Instability, Perturbations, and Substructure in Protoplanetary Disks

by

Jason Terry

(Under the Direction of Cassandra Hall)

Abstract

This work focuses on the analysis of protoplanetary disks, the sites of planet formation. Numerical, analytic, and machine learning techniques are applied to the topics of gravitational instability, planet-induced perturbations, and non-Keplerian motion in general. This is largely done through kinematic analysis: the quantification of the motion within the disk using line emission observations. Gravitational instability is a process that is thought to be of potential importance early in a disk's lifetime and may facilitate the formation of specific types of planets. I demonstrate that the signatures of gravitational instability contain a rich amount of information that can be difficult to disentangle. Planets in the disk may induce perturbations that are visible in observations. The identification of these perturbations is an important avenue for testing models of planet formation. I use machine learning techniques to find these perturbations and report a previously unidentified planet in the disk HD 142666. Finally, I demonstrate that unsupervised machine learning can be used to identify general non-Keplerian motion or anomalous regions in line emission observations. These results pave the way for a more complete understanding of the influences and effects of gravitational instability, demonstrates the ability of machine learning to identify important features, and introduces a novel method to apply machine learning to astronomical observations across subfields.

Index words: Astrophysics, Protoplanetary disks, Planet formation, Observations, Simulations, Hydrodynamics, Machine learning

Instability, Perturbations, and Substructure in Protoplanetary Disks

by

Jason Terry

B.S., University of Georgia, 2018
M.S., Brown University, 2019

A Dissertation Submitted to the Graduate Faculty of the
University of Georgia in Partial Fulfillment of the Requirements for the
Degree.

Doctor of Philosophy

Athens, Georgia

2024

Instability, Perturbations, and Substructure in Protoplanetary Disks

by

Jason Terry

Major Professor:	Cassandra Hall
Committee:	Cassandra Hall Inseok Song Phillip Stancil

Electronic Version Approved:

Ron Walcott
Dean of the Graduate School
The University of Georgia
May 2024

Dedication

This work is dedicated to Dr. Claude Terry, the professor of everything.

Acknowledgments

I would like to thank my advisor, Dr. Cassandra Hall, without whom this work would have been impossible and my time here much less enjoyable. I owe special thanks to my committee members, Dr. Phillip Stancil and Dr. Inseok Song, for their time, dedication, and help over the years. Thank you to my family for the tireless patience as I perpetually avoided real life. I am grateful for the invaluable help from Shan-Ho Tsai and other employees of the Georgia Advanced Computing Resource Center. Finally, I want to acknowledge Alaine for her unwavering encouragement and unconditional support. The road would have been more difficult and the end less rewarding without her by my side.

Contents

List of Figures

List of Tables

Chapter 1

Protoplanetary Disk Astrophysics

The last three decades have seen a revolution in the study of exoplanets and planet formation. Unprecedented observational capabilities, such as the Kepler space telescope, have revealed a diverse population of over 5,000 exoplanets (Figure 1.1). Most of these discovered exoplanets do not exist in stellar systems similar to ours, disallowing studying the the solar system as a representative example. As the sample size and variety of known planets increases, traditional theories of planet formation must be adapted to account for this diversity (Drążkowska et al., 2023). The interplay of gas, dust, light, gravity, magnetic fields, and chemistry necessary to quickly form planetary systems has thus far eluded a complete explanation able to recreate our observations (Drążkowska et al., 2023).

1.1 Planet formation

Traditional paradigms of planet formation have focused on bottom-up formation via the accretion of "cores" and "pebbles" (Lambrechts & Johansen, 2012; Ormel & Klahr, 2010; Safronov, 1972) and the direct formation of a planet through a collapse induced by gravitational instability (GI) (Kratter & Lodato, 2016). The bottom-up processes rely on the growth of mm-sized pebbles and km-sized planetesimals, ultimately resulting a planet. While questions remain (e.g., the role of dust, the formation of pebbles themselves, and the radial-dependence of efficiency), given the correct initial conditions, accretion methods have been shown to quickly form planets across a range of separations and masses (Johansen & Lambrechts, 2017). GI, on the other hand, tends to be transiently active at wide separations early in the disk lifetime and results in planets on the order of several Jupiter masses (Cadman, Hall, et al., 2020; Cadman

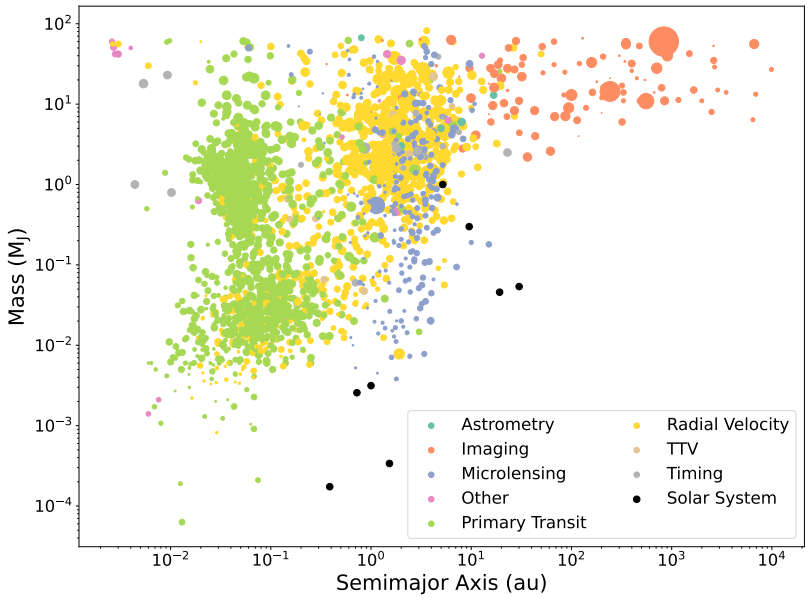

Figure 1.1: Confirmed exoplanets semi-major axis against mass ("NASA Exoplanet Archive", 2024). Points are colored by detection type, and their sizes are based on the host star mass.

et al., 2021; Clarke, 2009; D. Forgan et al., 2010; Nayakshin, 2010). Clearly, GI alone cannot explain the diversity of exoplanets discovered, but it cannot be completely ruled out as a mechanism responsible for some exoplanets (Marois et al., 2008; Vigan et al., 2017).

In order to fully understand the process of planet formation and make sufficiently general theories that can explain all observations, it is necessary to observe the process at all stages. Only by understanding the initial conditions and their ultimate influences can we hope to understand the resulting stellar systems.

Protoplanetary disks are the perfect laboratory for this investigation.

1.2 Protoplanetary disks

Protoplanetary accretion disks are the sites of planet formation. After an initial period of collapse of a molecular cloud, a disk in hydrostatic and rotational equilibrium forms around the protostar. The disk contains both gas and dust. Material within inner disk can accrete onto the protostar and planets can form within the disk itself. These processes require the transport of angular momentum throughout the disk.

Given the wide variety of planet masses and semimajor axes shown in Figure 1.1, there must also be a large range of possible protoplanetary disk masses and complex disk dynamics, including angular momentum transport. Mass is one of the most important fundamental parameters of the disk as it sets the mass budget of the resulting planets and is known to be an important factor in the disk's evolution (Longarini et al., 2021; Terry, Hall, Longarini, et al., 2022). The protostar's mass also influences the disk's dynamics, which is trivially seen by its explicit presence in Equation 1.2. Other varying factors, such as the magnetic field strength and orientation (Wurster, 2021; Wurster et al., 2016, 2021), ionization rate (Wurster et al., 2018), and host star's metallicity (Santos et al., 2004), have been found to influence disk behavior and the resulting system architecture. These complex environments evidently require careful study in order to understand the interplay and effects of the disk's physical and environmental characteristics.

Laminar Keplerian motion is good approximation for a disk without disturbance. This is due to the dominance of the gravitational potential of the central protostar. The angular velocity, Ω, in a Keplerian orbit is entirely determined by the mass of the protostar and the distance from it (Equation 1.2). In reality, the gas is partially supported by pressure, making the gas disk slightly sub-Keplerian.

$$\Omega_{\rm K} = \sqrt{\frac{GM_\star}{r^3}}$$

Angular momentum transport, necessary for accretion, requires the breaking of laminar Keplerian motion. Without dissipative effects, such as turbulence or viscosity, the material within the disk will simply continue to orbit as a laminar flow. By decreasing or redistributing angular momentum through mechanisms, such as magnetic torques or gravitational instability, gas is able to accrete onto the central protostar and form substructures within the disk itself. In a Keplerian disk, angular momentum increases with distance from the star (seen in Equation 1.2), so material that loses momentum will drift inwards whereas material that gains angular momentum will move outwards.

Due to the intractability of explicitly accounting for all sources of angular momentum transport in a disk, it is often necessary to rely on effective viscous disk theory that uses the α-prescription (Shakura & Sunyaev, 1973). In this model, angular momentum transport within the disk is approximated as a fluid viscosity, ν, that can be parameterized with a single dimensionless constant, α (Armitage, 2022; Shakura & Sunyaev, 1973).

$$\nu = \alpha c_s H,$$

where c_s is the sound speed, and $H \equiv \frac{c_s}{\Omega_K}$. Larger values of α imply higher efficiency of angular momentum transport (and faster cooling (Gammie, 2001)). A typically assumed value is $\alpha = 10^{-2}$ (W. K. M. Rice & Armitage, 2009), although it can vary by several orders of magnitude. At sufficiently high values of α ($\gtrsim 0.06$), the disk will tend to fragment into bound objects (W. K. M. Rice & Armitage, 2009). Smaller values ($\alpha \leq 10^{-4}$) are associated with very low viscosity, and those regions can become dead zones with little angular momentum transport (W. K. M. Rice & Armitage, 2009). The benefit of this model is that disk properties may be calculated without regards to the specifics of the angular momentum transfer mechanisms.

1.3 Observations

Observations are the most important method we have to study protoplanetary disks. While simulations have become extremely complex, e.g., global 3D magnetohydrodynamics (Price et al., 2018) with millions of particles, their results always require benchmarking against reality as seen in observations. The Atacama Large Millimetre/sub-Millimetre Array (ALMA) is one of the most

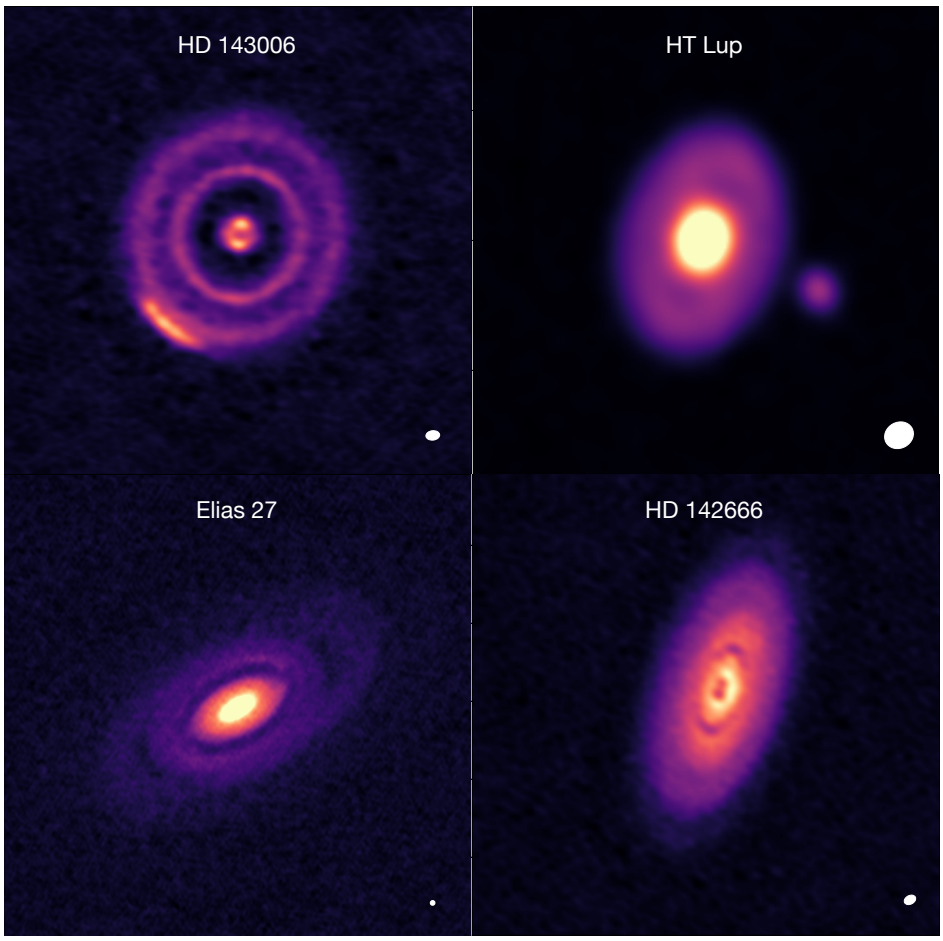

Figure 1.2: Subset of disks observed by ALMA at 1.25 mm for the Disk Substructures at High Angular Resolution Project (DSHARP) (Andrews et al., 2018). Upper left: HD 143006. Upper right: HT Lup. Lower left: Elias 27. Lower right: HD 142666.

powerful tools that has observed protoplanetary disks. It has produced unprecedented high-resolution observations that reveal the morphology of disks. Example continuum (1.25 mm) observations of disks are shown in Figure 1.2, revealing a diversity of appearances. The presence of many different morphological features, including substructures, (see, e.g., Andrews et al., 2018; J. Huang, Andrews, Dullemond, et al., 2018; J. Huang, Andrews, Pérez, et al., 2018) is further confirmation that the environments and dynamics of disks is extremely complex and varied.

An essential tool for observing protoplanetary disks is integral field spectroscopy (IFS). Rather than simply creating a flat image centered on a given wavelength, an IFS observation results in a three dimensional data cube with a spectrum at each pixel. IFS is used widely across astronomy, e.g., (Disk Dynamics Collaboration et al., 2020; Seitenzahl et al., 2018), but the wavelength range and spatial/spectral resolution of ALMA has given unprecedented insight into the environments of protoplanetary disks.

1.4 Kinematic analysis

The kinematic analysis of protoplanetary disks is a powerful tool to quantify and understand the motion of material within a disk in an effort to infer its conditions and any embedded objects (Disk Dynamics Collaboration et al., 2020). Line emission profiles from IFS observations are used to quantify the motion of a given region. Each wavelength ("channel") in the cube corresponds to a different velocity relative to the system and our line of sight. For an inclined disk ($i \neq 0°, 90°$), the radial component of material's motion relative to the observer will vary in different regions. Due to the Doppler shift, a given wavelength will therefore be observed with differing intensity across the disk, giving information about the relative motion of material within the disk. The observed intensity for that wavelength is called a "channel map." Examples of these maps can be seen in Figure 1.3 and Figure 1.4. The channel with the velocity of system itself is called the "systemic channel" and is often taken as a reference for all other channels. Disregarding the motion of the disk as a whole, the material in this channel is moving completely perpendicular to our line of sight and therefore has no additional radial velocity component.

A major purpose of this investigation into is to understand how processes and bodies affect the dynamics of the disk itself. Of particular interest is its ability to identify exoplanets embedded in the parent disk during or soon after their formation. This gives unique insight into the mechanisms behind planet formation. Importantly, kinematic analysis is most sensitive to disturbances at

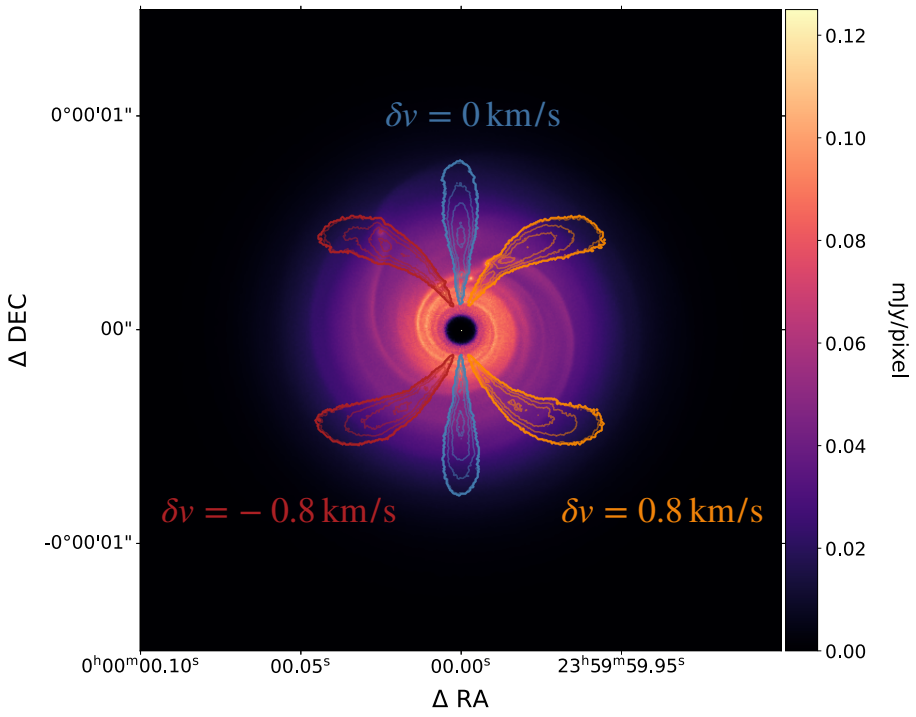

Figure 1.3: 1.25 mm continuum image overlaid with 3 different velocity channel maps. The velocities are reported relative to the system itself (systemic channel in blue).

tens of au (Disk Dynamics Collaboration et al., 2020). This is the region in which both GI and pebble accretion are thought to be efficient (Clarke, 2009; Johansen & Lambrechts, 2017). Kinematic analysis therefore gives us the ability to differentiate between the two mechanisms and understand the extent to which each plays a role in the process of planet formation.

1.5 Instabilities

An disk is undergoing an instability if a process enters a self-reinforcing, runaway regime. These instabilities are thought to be of tremendous importance for angular momentum transport, substructure formation, and planet formation (Armitage, 2022; Balbus & Hawley, 1991; Barraza-Alfaro et al., 2021; Drążkowska & Alibert, 2017; Drążkowska et al., 2023; Hall et al., 2017). Many potentially important instabilities have been put forth as influential processes in PPDs. These include gravitational instability (GI) (Boss, 1997), magnetorotational instability (MRI) (Balbus & Hawley, 1991), streaming instability (Youdin & Goodman, 2005), vertical shear instability (VSI) (R. P. Nelson et al., 2013), and the gravomagneto limit cycle instability (Armitage et al., 2001). While these all play potentially important roles in planet formation and disk evolution, GI is the main focus.

1.5.1 Gravitational instability

Gravitational instability has been put forth as a potential formation mechanism for some exoplanets, particularly massive planets (several M_J) at wide orbits (tens of au) (Boss, 1997, 1998; Cadman, Hall, et al., 2020; Cadman et al., 2021; Clarke, 2009; D. Forgan et al., 2010; Nayakshin, 2010). A global instability, it arises early in the disk's lifetime when the disk and star are of comparable mass (Hall et al., 2020). As the disk is cleared through accretion and ejection, the mass ratio ($q = M_{disk}/M_\star$) decreases, and the disk's self-gravity becomes negligible compared to the star's gravity. During this period, non-axisymmetric perturbations can generate spiral arms, which effectively transport angular momentum and may fragment into gravitationally bound objects (Durisen et al., 2007; Gammie, 2001; W. K. M. Rice, Armitage, Bate, & Bonnell, 2003; W. K. M. Rice, Armitage, Bonnell, et al., 2003). GI has been proposed as the cause of the spiral substructure seen in Elias 2-27 (Meru et al., 2017) and the formation mechanism of some observed massive objects (Cadman et al., 2021; Marois et al., 2008; Nayakshin et al., 2020; Tsukagoshi et al., 2019).

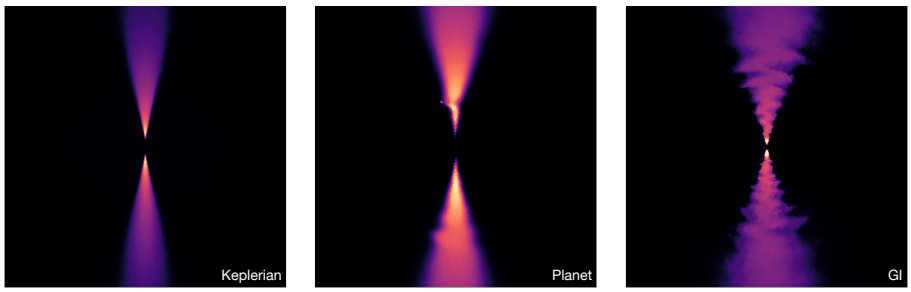

Figure 1.4: Systemic ($\delta v = 0\,\mathrm{km/s}$) channel map for a Keplerian disk (left), a disk with a planet (center), and a GI disk (right).

1.6 Perturbations

The identification of perturbations from Keplerian motion is at the heart of kinematic analysis. Bodies and processes all have characteristic imprints on the motion within the disk (Barraza-Alfaro et al., 2021; Bollati et al., 2021; Hall et al., 2020; Pinte et al., 2023; Teague et al., 2018). Planets, for example, tend to leave localized kinks (Bollati et al., 2021; Teague et al., 2018). GI disks, on the other hand, exhibit global perturbations (the "GI wiggle") across all azimuths and radii (Hall et al., 2020). This differs from Keplerian motion, which results in a "butterfly" pattern across the disk without any perturbations. Figure 1.4 shows example systemic channel maps for these three types of disk.

Analyzing the perturbations in the motion as shown in channel maps gives insight into the environment of the disk. For example, planets can be identified by their perturbations, and the location and extent of the perturbation can be used to infer the planet's semimajor axis and mass (Bollati et al., 2021; Pinte et al., 2018, 2019; Teague et al., 2018).

1.7 Substructure

Observations of disks, e.g., Figure 1.2, have shown that substructures are common and diverse. Substructures can give important insights into the physics and conditions within the disk. Spiral arms can be excited by gravitational instability (Hall et al., 2017). Gaps can be cleared by planets (ALMA Partnership et al., 2015). The streaming instability can produce long filaments (Youdin & Goodman, 2005).

These substructures are not only valuable probes of the disk, but they also have potential and significant ramifications for planet formation (Drążkowska et al., 2023; Toci et al., 2021; Youdin & Goodman, 2005; Zormpas et al., 2022). For example, the filaments created by the streaming instability are efficient regions for planetesimal formation. In general, substructures have been shown to decrease the rate of radial drift of dust (Toci et al., 2021). Radial drift is a result of the interaction of the Keplerian dust and slightly sub-Keplerian gas wherein the drag induced on the dust reduces its angular momentum, causing it to drift inwards. The "radial drift barrier" is a product of this that has caused significant problems for explaining planetesimal formation by dust growth because the dust drifts more quickly than it aggregates (Drążkowska et al., 2023). The decreased efficiency of radial drift in substructures therefore provides a possible mechanism for overcoming the drift barrier, facilitating planet formation (Toci et al., 2021).

1.8 Simulations

Simulations provide an essential tool to develop and test theories about disk evolution and planet formation. We can more fully understand observations when they are compared to simulations. They provide information on the influences of different physical processes, companion bodies, environmental conditions, and so on, giving us a "laboratory" that furthers our insights.

1.8.1 Smoothed-particle hydrodynamics

Simulations of disks frequently rely on hydrodynamics in which the gas and dust are treated as fluids (Benítez-Llambay & Masset, 2016; Price et al., 2018; Stone et al., 2020). Broadly, there are two general types of hydrodynamic methods: grid-based and smoothed-particle hydrodynamics (SPH). This work largely focuses on SPH, in particular, the open-source code `PHANTOM` (Price et al., 2018).

In SPH, the fluid is represented as a discrete set of Langrangian particles (Gingold & Monaghan, 1977; Lucy, 1977). Physical values are calculated through interpolation between particles, which approximately recreates a fluid. A strong benefit of this method is that individual Lagrangian particles may be tracked over time. This is in contrast to grid-based codes, which are Eulerian, and track values in a given region instead of a given particle. Furthermore, treating the fluid as particles allows the direct application of forces on those particles. This is particularly useful for self-gravity, which is trivially represented as a force in

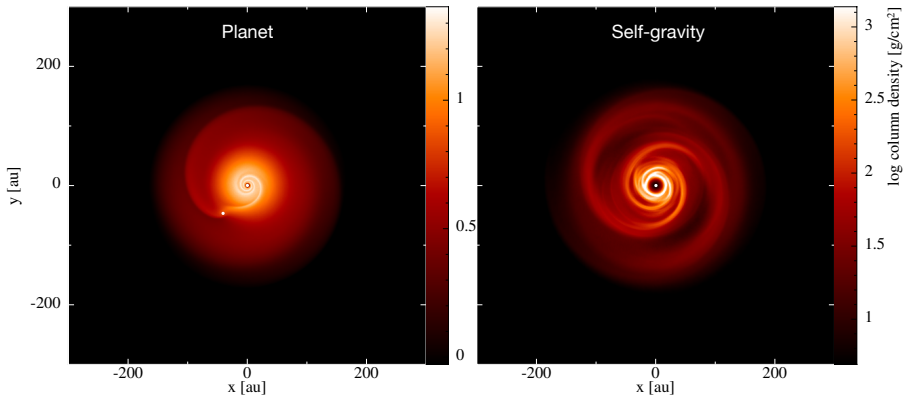

Figure 1.5: Density renderings of disks simulated with PHANTOM. Left: a disk with a planet. Right: a self-gravitating disk. Visualized with SPLASH (Price, 2007).

SPH. In grid-based codes, its inclusion requires solving the Poisson equation for the background gravitational potential, ($\nabla^2\Phi = 4\pi G\rho$), which is a time-consuming process. The ease of integrating self-gravity is a major reason for the focus on SPH. Additional benefits include the conservation of mass, energy, momentum, and angular momentum up to the integration technique used (PHANTOM uses leapfrog) (Price, 2010).

A particle in SPH moves with a local fluid velocity, $\mathbf{v_i}$, through a region with a density, ρ. The density for a given particle, ρ_i is easily computed using

$$\rho_i = \sum_j m_j W(|r_i - r_j|, h_i)$$

where W is the smoothing kernel, r_i is the particle's position, h_i is the particle's smoothing length, and m_j and r_j are the mass and position of the interacting particles, respectively. Typically, only those particles within some cutoff radius are considered. h_i is allowed to change with the density of the environment, $h_i = h_0(\frac{m_i}{\rho_i})^{1/3}$ using the assumption of equal-mass SPH particles. This gives the simulation variable resolution according to local conditions (Price et al., 2018).

Smoothing kernels are typically given in the form

$$W(h) \equiv \frac{\sigma}{h^\nu} f(q),$$

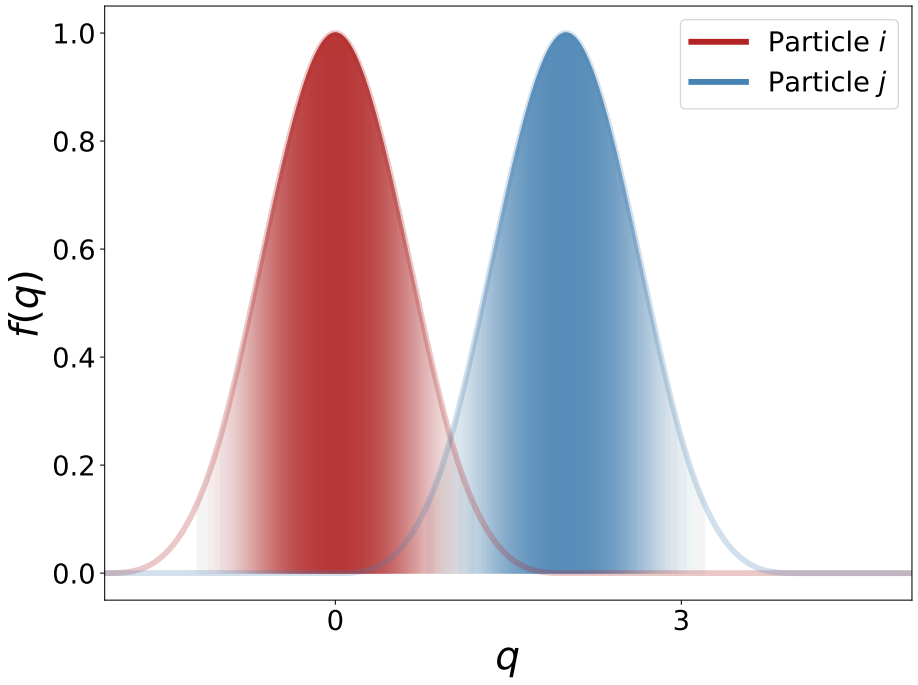

Figure 1.6: Schematic of interacting SPH particles according to Equation 1.8.1. The shading represents the kernel value, i.e., the extent of the interaction.

where σ is some normalization constant, ν is the number of spatial dimensions, and q is the distance between particles i and j in terms of smoothing length (i.e., $q \equiv |\mathbf{r_i} - \mathbf{r_j}|/h$) (Price, 2010). An tyical kernel is given in Equation 1.8.1

$$f(q) = \begin{cases} \frac{1}{4}(2-q)^3 - (1-q)^3, & 0 \leq q < 1 \\ \frac{1}{4}(2-q)^3, & 1 \leq q < 2 \\ 0, & q \geq 2 \end{cases}$$

(Monaghan, 1992). The interaction between particles is shown schematically in Figure 1.6.

The particle evolves according to the compressible hydrodynamics equations:

$$\frac{\mathrm{d}\mathbf{r_i}}{\mathrm{dt}} = \mathbf{v_i},$$

$$\frac{\mathrm{d}\rho_\mathrm{i}}{\mathrm{dt}} = -\rho(\nabla \cdot \mathbf{v_i}),$$

with

$$\frac{\mathrm{d}}{\mathrm{dt}} \equiv \frac{\partial}{\partial t} + \mathbf{v}_i \cdot \nabla.$$

These are the Lagrangian update of particle position, the continuity equation that enforces the conservation of mass, and the Lagrangian time derivative, respectfully (Price et al., 2018).

The particles with pressure, P, and specific internal energy, u, are then evolved with the compressible hydrodynamics equations:

$$\frac{\mathrm{d}\mathbf{v}}{\mathrm{dt}} = -\frac{\nabla P}{\rho} + \Pi_\mathrm{shock} + \mathbf{a}_\mathrm{ext}(\mathbf{r}, t) + \mathbf{a}_\mathrm{sink-gas}(\mathbf{r}, t) + \mathbf{a}_\mathrm{self-gravity}(\mathbf{r}, t),$$

$$\frac{\mathrm{d}u}{\mathrm{dt}} = -\frac{P}{\rho}(\nabla \cdot \mathbf{v}) + \Lambda_\mathrm{shock} - \frac{\Lambda_\mathrm{cool}}{\rho},$$

where the particle indices have been dropped (Price et al., 2018). The Π_shock and Λ_shock terms ensure the correct entropy change at shock fronts, and Λ_cool controls cooling. All $\mathbf{a}_{...}$ terms are accelerations due to external forces, the sink-gas interactions, and self-gravity, respectively, although they do not all have to be included. With the imposition of an equation of state, these equations are closed and simulating can proceed. Figure 1.5 shows example simulation results.

1.8.2 Synthetic observations

Synthetic observations are an essential tool for understanding protoplanetary disks as they give a direct comparison between simulations and observation. While SPH simulations give information on disk evolution, they are not capable of creating synthetic observations.

A common method for generating synthetic observations is radiative transfer. Since `MCFOST` is used in this work, it will be used as a descriptive example (Pinte et al., 2006, 2009). In `MCFOST`, monochromatic "photon packets" are propagated within the computed density and temperature structure. The photon packets interact with the disk, scattering according to some computed or stored optical properties, and eventually exiting the disk. Monte Carlo methods are used to simulate a photon packet's path throughout the disk. The final

photon packet distribution is used to generate synthetic observations (e.g., continuum maps, polarization maps, spectral energy distributions, and IFS cubes).

Continuum images and IFS cubes are the most important results in this project. Figure 1.3 shows a 1.25 mm continuum image overlaid with channel contours from a ^{13}CO (J 3 $\rightarrow$ 2) IFS cube, both made with `MCFOST`. Figure 1.4 shows single channels from synthetic IFS cubes made with `MCFOST`. Continuum images are useful for analyzing substructures, and IFS cubes are used for kinematic analysis.

