the continuum normalization (Section 3.3.2).

Figure 3.3 illustrates this process for an example galaxy in the sample, with equivalent figures for the full sample displayed in Appendix 3.7. We derive the low-ionization ISM covering fraction profiles $C_f(v)$ from these mean absorption profiles for the galaxies in our sample, shown in Figure 3.5.

3.4.2. Fitting $C_f(v)$. The mean ISM absorption profiles shown in Figure 3.5 encode the key observational results of this paper. From these profiles we can examine the typical outflow velocities, the maximum velocities with substantial gas covering fractions, and diversity within the sample, among other properties. For analysis purposes, it is useful to have an analytic form which captures the velocity structure of ISM absorption profiles. For quasar sightlines a Voigt profile is appropriate to describe distinct absorption components, but this is not suitable for galaxy spectra whose profiles represent a large number of interstellar clouds.

Although we adopt the weighted mean profile measurements shown in Figure 3.5 as the ground truth, we also fit two analytic functions to each profile. The first is a single Gaussian (hereafter SG) function of the form:

$$C_f(v) = A_{sg} \exp[(v - v_{sg})^2 / 2\sigma_{sg}^2] \tag{3.3}$$

where the subscripts indicate a single Gaussian (sg). This does not capture the clear asymmetries seen in most of the sample (Figure 3.5). It is nonetheless instructive since this fit captures the information equivalent of a low-resolution ($R \sim 300 - 1000$) spectrum, in which the absorption would be only marginally resolved. The second profile is a double Gaussian (hereafter DG) function of the following form:

$$C_f(v) = A_0 \exp[(v - v_0)^2 / 2\sigma_0^2] + A_1 \exp[(v - v_1)^2 / 2\sigma_1^2] \tag{3.4}$$

where (A_0, v_0, σ_0, A_1, v_1, σ_1) are the parameters to be fit. We adopt a convention that $v_0 < v_1$. The DG is relatively simple but versatile. We find that it yields a reasonable fit to the velocity substructure detected in our sample. The median residuals of the best fit DG model measured between the velocity range v_{99} and v_{01} are ~ 0.03 whereas for the SG they are $\sim 5\times$ higher. We therefore make use of the DG fits to derive kinematic properties such as the velocity centroid and width (Section 3.5.1). The SG fits are used mainly as an emulator of lower spectral resolution data.

To quantify the uncertainty in each parameter, we fit each weighted mean absorption profile with 250 realizations of the Basin-Hopping stochastic algorithm (Wales & Doye, 1997). For each realization we add random 1σ noise to $C_f(v)$ based on the error spectrum. The velocity centroids (v_{sg}, v_0, v_1) are allowed to vary from -700 to 500 $\mathrm{km\,s^{-1}}$, dispersions ($\sigma_{sg}, \sigma, \sigma_1$) from 50 to 700 $\mathrm{km\,s^{-1}}$, and absorption depth (A_{sg}, A_0, A_1) from 0.1 to 1. For each realization, all parameters are initialized to random values within the above ranges. These bounds are chosen based on the observed covering fraction profiles such that they sample the entire parameter space. We place an additional constraint -800 km s^{-1} $< (v_0 - v_1) < 0$ km s^{-1} when fitting the double Gaussian. This ensures that the same component (v_0) always captures the blueward absorption, which we will generally attribute to outflowing gas. We note that this approach is somewhat more general than that of e.g. Bordoloi et al. (2016) in which one component's centroid is fixed to represent the systemic component; here we do not require any component to exactly trace the systemic velocity.

Table 3.2. Definitions of velocity measurements presented in this paper. We note that v_{50} is the centroid velocity and is used interchangeably with v_{cent}.

Parameter-DG	Parameter-SG	Description
v_{05}	$v_{05,SG}$	Velocity at 5% absorption
$v_{50} = v_{cent}$	$v_{50,SG} = v_{cent,SG}$	Velocity at 50% absorption
v_{90}	$v_{90,SG}$	Velocity at 90% absorption
v_{95}	$v_{95,SG}$	Velocity at 95% absorption
v_{99}	$v_{99,SG}$	Velocity at 99% absorption
Δv_{90}	$\Delta v_{90,SG}$	$v_{95} - v_{05}$
$v_{05,V2}$	$v_{05,SG,V2}$	These quantities are
$v_{50,V2}$	$v_{50,SG,V2}$	calculated in the same
$v_{90,V2}$	$v_{90,SG,V2}$	way as described
$v_{95,V2}$	$v_{95,SG,V2}$	above but considering
$v_{99,V2}$	$v_{99,SG,V2}$	only absorption with $v < 0$

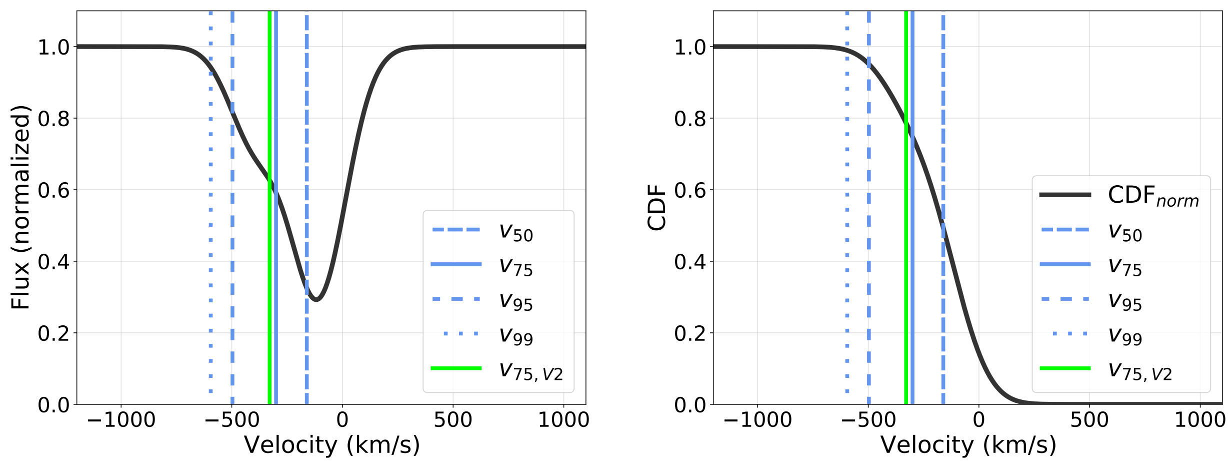

FIGURE 3.6. Illustration of the velocity measurements used in this work. *Left:* An example normalized flux profile analogous to those in Figure 3.5. *Right:* The normalized Cumulative Distribution Function (CDF) of the absorption profile shown in the left panel. This is essentially the CDF of the covering fraction $C_f(v)$ (Equation 3.4.2). Both panels show velocities v_{50}, v_{75}, v_{95}, and v_{99} which are defined as the 50%, 75%, 95% and 99% percentiles of absorption (as seen from the CDF). A subscript of V2 (e.g., $v_{75,V2}$ shown in green) indicates that only $v < 0$ was considered; essentially the CDF is normalized to zero at $v = 0$ for this case. Table 3.2 lists the entire set of quantities used.

The median fit values obtained at the end of 250 realizations are used to estimate all velocity measurements used in this paper. The 1σ standard deviation of each quantity is calculated as $\sigma \approx \mathrm{MAD}/0.675$ where MAD = Median Absolute Deviation. Unlike the mean and standard deviation which are easily affected by spurious outliers, the median value and MAD offer better quantifiable values to describe the fits. The resulting best-fit profiles are plotted in Figure 3.5 along with the data and observational uncertainties. Appendix 3.7 lists the fit parameters obtained for each of the targets along with the derived uncertainties.

In all targets we find that the DG fits are able to capture the broad asymmetric wings which are ubiquitously present in the absorption profiles. In addition, they also accommodate complex absorption profiles such as the Cosmic Eye which includes a strong redshifted component. The performance of the SG fits on the other hand varies heavily depending on the asymmetry of the profile. In some cases (e.g., CSWA103), they provide reasonably good fits whereas in more asymmetric cases (e.g., CSWA128) there are large residuals, especially at high velocities. Encouragingly, the residuals obtained for the DG fits are consistently centered around 0 with the standard deviation being generally compatible with the signal-to-noise of each spectrum, indicating a reasonable fit to the data.

3.5. Kinematic Features of the Gaseous ISM at Cosmic Noon

Having obtained a covering fraction profile including parametric fits for each galaxy, we now explore kinematic properties of the sample. We measure various standard quantities to facilitate comparison of these moderate resolution down-the-barrel results with other probes (including quasars, low resolution galaxy spectra, and theoretical simulations). To best compare with the literature we adopt two parallel lines of analysis: (a) considering the entire covering fraction profile; and (b) considering only the absorption at $v < 0$ (i.e., blueshifted) which we denote with a $_{V2}$ subscript. The latter is useful in comparison with theoretical studies which consider only outflowing gas. However we note that the $v < 0$ absorption still includes approximately half of the systemic component.

For both analyses we measure the velocity corresponding to the percentiles 5%, 75%, 95%, 99%, and 50% of absorption (denoted as v_{05}, v_{90}, etc.), larger percentiles being more blueshifted. Here $v_{50} = v_{cent}$ is the velocity centroid (v_{50} and v_{cent} are interchangeable). We also measure the velocity width $\Delta v_{90} = v_{05} - v_{95}$, spanning the 5–95 percentile of absorption. Table 3.2 lists all quantities used in our analysis, and we illustrate some of these for an example velocity profile in Figure 3.6. All quantities are calculated for both the SG and DG fits, with results given in the appendix. These quantities have been found useful to describe the kinematics in observational and simulation studies in the literature, and we adopt the same conventions for ease of comparison.

3.5.1. Bulk outflow motion of ISM gas. A visual inspection of the global covering fraction profiles obtained in Section 3.4 and Figure 3.5 indicates that the bulk motion of the gas in the ISM is outflowing, with blushifted velocity centroids ($v_{cent} < 0$). Quantifying the kinematic properties is an important step towards understanding the feedback processes which drive these outflows and impact the host galaxy evolution. In this section we probe the outflow velocity structure quantitatively in terms of v_{cent}, v_{75}, and v_{95}.

The centroid v_{cent} gives a measure of typical outflow velocities, which can readily be compared with other samples. The median absorption centroid and its sample standard deviation for the galaxies in the sample is -141 ± 111 km s^{-1} for the full profiles, and -216 ± 61 km s^{-1} if we consider only the velocities $v < 0$. This latter number is a lower limit to the purely outflowing gas component (as opposed to the total including systemic interstellar absorption). Figure 3.7 illustrates the histogram obtained for both these metrics. These values are similar to measurements from larger samples of $z \simeq 2$–3 galaxies at lower spectral resolution (e.g., -168 ± 16 km s^{-1} from Steidel et al. 2010 compared with our sample median -141 ± 25 km s^{-1}). The covering fraction at $v_{cent,V2}$ ranges from 18-95% for galaxies in this sample with a median of 50%, suggesting a patchy covering fraction of the outflowing gas with substantial variations within the sample.

The v_{75} and v_{95} values probe the high-velocity blueshifted tail of outflowing gas. The distributions of these values for the lensed sample are also shown in Figure 3.7 (lower panel). Compared to the centroid velocity ($v_{cent} = v_{50}$), we find that the median $v_{75} \approx 2 \times v_{cent}$ and $v_{95} \gtrsim 3 \times v_{cent}$. Thus we see clear signatures of outflows at >3 times the centroid velocity, with absolute v_{95} typically extending beyond 450 $\mathrm{km\,s^{-1}}$, although the covering fraction is smaller at larger absolute velocity.

3.5.2. Quantifying the asymmetry in absorption. Another significant visual feature of the covering fraction profiles is the asymmetry. The quantities $|v_{50} - v_{05}|$ and $|v_{50} - v_{95}|$ trace the extent of the gas present redward and blueward of the bulk outflowing gas velocity. The median $|v_{50} - v_{05}|$ and $|v_{50} - v_{95}|$ measured with the DG are 292 $\mathrm{km\,s^{-1}}$ ($\approx 1.7 \times |v_{50}|$) and 357 $\mathrm{km\,s^{-1}}$ ($\approx 2.1 \times |v_{50}|$) respectively, indicating a clear skewness on average with a shallower slope for the blueshifted velocity range. In comparison, a SG fit gives 340 $\mathrm{km\,s^{-1}}$ ($\approx 2 \times |v_{50}|$) for the same quantities, which are identical by symmetry of the single Gaussian. Figure 3.8 plots a histogram of the skewness ratio defined as

$$\text{Skewness Ratio} = \frac{|v_{50} - v_{95}|}{|v_{50} - v_{05}|} - 1 \tag{3.5}$$

where a positive Skewness Ratio indicates that the blue wing is more extended than the red wing.

16 out of the 20 galaxies in our sample have positive skewness (i.e. Skewness Ratio > 0). Looking at the covering fraction profiles of galaxies which have Skewness ratio < 0 (e.g., J1458), one can clearly see that they have an inverted skewed profile wherein the redshifted side has a shallower slope (e.g., Figure 3.8) which can give rise to a negative skewness (i.e. $|v_{50} - v_{95}| < |v_{05} - v_{50}|$). The origin of this skewness in the profile is an interesting but challenging question which we do not tackle in this paper, but in a simplistic sense, the different skewness ratios could be interpreted as the response of the ISM gas to a galactic wind captured either at different points in time or viewing angles. A key point is that such

details about the kinematic structure are not captured by the SG fits (nor by low resolution spectra). This illustrates the need for good spectral resolution to reveal the complex velocity structure of outflowing gas.

3.5.3. Width of absorption using Δv_{90}. The Δv_{90} diagnostic is commonly used in the literature for quasar absorption systems, and is usually defined as the velocity range spanning 5% to 95% of the total column density. However there are some key differences between quasar probes and our measurements. First, quasars probe the full line-of-sight through a halo (distances $-\infty$ to ∞) whereas our "down-the-barrel" galaxy spectra sample only half the halo (0 to ∞). Our spectra do not probe the redshifted outflowing gas on the far side of the galaxy, causing Δv_{90} to be smaller than for a background quasar at impact parameter $b = 0$. Sec ond, quasars probe a narrow "pencil beam" area which is prone to stochastic sampling of absorbing gas clouds (e.g., Marra et al., 2022), whereas our galaxy spectra encompass a much larger cross-sectional area of several kpc^2. Thus we may expect our galaxy spectra to be more representative of the gas covering fraction. Third, the absorption profiles from Section 3.4 are constructed from the strongest ISM lines, which are more sensitive to gas covering fraction as opposed to column density. In summary, the Δv_{90} values for our sample represent approximately the velocity width of covering fraction profiles through half of the host galaxy halos.

3.5.3.1. *Kinematics at different optical depths(τ).* To assess how well the absorption profiles from strong ISM lines trace the column density, we compare them with weaker ISM absorption lines whose apparent optical depth is $\tau \sim 0.1$–1. The low ion velocity profiles are typically constructed from the strongest ISM transitions, with $\tau \gtrsim 1$ (see appendix). We compare these with the Al II $\lambda1670$ and Fe II $\lambda1608$ lines which are often unsaturated ($\tau \sim 1$), as well as the optically thin ($\tau \ll 1$) transition Si II $\lambda1808$. Median velocity profiles for each of these lines are obtained by stacking the spectra from all objects in the sample with the relevant wavelength coverage. Figure 3.9 shows a plot of the median stacked profiles for these different ISM absorption features as function of optical depth.

Visually inspecting the profiles reveals a remarkable similarity in the kinematics probed by the different transitions, despite the varying optical depths. Assuming that the stack of strong lines traces the covering fraction at $\tau \gg 1$ (Equation 3.2), the maximum absorption depth suggests $\tau \simeq 1.5$ for Al II $\lambda1670$, $\tau \simeq 0.6$ for Fe II , and $\tau \simeq 0.1$ for Si II $\lambda1808$ (supporting an optically thin interpretation). The Si II $\lambda1808$ profile exhibits blueshifted absorption consistent with the stronger features, although at lower signal-to-noise ratio. We perform a DG fit to the strongest ISM absorption line profiles and Si II $\lambda1808$ line to derive the velocity centroid (v_{cent}) and Δv_{90} values, as described above. Figure 3.9 plots these quantities (lower panels). We find the velocity centroid is ≈ -160 km s^{-1} for all transitions. The Δv_{90} for the stronger low-ion transitions is ~ 630 km s^{-1}, including for Fe II which has apparent $\tau < 1$, whereas for Si II $\lambda1808$ it is $\sim 400 \pm 100$ km s^{-1}. Visually, this difference in Δv_{90} between the optically thin and thick lines likely arises from the higher outflow velocity regions, which may be affected by lower τ in addition to reduced signal-to-noise. Nonetheless the line widths are broadly similar across a range of optical depth, indicating that we can use the Δv_{90} measurements obtained from strong low ion transitions to compare with measurements based on optical depth from quasar sightlines, with the caveat that values based on optical depth may be lower by ~ 250 km s^{-1}.

3.5.3.2. *Comparison to quasar sightlines.* In this subsection, we compare the width of absorption measured using Δv_{90} and Equivalent Width (EW) as we step away from "down-the-barrel" observations to pencil beam quasar sightlines probing larger impact parameters. Figure 3.10 plots the Δv_{90} measurements as a function of redshift (z). These values are compared with various quasar surveys: XQ-100 (Berg et al., 2016), EUADP (Quiret et al., 2016), and "Dusty DLAs" with 2175 Å dust attenuation bumps (Ma et al., 2017). The galaxies from this work have Δv_{90} values ranging between 440 and 920 km s^{-1} with a median

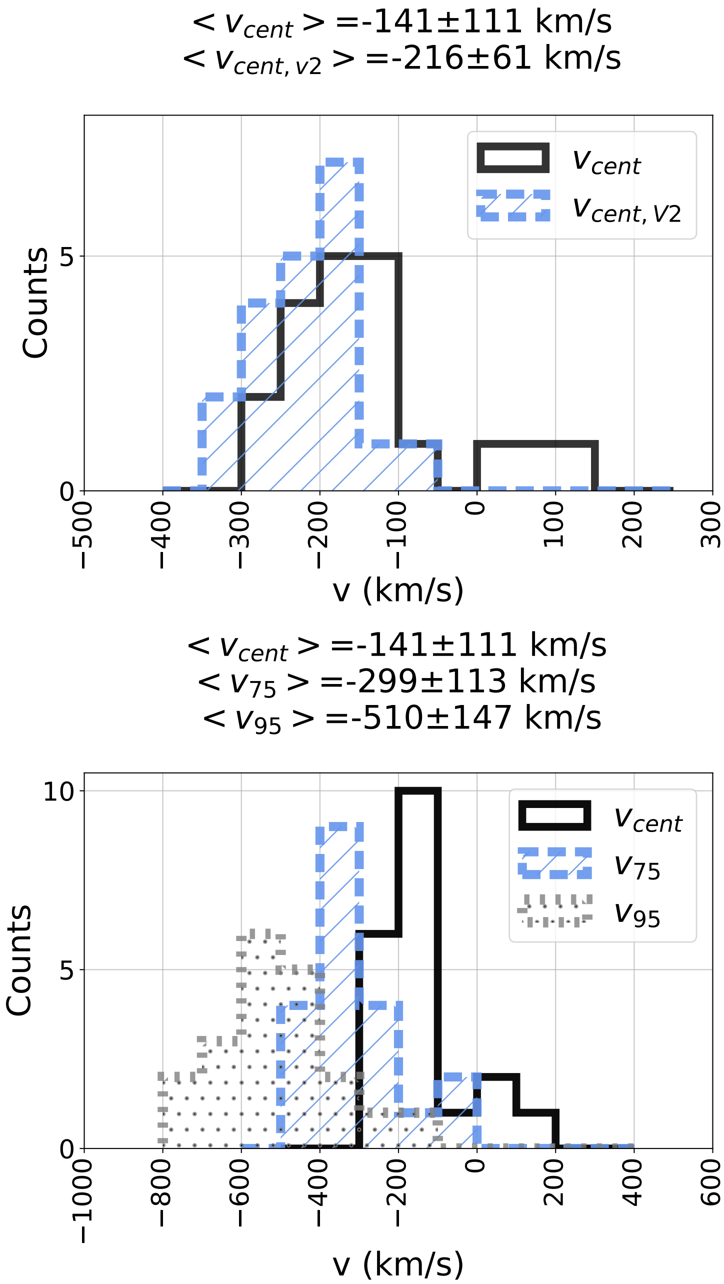

FIGURE 3.7. *Top:* Histogram of the absorption velocity centroid v_{cent} for the sample measured by considering (i) the entire profile (solid), and (ii) only $v < 0$ km s^{-1} (dashed). *Bottom:* Histogram of v_{cent}, v_{75}, and v_{95} values for all targets in the sample. The mean and sample standard deviation of each quantity are given above the plots. The aggregate sample shows typical centroids blueshifted by $\sim$150 km s^{-1} relative to the systemic velocity, with significant absorption extending to outflow velocities of 300–500 km s^{-1} or more.

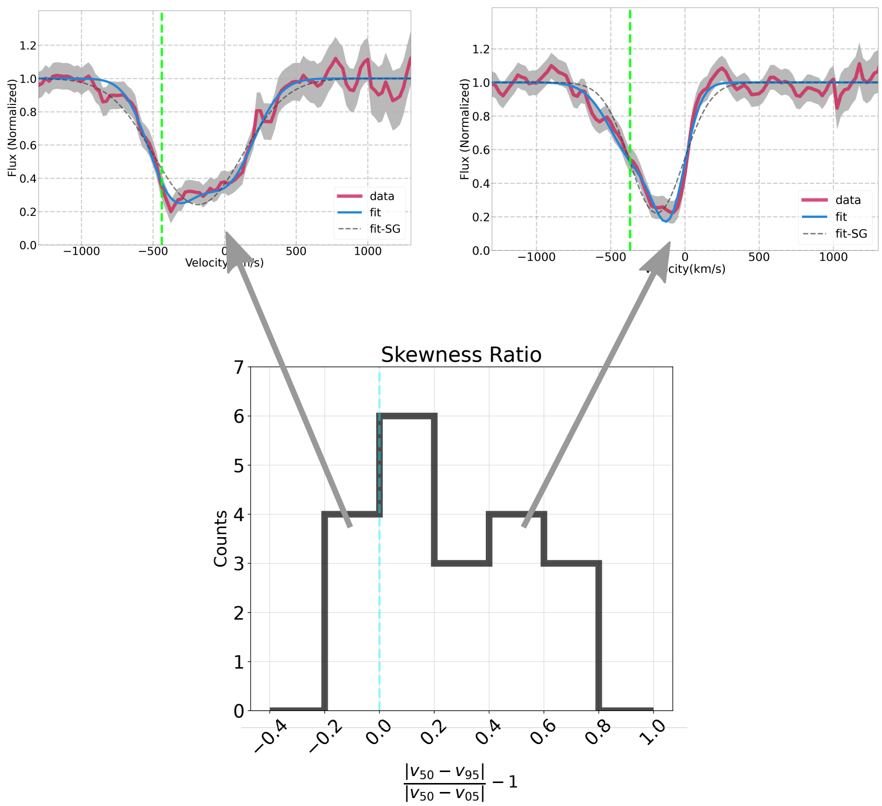

Figure 3.8. Histogram of skewness ratios. A value of 0 (denoted by cyan dashed line) indicates no skewness between the blue and redward absorptions, i.e., absorption which is symmetric about the centroid. A skewness value > 0 indicates that the slope of the blueward absorption is shallower than the redward absorption, while values < 0 correspond to shallower redward slopes. Examples of positively and negatively skewed profiles are shown at the top. Nearly all galaxies in the sample show positive skewness, indicating asymmetric profiles with a broad tail of blueshifted absorption from outflowing gas.

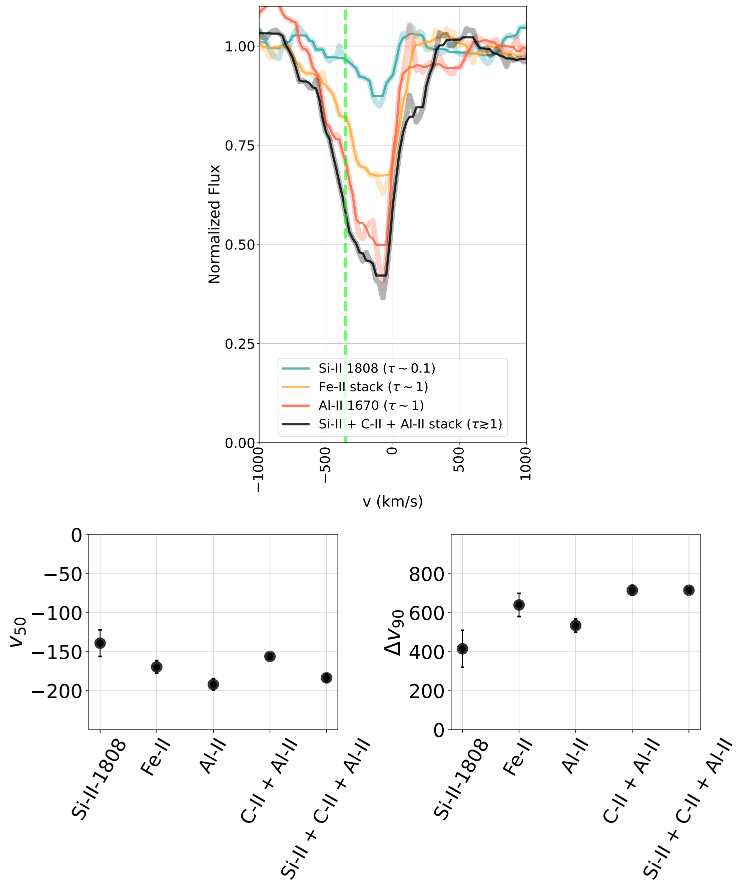

FIGURE 3.9. *Top:* Median stacked absorption profiles for lines with different optical depths (τ) : Si II λ1808 (green), Fe II λ1608 (orange), and Al II λ1670 (red). A joint stack of the strongest low ion transitions (Si II, C II, Al II) is shown in black. Bold lines are running medians of each profile whereas the lighter shade is the full-resolution data. Si II λ1808 probes optically thin gas ($\tau \ll 1$), while Fe II and Al II have apparent optical depths of order unity ($\tau \sim 1$) and the strong transitions have $\tau \gtrsim 1$. The green vertical line is the median $v_{75,V2} = -327$ km s^{-1} measured from this work shown as a reference point. Notably the profiles of different τ all show similar mean blueshifted velocities, indicating that velocity centroids are robust to the choice of absorption lines. *Bottom:* v_{50} (centroid) and Δv_{90} measurements for the four profiles shown above, with τ increasing toward the right. We determine v_{50} and Δv_{90} from a DG fit to the absorption profiles. The main results are unchanged if SG fits are used. v_{cent} is approximately constant across the range of optical depths, as can be seen visually in the top panel. In contrast, Δv_{90} is smaller for the optically thin lines which trace the bulk of the total column density, although the difference is comparable to the uncertainty. This indicates that high-velocity absorption seen in the strongest transitions is likely a small fraction of the total outflow mass.

of 630 $\mathrm{km\,s^{-1}}$. These galaxy values are ~ 6 times greater than those observed in the quasar absorption samples, falling near and beyond the largest values seen toward quasars. However, we caution that there are two main caveats in this comparison: (1) Δv_{90} for the low ions is likely overestimated by ~ 250 $\mathrm{km\,s^{-1}}$ compared to the optically thin lines (Section 3.5.3.1), and does not separate the systemic interstellar gas from outflowing and inflowing components. (2) Δv_{90} for the galaxies probes only one side of the galaxy (along our line-of-sight), such that it is smaller than would be observed toward a background source which would capture the highly redshifted outflowing gas on the far side of the galaxy. Despite these caveats, whose effects are in opposite directions, it is clear that the galaxy absorption profiles span velocity ranges comparable to the largest seen in quasar absorber systems at similar redshifts.

The large Δv_{90} values in our sample are likely driven by gas at smaller impact parameters than probed towards quasars. For most quasar absorbers, the host galaxy position and hence impact parameter is unfortunately unknown. This is particularly true for galaxies at higher redshifts since they are fainter and harder to identify from available imaging. Figure 3.11 (top panel) plots the Δv_{90} measurements from this work alongside those obtained from quasar sightlines of a DLA sample with known host location (Fynbo et al., 2013) as a function of impact parameter (b). As shown in the figure, these DLAs have lower Δv_{90} measurements even at modest impact parameters $b \sim 10$ kpc.

Dedicated surveys such as MAGIICAT (Nielsen et al., 2013a) and MEGAFLOW (Schroetter et al., 2016) have studied quasar absorption associated with known host galaxies. This provides information on trends with impact parameter, although the hosts are at lower redshifts than our sample. MAGIICAT galaxies have measurements of equivalent width (EW) of Mg II, a low-ion species with which we can directly compare. For a subset of our sample which has spectral coverage and good SNR for both Mg II and shorter-wavelength low ionization lines, we find that the Mg II profile closely traces the ISM absorption line profiles

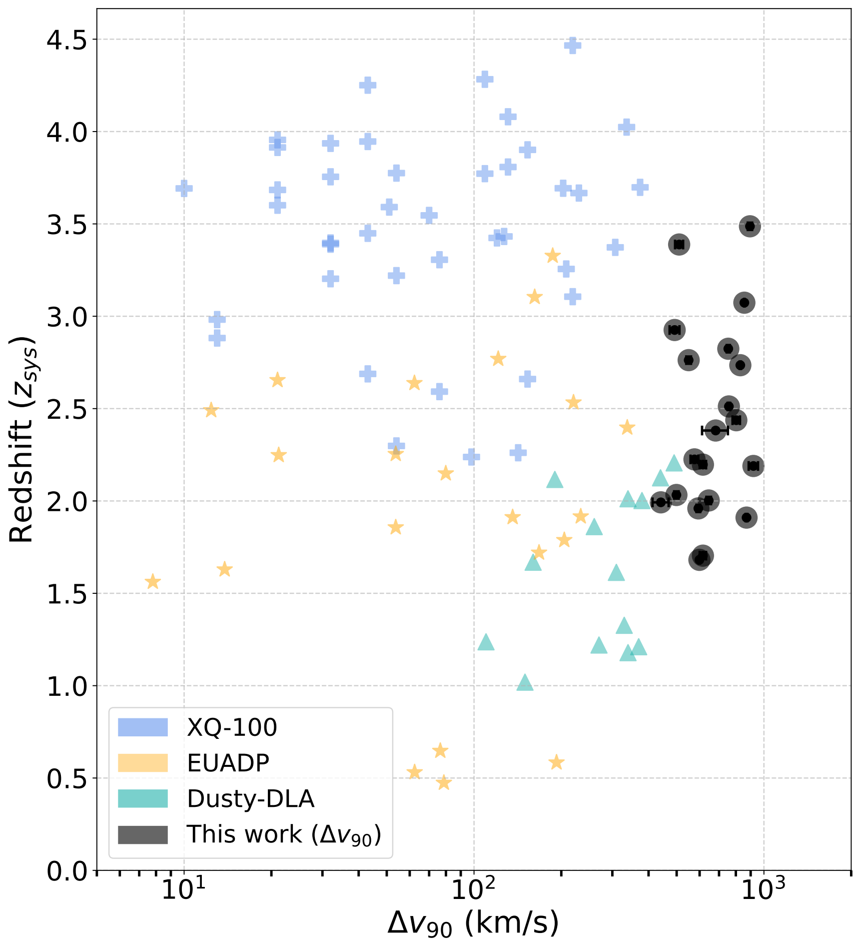

FIGURE 3.10. Comparison of Δv_{90} measured in this work with those obtained from surveys of background quasar sightlines: XQ-100 (Berg et al., 2016), EUADP (Quiret et al., 2016) and Dusty-DLAs (Ma et al., 2017). Compared to quasar absorber samples, this work probes gas at very low impact parameters ($b = 0$) and larger cross-sectional area, while sampling only the foreground region ($R = 0 \to \infty$ c.f. background quasars which probe $R = -\infty \to \infty$). Despite probing only half of the halo, the typical $\Delta v_{90} \simeq 600$ km s^{-1} for this work is considerably larger than for quasar absorber systems. Only the most extreme quasar systems have comparable Δv_{90} values. This suggests that the high-velocity gas which is ubiquitous in our down-the-barrel galaxy sightlines is located at small impact parameters which are extremely rare in quasar samples. DLAs at similar redshifts with known hosts have been found to probe impact parameters $b \lesssim 25$ kpc and $\Delta v_{90} \lesssim 350$ km s^{-1} (Figure 3.11; Fynbo et al., 2013).

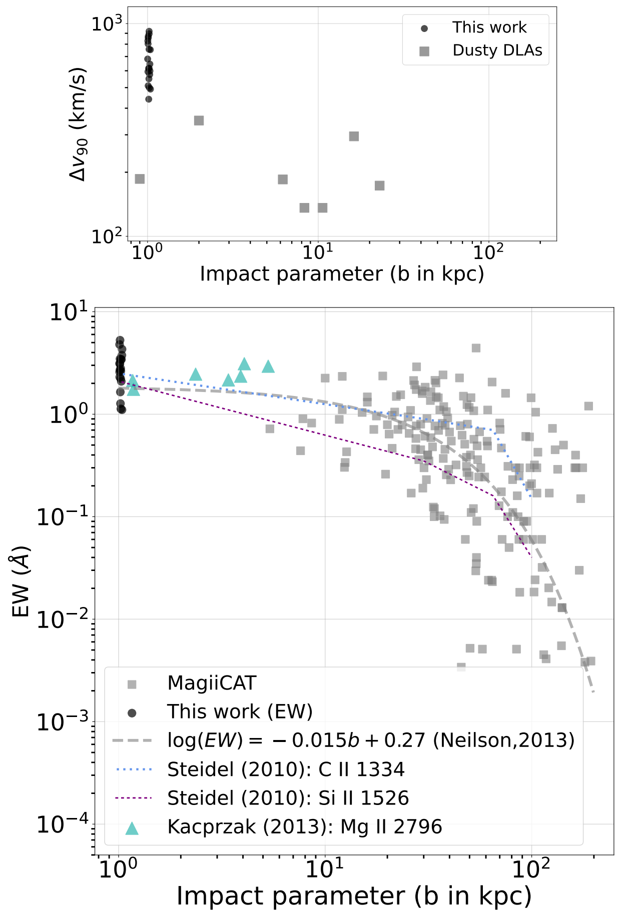

FIGURE 3.11. *Top*: Comparison of Δv_{90} as a function of impact parameter b for our sample (at b = 0 but plotted here at b = 1 kpc for clarity), and for DLAs at $z \sim 2$ probed by quasar sightlines (Fynbo et al., 2013). Even at impact parameters $\lesssim 25$ kpc, these DLAs have smaller Δv_{90} values than those obtained from "down-the-barrel" galaxy observations. *Bottom*: Comparison of absorption equivalent width (EW, in Å) as a function of impact parameter (b) for our sample and the MAGIICAT survey (Nielsen et al., 2013b) of Mg II absorbers in background quasar sightlines. The average EW clearly drops off by orders of magnitude as we move to higher impact parameters. The dashed line denotes the best fit to the MAGIICAT sample extrapolated to lower impact parameters, and it appears to be consistent with the EWs obtained from this work. We note that the large EW for our sample at small impact parameters is primarily due to the large velocity widths from outflows, whereas large EW for MAGIICAT galaxies might also arise from narrower absorption (smaller Δv_{90}) with higher covering fractions. EW values obtained from quasar sightlines at low impact parameters ($b \lesssim 5$kpc; Kacprzak et al., 2013) and from galaxy-galaxy pairs at $z \sim 2$ from Steidel et al. (2010) are shown in green points and dotted lines respectively, also showing a similar trend.

used in this work, including in the high-velocity wings. Therefore, we convert the low ion covering fraction to an expected EW of Mg II $\lambda 2796$ as follows:

$$EW(\text{in Å}) = \frac{W_{vel}}{c}\lambda_{2796} \tag{3.6}$$

where

$$W_{vel} = \sum C_f(v)\Delta v \tag{3.7}$$

and λ_{2796} corresponds to the rest frame wavelength of Mg II. This EW estimate assumes $\tau \gtrsim 1$ for Mg II $\lambda 2796$ absorption in our sample, which we expect based on the observed low ions. Here W_{vel} is effectively an equivalent width in units of velocity, calculated by summing the covering fraction profiles.

Figure 3.11 (bottom panel) compares the EW obtained for our sample, MAGIICAT (with typical $z \sim 0.4$), and stacks of $z \sim 2$ galaxy-galaxy pairs from Steidel et al. (2010). The galaxy-galaxy pairs are a useful comparison since they probe the cross-sectional area of a background galaxy, similar to down-the-barrel spectra. The galaxies from Steidel et al. (2010) have similar stellar mass and SFR as our sample, and those from MAGIICAT have similar stellar mass (Churchill et al., 2013). Our $z \sim 2$ sample at $b \simeq 0$ kpc spans EW = 1–5 Å, whereas MAGIICAT probes larger impact parameters and has a median EW of 0.43 Å. The width of absorption drops by orders of magnitude as b increases away from the galaxy. Extrapolating the trend line obtained from Nielsen et al. (2013b) for MAGIICAT ($EW \propto 10^{-0.015b}$) to lower impact parameters provides a good match to our sample average. This result is complimented by the galaxy-galaxy pairs which also show low ion EW decreasing similarly at higher impact parameters. We can see that the galaxy-pair trend line obtained for C II 1334 is a better match to the MAGIICAT sample, whereas Si II 1526 falls below this trend line. This may be due to lower optical depth of Si II 1526 compared to both C II 1334 and Mg II 2796.

Looking at the maximum EW obtained in the quasars, one can find some values which seem to have comparable EW to the high-z sample. These may be associated with orientation effects where the quasar sightline probes near the minor axis where we would expect outflows, or if the sightline incidentally passes through a fast moving cloud of gas. The background galaxy samples are likely to show smaller scatter because of the greater cross-sectional area probed. For our sample, the effective area varies for each source and is typically of order half the total cross-section of the galaxy (based on spectroscopic slit placement), or several square kpc (and $\gtrsim 1$ kpc^2 in all cases). Arc tomography studies probing $\sim$kpc^2 regions have indeed found smaller scatter than observed toward quasar sightlines (Mortensen et al., 2021a; Lopez et al., 2018a).

In summary, we find that the absorption width (EW and Δv_{90}) values obtained in this work at $b \approx 0$ are higher than those typically observed in quasar sightlines at larger impact parameter. However, extrapolating the trend in quasar absorption ($EW \propto 10^{-0.015b}$) to lower impact parameters offers reasonable agreement. We find similar agreement with measurements from galaxy-galaxy pairs at $z \sim 2$, suggesting a smooth decrease in absorption equivalent width with impact parameter with little redshift dependence. This indicates that the large EW and Δv_{90} in our down-the-barrel spectra arises from gas at small distances ($\lesssim 10$ kpc) from the host galaxy. Spatially resolved emission line studies mapping Fe II* and Mg II in one of the lensed galaxies in this work (RCSGA0327-G; Shaban et al., 2022) and other star-forming galaxies (e.g., Finley et al., 2017a; Burchett et al., 2021a) find similar spatial extent, further supporting a relatively small distance for the gas associated with down-the-barrel absorption.

3.5.4. Kinematics at intermediate and high ionization states. The warm ISM and CGM gas with $T \sim 10^4$ K is multiphase, with contributions from H I, H II, and a range of metal

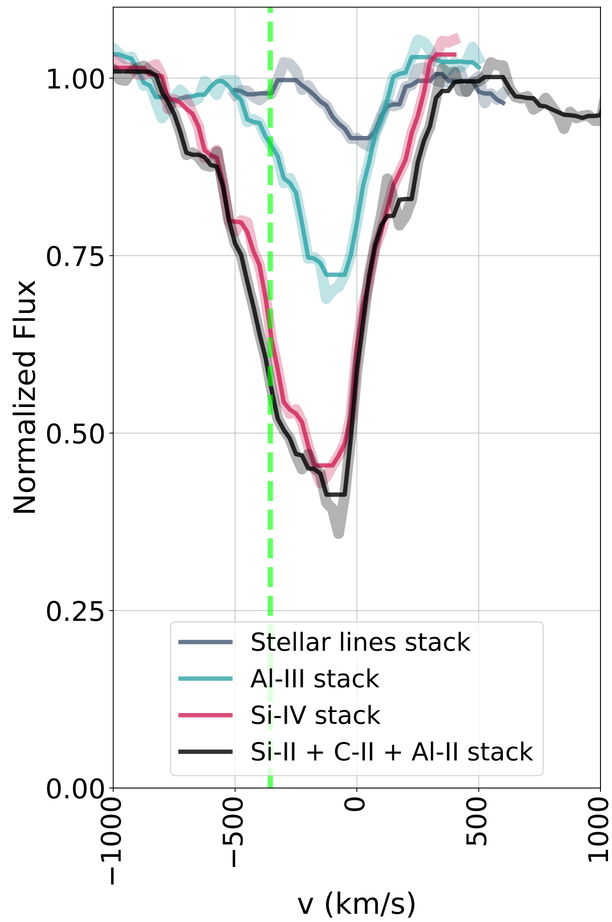

FIGURE 3.12. Median absorption profiles of gaseous species with different ionization potentials, compared to the stellar kinematics. We show stacks of stellar photospheric lines in dark blue, Al III $\lambda\lambda$1854,1862 in cyan, a stack of strong low ions (Si II, C II and Al II) in black, and Si IV $\lambda\lambda$1393,1402 in red. The green line shows the median $v_{75,V2} = -327$ km s^{-1} value measured from this work, for reference. Stellar absorption kinematics show a median $v = 0$ as expected, while the ISM profiles are clearly blueshifted due to prominent outflows (associated with the baryon cycle schematic illustrated in Figure 3.4). The low ions, Al III and Si IV all show nearly identical kinematics suggesting that these phases are co-spatial and powered by the same outflow mechanism. We note that the Al III transitions appear to have moderate optical depth $\tau \lesssim 1$, while other interstellar absorption profiles appear to be optically thick and thus trace the gas covering fraction.

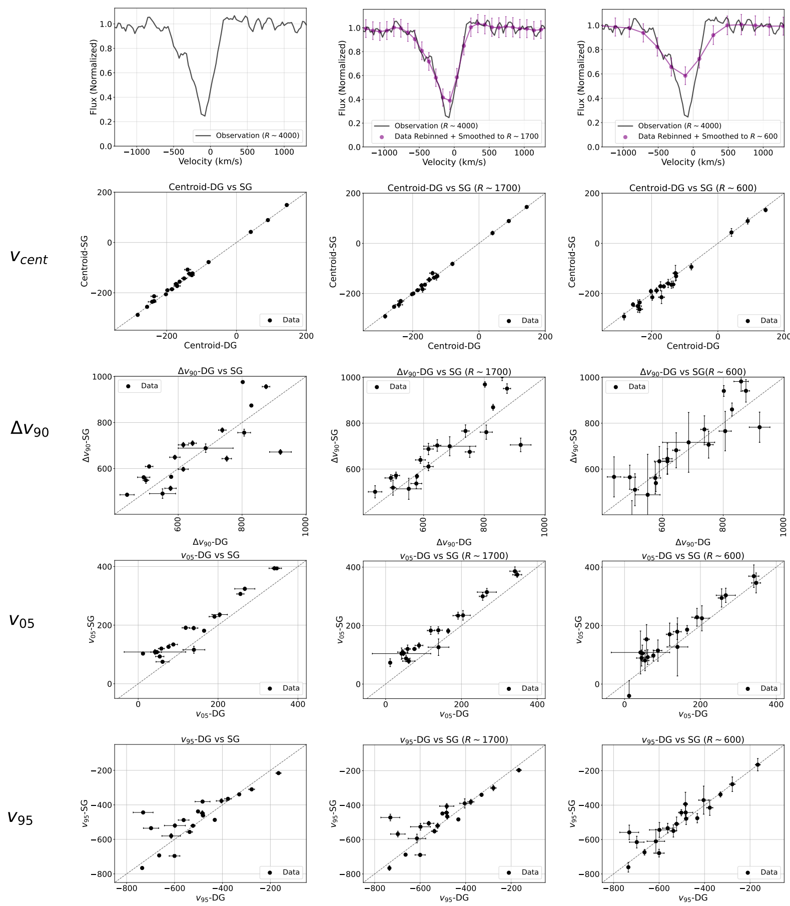

FIGURE 3.13. *Top panels*: Example of an absorption profile used for the SG fitting at different spectral resolution. The gray line denotes the profile in its native resolution, whereas purple lines show the profile after smoothing and rebinning. *Lower panels*: v_{cent}, Δv_{90}, v_{05} and v_{95} values obtained from double Gaussian (DG) fits compared with single Gaussian (SG) fits to the covering fraction profiles, at (i) the native resolution ($R \gtrsim 4000$; *Left*), (ii) $R \sim 1700$ (*Middle*) and (iii) $R \sim 600$ (*Right*). We describe the method used to transform our observed data to lower spectral resolution in Section 3.5.5.2. The SG fits are representative of the information content for lower spectral resolution data, whereas DG fits accurately capture the full velocity structure resolved in our sample. The black dashed line in each panel represents one-to-one correspondence between the two fits. The centroid velocity measurements of the SG fits agree well with those obtained from the DG, whereas the Δv_{90}, v_{05} and v_{95} measurements show larger scatter around the average linear trend. Some metrics are clearly biased in the SG fits, most clearly seen for v_{05} where the SG value is systematically larger. This scatter and bias is explained by asymmetry in the observed profiles, with most galaxies having a skewness ratio > 0 (see Section 3.5.2 for discussion), which the SG is unable to capture. Specifically, galaxies with higher skewness ratios are systematically more biased (see Figure 3.14).

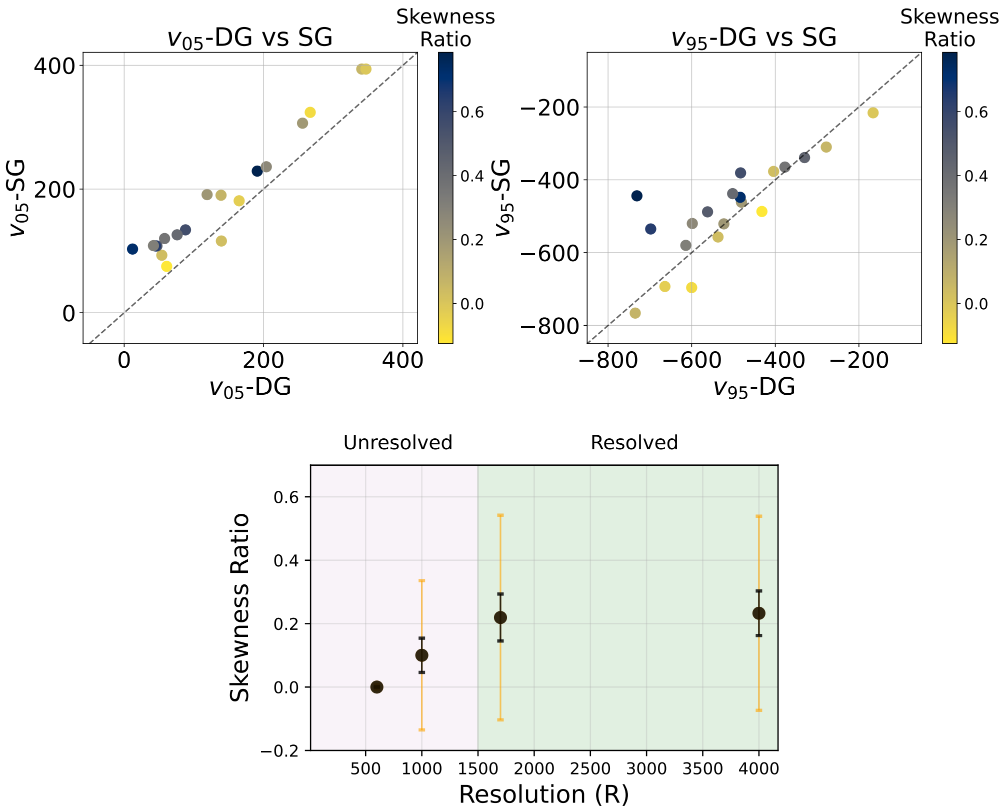

FIGURE 3.14. *Top:* v_{05} and v_{95} values obtained from double Gaussian (DG) fits compared with single Gaussian (SG) fits to the covering fraction profiles at $R \gtrsim 4000$, color coded by Skewness Ratio of the absorption profile (Section 3.5.2). v_{05} measurements tracing the redshifted velocities show a clear bias, whereas v_{95} which probes the high velocity outflowing gas has a lower bias but a larger scatter of 85 $\mathrm{km\,s^{-1}}$ (see Table 3.3). Absorption profiles which have higher Skewness Ratio values (e.g., J0004, RCSGA0327-G, CSWA128) show a clear bias in both v_{05} and v_{95} measurements. This suggests that velocity metrics other than the centroid (v_{50}) are not captured by symmetric fitting profiles, and thus are largely unreliable at low spectral resolution. *Bottom:* Average Skewness Ratio of absorption profiles in the sample, fit with a Double Gaussian (DG) after smoothing to different spectral resolution (R). The sample standard deviation and uncertainty in the mean are denoted in black and orange error bars respectively. At $R < 600$, the profiles uniformly appear symmetric (Skewness Ratio = 0) even when fit with a DG. We find that $R \gtrsim 1700$ is essential to recover the shape of the velocity profiles (e.g., skewness) and reduce biases that might be introduced due to lower resolution. This region is labeled as 'Resolved' indicating where the intrinsic profile skewness in our sample is recovered with a DG fit. This threshold corresponds to a FWHM spectral resolution of 4× smaller than the Δv_{90} line width, to adequately resolve the asymmetry.

thought to predominantly trace H I. Here we briefly examine species of different ionization potential in order to assess whether the low-ion results are applicable to other phases.

We construct median stacks of Al III $\lambda\lambda1854$, 1862, and Si IV $\lambda\lambda1393$, 1402 absorption lines with the same methodology as in Section 3.5.3.1. These span ionization potentials from 1–3.3 Rydberg. Figure 3.12 compares the stacked velocity profiles of these species along with the stacked low ions used in previous sections. An equivalent stack of stellar photospheric lines (Si III $\lambda1294$, Si III $\lambda1417$, S V $\lambda1501$, and N IV $\lambda1718$) is also plotted to show the stellar velocity range, which likely reflects that of the systemic (as opposed to outflowing) gas. We confirm that the stellar absorption is symmetric about $v = 0$ as expected. The kinematic structure of Si IV and Al III is similar to the low ion stack, suggesting that these species exist co-spatially. We note that Al III is unsaturated with $\tau \lesssim 1$, as the $\lambda1854$ line is clearly stronger than $\lambda1862$, whereas the low ions appear optically thick.

In summary, all of these ions – which are typically associated with $\sim 10^4$ K gas – exhibit similar kinematics. Chisholm et al. (2018) have also analyzed O VI for one of the lensed galaxies in this sample (CSWA38), and find that this hotter O VI phase is likely also co-spatial with the low ions, although with a different column density profile. We conclude from the similar absorption profiles that the various ions associated with $\sim 10^4$ K gas are likely co-spatial, tracing the same outflows.

3.5.5. Implications for low spectral resolution surveys. In this section we assess the extent to which lower-resolution spectra can accurately capture the kinematics of outflowing gas. There have been several large surveys of galaxies at $z \gtrsim 2$ which have characterized ISM absorption at lower spectral resolution $R \lesssim 1000$ (e.g., Shapley et al., 2003; Vanzella et al., 2009; Steidel et al., 2010; Weldon et al., 2022). In order to examine which kinematic properties can be reliably obtained with such data, we consider two scenarios below. First, we examine single Gaussian fits to the absorption profiles, which represents an idealized case. We then perform an equivalent analysis after smoothing and rebinning the data to mimic lower resolution surveys, with potentially detrimental effects from blending of adjacent spectral features. In practice, such blending may also affect the continuum normalization which would result in larger biases (i.e., worse performance) than the idealized case we consider herein

We note that the data used as the basis of comparison in this section has finite resolution $R \sim 4000$. Given the line widths, correcting for the instrument line spread function (LSF) has a small effect: Δv_{90} decreases by 8 $\mathrm{km\,s^{-1}}$ on average and v_{50} remains unchanged. As this is a small difference relative to the uncertainties, we report measurements directly from the $R \sim 4000$ spectra without correcting for the LSF. The true intrinsic line widths are thus $\sim$1% smaller than these reported values.

3.5.5.1. *Single Gaussian fit to $R \sim 4000$ data.* Single Gaussian (SG) fits to the ISM absorption profiles are described in Section 3.4.2 along with the resulting velocity metrics. Figure 3.13 (left panel) compares the quantities Δv_{90}, v_{cent}, v_{05}, and v_{95} obtained from the SG and Double Gaussian (DG) fits, both at the native $R \sim 4000$ spectral resolution. The mean offset and scatter between SG and DG fits for each metric are listed in Table 3.3. Velocity centroids show excellent agreement, with a mean offset $\langle v_{cent,DG} - v_{cent,SG} \rangle$ of only -4 ± 1 $\mathrm{km\,s^{-1}}$ and sample standard deviation of 10 $\mathrm{km\,s^{-1}}$.

The limitations of SG fits (and of low resolution spectra) are nonetheless apparent in higher-order velocity measurements. The v_{05} velocity in Figure 3.13 shows a clear bias (mean offset of -52 ± 4 $\mathrm{km\,s^{-1}}$) and substantial scatter (indicating error for individual objects) with SG fits. This bias is also evident in v_{95} and Δv_{90} which have a smaller mean offset but larger scatter. The bias is a consequence of the intrinsic asymmetry in observed line profiles which is not captured by a SG fit; a single Gaussian profile cannot recover the skewness (Section 3.5.2). Consequently we also find that absorption profiles with higher Skewness Ratios have larger biases in SG fits (Figure 3.14, *Top*).

These results demonstrate that asymmetric fitting profiles are essential to recover the covering fraction, skewness, and higher-order velocity measurements of absorption profiles. While a symmetric SG profile is able to recover accurate velocity centroids, the quantities describing both the blue- and red-shifted velocity extremes (such as v_{95} and v_{05}) are subject to large scatter and systematic biases. Consequently the spectral resolution must be sufficiently high to distinguish the asymmetric profile shapes.