Chapter 2

Gravitational Instability and Disk-to-Star Mass Ratio

*This work was originally published in J. P. Terry et al. MNRAS **510**, 1671–1679 (2022).*

Arguably the most fundamental parameter of a protoplanetary disk is its total mass, since it determines the total amount of mass available for planet formation. However, direct measurements of disk mass remain elusive. At the cool temperatures protoplanetary disks exist at, molecular hydrogen —H_2, which constitutes between ∼90-99% of the disk mass —cannot be observed. In order to circumvent this problem, disk masses are frequently estimated by converting continuum flux density at ∼mm wavelengths to a total dust mass. The total disk mass is then evaluated through the assumption of optically thin emission and a constant dust-to-gas ratio (Beckwith et al., 1990). While this ratio is canonically assumed to match that of the ISM (1:100), disk measurements show a variety of deviations. First, the gas disk, measured in ^{12}CO, typically extends somewhere between a factor of 2 and a factor of 4 out beyond the dust disk, as observed in the ∼mm continuum (Ansdell et al., 2016; Birnstiel & Andrews, 2014; Facchini et al., 2018; Panic, 2009; Pinte et al., 2016; Toci et al., 2021) due to inward radial drift of dust (Weidenschilling, 1977). Secondly, if significant grain growth has taken place, this would shift more emission to longer wavelengths, and would require observations at longer wavelengths to recover the dust mass (Ilee et al., 2020). It is therefore reasonable to assume that the dust:gas ratio is larger than the canonical ISM value in these regions.

An alternative method is to measure line emission from molecules such as CO and its isotopologues (Bergin & Williams, 2017; Miotello et al., 2014, 2016; Williams & Best, 2014) and convert to total gas mass (or surface density Miotello et al., 2018), through abundance ratios. Although generally thought to be more accurate than relying on dust-to-gas mass conversion, it is now understood that molecular abundances change in both space and time throughout a disk (Ilee et al., 2017; Quénard et al., 2018; K. Zhang et al., 2019), and so this method is likely subject to similar uncertainties driven by local changes to abundance ratios.

Recent near-infrared and sub-millimetre observations of protoplanetary disks (see, e.g., Andrews et al., 2018; Benisty et al., 2015; Benisty et al., 2021; J. Huang, Andrews, Dullemond, et al., 2018; J. Huang, Andrews, Pérez, et al., 2018; L. M. Pérez et al., 2016) have revealed prominent spiral structure in multiple disks, which may, in some cases, be due to gravitational instability (GI) (Cadman, Hall, et al., 2020; Chen et al., 2021; Dong et al., 2015; Hall et al., 2020; Hall et al., 2016; Meru et al., 2017; Veronesi et al., 2019). While it has been demonstrated that GI can be responsible for spiral morphology of some observed disks, it requires the disk-to-star mass ratio, q, be $\gtrsim 0.1$ for the spirals to be observable (Cossins et al., 2010; Dipierro et al., 2014; Dong et al., 2015; Hall et al., 2018, 2019; Kratter & Lodato, 2016). The presence of GI, by definition of its existence, therefore places constraints on the mass of a disk.

An interesting possibility with GI disks is fragmentation. Essentially, if a disk is sufficiently massive, and able to cool sufficiently quickly, then a region of that disk may fragment to form gravitationally bound objects (Gammie, 2001; W. K. M. Rice, Armitage, Bate, & Bonnell, 2003; W. K. M. Rice et al., 2005). Fragmentation has been proposed as a complementary planet formation pathway to the standard core accretion paradigm (Boss, 1997, 1998; Nayakshin, 2010), that could offer a plausible explanation for massive objects on wide orbits such as those around HR 8799 (Marois et al., 2008), the potential companion object in TW Hya (Nayakshin et al., 2020; Tsukagoshi et al., 2019), and the massive objects potentially forming in AB Aurigae (Cadman et al., 2021).

However, in general, if such objects regularly form they are likely to evade detection with instruments such as ALMA (Humphries et al., 2021). Analytical calculations (Rafikov, 2005), population synthesis models (D. H. Forgan, Hall, et al., 2018; D. Forgan & Rice, 2013) and hydrodynamical simulations (Hall et al., 2017) indicate that GI planet formation is most likely to result in objects $\geq 5\mathrm{M}_J$ at distances ≥ 50 au from their host star (K. Rice et al., 2015). Furthermore, it has recently been shown that less massive disks are more stable to GI (Haworth

et al., 2020), which may cause fragmentation to occur preferentially around the most massive objects (Cadman, Rice, Hall, et al., 2020; Ilee et al., 2018).

In either case, measuring the disk mass is crucial for determining the dynamical fate of the system - i.e., fragmentation (Gammie, 2001; W. K. M. Rice, Armitage, Bate, & Bonnell, 2003; W. K. M. Rice, Armitage, Bonnell, et al., 2003), quasi-steady GI (Lodato & Rice, 2004), episodic GI (Lodato & Rice, 2005) (see Kratter and Lodato, 2016 for a review of these topics) - and the total mass budget available for planet formation. Analysis of the disk rotation curve can give insight into this parameter. Recently, dynamical measurements of disk mass have been obtained through observing deviation from Keplerian behavior ($v_\phi = [\mathrm{GM_*/r}]^{1/2}$) in the rotation curve of Elias 2-27 (Veronesi et al., 2021).

Morphologically, it has been observed that there is an approximate relationship between the disk-to-star mass ratio, q, and the number of spiral arms, m, such that $m \sim 1/q$ (Cossins et al., 2009; Dong et al., 2015; Hall et al., 2019). However, simulated observations have shown that for an instrument such as ALMA, the correct number of spiral arms will not always be recovered from the observation (Dipierro et al., 2014; Hall et al., 2019), which depends not only on the resolution, but also on the arm to inter-arm contrast ratio (Hall et al., 2016). Therefore, this $m \sim 1/q$ method cannot be relied upon to infer total disk mass from observations.

Recently, Hall et al., 2020 presented a prediction for the kinematic signature of GI known as the "GI Wiggle". In individual line emission velocity channels, it is a distinctive "zig-zag" feature. When Keplerian rotation is subtracted from the intensity-weighted velocity, the wiggle appears as "interlocking fingers" (Hall et al., 2020). It is caused by sustained velocity perturbations throughout the disk, above and below the average azimuthal velocity inside and outside the disk spiral arms. Recent observations have determined that there is evidence of this feature in the Elias 2-27 system (Paneque-Carreño et al., 2021).

The aim in this work is to present, through numerical simulations, the relationship between the disk-to-star mass ratio, q, and the "strength" or "amplitude" of the GI Wiggle. Kinematic analysis is a promising avenue for extracting system properties, as deviations from Keplerian velocities within the disk can be linked to embedded objects (e.g. protoplanets) or physical processes (e.g. GI) that influence the system evolution (Bollati et al., 2021; Disk Dynamics Collaboration et al., 2020; Hall et al., 2020; Paneque-Carreño et al., 2021; Perez et al., 2015; Pinte et al., 2019, 2020; Teague et al., 2018; Wölfer et al., 2021). It has been shown by Hall et al., 2020 that, in contrast to local velocity fluctuations caused by the spiral wake of a protoplanet, (see, e.g., Pinte et al., 2018) GI-driven spirals result in global velocity perturbations ("GI wiggles") with a high degree of rota-

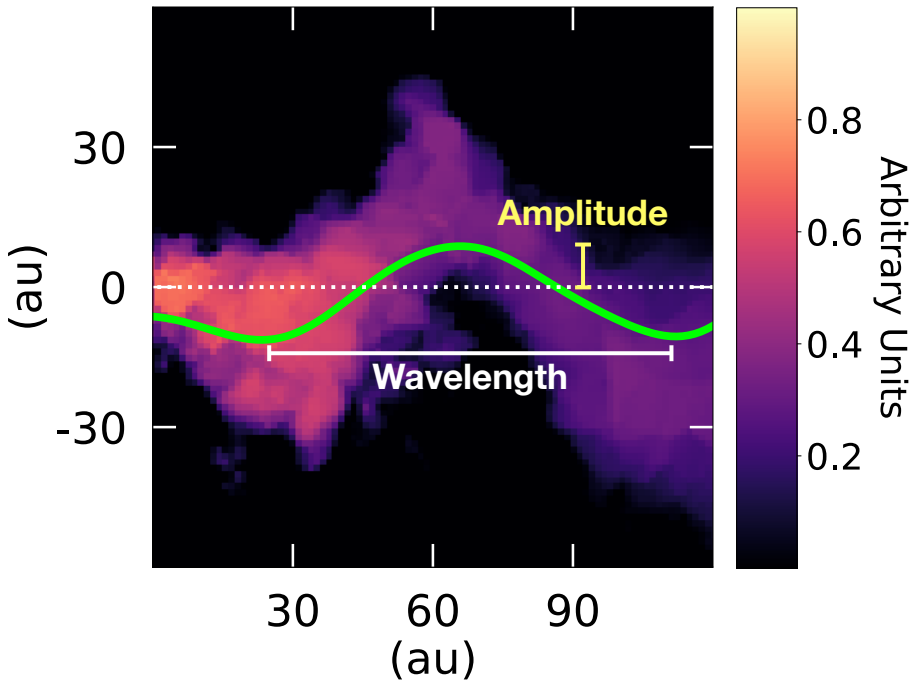

Figure 2.1: Illustration of the dominant amplitude (A_{wiggle}) and wavelength (λ_{wiggle}) of the GI wiggle overlaid on the synthetic line emission observation from which the signal was extracted. The systemic velocity channel is used ($v_{\text{obs}} = v_{\text{systemic}}$). For clarity, the displayed signal has been rotated to 90° relative to Figure 2.2 and now travels radially outwards from the star.

tional symmetry, and a more uniform quality than (for example), perturbations caused by a companion of larger planet mass (S. Pérez et al., 2018), or a vortex induced by a Rossby wave instability (P. Huang et al., 2018).

Another process, the vertical shear instability (VSI) (R. P. Nelson et al., 2013), may, like GI, be able to induce perturbations with a high degree of rotational symmetry (Barraza-Alfaro et al., 2021). However, these perturbations are an order of magnitude smaller than those induced by GI (∼0.06 km/s compared to ∼0.4 Barraza-Alfaro et al., 2021; Hall et al., 2020), so it should be possible to differentiate between them. Finally, a new and promising rotation curve technique demonstrated by Veronesi et al., 2021 would—when coupled with the presence of the GI-Wiggle —provide very strong evidence of GI over VSI.

I describe the wiggle in terms of the parameters wavelength and amplitude. Figure 2.1 shows these for a representative systemic velocity channel ($v_{\rm obs} = v_{\rm systemic}$). The perturbations corresponding to the amplitude are in the azimuthal direction, and the wavelength is measured radially from the central star. Figure 2.2 further illustrates this. I use these parameters to obtain the relationship between the wiggle and the disk-to-star mass ratio.

2.1 Methods

2.1.1 Hydrodynamical simulations

I performed simulations of three-dimensional self-gravitating protostellar disks composed of dust and gas using the `Phantom` Smoothed Particle Hydrodynamics (SPH) code (Price et al., 2018). Dust was modelled self-consistently with the gas, using the "one-fluid" method (Hutchison et al., 2018; Laibe & Price, 2014a, 2014b, 2014c), which in practice relies on the terminal velocity approximation (Youdin & Goodman, 2005). The back reaction of the dust onto the gas is included, and I use the flux-limited prescription of Ballabio et al., 2018. I allowed $\alpha_{\rm SPH}$ to vary between 0.1 and 1.0 and set $\beta_{\rm SPH} = 2.0$. The $\alpha_{\rm SPH}$ values correspond to Shakura-Sunyaev viscosities, $\alpha_{\rm SS}$, of $\alpha_{\rm SS} \approx 0.0025 - 0.025$. I assumed an initial dust-to-gas mass ratio of $\epsilon = 0.01$, and followed the spatial evolution of dust divided into 5 different size bins from 1 micron to 4 mm. The initial dust distribution was set to be perfectly mixed with gas. Even though I focus in this work on the synthetic CO emission, I include dust since it allows us to accurately determine the temperature structure of the disk in the radiative transfer calculation (see Section 2.1.2).

Seven simulations were performed in total, with disk-to-star mass ratios of $q = 0.15, 0.2, 0.25, 0.4, 0.5, 0.6$ and 0.75. As shown in Table 2.1, the number

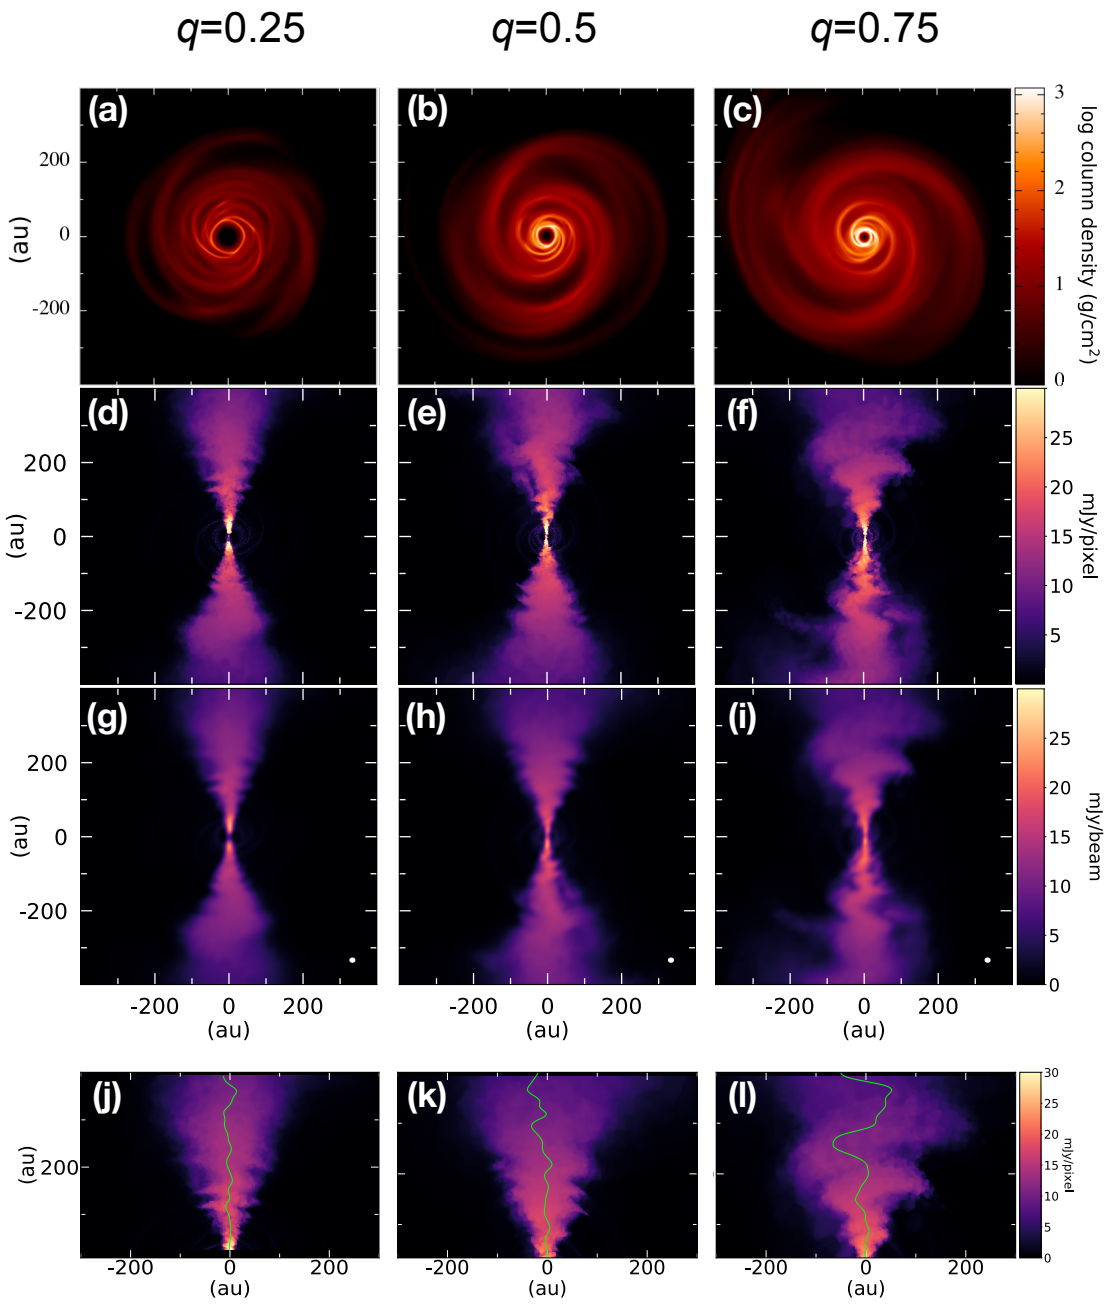

Figure 2.2: Hydrodynamical models of self-gravitating disks and associated GI wiggle. From left to right, q=0.25, q=0.5, q=0.75. Top Row: SPH log column density [g cm^{-2}]. Second Row: ^{13}CO line emission in mJy/pixel from systemic ($\Delta v \equiv v_{\rm obs} - v_{\rm systemic} = 0$ km/s) velocity channels from radiative transfer calculations on raw simulation results. Each pixel is 1.33 au by 1.33 au. Third Row: ^{13}CO line emission in mJy/beam from $\Delta v = 0$ km/s velocity channels after spatial and spectral convolution with beam size of 7.7 au by 5.6 au, indicated in the lower right corner. Bottom Row: Extracted signals overlaid on the $\Delta v = 0$ km/s velocity channel. Note that these are rotated relative to the signal in Figure 2.1.

Table 2.1: disk $R_{\rm outer}$, $R_{\rm res}$, and N (number of particles) for all q.

q	$R_{\rm outer}$ (au)	$R_{\rm res}$ (au)	N
0.15	153	32	500,000
0.2	174	31	500,000
0.25	186	28	750,000
0.4	213	24	1,000,000
0.5	216	23	1,250,000
0.6	234	21	1,250,000
0.75	237	13	1,500,000

of SPH particles was varied to maintain the relationship between scale height, H, and smoothing length, h, such that $h/H < 0.25$ for the majority of the disk, satisfying resolution requirements as outlined in Lodato and Clarke, 2011; A. F. Nelson, 2006. This has the additional effect of roughly maintaining the artificial viscosity as disk mass increases. See the Appendix for a more thorough discussion on the selection of the number of SPH particles for each q. The main simulation parameters are shown in Table 2.1.

The inner and outer radii of the disk were set to 10 and 300 au, respectively. The central star is represented as a sink particle (Bate et al., 1995) of mass $M_* = 0.6\,{\rm M}_\odot$ with an accretion radius of 1 au. The gas surface density profile is $\Sigma \propto R^{-1}$ and the sound speed profile is $c_{\rm s} \propto R^{-0.25}$, consistent with both observational results (J. Huang, Andrews, Pérez, et al., 2018; Paneque-Carreño et al., 2021; L. M. Pérez et al., 2016; Veronesi et al., 2021) and previous simulations of self-gravitating disks aimed at reproducing observations (D. H. Forgan, Ilee, & Meru, 2018; Hall et al., 2018; Meru et al., 2017; Tomida et al., 2017).

The disk was set such that it was initially stable to GI, with the Toomre parameter (Toomre, 1964)

$$Q = \frac{c_s \kappa}{\pi {\rm G} \Sigma}, \tag{2.1}$$

set to $Q \gtrsim 2$. Here, κ is the epicyclic frequency, G is the gravitational constant, c_s is the sound speed, and Σ is the surface density. In keplerian rotation, κ is simply Ω. In the regime where $Q \lesssim 1$, the disk becomes unstable. If cooling processes balance heating processes, GI will efficiently transport angular momentum throughout the disk on long timescales (Cossins et al., 2009; D. Forgan et al., 2010; Gammie, 2001). Alternatively, if the cooling rate is high, the disk may fragment to form gravitationally bound objects at the local Jeans mass (Gammie, 2001). For a recent review, see Kratter and Lodato, 2016. The disk

was allowed to cool through the simple β-cooling prescription (Gammie, 2001), where the cooling timescale is given by $t_c = \beta/\Omega$. In all simulations, $\beta = 15$, as in Hall et al., 2020. The resulting surface density is shown in the top row of Figure 2.2 for examples of the resulting column density.

2.1.2 Thermal disk structure and ^{13}CO channel map

The thermal disk structure and ^{13}CO $J = 3 \rightarrow 2$ channel map was computed using MCFOST, a Monte Carlo radiative transfer code (Pinte et al., 2006, 2009). ^{13}CO was chosen over ^{12}CO as it is less likely to be affected by foreground contamination. This is particularly important if the disk is still at least partially embedded, as is likely the case for young self-gravitating disks.

I assumed that the gas and dust temperatures were equal and used 10^8 photon packets to calculate dust temperature. As was done in the SPH simulations, the dust-to-gas ratio was 1:100. The density structure was obtained through direct Voronoi tesselation of the SPH particles, where each SPH particle corresponds to one MCFOST cell. Dust was assumed to be a mixture of silicates and carbon (Draine & Lee, 1984), and the optical properties were calculated using Mie theory (Andrews et al., 2009). The dust grain sizes vary between 0.03 μm and 4 mm, following a logarithmic distribution split into 100 bins. At any location in the model, the density of the dust of grain size a_i was obtained by interpolation from the SPH dust sizes. Any dust smaller than half the smallest SPH dust grain size (0.5 μm) was assumed to be perfectly coupled to the gas. The maximum size was set to be equal to the largest dust grain size in the SPH simulation (4 mm). Finally, the size distribution of the dust was normalised by integrating over all grain sizes assuming that the number density of dust grains, $n(a)$, was given by $\mathrm{d}n(a) \propto a^{-3.5}\,\mathrm{d}a$.

The ^{13}CO abundance relative to H_2 was set to 7×10^{-7} as in Hall et al., 2020. I set the stellar parameters to match those of the Elias 2-27 system as described in in Andrews et al., 2009: T_* = 3850 K, R_* = 2.3 $R_\odot$, and M_* = 0.6 $M_\odot$. The flux was synthetically simulated at a distance of 140 pc and an inclination of 30°. Two sets of channel maps were produced. The "raw" data, with no spectral or spatial convolution (examples in second row in Figure 2.2), and a set where the channels are spectrally and spatially convolved (examples in third row in Figure 2.2). The convolved channel maps were generated using Hanning-smoothing at a spectral resolution of 0.03 km/s. A turbulent velocity of 0.05 km/s was assumed. The maps were then spatially convolved with a Gaussian beam of size $0.11'' \times 0.07''$ and a position angle of $-38\,\mathrm{deg}$. This matches the expected spatial and spectral resolution necessary to kinematically detect a planet (Pinte et al., 2019), although it may be possible to detect such a signature

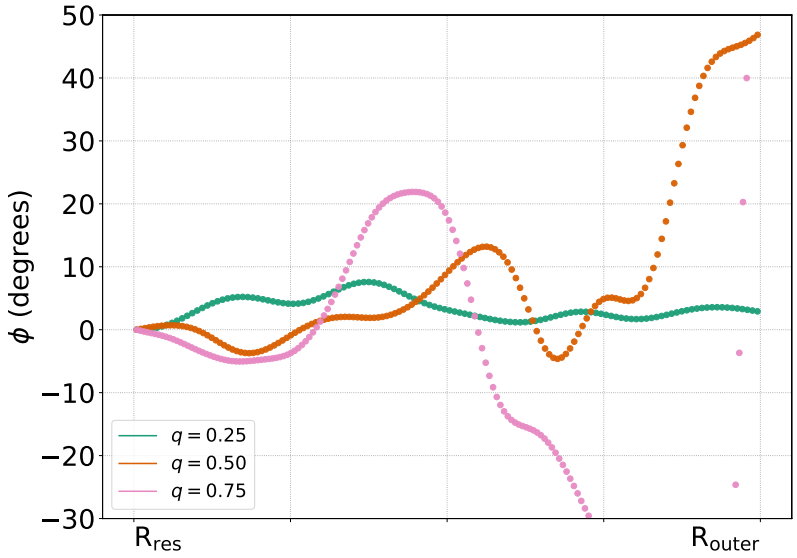

Figure 2.3: Smoothed GI wiggle signals from the $\Delta v = 0$ km/s channel in polar coordinates for selected mass ratios. Signals trace the spine of ^{13}CO emission and are extracted using Equation 2.3. All signals are between each disk's $R_{\rm res}$ and $R_{\rm outer}$ to ensure the disk is well-resolved at that distance.

at lower spatial and/or spectral resolution. The GI Wiggle is a somewhat larger feature than a planet-induced kink, and should therefore be readily detectable at this spatial and spectral resolution.

2.1.3 Velocity perturbations

An observed velocity field is be described by a decomposition into the azimuthal, radial, vertical, and systemic motion using

$$v_{\rm obs} = v_\phi {\rm sin}(i){\rm cos}(\phi) + v_{\rm r}{\rm sin}(i){\rm sin}(\phi) + v_{\rm z}{\rm cos}(i) + v_{\rm systemic}. \quad (2.2)$$

In the case of a purely Keplerian disk, $v_r = v_z = 0$, $v_\phi = \sqrt{GM_*/r}$, and the observed velocity field is the well known "butterfly pattern." When a self-gravitating spiral is present, the velocity field is perturbed. This perturbation results in a GI wiggle, which is a deviation from Keplerian rotation such that v_r and v_z become non-zero, and $v_\phi \neq [{\rm GM_*/r}]^{1/2}$.

When observed in line emission, this is a wave-like perturbation and therefore characterised by two properties: wavelength and amplitude, as illustrated in Figure 2.1. To extract wavelength and amplitude, I first define the signal, given by:

$$y_n = \frac{1}{I_n + \sum_{i=1}^{5}(I_i + I_{-i})}\left[I_n y + \sum_{i=1}^{5}\Big(I_i(y+i) + I_{-i}(y-i)\Big)\right], \tag{2.3}$$

where y_n is the signal (in au). For the n^{th} row in a given line emission velocity channel (e.g. a row in Figure 2.2 d, e or f), I locate the pixel (column in x_n), y, with the strongest emission, I_n. To minimise noise, I weight y_n by the intensities of the 5 pixels on either side. A given pixel that is i columns away from y in x_n has intensity I_i and is located at $y + i$. For examples of extracted signals, Figure 2.2 shows extracted signals overlaid on the line emission, and Figure 2.3 shows signals in polar coordinates. Note that the signals in all figures have been convolved with a Gaussian (i.e. Gaussian smoothing) for clarity.

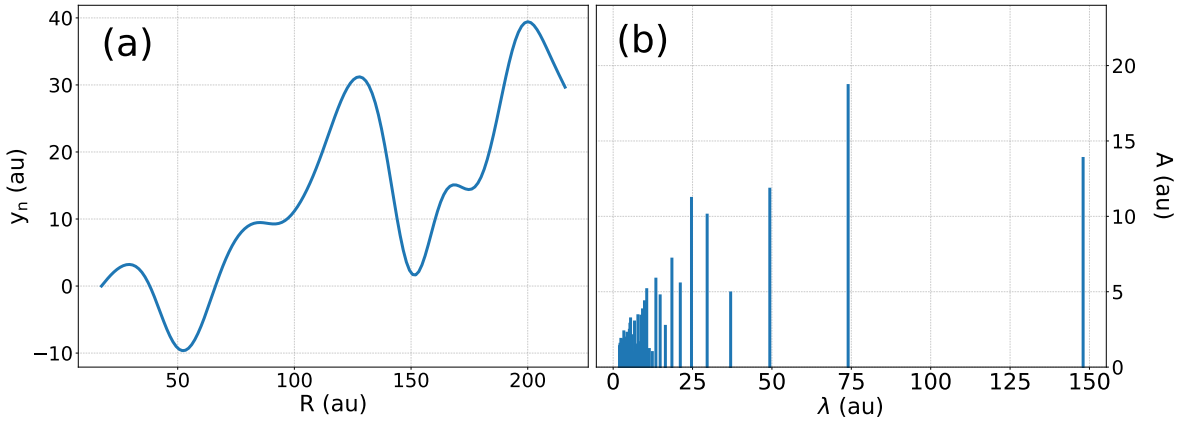

Figure 2.4: (a) Smoothed, intensity-weighted extracted signals from the $\Delta v = 0$ km/s channel for $q = 0.5$. (b) Amplitude vs wavelength after Fourier decomposition. For this channel, $\lambda = 74$ au is the dominant wavelength with an amplitude of A $= 18.7$ au.

Since the perturbation is approximately sinusoidal, it is well-described using a Fourier series. The extracted wiggle signal is decomposed using a Fourier transform into its power spectrum, $|A(\lambda)|^2$, for a given wavelength, λ. $|A(\lambda)|$ can then be interpreted as the amplitude of the component of the signal with that wavelength. Figure 2.4 shows this decomposition. I define the wave amplitude, A_{wiggle}, as max($|A(\lambda)|$) and λ_{wiggle} as the corresponding wavelength.

I only consider signals at radial distances, R, in the range $R_{\text{res}} \leq R \leq R_{\text{outer}}$, where R_{outer} is the outer radius containing 95% of the disk mass and R_{res} is the shortest radial distance from the central sink such that $h/H < 0.25$

for all $R_{\rm res} \leq R \leq R_{\rm outer}$. These values, along with the number of particles used for each simulation, can be found in Table 2.1. This method ensures that all extracted signals come from regions of the disk that are properly resolved. I apply this method over a range of velocity channels ($|\Delta v| = |v_{\rm obs} - v_{\rm systemic}| <$ 1 km/s) that covers the the majority of the disk. I do not consider higher velocity channels due to the difficult of extracting the signal. The average dominant wavelength and amplitude of the velocity perturbation is determined for a given q. Results are also normalised by $R_{\rm outer}$ to present a general result that can be extrapolated to other systems.

Determining $R_{\rm outer}$ consistently from real observations is, of course, challenging. One could use either the size of the continuum emission or that of line emission. The dust size measured from continuum emission is typically a factor 2-3 smaller than the gas size measured from ^{12}CO emission lines (Ansdell et al., 2018; Sanchis et al., 2021; Tazzari et al., 2021), due to a combination of radial drift and possibly substructure formation (Rosotti et al., 2019; Toci et al., 2021). CO-based measurements are generally more reliable than dust size measurements, but one should be careful that, due to optical thickness effects, the radius enclosing 95% of the ^{12}CO flux contains much more than 95% of the mass. Using the 68% CO flux radius may trace the bulk of the gas mass more accurately (Toci et al., 2021; Trapman et al., 2020).

As described in Section 2.1.2, two sets of channel maps are produced. The first is the "raw" data - i.e. the exact GI wiggle signal (second row in Figure 2.2). I first perform the analysis on the raw data (obtaining wavelength and amplitude), then repeat the analysis on the spatially and spectrally convolved channel maps (third row in Figure 2.2), to ensure that it would be possible to extract these properties from an observed system with current spatial and spectral resolution.

2.2 Results

The results show that there is a strong positive correlation between q and GI wiggle amplitude, shown in Figure 2.5. Plot (a) shows the values in au, and plot (b) shows the values normalised by $R_{\rm outer}$ so that they can be generalised to disks of any size. Blue lines with squares are wiggle wavelengths, and orange lines with circles are wiggle amplitudes. Solid lines are raw data, while dotted lines are spatially and spectrally convolved data. I quantify the results by performing an R^2 linear regression analyses, shown in Table 2.2 and Table 2.3.

When I normalize the results with respect to $R_{\rm outer}$, shown in Figure 2.5(b), the positive correlation between amplitude and q remains, although the R^2 values in Tables 2.2 and 2.3 show that it is less robust. However, the wavelength

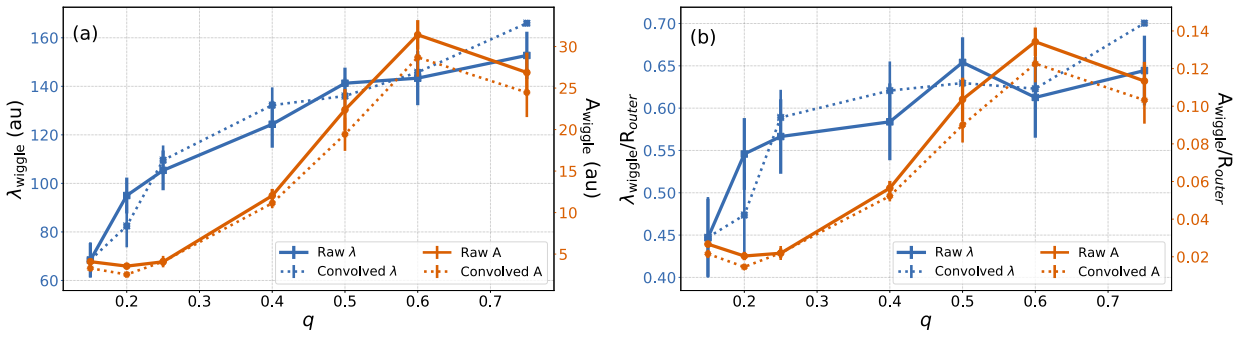

Figure 2.5: Amplitude (orange circles) and wavelength (blue squares) from the raw simulation data (solid line) and results that have been spatially and spectrally convolved (dotted line) plotted with standard errors, which are mostly smaller than the points. (a) Amplitude and wavelength in terms of au; (b) Amplitude and wavelength normalized by R_{outer}. Standard errors, $\sigma/\sqrt{n}$, for a given distribution are plotted.

Table 2.2: Linear fit results from A_{wiggle} and λ_{wiggle} as functions of q.

Value	Fit Slope: m (au)	Fit Intercept: b (au)	R^2
Raw Amplitude	50 ± 8	-5.4 ± 3.8	0.88
Raw Wavelength	129 ± 20	66.0 ± 9.3	0.89
Convolved Amplitude	48 ± 9	-6.6 ± 4.5	0.87
Convolved Wavelength	133 ± 20	68.7 ± 9.7	0.92

Table 2.3: Linear fit results from the R_{outer}-normalized A_{wiggle} and λ_{wiggle} as functions of q.

Value	Fit Slope, m	Fit Intercept, b	R^2
Raw Amplitude	0.200 ± 0.037	-0.013 ± 0.017	0.85
Raw Wavelength	0.27 ± 0.078	0.471 ± 0.035	0.70
Convolved Amplitude	0.196 ± 0.040	-0.021 ± 0.020	0.85
Convolved Wavelength	0.308 ± 0.088	0.467 ± 0.043	0.75

as a fraction of $R_{\rm outer}$ substantially weakens the correlation with q. While the relevant R^2 values in Table 2.3 are still large enough for a linear approximation to be acceptable, they are significantly lower than the equivalent values in Table 2.2.

2.3 Discussion

The positive correlation between q and $A_{\rm wiggle}$ is unambiguous and expected due to the physical origin of the wiggles. There is a correlation between the wiggles, in velocity space, and the spiral density waves, in physical space (Longarini et al., 2021). Given this, larger spirals in space will to some degree correspond with larger velocity perturbations. For example, more massive disks, with lower m-modes and therefore fewer spirals that dominate the morphology, are expected to have a larger wiggle amplitude. However, it is important to note that the amplitude of the wiggle is determined by how large a velocity shift the material in that spatial location has received, rather than the spatial scale of the spiral causing the wiggle.

The positive correlation between q and $\lambda_{\rm wiggle}$ is easy to understand. Cossins et al., 2009 showed that $kH \simeq 1$, where k is the radial wavenumber and H the height of the disk. The wavelength of the perturbation, i.e. the wavelength of the wiggle, is thus proportional to H that, in a self-gravitating disk, scales as the disk to star mass ratio q (Kratter & Lodato, 2016).

Figure 2.6 shows moment-1 maps with Keplerian rotation subtracted (with each pixel scaled by its distance from the central sink for clarity) for mass ratios of 0.25 and 0.75. These moment-1 maps shows the intensity-weighted velocity, which is calculated by Equation 2.4

$$\langle v \rangle = \frac{\int_{-\infty}^{\infty} v \mathrm{I}\left(v\right) \mathrm{d}v}{\int_{-\infty}^{\infty} \mathrm{I}\left(v\right) \mathrm{d}v}. \tag{2.4}$$

Both disks in Figure 2.6 exhibit the "interlocking finger" pattern predicted by Hall et al., 2020 as a characteristic imprint of GI. However, the more massive disk, $q = 0.75$, clearly has more pronounced perturbations, strongly suggesting that, when compared to the less massive disk, GI has a larger influence on its dynamics.

2.3.1 Limitations and considerations

This study comes with a number of limitations. I vary only mass ratio and number of SPH particles between runs, exploring only the GI wiggle's dependence on disk-to-star mass ratio, and not on other parameters. This is likely a

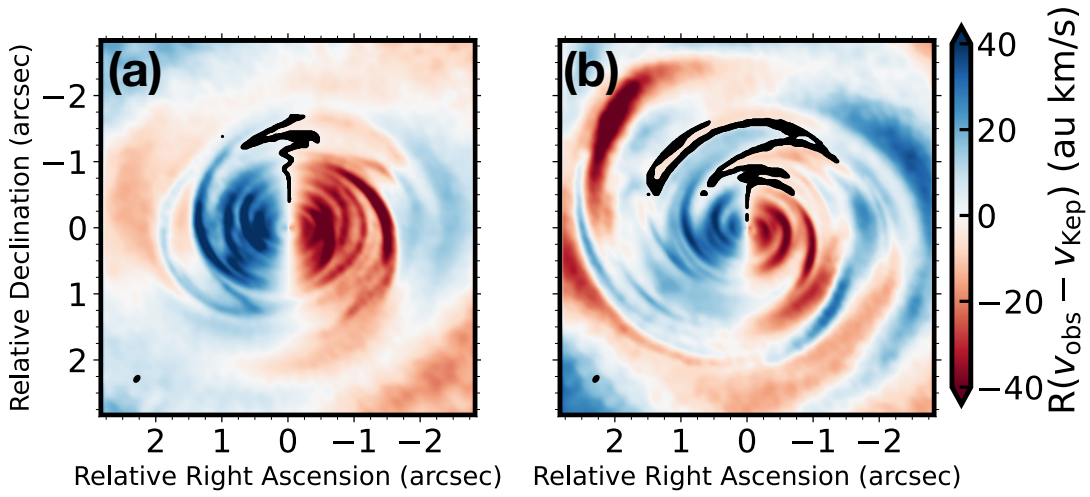

Figure 2.6: Moment-1 maps with Keplerian rotation subtracted and multiplied by the radius from the central sink, R (in au), for clarity. Using disk-to-star mass ratios of (a) $q = 0.25$, (b) $q = 0.75$. The "interlocking fingers" that are predicted signatures of GI are present although clearly stronger for the more massive disk. The black line indicates where $v_{\rm obs} = v_{\rm Kep}$.

complicated issue, since, for example, disks are more stable around low-mass stars (Cadman, Rice, Hall, et al., 2020; Haworth et al., 2020). Another issue is the cooling parameter, β, which I do not vary. Here, I have shown that for a constant β, $A_{\rm wiggle}$ depends on q. This work, therefore, is essentially assuming that most disks undergoing GI will have similar radiative properties. Since $\delta\Sigma \propto \beta^{-1/2}$, I expect that smaller β gives stronger perturbations, thus wiggles with bigger amplitude. Longarini et al., 2021 provide an analytical model for the GI-wiggle, confirming this trend. Even had I varied β, this would not accurately reflect reality, since heating and cooling are complicated processes best considered by a full radiative transfer treatment. However, assuming constant β allowed us to remove one variable from the analysis, which would not have been possible with full radiative transfer.

I do not include the effect of viscosity on the damping of the GI wiggle in this analysis, but I can provide an order of magnitude estimate of its effect. In a non-viscous, non-irradiated disk, the amplitude of the density perturbation due to gravitational instabilities is such that the effecting heating provided by the instability balances radiative cooling (Kratter & Lodato, 2016). The effective heating associated to GI can be parameterized through an effective $\alpha_{\rm GI}$,

proportional to the density perturbation:

$$\alpha_{\rm GI} = \frac{4}{9\gamma(\gamma-1)}\frac{1}{\beta} \approx \frac{4}{9\gamma(\gamma-1)}\left(\frac{\Delta\Sigma}{\Sigma}\right)^2, \tag{2.5}$$

where γ is the ratio of specific heats and in the last step I have used the fact that for a non-viscous self-regulated disk $(\Delta\Sigma/\Sigma)^2 \approx 1/\beta$ (Cossins et al., 2009). If an additional viscosity is present, parameterized by a Shakura-Sunyaev coefficient $\alpha_{\rm SS}$, thermal balance is then rewritten as:

$$\alpha_{\rm GI} + \alpha_{\rm SS} = \frac{4}{9\gamma(\gamma-1)}\left(\frac{\Delta\Sigma}{\Sigma}\right)^2 + \alpha_{\rm SS} = \frac{4}{9\gamma(\gamma-1)}\frac{1}{\beta}, \tag{2.6}$$

from which I obtain

$$\left(\frac{\Delta\Sigma}{\Sigma}\right)^2 = \frac{1}{\beta}\left(1 - \frac{9\gamma(\gamma-1)\alpha_{\rm SS}\beta}{4}\right). \tag{2.7}$$

Since the velocity perturbations are proportional to the density perturbations (Longarini et al., 2021), I expect a similar reduction also in the amplitude of the wiggle, when $\alpha_{\rm ss}$ is non-negligible. Note that Equation (2.7) is analogous to equation (15) in W. K. M. Rice et al., 2011, derived for irradiated disks, and in fact a similar reduction in the wiggle amplitude is also expected when one adds also irradiation as a source of disk heating.

In this work, I intended to determine if it was possible to characterise the GI wiggle from current ALMA observations. I therefore included dust in the simulations since the one-fluid method of Laibe and Price, 2014a, 2014b, 2014c only results in a very small slow-down ($\sim$ 5%) of the code, so it is possible to perform a much more realistic, observationally motivated calculation at very little extra computational cost. GI disks trap dust in their spiral arms (W. K. M. Rice et al., 2004), resulting in a varied radial dust distribution which affects the thermal disk structure in the radiative transfer calculation. This in turn can affect the intensity of molecular line emission, which is why I opted to use dust in the calculations. However, it would be interesting in future to vary both dust-to-gas ratio and dust grain distribution to explore what effect, if any, it has on the GI wiggle.

An important point to note is that the resolution used in this analysis represents some of the highest spectral and spatial resolutions found in observations to date. The convolved results had a spectral resolution of 0.03 km/s after Hanning-smoothing and were spatially convolved with a Gaussian beam of size $0.11'' \times 0.07''$. While this resolution has been used on planet-containing

disks (Pinte et al., 2019), these objects are far below the mass threshold required for GI to be active. The current spatial and spectral resolution of observations of the best known GI candidate, Elias 2-27 (Paneque-Carreño et al., 2021; L. M. Pérez et al., 2016; Veronesi et al., 2021), is roughly a factor 3 below the resolution I use here. Further ALMA observations are therefore required to apply the method I describe here.

2.4 Conclusions

I used numerical simulations to determine a positive linear relationship between the amplitude of the GI Wiggle and disk-to-star mass ratio q, for a constant cooling parameter, β. The best fit relationship using $\beta = 15$ is $A_{\rm wiggle} = 50q - 5.4$. The R^2 value from this fit is 0.88, suggesting that a linear fit is appropriate. A similar linear regression on the GI wavelength gave $\lambda_{\rm wiggle} = 129q + 66$ with an R^2 value of 0.89. This also indicates that approximating the GI wavelength as a linear function of the mass ratio is also valid. I present a heuristic argument based both on previous findings and physical reasoning to support these numerical results.

The results hold for convolution to a spectral resolution of 0.033 km/s, and spatial convolution using a Gaussian beam of size $0.11'' \times 0.07''$. This indicates that determination of wiggle wavelength and amplitude from ALMA observations with maximum resolution is immediately possible. I therefore suggest that these results can be used to constrain disk mass in systems that contain the GI Wiggle.

Chapter 3

Gravitational Instability in the Presence of Stars

Sufficiently massive and cold protoplanetary accretion disks may undergo a period where the disk's self-gravity is important (Kratter & Lodato, 2016). Under these conditions, typically early in the disk's lifetime before significant mass has been cleared, gravitational instability (GI) may be triggered, qualitatively changing angular momentum transport, facilitating accretion, and forming spiral arms (Kratter & Lodato, 2016). GI has been proposed to as a possible mechanism for stellar binary companions, brown dwarfs, and massive gas giants (Adams et al., 1989; Boss, 1997). The potential importance of GI in the resulting stellar system therefore demands a thorough study of self-gravitating protoplanetary disks.

Previous works have found that GI causes global velocity perturbations from Keplerian motion ($\Omega_{Kep} = \sqrt{GM_\star/r^3}$) (Hall et al., 2020). The "GI wiggle" seen in line emission profiles is an important signature of GI and correlates with the both cooling rate and disk-to-star mass ratio (Longarini et al., 2021; Terry, Hall, Longarini, et al., 2022). A lower cooling rates, parameterized by $\beta_{cool} = t_{cool}\Omega_{Kep}$ (Gammie, 1996), or a high disk-to-star mass ratios make the GI wiggle more pronounced (Longarini et al., 2021; Terry, Hall, Longarini, et al., 2022).

These efforts demonstrated that GI gives a clear signature that contains essential information about the disk's properties. However, the works assumed both a constant stellar mass and β_{cool}. β_{cool} is more realistically calculated according to local disk conditions rather than taken as a constant across the entire disk. Disk conditions explicitly depend on stellar mass through Keplerian fre-

quency terms (similarly for β_{cool}) and implicitly enters if stellar irradiation is included. I therefore calculate $\beta_{cool} = \beta_{cool}(r, M_\star)$ to give a more realistic treatment of the stability of irradiated disks across radii and stellar masses.

The chapter is arranged as follows: Section 3.1 gives an overview of the methods used to self-consistently calculate the cooling and quantify the GI wiggle. Section 3.2 presents the results. Section 3.3 reviews the results, and Section 3.4 offers the conclusions.

3.1 Methods

3.1.1 Disk models

One-dimensional disk profiles are calculated self-consistently under the assumption of self-regulating viscous equilibrium(Clarke, 2009). Q is the Toomre parameter given in Equation 3.1 (Toomre, 1964). Disks are gravitationally unstable when $Q \lesssim 1$. For a given accretion rate, $\dot{M}$, the profiles are calculated according to Equation 3.2 using $Q = 1$.

$$Q = \frac{c_s \kappa}{\pi G \Sigma}, \tag{3.1}$$

where $\kappa = \Omega$ in a Keplerian disk.

$$\dot{M} = 3\pi\nu\Sigma. \tag{3.2}$$

In the $\alpha-$prescription (Shakura & Sunyaev, 1973),

$$\nu = \alpha c_s H, \tag{3.3}$$

where H is the pressure scale height,

$$H = c_s/\Omega. \tag{3.4}$$

In the case of $Q \approx 1$, α can be directly expressed (Gammie, 2001)

$$\alpha = \frac{4}{9\gamma(\gamma-1)t_{cool}\Omega}, \tag{3.5}$$

where γ is the specific heat ratio ($\gamma = 7/5$ in this work).

The cooling time, t_{cool}, depends on an effective temperature (Menou & Goodman, 2004)

$$T_e^4 = \frac{4}{3}\frac{T_c^4 - T_{irr}^4}{\left(\tau/2 + 1/\sqrt{3} + 1/(3\tau)\right)} \tag{3.6}$$

The disk's midplane temperature, T_c, is given by

$$T_c = \frac{\mu m_p c_s^2}{\gamma k_B}, \tag{3.7}$$

where $\mu = 2.353$ and m_p is the proton mass. A temperature floor of 2.726 K is enforced.

The irradiation temperature, T_{irr}, depends on the stellar luminosity and local conditions (Menou & Goodman, 2004)

$$T_{irr}^4 = \frac{L_\star(1-\epsilon)}{\sigma 4\pi r^2}\frac{H}{r}\left(\frac{\mathrm{d}\log H}{\mathrm{d}\log r} - 1\right), \tag{3.8}$$

where $\epsilon = 1/2$ and $\frac{\mathrm{d}\log H}{\mathrm{d}\log r} = 9/7$, implying no self-shadowing. I use Siess et al., 2000 isochrones for a 1 Myr star.

The optical depth, τ, is given in terms of the Rosseland mean opacity, κ, (Lin & Papaloizou, 1985)

$$\tau \approx H\kappa\rho_{mid} \tag{3.9}$$

with the midplane density (W. K. M. Rice & Armitage, 2009),

$$\rho_{mid} = \frac{\Sigma}{2H}. \tag{3.10}$$

. The full equations for κ are given in the Appendix A.

The disk's internal energy per unit area, U, is (W. K. M. Rice & Armitage, 2009)

$$U = \frac{c_s^2\Sigma}{\gamma(\gamma-1)}. \tag{3.11}$$

The cooling rate is (W. K. M. Rice & Armitage, 2009)

$$\Lambda = 2\sigma T_e^4. \tag{3.12}$$

Finally, the cooling time, t_{cool}, is (W. K. M. Rice & Armitage, 2009)

$$t_{cool} = U/\Lambda = \beta_{cool}/\Omega_{kep}. \tag{3.13}$$

To account for expected differences between stellar masses, the minimum and maximum radial extent of the disk is determined by scaling given by Burn

et al., 2021 and Andrews et al., 2010, respectively

$$R_{in} = 0.1\mathrm{au}\left(\frac{M_\star}{M_\odot}\right)^{1/3}, \tag{3.14}$$

$$R_{out} = 2R_{crit} = 2\left(\frac{M_{disk}}{5\times10^{-4}M_\odot}\right)^{1/1.6}. \tag{3.15}$$

This directly encodes both the stellar mass and disk-to-star mass ratio into the used profile.

3.1.2 Velocity structure

The calculated cooling rate profiles are used to determine $\beta_{cool}(r)$. Following Longarini et al., 2021, the β_{cool} are used to calculate the velocity structure of the disk. Longarini et al., 2021 showed that the radial and azimuthal velocity perturbations due to GI (δu_r and δu_ϕ, respectively) are both proportional to the inverse square root of β_{cool}. That is,

$$\delta u_{r,\phi} \propto \beta_{cool}^{-1/2}. \tag{3.16}$$

I quantify the effect of GI on the disk through the GI wiggle (Hall et al., 2020; Longarini et al., 2021; Terry, Hall, Longarini, et al., 2022). The wiggle, $w(r)$, is extracted using the intensity-weighted average velocity, i.e., the moment-1 map

$$M_1 = \frac{\int_{-\infty}^{\infty} v\mathrm{I}\left(v\right)\mathrm{d}v}{\int_{-\infty}^{\infty}\mathrm{I}\left(v\right)\mathrm{d}v}. \tag{3.17}$$

The wiggle at a given radius, r, is the angle, ϕ, where the moment-1 map is zero. That is,

$$w(r) = \phi \quad \text{s.t.} \quad M_1(\phi, r) = 0. \tag{3.18}$$

It is essential to use the systemic channel (i.e, $M_1 = v_{obs} = 0$ km/s) because its azimuthal location is independent of stellar mass. No other velocity channels are azimuthally coincident with one another across stellar mass due to the explicit mass term in Ω_{Kep}. This would prohibit any direct comparisons between wiggle properties.

I compare the results with those obtained from a constant β_{cool} treatment, as per Longarini et al., 2021.

3.2 Results

3.2.1 Disk profiles

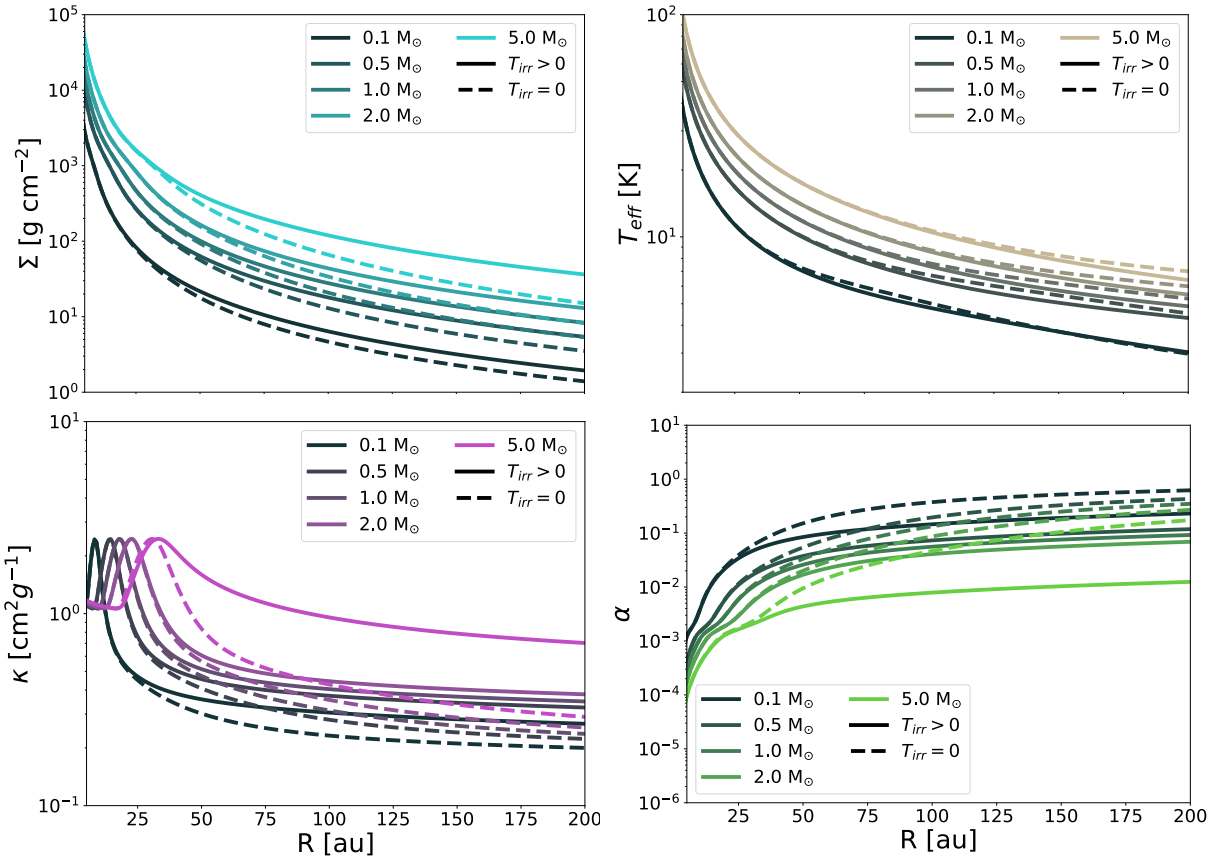

Figure 3.1: Self-consistent profiles calculated according to Section 3.1.1 with an accretion rate of $\dot{M} = 10^{-6}$M$_\odot$/yr. Upper left: surface density. Upper right: effective temperature (Equation 3.6) Lower left: opacity (Appendix A) Lower right: α (Equation 3.5) Solid lines: $T_{irr} = 0$. Dashed lines: $T_{irr} > 0$.

Figure 3.1 shows selected values from the self-consistent profiles calculated using the methods given in Section 3.1.1 using an accretion rate of $\dot{M} = 10^{-6}$M$_\odot$/yr. Figure 3.2 shows the irradiation temperature (Equation 3.8) over mass ratios and stellar masses. Figure 3.3 shows the resulting cooling times and β_{cool}. Figure 3.4 shows the value of β_{cool} evaluated at the critical radius ($R_{crit} = R_{out}/2$) with irradiation turned off (left) and on (right).

3.2.2 GI wiggles

Figure 3.5 shows calculated moment-1 maps using the self-consistent profiles. Figure 3.6 and Figure 3.7 shows extracted wiggles across methods and stellar

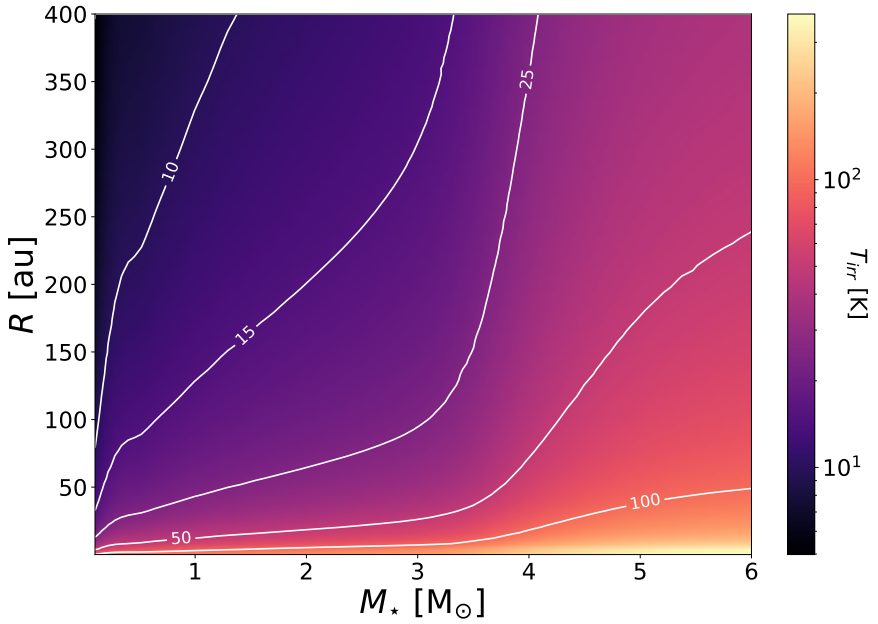

Figure 3.2: Irradiation temperature over disk-to-star mass ratio and stellar mass (Equation 3.8).

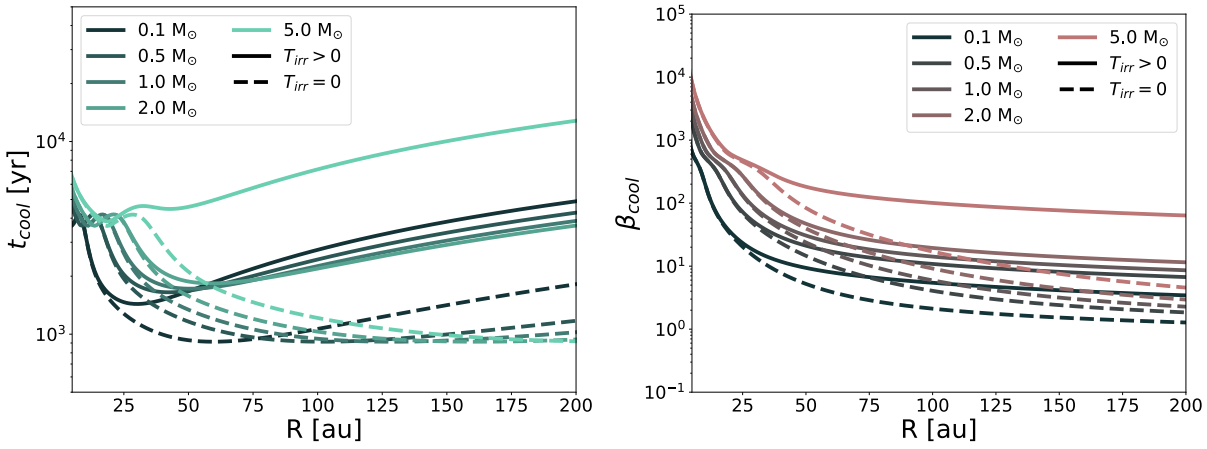

Figure 3.3: Left: Cooling time (Equation 3.13). Right: β_{cool}. $\dot{M} = 10^{-6} M_\odot$/yr.

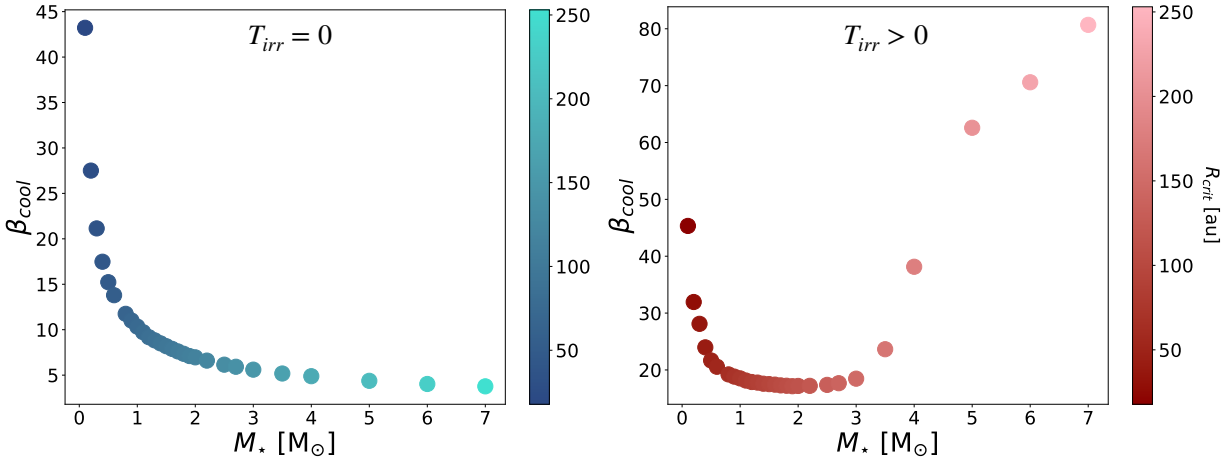

Figure 3.4: β_{cool} at $R_{crit} = R_{out}/2$ for each stellar mass assuming $M_{disk}/M_\star = 0.5$. Left: $T_{=irr} = 0$. Right: $T_{irr} > 0$.

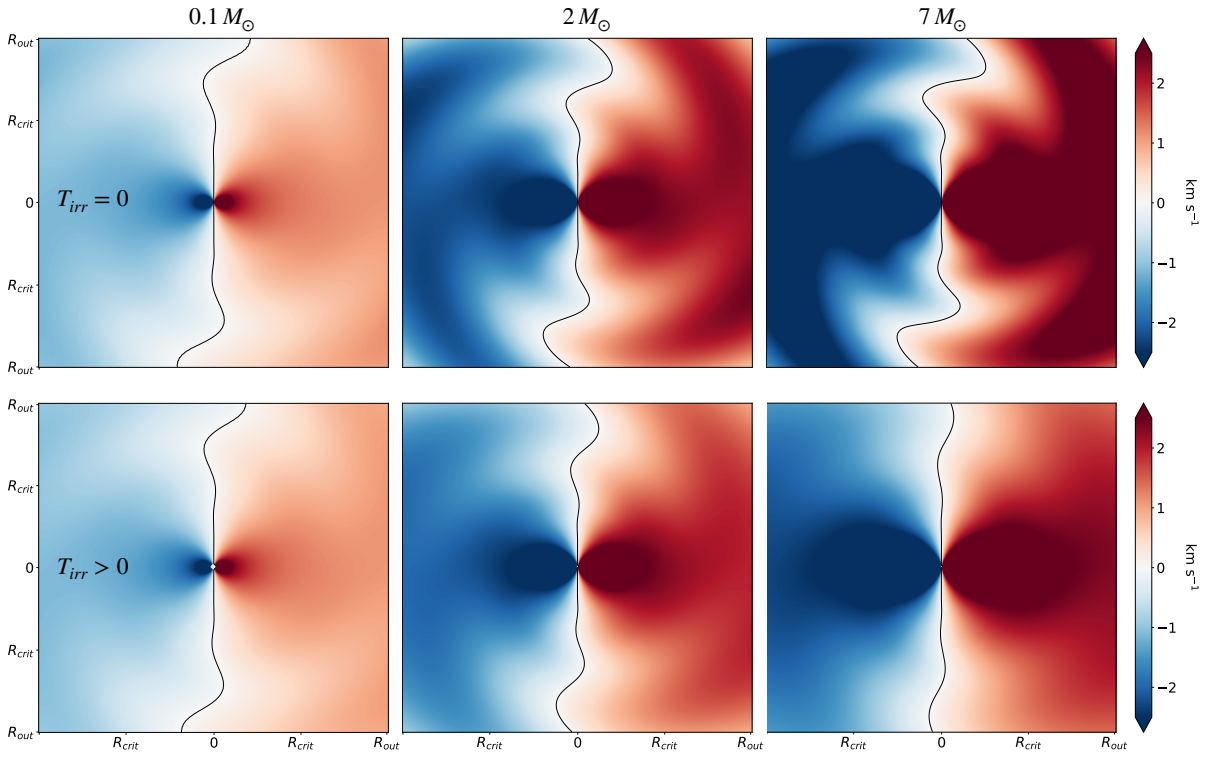

Figure 3.5: Moment-1 maps with $M_{disk}/M_\star = 0.5$. Top: $T_{irr} = 0$ K. Bottom: $T_{irr} > 0$ K. Left: $M_\star = 0.1\,M_\odot$. Center: $M_\star = 2\,M_\odot$. Right: $M_\star = 7\,M_\odot$. The black line shows the GI wiggle for each disk.

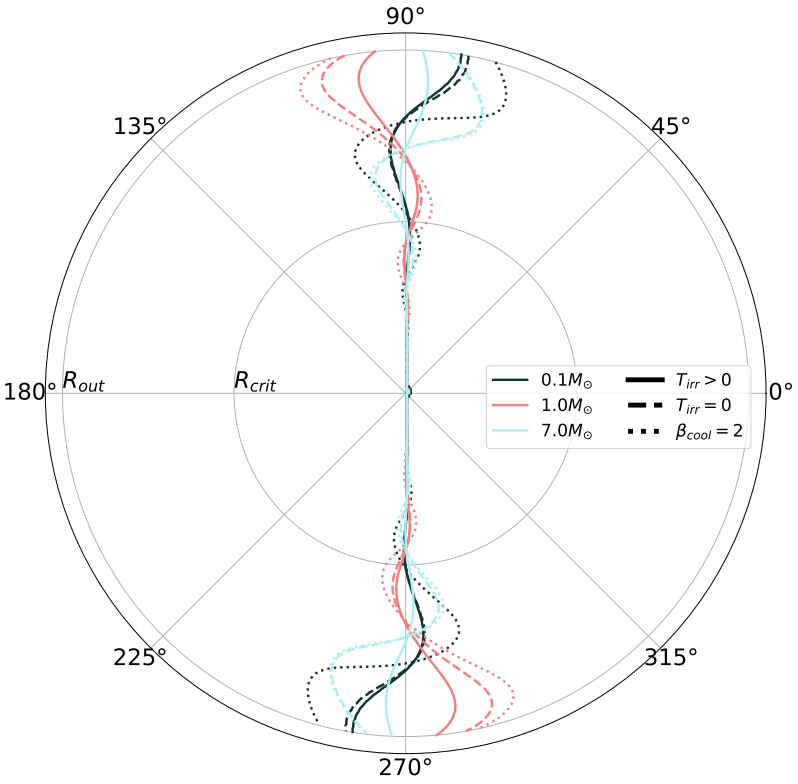

Figure 3.6: Extracted GI wiggles over masses and methods. Black line: $M_\star - 0.1$ $M_\odot$ Pink line: $M_\star - 1$ $M_\odot$. Blue line: $M_\star - 7$ $M_\odot$. Solid lines: $T_{irr} > 0$ K. Dashed lines: $T_{irr} = 0$ K. Dotted lines: Longarini et al., 2021 method with constant $\beta_{cool} = 2$.