3.5.5.2. *Profile fits at lower spectral resolution.* We now consider the quantitative effects of fitting to data of lower spectral resolution, where the intrinsic asymmetry of absorption profiles is less apparent. We smooth the absorption profiles to a spectral resolution of $R \simeq$ 1700, 1000, and 600 via convolution with a Gaussian kernel (of σ_{smooth} = 75, 125, and 200 $\mathrm{km\,s^{-1}}$ respectively). The smoothed spectra are also rebinned to σ_{smooth} per spectral pixel. The set of R is chosen to span an illustrative range, with the lowest resolution being comparable to large $z \gtrsim 2$ galaxy samples observed with Keck/LRIS and VLT/FORS2.

Figure 3.13 (top row) shows the rebinned and smoothed absorption profile of an example target spanning the range of resolutions considered here. We fit the smoothed and rebinned data with a SG and DG profile following the same methods as for $R \sim 4000$ (Section 3.4.2). Parameters from the SG and DG fits are then corrected for the effect of smoothing (i.e., deconvolved from the smoothing kernel). Mathematically this can be expressed as:

$$v_{deconv} = v_{fit} \tag{3.8}$$

$$\sigma_{deconv} = s \times \sigma_{fit} \tag{3.9}$$

TABLE 3.3. Performance of a single Gaussian (SG) fit to absorption profiles at different spectral resolutions compared to a double Gaussian (DG) fit at higher $R \sim 4000$. The mean offset (e.g., measured as $\langle v_{cent,DG} - v_{cent,SG} \rangle$ for v_{cent}) and the sample standard deviation (σ; e.g., of $v_{cent,DG} - v_{cent,SG}$) are given for the velocity metrics Δv_{90}, v_{cent}, v_{05}, and v_{95}. This sample σ represents the typical error introduced by not resolving deviations from a symmetric Gaussian profile, which varies for individual objects depending on their actual profile shape, such as skewness (Figure 3.14). The skewness ratio (discussed in Section 3.5.2) as measured by a single Gaussian fit is 0 by symmetry, whereas it varies between -0.2 to 1.0 for the double Gaussian.

Quantity	$R \sim 4000$		$R \sim 1700$		$R \sim 1000$		$R \sim 600$	
measured	Mean offset	Sample σ	Mean offset	Sample σ	Mean offset	Sample σ	Mean offset	Sample σ
	$\mathrm{km\,s^{-1}}$	$\mathrm{km\,s^{-1}}$	$\mathrm{km\,s^{-1}}$	$\mathrm{km\,s^{-1}}$	$\mathrm{km\,s^{-1}}$	$\mathrm{km\,s^{-1}}$	$\mathrm{km\,s^{-1}}$	$\mathrm{km\,s^{-1}}$
v_{cent}	-4±1	10	0±2	11	5±2	13	8±4	21
Δv_{90}	-20±6	97	-21±8	88	-16±10	84	-14±17	104
v_{05}	-52±4	40	-48±5	39	-40±6	41	-34±10	55
v_{95}	-28±5	85	-22±6	77	-20±6	69	-15±10	70

$$A_{deconv} = \frac{A_{fit}}{s} \tag{3.10}$$

$$\text{where } s = \sqrt{1 - \left(\frac{\sigma_{smooth}}{\sigma_{fit}}\right)^2}, \tag{3.11}$$

with σ_{fit} being the best-fit Gaussian velocity dispersion to the smoothed profile. The value of σ_{deconv} can then be compared directly to the velocity dispersion obtained at higher resolution. The median scaling factor to correct for instrument resolution is $s = 0.92$ at $R \sim 1700$ and $s = 0.64$ at $R \sim 600$. In other words, the intrinsic line profiles are broadened by a factor $\frac{1}{s} \simeq 1.6$ at $R \sim 600$, which can be reasonably corrected in most cases. Velocity metrics are measured from this scaled velocity profile using the same methods described earlier (Figure 3.6). We note that at $R \sim 600$, the intrinsic Δv_{90} absorption profile widths of our targets are sampled with only $\lesssim$1.5 independent FWHM spectral elements, while objects in our sample with the smallest widths (e.g., J1527) are effectively unresolved. The middle and right columns in Figure 3.13 compare the v_{cent}, Δv_{90}, v_{05}, and v_{95} values obtained from the SG fits at different spectral resolution to the DG fits (at $R \sim 4000$). Table 3.3 summarizes the mean offset and sample standard deviation for each quantity at different R.

We find that the results of SG fits are generally unaffected by degraded spectral resolution, agreeing within 1σ of the SG fits to $R \sim 4000$ data (Section 3.5.5.1). This is expected since the spectral resolution has been corrected using precise knowledge of the smoothing kernel. We thus obtain approximately the same systematic bias and scatter in SG fits to lower-R data as for the case of $R \sim 4000$.

The performance of SG fits discussed here should be taken as an optimal scenario given the good signal-to-noise ratio (SNR) of the lensed galaxy sample. Typical survey data will have larger statistical uncertainty. SNR is not necessarily a limiting factor however, as line width metrics are limited by the intrinsic scatter found between DG and SG fits (e.g., $\sim$100 km s^{-1} scatter for Δv_{90} seen in Figure 3.13 and Table 3.3).

Our analysis has demonstrated that quantifying the asymmetric structure of absorption profiles is necessary for accurately measuring quantities such as the maximum outflow velocity (e.g., v_{95}). This in turn is crucial for establishing galaxy scaling relations with outflow velocity, and comparing to feedback models (as we discuss in Section 3.6). We now quantify the resolution needed to recover the full asymmetric covering fraction profile structure of our sample. We make use of the skewness ratio defined in Section 3.5.2 as a reliable measure of this asymmetry. A skewness ratio of 0 indicates a symmetric profile, whereas most of the galaxies in our sample (80%) have skewness ratios > 0. Figure 3.14 (*bottom*) plots the skewness ratio obtained by a DG fit to the rebinned and smoothed absorption profile at different resolution (R). We find that at $R \sim 1700$, the shape of the profile is largely recovered: the mean and sample standard deviation in skewness ratio is 0.22 ± 0.32 compared to 0.23 ± 0.30 for the $R \sim 4000$ data. We also find that at $R \sim 1700$, other velocity metrics ($\Delta v_{90}, v_{95}, v_{05}$) have a mean offset of ~ 0 km s^{-1} and a modest sample scatter of ~ 50 km s^{-1} cf. $R \sim 4000$ measurements.

However, at $R \sim 600$, the skewness ratios are uniformly near zero indicating that the diversity and asymmetry of absorption profile shapes is not recovered for any of our targets at such low resolution. At the intermediate $R \sim 1000$, the average recovered skewness ratio is approximately half that of the high-resolution data. Individual galaxies with narrower profiles will have worse results at degraded resolution. In other words, the R required to distinguish asymmetric structure depends on the profile width. In this case the threshold $R \gtrsim 1700$ corresponds to sampling the average Δv_{90} with $\simeq$4 independent FWHM resolution elements. For samples with different gas kinematics, the required resolution should scale as the inverse of the profile width (e.g., $R \propto 1/\Delta v_{90}$).

To summarize, lower resolution data are sufficient to recover v_{cent}, while higher resolution $R \gtrsim 1700$ is required to recover the full asymmetric covering fraction profile structure and outflow velocity metrics for our sample (e.g., $v_{95}, \Delta v_{90}$). The threshold R required for reliable results will vary with the intrinsic profile width, which effectively corresponds to the gas outflow velocity. This analysis also demonstrates that the well-resolved profile shapes of our sample (Figure 3.5) can provide guidance for trade studies of spectral resolution and SNR for future surveys, which may be optimized for different scientific goals.

3.6. Trends with Galaxy properties

In order to understand the feedback effects of galactic outflows, we seek to compare outflow properties with galaxy demographics such as stellar mass and star formation rate (SFR). We necessarily restrict this analysis to the subset of the lensed sample with suitable ancillary data. In particular, for accurate stellar population properties, we require photometry at infrared observed wavelengths, as well as a lens model to correct for magnification by the foreground deflector galaxy.

Out of the 20 targets in our sample, 12 have reliable stellar mass and SFR measurements (and one more has SFR only). Masses and SFRs for six CASSOWARY targets are reported by Mainali et al. (2023a), while measurements for other sources are compiled from the literature. Table 3.4 lists the adopted stellar mass, SFR, and lensing magnification (μ) values along with the original references. All stellar population parameters are scaled to the Chabrier (2003) IMF where necessary. The stellar masses span $\log \mathrm{M}_*/\mathrm{M}_\odot = 9.1 - 10.8$ and the SFRs range from 10–210 $\mathrm{M}_\odot\ \mathrm{yr}^{-1}$, which are typical of moderately massive star forming main-sequence galaxies at these redshifts (e.g., Speagle et al., 2014).

One caveat in comparing the galaxy properties to the outflow properties is that the inferred SFR and stellar mass are global galaxy properties, whereas outflows may vary across different star-forming clumps (e.g., Bordoloi et al., 2014). In the following section, we assume

Objid	$\log\left(\frac{M_*}{\mathrm{M}_\odot}\right)$	SFR	μ	Ref
CSWA2	$9.1^{+0.3}_{-0.3}$	32^{+23}_{-13}	8.4	A0
RCSGA0327-G	$9.80^{+0.05}_{-0.05}$	40^{+10}_{-10}	17.2 ± 1.4	A5,A6
CSWA38	$9.8^{+0.2}_{-0.2}$	$10^{+0.2}_{-0.2}$	7.5 ± 1.5	A7
8oclock	$9.90^{+0.12}_{-0.13}$	162^{+124}_{-95}	5 ± 1	A4
Horseshoe	$9.9^{+0.2}_{-0.3}$	210^{167}_{-167}	10.3 ± 5.0	A1
J1527	$9.9^{+0.3}_{-0.4}$	116^{+86}_{-60}	15	A0
CSWA128	$9.9^{+0.1}_{-0.1}$	11.69^{+2}_{-1}	10	A0
Clone	$10.1^{+0.2}_{-0.2}$	68^{+24}_{-44}	13.1 ± 0.7	A1
CSWA103	$10.4^{+0.1}_{-0.2}$	23^{+18}_{-7}	4.7	A0
CSWA19	$10.5^{+0.1}_{-0.1}$	27^{+10}_{-5}	6.5	A0
CosmicEye	$10.76^{+0.07}_{-0.08}$	$37.6^{+4.3}_{-4.3}$	3.69 ± 0.12	A2
CSWA40	$10.8^{+0.2}_{-0.2}$	169^{+146}_{-66}	3.2	A0
J1429	-	90	8.8	A3

TABLE 3.4. Stellar mass, SFR (in units of $\mathrm{M}_\odot$ yr^{-1}), and magnification values for the targets presented in this paper. All values have been scaled to a Chabrier (2003) IMF. References are as follows. A0: Mainali et al. (2023a) A1: Jones et al. (2013a) A2: Richard et al. (2011) A3: Marques-Chaves et al. (2017) A4: Dessauges-Zavadsky et al. (2011) A5: Wuyts et al. (2014c) A6: Wuyts et al. (2010) A7: Solimano et al. (2022)

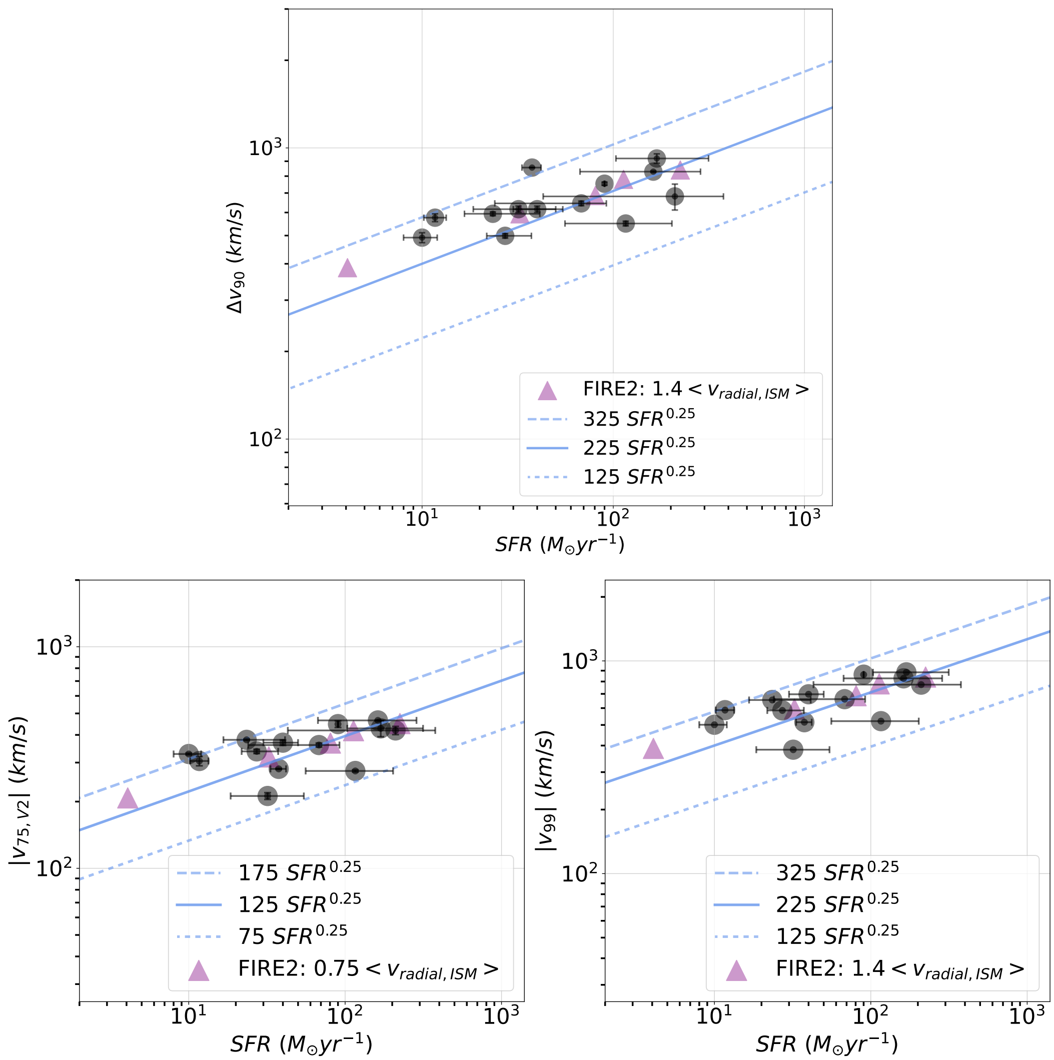

FIGURE 3.15. Plot of $|v_{75,V2}|$, $|v_{99}|$ and Δv_{90} versus SFR for the galaxies in our sample which have reliable SFR measurements in the literature. There is a clear correlation in that galaxies which have high SFR values also tend to have higher outflow velocities. Trend lines of the form $v \propto (SFR)^{0.25}$ corresponding to a momentum-driven wind scenario are shown in the figure for reference. We can see that $v \approx 225(SFR)^{0.25}$ km s^{-1} is able to capture the $|v_{99}|$ and Δv_{90} velocity measurements from this work, while $125(SFR)^{0.25}$ km s^{-1} better describes $|v_{75,V2}|$. This supports a positive correlation of outflow velocity and SFR which is close to the expected scaling relation for momentum-driven outflows. We note that this result holds for various ways of defining the outflow velocity as shown in each panel. Mass weighted radial velocity values measured in the ISM of FIRE-2 simulated galaxies at similar redshifts are denoted in purple, and are scaled by a constant factor to match the approximate average trend of the lensed galaxies.

that the global galaxy averaged outflow properties are sufficiently captured by our slit spectra, which probe several square kpc in the source plane.

3.6.1. Galaxies with high SFR also have high outflow velocities and absorption widths.

One of the key trends we want to explore is whether higher outflow velocities – traced by $v_{75,V2}$ and Δv_{90} values for example – correlate with higher SFR in the host galaxies. Such a correlation may be naturally expected since the outflows are driven by energy and momentum released by star formation. We consider a simple power law response of the following form: $v = v_0 SFR^{\alpha}$ and $\Delta v = \Delta v_0 SFR^{\alpha}$. Physically, the value for α is determined by the mode of feedback. For example, Murray et al. (2005) find that in galactic winds primarily driven through momentum injection from supernovae, the luminosity (L) scales with the galaxy velocity dispersion (σ) as $L \propto \sigma^4$ whereas an energy-driven wind would follow $L \propto \sigma^5$. They also find that for starburst galaxies at high-z, momentum driven winds are more favorable, as have other studies (e.g., Davé et al., 2011). Therefore, taking the SFR to be a tracer of luminosity and the outflow velocity to roughly scale linearly with the galaxy dispersion (e.g., Cicone et al., 2016), we might expect the kinematics of the outflowing gas also to scale as $SFR^{0.25}$ (i.e. $v \propto SFR^{0.25}$ and $\Delta v \propto SFR^{0.25}$).

Figure 3.15 plots the SFR versus the outflow velocity (v_{99}, $v_{75,V2}$, and Δv_{90} metrics) along with a power law scaling relation motivated by the momentum-driven wind scenario. As we can see from the figure, a power law fit is a reasonably good description of our measurements. Notably, we find a similar power law correlation for all three outflow velocity metrics, which primarily differ in normalization as expected. We overlay scaled mass-weighted radial outflow velocities (e.g., $v = 1.4 \langle v_{rad,ISM} \rangle$) obtained in FIRE-2 simulated galaxies at $z = 2 - 4$ (discussed further in Section 3.6.2), which also show reasonable agreement with this power law correlation. We find that among the different metrics tested, Δv_{90} correlates well with SFR, having a Spearman correlation coefficient of 0.7 and p-value = 0.007. The best fit power law between Δv_{90} and SFR for our sample is given by $(389 \pm 77)\, SFR^{0.13 \pm 0.05}$. The slope of the best-fit relation is somewhat shallower than expected for momentum-driven winds at $\sim 2\sigma$ significance, although this power-law slope is consistent with the relation between SFR and velocity FWHM (with $\alpha = 0.19$) found at low redshift by Xu et al. (2022a) for their sample which probes a larger dynamic range in SFR. We note that Xu et al. (2022a) also find a marginally steeper slope ($\alpha = 0.22$) in SFR versus outflow velocity, and find good overall agreement with a momentum driven wind scenario.

Various studies of galaxy outflow velocities and their scaling relations, spanning a wide range of redshift and galaxy properties, have found that the power law coefficient α (where $v \propto SFR^{\alpha}$) ranges from $\alpha \in (0.03-0.35)$ (e.g., Martin, 2005; Rupke et al., 2005; Weiner et al., 2009; Steidel et al., 2010; Martin et al., 2012; Erb et al., 2012; Bordoloi et al., 2014; Chisholm et al., 2015, 2016; Heckman & Borthakur, 2016; Sugahara et al., 2017). A key challenge is that different studies employ inhomogeneous data and analysis techniques, including the definition of outflow velocity (such as the maximum outflow velocity v_{max}, ISM velocity centroid v_{IS}, and others). Our analysis suggests that Δv_{90} is potentially of broad use for comparison, as it is more readily measured than quantities such as v_{99}, while it correlates well with other metrics and is reasonably robust to spectral resolution effects (Section 3.5.5).

3.6.2. Comparison to cosmological simulations. In this section, we focus on comparing our observational results with predictions for outflow properties obtained from recent cosmological simulations which incorporate stellar feedback. Such comparisons are a valuable test of feedback models used in these simulations, and we also highlight pathways which would be beneficial for future investigation. In particular, we compare observations with results from two sets of simulations: TNG50 (Nelson et al., 2019) and FIRE-2 (Pandya et al., 2021). These were chosen due to the availability of suitable outflow velocity metrics.

One challenge in comparing with simulations is that the radial distribution of gas responsible for the absorption in our sample is unknown. The observational data probe the total projected velocity of gas along the line-of-sight only on the near-side of a target galaxy. This

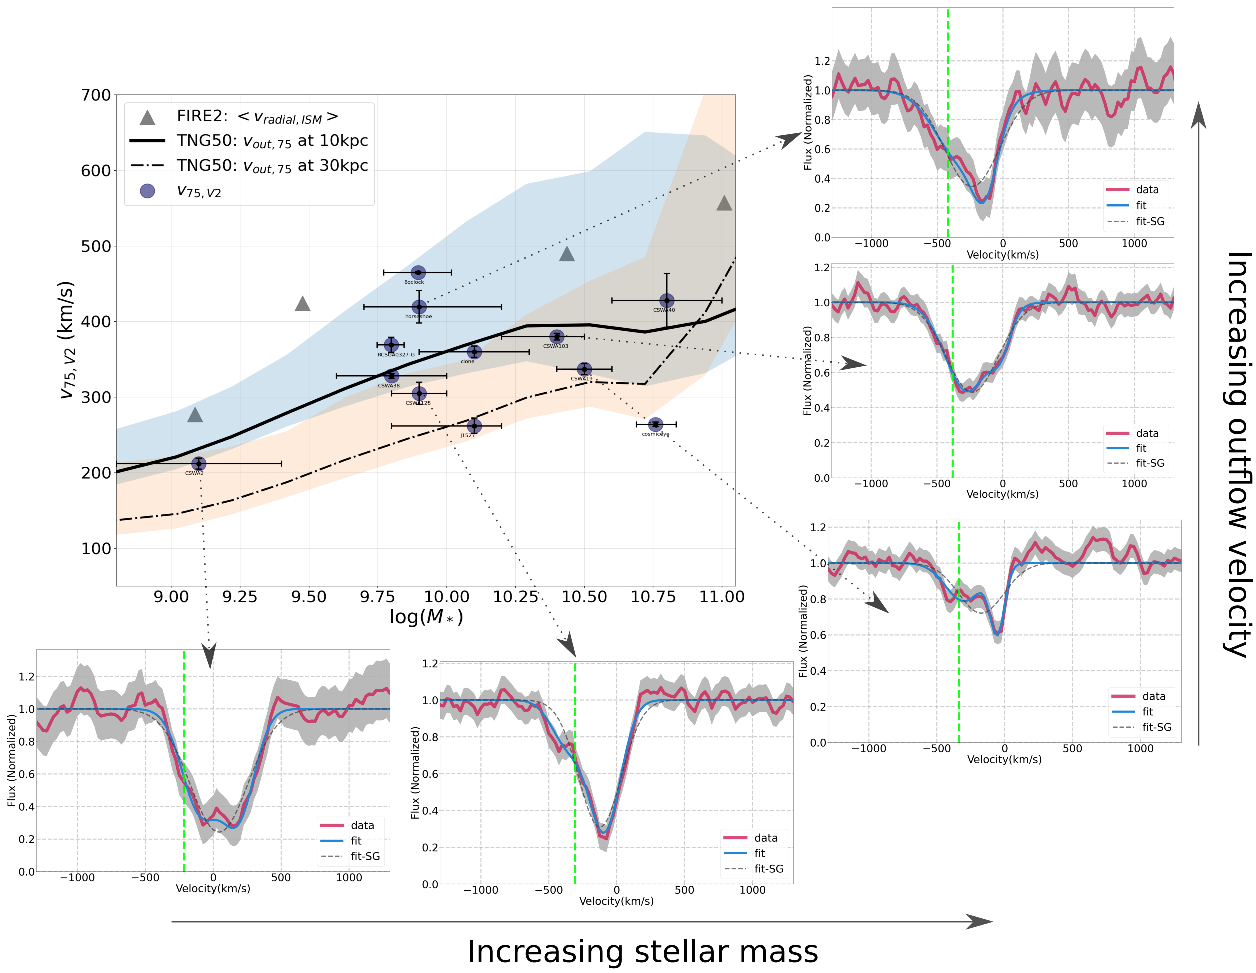

FIGURE 3.16. *Top left:* Comparison of measured outflow velocities for the lensed galaxies ($v_{75,V2}$: black points with error bars) with 75th percentile outflow velocities in the TNG50 simulation and mass-weighted radial velocity from FIRE-2. The solid and dotted black trend lines from TNG50 are the median values at r = 10 and 30 kpc, while shaded regions show the scatter (16–84 percentile range). The observations and TNG50 simulations appear generally comparable in terms of the overall trends and scatter (see Section 3.6.1). *Surrounding panels:* Velocity profiles for several of the lensed galaxies are shown to illustrate the velocity structure across the range of properties spanned by the sample (see Figure 3.5 for the full sample), with $v_{75,V2}$ values denoted by vertical green dashed lines. A key issue for comparison is that the observations include absorption from interstellar (systemic) and recycling gas, whereas TNG50 and other simulations can isolate purely outflowing material. We also note that the TNG50 results include highly ionized gas, although our analysis shows no significant difference in the velocity structure of low and intermediate ion species (see Section 3.5.4). The FIRE-2 radial outflow velocities are comparable to those obtained from the 75th percentile outflow velocities from TNG50 at 10kpc.

does not necessarily correspond with the metrics used in theoretical analysis or reported by simulations, where full 3-D spatial and velocity information is available. For example, TNG50 and FIRE-2 are able to examine gas outflow velocities as a function of radius from the host galaxy. To provide better context for comparison, we thus first consider the likely radial distribution and dynamical timescale of absorbing gas in the observed galaxy sample. Following Jones et al. (2018), we expect that the majority of absorption occurs within at most a few tens of kpc from the host galaxy. If we assume that outflowing gas starts at radius $r = 0$ and is driven at constant velocity, then its radial distance after a time t is

$$r = \frac{v}{-150\ \mathrm{km\,s^{-1}}} \frac{t}{100\ \mathrm{Myr}} \times 15\ \mathrm{kpc}. \tag{3.12}$$

Given typical velocities $\sim$150 $\mathrm{km\,s^{-1}}$ and $r \lesssim 50$ kpc (and quite possibly much smaller r) deduced from comparison to quasar sightlines, this implies the gas seen in absorption was launched $\lesssim 300$ Myr ago.

In Figure 3.16 we compare our measured outflow velocities with TNG50 simulated galaxies at $z = 2$ and FIRE-2 galaxies in the $z = 2{-}4$ bin, as a function of stellar mass. Specifically, we compare $v_{75,V2}$ values from this work with the 75 percentile mass-weighted velocities at different radii in TNG50 simulations. We expect these to be comparable, although they are not strictly identical measures. The observations are well bounded by TNG50 values for $r \lesssim 30$ kpc as shown in Figure 3.16, indicating reasonable agreement between the data and simulations. FIRE-2 galaxies, on the other hand, have measurements of mass-weighted radial velocity ($\langle v_{radial,ISM} \rangle$) for gas in the radius range $r = 0.1 - 0.2 r_{vir}$ with typical $r_{vir} \gtrsim 150$ kpc. Encouragingly, these values are also comparable to those seen in observations and those from TNG50 at $r = 10$ kpc. Therefore, the feedback prescriptions used in the TNG50 and FIRE-2 simulations yield outflow velocities comparable to those seen in observations. We discuss prospects for future work in this direction in Section 3.6.2.2

3.6.2.1. *Enrichment of the CGM/IGM via outflows.* A key question for galaxy formation is the amount of outflowing material which is able to escape a galaxy's gravitational potential, as opposed to remaining in the CGM and potentially recycling back into the galaxy, and how this varies with galaxy mass. To address this, we compare our measured outflow velocity profiles with estimated escape velocities of the sample.