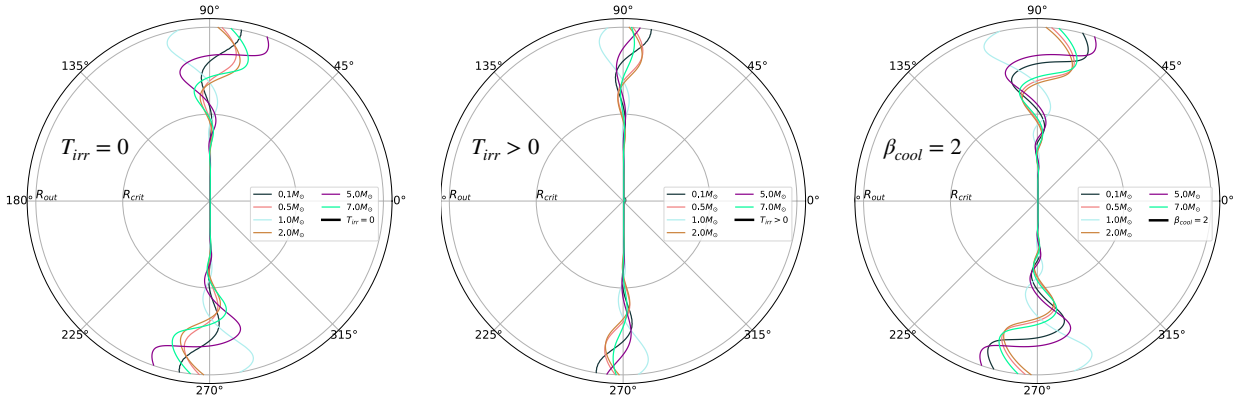

Figure 3.7: Extracted GI wiggles over stellar masses with $M_{disk}/M_\star = 0.5$. Left: $T_{irr} = 0$ K. Center: $T_{irr} > 0$ K. Right: Longarini et al., 2021 method with constant $\beta_{cool} = 2$.

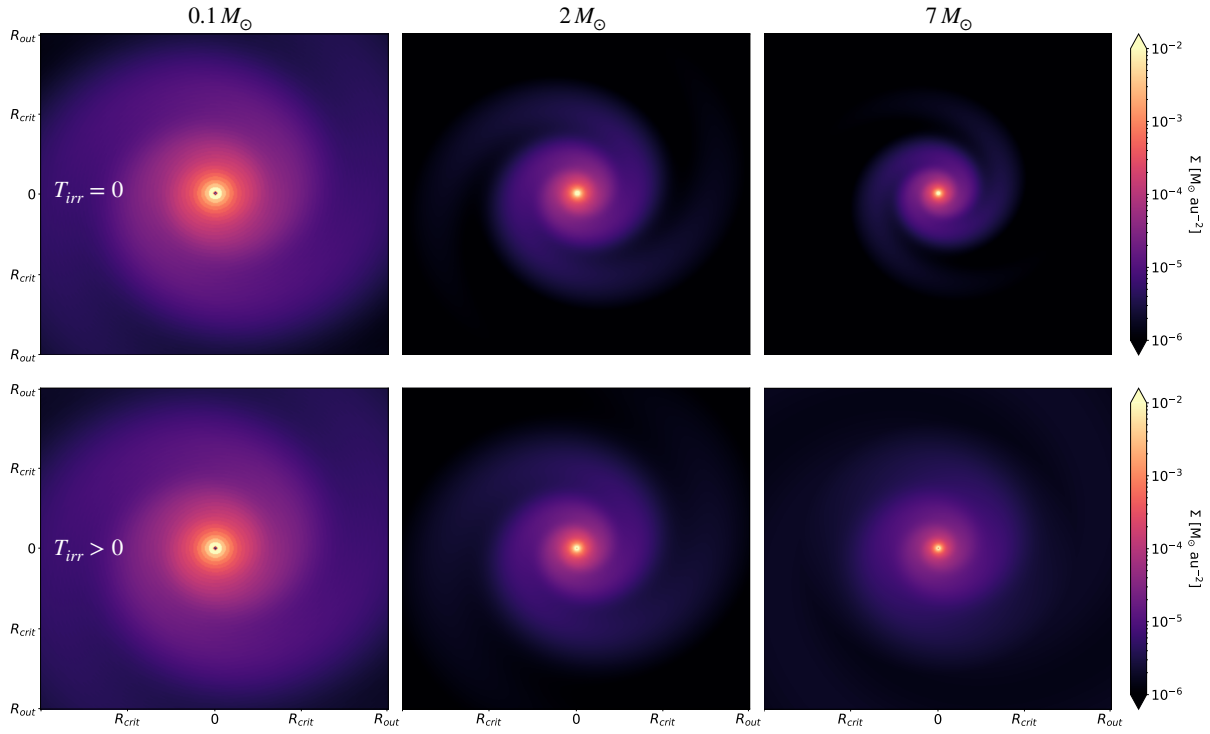

Figure 3.8: Calculated surface density with $M_{disk}/M_\star = 0.5$. Top: $T_{irr} = 0$ K. Bottom: $T_{irr} > 0$ K. Left: $M_\star = 0.1\,M_\odot$. Center: $M_\star = 2\,M_\odot$. Right: $M_\star = 7\,M_\odot$.

masses. Figure 3.8 shows the calculated surface density profile for $M_{disk}/M_\star = 0.5$.

3.3 Discussion

3.3.1 Calculated disk properties

The results in Section 3.2.1 give a clear indication that irradiation can have a significant impact on disk structure and stability. Figure 3.1 shows that irradiated disks (dashed lines) have a higher surface density, a lower effective temperature, a higher opacity, and less efficient angular momentum transport than non-irradiated disks (solid lines). These differences are particularly visible for massive stars and large radii.

The cumulative effect of these differences is increased stability around irradiated disks. Figure 3.3 shows that, for all masses, β_{cool} is larger when irradiation is included. The difference increases with radius and mass. This is in line with previous work. D. Forgan and Rice, 2013 found that irradiation can suppress fragmentation, which typically occurs at lower values of β_{cool}.

Figure 3.4 gives further insight the irradiation-induced stability. The left and right panels show β_{cool} at the critical radius for each mass without and with irradiation turned on, respectively. Turning on irradiation increases β_{cool} for all masses, although the effect is minimal for low mass stars. Additionally, the intensity of the irradiation from high mass stars ($\gtrsim 2.5 M_\odot$) overcomes the trend of β_{cool} decreasing as mass increases. The disks around the most massive stars have the highest β_{cool}. Therefore, while irradiated disks become more prone to fragmentation as mass increases when $M_\star \lesssim 2.5\, M_\odot$, they become less likely to fragment for stars above this mass threshold. This non-monotonic behavior cannot be explained by either the explicit inclusion of $\sqrt{M_\star}$ in several of the formulas in Section 3.1.1 or the scaling of R_{crit}. These terms are constant between the $T_{irr} = 0$ and $T_{irr} > 0$ cases. The only explicit inclusion of $M_\star$ in T_{irr} (Equation 3.8) is through the scale height, H. However, $\sqrt{M_\star}$ appears in the denominator of that term, so the only explanation for increasing β_{cool} is that the positive scaling with stellar luminosity overcomes the negative scaling from Ω_{Kep}. This behavior is therefore purely a consequence of the stellar irradiation.

Previous simulations done by Cadman, Rice, Hall, et al., 2020 showed that, for a given disk-to-star mass ratio, irradiated disks are more likely to fragment as stellar mass increases. Since the stellar masses considered in that work do not exceed 2 $M_\odot$, the results are compatible. Comparable simulations for stellar masses in the regime in which β_{cool} increases with stellar mass ($M_\star \gtrsim 2.5\, M_\odot$)

would give further insight into the validity of the model. Of particular interest is the most massive stars, $M_\star \gtrsim 4M_\odot$, which have higher β_{cool} values than even the lowest mass stars.

3.3.2 GI wiggles

The results in Section 3.2.2 further display the stabilizing effects of irradiation relative to both the $T_{irr} = 0$ K method and the Longarini et al., 2021 constant-β_{cool} method. In line with Figure 3.4, Figure 3.5 shows that the GI wiggle tends to increase with stellar mass (left to right), but this trend is reversed for the most massive stars when $T_{irr} > 0$ K (bottom row).

The extracted GI wiggles (Figure 3.6 and Figure 3.7) also show this. The $T_{irr} = 0$ K case is very similar to the $\beta_{cool} = 2$ case in that the wiggles increase in prominence as stellar mass increases. The $T_{irr} > 0$ K case has a less distinct wiggle than the other two methods, and the trend begins to once again reverse for high masses.

Morphologically, this is reflected in the spiral structures shown in Figure 3.8. The spirals are most similar between $T_{irr} = 0$ K and $T_{irr} > 0$ K for $M_\star = 0.1$ $M_\odot$, and they are still fairly similar at 1 $M_\odot$. However, at $M_\star = 7$ $M_\odot$, the spiral structure has almost disappeared for the irradiated disk but is still very clear when irradiation is not included. This is completely in line with expectations as it has been shown (Hall et al., 2020) that the GI wiggle and spiral arms are coincident with one another. Irradiation therefore changes both the disk's morphology and motion.

These results demonstrate the importance of considering both stellar mass and irradiation when analysing the GI wiggle. Irradiation qualitatively changes the relationship between stellar mass and the GI wiggle by creating a non-monotonic trend that is in contrast to both a constant β_{cool} and $T_{irr} = 0$ K treatment. While this demonstrates the rich information contained within the wiggle, it shows that extracting this crucial information is even harder than previously shown.

Importantly, this analysis was done at a single Toomre Q value at a single accretion rate with a constant stellar age. All three of these variables will likely further influence the GI wiggle's properties. These values are reasonable, I are considering unstable disks, which tend to have high accretion rates and are expected to be most present in young systems. Regardless, there is a much wider range of parameter space that can be included to more fully understand the GI wiggle. Additional equations of state and the inclusion of self-shadowing may also lead to new insights. Global 3D simulations of high mass ($M_\star \gtrsim 4$ $M_\odot$) irradiated disks should also be considered.

3.4 Conclusions

This work demonstrates the importance of irradiation and stellar mass in the effects of GI. The differences are morphological and kinematic. Ignoring irradiation qualitatively changes the relationship of GI effects and stellar mass. Irradiation stabilizes relative to no irradiation in all cases. Perhaps more importantly, it creates a non-monotonic relationship that stabilizes the disks around very massive stars relative to lower mass stars. Intermediate mass stars tend to have the most prominent GI wiggles when irradiation is included.

Further exploration throughout the large parameter space is necessary, and thorough simulations may provide additional information. A more complex understanding of the extent to which important system properties influence the morphological and kinematic effects of GI will be essential if I are to effectively analyse potentially GI systems, such as Elias 2-27 (Paneque-Carreño et al., 2021; L. M. Pérez et al., 2016; Veronesi et al., 2021). The transient nature of GI makes any observation crucial to theoretical insights, so it is important that analysis be as informed as possible. This work motivates the explicit inclusion of both stellar mass and irradiation in those considerations.

Chapter 4

Detecting Planets with Machine Learning

*This work was originally published in J. P. Terry et al. 2022 ApJ **941** 192 and J. P. Terry et al. 2023 ApJ **947** 60.*

In recent years, the kinematic analysis of protoplanetary disks has proven important in demonstrating that ring and gap-like features are indeed due to the presence of unseen planets (Pinte et al., 2018; Teague et al., 2018). The spiral wake from planets causes localized deviations from Keplerian motion known as "kinks." Other mechanisms, such as gravitational instability (Hall et al., 2020), vertical shear instability (Barraza-Alfaro et al., 2021), and magnetorotational instability (Simon et al., 2015), may leave similar kinematic imprints.

Recent kinematic observations of the HD 97048 system revealed a localized velocity deviation signaling the presence of a forming exoplanet (Pinte et al., 2019) too faint to be observed directly at ∼mm wavelengths. To estimate the location of the forming exoplanet, several computationally expensive hydrodynamics simulations were run, synthetic observations were generated and compared "by-eye" to real observations to obtain the best fit.

Machine learning offers an alternative route. The use of machine learning in astronomy has exploded in the last decade (Alexander et al., 2020; Jo & Kim, 2019; Möller & de Boissière, 2020). The study of protoplanetary disks largely relies on images, so machine learning techniques for computer vision are perfectly suited for application to these kinematics observations. Computer vision is a field of AI that focuses on image and video analysis. From simple transcription of hand-written digits (Cireşan et al., 2011) to locating and classifying objects pixel-by-pixel (Minaee et al., 2020), computer vision has a wide range of uses. The introduction of advanced architectures, such as residual (He et al.,

2015) and inception (Szegedy et al., 2014) blocks, along with qualitatively new algorithms, such as the vision transformer (Dosovitskiy et al., 2020; Vaswani et al., 2017), has led to computer vision models exceeding human experts at times (Zhou et al., 2021).

Given the applicability and strong performance of computer vision techniques, I have applied them to the kinematic analysis of protoplanetary disks.

Previous work has successfully applied deep learning to the problem of inferring properties of planets within protoplanetary disks (Auddy et al., 2021, 2022) using models trained on simulated data, but these efforts do not use kinematic information and instead rely on gaps as the signatures. I demonstrate that a trained network can detect the presence of planets in large datasets much faster than a human, recover all previously identified planets, and possibly identify ones that have been missed in previous analyses. Inferring other properties, such as the number of planets and their masses, may be within the reach of machine learning techniques as well, which I leave to future work. This paper presents a proof of concept application of a relatively simple machine algorithm applied to protoplanetary accretion disks.

4.1 Methods

Training a model to predict the presence of a planet requires a large amount of data spanning a large parameter space. I run 1,000 simulations in this work. The parameter space is sampled using a Latin Hypercube (LHC) (McKay et al., 1979), which accounts for previously sampled points to evenly distribute values across the entire parameter space. I sample values from observed ranges inferred from disk surveys (Andrews et al., 2009, 2018; J. Huang, Andrews, Dullemond, et al., 2018) and widely accepted simulational parameters (Pinte et al., 2018, 2019, 2020).

Table 4.1 shows the variables sampled and their minimum and maximum value. Approximately one third of the simulations have no planets in to address any class imbalance (i.e., not having an equal number of “at least one planet" class and “no planet" class simulations) while ensuring diversity in the number of planets. After sampling the stellar mass for use in the simulation, stellar temperatures and radii, both of which are used in the radiative transfer calculations, are determined using isochrones calculated by Siess et al., 2000.

Table 4.1: Simulation parameter ranges for the LHC sampling.

Parameter	Variable	Minimum Value	Maximum Value
Number of planets	N	0	4
Disk Inner Radius (au)	R_{out}	150	250
Disk Outer Radius	R_{in}	$0.1R_{out}$	$0.1R_{out}$
Planetary Mass ($M_\oplus$)	M_{planet}	5	120
Stellar Mass ($M_\odot$)	$M_\star$	0.5	2.0
Planet Semimajor Axis (au)	a	$R_{in} + 0.2R_{out}$	$R_{in} + 0.8R_{out}$
Mass Ratio	$M_{\rm disc}$ / $M_\star$	5×10^{-4}	10^{-2}
Disk Viscosity	$\alpha_{\rm SS}$	5×10^{-4}	2.5×10^{-3}
Surface Density Profile ($\Sigma = \Sigma_0 R^{-p}$)	p	0.1	1.0
Sound Speed Profile ($c_s = c_{s,0} R^{-q}$)	q	0.1	0.75
Stellar Age (Myr)	t	0.5	5.0
Inclination (degrees)	i	10	80
Azimuth (degrees)	ϕ	0	360
Distance (pc)	d	100	200
Spectral Resolution (km/s)	Δv	0.01	0.1
Spatial Resolution (arcseconds)	$''$	0.025	0.25
RMS Noise (%)	δ	0.05	20

4.1.1 Hydrodynamical simulations

I run 1,000 3D smoothed particle hydrodynamics (SPH) simulations using the `PHANTOM` (Price et al., 2018) code. System properties, such as stellar mass, disk properties, observational resolution etc. were LHC-sampled from Table 4.1. A significant fraction (25%) of these simulations were withheld for testing. Each simulation uses 10^6 SPH particles and evolves for 100 orbits. A given planet's accretion radius is set to 1/4 of its Hill radius. The surface density profile is set to $\Sigma \propto R^{-p}$, and the sound speed is set to $c_{\rm s} \propto R^{-q}$.

4.1.2 Velocity channel maps

The simulation results are used to create velocity channel maps in only ^{13}CO to keep the size of the parameter space manageable, although ^{12}CO and C^{18}O are also reasonable choices. The channel maps are created using the `MCFOST` radiative transfer code (Pinte et al., 2006, 2009). For each simulation, multiple points between 10 and 100 orbits are chosen randomly. Since one third of the simulations have no planet, six snapshots are chosen from each of those simulations while only two snapshots are chosen for simulations with planets. This

eliminates the class imbalance. From each snapshot, a velocity channel cube is made using either the ^{13}CO J 2→1 (220.4 GHz) or 3→2 (330.6 GHz) transition, which is randomly sampled. These transitions are used because they were also used to identify the planets in HD 163296 and HD 97048 (Pinte et al., 2018, 2019). The final result is a velocity cube, where each velocity channel is an image with 600x600 pixels at a resolution of 1 au/pixel. This velocity cube is used as the model input. Channel maps are calculated using 10^8 photon packets. The dust-to-gas ratio is 1:100, and dust is composed of carbon and silicates (Draine & Lee, 1984). Stellar radii and temperatures are calculated with isochrones given by Siess et al., 2000 using the LHC-sampled masses and ages. A spherical beam is used.

4.1.3 Simulated observations

The velocity channel maps are convolved spatially and spectrally, and noise is added by LHC sampling from Table 4.1. RMS noise is added independently for each pixel, which more faithfully replicates observational effects such as hot pixels. This method results in an average RMS of $\sim 10\%$ for each image.

Except for the specific values, this process largely follows that done in Terry, Hall, Longarini, et al., 2022. Figures 4.1 and 4.2 show the results of convolving selected velocity channels for a disk with and without a planet, respectively. Figure 4.1 shows that while the perturbation is strongest in the velocity channel that covers the planet (rightmost column), the perturbation is not necessarily completely localized and can be visible in distant regions of the disk. This is an expected behavior. The deviation is strongest at the location of the planet, but the spiral wake leaves further imprints as it disperses throughout the disc.

By itself, a single velocity channel is poorly suited for this method. Only a few channels display kinks, and channels with kinks can neighbor the exact channel containing the planet. I therefore stack all C velocity channels into a single image of dimension (H, W, C), where H and W are the height and width of the image (600x600 pixels for these purposes).

The number of velocity channels in an observation varies, but mostly fall in the range $\sim 40 - 80$ channels that cover the disk (Andrews et al., 2018; Öberg et al., 2021). I therefore use C=47, 61, and 75 channels for each observation. The number of channels must be odd to cover the disk symmetrically and still retain the systemic ($\delta v = 0$) channel. The input velocity of the simulated channels ranges from $|\delta v| = |v_{systemic} - v| \leq 2.5$ km/s to fully cover the disk.

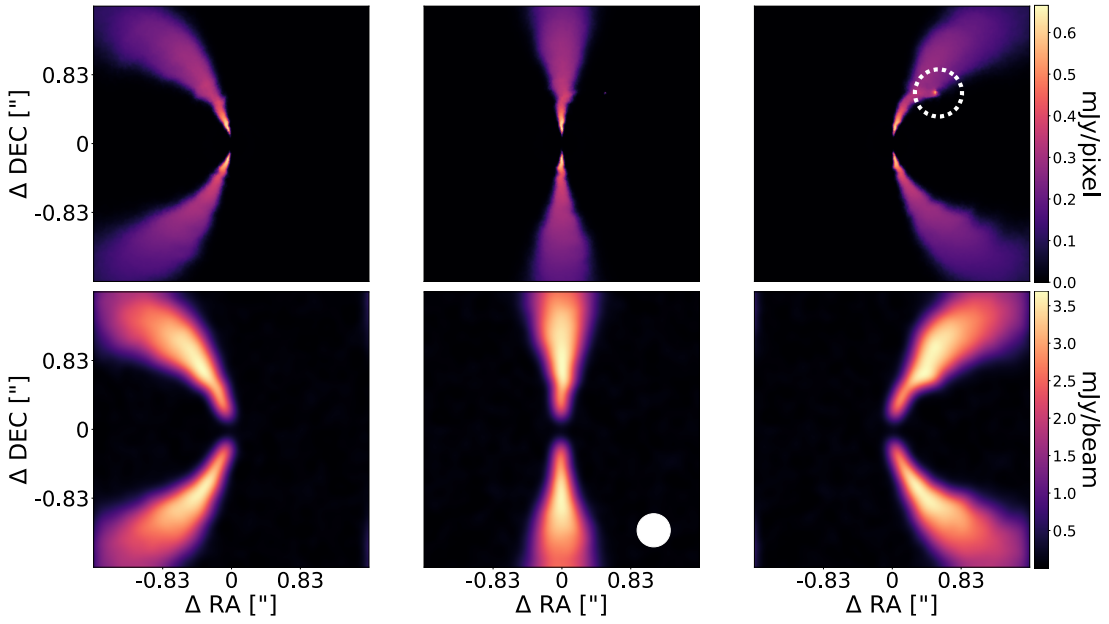

Figure 4.1: Example raw (top row) and convolved, noisy (bottom row) channel maps in a disk with a planet present. The planet (circled in white) is visible as a kink in the right column. The opposite velocity channel is shown in the left column, and the systemic channel is shown in the middle column. The beam size is indicated in the bottom middle image. This disk is one of the smallest and farthest simulated and is observed with some of the worst spatial resolutions, which is why the beam is so large.

4.1.4 Machine learning models

The input for all models is a (60ox60oxC) image, and their output is a two-component vector with softmax activation (Goodfellow et al., 2016). The softmax function maps the output vector such that the sum of the components equals one. Each component can be interpreted as the predicted probability that the input belongs to the corresponding class (i.e., "contains at least one planet" class or "does not contain a planet" class). An Adam optimizer (Kingma & Ba, 2014) is used to add momentum, which remembers previous gradients, and gradient scaling to a simple gradient descent algorithm. This helps escape local minima and improves the speed and range of convergence. The cross-entropy loss (Hinton et al., 1995) penalizes confident predictions that are incorrect. Training stops when there is no improvement in the validation loss after eight complete epochs, i.e. eight sweeps through the training data.

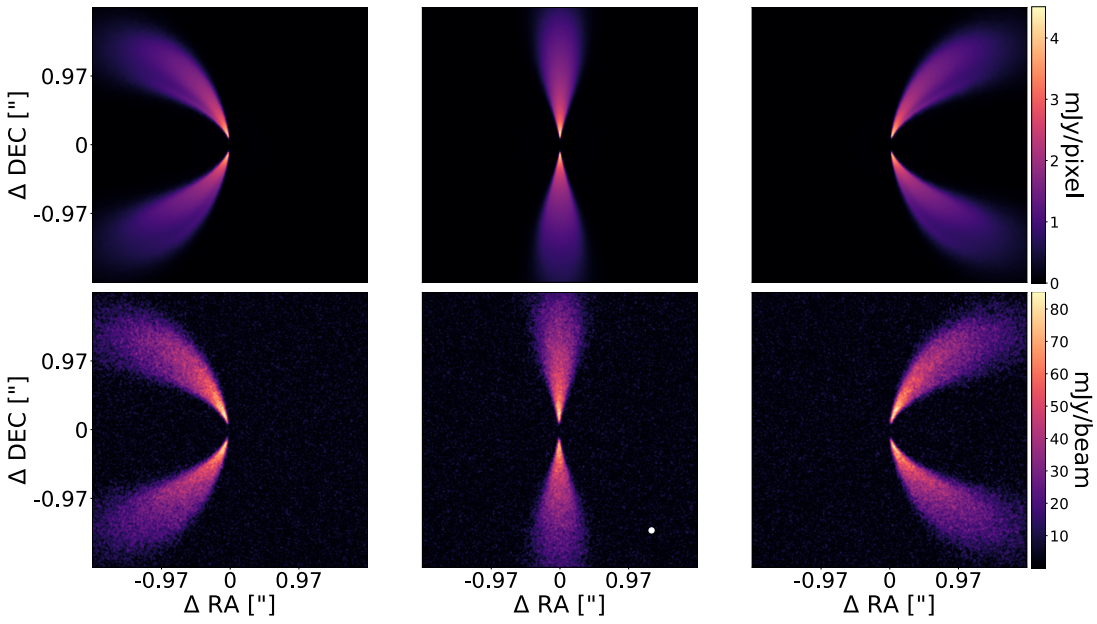

Figure 4.2: Example raw (top row) and convolved, noisy (bottom row) channel maps in a disk without a planet present. The beam size is indicated in the bottom middle image.

Multiple architectures are used to compare performance on two tasks: predicting if the system hosts at least one planet and determining its azimuthal location.

I use two different models based on PyTorch (Paszke et al., 2019) implementations: EfficientNetV2 (Tan & Le, 2021) and RegNet (Xu et al., 2021). Neither model uses the default hyperparameters or pre-trained weights.

EfficientNetV2 is an updated version of EfficientNet (Tan & Le, 2019) that maximizes performance and minimizes the number of trainable parameters and training time. RegNet, on the other hand, seeks to maximize performance by adding variants of convolutional recurrent neural networks to ResNet (He et al., 2015). This can extract greater spatial information at the cost of significantly more trainable parameters.

I perform Bayesian hyperparameter tuning using `WANDB` (Biewald, 2020) to find separate sets of hyperparameters that minimize the validation loss for C=47, 61, and 75. The default parameters for these sweeps are based on default versions "EfficientNetV2 S" and "RegNetY 16GF."

Each architecture is evaluated for accuracy and area under the receiver operating characteristic curve (AUC) (Hanley & McNeil, 1982). Accuracy in this context is defined as whether a model correctly predicted that a system with planet(s) has at least one planet, and similarly for systems with no planets. I choose AUC because it gives a robust measure of the performance of the model across different decision thresholds, i.e. the softmax threshold required to claim the presence of a planet. Substantial confidence is required to claim a planet detection, but setting the decision threshold too high can lead to false negatives. False positives may be more acceptable because further analysis can overrule them, whereas a false negative may lead to no further analysis and result in the planet not being detected. The optimum threshold for this boundary is open for discussion by the community. I show multiple thresholds (50%, 75%, 90% and 95%) in Figures 4.3 and 4.4 to demonstrate how the threshold for claiming the presence of at least one planet affects the results .

While some protoplanetary disks exhibit gaps in continuum that are highly suggestive of a planet, this does not help one find the azimuthal location. Planet-induced non-Keplerian kinks tend to be strongest near the planet, so the model must determine if and where the kink is in the data in order to give information on the azimuthal location.

One simple approach to this is to look at the activations inside the model itself. The activation strengths highlight regions which the model finds particularly informative for classification. EfficientNetV2 uses SiLU activations (Elfwing et al., 2017), and RegNet uses ReLU (Glorot et al., 2011).I use mean-subtracted activations to achieve this. For a given pixel with an activation value, x, the mean-subtracted activation, x', is defined as $x' = |x - \langle x \rangle|$, where $\langle x \rangle$ is the mean activation over the entire image. Activations occur most strongly at the location of the strongest deviation from Keplerianity, i.e. the kink.

4.2 Results and discussion

4.2.1 Machine learning performance

I train six models in total: an EfficientNetV2 and a RegNet for 47, 61, and 75 velocity channels. Using a typical decision threshold of 0.5, each model, other than the 47-channel RegNet, has an accuracy of at least 93%. When this decision threshold is increased to 0.95, four models have an accuracy of at least 92%. Five models have an AUC of at least 0.98. This confirms that the models have learned to correctly detect the presence of planets using kinematic signatures.

Table 4.2: Final model descriptions and metrics. 95% confidence intervals are calculated by bootstrapping each metri
times using random selections of 80% of the test data. The percentages next to the accuracy labels denote the decision thr

Model	Parameters (millions)	Accuracy (50%) (%)	Accuracy (95%) (%)	AUC
EN47	20.2	97 ± 0.5	96 ± 0.5	0.99 ± 0.002
EN61	20.2	97 ± 0.5	94 ± 0.7	0.99 ± 0.003
EN75	20.2	93 ± 0.7	88 ± 0.9	0.98 ± 0.003
RN47	51.0	78 ± 1.1	65 ± 1.3	0.86 ± 0.010
RN61	62.8	98 ± 0.4	96 ± 0.6	$>0.99 \pm 0.001$
RN75	114	95 ± 0.6	92 ± 0.7	0.98 ± 0.003

Table 4.2 shows the number of parameters as well as relevant performance measures for each final model. All EfficientNetV2 models converge to a similar number of trainable parameters, whereas the number of parameters for the RegNets increases sharply with the number of channels. RegNet is a larger and more complicated architecture, so this is expected. The performance of the RegNet models varies widely compared to the EfficientNetV2 models. The accuracy of EfficientNetV2 models varies by at most 8%, whereas RegNet varies by as much as 31%. By all performance metrics, while the average EfficientNetV2 outperforms the average RegNet, the 61-channel RegNet outperforms all other models.

Figure 4.3 shows the confusion matrix, which describes the distribution of correct and incorrect predictions, for all models applied to the withheld test set. Panel (a) shows the results using a 50% decision threshold, i.e. a planet is predicted if the confidence is over 50%. Panels (b), (c), and (d) show the same for cutoffs of 75%, 90%, and 95%. Figure 4.4 gives several relevant metrics for model performance and all ROC curves. Claiming a planet is present should come with high confidence, so a 50% decision threshold isn't necessarily what should be used in practice. It is encouraging that, for all other models than the 47-channel RegNet, there are no qualitative changes in the results when a higher decision threshold is used; the main effect is simply increasing the number of false negatives and decreasing the number of false positives.

If there are multiple planets in a disk, then multiple kinks may be present. It is therefore prudent to check that the quality of the predictions does not depend on the number of planets in the disk. I check the possibility of this dependence by calculating the accuracy for each model for disks with a given number of planets. Figure 4.5 gives the results of these calculations. The p-

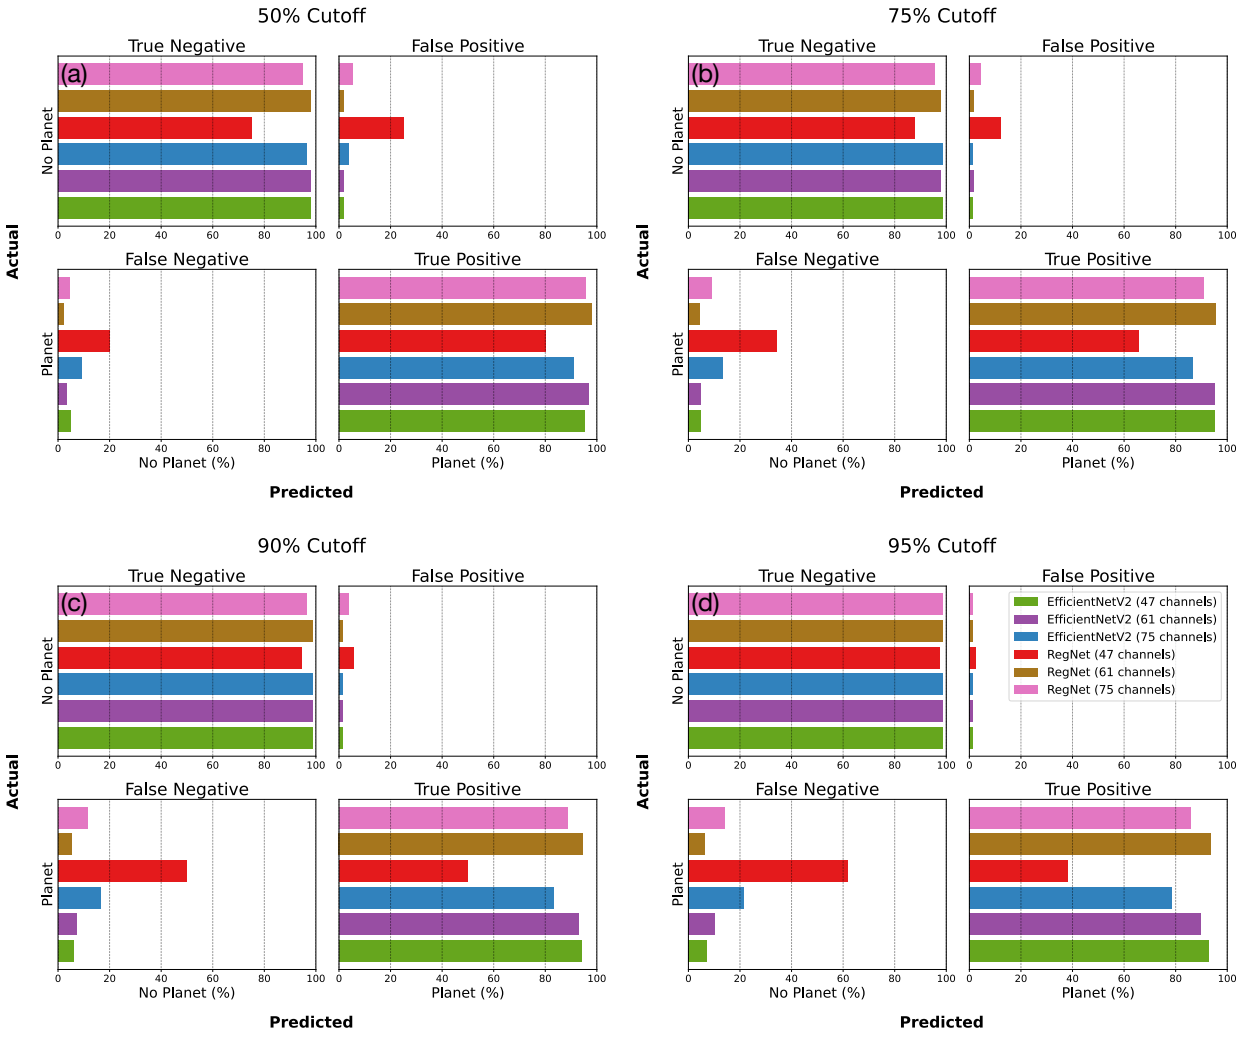

Figure 4.3: Confusion matrix from withheld test set for all models. The x axis is the percent of disks in that class that were predicted accordingly. (a): Using a 50% decision threshold. (b): Using a 75% decision threshold. (c): Using a 90% decision threshold. (d): Using a 95% decision threshold. The top left and bottom right blocks of each figure show the counts of true negatives and true positives, respectively. The top right and bottom left blocks of a figure show the false positives and false negatives, respectively. A perfectly accurate model will have no entries in the top right and bottom left.

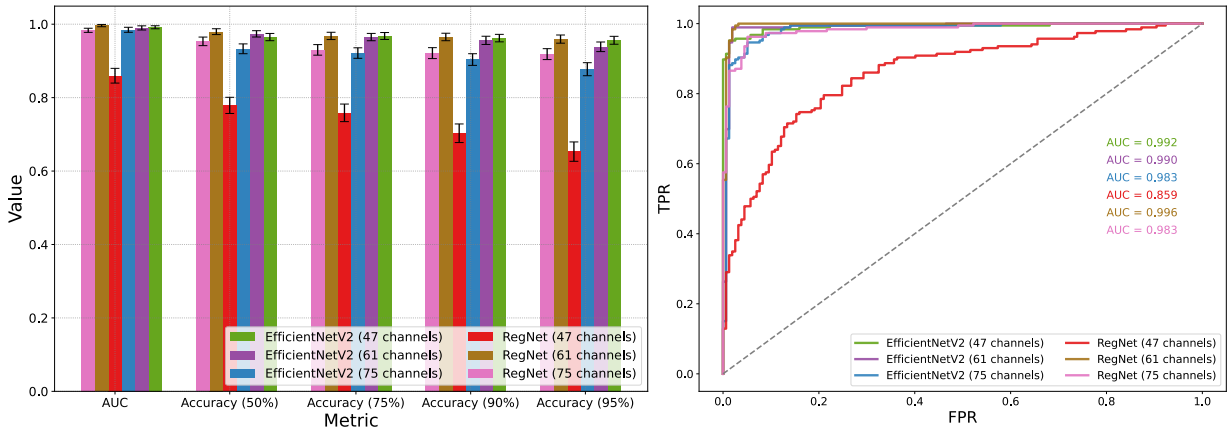

Figure 4.4: Left: Various metrics calculated from the withheld test set for all models. Right: Corresponding ROC curves. Error bars show 95% confidence intervals that are calculated by bootstrapping each metric 1,000 times using random selections of 80% of the test data. The percentages next to the accuracy labels denote the decision threshold.

values from Pearson correlation tests shown in Figure 4.6 demonstrate that only the 47-channel RegNet using a decision threshold of 50% has a statistically significant relationship between accuracy and the number of planets ($p < 0.05$). While the 47-channel RegNet using a decision threshold of 95% does not have a statistically significant relationship, it is clear that the performance is worse when there are planets, i.e. there are many false negatives. This is in line with this model's relatively low AUC and Figure 4.3. The performance of all other models does not differ significantly with the number of planets.

4.2.2 Simulation activations

Many disks do not show obvious signs of planets, either in continuum or line emission. Inspecting the activations of the disks helps overcome this. Figure 4.7 shows the result from a disk that is difficult to classify by eye. The models predict the presence of an exoplanet with high confidence, activating on a subset of channels that are deemed important by the model. This gives information on the azimuthal location, and the radial location may be inferred by either the location with the strongest kink or coincident gaps in continuum images.

Panels (c) and (d) in Figure 4.7 show the raw and convolved velocity channels, respectively. If one were simply inspecting the velocity channels by eye, it would be trivially easy to miss the kink even in the unconvolved data. The convolved, noisy data shows essentially no sign of the planets. Despite lacking any clear signatures that can be observed visually, the models correctly classify

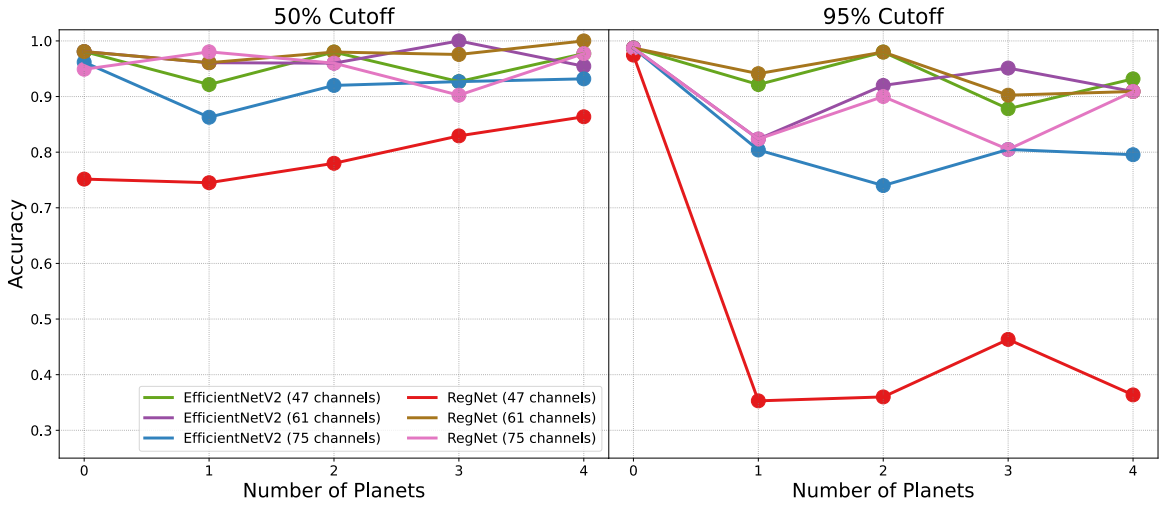

Figure 4.5: Accuracy of all models for a given number of planets. Left: 50% decision threshold. Right: 95% decision threshold

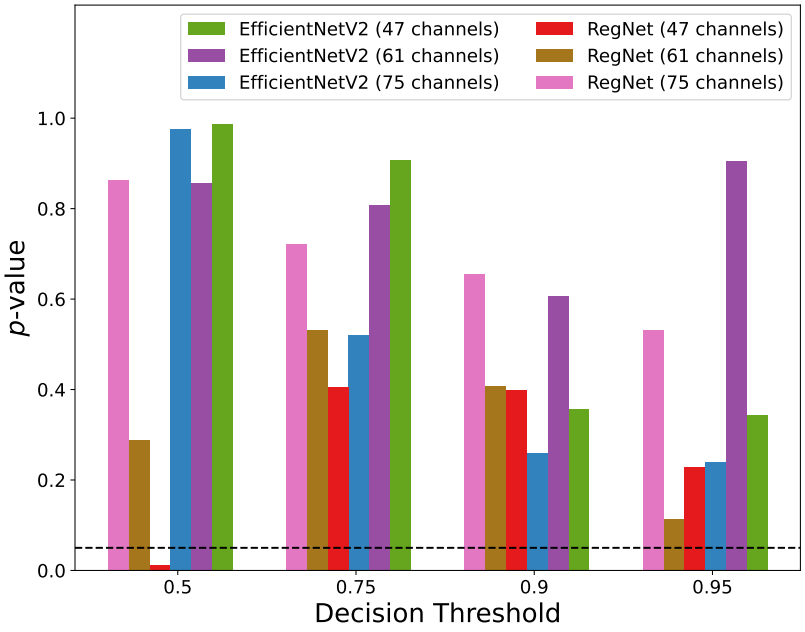

Figure 4.6: p-values from testing the correlation between accuracy and the number of planets at different decision thresholds. The black line is $p = 0.05$, which is the standard cutoff for statistical significance.

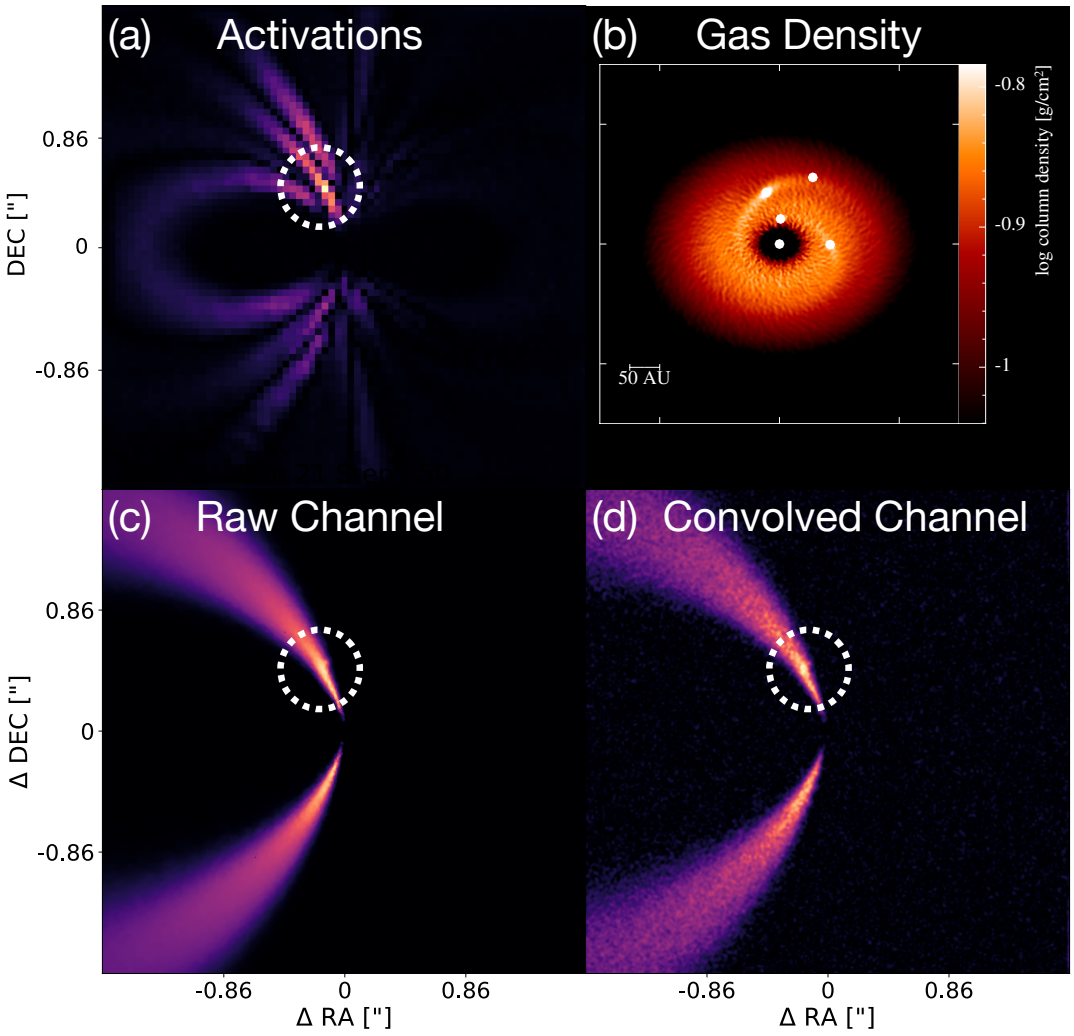

Figure 4.7: (a): Mean-subtracted activations using the 47-channel EfficientNetV2. The innermost planet has M=0.28M_J, the second planet has M=0.20M_J, the third planet has M=0.36M_J, and the outermost planet has has M=0.02M_J. (b): Disk column density after an azimuthal rotation of $\phi = 131°$ and inclined by $i = 36°$. (c): Unconvolved velocity channel with the most obvious kink. (d) Corresponding convolved velocity channel. Based on the kink, the approximate location of the planet causing the most significant kink is in the white circle. The planets were correctly predicted with a confidence of $> 99\%$.

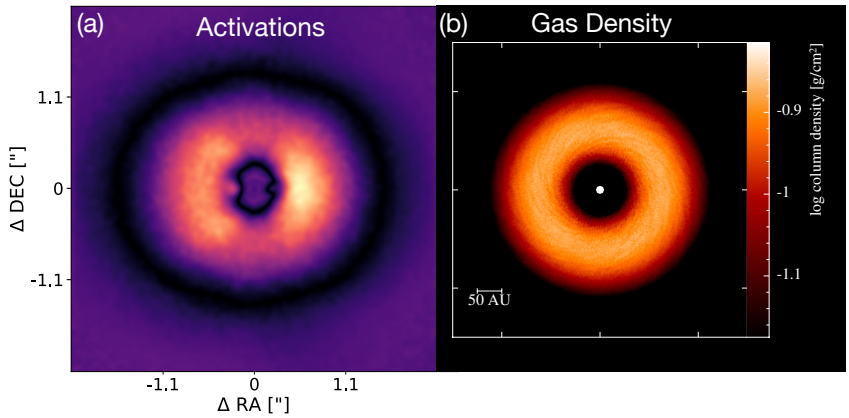

Figure 4.8: (a): Typical example of mean-subtracted activations for a disk with no planet. (b): Disk column density after an azimuthal rotation of $\phi = 153°$ and inclined by $i = 13°$.

the disks and locate the planets. This demonstrates that the models are able to accurately find planets that the by-eye method may fail to find.

Figure 4.8 shows a typical activation for a disk with no planet. The entire disk is activated equally, indicating that there is no localized region that exhibits behaviors indicative of a planet.

4.2.3 Application to real observations

Next, I demonstrate the effectiveness of the model by applying it to real telescope observations. Kinematic detections of planets have been demonstrated for many disks (Pinte et al., 2020), but here I focus on the HD 163296 and HD 97048 disks as a proof of concept (Pinte et al., 2018, 2019). Using the same data as Pinte et al., 2018 and Pinte et al., 2019, I demonstrate that the models replicate the prediction and estimated location of the forming exoplanet. Figure 4.9 (a) shows the HD 97048 in continuum, and panel (b) shows the velocity channel with the planet's signature circled.

Table 4.3 shows the softmax values for the observations. For HD 97048, all RegNet models correctly predict a planet with $> 95\%$ confidence. Two EfficientNetV2 models correctly predict a planet with $\geq 85\%$ confidence. Figure 4.9 shows that the velocity channels that contain the planet are strongly activated. The activation in Figure 4.9 (d) is strong on the planet itself rather than simply the entire velocity channel.

Table 4.3: Prediction confidences for all models tested on HD 97048 and HD 163296.

Model	HD 97048 Softmax	HD 163296 Softmax
EN47	85%	37%
EN61	89%	69%
EN75	64%	73%
RN47	$> 99\%$	74%
RN61	96%	$> 99\%$
RN75	$> 99\%$	$> 99\%$

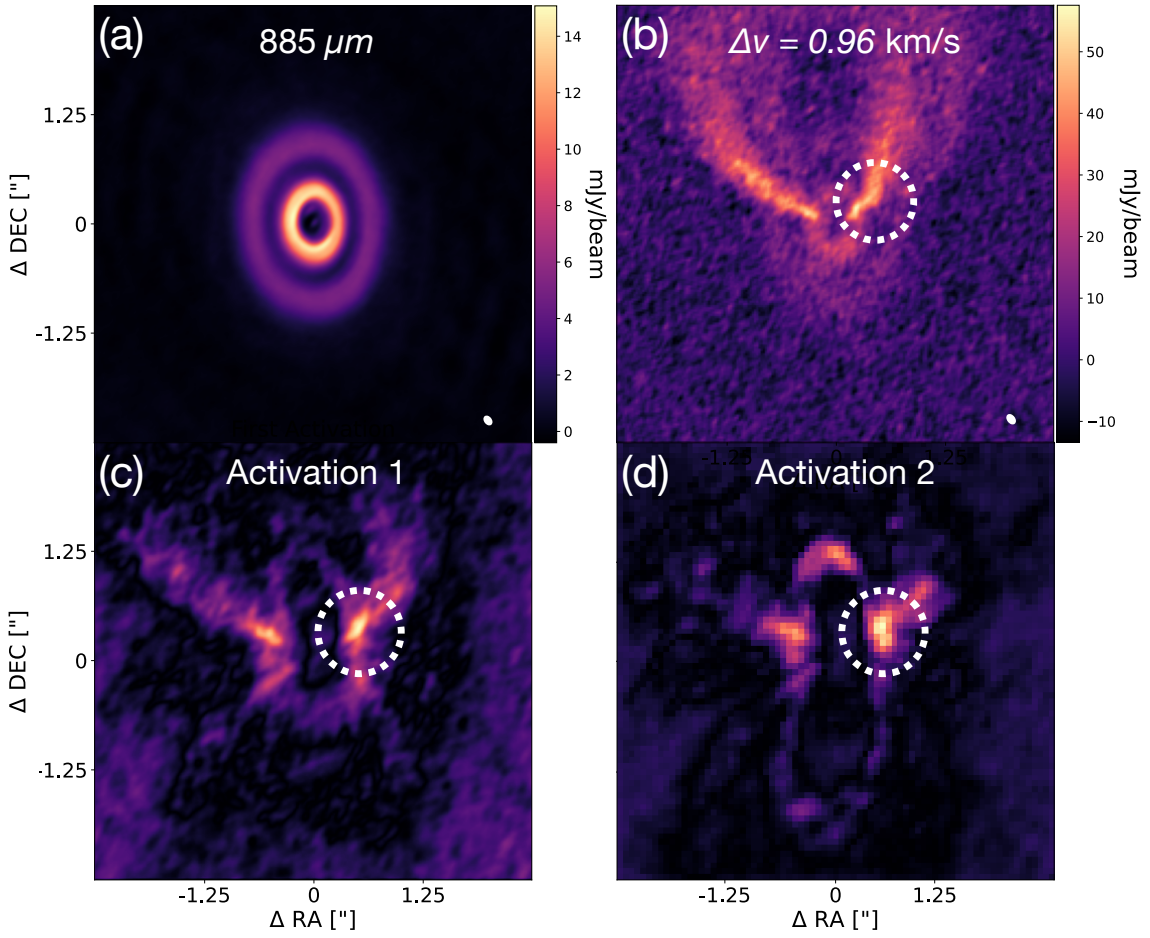

Figure 4.9: (a): HD 97048 at 885 μm. (b): $\Delta v = 0.96$ km/s channel. Bottom: Mean-subtracted activations from two different layers for HD 97048 using the 47-channel RegNet. The estimated planet location is given inside the white circle. The beams for the observations are given in the bottom right corner.

The results for HD 163296 are not as confident, but the two best-performing models predict the presence of at least one planet with over 99% confidence.

Both classes of model perform better when there are more input channels, a trend not seen for HD 97048. A possible explanation is the fact that there are about twice as many velocity channels that cover HD 163296 when compared to HD 97048. While HD 97048 has roughly 75 channels covering the disk, HD 163296 has roughly 150. All models applied to HD 163296 must therefore combine significantly more velocity channels, which may decrease its ability to extract information from the images. The activation structure tends to be less informative, perhaps for the similar reasons, so it is omitted. Despite this limitation, the high confidence of some of the models should encourage a human observer to seriously scrutinize the kinematic structure.

4.2.4 Limitations and future work

This work comes with several limitations. All simulations were run with SPH, so I did not demonstrate that the methods would also be applicable to data from grid-based codes. However, there is typically a qualitative agreement between results from SPH and grid-based methods (Agertz et al., 2007; de Val-Borro et al., 2006). A good direction for future work would be to train on simulation data from a variety of hydrodynamics codes, and check that the same results can be obtained.

Any differences in performance of the models between the codes, e.g. if the models consistently performed better on SPH data, would be difficult to interpret and the cause could be hard to identify. It is possible that the method of domain adaptation (Ben-David et al., 2010; Ćiprijanović et al., 2021), which has already found success in astronomy (Alexander et al., 2023; Ćiprijanović et al., 2022; Vilalta et al., 2019), could be used to encourage the models to overcome any differences between datasets. This is an avenue that is ripe for exploration in future work.

In order to keep the parameter space manageable and demonstrate a proof of concept in this work, I therefore trained on data produced by only one hydrodynamics code (`PHANTOM`).

The absence of true image segmentation (Minaee et al., 2020) or object detection (Zhao et al., 2018) during training is another limitation that requires addressing in future work. While I present results in the form of activation structures that suggest the location(s) of the planet(s), these activations do not always perfectly coincide with the planets, nor capture all planets. User analysis is required in order to pinpoint the radial (and sometimes azimuthal) location(s) of the planets even with the information given by the activation structures. The

main utility of the presented models is therefore to give a user an indication as to whether a planet is present and motivate and direct the analysis of the velocity channels. Segmentation and object detection would address these shortcomings by precisely pinpointing the location(s) of the planet(s) without any necessary further analysis.

The assertion of a vertically isothermal disk is a simplifying assumption that is commonly used in simulations of protoplantary discs (see, e.g., Dipierro et al., 2015; Pinte et al., 2019) . It has been demonstrated that the dispersion relation and morphology of the planet's signature can change significantly if the disk is vertically stratified (Juhász & Rosotti, 2018; Ogilvie & Lubow, 2002).

However, the presence of any disturbance is more important than its specific morphology. There does exist the possibility that vertical stratification dampens the signatures to an extent that the models' performances appreciably decrease, but the results demonstrate that even extremely subtle signals can be found by the models. Regardless, a more realistic treatment of the thermal structure may be useful in future works.

4.3 HD 142666

I apply the previously trained models to ALMA data of the HD 142666 system. I inspect the softmax values and activation structures to gain insight into whether a planet might be in the system and, if so, where its signature is the strongest.

4.3.1 Observational data

This chapter makes use of the following ALMA data: ADS/JAO.ALMA #2016.1.00484.L and ADS/JAO.ALMA#2016.1.00826.S. ALMA is a partnership of ESO (representing its member states), NSF (USA) and NINS (Japan), together with NRC (Canada), MOST and ASIAA (Taiwan), and KASI (Republic of Korea), in cooperation with the Republic of Chile. The Joint ALMA Observatory is operated by ESO, AUI/NRAO and NAOJ.

The HD 142666 data was taken from the DSHARP catalogue (Andrews et al., 2018; J. Huang, Andrews, Dullemond, et al., 2018). Data includes ^{12}CO line emission ($J = 2 \rightarrow 1$) and 1.25 mm continuum images. The system was imaged with a beam with FWHM of 77×61 mas ($\approx 11 \times 9$ au) with an RMS noise of 1.3 mJy beam^{-1}; channels have a 0.35 km s^{-1} resolution (Andrews et al., 2018). Figure 4.10 shows selected velocity channels overlaid on the continuum.

The image was cropped to focus on the disk, and a subset of velocity channels was used. The channels were reshaped to 600×600 pixels and normalized such that all pixel values were between 0 and 1.

4.3.2 Hydrodynamical simulations of HD142666

I run a suite of SPH simulations using `PHANTOM`, varying the mass of the embedded planet between 1 and 5 M_J. For each simulation, I create channel maps using `MCFOST` in the same way that the original training data was made. The kink is approximately 75 au from the center of the disk, so I place a planet at this distance.

System parameters are taken from Andrews et al., 2018; J. Huang, Andrews, Dullemond, et al., 2018; Rubinstein et al., 2018. The stellar mass, temperature, and radius are 2.0 $M_\odot$, 7500 K, and 2.2 $R_\odot$, respectively. The disk has a mass of 0.0533 $M_\odot$, an inner radius of 1.3 au, and an outer radius of 150 au. The system is inclined at 62 degrees with a position angle of 162 and an azimuth of 72 degrees. It is located 148 pc from Earth.

The SPH outputs are used to create line emission maps to mimic ALMA capabilities. These calculations are done using the `MCFOST` radiative transfer code (Pinte et al., 2006, 2009). Each calculation uses 10^8 photon packets and includes carbon/silicate dust (Draine & Lee, 1984) with a dust-to-gas ratio of 1:100. The resulting outputs were convolved spatially and spectrally to match the observed line emission resolution.

4.3.3 HD142666 results and discussion

Figure 4.11 shows that HD 142666 has a strong, localized kink that is detected by the ML models. The kink is particularly visible in the upper middle ($\Delta v = -1.75$ km/s) channel. The lower row shows activation structures that roughly correspond to the above channels.

The average softmax value is over 0.84, which means that the models predict the probability that the input for HD 142666 contains a planet to be over 84%. This prompts further scrutiny of the activations, which I use to determine the most probable channel that contains the kink.

The strength and localization of the newly identified kink are reminiscent of the kinks in HD 163296 and HD 97048. As with HD 163296, the kink in the gas is outside of the radial extent of the continuum disk. Both of these disks were found to host planets after SPH simulations containing a planet recreated the kinematic structure observed in CO observations (Pinte et al., 2018, 2019; Teague et al., 2018). I apply this same method to HD 142666 to demonstrate

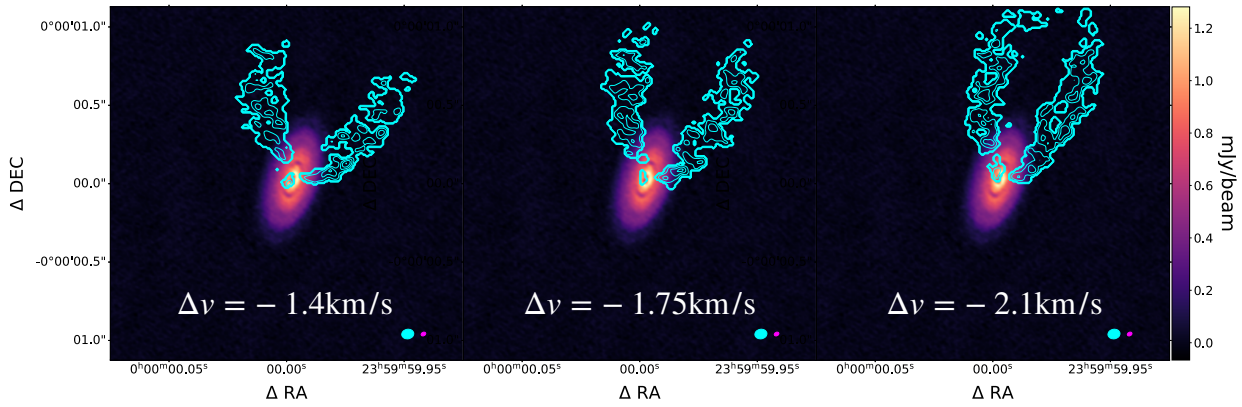

Figure 4.10: Line emission overlaid on continuum. Left: $\Delta v = -1.4$ km/s channel. Middle: $\Delta v = -1.75$ channel. Right: $\Delta v = -2.1$ channel. The continuum beam is in magenta, and the line emission beam is in cyan.

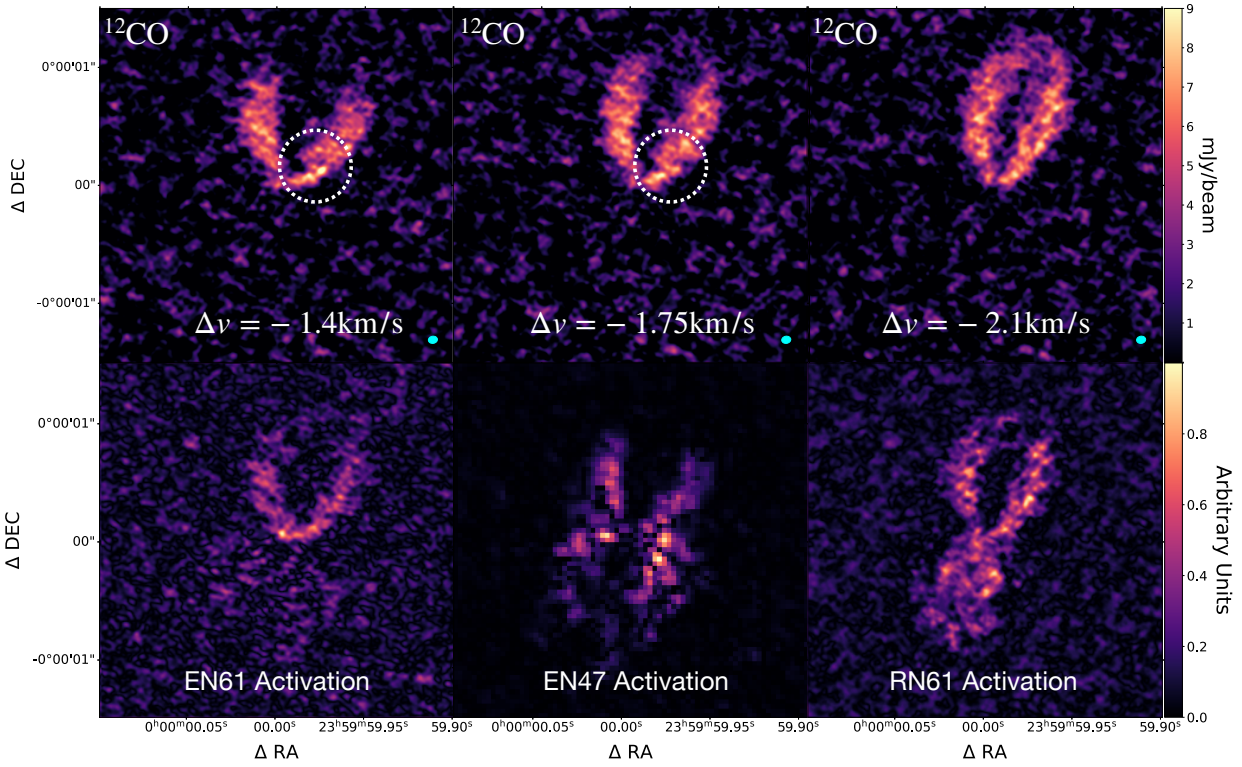

Figure 4.11: HD 142666 structure (^{12}CO: $J = 2 \rightarrow 1$) and activations. Upper left: $\Delta v = -1.4$ km/s channel with kink circled in white. Upper middle: $\Delta v = -1.75$ channel with kink circled in white. Upper right: $\Delta v = -2.1$ channel. Bottom row: selected mean-subtracted activations that roughly correspond to the channels in the upper row. Activations are from three different models (EN61, EN47, and RN61, respectively). Line emission beams are the cyan ellipses in the lower right of the upper row panels.