The escape velocity v_{esc} is related to the rotational velocity of a galaxy ($v_{rot,max}$) and the virial radius (r_{vir}). In the case of an isolated galaxy with a truncated isothermal sphere mass distribution, the relation is

$$v_{esc}(r) = v_{rot,max}\sqrt{1 + \log\left(\frac{r_{vir}}{r}\right)} \tag{3.13}$$

for gas at radius r (Veilleux et al., 2005b). We estimate rotation velocities $v_{rot,max} \simeq 150 - 200$ km s^{-1} for the lensed sample based on the width of stellar photospheric features (Figure 3.12; Section 3.5.4), which is also supported by rotation curves of galaxies with similar redshift and stellar mass (e.g., Wisnioski et al., 2015; Förster Schreiber et al., 2018). Assuming $r \sim 0.1 - 0.2\, r_{vir}$, the escape velocity for these galaxies is 200-300 km s^{-1} from Equation 3.13. The mean 75% outflow velocity seen in our lens sample is $|v_{75}| \sim 300$ km s^{-1} (Section 3.5.1), suggesting that $\sim 25\%$ of the gas absorption profile has sufficient velocity to escape into the IGM. However, this simple analysis does not account for the interaction of outflows with the ambient CGM and the role of environment, such that the actual amount of gas exceeding the escape velocity may be smaller.

Figure 3.17 shows the escape velocity of gas at $0.1r_{vir}$ and at the halo radius (r_{vir}) obtained in the FIRE-2 simulations, compared to outflow velocities measured for the lensed sample (specifically $v_{99,V2}$, $v_{75,V2}$, and $v_{50,V2}$ corresponding to the 99, 75, and 50 percentiles of absorption blueward of systemic velocity). The $v_{99,V2}$ and $v_{75,V2}$ metrics trace the faster moving outflowing gas seen in absorption, whereas $v_{50,V2}$ traces the bulk motion of gas (Section 3.5.1). From the figure, it is clear that the $v_{99,V2}$ values are higher than those needed to escape the gravitational potential of the simulated galaxies and their halos, whereas the gas at $v_{75,V2}$ velocities would be able to escape only if the absorbing gas is located at large

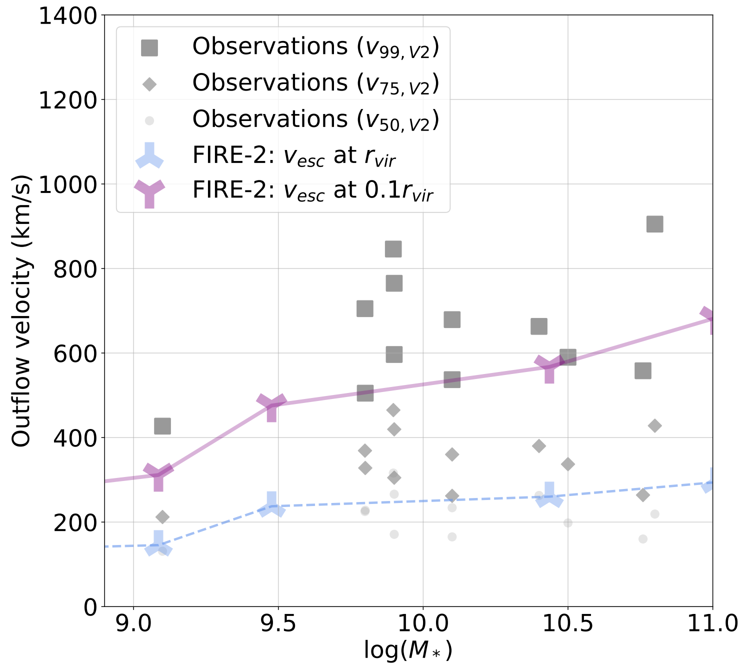

FIGURE 3.17. Comparison of the escape velocity v_{esc} obtained in FIRE-2 simulations with outflow velocities measured from the lensed sample. The v_{99} metric corresponds approximately to the largest velocity at which outflowing gas is detected in absorption. The simulations show a trend of $v_{\mathtt{esc}}$ increasing with mass, and decreasing with radius, as expected. The measured $v_{99,V2}$ values are comparable or larger than v_{esc} even at small radii ($\sim 0.1r_{vir}$), such that the highest velocity gas is able to escape the galaxies' gravitational potential. However $v_{50,V2}$ is typically below the escape velocity even at the virial radius, indicating that the majority of outflowing gas will remain gravitationally bound.

radii ($> 0.1 r_{vir}$). On the other hand, the mean outflow velocity centroid for the sample is $|v_{cent,V2}| = 188$ km s^{-1} which is below the escape velocity even at r_{vir}.

Based on this analysis, the majority of the $T \sim 10^4$ K outflowing gas, although moving at over a hundred km s^{-1}, appears to be bound within the halo and/or ISM of the galaxy (i.e. it is recycling gas; Figure 3.4). The fastest moving gas seen in absorption ($v > v_{75,V2}$) is capable of escaping into the CGM/IGM, enriching it with heavy metals, but is subject to deceleration from interactions with gas and dust along its path. This is consistent with results from Rudie et al. (2019), who find that 70% of the galaxies with detected metal absorption in the CGM also have unbounded metal-enriched gas capable of escaping the halo.

3.6.2.2. *Spatial distribution of the ISM gas.* Finally, we return our attention to the spatial distribution of the ISM gas around a galaxy. This is an essential quantity for determining outflow rates, mass loading factors, and whether outflowing gas will become unbound and escape into the IGM. However it is challenging to determine, as the observed absorption profiles do not directly depend on galactocentric radius. As discussed in Section 3.5.3.2, we can place constraints on the radius of outflowing gas seen in absorption based on comparison with background sightline samples at different impact parameters. The large absorption velocities and equivalent widths seen in our sample indicate the bulk of outflowing gas is at relatively small radius (conservatively within a few tens of kpc). Here we briefly consider prospects for future work.

Considering the encouraging comparison with simulations, a promising approach is to compare measured outflow velocity profiles with "mock spectra" generated from simulations where the spatial distribution of gas is known. This could be useful to assess the likely radial distribution of gas seen in absorption, and as a further test of feedback prescriptions used in simulations. Simulations can also be used to disentangle the outflowing, systemic, and recycling gas components and assess their relative contributions to the total absorption profile. Tools such as TRIDENT (Hummels et al., 2017) and FOGGIE (Peeples et al., 2019) are promising for such analyses. However, a challenge for such work is to self-consistently model the incident spectra and ionization state of the gas; in this case the host galaxy stellar emission may dominate over the extragalactic UV background. Finally, the technique of arc tomography (in which lensed arcs are used to spatially map CGM gas of lower-z galaxies in absorption) has recently proven to be highly effective (e.g., Lopez et al., 2018b, 2020; Mortensen et al., 2021a). While current studies are limited to $z < 1$, expanding to higher redshifts with multiple-arc systems is a promising future avenue. Strong lensing galaxy clusters such as the Hubble Frontier Fields (e.g., Mahler et al., 2018) may prove valuable for such analyses.

3.7. Summary and Conclusions

In this paper, we have used moderate resolution spectra ($R \gtrsim 4000$) to characterize the ISM and outflowing gas in a sample of 20 strongly lensed galaxies at $z = 1.5 - 3.5$ observed "down-the-barrel." We construct the covering fraction profile (C_f) of absorbing gas, and measure various metrics of the gas kinematics. In this work, we examine the outflow velocities (parameterized by v_{50}, v_{75}, etc.), width of absorption (Δv_{90}), skewness of absorption profiles, and optical depth (τ) of absorbing gas. We also explore the relations between outflowing gas kinematics and the host galaxy properties (e.g., M_* and SFR), and compare them with those obtained in cosmological simulations. We demonstrate the importance of having good spectral resolution in studies of outflowing gas by considering which of our results can be accurately recovered from lower resolution spectra ($R \lesssim 1000$), and which results would be biased. Below we summarize the main properties of the absorbing gas kinematics found from this work:

(1) The low ionization gas is characterized by a diverse range of covering fraction profiles (Figure 3.5; Sections 3.3, 3.4). The profiles are asymmetric, typically with a steep ingress at redshifted velocities and a shallow egress at blueshifted (outflowing)

velocities. 80% of the sample exhibits this skewness toward blueshifted velocities (Figure 3.8). A double Gaussian fit is sufficient to capture the structure of ISM absorption kinematics as measured at $R \simeq 4000$ and SNR $\simeq 10$ for the full sample.

(2) We observe ubiquitous outflows with a typical median velocity of $v_{50} \simeq -150\,\mathrm{km\,s^{-1}}$, with the extent of detected absorption reaching $3\times$ this median value in most cases (~ -500 $\mathrm{km\,s^{-1}}$; Section 3.5). The typical width of absorption profiles is $\Delta v_{90} \simeq 600$ $\mathrm{km\,s^{-1}}$, which is around 6 times larger than in typical DLA systems at similar redshifts probed by quasar spectra. Given the large absorption widths, it is likely that our down-the-barrel spectra are predominantly probing gas close to the center of the host galaxies (within a few tens of kpc or ~10% of the virial radius), whereas quasar absorption systems typically sample larger impact parameters. We note that our Δv_{90} values are measured for strong transitions which probe the gas covering fraction. Stacks of optically thin transitions suggest that the column density profile width is likely smaller ($\Delta v_{90} \sim 400$ $\mathrm{km\,s^{-1}}$; Figure 3.9), although still very large compared to quasar DLA systems.

(3) The lensed sample spans more than an order of magnitude in stellar mass and SFR, allowing us to examine scaling relations with outflow properties along the star forming main sequence at these redshifts (Section 3.6). We observe a positive correlation of outflow velocities and absorption widths (Δv_{90}) with both SFR and stellar mass, although the correlations are of modest significance within this sample. Among the metrics tested, Δv_{90} correlates well with SFR with a Spearman coefficient of 0.7 at 2.7σ significance (p-value = 0.007). We compare these measured trends in outflow velocity with the TNG50 and FIRE-2 cosmological simulations, and find reasonable agreement, which is encouraging for future work using simulations to help interpret outflow properties. The observed scaling relations are consistent at the 2σ level with expectations for momentum-driven outflows.

(4) To assess which kinematic properties can be recovered from low-resolution spectra, we compare results from the well-resolved velocity profiles with quantities derived from a single Gaussian fit (Figure 3.13; Section 3.5.5), both at $R \sim 4000$ and at degraded resolution (down to $R \sim 600$). A single Gaussian is appropriate for the information content of marginally-resolved spectra, and applying such fits at different R allow us to assess possible biases. We find that for single Gaussian fits, velocity centroids are largely reliable, having a mean difference $\langle v_{cent,DG} - v_{cent,SG} \rangle = 8 \pm 4\ \mathrm{km\,s^{-1}}$ and a scatter of only $\pm 21\ \mathrm{km\,s^{-1}}$ (1σ) at $R \sim 600$. Centroid measurements are nonetheless more precise and have lower scatter with increasing spectral resolution (Table 3.3). Velocity widths such as Δv_{90} are affected by large scatter with single Gaussian fits and require caution to avoid bias. Velocity metrics which are sensitive to the asymmetry in the absorption profile, such as v_{95} or other indicators of "maximum" outflow velocity, show a large scatter and clear bias even at $R \sim 4000$ when fit with a symmetric Gaussian profile, illustrating that such metrics are only reliable when the resolution and measurement method is sufficient to capture asymmetric structure. We find that $R \gtrsim 1700$ is needed to adequately capture the shape (e.g., skewness) of the absorption profiles in our sample. This corresponds to a FHWM resolution element $\lesssim \frac{\Delta v_{90}}{4}$. These results highlight the important role that spectral resolution plays in inferring key outflow properties.

This work represents the largest sample to date of well-resolved velocity profiles of gas outflows driven by star forming galaxies at cosmic noon ($z \sim 2$–3). We have robustly characterized the typical outflow kinematics and diversity among the galaxy population, with $\sim$10 independent resolution elements across the velocity profiles. While such analysis is currently practical only for galaxies which are highly magnified by gravitational lensing, this sample provides context for interpreting outflow properties from far larger existing samples of high-redshift galaxies with lower spectral resolution. For example, our findings that the

v_{50} and Δv_{90} metrics can be robustly recovered at low spectral resolution validate their use to characterize outflow scaling relations across larger samples and broader dynamic range than in this work. Moreover, these results can inform the optimal spectral resolution to be used for $z > 2$ galaxy surveys with upcoming 30-meter class extremely large telescopes (ELTs).

A promising avenue for future work is to explore spatially resolved outflow structure, along with the local conditions which launch strong galactic winds. As an immediate next step, some targets from this work are being followed up using the Keck Cosmic Web Imager to spatially map these ISM lines. Some will also be part of the galaxy evolution Key Science Program with KAPA (Keck All-sky Precision Adaptive Optics; Wizinowich et al., 2020) which will provide kinematic maps of the nebular emission at ∼100 parsec resolution, providing a detailed view of the star formation morphology and ionized gas kinematics. Combining spatially resolved galaxy structure with spatially+spectrally resolved outflow properties will provide greater insight into the physical process responsible for the feedback which regulates galaxy formation.

Appendix: Absorption profiles and best-fit parameters for the lensed sample

The appendix material corresponding to this chapter can be accessed in the digital edition of the respective journal publication.

CHAPTER 4

Spatially Resolved Galactic Winds at Cosmic Noon: Outflow Kinematics and Mass Loading in a Lensed Star-Forming Galaxy at $z = 1.87$

Submitted as Vasan G.C. et al. (2024) to The Astrophysical Journal (eprint available on arXiv: 2402.00942).

4.1. Abstract

We study the spatially resolved outflow properties of CSWA13, an intermediate mass ($M_* = 10^9$ M$_\odot$), gravitationally lensed star-forming galaxy at $z = 1.87$. We use Keck/KCWI to map outflows in multiple rest-frame ultraviolet ISM absorption lines, along with fluorescent Si II* emission, and nebular emission from C III] tracing the local systemic velocity. The spatial structure of outflow velocity mirrors that of the nebular kinematics, which we interpret to be a signature of a young galactic wind that is pressurizing the ISM of the galaxy but is yet to burst out. From the radial extent of Si II* emission, we estimate that the outflow is largely encapsulated within 3.5 kpc. We explore the geometry (e.g., patchiness) of the outflow by measuring the covering fraction at different velocities, finding that the maximum covering fraction is at velocities $v \simeq -150$ km s^{-1}. Using the outflow velocity (v_{out}), radius (R), column density (N), and solid angle (Ω) based on the covering fraction, we measure the mass loss rate $\log \dot{m}_{out}/(\mathrm{M}_\odot \mathrm{yr}^{-1}) = 1.73 \pm 0.23$ and mass loading factor $\log \eta = 0.04 \pm 0.34$ for the low-ionization outflowing gas in this galaxy. These values are relatively large and the bulk of the outflowing gas is moving with speeds less than the escape velocity of the galaxy halo, suggesting that the majority of outflowing mass will remain in the circumgalactic medium and/or recycle back into the galaxy. The results support a picture of high outflow

rates transporting mass and metals into the inner circumgalactic medium, providing the gas reservoir for future star formation.

4.2. Introduction

Galaxies self-regulate their growth across cosmic time through processes of gas outflows, inflows, and recycling (e.g., Davé et al., 2011; Lilly et al., 2013). As stars form and evolve in a galaxy, they inject energy and momentum into the surrounding gas through feedback processes such as stellar winds and supernovae, which in turn redistribute and enrich the interstellar and circumgalactic medium (ISM and CGM) with metals (e.g., Péroux & Howk, 2020). Surveys of high redshift $z \gtrsim 2$ star-forming galaxies (e.g., Shapley et al., 2003; Steidel et al., 2010) have detected near-ubiquitous outflow signatures in absorption, and the prevalence of a metal-enriched CGM has been established to $z \sim 6$ and beyond using background quasars (e.g., Becker et al., 2009). The primary mechanism attributed to the outward transport of gas and metals is galactic-scale outflows, i.e., gas being expelled across the entire galaxy.

Theoretical work has suggested that galactic outflows in intermediate mass galaxies at $z \sim 2$ have typical mass loading factors $\eta \sim 1$–10 (where $\eta = \frac{\dot{M}_{out}}{\mathrm{SFR}}$, $\dot{M}_{out}$ is the mass outflow rate, and SFR is the star formation rate; e.g., Muratov et al., 2015, 2017; Nelson et al., 2019; Pandya et al., 2021). Chemical evolution analysis supports similarly high mass loading (e.g., Sanders et al., 2021). However, directly measuring the gas mass loss rates (and mass loading factors) from galaxies has been challenging due to the need for spatial information. Characterizing galaxies at high redshifts requires sophisticated instruments and long integration times on 8–10m class telescopes. Even with such facilities, the ability to conduct spatially resolved observations is limited to the brightest galaxies at high redshifts. Gravitational lensing, whereby distant background galaxies are magnified by massive galaxies along the line of sight (e.g., Schneider et al., 1992; Narayan & Bartelmann, 1996; Treu, 2010, and references therein), offers a promising way to carry out resolved studies at high redshifts (e.g., Jones et al., 2013c; Bordoloi et al., 2014; Leethochawalit et al., 2016c; Spilker et al., 2020; Shaban et al., 2022). Dedicated searches for strong lens systems have resulted in substantial and growing samples (e.g., Stark et al., 2013; Tran et al., 2022b).

Tens of lens systems have now been followed up with deep slit spectroscopy to characterize their outflow properties (e.g., outflow velocity), taking advantage of the lensing magnification (Jones et al., 2018; Rigby et al., 2018; Vasan G. C. et al., 2023). Recently, integral field spectroscopic (IFS) observations have shown great promise in spatially resolving the outflows in high-redshift galaxies, including direct measurements of their spatial extent (Finley et al., 2017b; Burchett et al., 2021b; Shaban et al., 2022, 2023). In the case study presented in this paper, we seek to more robustly establish the mass loss rate ($\dot{M}_{out}$) and mass loading factor (η) of the low-ionization gas phase by combining IFS measurements of H I column density, outflow velocity, radial extent and the spatial structure of outflowing gas.

Our target is CSWA13, a bright star-forming galaxy discovered as part of the Cambridge And Sloan Survey Of Wide ARcs in the skY (CASSOWARY; Belokurov et al., 2009). The lensing nature of this system was spectroscopically confirmed by Stark et al. (2013) with redshifts of $z_d = 0.41$ for the foreground deflector galaxy and $z_s = 1.87$ for the background bright arc. An analysis of the stellar populations and UV nebular emission of this target is presented by Mainali et al. (2023b). Here, we use integral field spectroscopy from the Keck Cosmic Web Imager (KCWI; Morrissey et al., 2018) at the W. M. Keck Observatory to spatially map the outflowing gas from this galaxy.

This paper is organized as follows. Section 4.3 describes the KCWI observations and data reduction. Section 4.4 describes the lens model used to obtain accurate intrinsic properties of the source galaxy. The spectroscopic analysis methodology is described in Section 4.5. We discuss the results in Section 4.6 with conclusions in Section 4.7. Throughout this work, we assume a flat ΛCDM cosmology with $H_0 = 70$ km s^{-1} Mpc^{-1} and $\Omega_\mathrm{m} = 0.3$.

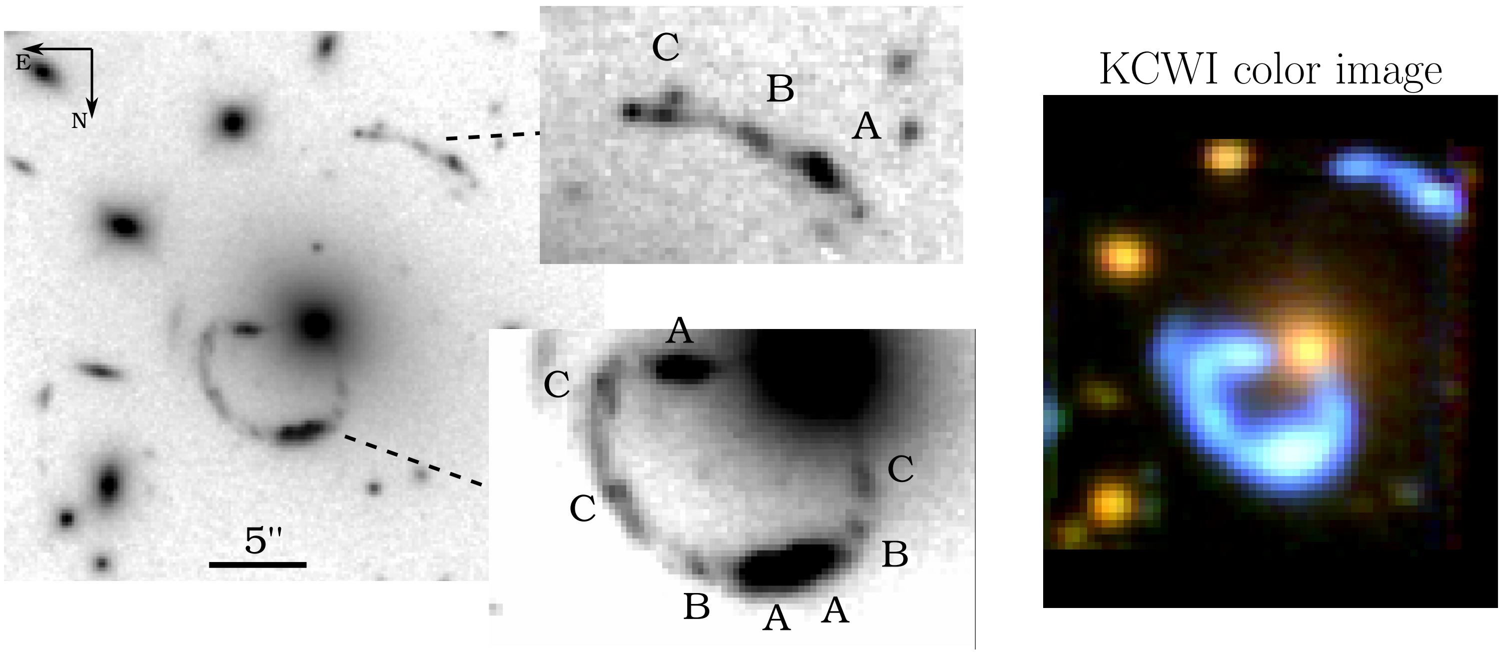

FIGURE 4.1. *Left:* HST near-infrared image of CSWA13 in the F140W filter, which probes rest-frame optical wavelengths (B/g-band) at $z \simeq 1.87$. East is left and North is down. The multiple images of the three bright distinct star-forming clumps of the galaxy are denoted as A, B and C. Identification of these multiply imaged regions is also confirmed by the KCWI spectra, which show distinct velocity profiles for each source-plane region (see Section 4.5.1). *Right:* KCWI color image centered on the main arc and deflector. The color channels were constructed by summing three broad wavelength regions (B: 3530–3930 Å, G: 4230–4630 Å, R: 4930–5330 Å) of the datacube. The lensed galaxy images form prominent bright blue arcs.

4.3. KCWI observations

We observed CSWA13 with KCWI on 06 April 2022 with the BL grating and Medium slicer configuration. The total wavelength coverage is 3229–5825 Å corresponding to 1127–2032 Å in the $z = 1.87$ galaxy rest frame, with spectral resolution of 2.4 Å FWHM ($R \simeq 1800$). The total exposure time was 2 hours, with six exposures of 1200 seconds each. Three exposures were taken at each of two orthogonal position angles (PA) of 0 and 90 degrees, both of which covered the entire lens system (Figure 4.1). The observing conditions ranged from clear to thin clouds with seeing of $\sim 1 - 1\farcs3$ FWHM. For sky subtraction, a nearby blank sky reference area was observed for 300 seconds with each set of PA exposures.

The data were reduced using the IDL version of the KCWI data reduction pipeline (KDERP-v1.0.2)[1]. We initially ran stages 1–4 of the pipeline, which perform bias, gain, and dark current corrections as well as cosmic ray rejection. We use the observed sky frame for sky subtraction in stage 5 of the pipeline, as there is insufficient blank sky background in the lensing field of CSWA13 itself. Flux calibrated and DAR (differential atmospheric refraction) corrected datacubes are obtained by running stages 6–8. Flux calibration was carried out using observations of the standard star BD26D2606 taken on 20 June 2020 using the same instrumental setup. The flux calibration is uncertain due to the non-photometric conditions, but the analysis and results of this paper do not require absolute flux calibration. We correct for any nonzero residual background in each wavelength slice in the following way: we consider the median flux within a 48 Å wavelength bin centered on each slice, mask out regions containing detected sources, fit the resultant background profile with a 2D first-order polynomial, and remove its contribution from the 2D slice. This is similar to the procedure described in Mortensen et al. (2021b), which was applied to a narrow wavelength range, while here we are interested in features that span the entire wavelength range of the KCWI spectra. The processed datacubes were then resampled to a common grid, aligned, averaged,

[1] `www.github.com/Keck-DataReductionPipelines/KcwiDRP`

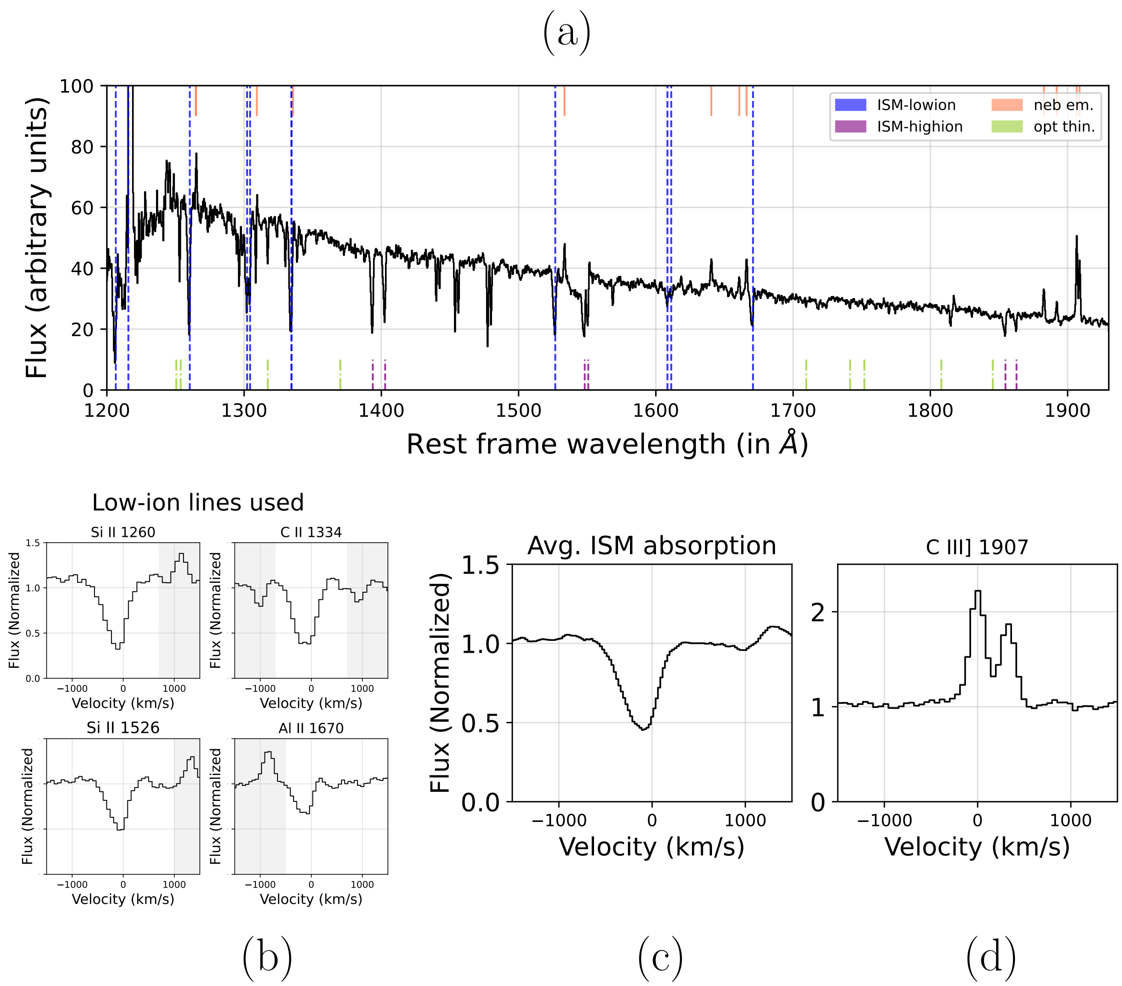

FIGURE 4.2. (a): Flux-calibrated KCWI spectrum of CSWA13 obtained by summing the flux from all spaxels covered by the arc. Prominent low ionization (e.g., Si II λ1260, λ1526, Fe II λ1608, Al II λ1670), high ionization (e.g., C IV $\lambda\lambda$1549,1551) and optically thin absorption lines (e.g., Si II λ1808, Ni II λ1317, λ1370, λ1709, λ1741) are marked in blue, purple, and green respectively. Fine structure emission from Si II* and nebular emission lines (e.g., C III]) are marked in coral. (b): Velocity profiles of various low-ion ISM absorption lines (from Si II, C II and Al II) used in this study. The gray shaded regions are masked and not used for ISM absorption analysis due to the presence of other significant features. For example, C II λ1334 is flanked by absorption lines from intervening systems which are masked out. Similarly, for Al II the region affected by nebular O III] λ1666 emission is also excluded from the absorption line analysis. Panel (c) shows the average ISM absorption profile obtained from combining the different low-ionization absorption profiles shown in panel (b). We use this combined absorption profile to probe the outflow kinematics across the arc (described in Section 4.5.2). The lower-right panel (d) shows strong C III] nebular emission which we use to trace the systemic velocity field across the arc.

and rebinned to a spatial pixel size of $0\farcs3\times0\farcs3$. We also converted the pipeline output from air to vacuum wavelengths. We use the 1σ standard deviation of flux from blank regions of the sky to estimate the uncertainty of our spectra.