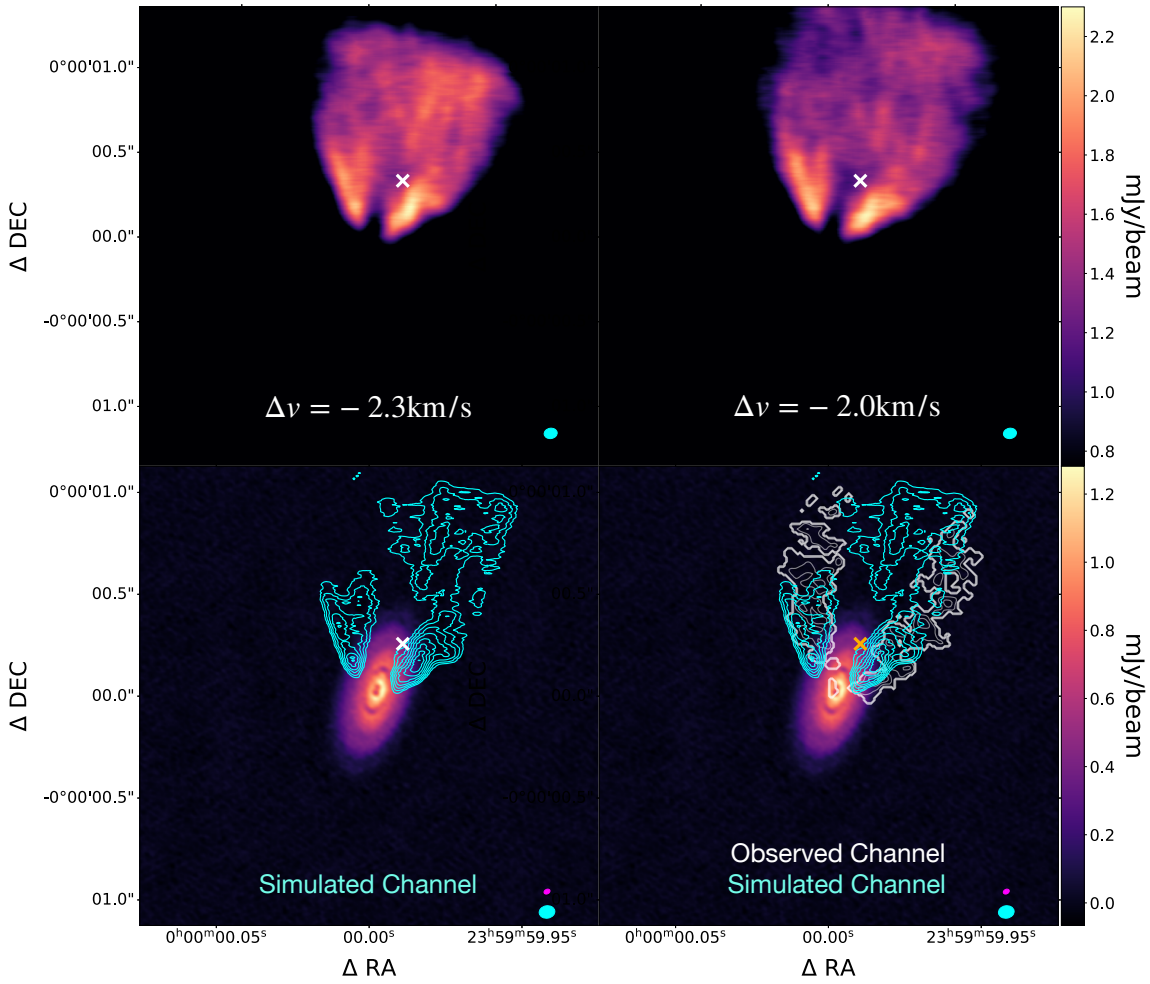

Figure 4.12: HD 142666 simulation results. Upper left: $\Delta v = -2.3$ km/s channel from the simulation (convolved beam in lower right). Upper right: $\Delta v = -2.0$ km/s channel from the simulation (convolved beam in lower right). Lower left: observed continuum overlaid with contours of simulated $\Delta v = -2.3$ km/s channel. Lower right: observed continuum overlaid with simulated (cyan) and observed (white) channels. Continuum beam is in magenta, and the line emission (simulated and observed) beam is in cyan. The system includes a 5 M_J planet at 75 au. The simulated channels have the continuum and background subtracted for clarity. The planet's location is indicated with an x.

that the kink identified by the models is consistent with kinks identified by conventional means in HD 163296 and HD 97048.

I found that a simulation of a protoplanetary disk with a 5 M_J planet most accurately reproduced the observation. Figure 4.12 shows the results. A localized kink in the vicinity of the planet is clear in the upper left panel Figure 4.12 ($\Delta v = -2.3$ km/s). This kink is visible to a lesser extent in the $\Delta v = -2.0$ km/s in the upper right panel of Figure 4.12, which is also the case in Figure 4.10. There is strong agreement between this feature and the non-Keplerian channel identified by the models: both display a kink of approximately the same shape and size at approximately the same radial location. This can be seen in the lower left and right panels of Figure 4.12. Note that the simulation and observation do not display the strongest kink in the same velocity channel. This is simply a relic of the finite temporal resolution of the simulation, which makes it extremely unlikely that the simulation will be saved when the planet is exactly coincident with the observation. The temporal resolution of the simulation was increased to mitigate this effect, but it persists to some extent.

I conclude that HD 142666 hosts a planet.

I note that the conclusion is confirmed using the same methods described by Pinte et al., 2018, 2019; Teague et al., 2018. However, what is new about this approach is that the non-Keplerian motion was first identified by ML models, highlighting a protoplanet candidate that had previously been missed upon visual analysis. Verification of the evidence is still done using the same methodology as previous works (Pinte et al., 2018, 2019). I strongly advocate that this should always be done for any potential discovery.

Networks such as PGNets (S. Zhang et al., 2022) and DPNNet-2.0 (Auddy et al., 2021) offer a potentially fruitful route that would increase the accuracy and speed of the analysis of disks and channels highlighted by the models. These networks are designed to infer planetary mass from continuum images. One could use the models to determine if it is likely that a disk hosts a planet and, if so, feed the corresponding continuum images into the secondary networks. The predicted planet mass could then be used as a starting point for followup simulations rather than simply starting an uninformed parameter sweep. This would speed up the verification step of the procedure. Such a pipeline could be useful to explore in future works

Chapter 5

Anomaly Detection in Protoplanetary Disks

Protoplanetary disk astronomy is undergoing a revolution. The commissioning of observatories, such as The Atacama Large Millimeter/submillimeter Array (ALMA), has revolutionized our understanding of protoplanetary disks (PPDs) and planet formation. Observations of systems such as HL Tau (ALMA Partnership et al., 2015) and large surveys including Molecules with ALMA at Planet-Forming Scales (MAPS) (Öberg et al., 2021) and the Disk Substructure at High Angular Resolution Project (DHSARP) (Andrews et al., 2018) have revealed a rich variety of features which challenge our traditional understanding of PPD evolution and planet formation Boehler et al., 2018; Cridland et al., 2019; Dipierro et al., 2015; Dong et al., 2018; D. H. Forgan, Ilee, and Meru, 2018; Hall et al., 2018; Meru et al., 2017; Okuzumi et al., 2016; Pinte et al., 2016, 2018, 2019; Shen et al., 2020; Teague et al., 2018. An essential tool is integral field spectroscopy (IFS). These observations measure the spectrum at each spatial pixel to create a data cube across wavelengths. This allows the quantification of motion within a disk, kinematic analysis (Pinte et al., 2023), can identify bodies and physical processes (Fedele et al., 2023; Hall et al., 2020; Pinte et al., 2018, 2019, 2023; Teague et al., 2018; Terry et al., 2023).

At the same time, advanced magnetohydrodynamics (MHD) codes have taken advantage of increased computational power to perform high-fidelity simulations with unprecedented resolution and physics (Benítez-Llambay & Masset, 2016; Mignone et al., 2007, 2012; Price et al., 2018; Stone et al., 2020). These simulations provide crucial information for the interpretation of observations (Dipierro et al., 2015; Hall et al., 2020; Meru et al., 2017; Pinilla et al., 2012; Teague et al., 2018).

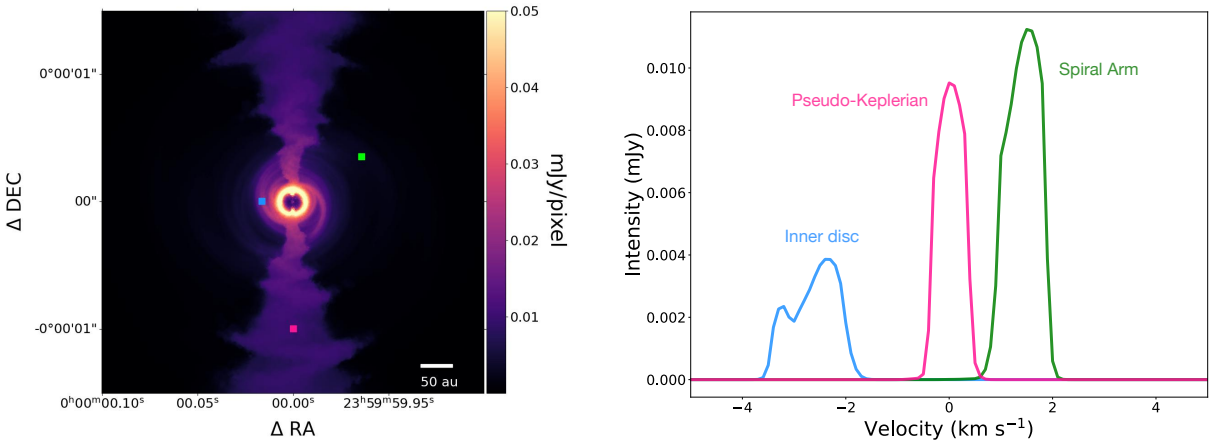

Figure 5.1: Schematic of the difference between computer vision inputs (left) and sequence-to-sequence inputs (right) for a self-gravitating disk. The left shows a single velocity channel image. The right shows the spectra corresponding to the highlighted locations. Pink: Keplerian region. Green: spiral arm. Blue: inner disk.

Separately, the proliferation of machine learning (ML) and artificial intelligence (AI) in astronomy (Alexander et al., 2020; Auddy et al., 2021; Ćiprijanović et al., 2021, 2023; Jo & Kim, 2019; Lochner et al., 2016; Möller & de Boissière, 2020; Terry, Hall, Abreau, & Gleyzer, 2022; Terry et al., 2023; Vilalta et al., 2019; S. Zhang et al., 2022) has led to an increased understanding of the power and potential of ML.

Most previous ML applications in PPD astronomy have relied on computer vision (CV). This is the class of algorithms for image analysis (Voulodimos et al., 2018), e.g., a convolutional neural network (CNN). Astronomical observations are often image-based. For example, images centered on a single wavelength, such as the continuum, are essentially grayscale images. An IFS data cube is simply an image with an arbitrary number of input channels. This makes CV a natural choice for observational analysis.

However, these traditional computer vision methods have limitations. They employ feature engineering to handcraft algorithms that extract features like edges, corners, blobs, and textures. Traditional methods have a localized focus that analyzes small sections of the image independently to identify specific features. CV is often specialized, highly tailored to perform specific sub-tasks like object detection, image segmentation, and classification.

Transformer-based image analysis, on the other hand, is more generalizable. Transformers rely on the self-attention mechanism to analyze images as a se-

quence of patches and can identify relationships between different parts of the image globally, even if patches are distant from each other. The transformer model can learn to identify novel, heretofore unseen features that differ from the Keplerian images in the training data. This eliminates the need for manual feature engineering. Transformers can adapt flexibly to various computer vision tasks through fine-tuning and adapt to new target domain through "domain adaptation" that doesn't require fundamentally different algorithms for each task.

However, CV has some significant drawbacks. One is the sheer amount of data necessary to create robust models. For example, Terry, Hall, Abreau, and Gleyzer, 2022 simulated 1,000 protoplanetary disks. Effective, general ML models require diverse training sets, and this is typically accomplished through generating synthetic data through simulations (Auddy et al., 2021; Jo & Kim, 2019; Terry, Hall, Abreau, & Gleyzer, 2022; S. Zhang et al., 2022). Synthetic training sets allow an arbitrary amount of data with relevant values (e.g., mass, number of planets, etc.) known beforehand. This combination allows a dataset that is sufficient to train models for a variety of tasks that would be impossible if trained on real (not synthetic) observations.

Creating a diverse set of simulations can be time-consuming. Hundreds or thousands of simulations may be necessary to fully capture the diversity of expected physical systems (see, e.g., Terry, Hall, Abreau, and Gleyzer, 2022). Simply generating the training data can take months.

Another complication with CV models is that they require a fixed input architecture that generally has a certain number of channels, N, and height and width, H and W, respectively. For example, a grayscale image has dimensions $(1, H, W)$ whereas an RGB image has dimensions $(3, H, W)$. The input dimension is an intrinsic, unchangeable aspect of the model architecture. Unfortunately for observational astronomy, observations do not all have the same dimensions or number of spectral channels. There is no guarantee that two observations will have the same number of pixels covering the area of interest.Even worse, for IFS data, there is no guarantee that there will even be the same number of channels between observations.

This necessitates data augmentation for model training and deployment, such as trimming or padding the number of pixels or channels, which can hurt performance by introducing superfluous information or obscuring important information.

I propose a different paradigm for ML analyses of IFS data. Rather than relying on CV algorithms, I use sequential models, such as recurrent neural networks (RNN) (Tsoi, 1997), designed to work effectively on input data with

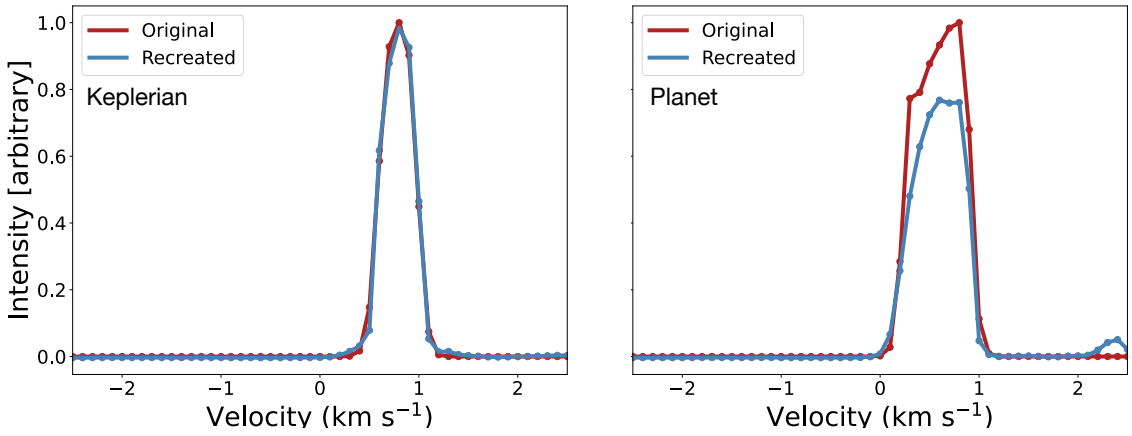

Figure 5.2: Example outputs for a Keplerian disk (left) and the region surrounding a planet (right). Red: input spectrum. Blue: output spectrum.

an inherent order. They are the natural choice for problems such as natural language processing and time-series data. While they perform typical ML tasks, e.g., classification and regression, they can also output entire sequences, such as language translations. These models are known as "sequence-to-sequence" (Seq2Seq) models. I focus on Seq2Seq models in this specific application, but sequential models can, in general, be used for a variety of tasks, such as classification or regression.

A particularly important innovation is the transformer architecture (Vaswani et al., 2017). Designed to improve upon traditional sequential methods, it has quickly become a cornerstone of AI and forms the basis of GPT (Bubeck et al., 2023). Transformers use multi-head self-attention to increase performance and parallelization when compared to previous architectures. Their unprecedented power for analyzing sequences has made them an indispensable tool for language modeling, signal processing, and even image analysis (Han et al., 2023).

I focus on PPD anomaly detection using IFS observations. AI-driven anomaly detection highlights atypical dataset features (Chalapathy & Chawla, 2019). This method is commonly used method for noisy datasets such as medical image analysis (van Hespen et al., 2021). It has already found use in astronomy as well. For example Alexander et al., 2020 showed that unsupervised CV-based anomaly detection can effectively highlight substructure in observations of strong lensing phenomena. When trained on only non-anomalous data, in-

ference on anomalous data reveals increased errors in regions with substructure. Mapping the errors across the observed field reveals the substructures of interest.

While this method is computational, its aim is qualitatively similar to that of much kinematic analysis: identify regions with features that don't conform to expectations. Human-driven PPD anomaly detection has identified planets within IFS datasets (Fedele et al., 2023; Pinte et al., 2018, 2019; Teague et al., 2018). Recent works have shown that basic CV models can also detect rudimentary PPD anomalies (Terry, Hall, Abreau, & Gleyzer, 2022; Terry et al., 2023).

Despite previous successes, both human-driven and conventional CV anomaly detection methods suffer from serious drawbacks. Human-driven methods are time consuming, and subtle signatures are easily missed. CV methods have been shown to identify previously unreported anomalies (Terry et al., 2023), but these models rely on a rigid, inflexible architecture described previously and were not explicitly trained to detect anomalies.

Sequential models offer a much simpler route towards anomaly detection. Rather than using thousands of simulations for training (Auddy et al., 2021; Terry, Hall, Abreau, & Gleyzer, 2022; S. Zhang et al., 2022), I train this model on only four different synthetic IFS observations. Since highlighting non-Keplerian regions is a major aim of kinematic analysis, all four disk simulations are purely Keplerian. The variation in the shape of purely Keplerian spectra is insignificant between disk parameters, so the main point of multiple simulations is to familiarize the sequential model with a broad range of system and observational variations, such as inclination and resolution. Model performance did not improve with a larger dataset.

These simulations can be done in a day, so the data generation time is trivially short. Even with this relatively tiny dataset, the models have access to more training data points. In CV, a single simulation may yield a few synthetic observations, each of which only counts as one data point for the purposes of training. In the sequential paradigm, a single $(N \times H \times W)$ data cube is now $H \times W$ spectra of length N. Keplerian spectra have little variation other than the peak frequency, so this is more than adequate to train a model to understand Keplerian spectra.

I propose to replace CV algorithms with sequential models, including transformers, for the application of AI to IFS observations, specifically anomaly detection (Chalapathy & Chawla, 2019) in this case. Rather than considering an IFS data cube to be a series of images, viewing them as a matrix of sequences unlocks the potential to use sequential models. This is shown schematically in

Figure 5.1. Given their suitability for a wide variety of datasets and problems, transformers are a natural choice for this task and come with significant benefits over traditional CV methods. I demonstrate they can effectively analyze observations with arbitrary input dimensions after being trained on only a handful of simulations. The results showcase the significant advantages of transformer models for spectral processing using protoplanetary disks as a demonstration case. These models offer faster training, robustness, generalizability, and flexibility compared to legacy computer vision approaches. With the ubiquity of IFS data across astronomical subfields (e.g. Seitenzahl et al., 2018), this work paves the way for groundbreaking discoveries in spectral analysis for the broader astronomy community.

5.1 Methods

5.1.1 Simulations

The training data consists of several synthetic IFS observations of disks simulated with smoothed-particle hydrodynamics (SPH). Each cube is of dimensions $(N \times H \times W)$, where N is the length of the spectra and H and W are the spatial height and width of the observations, respectively. A single cube therefore has $H \times W$ spectra that can be used for training. I train on only 4 simulations and still have a larger training set containing 86,4000 unique spectra.

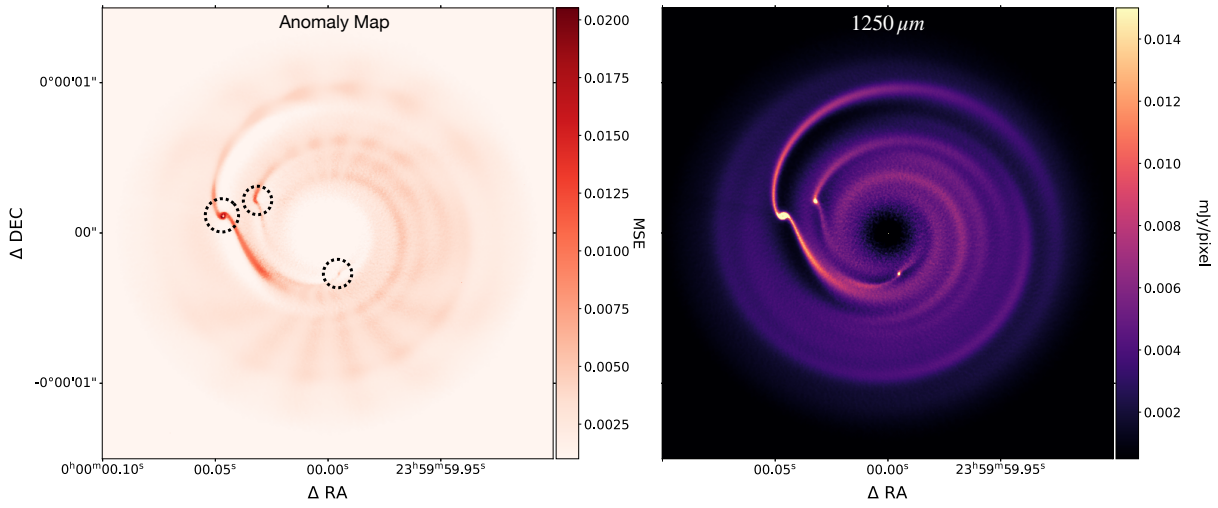

Figure 5.3: Mean-squared error map for a disk containing planets (left) and the corresponding 1250 μm continuum image (right). The three planets and spiral wakes are circled in black.

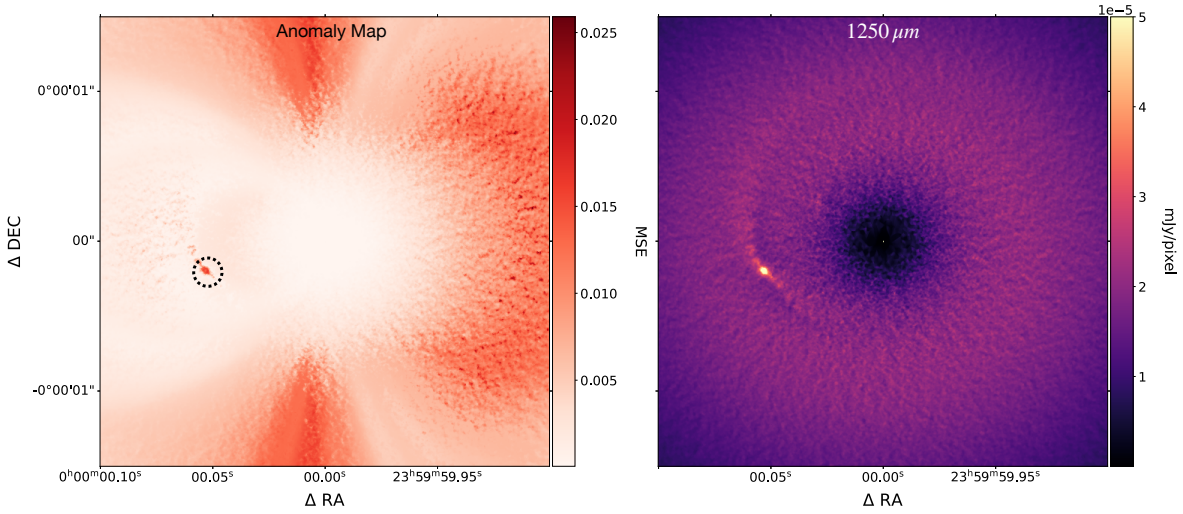

Figure 5.4: Mean-squared error map for a disk containing planets (left) and the corresponding 1250 μm continuum image (right). The planet and spiral wake is circled in black.

Smoothed-Particle Hydrodynamics

As discussed above, the paradigm focusing on spectra, instead of images, needs much less data. Only 4 training simulations were needed in total. `PHANTOM`, an open-source SPH code was used for the simulations (Price et al., 2018). Each simulation used 10^6 particles, following the procedure outlined in Terry, Hall, Abreau, and Gleyzer, 2022.

Training data generation for anomaly detection is trivial. The entire point is to locate non-Keplerian regions, and this is achieved by training the model to recreate Keplerian spectra. Therefore, only Keplerian simulations are used for training. These simulations take less than a day.

The testing data consists of disks with planets, and disks undergoing GI.

Table 5.1 gives information regarding all training and testing data.

Radiative Transfer

The SPH simulations are passed through `MCFOST`, an open-source radiative transfer code (Pinte et al., 2006, 2009), to create synthetic line emission observations of ^{13}CO $J = 3 \rightarrow 2$ in the form of IFS data cubes. `MCFOST` uses Voronoi tesselation of the simulation's particles to obtain the density structure. Parameters are similar to Terry, Hall, Abreau, and Gleyzer, 2022 and Terry, Hall, Longarini, et al., 2022, but I highlight some important aspects here. 10^8 photon packets were used. Dust was assumed to be a mixture of carbon and

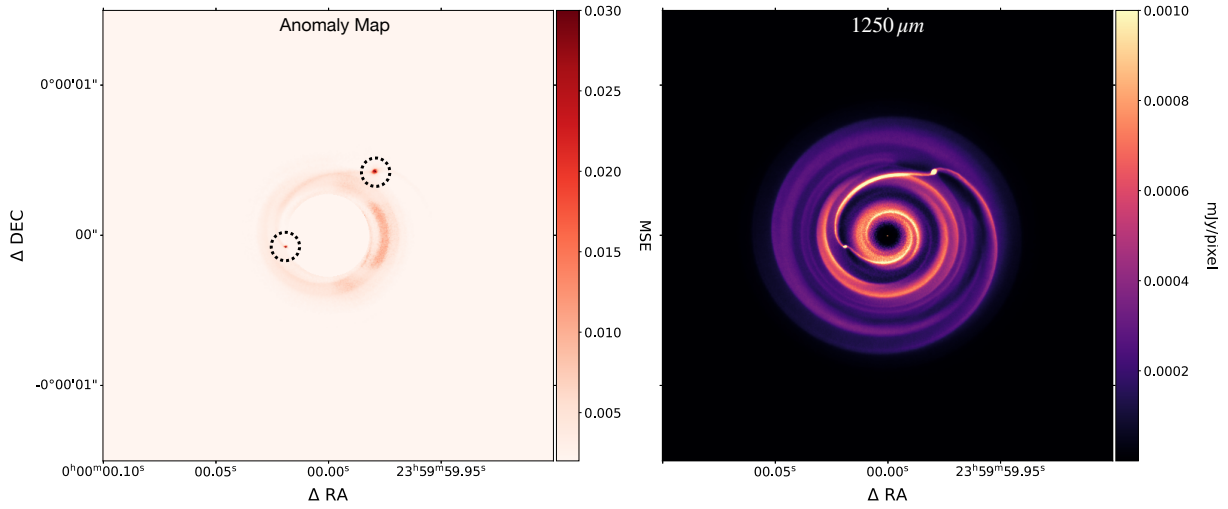

Figure 5.5: Mean-squared error map for a disk containing planets (left) and the corresponding 1250 μm continuum image (right). The two planets are circled in black.

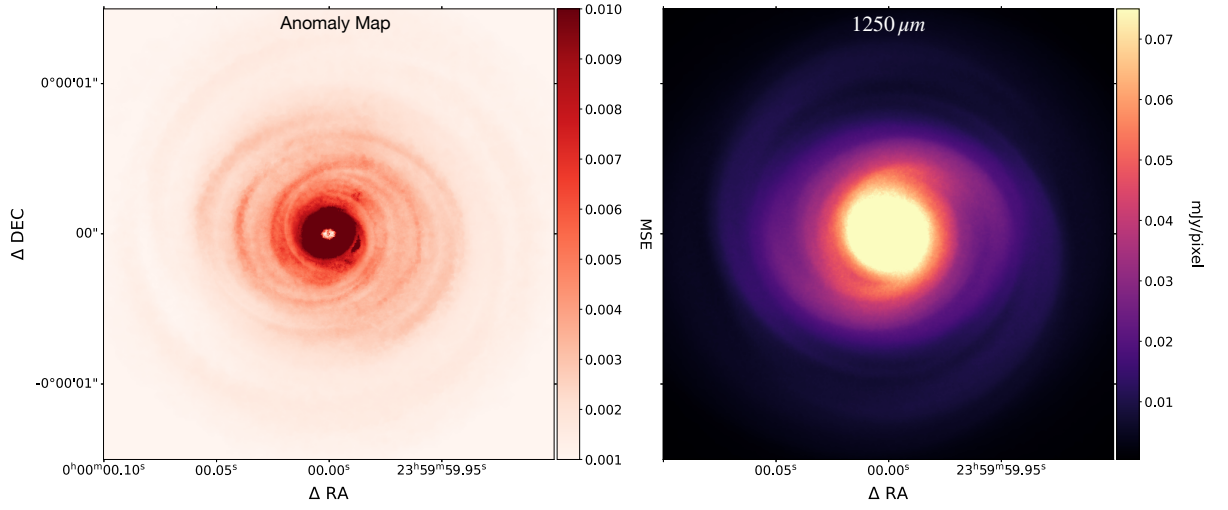

Figure 5.6: Mean-squared error map for a self-gravitating disk (left) and the corresponding 1250 μm continuum image (right).

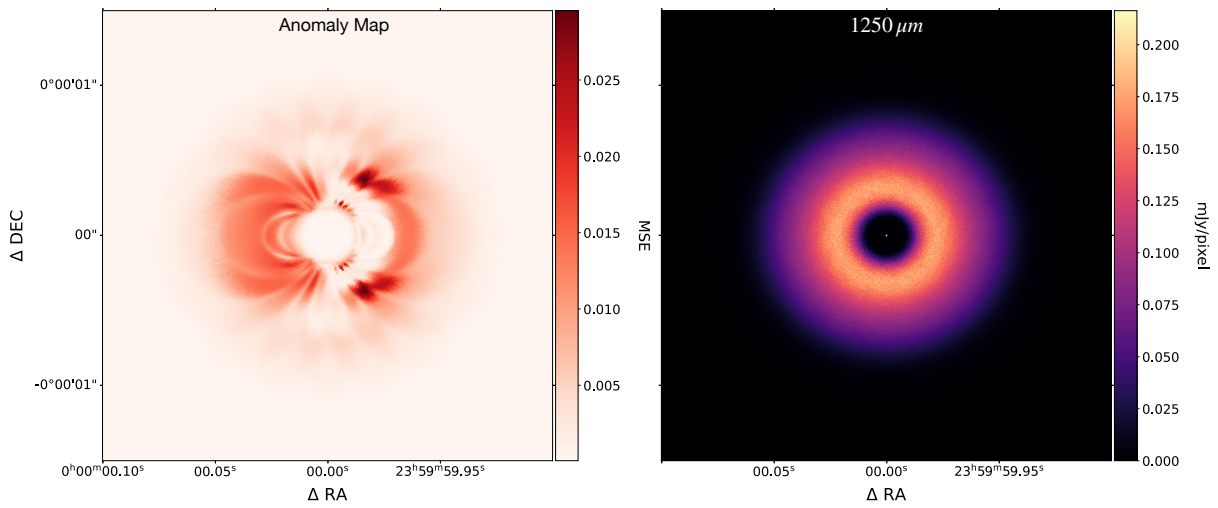

Figure 5.7: Mean-squared error map for a Keplerian disk (left) and the corresponding 1250 μm continuum image (right).

Table 5.1: Simulation setup information.

Setup	Environment	Planets	$M_{disk}/M_\star$ (%)	Resolution (km/s)
Train 1	Keplerian	0	0.3	0.2
Train 2	Keplerian	0	0.9	0.1
Train 3	Keplerian	0	0.4	0.1
Train 4	Keplerian	0	0.4	0.1
Test 1	Planet	3	0.7	0.2
Test 2	Planet	4	0.4	0.133
Test 3	Planet	2	0.9	0.1
Test 4	GI	0	50	0.1

silicates (Draine & Lee, 1984) with a dust-to-gas ratio of 0.01. The ^{13}CO-to-H_2 ratio was assumed to be 7×10^{-7} (Hall et al., 2020; Terry, Hall, Longarini, et al., 2022).

The result of the radiative transfer is an IFS data cube of $(N \times H \times W)$, i.e., spectra of length N for all $H \times W$ pixels. The data was continuum-subtracted using the $|\delta v| = 10\,km/s$ channels and only included channels $-3\,km/s \leq \delta v \leq 3\,km/s$. The resulting cube was normalized such that the minimum value is zero and the maximum value is one. An important benefit of the sequential architectures is that the input spectrum does not need to have a predefined length (i.e., N is a variable). To take advantage of this and help the models generalize, I create synthetic observations with several different spectra lengths (between 51 and 101 channels), corresponding to different spectral resolutions.

5.1.2 Anomaly Detection

I train a transformer to recreate an input Keplerian spectrum. The input to the model is the spectrum, and the output is the attempted recreation of the spectrum. The spectrum is first "encoded" into a lower dimensional space then "decoded" back to the original spectrum's dimension ("encoder-decoder model"). Mean-squared error loss was minimized using `WANDB` for a Bayesian hyperparameter sweep (Biewald, 2020). The final transformer model lacks inherent knowledge of non-Keplerian dynamics, having solely been trained on data pertaining to Keplerian disks. The trained transformer model lacks the ability to directly infer non-Keplerian motion and will therefore perform worse when recreating non-Keplerian spectra. Mapping the reconstruction error for each pixel gives an anomaly map that should highlight regions of interest, e.g., non-Keplerian regions.