Figure 4.1 shows a color composite image generated using HumVI (Marshall et al., 2015) from the reduced KCWI data cube by summing three broadband wavelength regions centered at 3730, 4430 and 5130 Å, each with a width of 400 Å. Multiple images of the arc are clearly identifiable by their blue color in the KCWI data, as well as prominent spectral features. Figure 4.1 *(Left panel)* also shows the Hubble Space Telescope (HST) near-infrared image obtained with WFC3-IR using the F140W filter. We identify three distinct star-forming complexes in the HST image (A, B, and C), which are multiply imaged across the entire arc. The integrated KCWI spectrum from summing spaxels containing the arc is shown in Figure 4.2(a) for the rest-frame wavelength range 1200–1930 Å. Metal absorption lines from three intervening galaxy systems at $z = 1.67, 1.69$, and 1.74 are also detected in the arc spectra, with prominent C IV and other features. We mask their contribution in all the analyses presented in the rest of this paper.

4.4. Lens Model

Gravitational lens modeling is necessary to reconstruct the galaxy's source plane morphology and intrinsic properties. We use the lens modeling software LENSTRONOMY (Birrer & Amara, 2018; Birrer et al., 2021) to build the lens model. We adopt the conjugate point modeling method which is commonly used in cluster-scale lens modeling (see Kneib & Natarajan, 2011, and references therein), with multiple image positions for the A, B, and C components of the source (Figure 4.1). Our KCWI data confirm the multiple image nature of components A, B, and C, for example via their distinct velocity structure and Lyα profiles. The lens model is optimized with the nested sampler DYNESTY (Speagle, 2019). Our lens model consists of an elliptical Navarro–Frenk–White (NFW; Navarro et al., 1996, 1997) halo profile for the lensing galaxy group, a double Chameleon profile for the stellar mass distribution of the central galaxy (Dutton et al., 2011), a point mass to account for the supermassive black hole (SMBH) at the center of this galaxy, eight singular isothermal ellipsoid profiles to model the eight brightest galaxies within 20 arcsec of the central galaxy (Kassiola & Kovner, 1993), a residual shear field and a residual flexion field.

The ellipticity of the projected NFW halo is parameterized in the convergence or the surface density and not in the potential (Oguri, 2021). The Chameleon profile is a combination of two non-singular isothermal ellipsoids that provide a good approximation to the Sérsic profile within $0.5-3R_{\rm eff}$, where $R_{\rm eff}$ is the half-light radius (Sérsic, 1968; Dutton et al., 2011). We find that a superposition of two Sérsic profiles is necessary to fit the light distribution of the central galaxy well (Claeskens et al., 2006; Suyu et al., 2013). We convert the best-fit parameters of the double Sérsic profile into the parameters of the double Chameleon profile. The pre-fitted parameters determining the angular and radial shape of the double Chameleon profile are fixed during the lens model optimization, and only the mass-to-light ratio is free. We fix the NFW halo mass to be $10^{14}M_{\odot}$ and impose a Gaussian prior on the concentration parameter ($c_{200} = 5.0 \pm 0.8$) following the results of Newman et al. (2015) for group-scale halos. The ellipticity and centroid of the NFW halo are fixed to be the same as the ellipticity and centroid, respectively, of the central galaxy's light distribution. Thus, the concentration is the only free parameter for the NFW profile. We impose a prior on the stellar mass to SMBH mass relation using the results from Li et al. (2023).

For the eight nearby galaxies included explicitly in the model, we use the PHOTUTILS package to measure aperture photometry and ellipticity. We fix the centroid and ellipticity of these galaxies and only allow the Einstein radii of each galaxy as free parameters. The residual shear (also called "external shear" in the literature) has two free parameters (i.e., the shear magnitude $\gamma_{\rm shear}$ and angle $\phi_{\rm shear}$) and the residual flexion field has four parameters. These residual shear and flexion fields account for both the "internal" angular structure of the central lensing galaxy that is not fully captured by the ellipticity parametrization of

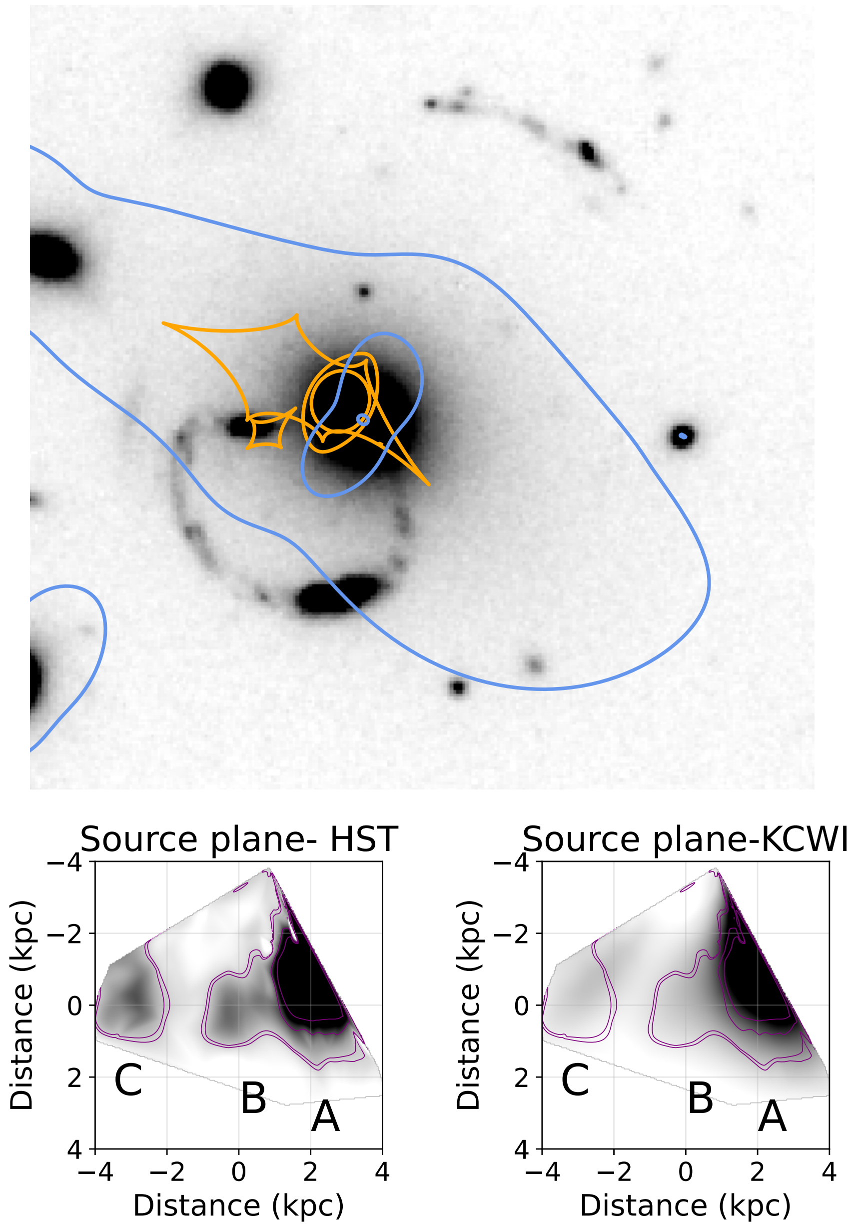

FIGURE 4.3. *Top*: Illustration of the best-fit lens model (described in Section 4.4). The critical and caustic curves are shown in blue and orange respectively. *Bottom*: Source plane reconstructed maps of the HST and KCWI continuum images, revealing the clumpy galaxy morphology of CSWA13. Purple lines denote the contours obtained from the source plane HST imaging. The three distinct star-forming regions A, B and C span a physical distance of ~ 8 kpc in the source plane.

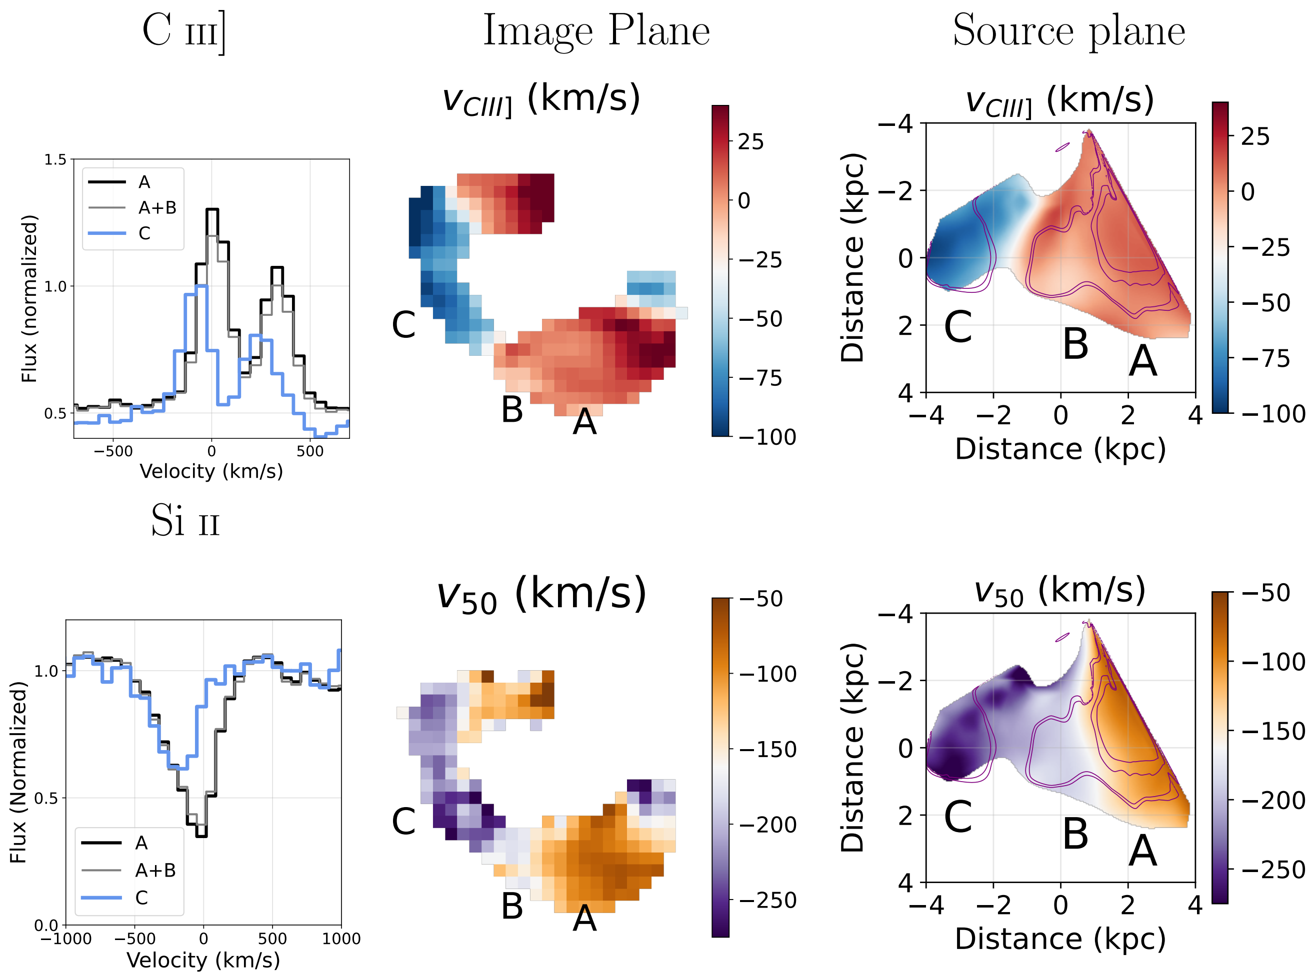

FIGURE 4.4. *Left*: Zoom-in of the C III]$\lambda\lambda 1907, 09$ emission and Si II $\lambda 1526$ ISM absorption line profile in regions A, A+B and C. *Middle*: Spatial maps of the centroid velocity of nebular C III] emission which traces the young stars, and the outflow velocity v_{50}. The outflow velocity maps are measured with respect to a systemic redshift of $z_s = 1.865947$ obtained from a galaxy-integrated spectrum. A similar velocity structure is apparent in both, suggesting that the outflowing gas is closely associated with the stars within each $\sim$ kpc spatial resolution element. The fold-symmetry of the velocity fields in the image plane is a result of the multiple image lensing configuration (Figures 4.1, 4.3). *Right*: Source plane maps of the centroid velocity of nebular C III] emission and outflow velocity v_{50}. The purple lines show contours from HST imaging, as in Figure 4.3. The ISM absorption profiles show significant spatial variation, and we can clearly see that the outflow velocity mirrors the systemic velocity structure along the major axis. We discuss the quantitative comparison of outflow and nebular velocity and its implications in Section 4.6.1.

the NFW and Chameleon profiles, and also "external" contributions from mass structure beyond the galaxies, which are explicitly accounted for in our model. The residual flexion field is also necessary to correctly reproduce the atypical image configuration of the multiple components in this system. This residual flexion is plausible since many group-member galaxies are not accounted for in our lens model except for the brightest eight. Our model has 17 free parameters in total, and there are 24 data points from 4 positions for each of the A, B, and C image sets. The lens model is optimized by minimizing the total distances between the mapped positions of each image set on the source plane.

Figure 4.3 illustrates the best-fit lens model with the caustic and critical curves. The lower panels of Figure 4.3 show the intrinsic galaxy morphology from HST and KCWI continuum images, reconstructed in the $z = 1.87$ source plane. The three star-forming complexes (A, B and C) span a physical distance of ~ 8 kpc in the source plane, using our fiducial cosmology. We find that the mean areal magnification of the southwestern counter-image of the galaxy is $|\mu| = 12 \pm 3$. We obtain a magnification corrected stellar mass and SFR of $\log_{10}(M_*/\mathrm{M}_\odot) = 9.00 \pm 0.32$ and $\log_{10}(\mathrm{SFR}/(\mathrm{M}_\odot\ \mathrm{yr}^{-1})) = 1.71 \pm 0.21$ after scaling the BEAGLE outputs presented in Mainali et al. (2023b).

4.4.1. Galaxy morphology. CSWA13 is a moderately dusty ($\tau_V \sim 0.5$) galaxy with specific SFR placing it above the star-forming main sequence at $z \sim 2$ (e.g., Whitaker et al., 2012). The spatially resolved kinematics from this work uniquely allows us to tie the observed morphology with the kinematics of the gas around the galaxy. The clumpy morphology of CSWA13 (Figure 4.3) resembles a tadpole or a chain galaxy (e.g., Cowie et al., 1995; van den Bergh et al., 1996; Elmegreen et al., 2005; Förster Schreiber et al., 2009; Elmegreen & Elmegreen, 2010) with a bright head (region A) and a tail (regions B and C). These tadpoles are common at $z \sim 2$, comprising $\sim$10% of the galaxies in the Hubble Ultra Deep Field (HUDF) and $\sim$44% if these tadpoles are indeed edge-on projections of the clump-cluster (Elmegreen et al., 2004) and chain galaxies. While the nebular kinematics for a handful

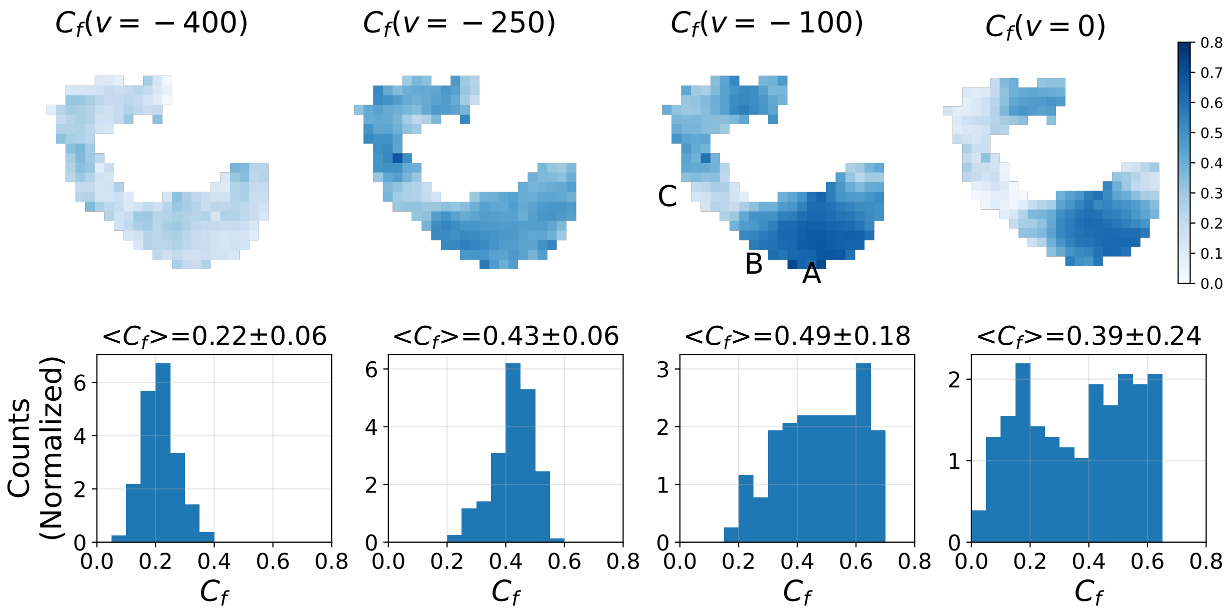

FIGURE 4.5. *Top*: Spatial map of the covering fraction C_f of the low ionization gas at outflow velocities of $v = -400$ to -100 km s^{-1} along with the systemic velocity ($v = 0$ km s^{-1}). *Bottom*: Distribution of C_f in each velocity bin. Gas at higher outflow velocities is more uniformly distributed across the entire galaxy, for example with a mean $C_f(v = -250) = 0.43$ with a relatively small scatter of 0.06. However, at systemic velocities it is more heterogenous, with a substantial variation in the covering fraction of gas with region A and B having $C_f \sim 0.6$ compared to $C_f \sim 0.1$ in region C. We also find that regions A+B have blueshifted Lyα absorption and redshifted emission (Figure 4.6) whereas region C has broad Lyα emission extending to blueshifted velocities. Therefore, regions A+B likely have a higher column density of slow-moving gas at systemic velocities seen in metal absorption transitions as well as Lyα, with a paucity of gas at $v \sim 0$ in region C.

of tadpole galaxies at $z \sim 2$ were measured as part of the SINS survey (Förster Schreiber et al., 2009) and show similar nebular kinematics as CSWA13, this paper is the first study to simultaneously map nebular and outflow kinematics in a galaxy of this type.

4.5. Spatially resolved ISM gas kinematics

Our overall aim is to quantify the key properties of outflows in CSWA13 such as the spatial extent, mass loss rate, and mass loading factor. This in turn requires measuring

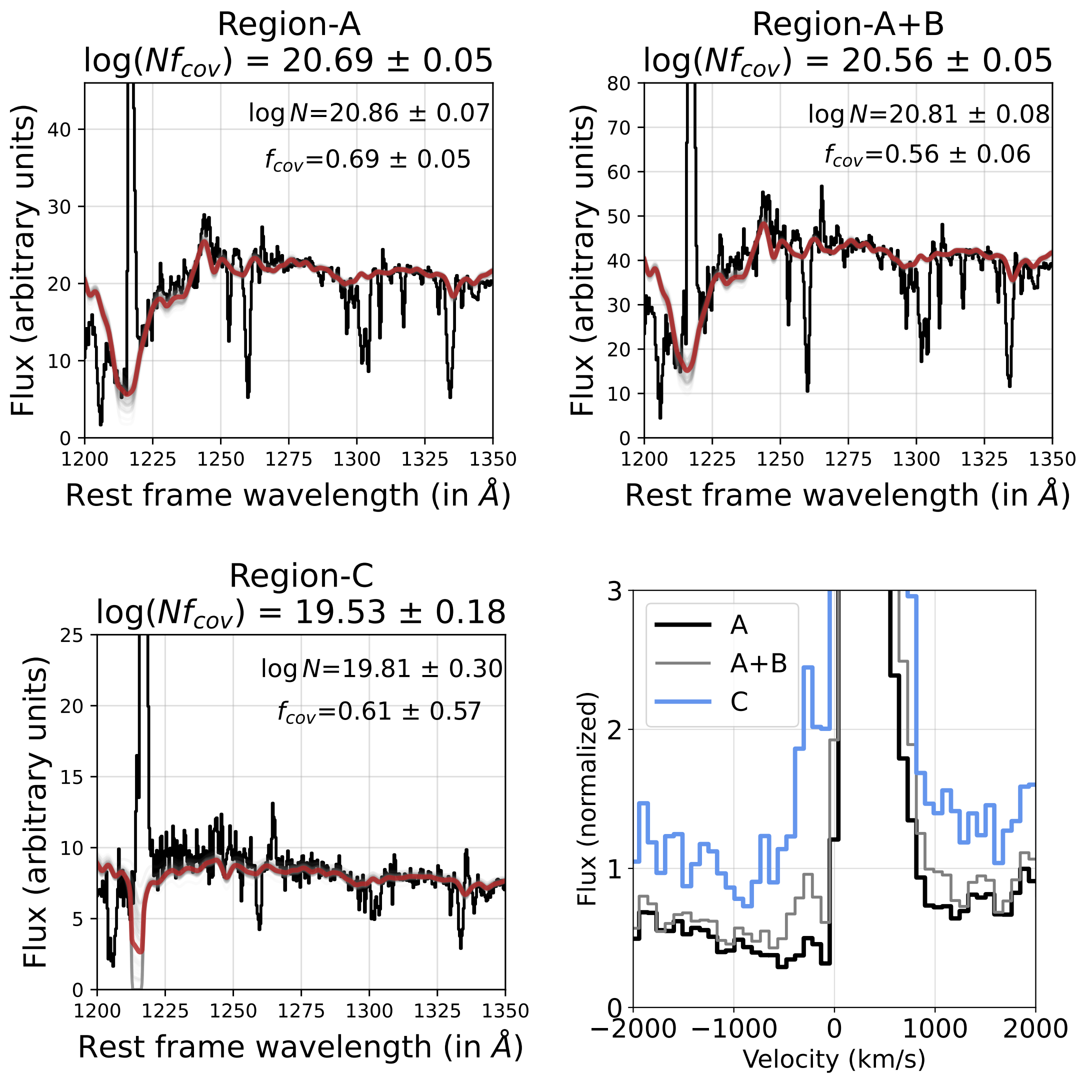

FIGURE 4.6. Spectra of regions A, A+B and C (black) along with the best-fit Voigt and stellar population fit (described in Section 4.5.3) shown in red. The gray lines show the 100 independent realizations used to obtain the best fit. *Bottom Right*: Close-up of the Lyα profile showing differences in blueshifted emission ($v < 0$ km s^{-1}) from each region. Region A+B has higher column density and covering fraction ($\log N \sim 21$; $f_{cov} \sim 0.6$) compared to Region C. This is complemented by our spatial maps of covering fraction ($C_f(v)$) obtained by independently analyzing the ISM metal absorption lines (Figure 4.5) which show a similar variation in C_f across the galaxy.

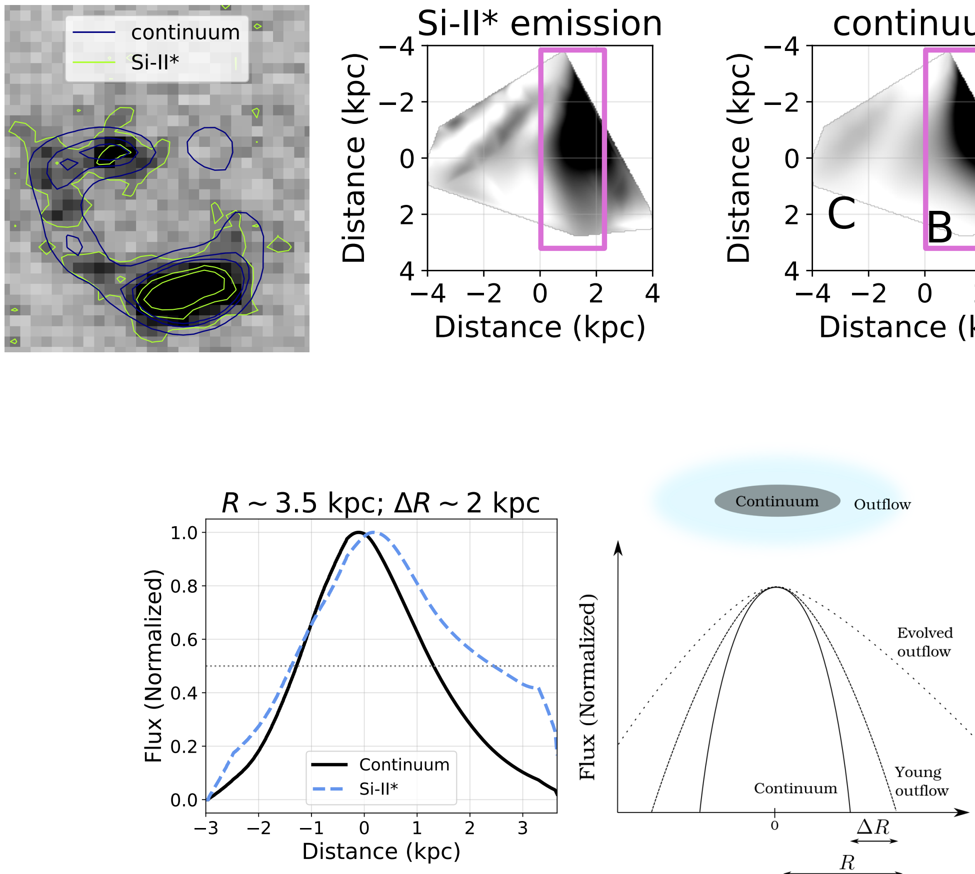

FIGURE 4.7. *Top Left*: Image plane Si II* λ1533 emission line map obtained from a single Gaussian fit (described in Section 4.5.4). Contours of the continuum and Si II* emission are shown in blue and green respectively. We find that the fluorescent emission is spatially extended compared to the continuum in all the three fold images of the arc. *Top Right*: Source plane continuum and Si II* λ1533 emission line maps. The purple box shows the position of a pseudo slit used to extract the continuum and emission line flux. It is centered on regions A+B which show the highest column density in absorption, and aligned roughly along the kinematic minor axis (see Figure 4.4). *Bottom Left*: Plot of extracted continuum and Si II* flux from the pseudo slit as a function of distance, showing larger spatial extent of fluorescent Si II* compared to the stellar continuum. *Bottom Right*: Illustration of two different outflow scenarios (young outflow, evolved outflow) in this phase space. R measures the radial extent of the outflowing gas from the center of the galaxy. ΔR in the spherical shell model represents the thickness of the shell that the outflow is enclosed in. Here we assume that the center of the galaxy is at the center of pseudo slit and define $\Delta R = R - R_{cont,50}$ where $R_{cont,50}$ is the radii where the flux reaches 50%. Based on this nomenclature, we estimate $R \sim 3.5$ kpc and $R/\Delta R \sim 2$.

various physical quantities which we address in this section: spatial maps of outflow velocity (v_{out}), covering fraction and solid angle (Ω), H I column density (N), and the radial extent of outflowing gas (R). Ultimately the spatially resolved KCWI spectroscopy provides a comprehensive view of the warm ($T \sim 10^4$ K) outflowing gas, including the mass loss rate (e.g., $\dot{M} \propto v_{out}\Omega N R$), as we will discuss in Section 4.6.

4.5.1. Systemic redshift and velocity field. For spatial analysis of the outflow kinematics, we require not only the systemic redshift but rather the velocity field of young stars in the galaxy. We obtain this using the nebular C III]$\lambda\lambda 1907, 09$ emission line doublet, which we fit with a double Gaussian function assuming the same redshift and velocity dispersion for both lines. For the integrated spectrum, we obtain a systemic redshift $z_s = 1.865947 \pm 0.000031$. The best-fit velocity centroid in each spaxel relative to this systemic redshift is shown in Figure 4.4, revealing a coherent velocity shear of $\sim$150 km s^{-1} seen consistently in all four multiple images.

A fit to the stellar photospheric absorption lines Si III λ1294,1298 yields a systemic redshift of $z_s = 1.865687 \pm 0.000307$, which agrees with C III] within the measurement uncertainties. This supports the use of C III] to trace the velocity field; we use the emission lines due to their higher signal-to-noise ratio relative to stellar absorption. We note that the difference between stellar and nebular systemic redshift corresponds to a $\sim 2\%$ change in the measured outflow velocity (v_{50}), which does not affect our results. A detailed analysis of the stellar kinematics will be presented in a future work (Rhoades et al. in prep). We use the redshift obtained from the integrated spectrum for measuring the outflow velocities (v_{50}) in Section 4.5.2 and discuss the spatially resolved outflow kinematics (i.e., $v_{50} - v_{sys}$) further in Section 4.6.1.

4.5.2. Outflow velocity and covering fraction. We follow the methodology described in Vasan G. C. et al. (2023) to analyze the ISM absorption profiles and derive outflow velocity metrics. For each spaxel, we obtain a covering fraction profile (C_f) by normalizing the spectra and combining the flux from Si II λ1260, C II λ1334, Si II λ1526 and Al II λ1670 using an inverse variance weighted average. We mask regions in the spectra which are affected by intervening absorption systems (e.g., -1000 and 1000 $\mathrm{km\,s^{-1}}$ from C II) and nebular emission features (e.g., O III] blueward of Al II). Figure 4.2's bottom panels demonstrate our methodology applied to the integrated galaxy spectrum. We note that the low-ion ISM absorption lines seen in Figure 4.2 as well as in individual spaxels show a clear asymmetric profile, with a blueshifted velocity centroid and an absorption wing extending to outflow velocities $v \gtrsim 250$ $\mathrm{km\,s^{-1}}$. Thus, we use a double Gaussian function to fit the resulting mean absorption line profile, which adequately captures the skewness apparent in the absorption line (Vasan G. C. et al., 2023). The covering fraction in each spaxel can be parameterized in the following form for ISM absorption lines with high optical depths ($\tau \gg 1$):

$$\frac{I}{I_0}(v) = 1 - C_f(v) \tag{4.1}$$

$$C_f(v) = C_{f,G1}(v) + C_{f,G2}(v) \tag{4.2}$$

where $C_{f,G1}$ and $C_{f,G2}$ are Gaussian functions which capture the faster- and slower-moving velocity components respectively.

From the fitted profiles, we measure the velocity centroid v_{50} (defined as the 50th percentile of absorption equivalent width) which traces the bulk outflow motion in the galaxy. The v_{50} metric is also robust to resolution and blending effects (as discussed in Vasan G. C. et al., 2023). Figure 4.4 shows the v_{50} map obtained for all spaxels in the arc which have a continuum signal-to-noise ratio (SNR) > 4 per pixel at representative wavelengths, enabling reliable fits to the ISM absorption. The v_{50} maps show that the bulk outflow velocity of the ISM gas varies across different star-forming complexes in the galaxy, from $|v_{50}| \lesssim 100$ $\mathrm{km\,s^{-1}}$ in region A to $|v_{50}| \gtrsim 200$ $\mathrm{km\,s^{-1}}$ in region C, relative to the adopted $z_s = 1.865947$. The spatial structure of v_{50} from outflows is however quite similar to the nebular emission velocity (Section 4.5.1). The outflow velocity relative to *local* systemic redshift is thus relatively constant. We discuss the implications of this in detail in Section 4.6.1.

Figure 4.4 (left panel) also shows the ISM absorption profile obtained from regions A, A+B and C revealing that region C has comparable absorption at high outflow velocities ($|v| \gtrsim 250$ km s^{-1}) but lower covering fraction at $v \sim 0$ relative to the other regions. This suggests that spatial variation in the observed outflow kinematics in this galaxy is due to the paucity of slower moving gas at systemic velocities, i.e., lower covering fraction at $v \sim 0$. We quantify this further by considering spatial maps of the covering fraction across the arc. Figure 4.5 shows the spatial map of the best-fit covering fraction profile (C_f) at different velocities along with a histogram of its spatial variation. We find that the gas at higher outflow velocities has a relatively uniform covering fraction, with mean $C_f = 0.22$ and 1σ spatial scatter of 0.06 at $v = -400$ km s^{-1}. At the systemic velocity, the covering fraction varies with mean and spatial scatter $C_f = 0.39 \pm 0.24$. Specifically, at $v \sim 0$, region A has a $C_f \sim 0.6$ compared to $C_f \sim 0.1$ in region C.

We can summarize the variations in outflow velocity as having a mean and spatial scatter of $v_{out} = -144 \pm 79$ km s^{-1}. The covering fraction varies both spatially and spectrally, with typical value and scatter of approximately $C_f = 0.4 \pm 0.2$.

4.5.3. Column density of H I. The low-ionization phase of outflowing gas is dominated by H I, whose column density we can measure directly from Lyα absorption. This provides important information on the total outflowing gas mass. We use a linear combination of Starburst99 templates (Leitherer et al., 1999) and a Voigt profile to simultaneously fit the Lyα absorption as well as the stellar continuum (Chisholm et al., 2019; Hu et al., 2023). During fitting, we mask out the strong stellar and interstellar features as well as the Lyα emission region, and fix the velocity centroid of the Lyα absorption to be the same as the v_{50} obtained from ISM absorption lines. We allow the extinction $E(B-V)$, Doppler parameter

b, column density N_{HI}, and covering fraction $f_{\rm cov}$ to be free parameters. The best fit values are obtained from taking the median and Median Absolute Deviation (MAD) from running 100 independent realizations. Here we refer to covering fraction as f_{cov} in the context of Lyα Voigt profile fitting, where f_{cov} can be thought of as the typical value for high column density gas. The covering fraction $C_f(v)$ derived from metal absorption lines is used when we are considering the velocity structure.

Figure 4.6 shows the resulting best fit to the spectrum of regions A, A+B and C. In the brightest region A+B of the galaxy, we obtain a value of $\log(N_{HI}) = 20.81 \pm 0.08$ for the column density of the H I gas and a covering fraction of $f_{cov} = 0.56 \pm 0.06$. This gives us a mean column density along the line of sight as $\log(N_{mean}) = \log(N_{HI} \times f_{cov}) = 20.56 \pm 0.05$. This is characteristic of damped Lyα absorption systems and is comparable to the column densities seen in other star-forming lensed galaxies at similar redshift (e.g., Jones et al., 2018). In contrast, region C does not show such strong Lyα damping wings and we find an order of magnitude lower N_{mean} compared to region A+B. We note that the f_{cov} values obtained from this fitting routine are consistent with the C_f obtained from independently fitting the ISM absorption lines. Based on these findings, we use the N_{mean} (Table 4.1) obtained from region A+B as the dominant outflow mass component for estimating the mass loss rate and mass loading factor, discussed further in Section 4.6.3.

4.5.4. Si II* emission line map. In this section, we examine fluorescent Si II* emission to establish the spatial profile of outflowing gas. Si II* emission predominantly occurs when a Si^+ ion absorbs a photon from the ground state and subsequently decays to an excited fine structure ground state, producing a photon with slightly lower energy (longer wavelength) than the one originally absorbed. These Si II* transitions appear to be optically thin in our target, such that the fluorescent emission directly traces the spatial distribution of *absorbing* gas (e.g., Jones et al., 2012; Prochaska et al., 2011), which is dominated by the outflowing component. We use the Si II* λ1533 line (Figure 4.2) which is detected across the entire arc and free from intervening absorption. However, it falls within the broad C IV stellar wind feature. Therefore, we fit a small region around Si II* $\lambda 1533$ with a Gaussian emission line profile combined with a linear continuum to account for the slope of the stellar wind feature.

Figure 4.7 (top panel) shows the resulting spatial map of fine structure emission flux in both the image and source plane, with the underlying continuum removed. We find that Si II* emission is patchy but spatially extends across the entire galaxy. The emission is strongest around regions A and B, which have both stronger continuum emission and higher H I column density than region C (Section 4.5.3). Together with the velocity measured from corresponding absorption lines, this indicates that the bulk of outflowing mass is associated with regions A and B.

Figure 4.7 also shows the continuum and Si II* spatial profiles extracted from a pseudo-slit through region A+B, probing the minor axis of the galaxy. The emission is well detected to a radial distance of $R \sim 3.5$ kpc. Si II* is more extended than the stellar continuum but with a rapidly declining flux profile. This suggests that the majority of outflow column density seen in absorption arises from gas confined to small radii (and impact parameters), $R \lesssim 5$ kpc, which is supported by observations from galaxy-galaxy pairs and quasar sightlines (e.g., Steidel et al., 2010; Nielsen et al., 2013b; Vasan G. C. et al., 2023).

The spatial extent of outflows as probed in emission by Si II*, in combination with the kinematics and column density discussed previously, allows a direct measurement of outflow mass loss rate in the low-ionization phase. We discuss the mass loss and its implications in the following section.

4.6. Results and Discussion

4.6.1. Absorption traces recently-launched outflows. Comparing the spatial maps of outflow velocity and nebular emission kinematics allows us to examine whether the outflows are associated with local launching sites, or galaxy-wide winds. We consider three example scenarios in the evolution of a galactic-scale wind: (a) young outflow, where the gas is still

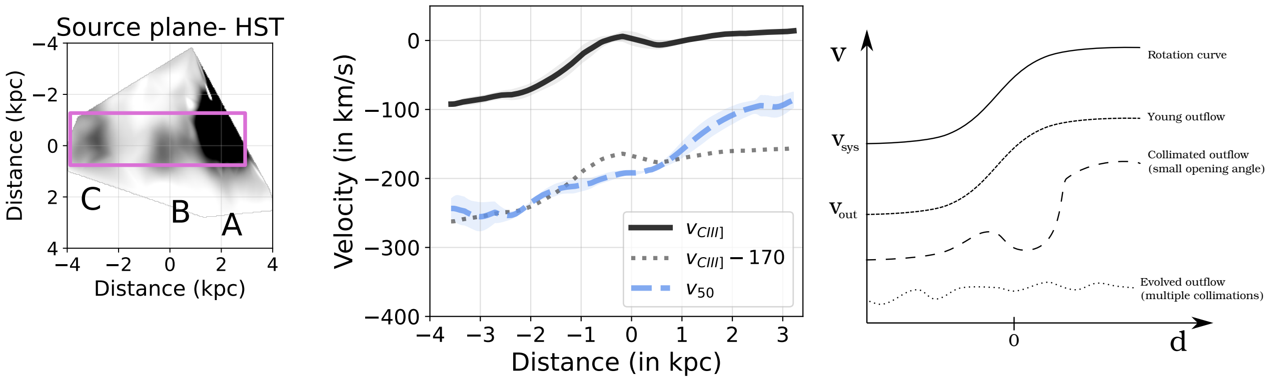

FIGURE 4.8. *Left:* Source plane HST image with a pseudo-slit (purple) along the major axis, used to characterize the outflow velocity structure. *Center:* Systemic and outflow velocity as a function of distance measured along the pseudo-slit. The black and blue lines represent the median systemic and outflow velocities, respectively. The shaded regions show the 1σ scatter from collapsing the slit. The gray dashed line shows the $v_{CIII]}$ median line offset by $-170\,\mathrm{km\,s^{-1}}$, showing an approximately constant outflow velocity relative to the local systemic redshift. *R*ight: Schematic illustration of different stages in the evolution of a galactic scale wind (with velocity v) as a function of distance from the center of a galaxy (d). The rotation curve of the galaxy tracing the motion of the stars is denoted by v_{sys}. With respect to this systemic velocity (v_{sys}), three different stages of an outflow are shown: (a) young outflow ($v_{out} = v_{sys} + \text{constant}$), (b) collimated outflow ($v_{out} = v_{sys} + \text{constant} + f(r)$), (c) evolved outflow ($v_{out} = \text{constant}$). Each of these scenarios can be distinguished observationally by measuring $v_{out} - v_{sys}$ versus d. Based on our schematic, CSWA13 follows the young outflow case, with outflows mirroring the nebular kinematics.

located close to its launching site; (b) collimated outflow, launched from the central regions of a galaxy with a modest opening angle; and (c) evolved outflow, with coherent motion at distances larger than the galaxy stellar radius.

In the early stages of outflow, feedback processes (e.g., radiation pressure and supernovae) drive the ISM outward from regions of recent star formation. If the initial launching velocity is relatively uniform across different regions in a galaxy, then we would expect the galaxy's rotation curve (v_{sys}, defined as the local velocity of stars and H II regions) to be imprinted in the outflow: $v_{out} \approx v_{sys} + \text{constant}$. We consider this scenario where the ISM is being pressurized by feedback but is yet to burst out of the confines of the galaxy as a 'young outflow.'

As momentum continues to build up in the ISM, the gas can be preferentially launched from regions of low density (e.g., orthogonal to a galactic disk) resulting in a collimated biconical outflow. In this case, the outflow velocity would be higher in the region of collimation compared to the rest of the galaxy, i.e., $v_{out} = v_{sys} + \text{constant} + f(r)$ where $f(r)$ is a function of galactic radius. After a sufficient time, the wind may travel well beyond the galaxy's stellar radius and mix with the CGM, with any signal of the initial launching momentum being mixed such that the outflow velocity appears relatively uniform in down-the-barrel sightlines: $v_{out} = \text{constant}$. As galaxies evolve in time with multiple feedback episodes, we would expect the relation between the systemic and outflow velocity to be a combination of all of these scenarios.

Figure 4.8 (right) shows an illustration of these three scenarios. The observed outflow velocity profile of CSWA13 closely follows that of the nebular velocity (v_{sys}) with a constant offset, as shown in the top right panel of Figure 4.8. This is also evident in the spatial maps discussed earlier (Figure 4.4) which show similar gradients. We thus conclude that the outflows seen in absorption are dominated by the 'young outflow' scenario. Similar spatially resolved studies, although at coarser resolution, have found that in a $z = 4.9$ arc (Swinbank et al., 2009b) the outflows mirror the nebular emission similarly to CSWA13, whereas in the cosmic horseshoe ($z \sim 2.4$; James et al., 2018b) the velocity follows the evolved outflow case.

We find a median and 1σ scatter in velocity centroid of $v_{50} = -144 \pm 79$ km s^{-1} from spatially resolved regions in CSWA13 (Section 4.5.2). This is similar to down-the-barrel integrated absorption profiles at $z \sim 2$ (mean $v_{50} = -141$ and sample scatter $\sim$100 km s^{-1}) from Vasan G. C. et al. (2023) and Steidel et al. (2010). However, when we consider the local outflow velocity relative to the star formation regions (i.e., $v_{50} - v_{CIII]}$), the scatter is only 41 km s^{-1}. This spatial variation in outflow velocity is similar to the $\pm$40 km s^{-1} differences found in a lensed $z \sim 1.7$ galaxy by Bordoloi et al. (2016). Thus we find that the outflow velocity is closely connected to the bulk motion of star-forming gas in the galaxy.

4.6.1.1. *Outflows are encapsulated within the continuum.* Based on our previous discussion, if the outflows are indeed young then we expect a small radial extent. In this case, the fluorescent fine structure emission should be closely connected to the stellar morphology traced by continuum emission. We can estimate a characteristic timescale of the ongoing star formation as $\mathrm{sSFR}^{-1} \simeq 25$ Myr (Mainali et al., 2023b). A galactic wind launched ∼25 Myr ago with constant velocity $-150\ \mathrm{km\,s^{-1}}$ would travel a radial distance of ~ 4 kpc. In contrast, for an older evolved outflow we would expect fine structure emission to arise at larger radial galactocentric distances. From the 2D maps of continuum and fluorescent Si II* emission (Figure 4.7; described in Section 4.5.4), the Si II* is detected across the spatial extent of the galaxy and is patchier. This suggests that the radial extent of the outflowing gas is comparable to the projected size of the galaxy which is $\lesssim 8$ kpc.

One might also expect that turbulent young outflows at close radial distances would entrain the ambient gas surrounding the H II regions (e.g., McKee & Ostriker, 1977) resulting in broad emission lines. Visually inspecting the C III] and O III] nebular emission lines in region A+B where the SNR is high, we find that the nebular lines show a clear blueshifted wing component, with the velocity centroid of the blueshifted component similar to that seen in the absorption. This suggests that the outflow emission originates from the same regions as the absorption and supports the idea of a multi-phased wind being launched across the galaxy. This may also enable thorough mixing of the ambient ISM.

4.6.2. Outflowing gas is inhomogeneous. It is clear from the previous subsection that the outflows detected ubiquitously across the galaxy are likely young and radially confined within $\lesssim 8$ kpc. In this subsection, we explore the 'geometry' of the outflow as revealed by the variation in covering fraction $C_f(v)$ of the ISM gas at different velocities. $C_f(v)$ is measured from ISM absorption profiles in each spaxel (using double Gaussian fits; Section 4.5.2).

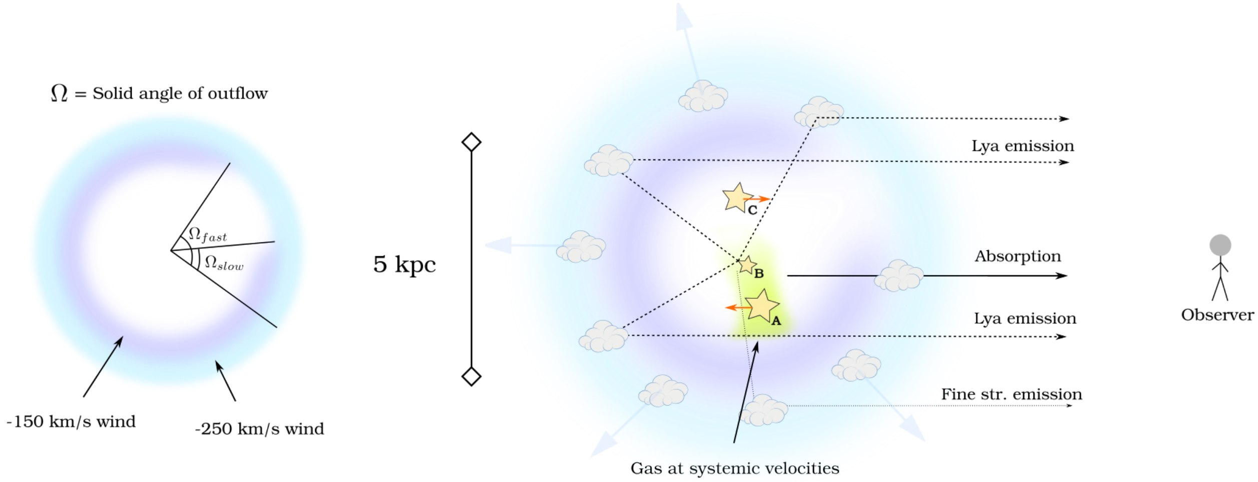

FIGURE 4.9. Illustration of the observed low-ionization outflow kinematics in CSWA13 using a spherical shell geometry. Gas at systemic velocities ($v \sim 0$) is shown in green. The blue and purple regions denote the azimuthal variation of the outflowing ISM gas at $v = -150$ and $v = -250$ km s^{-1}, with solid angles $\Omega_{\rm slow}$ and $\Omega_{\rm fast}$ respectively. Our observations of the covering fraction of outflowing gas (Figure 4.5) suggest that gas at higher outflow velocities is more homogenous compared to the slower moving gas. The receding side of the outflow causes backscattering of Lyα photons from H I atoms, resulting in redshifted Lyα emission which we detect across the entire galaxy (Figure 4.6). We also observe blueshifted Lyα photons leaking along the line-of-sight of region C in the galaxy from forward scattering, which is likely due to low column density of neutral gas at systemic velocities. The orange arrows denote kinematics of the nebular gas (Section 4.5.1) in the galaxy.

Assuming our line-of-sight is representative, the covering fraction is related to the outflow solid angle Ω as $C_f(v) = \Omega(v)/4\pi$, wherein a spherically symmetric outflow would correspond to C_f = constant in all spatial regions.

Figure 4.5 plots the covering fraction (C_f) at $v = -400, -250, -100$ and 0 km s^{-1} which correspond roughly to 2.7×, 1.7× and 0.7× the median outflow velocity (v_{50}) of the galaxy. We note that the data plotted here is not corrected for the instrumental line spread function. This has the effect of decreasing C_f at the tails of the distribution and increasing it near the v_{50} velocities. Nonetheless, the velocity channels in Figure 4.5 are nearly independent, and the spatial variation in C_f is robust to spectral resolution effects. At slower outflow velocities ($v \gtrsim -100$km s^{-1}), we can clearly see that the covering fraction of gas in regions A and B of the galaxy is significantly larger by up to $\sim 6\times$ than in region C. However, at faster outflow velocities $v \lesssim -250\ \mathrm{km\,s^{-1}}$, the covering fraction is approximately uniform across the galaxy. This suggests that the high-velocity gas is more homogeneous compared to the slower moving gas. Additionally, we do not find any significant difference between the low-ionization absorption profiles and the intermediate-ionization species (e.g., Al III, C IV) suggesting that the covering fraction does not vary significantly between these ionization states. We note however that the mean covering fraction measured at different outflow velocities is $\overline{C_f} \sim 0.4$ (corresponding to a solid angle of $\Omega = 1.4\pi$ steradians if our sightline is representative), indicating that the overall geometry of the outflow is patchy and asymmetric.

Figure 4.9 illustrates our findings on the spatial and velocity-resolved $C_f(v)$ with a simple spherical shell outflow schematic. We find further evidence of the patchiness of outflowing ISM gas from variations in Lyα absorption across different regions of the galaxy (e.g., Figure 4.6). Region A has damped Lyα absorption and exclusively redshifted emission, whereas region C does not show clear damping wings and exhibits a blueshifted emission component. This indicates a significantly lower column density of H I and associated low-ionization gas toward region C, along with the lower covering fraction (Figure 4.5).

4.6.3. Mass loss rate is comparable to the star formation rate. We now use the measurements of outflow geometry, velocity, and column density to estimate the mass loss rate and mass loading factor. Thanks to the spatial and spectral resolution of KCWI data, we can largely avoid systematic uncertainties arising from low spectral resolution, unknown radial distribution of outflowing gas, geometry, and/or outflow velocity relative to systemic, which have presented challenges for earlier efforts (e.g., Pettini et al., 2000; Chisholm et al., 2017).