Not only is this method easy to prepare data for, but it is also completely unsupervised: there are no explicit information given about the system during training. The only goal is to recreate the input. This attractive quality makes it agnostic towards any imposed model or system: an anomaly will be highlighted regardless of its origin. Any non-Keplerian anomaly can be detected whether it's caused by a planet, gravitational instability (GI), vertical shear instability, or any other mechanism (Barraza-Alfaro et al., 2021; Bollati et al., 2021; Disk Dynamics Collaboration et al., 2020; Hall et al., 2020; Pinte et al., 2023).

5.1.3 Domain Adaptation

The previous method is applied to raw simulation outputs, i.e., no observational effects are added. While this offers a proof of concept for sequence-to-sequence

models, it is not practical for deployment on observations. It is crucial to overcome this barrier.

Domain adaptation (DA) is a machine learning method designed to transfer models between datasets to improve performance on a target domain with limited data by leveraging knowledge from a different, but related, source domain with abundant labeled data (Ben-David et al., 2010; Ćiprijanović et al., 2021; Tzeng et al., 2017) . DA has already found success in astronomy (Alexander et al., 2023; Chalapathy & Chawla, 2019; Ćiprijanović et al., 2023). It has been used to apply models to different telescopes and, most relevant to this project, apply models trained on simulations to observational datasets.

The basic idea behind DA is to train a model to perform a task using a given set of data (the "source" dataset). This model is then retrained to perform the same task on a different set of data ("target" data). The model is able to use the information learned from the original training to precisely adapt itself to perform well on the target dataset. In the context of simulation-to-observation DA, I take the source dataset to be the raw simulation outputs. I then add noise and convolve the data, both spatially and spectrally, in accordance with Terry, Hall, Abreau, and Gleyzer, 2022 to produce the dirty "target" dataset.

I use adversarial discriminative domain adaptation (ADDA) (Tzeng et al., 2017) for this purpose. ADDA is categorized as an unsupervised domain adaptation technique because it doesn't require any labeled data in the target domain. ADDA is an adversarial method that encourages the retrained network to encode the target dataset into the same latent space as the source dataset. In this way, the model is able to learn how to represent the target data similarly to the source data then feed that representation through the same decoder used with the source dataset. The model therefore learns how to encode the "hard" target problem in such a way that decoding it is equivalent to decoding the "easy" source data. This has been demonstrated to outperform models that have simply been trained on the target dataset (Alexander et al., 2023).

I can represent ADDA as follows: take x_s to be the raw simulation spectrum, and $\hat{x}_s$ is the attempted recreation. x_t and $\hat{x}_t$ are corresponding target input and recreated spectra. The original encoder-decoder model is therefore $\hat{x}_s = D_s(E_s(x_s))$, where the source encoder and decoder are represented as E_s and D_s, respectively. A new encoder for the target data, E_t is trained along with a binary classifier, C, that determines whether an encoded spectrum belongs to the source or target set. C and E_t are adversarially trained such that $E_s(x_s)$ is indistinguishable from $E_t(x_t)$ (as judged by C). The target data is then used to reproduce a source spectrum by applying $\hat{x}_t = D_s(E_t(x_t))$.

5.2 Results and Discussion

Final model hyperparameters (i.e., the parameters describing the model's architecture) after the WANDB sweep are shown in Table 5.2. The resulting model had a test loss over all Keplerian simulations of 0.012, indicating the average spectrum is recreated to 1.2% accuracy.. Figure 5.2 gives an example of the model's predictions for a Keplerian disk (left) and the region surrounding a planet (right). As can be seen, the shoulder associated with the planet's perturbation is not captured by the model, leading to a larger error than the Keplerian region (left). This shoulder is the source of the kink that would be seen in a velocity channel map since it is an increase in emission away from the Keplerian peak.

The slight bump after 2 km/s in the right panel of Figure 5.2 is not uncommon in erroneous predictions of non-Keplerian regions. It is difficult to explain this fully. A possible reason for this behavior is the model optimizing for the area under the curve when compared to the Keplerian case. The increased area due to the non-Keplerian shoulder may be compensated for by simply adding flux to high-velocity channels. This hypothesis is difficult to verify, and not all models exhibited this behavior.

Example maps of these errors and corresponding 1250 μm continuum observations are shown in Figure 5.3, Figure 5.4, Figure 5.5, Figure 5.6, and Figure 5.7. Substructures, including rings, gaps, spiral wakes, and spiral arms, are highlighted in all non-Keplerian cases (Figure 5.3, Figure 5.4, and Figure 5.5, Figure 5.6).

Figure 5.3, Figure 5.4, and Figure 5.5, are the result of applying the model to disks that include planets. As is clear, the models strongly highlight the planets and their spiral wakes. Excitingly, the model is able to capture multiple planets. This is an improvement over the results of Terry, Hall, Abreau, and Gleyzer, 2022. That work only predicted if a disk had at least one planet, whereas this clearly shows that there are multiple planets. These maps also localize the planets much more effectively and precisely than the activation method used in Terry, Hall, Abreau, and Gleyzer, 2022. In a sense, this method is a powerful combination that can regress on the number of planets and segment the image such that planets are highlighted. It is further encouraging that a planet is clearly identified in Figure 5.4. Unlike the other disks, this system does not have clear rings and gaps, yet the planet is still clearly seen in the anomaly map.

Figure 5.6 shows the same results for a self-gravitating disk. In the continuum, the spiral arms are not apparent due to the overwhelming brightness of the inner disk. The error map, on the other hand, precisely shows the spiral structure. This encouraging result demonstrates that substructures may be

Table 5.2: Final hyperparameters from WANDB sweep.

Parameter	Value
Dropout	0.169
Initial learning rate	0.00044
Learning rate scheduler	Step decrease
Learning rate decrease factor	0.5
Learning rate step frequency	5 epochs
Number of attention heads	16
Number of transformer layers	5
Transformer hidden dimension	16
Total trainable parameters	44,421

recovered using anomaly detection despite not clearly being visible in other observations.

Figure 5.7 shows the error map and continuum for a disk without a planet or self-gravity, i.e., a purely Keplerian disk. As expected, there is little coherent structure or bright object in the map, indicating a lack of anomalies.

An additional benefit of the sequential approach is the simplicity of the models themselves. The final model only has 44,421 trainable parameters (Table 5.2). In contrast, the models used in Terry, Hall, Abreau, and Gleyzer, 2022 had between 20 million and 114 million trainable parameters. This significantly speeds up training to the extent that typical consumer resources can be used.

The dataset used also offers improvements over previous CV approaches. The massive increase in available training data points per simulation (600^2 per simulation step), output dramatically decreases the number of necessary simulations. Terry, Hall, Abreau, and Gleyzer, 2022 used 1,000 simulations to train their models, resulting in 1,680 unique training points (excluding validation and testing data).

5.2.1 Domain Adaptation Application

I apply the ADDA model to observations from the DSHARP catalogue (Andrews et al., 2018; J. Huang, Andrews, Dullemond, et al., 2018). As Figure 5.8 shows, the improvement is qualitative. The left shows the map from the models that were trained on only the source dataset. The center panel shows the same map after retraining with ADDA. The right shows an ALMA ^{12}CO: $J = 2 \rightarrow 1$ observation of the disk. Terry et al., 2023 identified a planet in the circled region. ADDA therefore takes a model that produces a map with little

coherent structure (left of Figure 5.8) and produces a new model that strongly identifies that channel and region in containing a planet. This demonstrates the viability of using ADDA to allow the successful application of sequential models in observational data.

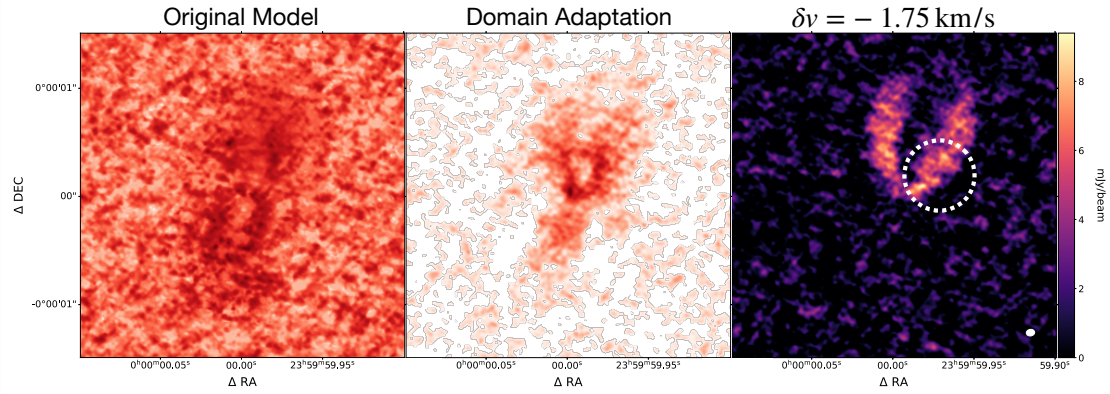

Figure 5.8: Application of ADDA to HD 142666. Left: map made using model without domain adaptation. Center: map made using ADDA model. Right: ^{12}CO: $J = 2 \rightarrow 1$ channel with a previously located planet's location circled.

5.3 Conclusions

These results demonstrate the clear ability of sequence-to-sequence models to precisely identify anomalous regions in IFS observations of protoplanetary disks. These results are a strong endorsement of this paradigm over the traditional computer vision methods. Reasons include:

1. The sequence-to-sequence method's ability to localize an arbitrary number of anomalies without regard to their source facilitates a thorough analysis of data. With planets, this allows the identification of any number of planets within the disk, an improvement on Terry, Hall, Abreau, and Gleyzer, 2022. Subtle signatures, such as spiral substructures, that are difficult to see in observations can be clearly seen.

2. Orders of magnitude fewer training simulations are necessary. Whereas Terry, Hall, Abreau, and Gleyzer, 2022 used 1,000 different simulations, this work only relies on three Keplerian simulations. This massive overhead reduction allows the simple and quick generation of the training dataset.

3. Whereas computer vision models require a fixed input dimension, sequence-to-sequence models can handle data of any shape. No padding or trim-

ming is necessary, allowing any observation to be directly fed into the models.

While this work is focused on anomaly detection in protoplanetary disks, the general framework transcends this particular area and problem. Sequential models can be used on any IFS data, including protoplanetary disks, galaxies, supernova remnants, and so on. Any field with IFS data can benefit from applying sequential models as a supplement to or replacement for traditional computer vision algorithms. With this work, I aim to encourage the application of these methods across disciplines and problems.

Chapter 6

Conclusion

This work demonstrates the crucial role of instabilities, perturbations, and substructures in protoplanetary disk evolution and planet formation. Simulations, kinematic analysis, observational analysis, and machine learning give insight into the complex physics, causes, and long-term implications.

The gravitational instability can qualitatively change disk evolution by creating spiral substructures and transporting angular momentum. These results indicate that its hallmark signature, the GI wiggle, contains information about the disk's mass, the disk's cooling, and the central star's properties.

Embedded protoplanets cause local deviations from Keplerian motion, which may be detected in observations. This work has shown the machine learning can effectively identify these perturbations, leading to the discovery of a previously unreported exoplanet in the disk HD 142666. Additionally, machine learning algorithms created for sequence analysis can also highlight any observational signature of an instability, perturbation, or substructure, regardless of its source. This novel method opens a new avenue for observational analysis across subfields in astronomy.

This dissertation lays the groundwork for further exploration into the nuances of gravitational instability and the application of machine learning in protoplanetary disk astronomy. Its results are timely. As the next generation of telescopes become operational, the breadth and quality of observational data will dramatically increase. This will provide new opportunities for more precise tests of predictions and the identification of important objects or processes.

Given this approaching leap, this work should be expanded upon in the interim. Of particular interest is the interplay of magnetism and gravitational instability. The role of magnetic fields in protoplanetary disks and planet formation in general is far from settled. A full understanding of their role can only come if they are included in the consideration of other essential processes, such

as the gravitational instability. A particularly understudied area is the inclusion of non-ideal magnetic effects, which likely plays a key role in the overall effects of magnetism. Understanding the physical and observational manifestation of this more realistic treatment may be of critical importance for the overall study of planet formation.

Separately, the new observations and explosion of machine learning capabilities presents an opportunity for the development and deployment of state-of-the-art methods. These have the potential to provide supplementary, new, or superior analysis techniques that will more thoroughly exploit the wealth of data in an observation. The future of the rapidly developing field of machine learning is difficult to predict, but it is almost certain that it will hold an important place in protoplanetary disk astronomy.

Protoplanetary disk astronomy seems to be in a continually improving golden age. The future of the field is promising. I hope that this work has made it a bit brighter.

Appendix A

Opacity Calculations

The Rosseland mean opacities used are calculated from the power-law fit (Equation 66) from Lin and Papaloizou, 1985

$$\kappa = \left[\left(\kappa_1^{-2} + \kappa_2^{-2} \right)^{-2} + \left(\kappa_3 / (1 + 10^{22} T^{-10}) \right)^4 \right]^{1/4}, \tag{A.1}$$

where

$$\kappa_1 = 2 \times 10^{-4}\, T^2 \qquad T < 170\mathrm{K}, \tag{A.2}$$

$$\kappa_2 = 2 \times 10^{16}\, T^{-7} \qquad 170\mathrm{K} < T < 210\mathrm{K}, \tag{A.3}$$

$$\kappa_3 = 5 \times 10^{-3}\, T \qquad 210\mathrm{K} < T. \tag{A.4}$$

Bibliography

Adams, F. C., Ruden, S. P., & Shu, F. H. (1989). Eccentric Gravitational Instabilities in Nearly Keplerian Disks. *ApJ*, *347*, 959. https://doi.org/10.1086/168187

Agertz, O., Moore, B., Stadel, J., Potter, D., Miniati, F., Read, J., Mayer, L., Gawryszczak, A., Kravtsov, A., Nordlund, Å., Pearce, F., Quilis, V., Rudd, D., Springel, V., Stone, J., Tasker, E., Teyssier, R., Wadsley, J., & Walder, R. (2007). Fundamental differences between SPH and grid methods. *MNRAS*, *380*(3), 963–978. https://doi.org/10.1111/j.1365-2966.2007.12183.x

Alexander, S., Gleyzer, S., Parul, H., Reddy, P., Tidball, M., & Toomey, M. W. (2023). Domain Adaptation for Simulation-based Dark Matter Searches with Strong Gravitational Lensing. *ApJ*, *954*(1), Article 28, 28. https://doi.org/10.3847/1538-4357/acdfc7

Alexander, S., Gleyzer, S., Parul, H., Reddy, P., Toomey, M. W., Usai, E., & Von Klar, R. (2020). Decoding Dark Matter Substructure without Supervision. *arXiv e-prints*, Article arXiv:2008.12731, arXiv:2008.12731. https://doi.org/10.48550/arXiv.2008.12731

ALMA Partnership, Brogan, C. L., Pérez, L. M., Hunter, T. R., Dent, W. R. F., Hales, A. S., Hills, R. E., Corder, S., Fomalont, E. B., Vlahakis, C., Asaki, Y., Barkats, D., Hirota, A., Hodge, J. A., Impellizzeri, C. M. V., Kneissl, R., Liuzzo, E., Lucas, R., Marcelino, N., … Tatematsu, K. (2015). The 2014 ALMA Long Baseline Campaign: First Results from High Angular Resolution Observations toward the HL Tau Region. *ApJL*, *808*(1), Article L3, L3. https://doi.org/10.1088/2041-8205/808/1/L3

Andrews, S. M., Huang, J., Pérez, L. M., Isella, A., Dullemond, C. P., Kurtovic, N. T., Guzmán, V. V., Carpenter, J. M., Wilner, D. J., Zhang, S., Zhu, Z., Birnstiel, T., Bai, X.-N., Benisty, M., Hughes, A. M., Öberg, K. I., & Ricci, L. (2018). The Disk Substructures at High Angular Resolution Project (DSHARP). I. Motivation, Sample, Calibration, and

Overview. *ApJL*, *869*(2), Article L41, L41. https://doi.org/10.3847/2041-8213/aaf741

Andrews, S. M., Wilner, D. J., Hughes, A. M., Qi, C., & Dullemond, C. P. (2009). PROTOPLANETARY DISK STRUCTURES IN OPHIUCHUS. *The Astrophysical Journal*, *700*(2), 1502–1523. https://doi.org/10.1088/0004-637x/700/2/1502

Andrews, S. M., Wilner, D. J., Hughes, A. M., Qi, C., & Dullemond, C. P. (2010). Protoplanetary Disk Structures in Ophiuchus. II. Extension to Fainter Sources. *ApJ*, *723*(2), 1241–1254. https://doi.org/10.1088/0004-637X/723/2/1241

Ansdell, M., Williams, J. P., Trapman, L., van Terwisga, S. E., Facchini, S., Manara, C. F., van der Marel, N., Miotello, A., Tazzari, M., Hogerheijde, M., Guidi, G., Testi, L., & van Dishoeck, E. F. (2018). ALMA Survey of Lupus Protoplanetary Disks. II. Gas Disk Radii. *ApJ*, *859*(1), Article 21, 21. https://doi.org/10.3847/1538-4357/aab890

Ansdell, M., Williams, J. P., van der Marel, N., Carpenter, J. M., Guidi, G., Hogerheijde, M., Mathews, G. S., Manara, C. F., Miotello, A., Natta, A., Oliveira, I., Tazzari, M., Testi, L., van Dishoeck, E. F., & van Terwisga, S. E. (2016). ALMA Survey of Lupus Protoplanetary Disks. I. Dust and Gas Masses. *ApJ*, *828*(1), Article 46, 46. https://doi.org/10.3847/0004-637X/828/1/46

Armitage, P. J. (2022). Lecture notes on accretion disk physics. *arXiv e-prints*, Article arXiv:2201.07262, arXiv:2201.07262. https://doi.org/10.48550/arXiv.2201.07262

Armitage, P. J., Livio, M., & Pringle, J. E. (2001). Episodic accretion in magnetically layered protoplanetary discs. *MNRAS*, *324*(3), 705–711. https://doi.org/10.1046/j.1365-8711.2001.04356.x

Auddy, S., Dey, R., Lin, M.-K., Carrera, D., & Simon, J. B. (2022). Using Bayesian Deep Learning to Infer Planet Mass from Gaps in Protoplanetary Disks. *ApJ*, *936*(1), Article 93, 93. https://doi.org/10.3847/1538-4357/ac7a3c

Auddy, S., Dey, R., Lin, M.-K., & Hall, C. (2021). DPNNet-2.0. I. Finding Hidden Planets from Simulated Images of Protoplanetary Disk Gaps. *ApJ*, *920*(1), Article 3, 3. https://doi.org/10.3847/1538-4357/ac1518

Balbus, S. A., & Hawley, J. F. (1991). A Powerful Local Shear Instability in Weakly Magnetized Disks. I. Linear Analysis. *ApJ*, *376*, 214. https://doi.org/10.1086/170270

Ballabio, G., Dipierro, G., Veronesi, B., Lodato, G., Hutchison, M., Laibe, G., & Price, D. J. (2018). Enforcing dust mass conservation in 3D sim-

ulations of tightly coupled grains with the PHANTOM SPH code. *MNRAS*, *477*(2), 2766–2771. https://doi.org/10.1093/mnras/sty642

Barraza-Alfaro, M., Flock, M., Marino, S., & Pérez, S. (2021). Observability of the vertical shear instability in protoplanetary disk CO kinematics. *A&A*, *653*, Article A113, A113. https://doi.org/10.1051/0004-6361/202140535

Bate, M. R., Bonnell, I. A., & Price, N. M. (1995). Modelling accretion in protobinary systems. *MNRAS*, *277*(2), 362–376. https://doi.org/10.1093/mnras/277.2.362

Beckwith, S. V. W., Sargent, A. I., Chini, R. S., & Guesten, R. (1990). A Survey for Circumstellar Disks around Young Stellar Objects. *AJ*, *99*, 924. https://doi.org/10.1086/115385

Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., & Vaughan, J. W. (2010). A theory of learning from different domains. *Machine Learning*, *79*(1), 151–175. https://doi.org/10.1007/s10994-009-5152-4

Benisty, M., Juhasz, A., Boccaletti, A., Avenhaus, H., Milli, J., Thalmann, C., Dominik, C., Pinilla, P., Buenzli, E., Pohl, A., Beuzit, J. .-., Birnstiel, T., de Boer, J., Bonnefoy, M., Chauvin, G., Christiaens, V., Garufi, A., Grady, C., Henning, T., … Pueyo, L. (2015). Asymmetric features in the protoplanetary disk MWC 758. *A&A*, *578*, Article L6, L6. https://doi.org/10.1051/0004-6361/201526011

Benisty, M., Bae, J., Facchini, S., Keppler, M., Teague, R., Isella, A., Kurtovic, N. T., Pérez, L. M., Sierra, A., Andrews, S. M., Carpenter, J., Czekala, I., Dominik, C., Henning, T., Menard, F., Pinilla, P., & Zurlo, A. (2021). A Circumplanetary Disk around PDS70c. *ApJL*, *916*(1), Article L2, L2. https://doi.org/10.3847/2041-8213/ac0f83

Benítez-Llambay, P., & Masset, F. S. (2016). FARGO3D: A New GPU-oriented MHD Code. *ApJs*, *223*(1), Article 11, 11. https://doi.org/10.3847/0067-0049/223/1/11

Bergin, E. A., & Williams, J. P. (2017). The Determination of Protoplanetary Disk Masses. In M. Pessah & O. Gressel (Eds.), *Formation, evolution, and dynamics of young solar systems* (p. 1, Vol. 445). Springer. https://doi.org/10.1007/978-3-319-60609-5_1

Biewald, L. (2020). Experiment tracking with weights and biases [Software available from wandb.com]. https://www.wandb.com/

Birnstiel, T., & Andrews, S. M. (2014). On the Outer Edges of Protoplanetary Dust Disks. *ApJ*, *780*(2), Article 153, 153. https://doi.org/10.1088/0004-637X/780/2/153

Boehler, Y., Ricci, L., Weaver, E., Isella, A., Benisty, M., Carpenter, J., Grady, C., Shen, B.-T., Tang, Y.-W., & Perez, L. (2018). The Complex Morphology of the Young Disk MWC 758: Spirals and Dust Clumps around a Large Cavity. *ApJ*, *853*(2), Article 162, 162. https://doi.org/10.3847/1538-4357/aaa19c

Bollati, F., Lodato, G., Price, D. J., & Pinte, C. (2021). The theory of kinks — I. A semi-analytic model of velocity perturbations due to planet-disc interaction. *MNRAS*. https://doi.org/10.1093/mnras/stab1145

Boss, A. P. (1997). Giant planet formation by gravitational instability. *Science*, *276*, 1836–1839. https://doi.org/10.1126/science.276.5320.1836

Boss, A. P. (1998). Formation of Extrasolar Giant Planets: Core Accretion or Disk Instability? *Earth Moon and Planets*, *81*(1), 19–26.

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Tulio Ribeiro, M., & Zhang, Y. (2023). Sparks of Artificial General Intelligence: Early experiments with GPT-4. *arXiv e-prints*, Article arXiv:2303.12712, arXiv:2303.12712. https://doi.org/10.48550/arXiv.2303.12712

Burn, R., Schlecker, M., Mordasini, C., Emsenhuber, A., Alibert, Y., Henning, T., Klahr, H., & Benz, W. (2021). The New Generation Planetary Population Synthesis (NGPPS). IV. Planetary systems around low-mass stars. *A&A*, *656*, Article A72, A72. https://doi.org/10.1051/0004-6361/202140390

Cadman, J., Hall, C., Rice, K., Harries, T. J., & Klaassen, P. D. (2020). The observational impact of dust trapping in self-gravitating discs. *MNRAS*, *498*(3), 4256–4271. https://doi.org/10.1093/mnras/staa2596

Cadman, J., Rice, K., & Hall, C. (2021). AB Aurigae: possible evidence of planet formation through the gravitational instability. *MNRAS*, *504*(2), 2877–2888. https://doi.org/10.1093/mnras/stab905

Cadman, J., Rice, K., Hall, C., Haworth, T. J., & Biller, B. (2020). Fragmentation favoured in discs around higher mass stars. *MNRAS*, *492*(4), 5041–5051. https://doi.org/10.1093/mnras/staa187

Chalapathy, R., & Chawla, S. (2019). Deep Learning for Anomaly Detection: A Survey. *arXiv e-prints*, Article arXiv:1901.03407, arXiv:1901.03407. https://doi.org/10.48550/arXiv.1901.03407

Chen, E., Yu, S.-Y., & Ho, L. C. (2021). Global Spiral Density Wave Modes in Protoplanetary Disks: Morphology of Spiral Arms. *ApJ*, *906*(1), Article 19, 19. https://doi.org/10.3847/1538-4357/abc7c5

Ćiprijanović, A., Kafkes, D., Downey, K., Jenkins, S., Perdue, G. N., Madireddy, S., Johnston, T., Snyder, G. F., & Nord, B. (2021). DeepMerge - II.

Building robust deep learning algorithms for merging galaxy identification across domains. *MNRAS*, *506*(1), 677–691. https://doi.org/10.1093/mnras/stab1677

Ćiprijanović, A., Lewis, A., Pedro, K., Madireddy, S., Nord, B., Perdue, G. N., & Wild, S. M. (2023). DeepAstroUDA: semi-supervised universal domain adaptation for cross-survey galaxy morphology classification and anomaly detection. *Machine Learning: Science and Technology*, *4*(2), Article 025013, 025013. https://doi.org/10.1088/2632-2153/acca5f

Ćiprijanović, A., Kafkes, D., Snyder, G., Sánchez, F. J., Perdue, G. N., Pedro, K., Nord, B., Madireddy, S., & Wild, S. M. (2022). DeepAdversaries: examining the robustness of deep learning models for galaxy morphology classification. *Machine Learning: Science and Technology*, *3*(3), Article 035007, 035007. https://doi.org/10.1088/2632-2153/ac7f1a

Cireşan, D. C., Meier, U., Masci, J., Gambardella, L. M., & Schmidhuber, J. (2011). High-Performance Neural Networks for Visual Object Classification. *arXiv e-prints*, Article arXiv:1102.0183, arXiv:1102.0183.

Clarke, C. J. (2009). Pseudo-viscous modelling of self-gravitating discs and the formation of low mass ratio binaries. *MNRAS*, *396*(2), 1066–1074. https://doi.org/10.1111/j.1365-2966.2009.14774.x

Cossins, P., Lodato, G., & Clarke, C. J. (2009). Characterizing the gravitational instability in cooling accretion discs. *MNRAS*, *393*(4), 1157–1173. https://doi.org/10.1111/j.1365-2966.2008.14275.x

Cossins, P., Lodato, G., & Testi, L. (2010). Resolved images of self-gravitating circumstellar discs with ALMA. *MNRAS*, *407*(1), 181–188. https://doi.org/10.1111/j.1365-2966.2010.16934.x

Cridland, A. J., Pudritz, R. E., & Alessi, M. (2019). Physics of planet trapping with applications to HL Tau. *MNRAS*, *484*(1), 345–363. https://doi.org/10.1093/mnras/stz008

de Val-Borro, M., Edgar, R. G., Artymowicz, P., Ciecielag, P., Cresswell, P., D'Angelo, G., Delgado-Donate, E. J., Dirksen, G., Fromang, S., Gawryszczak, A., Klahr, H., Kley, W., Lyra, W., Masset, F., Mellema, G., Nelson, R. P., Paardekooper, S. .-., Peplinski, A., Pierens, A., … Speith, R. (2006). A comparative study of disc-planet interaction. *MNRAS*, *370*(2), 529–558. https://doi.org/10.1111/j.1365-2966.2006.10488.x

Dipierro, G., Lodato, G., Testi, L., & de Gregorio Monsalvo, I. (2014). How to detect the signatures of self-gravitating circumstellar discs with the Atacama Large Millimeter/sub-millimeter Array. *MNRAS*, *444*(2), 1919–1929. https://doi.org/10.1093/mnras/stu1584

Dipierro, G., Price, D., Laibe, G., Hirsh, K., Cerioli, A., & Lodato, G. (2015). On planet formation in HL Tau. *MNRAS*, *453*(1), L73–L77. https://doi.org/10.1093/mnrasl/slv105

Disk Dynamics Collaboration, Armitage, P. J., Bae, J., Benisty, M., Bergin, E. A., Casassus, S., Czekala, I., Facchini, S., Fung, J., Hall, C., Ilee, J. D., Keppler, M., Kuznetsova, A., Le Gal, R., Loomis, R. A., Lyra, W., Manger, N., Perez, S., Pinte, C., … Zhang, K. (2020). Visualizing the Kinematics of Planet Formation. *arXiv e-prints*, Article arXiv:2009.04345, arXiv:2009.04345.

Dong, R., Hall, C., Rice, K., & Chiang, E. (2015). Spiral Arms in Gravitationally Unstable Protoplanetary Disks as Imaged in Scattered Light. *ApJL*, *812*(2), Article L32, L32. https://doi.org/10.1088/2041-8205/812/2/L32

Dong, R., Liu, S.-y., Eisner, J., Andrews, S., Fung, J., Zhu, Z., Chiang, E., Hashimoto, J., Liu, H. B., Casassus, S., Esposito, T., Hasegawa, Y., Muto, T., Pavlyuchenkov, Y., Wilner, D., Akiyama, E., Tamura, M., & Wisniewski, J. (2018). The Eccentric Cavity, Triple Rings, Two-armed Spirals, and Double Clumps of the MWC 758 Disk. *ApJ*, *860*(2), Article 124, 124. https://doi.org/10.3847/1538-4357/aac6cb

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Houlsby, N. (2020). An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *arXiv e-prints*, Article arXiv:2010.11929, arXiv:2010.11929.

Draine, B. T., & Lee, H. M. (1984). Optical Properties of Interstellar Graphite and Silicate Grains. *ApJ*, *285*, 89. https://doi.org/10.1086/162480

Drążkowska, J., & Alibert, Y. (2017). Planetesimal formation starts at the snow line. *A&A*, *608*, Article A92, A92. https://doi.org/10.1051/0004-6361/201731491

Drążkowska, J., Bitsch, B., Lambrechts, M., Mulders, G. D., Harsono, D., Vazan, A., Liu, B., Ormel, C. W., Kretke, K., & Morbidelli, A. (2023, July). Planet Formation Theory in the Era of ALMA and Kepler: from Pebbles to Exoplanets. In S. Inutsuka, Y. Aikawa, T. Muto, K. Tomida, & M. Tamura (Eds.), *Protostars and planets vii* (p. 717, Vol. 534). https://doi.org/10.48550/arXiv.2203.09759

Durisen, R. H., Boss, A. P., Mayer, L., Nelson, A. F., Quinn, T., & Rice, W. K. M. (2007, January). Gravitational Instabilities in Gaseous Protoplanetary Disks and Implications for Giant Planet Formation. In B. Reipurth, D. Jewitt, & K. Keil (Eds.), *Protostars and planets v* (p. 607).

Elfwing, S., Uchibe, E., & Doya, K. (2017). Sigmoid-Weighted Linear Units for Neural Network Function Approximation in Reinforcement Learning. *arXiv e-prints*, Article arXiv:1702.03118, arXiv:1702.03118.

Facchini, S., Juhász, A., & Lodato, G. (2018). Signatures of broken protoplanetary discs in scattered light and in sub-millimetre observations. *MNRAS*, *473*(4), 4459–4475. https://doi.org/10.1093/mnras/stx2523

Fedele, D., Bollati, F., & Lodato, G. (2023). Kinematics signature of a giant planet in the disk of AS 209. *A&A*, *672*, Article A125, A125. https://doi.org/10.1051/0004-6361/202244486

Forgan, D. H., Hall, C., Meru, F., & Rice, W. K. M. (2018). Towards a population synthesis model of self-gravitating disc fragmentation and tidal downsizing II: the effect of fragment-fragment interactions. *MNRAS*, *474*(4), 5036–5048. https://doi.org/10.1093/mnras/stx2870

Forgan, D., & Rice, K. (2013). Towards a population synthesis model of objects formed by self-gravitating disc fragmentation and tidal downsizing. *MNRAS*, *432*(4), 3168–3185. https://doi.org/10.1093/mnras/stt672

Forgan, D., Rice, K., Cossins, P., & Lodato, G. (2010). The nature of angular momentum transport in radiative self-gravitating protostellar discs. *Monthly Notices of the Royal Astronomical Society*, *410*(2), 994–1006. https://doi.org/10.1111/j.1365-2966.2010.17500.x

Forgan, D. H., Ilee, J. D., & Meru, F. (2018). Are Elias 2-27's Spiral Arms Driven by Self-gravity, or by a Companion? A Comparative Spiral Morphology Study. *ApJL*, *860*(1), Article L5, L5. https://doi.org/10.3847/2041-8213/aac7c9

Gammie, C. F. (1996). Layered Accretion in T Tauri Disks. *ApJ*, *457*, 355. https://doi.org/10.1086/176735

Gammie, C. F. (2001). Nonlinear Outcome of Gravitational Instability in Cooling, Gaseous Disks. *ApJ*, *553*(1), 174–183. https://doi.org/10.1086/320631

Gingold, R. A., & Monaghan, J. J. (1977). Smoothed particle hydrodynamics: theory and application to non-spherical stars. *MNRAS*, *181*, 375–389. https://doi.org/10.1093/mnras/181.3.375

Glorot, X., Bordes, A., & Bengio, Y. (2011). Deep sparse rectifier neural networks. *AISTATS*.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning* [http://www.deeplearningbook.org]. MIT Press.