If the outflowing gas is described by a spherical shell geometry, with average radial distance R and width ΔR, then we can estimate the mass loss rate through the shell as follows:

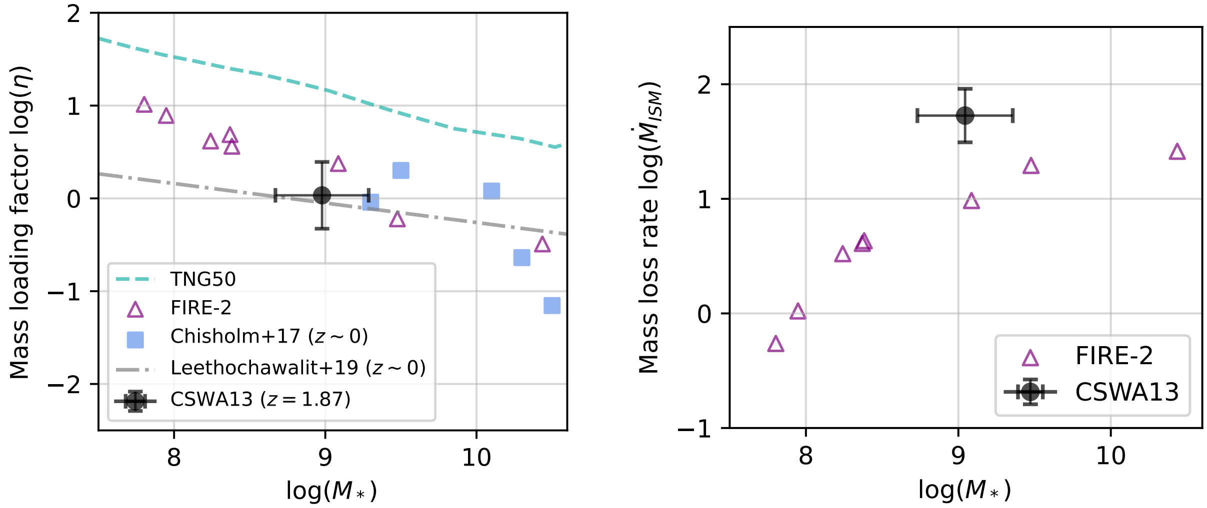

FIGURE 4.10. Comparison of mass loading factor ($\eta = \dot{M}_{out}/\mathrm{SFR}$; left) and mass loss rate ($\dot{M}_{out}$; right) obtained for CSWA13 from this work (black point) with cosmological simulations (TNG50, FIRE-2) and galaxies in the low redshift universe. The average η obtained from the FIRE-2 (Pandya et al., 2021) simulations at $z \sim 2$ are marked in purple triangles measured at a radial distance of $R = 0.1 - 0.2\ R_{vir}$. The green dashed line denotes the median η value obtained from simulated $z = 2$ galaxies in TNG50 (Nelson et al., 2019) at $R = 10$ kpc. The η measured in local star-forming regions from Chisholm et al. (2017) are shown in blue. The gray dashed line shows the best fit from low-redshift quiescent galaxies (Leethochawalit et al., 2019).

$$\begin{aligned}
\dot{M}_{out}(v_{50}) &= (4\pi)(\mu m_p) \times (f_{\mathrm{cov}} N_{HI}) \times v_{50} \times R \times \frac{R}{\Delta R} \\
&= (77\mu) \times \left(\frac{f_{\mathrm{cov}} N_{HI}}{10^{21}\ \mathrm{cm}^{-2}}\right) \times \\
&\left(\frac{v_{50}}{-150\,\mathrm{km\,s^{-1}}}\right) \times \left(\frac{R}{5\ \mathrm{kpc}}\right) \times \left(\frac{R}{\Delta R}\right) \mathrm{M_{\odot} yr^{-1}}
\end{aligned} \tag{4.3}$$

where $\mu \approx 1.4$ is the average mass per hydrogen atom in the outflow (mainly accounting for hydrogen and helium), f_{cov} is the mean covering fraction of the outflowing ISM gas over the entire solid angle of 4π steradians, N_{HI} is the average column density obtained from the integrated spectrum of the galaxy, and v_{50} is the median outflow velocity. The negative sign in velocity indicates that the gas is outflowing from the galaxy. This is similar to the formalism used in Jones et al. (2018) and Pettini et al. (2000), with the normalization in Equation 4.3 chosen such that each term in the parentheses is ~ 1 for our target. Rewriting Equation 4.3 in logarithmic units, with $\mu = 1.4$ and $N_{\rm mean} = f_{\rm cov} \times N_{HI}$ (discussed in Section 4.5.3), we obtain

$$\begin{aligned}\log \frac{\dot{M}_{out}(v_{50})}{\mathrm{M_\odot yr^{-1}}} &= 2.033 + \log\left(\frac{N_{\rm mean}}{10^{21}\ \mathrm{cm^{-2}}}\right) \\ &+ \log\left(\frac{v_{50}}{-150\,\mathrm{km\,s^{-1}}}\right) + \log\left(\frac{R}{5\ \mathrm{kpc}}\right) \\ &+ \log\left(\frac{R}{\Delta R}\right).\end{aligned} \tag{4.4}$$

If the ratios in each term are of order unity, then the inferred mass loss rate is $\dot{M}_{out} \sim 100\,\mathrm{M_\odot yr^{-1}}$.

Throughout this paper (Section 4.5.2, 4.5.3, 4.5.4), we have used the spatially resolved data to constrain each term in Equation 4.4. We summarize these values in Table 4.1. The median and 1σ spatial scatter of the covering fraction of the gas is $\overline{C_f} = 0.4 \pm 0.2$ outflowing at $v_{50} = -144 \pm 79\ \mathrm{km\,s^{-1}}$. Based on the fluorescent emission, we find that the wind is confined to a radius similar to that seen in the stellar continuum ($R \sim$3.5 kpc) with the thickness of the shell $\Delta R \sim$2 kpc. The mean H I column density of gas in the outflow is $\sim 10^{21}\ \mathrm{cm^{-2}}$ with region A+B having the dominant contribution. Based on these measurements, we estimate the mass loss rate for the low ionization phase as $\log(\dot{M}_{out}/(\mathrm{M_\odot yr^{-1}})) = 1.73 \pm 0.23$. The derived mass loss rate is similar to predictions from FIRE-2 simulations at $z \sim 2$ measured at a radial distance of $R = 0.1 - 0.2\ R_{vir}$ given the stellar mass of CSWA13 (Figure 4.10). We note that this value for the low-ionization outflowing gas represents a lower limit on the total mass loss rate of the galaxy, which likely has contributions from other ionization phases (i.e., ionized and molecular hydrogen). For example, analysis of the FIRE simulations by Muratov et al. (2017) suggests that $\sim 70\%$ of the outflowing and circumgalactic gas is in the low-ionization phase for galaxies with similar mass and redshift as CSWA13. Including the contribution from other phases would therefore increase our estimate of the *total* mass loss rate $\dot{M}_{out,\ total}$ by ~ 0.2 dex. The contribution of different ionization states can also be quantified with column densities of various metal ions measured from their absorption lines (e.g., Chisholm et al., 2016; Jones et al., 2018), although this is beyond the scope of this paper.

Despite the measured high mass loss rate in CSWA13, the bulk of the ISM gas entrained in the outflow is likely unable to escape its gravitational potential well. We find that the outflow velocity ($v_{50} \simeq -150\ \ \mathrm{km\,s^{-1}}$) is lower than the escape velocity estimated from the stellar mass, via the stellar-to-halo mass relation (e.g., Behroozi et al., 2019) or estimated escape velocities in FIRE-2 simulated galaxies with similar stellar mass (Vasan G. C. et al., 2023). This suggests that the gas launched via outflows from the ISM will remain bound within the halo and/or recycle back at later times. Given a constant mass loss rate, CSWA13 would need only ~20 Myr to enrich its CGM with a gas mass comparable to its stellar mass. The large outflow rate, if sustained, is thus capable of creating a metal-enriched circumgalactic gas reservoir which can in turn sustain future star formation via recycling.

We now turn to the *efficiency* of stellar feedback in driving these powerful outflows. This is quantified by the mass loading factor ($\eta = \frac{\dot{M}_{out}}{\mathrm{SFR}}$), defined as the ratio of the mass loss rate of outflowing gas to the SFR of the galaxy. Cosmological simulations such as TNG50 (Nelson et al., 2019) and FIRE-2 (Pandya et al., 2021) predict that galaxies at $z \sim 2$ are highly efficient at driving outflows with typical mass loading factors ranging from $\log \eta \sim$ 0–1.7 (factors $\eta \sim$ 1–50) in the stellar mass range $\log M_*/\mathrm{M}_\odot =$ 8–10 (Figure 4.10). For CSWA13, using our resolved observations, we measure the mass loading factor for the low-ionization gas phase as $\log \eta = 0.04 \pm 0.34$. This efficiency value is similar to predictions from FIRE-2, nearby star-forming galaxies (Chisholm et al., 2017), low-redshift quiescent galaxies (Leethochawalit et al., 2019) as well as lensed quiescent galaxies at $z \sim 1$ (Zhuang

Quantity		Measured value	Reference
Stellar Mass	$\log M_*^a$	9.00 ± 0.32 M$_\odot$	Section 4.4
Star Formation Rate	$\log \mathrm{SFR}^a$	1.71 ± 0.21 M$_\odot$ yr^{-1}	Section 4.4
Outflow velocity Centroid	v_{50}	-144 ± 79 km s^{-1}	Section 4.5.2, Figure 4.4
Mean Column Density	$\log(N_{\mathrm{mean}})^b$	20.56 ± 0.05 cm^{-2}	Section 4.5.3, Figure 4.6
Radial extent of outflowing gas	R	3.5 kpc	Section 4.5.4, Figure 4.7
Thickness of radial shell	ΔR	2 kpc	Section 4.5.4, Figure 4.7
Mass loss rate	$\log \dot{M}_{out}^b$	1.73 ± 0.23 M$_\odot$ yr^{-1}	Section 4.6.3, Figure 4.10
Mass loading factor	$\log \eta^c$	0.04 ± 0.34	Section 4.6.3, Figure 4.10

TABLE 4.1. Summary of measured quantities. a - Corrected for lensing magnification $|\mu|$.
b - Here $N_{\mathrm{mean}} = N_{HI} \times f_{cov}$ where f_{cov} is obtained from fitting the Lyα profile, as opposed to $C_f(v)$ obtained from the absorption profile of metal ion transitions.
c - Assuming a spherical geometry with $\frac{R}{\Delta R} = 1.75$

.

et al., 2023), but it is an order of magnitude lower than those predicted by TNG50 at similar redshifts (Figure 4.10). We note that this is not strictly a direct comparison as the methods employed to observationally estimate η differ, and the simulation values correspond to a fixed radii (e.g., R = 10 kpc for TNG50) and thickness (e.g., 0.1 – 0.2 R_{vir} for FIRE-2) and can make use of full spatial and kinematic information as opposed to down-the-barrel observations of a cylindrical sightline. Nevertheless, our spatially resolved observations serve as an excellent test of different feedback prescriptions, and we view further direct comparison with simulations as a promising prospect.

4.6.4. Spatial variation in outflow properties. We have demonstrated significant spatial variation in the outflow properties of CSWA13, with the higher surface brightness regions A+B also having stronger outflows (i.e., larger mass loss rates) compared to region C. We also observe variation in the velocity structure, with larger effective outflow velocity in region C. One might expect a correlation of higher star formation densities leading to higher outflow velocities (e.g., Heckman, 2002; Cicone et al., 2016), in contrast to our results. However, the higher outflow velocity in region C is driven by a lower C_f at low velocities. This may indicate less mass loading of the ambient ISM from region C, resulting in higher velocity from momentum conservation. This is supported by Lyα measurements indicating a lower H I column density and mass loading factor toward region C.

The complex morphology of CSWA13, aided by gravitational lensing, illustrates the value of spatially resolved information for characterizing gas outflows. We have found order-of-magnitude variation in the total column density toward different regions of the galaxy, with resolved spectroscopy pinpointing regions A+B as the dominant outflow launching sites. We also observe variation in C_f and Lyα emission profiles. The lower covering fraction and blueshifted Lyα emission component in region C may be particularly interesting in terms of understanding how ionizing photons escape from galaxies, as these signatures are indicative of significant ionizing escape fractions (e.g., Verhamme et al., 2008; Jones et al., 2013c;

Leethochawalit et al., 2016d). The rich variations revealed in this galaxy clearly demonstrate inhomogeneous outflow properties, and the value of spatially resolved information.

4.7. Conclusions

In this paper, we have investigated the spatially resolved outflow properties and kinematics of a $z \sim 2$ gravitationally lensed star-forming galaxy (CSWA13) using Keck/KCWI. We map outflows in multiple ultraviolet ISM absorption lines, along with fluorescent Si II* emission tracing the outflow spatial structure, and nebular emission from C III] tracing the systemic redshift and velocity structure. We summarize our key findings below.

(1) The spatial structure of outflow velocity resembles that of the nebular kinematics, which we interpret to be a signature of a young galactic wind that is pressurizing the ISM of the galaxy.

(2) From the radial extent of Si II* emission, we estimate that the outflow is largely encapsulated within 3.5 kpc. We explore the geometry (e.g., patchiness) of the outflow by measuring the covering fraction at different velocities, finding that the maximum covering fraction is at velocities $v \sim -150$ km s^{-1}. We find significant variation in the outflow covering fraction near this peak velocity, with lower but more uniform covering fraction in the higher-velocity gas.

(3) We calculate the mass loss rate and mass loading factors from measurements of the outflow velocity, radius, column density, and covering fraction for the low-ionization outflowing gas in CSWA13. The mass loss rate ($\log \dot{M}_{out}/(\mathrm{M_{\odot} yr^{-1}}) = 1.73 \pm 0.23$) is comparable to the star formation rate ($\log \mathrm{SFR}/(\mathrm{M_{\odot} yr^{-1}}) = 1.71 \pm 0.21$) resulting in a mass loading factor $\log \eta \sim 0.04 \pm 0.34$ in the galaxy, indicating efficient coupling of stellar feedback to drive outflowing mass that is likely to remain in the inner circumgalactic medium or be recycled back into the galaxy. This low-ionization outflow rate is a lower limit on the total mass loss rate of the galaxy, although the low ionization phase is likely the dominant contributor. Based on theoretical predictions,

we estimate that the total outflow rate is $\sim$ 0.2 dex higher with all ionization phases included (Section 4.6.3). We compare our measurement with cosmological simulations, finding that the mass loading factor agrees with predictions from FIRE-2 but is lower by an order of magnitude than those seen in TNG50.

(4) The outflow properties of CSWA13 exhibit significant spatial variation, with the higher surface brightness regions A+B being the dominant launching site of strong outflows (i.e., larger column density and mass loss rate) compared to the lower surface brightness region C. We also observe variation in the velocity structure, with larger effective outflow velocity in region C. Spatially resolved data aided by gravitational lensing is important for capturing the rich variations in inhomogeneous outflow properties in high-redshift galaxies such as CSWA13.

Overall, these findings support a picture in which outflows observed ubiquitously in early star-forming galaxies such as CSWA13 are responsible for transporting large amounts of mass and metals into the inner circumgalactic medium. This process provides a gas reservoir to sustain star formation at lower redshifts. This work represents early results of our ongoing efforts to spatially resolve outflow (and systemic) kinematics and composition in lensed galaxies at cosmic noon ($z \simeq 2-3$), and demonstrates the power of sensitive rest-frame ultraviolet IFS to characterize the effects of feedback on the ISM and CGM at these redshifts. An enlarged sample will help to demonstrate scatter in the population and scaling relations with galaxy mass and other properties. Ultimately the methods used herein represent a path toward establishing the cosmic history of baryon cycling and providing a benchmark for comparison with theoretical models of feedback and galactic outflows.

CHAPTER 5

Conclusion and future directions

In this chapter, I provide a brief summary of the work presented in this thesis and outline exciting projects that I intend to work on as a postdoctoral researcher in the upcoming years.

5.1. New insights on star formation feedback and galactic outflows using gravitational lensing

Chapter 2 evaluated the performance of different CNN learning approaches and data augmentations on their ability to efficiently find gravitational lens candidates. Using state-of-the-art semi-supervised learning approaches on deep learning architecture, along with a training dataset consisting of simulated lenses and survey image non-lenses, this work demonstrated that we can greatly reduce the human effort required to find lensed candidates from imaging surveys. Applying this approach to the entire Deep Lens Survey (DLS) survey, and visually inspecting the top ~ 2500 lens candidates, we found 9 Grade-A and 13 Grade-B lensed candidates (22 lensed candidates in total). The lensed nature of 2 Grade-A candidates were confirmed with spectroscopy and high-resolution imaging, demonstrating that our methods are successful. The number of lenses found in the Deep Lens Survey (DLS) corresponds to $\sim 10\times$ higher sky density of lenses per deg^2 compared to the shallower DES/DECaLS survey imaging. This supports predictions that vast numbers of lens systems ($\gtrsim 10^5$) will be detectable in the upcoming generation of all-sky surveys (such as Rubin/LSST, Roman, and Euclid) which will survey the sky at high angular resolution and sensitivity.

Magnification from strong lensing enables observations with good spectral resolution, thus allowing us to characterize the complex kinematic structure of ISM gas. Chapter 3 represents the largest sample to date of well-resolved velocity profiles of gas outflows driven by star forming galaxies at cosmic noon. This work used moderate resolution ($R \gtrsim 4000$) 'down-the-barrel' ISM absorption line spectra to characterize the ISM and outflowing gas in a sample of 20 strongly lensed galaxies, probing the absorption lines with ~10 independent spectral resolution elements. We examined the covering fraction profile (C_f) and outflow velocities (parameterized by v_{50}, v_{75}, etc.), width of absorption (Δv_{90}), skewness of absorption profiles, and optical depth (τ) of absorbing gas. We observed ubiquitous outflows with a typical median velocity of $v_{50} \simeq -150$ km s^{-1}, with the extent of detected absorption reaching 3× this median value in most cases. We also explored scaling relations between outflowing gas kinematics and the host galaxy properties (e.g., M_* and SFR), finding that Δv_{90} correlates well with SFR and that the observed scaling relations are consistent with theoretical expectations for momentum-driven outflows. We demonstrated the importance of having good spectral resolution in studies of outflowing gas by showing that $R \gtrsim 1700$ is required to recover the full asymmetric covering fraction profile structure. Outflow velocity metrics such as maximum velocity and Δv_{90} that are crucial for establishing galaxy scaling relations show a clear bias and substantial scatter at lower spectral resolution. We compared these measured trends in outflow velocity with the TNG50 and FIRE-2 cosmological simulations, and find reasonable agreement, which is encouraging for future work using simulations to help interpret outflow properties.

Spatially resolved data aided by gravitational lensing is important for capturing the rich variations in inhomogeneous outflow properties in high-redshift galaxies. Using the power of Integral Field Spectroscopy (IFS), Chapter 4 investigated the spatially resolved outflow properties and kinematics of a $z \sim 2$ gravitationally lensed star-forming galaxy (CSWA13) using Keck/KCWI. We mapped outflows in multiple ultraviolet ISM absorption lines, along with fluorescent Si II* emission tracing the outflow spatial structure, and nebular emission from C III] tracing the systemic redshift and velocity structure. The outflow properties of CSWA13 exhibit significant spatial variation, with the higher surface brightness regions being the dominant launching site of strong outflows (i.e., larger column density and mass loss rate) compared to the lower surface brightness regions. The Si-II* emission shows larger spatial extent compared to the stellar continuum and from its radial extent, we estimate that the outflow is largely encapsulated within 3.5 kpc. We calculated the mass loss rate and mass loading factors from measurements of the outflow velocity, radius, column density, and covering fraction for the low-ionization outflowing gas in CSWA13. Overall, these findings support a picture in which outflows observed ubiquitously in early star-forming galaxies such as CSWA13 are responsible for transporting large amounts of mass and metals into the inner circumgalactic medium. This process provides a gas reservoir to sustain star formation at lower redshifts. An enlarged sample will help to demonstrate scatter in the population and scaling relations with galaxy mass and other properties. Ultimately, the methods used herein represent a path toward establishing the cosmic history of baryon cycling, and providing a benchmark for comparison with theoretical models of feedback and galactic outflows.

5.2. Future work

Chapter 4 has demonstrated the power of sensitive rest-frame ultraviolet IFS to characterize the effects of feedback by measuring the mass loss rate and mass loading factor. *All* current models of galaxy evolution agree that feedback via outflows is an essential component of galaxy formation. However, the actual mass loss rates are largely unknown, with the prevailing uncertainties typically of an order of magnitude. For example, the FIRE (e.g., Muratov et al., 2015, 2017; Pandya et al., 2021) and TNG50 (Nelson et al., 2019) cosmological simulations at $z \sim 2$ predict large outflow mass loading factors ranging from $\eta \gtrsim 100$ at low mass ($M_h \sim 10^9 \mathrm{M}_\odot$) to of order unity at higher masses ($M_h \sim 10^{11} \mathrm{M}_\odot$). The high η values predicted for low-mass and high-redshift galaxies have yet to be confirmed by observations. Outflow scaling relations have been studied in galaxies at $z \sim 0$ using HST/Cosmic Origins Spectrograph reaching down to $M_* \sim 10^6$ $\mathrm{M}_\odot$ (Xu et al., 2022b) with a measured $\eta \sim 10$ for the low mass galaxies and $\eta \lesssim 1$ for the intermediate-mass galaxies. As part of this thesis, I have obtained suitable IFS data to measure the outflow velocity, spatial extent of outflowing gas, gas column densities and eventually the mass loss rate and mass loading factor in a sample of $\gtrsim 10$ lensed galaxies at $z \sim 2$ that span an order of magnitude in their properties such as SFR and stellar masses. Chapter 4 represents the first results using such spatially resolved IFS measurements to establish the mass loading factor in a lensed galaxy with a stellar mass of $M_* \sim 10^9$ $M_\odot$. Over the course of the next few years, I plan to use the methodology described in Chapter 4 to self-consistently establish these quantities in the lensed galaxy sample. Additionally, I will utilize down-the-barrel 'mock spectra' of $z \sim 2$ galaxies generated from simulations using different feedback prescriptions to directly compare measurements of outflowing gas such as mass loss rate and outflow velocity with observations. This could also be useful to assess the likely radial distribution of gas seen in absorption, and as a further test of feedback prescriptions used in simulations. Cosmological simulations (e.g., Pandya et al., 2021) also predict that the warm $\sim 10^4$K gas phase dominates the mass loading fraction in intermediate-mass and dwarf galaxies, carrying most of the momentum, energy and metals. Conducting a study similar to Chapter 4 on intermediate/high-ionization absorption lines such as Si IV and O VI (e.g., Chisholm et al., 2018) would provide useful constraints on the relative contributions of the gas phases to the mass loading.

A complimentary approach to test different prescriptions of feedback is to characterize the spatially resolved maps of galaxies gas-phase metallicity ('metallicity map') and their radial gradients ('metallicity gradient'). For example, stronger feedback leads to more mixing of heavy elements in the ISM and CGM, resulting in flatter metallicity gradients. This represents a major science goal of KAPA (Keck All-sky Precision Adaptive Optics), an ongoing major upgrade to the Keck I adaptive optics system with completion expected in mid-2025, for which I am a science team member. Specifically, one of KAPA's Key Science Programs aims to measure metallicity gradients for 40 lensed galaxies, along with their ISM kinematics and resolved star formation. Lensed galaxies with AO give the best possible spatial resolution for understanding the structure of distant galaxies, along with spectral resolution to map kinematics and gas metallicity from nebular emission lines. Resolved measurements of metallicity gradients in galaxies at sub-kpc scales will help further understand the physical processes that drive outflows in galaxies at Cosmic Noon and offer observational constraints to sub-grid physics models used in the current generation of simulations. In addition to the metallicity maps that I will obtain from KAPA, as part of my postdoctoral position, I will be analyzing the nebular kinematics and metallicity of lensed galaxies observed using JWST/NIRSpec.

Finally, the technique of arc tomography offers a unique possibility to probe the structure of the circumgalactic medium (CGM), which represents a key component in the baryon cycle. In this method, bright lensed arcs are used as background lights against which absorption from CGM gas can be spatially mapped. During my graduate studies, I worked on one of the first such studies using this technique (Mortensen, Vasan G.C. et al. 2021), wherein two lensed arcs at $z = 2.92$ were used to spatially probe the CGM of a $z = 0.77$ galaxy. We continuously mapped the spatial and kinematic distribution of the CGM gas using Mg II absorption out to 30 kpc, finding that the CGM gas is mainly dispersion-supported, anisotropic and optically thick, with the absorption strength decreasing with increasing impact parameter. This provides evidence of a reservoir of metal-enriched, dispersion-supported gas recycling through the CGM, serving as the fuel for future star formation and perpetuating the baryon cycle in this galaxy. In the sample of Keck/KCWI observations that I have obtained to probe the mass loss rates and mass loading factors, I have identified numerous cases of multiple-arc systems at high redshifts or other intervening absorbers at lower redshifts that are optimal for arc tomography. I am currently conducting the first such tomographic study to probe the CGM in the high redshift universe (Vasan G.C. et al., in prep) using these observations and

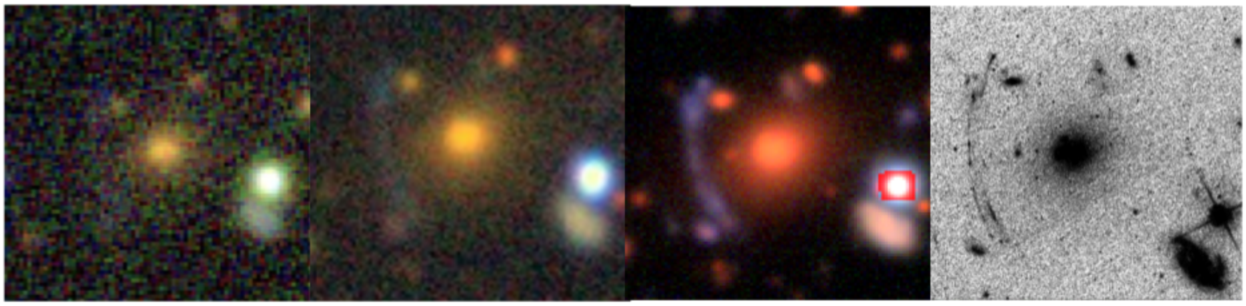

FIGURE 5.1. *(Left)*: Comparison of the image quality of a lensed galaxy found using Machine Learning techniques (Keerthi Vasan et al., 2023) as seen in different ground-based surveys (SDSS, DECaLS and DLS). The depth and sensitivity of the images increases from left to right. *(Right)*: Diffraction-limited space-based (HST) observations of the same lensed galaxy showing a typical lensing morphology with numerous multiply imaged star-forming regions. Euclid and the soon-to-be-launched Roman missions will image thousands of lensed galaxies at diffraction-limited angular resolutions and good depths. Additionally, these campaigns will be complemented by wide-area deep imaging from telescopes such as Rubin Observatory, and the availability of sensitive instruments on the upcoming ELTs. Our understanding of galaxy evolution will immensely benefit from the detailed explorations that these upcoming missions will enable.

aim to extend this technique to a larger sample of galaxies across cosmic time to statistically probe the CGM gas surrounding galaxies in the future.

This is an absolutely exciting time to be working on strong gravitationally lensed galaxies. The number of known lensed galaxies has significantly increased due to state-of-the-art machine learning methods such as those described in Chapter 2. Using these methods on the diffraction-limited, resolved, deep and multi-band observations from wide-area survey missions such as Rubin/LSST, Euclid and Roman will dramatically increase the number of known highly-magnified systems. A large sample of lensed galaxies discovered from these surveys will probe fainter sources (low-mass/high-z) and smaller Einstein Radii (low-mass deflectors) which will allow us to study galaxies across a wide range of galaxy properties at spectacular angular resolution. Figure 5.1 compares the image quality of a lensed galaxy observed with multi-band, ground-based imaging survey at good depth (e.g., DLS) and diffraction-limited space-based telescope (HST) with moderate exposure time demonstrating the significant gain we will derive from upcoming wide-area surveys. Chapters 3 and 4 have clearly demonstrated the power of good spectral resolution and sensitive IFS observations to characterize the effects of feedback on the ISM and CGM at high redshifts. The upcoming sensitive IFS instruments equipped with AO capabilities on 30m class telescopes which are currently being built, will have better light-gathering capability and angular resolution compared to the current ground-based telescopes. Our understanding of the formation and evolution of galaxies will benefit immensely from well resolved observations of lensed (and unlensed) galaxies using the numerous IFS instruments that will be on-board these telescopes.

Bibliography

Adebayo J., Gilmer J., Muelly M., Goodfellow I., Hardt M., Kim B., 2018, in Proceedings of the 32nd International Conference on Neural Information Processing Systems. NIPS'18. Curran Associates Inc., Red Hook, NY, USA, p. 9525–9536

Alard C., 2006, arXiv e-prints, pp astro–ph/0606757

Arjovsky M., Chintala S., Bottou L., 2017, in Precup D., Teh Y. W., eds, Proceedings of Machine Learning Research Vol. 70, Proceedings of the 34th International Conference on Machine Learning. PMLR, International Convention Centre, Sydney, Australia, pp 214–223, `http://proceedings.mlr.press/v70/arjovsky17a.html`

Ascaso B., Wittman D., Dawson W., 2014, MNRAS, 439, 1980

Becker G. D., Rauch M., Sargent W. L. W., 2009, ApJ, 698, 1010

Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, MNRAS, 488, 3143

Belfiore F., Vincenzo F., Maiolino R., Matteucci F., 2019, MNRAS, 487, 456

Belokurov V., Evans N. W., Hewett P. C., Moiseev A., McMahon R. G., Sanchez S. F., King L. J., 2009, MNRAS, 392, 104

Berg T. A. M., et al., 2016, MNRAS, 463, 3021

Berthelot D., Carlini N., Goodfellow I., Papernot N., Oliver A., Raffel C. A., 2019, in Wallach H., Larochelle H., Beygelzimer A., d'Alché-Buc F., Fox E., Garnett R., eds, , Advances in Neural Information Processing Systems 32. Curran Associates, Inc., pp 5049–5059

Berthelot D., Roelofs R., Sohn K., Carlini N., Kurakin A., 2021, arXiv preprint arXiv:2106.04732

Bertin E., Arnouts S., 1996, Astronomy and Astrophysics Supplement, 117, 393

Birrer S., Amara A., 2018, Physics of the Dark Universe, 22, 189

Birrer S., et al., 2021, JOSS, 6, 3283

Bolton A. S., Burles S., Koopmans L. V. E., Treu T., Moustakas L. A., 2006, ApJ, 638, 703

Bolton A. S., Burles S., Koopmans L. V. E., Treu T., Gavazzi R., Moustakas L. A., Wayth R., Schlegel D. J., 2008, ApJ, 682, 964

Bordoloi R., et al., 2014, ApJ, 794, 130

Bordoloi R., Rigby J. R., Tumlinson J., Bayliss M. B., Sharon K., Gladders M. G., Wuyts E., 2016, Monthly Notices of the Royal Astronomical Society, 458, 1891

Bouché N., et al., 2010, ApJ, 718, 1001

Bradač M., Schneider P., Steinmetz M., Lombardi M., King L. J., Porcas R., 2002, A&A, 388, 373

Brownstein J. R., et al., 2012, ApJ, 744, 41

Bruzual G., Charlot S., 2003, MNRAS, 344, 1000

Bullock J. S., Boylan-Kolchin M., 2017, AR&AA, 55, 343

Burchett J. N., Rubin K. H. R., Prochaska J. X., Coil A. L., Vaught R. R., Hennawi J. F., 2021a, ApJ, 909, 151

Burchett J. N., Rubin K. H. R., Prochaska J. X., Coil A. L., Vaught R. R., Hennawi J. F., 2021b, ApJ, 909, 151

Cañameras R., et al., 2020, A&A, 644, A163

Chabrier G., 2003, PASP, 115, 763

Chan T. K., Kereš D., Hopkins P. F., Quataert E., Su K. Y., Hayward C. C., Faucher-Giguère C. A., 2019, MNRAS, 488, 3716

Chattopadhay A., Sarkar A., Howlader P., Balasubramanian V. N., 2018, in 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp 839–847, doi:10.1109/WACV.2018.00097

Chiba M., 2002, ApJ, 565, 17

Chisholm J., Tremonti C. A., Leitherer C., Chen Y., Wofford A., Lundgren B., 2015, ApJ, 811, 149

Chisholm J., Tremonti C. A., Leitherer C., Chen Y., Wofford A., 2016, MNRAS, 457, 3133

Chisholm J., Tremonti C. A., Leitherer C., Chen Y., 2017, MNRAS, 469, 4831

Chisholm J., Bordoloi R., Rigby J. R., Bayliss M., 2018, MNRAS, 474, 1688

Chisholm J., Rigby J. R., Bayliss M., Berg D. A., Dahle H., Gladders M., Sharon K., 2019, ApJ, 882, 182

Churchill C. W., Trujillo-Gomez S., Nielsen N. M., Kacprzak G. G., 2013, ApJ, 779, 87

Cicone C., Maiolino R., Marconi A., 2016, A&A, 588, A41

Cimatti A., Fraternali F., Nipoti C., 2019, Introduction to Galaxy Formation and Evolution: From Primordial Gas to Present-Day Galaxies. Cambridge University Press

Claeskens J.-F., Sluse D., Riaud P., Surdej J., 2006, A&A, 451, 865

Collett T. E., 2015, ApJ, 811, 20

Cowie L. L., Hu E. M., Songaila A., 1995, AJ, 110, 1576

Davé R., Oppenheimer B. D., Finlator K., 2011, MNRAS, 415, 11

Dessauges-Zavadsky M., Christensen L., D'Odorico S., Schaerer D., Richard J., 2011, A&A, 533, A15

Di Matteo T., Springel V., Hernquist L., 2005, Nature, 433, 604

Diehl H. T., et al., 2009, ApJ, 707, 686

Du X., et al., 2018, ApJ, 860, 75

Dutton A. A., et al., 2011, MNRAS, 417, 1621

Elmegreen B. G., Elmegreen D. M., 2010, ApJ, 722, 1895

Elmegreen D. M., Elmegreen B. G., Hirst A. C., 2004, ApJL, 604, L21

Elmegreen D. M., Elmegreen B. G., Rubin D. S., Schaffer M. A., 2005, The Astrophysical Journal, 631, 85

Erb D. K., Quider A. M., Henry A. L., Martin C. L., 2012, ApJ, 759, 26

Erhan D., Courville A., Bengio Y., Vincent P., 2010, in Teh Y. W., Titterington M., eds, Proceedings of Machine Learning Research Vol. 9, Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. PMLR, Chia Laguna Resort, Sardinia, Italy, pp 201–208, http://proceedings.mlr.press/v9/erhan10a.html

Ertl S., Schuldt S., Suyu S. H., Schmidt T., Treu T., Birrer S., Shajib A. J., Sluse D., 2023, A&A, 672, A2

Fabian A. C., 2012, AR&AA, 50, 455

Fassnacht C. D., Moustakas L. A., Casertano S., Ferguson H. C., Lucas R. A., Park Y., 2004, ApJL, 600, L155

Fernando Z. T., Singh J., Anand A., 2019, in Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, doi:10.1145/3331184.3331312, https://doi.org/10.1145%2F3331184.3331312

Fielding D., Quataert E., Martizzi D., 2018, MNRAS, 481, 3325

Finley H., et al., 2017a, A&A, 605, A118

Finley H., et al., 2017b, A&A, 605, A118

Förster Schreiber N. M., et al., 2009, ApJ, 706, 1364

Förster Schreiber N. M., et al., 2018, ApJS, 238, 21

Frenk C. S., White S. D. M., 2012, Annalen der Physik, 524, 507

Frye B., Broadhurst T., Benítez N., 2002, ApJ, 568, 558

Fynbo J. P. U., et al., 2013, MNRAS, 436, 361

Garvin E. O., Kruk S., Cornen C., Bhatawdekar R., Cañameras R., Merín B., 2022, arXiv e-prints, p. arXiv:2207.06997

Gavazzi R., Marshall P. J., Treu T., Sonnenfeld A., 2014, ApJ, 785, 144

Gilman D., Birrer S., Treu T., Nierenberg A., Benson A., 2019, MNRAS, 487, 5721

Goodfellow I. J., Pouget-Abadie J., Mirza M., Xu B., Warde-Farley D., Ozair S., Courville A., Bengio Y., 2014, in Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2. NIPS'14. MIT Press, Cambridge, MA, USA, p. 2672–2680

Gu J., et al., 2018, Pattern Recognition, 77, 354

Gulrajani I., Ahmed F., Arjovsky M., Dumoulin V., Courville A. C., 2017, in Guyon I., Luxburg U. V., Bengio S., Wallach H., Fergus R., Vishwanathan S., Garnett R., eds, Vol. 30, Advances in Neural Information Processing Systems. Curran Associates, Inc., https://proceedings.neurips.cc/paper/2017/file/892c3b1c6dccd52936e27cbd0ff683d6-Paper.pdf

Hayward C. C., Hopkins P. F., 2017, MNRAS, 465, 1682

He K., Zhang X., Ren S., Sun J., 2016a, in European conference on computer vision. pp 630–645

He K., Zhang X., Ren S., Sun J., 2016b, in Proceedings of the IEEE conference on computer vision and pattern recognition. pp 770–778

Heckman T. M., 2002, in Mulchaey J. S., Stocke J. T., eds, Astronomical Society of the Pacific Conference Series Vol. 254, Extragalactic Gas at Low Redshift. p. 292 (arXiv:astro-ph/0107438)

Heckman T. M., Borthakur S., 2016, ApJ, 822, 9

Heckman T. M., Armus L., Miley G. K., 1990, ApJS, 74, 833

Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, MNRAS, 445, 581

Hu W., et al., 2023, arXiv e-prints, p. arXiv:2307.04911

Huang X., et al., 2020, The Astrophysical Journal, 894, 78

Hummels C. B., Smith B. D., Silvia D. W., 2017, ApJ, 847, 59

Ioffe S., Szegedy C., 2015, in International conference on machine learning. pp 448–456

Ivezić Ž., et al., 2019, ApJ, 873, 111

Jacobs C., Glazebrook K., Collett T., More A., McCarthy C., 2017, MNRAS, 471, 167

Jacobs C., et al., 2019a, ApJS, 243, 17

Jacobs C., et al., 2019b, MNRAS, 484, 5330

James B. L., Auger M., Pettini M., Stark D. P., Belokurov V., Carniani S., 2018a, MNRAS, 476, 1726

James B. L., Auger M., Pettini M., Stark D. P., Belokurov V., Carniani S., 2018b, MNRAS, 476, 1726

Jones T., Stark D. P., Ellis R. S., 2012, ApJ, 751, 51

Jones T., Ellis R. S., Richard J., Jullo E., 2013a, ApJ, 765, 48

Jones T., Ellis R. S., Richard J., Jullo E., 2013b, ApJ, 765, 48

Jones T. A., Ellis R. S., Schenker M. A., Stark D. P., 2013c, ApJ, 779, 52

Jones T., Stark D. P., Ellis R. S., 2018, ApJ, 863, 191

Kacprzak G. G., Cooke J., Churchill C. W., Ryan-Weber E. V., Nielsen N. M., 2013, ApJL, 777, L11

Kassiola A., Kovner I., 1993, ApJ, 417, 450

Keerthi Vasan G. C., Sheng S., Jones T., Choi C. P., Sharpnack J., 2023, MNRAS, 524, 5368

King A., 2003, ApJL, 596, L27

Kingma D. P., Welling M., 2014, Auto-Encoding Variational Bayes (`arXiv:1312.6114`)

Kneib J.-P., Natarajan P., 2011, The Astronomy and Astrophysics Review, 19, 47

Koekemoer A. M., et al., 2007, ApJS, 172, 196

Kolb E. W., Turner M. S., 1990, The early universe. Vol. 69

Koopmans L. V. E., Treu T., Bolton A. S., Burles S., Moustakas L. A., 2006, ApJ, 649, 599

Kormann R., Schneider P., Bartelmann M., 1994, A&A, 284, 285

Krizhevsky A., 2009a, Technical report, Learning multiple layers of features from tiny images

Krizhevsky A., 2009b.

Krizhevsky A., Sutskever I., Hinton G. E., 2012, in Pereira F., Burges C. J. C., Bottou L., Weinberger K. Q., eds, Vol. 25, Advances in Neural Information Processing Systems. Curran Associates, Inc., https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf

Kubo J. M., Dell'Antonio I. P., 2008, MNRAS, 385, 918

LSST Science Collaboration et al., 2009, arXiv e-prints, p. arXiv:0912.0201

Laine S., Aila T., 2017, ArXiv, abs/1610.02242

Laureijs R., et al., 2011, arXiv e-prints, p. arXiv:1110.3193

LeCun Y., Boser B., Denker J. S., Henderson D., Howard R. E., Hubbard W., Jackel L. D., 1989, Neural Computation, 1, 541

Lee D.-H., 2013, ICML 2013 Workshop : Challenges in Representation Learning (WREPL)

Leethochawalit N., Jones T. A., Ellis R. S., Stark D. P., Richard J., Zitrin A., Auger M., 2016a, ApJ, 820, 84

Leethochawalit N., Jones T. A., Ellis R. S., Stark D. P., Richard J., Zitrin A., Auger M., 2016b, ApJ, 820, 84

Leethochawalit N., Jones T. A., Ellis R. S., Stark D. P., Richard J., Zitrin A., Auger M., 2016c, The Astrophysical Journal, 820, 84

Leethochawalit N., Jones T. A., Ellis R. S., Stark D. P., Zitrin A., 2016d, ApJ, 831, 152

Leethochawalit N., Kirby E. N., Ellis R. S., Moran S. M., Treu T., 2019, ApJ, 885, 100

Lehnert M. D., Heckman T. M., Weaver K. A., 1999, ApJ, 523, 575

Leitherer C., et al., 1999, ApJS, 123, 3

Li R., et al., 2020, ApJ, 899, 30

Li J. I.-H., et al., 2023, The Sloan Digital Sky Survey Reverberation Mapping Project: The Black Hole Mass−Stellar Mass Relations at $0.2 \lesssim z \lesssim 0.8$ (arXiv:2023arXiv230104177L), doi:10.48550/arXiv.2301.04177, https://ui.adsabs.harvard.edu/abs/2023arXiv230104177L

Lilly S. J., Carollo C. M., Pipino A., Renzini A., Peng Y., 2013, ApJ, 772, 119

Litjens G., et al., 2017, Medical Image Analysis, 42, 60

Lopez S., et al., 2018a, Nature, 554, 493

Lopez S., et al., 2018b, Nature, 554, 493

Lopez S., et al., 2020, MNRAS, 491, 4442

Lupton R., Blanton M. R., Fekete G., Hogg D. W., O'Mullane W., Szalay A., Wherry N., 2004, PASP, 116, 133

Ma J., Ge J., Zhao Y., Prochaska J. X., Zhang S., Ji T., Schneider D. P., 2017, MNRAS, 472, 2196

Madau P., Dickinson M., 2014, AR&AA, 52, 415

Madireddy S., Li N., Ramachandra N., Butler J., Balaprakash P., Habib S., Heitmann K., 2019, arXiv e-prints, p. arXiv:1911.03867

Mahler G., et al., 2018, MNRAS, 473, 663

Mainali R., Stark D. P., Jones T., Ellis R. S., Hezaveh Y. D., Rigby J. R., 2023a, MNRAS, 520, 4037

Mainali R., Stark D. P., Jones T., Ellis R. S., Hezaveh Y. D., Rigby J. R., 2023b, Monthly Notices of the Royal Astronomical Society, 520, 4037

Marques-Chaves R., et al., 2017, ApJL, 834, L18

Marra R., Churchill C. W., Kacprzak G. G., Nielsen N. M., Trujillo-Gomez S., Lewis J. G., 2022, arXiv e-prints, p. arXiv:2202.12228

Marshall P., Sandford C., More A., Buddelmeijerr H., 2015, HumVI: Human Viewable Image creation (ascl:1511.014)

Martin C. L., 2005, ApJ, 621, 227

Martin C. L., Shapley A. E., Coil A. L., Kornei K. A., Bundy K., Weiner B. J., Noeske K. G., Schiminovich D., 2012, ApJ, 760, 127

McKee C. F., Ostriker J. P., 1977, ApJ, 218, 148

Miranda M., Macciò A. V., 2007, MNRAS, 382, 1225

Miyato T., Maeda S., Koyama M., Ishii S., 2019, IEEE Transactions on Pattern Analysis and Machine Intelligence, 41, 1979

More A., et al., 2016, MNRAS, 455, 1191

Morrissey P., et al., 2018, The Astrophysical Journal, 864, 93

Mortensen K., Keerthi Vasan G. C., Jones T., Faucher-Giguère C.-A., Sanders R. L., Ellis R. S., Leethochawalit N., Stark D. P., 2021a, ApJ, 914, 92

Mortensen K., Keerthi Vasan G. C., Jones T., Faucher-Giguère C.-A., Sanders R. L., Ellis R. S., Leethochawalit N., Stark D. P., 2021b, ApJ, 914, 92

Moustakas L. A., et al., 2007, ApJL, 660, L31

Muratov A. L., Kereš D., Faucher-Giguère C.-A., Hopkins P. F., Quataert E., Murray N., 2015, MNRAS, 454, 2691

Muratov A. L., et al., 2017, MNRAS, 468, 4170

Murray N., Quataert E., Thompson T. A., 2005, ApJ, 618, 569

Murray N., Ménard B., Thompson T. A., 2011, ApJ, 735, 66

Naab T., Ostriker J. P., 2017, AR&AA, 55, 59

Narayan R., Bartelmann M., 1996, arXiv e-prints, pp astro–ph/9606001

Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563

Navarro J. F., Frenk C. S., White S. D. M., 1997, ApJ, 490, 493

Nelson D., et al., 2019, MNRAS, 490, 3234

Netzer Y., Wang T., Coates A., Bissacco A., Wu B., Ng A. Y., 2011, in NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011. `http://ufldl.stanford.edu/housenumbers/nips2011_housenumbers.pdf`

Newman A. B., Ellis R. S., Treu T., 2015, ApJ, 814, 26

Nielsen N. M., Churchill C. W., Kacprzak G. G., Murphy M. T., 2013a, ApJ, 776, 114

Nielsen N. M., Churchill C. W., Kacprzak G. G., 2013b, ApJ, 776, 115

Oguri M., 2010, Publications of the Astronomical Society of Japan, 62, 1017

Oguri M., 2021, PASP, 133, 074504

Oguri M., Marshall P. J., 2010, MNRAS, 405, 2579

Pandya V., et al., 2021, arXiv e-prints, p. arXiv:2103.06891

Paraficz D., et al., 2016, A&A, 592, A75

Peebles P. J. E., 1980, The large-scale structure of the universe

Peeples M. S., et al., 2019, ApJ, 873, 129

Péroux C., Howk J. C., 2020, AR&AA, 58, 363

Petrillo C. E., et al., 2019, MNRAS, 484, 3879

Pettini M., Steidel C. C., Adelberger K. L., Dickinson M., Giavalisco M., 2000, ApJ, 528, 96

Pettini M., Rix S. A., Steidel C. C., Adelberger K. L., Hunt M. P., Shapley A. E., 2002a, ApJ, 569, 742

Pettini M., Rix S. A., Steidel C. C., Adelberger K. L., Hunt M. P., Shapley A. E., 2002b, ApJ, 569, 742

Pourrahmani M., Nayyeri H., Cooray A., 2018, ApJ, 856, 68

Prochaska J. X., Kasen D., Rubin K., 2011, ApJ, 734, 24

Quider A. M., Pettini M., Shapley A. E., Steidel C. C., 2009, MNRAS, 398, 1263

Quider A. M., Shapley A. E., Pettini M., Steidel C. C., Stark D. P., 2010, MNRAS, 402, 1467

Quinonero-Candela J., Sugiyama M., Schwaighofer A., Lawrence N. D., 2008, Dataset shift in machine learning. Mit Press

Quiret S., et al., 2016, MNRAS, 458, 4074

Recchia S., Blasi P., Morlino G., 2016, MNRAS, 462, 4227

Richard J., Jones T., Ellis R., Stark D. P., Livermore R., Swinbank M., 2011, MNRAS, 413, 643

Rigby J. R., et al., 2018, AJ, 155, 104

Rojas K., et al., 2022, A&A, 668, A73

Rubin K. H. R., Prochaska J. X., Koo D. C., Phillips A. C., Martin C. L., Winstrom L. O., 2014, ApJ, 794, 156

Rudie G. C., Steidel C. C., Pettini M., Trainor R. F., Strom A. L., Hummels C. B., Reddy N. A., Shapley A. E., 2019, ApJ, 885, 61

Rupke D. S., Veilleux S., Sanders D. B., 2005, ApJS, 160, 87

Sahu N., et al., 2024, arXiv e-prints, p. arXiv:2405.15427

Sanders R. L., et al., 2021, ApJ, 914, 19

Schmidt S. J., Thorman P., 2013, MNRAS, 431, 2766

Schneider P., Ehlers J., Falco E. E., 1992, Gravitational Lenses, doi:10.1007/978-3-662-03758-4.

Schroetter I., et al., 2016, ApJ, 833, 39

Seidel G., Bartelmann M., 2007, A&A, 472, 341

Selvaraju R. R., Cogswell M., Das A., Vedantam R., Parikh D., Batra D., 2017, in Proceedings of the IEEE International Conference on Computer Vision (ICCV).

Shaban A., et al., 2022, The Astrophysical Journal, 936, 77

Shaban A., et al., 2023, Monthly Notices of the Royal Astronomical Society, 526, 6297

Shajib A. J., et al., 2022, arXiv e-prints, p. arXiv:2210.10790

Shapley A. E., Steidel C. C., Pettini M., Adelberger K. L., 2003, ApJ, 588, 65

Sheinis A. I., Bolte M., Epps H. W., Kibrick R. I., Miller J. S., Radovan M. V., Bigelow B. C., Sutin B. M., 2002, PASP, 114, 851

Sheng S., C K. V. G., Choi C. P., Sharpnack J., Jones T., 2022, arXiv e-prints, p. arXiv:2210.11681

Sohn K., et al., 2020, in Larochelle H., Ranzato M., Hadsell R., Balcan M., Lin H., eds, Vol. 33, Advances in Neural Information Processing Systems. Curran Associates, Inc., pp 596–608, `https://proceedings.neurips.cc/paper/2020/file/`

06964dce9addb1c5cb5d6e3d9838f733-Paper.pdf

Solimano M., et al., 2022, arXiv e-prints, p. arXiv:2206.02949

Somerville R. S., Davé R., 2015a, AR&AA, 53, 51

Somerville R. S., Davé R., 2015b, AR&AA, 53, 51

Sonnenfeld A., Treu T., Gavazzi R., Suyu S. H., Marshall P. J., Auger M. W., Nipoti C., 2013, ApJ, 777, 98

Sonnenfeld A., et al., 2018, Publications of the Astronomical Society of Japan, 70, S29

Speagle J. S., 2019, arXiv e-prints, p. arXiv:1904.02180

Speagle J. S., Steinhardt C. L., Capak P. L., Silverman J. D., 2014, ApJS, 214, 15

Spergel D., et al., 2015, arXiv e-prints, p. arXiv:1503.03757

Spilker J. S., et al., 2020, ApJ, 905, 85

Springenberg J. T., Dosovitskiy A., Brox T., Riedmiller M. A., 2015, CoRR, abs/1412.6806

Stark D. P., Swinbank A. M., Ellis R. S., Dye S., Smail I. R., Richard J., 2008, Nature, 455, 775

Stark D. P., et al., 2013, MNRAS, 436, 1040

Steidel C. C., Erb D. K., Shapley A. E., Pettini M., Reddy N., Bogosavljević M., Rudie G. C., Rakic O., 2010, ApJ, 717, 289

Stein G., Harrington P., Blaum J., Medan T., Lukic Z., 2021, arXiv e-prints, p. arXiv:2110.13151

Sugahara Y., Ouchi M., Lin L., Martin C. L., Ono Y., Harikane Y., Shibuya T., Yan R., 2017, ApJ, 850, 51

Sugahara Y., Ouchi M., Harikane Y., Bouché N., Mitchell P. D., Blaizot J., 2019, ApJ, 886, 29

Suyu S. H., et al., 2013, ApJ, 766, 70

Swinbank A. M., et al., 2009a, MNRAS, 400, 1121

Swinbank A. M., et al., 2009b, MNRAS, 400, 1121

Sérsic J. L., 1968, Atlas de Galaxias Australes. http://adsabs.harvard.edu/abs/1968adga.book.....S

Tarvainen A., Valpola H., 2017, in Guyon I., Luxburg U. V., Bengio S., Wallach H., Fergus R., Vishwanathan S., Garnett R., eds, , Advances in Neural Information Processing Systems 30. Curran Associates, Inc., pp 1195–1204

Tran K.-V. H., et al., 2022a, arXiv e-prints, p. arXiv:2205.05307

Tran K.-V. H., et al., 2022b, arXiv e-prints, p. arXiv:2205.05307

Tran K.-V. H., et al., 2022c, The Astronomical Journal, 164, 148

Treu T., 2010, AR&AA, 48, 87

Tsujimoto T., Nomoto K., Yoshii Y., Hashimoto M., Yanagida S., Thielemann F. K., 1995, MNRAS, 277, 945

Tumlinson J., Peeples M. S., Werk J. K., 2017, AR&AA, 55, 389

Vanzella E., et al., 2009, ApJ, 695, 1163

Vasan G. C. K., et al., 2023, The Astrophysical Journal, 959, 124

Veilleux S., Cecil G., Bland-Hawthorn J., 2005a, AR&AA, 43, 769

Veilleux S., Cecil G., Bland-Hawthorn J., 2005b, AR&AA, 43, 769

Verhamme A., Schaerer D., Atek H., Tapken C., 2008, A&A, 491, 89

Wales D. J., Doye J. P. K., 1997, The Journal of Physical Chemistry A, 101, 5111

Weiner B. J., et al., 2009, ApJ, 692, 187

Weldon A., et al., 2022, arXiv e-prints, p. arXiv:2203.09539

Whitaker K. E., van Dokkum P. G., Brammer G., Franx M., 2012, ApJL, 754, L29

Wilson J. C., et al., 2004, in Moorwood A. F. M., Iye M., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 5492, Ground-based Instrumentation for Astronomy. pp 1295–1305, doi:10.1117/12.550925

Wisnioski E., et al., 2015, ApJ, 799, 209

Wittman D. M., et al., 2002, in Tyson J. A., Wolff S., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 4836, Survey and Other Telescope Technologies and Discoveries. pp 73–82 (arXiv:astro-ph/0210118), doi:10.1117/12.457348

Wittman D., Margoniner V. E., Tyson J. A., Cohen J. G., Becker A. C., Dell'Antonio I. P., 2003, ApJ, 597, 218

Wittman D., Dell'Antonio I. P., Hughes J. P., Margoniner V. E., Tyson J. A., Cohen J. G., Norman D., 2006, ApJ, 643, 128

Wizinowich P., et al., 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. p. 114480E, doi:10.1117/12.2560017

Wuyts E., et al., 2010, ApJ, 724, 1182

Wuyts E., Rigby J. R., Gladders M. D., Sharon K., 2014a, ApJ, 781, 61

Wuyts E., Rigby J. R., Gladders M. D., Sharon K., 2014c, ApJ, 781, 61

Wuyts E., Rigby J. R., Gladders M. D., Sharon K., 2014b, The Astrophysical Journal, 781, 61

Xu X., et al., 2022a, arXiv e-prints, p. arXiv:2204.09181

Xu X., et al., 2022b, ApJ, 933, 222

Zhang D., 2018, Galaxies, 6, 114

Zhou B., Khosla A., A. L., Oliva A., Torralba A., 2016, CVPR

Zhuang Z., et al., 2023, ApJ, 948, 132

van den Bergh S., Abraham R. G., Ellis R. S., Tanvir N. R., Santiago B. X., Glazebrook K. G., 1996, AJ, 112, 359