Hall, C., Dong, R., Teague, R., Terry, J., Pinte, C., Paneque-Carreño, T., Veronesi, B., Alexander, R. D., & Lodato, G. (2020). Predicting the Kinematic

Evidence of Gravitational Instability. *ApJ*, *904*(2), Article 148, 148. https://doi.org/10.3847/1538-4357/abac17

Hall, C., Dong, R., Rice, K., Harries, T. J., Najita, J., Alexander, R., & Brittain, S. (2019). The Temporal Requirements of Directly Observing Self-gravitating Spiral Waves in Protoplanetary Disks with ALMA. *ApJ*, *871*(2), Article 228, 228. https://doi.org/10.3847/1538-4357/aafac2

Hall, C., Forgan, D., & Rice, K. (2017). Identifying and analysing protostellar disc fragments in smoothed particle hydrodynamics simulations. *MNRAS*, *470*(3), 2517–2538. https://doi.org/10.1093/mnras/stx1244

Hall, C., Forgan, D., Rice, K., Harries, T. J., Klaassen, P. D., & Biller, B. (2016). Directly observing continuum emission from self-gravitating spiral waves. *MNRAS*, *458*(1), 306–318. https://doi.org/10.1093/mnras/stw296

Hall, C., Rice, K., Dipierro, G., Forgan, D., Harries, T., & Alexander, R. (2018). Is the spiral morphology of the Elias 2-27 circumstellar disc due to gravitational instability? *MNRAS*, *477*(1), 1004–1014. https://doi.org/10.1093/mnras/sty550

Han, K., Wang, Y., Chen, H., Chen, X., Guo, J., Liu, Z., Tang, Y., Xiao, A., Xu, C., Xu, Y., Yang, Z., Zhang, Y., & Tao, D. (2023). A survey on vision transformer. *IEEE Transactions on Pattern Analysis &; Machine Intelligence*, *45*(01), 87–110. https://doi.org/10.1109/TPAMI.2022.3152247

Hanley, J. A., & McNeil, B. J. (1982). The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, *143 1*, 29–36.

Haworth, T. J., Cadman, J., Meru, F., Hall, C., Albertini, E., Forgan, D., Rice, K., & Owen, J. E. (2020). Massive discs around low-mass stars. *MNRAS*, *494*(3), 4130–4148. https://doi.org/10.1093/mnras/staa883

He, K., Zhang, X., Ren, S., & Sun, J. (2015). Deep Residual Learning for Image Recognition. *arXiv e-prints*, Article arXiv:1512.03385, arXiv:1512.03385.

Hinton, G. E., Dayan, P., Frey, B. J., & Neal, R. M. (1995). The "wake-sleep" algorithm for unsupervised neural networks. *Science*, *268*(5214), 1158–1161. Retrieved July 18, 2022, from http://www.jstor.org/stable/2888376

Huang, J., Andrews, S. M., Dullemond, C. P., Isella, A., Pérez, L. M., Guzmán, V. V., Öberg, K. I., Zhu, Z., Zhang, S., Bai, X.-N., Benisty, M., Birnstiel, T., Carpenter, J. M., Hughes, A. M., Ricci, L., Weaver, E., & Wilner, D. J. (2018). The Disk Substructures at High Angular Resolution Project (DSHARP). II. Characteristics of Annular Substructures. *ApJL*, *869*(2), Article L42, L42. https://doi.org/10.3847/2041-8213/aaf740

Huang, J., Andrews, S. M., Pérez, L. M., Zhu, Z., Dullemond, C. P., Isella, A., Benisty, M., Bai, X.-N., Birnstiel, T., Carpenter, J. M., Guzmán, V. V., Hughes, A. M., Öberg, K. I., Ricci, L., Wilner, D. J., & Zhang, S. (2018). The Disk Substructures at High Angular Resolution Project (DSHARP). III. Spiral Structures in the Millimeter Continuum of the Elias 27, IM Lup, and WaOph 6 Disks. *ApJL*, *869*(2), Article L43, L43. https://doi.org/10.3847/2041-8213/aaf7a0

Huang, P., Isella, A., Li, H., Li, S., & Ji, J. (2018). Identifying Anticyclonic Vortex Features Produced by the Rossby Wave Instability in Protoplanetary Disks. *ApJ*, *867*(1), Article 3, 3. https://doi.org/10.3847/1538-4357/aae317

Humphries, J., Hall, C., Haworth, T. J., & Nayakshin, S. (2021). Searching for wide-orbit gravitational instability protoplanets with ALMA in the dust continuum. *MNRAS*, *502*(1), 953–968. https://doi.org/10.1093/mnras/staa2411

Hutchison, M., Price, D. J., & Laibe, G. (2018). MULTIGRAIN: a smoothed particle hydrodynamic algorithm for multiple small dust grains and gas. *MNRAS*, *476*(2), 2186–2198. https://doi.org/10.1093/mnras/sty367

Ilee, J. D., Cyganowski, C. J., Brogan, C. L., Hunter, T. R., Forgan, D. H., Haworth, T. J., Clarke, C. J., & Harries, T. J. (2018). G11.92-0.61 MM 1: A Fragmented Keplerian Disk Surrounding a Proto-O Star. *ApJL*, *869*(2), Article L24, L24. https://doi.org/10.3847/2041-8213/aaeffc

Ilee, J. D., Forgan, D. H., Evans, M. G., Hall, C., Booth, R., Clarke, C. J., Rice, W. K. M., Boley, A. C., Caselli, P., Hartquist, T. W., & Rawlings, J. M. C. (2017). The chemistry of protoplanetary fragments formed via gravitational instabilities. *MNRAS*, *472*(1), 189–204. https://doi.org/10.1093/mnras/stx1966

Ilee, J. D., Hall, C., Walsh, C., Jiménez-Serra, I., Pinte, C., Terry, J., Bourke, T. L., & Hoare, M. (2020). Observing protoplanetary discs with the Square Kilometre Array - I. Characterizing pebble substructure caused by forming planets. *MNRAS*, *498*(4), 5116–5127. https://doi.org/10.1093/mnras/staa2699

Jo, Y., & Kim, J.-h. (2019). Machine-assisted semi-simulation model (MSSM): estimating galactic baryonic properties from their dark matter using a machine trained on hydrodynamic simulations. *MNRAS*, *489*(3), 3565–3581. https://doi.org/10.1093/mnras/stz2304

Johansen, A., & Lambrechts, M. (2017). Forming Planets via Pebble Accretion. *Annual Review of Earth and Planetary Sciences*, *45*(1), 359–387. https://doi.org/10.1146/annurev-earth-063016-020226

Juhász, A., & Rosotti, G. P. (2018). Spiral arms in thermally stratified protoplanetary discs. *MNRAS*, *474*(1), L32–L36. https://doi.org/10.1093/mnrasl/slx182

Kingma, D. P., & Ba, J. (2014). Adam: A Method for Stochastic Optimization. *arXiv e-prints*, Article arXiv:1412.6980, arXiv:1412.6980.

Kratter, K., & Lodato, G. (2016). Gravitational Instabilities in Circumstellar Disks. *ARA&A*, *54*, 271–311. https://doi.org/10.1146/annurev-astro-081915-023307

Laibe, G., & Price, D. J. (2014a). Dust and gas mixtures with multiple grain species - a one-fluid approach. *MNRAS*, *444*(2), 1940–1956. https://doi.org/10.1093/mnras/stu1367

Laibe, G., & Price, D. J. (2014b). Dusty gas with one fluid. *MNRAS*, *440*(3), 2136–2146. https://doi.org/10.1093/mnras/stu355

Laibe, G., & Price, D. J. (2014c). Dusty gas with one fluid in smoothed particle hydrodynamics. *MNRAS*, *440*(3), 2147–2163. https://doi.org/10.1093/mnras/stu359

Lambrechts, M., & Johansen, A. (2012). Rapid growth of gas-giant cores by pebble accretion. *A&A*, *544*, Article A32, A32. https://doi.org/10.1051/0004-6361/201219127

Lin, D. N. C., & Papaloizou, J. (1985, January). On the dynamical origin of the solar system. In D. C. Black & M. S. Matthews (Eds.), *Protostars and planets ii* (pp. 981–1072).

Lochner, M., McEwen, J. D., Peiris, H. V., Lahav, O., & Winter, M. K. (2016). Photometric Supernova Classification with Machine Learning. *ApJS*, *225*(2), Article 31, 31. https://doi.org/10.3847/0067-0049/225/2/31

Lodato, G., & Rice, W. K. M. (2004). Testing the locality of transport in self-gravitating accretion discs. *MNRAS*, *351*(2), 630–642. https://doi.org/10.1111/j.1365-2966.2004.07811.x

Lodato, G., & Rice, W. K. M. (2005). Testing the locality of transport in self-gravitating accretion discs - II. The massive disc case. *MNRAS*, *358*(4), 1489–1500. https://doi.org/10.1111/j.1365-2966.2005.08875.x

Lodato, G., & Clarke, C. J. (2011). Resolution requirements for smoothed particle hydrodynamics simulations of self-gravitating accretion discs. *MNRAS*, *413*(4), 2735–2740. https://doi.org/10.1111/j.1365-2966.2011.18344.x

Longarini, C., Lodato, G., Toci, C., Veronesi, B., Hall, C., Dong, R., & Patrick Terry, J. (2021). Investigating Protoplanetary Disk Cooling through Kinematics: Analytical GI Wiggle. *ApJL*, *920*(2), Article L41, L41. https://doi.org/10.3847/2041-8213/ac2df6

Lucy, L. B. (1977). A numerical approach to the testing of the fission hypothesis. *Astronomical Journal*, *82*, 1013–1024. https://doi.org/10.1086/112164

Marois, C., Macintosh, B., Barman, T., Zuckerman, B., Song, I., Patience, J., Lafrenière, D., & Doyon, R. (2008). Direct Imaging of Multiple Planets Orbiting the Star HR 8799. *Science*, *322*(5906), 1348. https://doi.org/10.1126/science.1166585

McKay, M. D., Beckman, R. J., & Conover, W. J. (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, *21*(2), 239–245. Retrieved July 18, 2022, from http://www.jstor.org/stable/1268522

Menou, K., & Goodman, J. (2004). Low-Mass Protoplanet Migration in T Tauri α-Disks. *ApJ*, *606*(1), 520–531. https://doi.org/10.1086/382947

Meru, F., Juhász, A., Ilee, J. D., Clarke, C. J., Rosotti, G. P., & Booth, R. A. (2017). On the Origin of the Spiral Morphology in the Elias 2-27 Circumstellar Disk. *ApJL*, *839*(2), Article L24, L24. https://doi.org/10.3847/2041-8213/aa6837

Mignone, A., Bodo, G., Massaglia, S., Matsakos, T., Tesileanu, O., Zanni, C., & Ferrari, A. (2007). PLUTO: A Numerical Code for Computational Astrophysics. *ApJs*, *170*(1), 228–242. https://doi.org/10.1086/513316

Mignone, A., Zanni, C., Tzeferacos, P., van Straalen, B., Colella, P., & Bodo, G. (2012). The PLUTO Code for Adaptive Mesh Computations in Astrophysical Fluid Dynamics. *ApJs*, *198*(1), Article 7, 7. https://doi.org/10.1088/0067-0049/198/1/7

Minaee, S., Boykov, Y., Porikli, F., Plaza, A., Kehtarnavaz, N., & Terzopoulos, D. (2020). Image Segmentation Using Deep Learning: A Survey. *arXiv e-prints*, Article arXiv:2001.05566, arXiv:2001.05566.

Miotello, A., Bruderer, S., & van Dishoeck, E. F. (2014). Protoplanetary disk masses from CO isotopologue line emission. *A&A*, *572*, Article A96, A96. https://doi.org/10.1051/0004-6361/201424712

Miotello, A., Facchini, S., van Dishoeck, E. F., & Bruderer, S. (2018). Probing the protoplanetary disk gas surface density distribution with ^{13}CO emission. *A&A*, *619*, Article A113, A113. https://doi.org/10.1051/0004-6361/201833595

Miotello, A., van Dishoeck, E. F., Kama, M., & Bruderer, S. (2016). Determining protoplanetary disk gas masses from CO isotopologues line observations. *A&A*, *594*, Article A85, A85. https://doi.org/10.1051/0004-6361/201628159

Möller, A., & de Boissière, T. (2020). SuperNNova: an open-source framework for Bayesian, neural network-based supernova classification. *MNRAS*, *491*(3), 4277–4293. https://doi.org/10.1093/mnras/stz3312

Monaghan, J. J. (1992). Smoothed particle hydrodynamics. *ARA&A*, *30*, 543–574. https://doi.org/10.1146/annurev.aa.30.090192.002551

Nasa exoplanet archive [Accessed: 2024-02-13]. (2024).

Nayakshin, S. (2010). Formation of planets by tidal downsizing of giant planet embryos. *MNRAS*, *408*(1), L36–L40. https://doi.org/10.1111/j.1745-3933.2010.00923.x

Nayakshin, S., Tsukagoshi, T., Hall, C., Vazan, A., Helled, R., Humphries, J., Meru, F., Neunteufel, P., & Panic, O. (2020). TW Hya: an old protoplanetary disc revived by its planet. *MNRAS*, *495*(1), 285–304. https://doi.org/10.1093/mnras/staa1132

Nelson, A. F. (2006). Numerical requirements for simulations of self-gravitating and non-self-gravitating discs. *MNRAS*, *373*(3), 1039–1073. https://doi.org/10.1111/j.1365-2966.2006.11119.x

Nelson, R. P., Gressel, O., & Umurhan, O. M. (2013). Linear and non-linear evolution of the vertical shear instability in accretion discs. *MNRAS*, *435*(3), 2610–2632. https://doi.org/10.1093/mnras/stt1475

Öberg, K. I., Guzmán, V. V., Walsh, C., Aikawa, Y., Bergin, E. A., Law, C. J., Loomis, R. A., Alarcón, F., Andrews, S. M., Bae, J., Bergner, J. B., Boehler, Y., Booth, A. S., Bosman, A. D., Calahan, J. K., Cataldi, G., Cleeves, L. I., Czekala, I., Furuya, K., … Zhang, K. (2021). Molecules with ALMA at Planet-forming Scales (MAPS). I. Program Overview and Highlights. *ApJs*, *257*(1), Article 1, 1. https://doi.org/10.3847/1538-4365/ac1432

Ogilvie, G. I., & Lubow, S. H. (2002). On the wake generated by a planet in a disc. *MNRAS*, *330*(4), 950–954. https://doi.org/10.1046/j.1365-8711.2002.05148.x

Okuzumi, S., Momose, M., Sirono, S.-i., Kobayashi, H., & Tanaka, H. (2016). Sintering-induced Dust Ring Formation in Protoplanetary Disks: Application to the HL Tau Disk. *ApJ*, *821*(2), Article 82, 82. https://doi.org/10.3847/0004-637X/821/2/82

Ormel, C. W., & Klahr, H. H. (2010). The effect of gas drag on the growth of protoplanets. Analytical expressions for the accretion of small bodies in laminar disks. *A&A*, *520*, Article A43, A43. https://doi.org/10.1051/0004-6361/201014903

Paneque-Carreño, T., Pérez, L. M., Benisty, M., Hall, C., Veronesi, B., Lodato, G., Sierra, A., Carpenter, J. M., Andrews, S. M., Bae, J., Henning, T.,

Kwon, W., Linz, H., Loinard, L., Pinte, C., Ricci, L., Tazzari, M., Testi, L., & Wilner, D. (2021). Spiral Arms and a Massive Dust Disk with Non-Keplerian Kinematics: Possible Evidence for Gravitational Instability in the Disk of Elias 2-27. *ApJ*, *914*(2), Article 88, 88. https://doi.org/10.3847/1538-4357/abf243

Panic, O. (2009, October). *High angular resolution studies of protoplanetary discs* [Doctoral dissertation, -].

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., … Chintala, S. (2019). PyTorch: An Imperative Style, High-Performance Deep Learning Library. *arXiv e-prints*, Article arXiv:1912.01703, arXiv:1912.01703.

Perez, S., Dunhill, A., Casassus, S., Roman, P., Szulágyi, J., Flores, C., Marino, S., & Montesinos, M. (2015). Planet Formation Signposts: Observability of Circumplanetary Disks via Gas Kinematics. *ApJL*, *811*(1), Article L5, L5. https://doi.org/10.1088/2041-8205/811/1/L5

Pérez, L. M., Carpenter, J. M., Andrews, S. M., Ricci, L., Isella, A., Linz, H., Sargent, A. I., Wilner, D. J., Henning, T., Deller, A. T., Chandler, C. J., Dullemond, C. P., Lazio, J., Menten, K. M., Corder, S. A., Storm, S., Testi, L., Tazzari, M., Kwon, W., … Mundy, L. G. (2016). Spiral density waves in a young protoplanetary disk. *Science*, *353*(6307), 1519–1521. https://doi.org/10.1126/science.aaf8296

Pérez, S., Casassus, S., & Benítez-Llambay, P. (2018). Observability of planet-disc interactions in CO kinematics. *MNRAS*, *480*(1), L12–L17. https://doi.org/10.1093/mnrasl/sly109

Pinilla, P., Birnstiel, T., Ricci, L., Dullemond, C. P., Uribe, A. L., Testi, L., & Natta, A. (2012). Trapping dust particles in the outer regions of protoplanetary disks. *A&A*, *538*, Article A114, A114. https://doi.org/10.1051/0004-6361/201118204

Pinte, C., Dent, W. R. F., Ménard, F., Hales, A., Hill, T., Cortes, P., & de Gregorio-Monsalvo, I. (2016). Dust and Gas in the Disk of HL Tauri: Surface Density, Dust Settling, and Dust-to-gas Ratio. *ApJ*, *816*(1), Article 25, 25. https://doi.org/10.3847/0004-637X/816/1/25

Pinte, C., Harries, T. J., Min, M., Watson, A. M., Dullemond, C. P., Woitke, P., Ménard, F., & Durán-Rojas, M. C. (2009). Benchmark problems for continuum radiative transfer. High optical depths, anisotropic scattering, and polarisation. *A&A*, *498*(3), 967–980. https://doi.org/10.1051/0004-6361/200811555

Pinte, C., Ménard, F., Duchêne, G., & Bastien, P. (2006). Monte Carlo radiative transfer in protoplanetary disks. *A&A*, *459*(3), 797–804. https://doi.org/10.1051/0004-6361:20053275

Pinte, C., Price, D. J., Ménard, F., Duchêne, G., Christiaens, V., Andrews, S. M., Huang, J., Hill, T., van der Plas, G., Perez, L. M., Isella, A., Boehler, Y., Dent, W. R. F., Mentiplay, D., & Loomis, R. A. (2020). Nine Localized Deviations from Keplerian Rotation in the DSHARP Circumstellar Disks: Kinematic Evidence for Protoplanets Carving the Gaps. *ApJL*, *890*(1), Article L9, L9. https://doi.org/10.3847/2041-8213/ab6dda

Pinte, C., Price, D. J., Ménard, F., Duchêne, G., Dent, W. R. F., Hill, T., de Gregorio-Monsalvo, I., Hales, A., & Mentiplay, D. (2018). Kinematic Evidence for an Embedded Protoplanet in a Circumstellar Disk. *ApJL*, *860*(1), Article L13, L13. https://doi.org/10.3847/2041-8213/aac6dc

Pinte, C., Teague, R., Flaherty, K., Hall, C., Facchini, S., & Casassus, S. (2023, July). Kinematic Structures in Planet-Forming Disks. In S. Inutsuka, Y. Aikawa, T. Muto, K. Tomida, & M. Tamura (Eds.), *Protostars and planets vii* (p. 645, Vol. 534). https://doi.org/10.48550/arXiv.2203.09528

Pinte, C., van der Plas, G., Ménard, F., Price, D. J., Christiaens, V., Hill, T., Mentiplay, D., Ginski, C., Choquet, E., Boehler, Y., Duchêne, G., Perez, S., & Casassus, S. (2019). Kinematic detection of a planet carving a gap in a protoplanetary disk. *Nature Astronomy*, *3*, 1109–1114. https://doi.org/10.1038/s41550-019-0852-6

Price, D. J. (2007). splash: An Interactive Visualisation Tool for Smoothed Particle Hydrodynamics Simulations. *PASA*, *24*(3), 159–173. https://doi.org/10.1071/AS07022

Price, D. J. (2010). Smoothed Particle Magnetohydrodynamics - IV. Using the vector potential. *MNRAS*, *401*(3), 1475–1499. https://doi.org/10.1111/j.1365-2966.2009.15763.x

Price, D. J., Wurster, J., Tricco, T. S., Nixon, C., Toupin, S., Pettitt, A., Chan, C., Mentiplay, D., Laibe, G., Glover, S., Dobbs, C., Nealon, R., Liptai, D., Worpel, H., Bonnerot, C., Dipierro, G., Ballabio, G., Ragusa, E., Federrath, C., … Lodato, G. (2018). Phantom: A Smoothed Particle Hydrodynamics and Magnetohydrodynamics Code for Astrophysics. *PASA*, *35*, Article e031, e031. https://doi.org/10.1017/pasa.2018.25

Quénard, D., Ilee, J. D., Jiménez-Serra, I., Forgan, D. H., Hall, C., & Rice, K. (2018). The Fate of Formamide in a Fragmenting Protoplanetary Disk. *ApJ*, *868*(1), Article 9, 9. https://doi.org/10.3847/1538-4357/aae4dd

Rafikov, R. R. (2005). Can Giant Planets Form by Direct Gravitational Instability? *ApJL*, *621*(1), L69–L72. https://doi.org/10.1086/428899

Rice, K., Lopez, E., Forgan, D., & Biller, B. (2015). Disc fragmentation rarely forms planetary-mass objects. *MNRAS*, *454*(2), 1940–1947. https://doi.org/10.1093/mnras/stv1997

Rice, W. K. M., Armitage, P. J., Bate, M. R., & Bonnell, I. A. (2003). The effect of cooling on the global stability of self-gravitating protoplanetary discs. *MNRAS*, *339*(4), 1025–1030. https://doi.org/10.1046/j.1365-8711.2003.06253.x

Rice, W. K. M., Armitage, P. J., Bonnell, I. A., Bate, M. R., Jeffers, S. V., & Vine, S. G. (2003). Substellar companions and isolated planetary-mass objects from protostellar disc fragmentation. *MNRAS*, *346*(3), L36–L40. https://doi.org/10.1111/j.1365-2966.2003.07317.x

Rice, W. K. M., Armitage, P. J., Mamatsashvili, G. R., Lodato, G., & Clarke, C. J. (2011). Stability of self-gravitating discs under irradiation. *MNRAS*, *418*(2), 1356–1362. https://doi.org/10.1111/j.1365-2966.2011.19586.x

Rice, W. K. M., & Armitage, P. J. (2009). Time-dependent models of the structure and stability of self-gravitating protoplanetary discs. *MNRAS*, *396*(4), 2228–2236. https://doi.org/10.1111/j.1365-2966.2009.14879.x

Rice, W. K. M., Lodato, G., & Armitage, P. J. (2005). Investigating fragmentation conditions in self-gravitating accretion discs. *MNRAS*, *364*(1), L56–L60. https://doi.org/10.1111/j.1745-3933.2005.00105.x

Rice, W. K. M., Lodato, G., Pringle, J. E., Armitage, P. J., & Bonnell, I. A. (2004). Accelerated planetesimal growth in self-gravitating protoplanetary discs. *MNRAS*, *355*(2), 543–552. https://doi.org/10.1111/j.1365-2966.2004.08339.x

Rosotti, G. P., Tazzari, M., Booth, R. A., Testi, L., Lodato, G., & Clarke, C. (2019). The time evolution of dusty protoplanetary disc radii: observed and physical radii differ. *MNRAS*, *486*(4), 4829–4844. https://doi.org/10.1093/mnras/stz1190

Rubinstein, A. E., Macías, E., Espaillat, C. C., Zhang, K., Calvet, N., & Robinson, C. (2018). A Cavity of Large Grains in the Disk around the Group II Herbig Ae/Be Star HD 142666. *ApJ*, *860*(1), Article 7, 7. https://doi.org/10.3847/1538-4357/aabfba

Safronov, V. S. (1972). *Evolution of the protoplanetary cloud and formation of the earth and planets.*

Sanchis, E., Testi, L., Natta, A., Facchini, S., Manara, C. F., Miotello, A., Ercolano, B., Henning, T., Preibisch, T., Carpenter, J. M., de Gregorio-Monsalvo, I., Jayawardhana, R., Lopez, C., Mužić, K., Pascucci, I.,

Santamaría-Miranda, A., van Terwisga, S., & Williams, J. P. (2021). Measuring the ratio of the gas and dust emission radii of protoplanetary disks in the Lupus star-forming region. *A&A*, *649*, Article A19, A19. https://doi.org/10.1051/0004-6361/202039733

Santos, N. C., Israelian, G., & Mayor, M. (2004). Spectroscopic [Fe/H] for 98 extra-solar planet-host stars. Exploring the probability of planet formation. *A&A*, *415*, 1153–1166. https://doi.org/10.1051/0004-6361:20034469

Seitenzahl, I. R., Vogt, F. P. A., Terry, J. P., Ghavamian, P., Dopita, M. A., Ruiter, A. J., & Sukhbold, T. (2018). Integral Field Spectroscopy of Supernova Remnant 1E0102-7219 Reveals Fast-moving Hydrogen and Sulfur-rich Ejecta. *ApJL*, *853*(2), Article L32, L32. https://doi.org/10.3847/2041-8213/aaa958

Shakura, N. I., & Sunyaev, R. A. (1973). Black holes in binary systems. Observational appearance. *A&A*, *24*, 337–355.

Shen, B.-T., Tang, Y.-W., & Koch, P. M. (2020). Spiral-arm Substructures in the Asymmetrical Dust Rings of the Circumstellar Disk MWC 758. *ApJ*, *904*(2), Article 125, 125. https://doi.org/10.3847/1538-4357/abbd9f

Siess, L., Dufour, E., & Forestini, M. (2000). An internet server for pre-main sequence tracks of low- and intermediate-mass stars. *A&A*, *358*, 593–599.

Simon, J. B., Hughes, A. M., Flaherty, K. M., Bai, X.-N., & Armitage, P. J. (2015). Signatures of MRI-driven Turbulence in Protoplanetary Disks: Predictions for ALMA Observations. *ApJ*, *808*(2), Article 180, 180. https://doi.org/10.1088/0004-637X/808/2/180

Stone, J. M., Tomida, K., White, C. J., & Felker, K. G. (2020). The Athena++ Adaptive Mesh Refinement Framework: Design and Magnetohydrodynamic Solvers. *ApJs*, *249*(1), Article 4, 4. https://doi.org/10.3847/1538-4365/ab929b

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2014). Going Deeper with Convolutions. *arXiv e-prints*, Article arXiv:1409.4842, arXiv:1409.4842.

Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. *arXiv e-prints*, Article arXiv:1905.11946, arXiv:1905.11946.

Tan, M., & Le, Q. V. (2021). EfficientNetV2: Smaller Models and Faster Training. *arXiv e-prints*, Article arXiv:2104.00298, arXiv:2104.00298.

Tazzari, M., Testi, L., Natta, A., Williams, J. P., Ansdell, M., Carpenter, J. M., Facchini, S., Guidi, G., Hogherheijde, M., Manara, C. F., Miotello, A.,

& van der Marel, N. (2021). The first ALMA survey of protoplanetary discs at 3 mm: demographics of grain growth in the Lupus region. *MNRAS*, *506*(4), 5117–5128. https://doi.org/10.1093/mnras/stab1912

Teague, R., Bae, J., Bergin, E. A., Birnstiel, T., & Foreman-Mackey, D. (2018). A Kinematical Detection of Two Embedded Jupiter-mass Planets in HD 163296. *ApJL*, *860*(1), Article L12, L12. https://doi.org/10.3847/2041-8213/aac6d7

Terry, J. P., Hall, C., Abreau, S., & Gleyzer, S. (2022). Locating Hidden Exoplanets in ALMA Data Using Machine Learning. *ApJ*, *941*(2), Article 192, 192. https://doi.org/10.3847/1538-4357/aca477

Terry, J. P., Hall, C., Abreau, S., & Gleyzer, S. (2023). Kinematic Evidence of an Embedded Protoplanet in HD 142666 Identified by Machine Learning. *ApJ*, *947*(2), Article 60, 60. https://doi.org/10.3847/1538-4357/acc737

Terry, J. P., Hall, C., Longarini, C., Lodato, G., Toci, C., Veronesi, B., Paneque-Carreño, T., & Pinte, C. (2022). Constraining protoplanetary disc mass using the GI wiggle. *MNRAS*, *510*(2), 1671–1679. https://doi.org/10.1093/mnras/stab3513

Toci, C., Rosotti, G., Lodato, G., Testi, L., & Trapman, L. (2021). On the secular evolution of the ratio between gas and dust radii in protoplanetary discs. *MNRAS*, *507*(1), 818–833. https://doi.org/10.1093/mnras/stab2112

Tomida, K., Machida, M. N., Hosokawa, T., Sakurai, Y., & Lin, C. H. (2017). Grand-design Spiral Arms in a Young Forming Circumstellar Disk. *ApJL*, *835*(1), Article L11, L11. https://doi.org/10.3847/2041-8213/835/1/L11

Toomre, A. (1964). On the gravitational stability of a disk of stars. *ApJ*, *139*, 1217–1238. https://doi.org/10.1086/147861

Trapman, L., Ansdell, M., Hogerheijde, M. R., Facchini, S., Manara, C. F., Miotello, A., Williams, J. P., & Bruderer, S. (2020). Constraining the radial drift of millimeter-sized grains in the protoplanetary disks in Lupus. *A&A*, *638*, Article A38, A38. https://doi.org/10.1051/0004-6361/201834537

Tsoi, A. C. (1997). Recurrent neural network architectures: An overview. *International School on Neural Networks, Initiated by IIASS and EMFCSC*, 1–26.

Tsukagoshi, T., Muto, T., Nomura, H., Kawabe, R., Kanagawa, K. D., Okuzumi, S., Ida, S., Walsh, C., Millar, T. J., Takahashi, S. Z., Hashimoto, J., Uyama, T., & Tamura, M. (2019). Discovery of An au-scale Excess in Millimeter Emission from the Protoplanetary Disk around TW Hya. *ApJL*, *878*(1), Article L8, L8. https://doi.org/10.3847/2041-8213/ab224c

Tzeng, E., Hoffman, J., Saenko, K., & Darrell, T. (2017). Adversarial Discriminative Domain Adaptation. *arXiv e-prints*, Article arXiv:1702.05464, arXiv:1702.05464. https://doi.org/10.48550/arXiv.1702.05464

van Hespen, K. M., Zwanenburg, J. J. M., Dankbaar, J. W., Geerlings, M. I., Hendrikse, J., & Kuijf, H. J. (2021). An anomaly detection approach to identify chronic brain infarcts on MRI. *Scientific Reports*, *11*, Article 7714, 7714. https://doi.org/10.1038/s41598-021-87013-4

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention Is All You Need. *arXiv e-prints*, Article arXiv:1706.03762, arXiv:1706.03762.

Veronesi, B., Lodato, G., Dipierro, G., Ragusa, E., Hall, C., & Price, D. J. (2019). Multiwavelength observations of protoplanetary discs as a proxy for the gas disc mass. *MNRAS*, *489*(3), 3758–3768. https://doi.org/10.1093/mnras/stz2384

Veronesi, B., Paneque-Carreño, T., Lodato, G., Testi, L., Pérez, L. M., Bertin, G., & Hall, C. (2021). A Dynamical Measurement of the Disk Mass in Elias 227. *ApJL*, *914*(2), Article L27, L27. https://doi.org/10.3847/2041-8213/abfe6a

Vigan, A., Bonavita, M., Biller, B., Forgan, D., Rice, K., Chauvin, G., Desidera, S., Meunier, J. .-., Delorme, P., Schlieder, J. E., Bonnefoy, M., Carson, J., Covino, E., Hagelberg, J., Henning, T., Janson, M., Lagrange, A. .-., Quanz, S. P., Zurlo, A., … Thalmann, C. (2017). The VLT/NaCo large program to probe the occurrence of exoplanets and brown dwarfs at wide orbits. IV. Gravitational instability rarely forms wide, giant planets. *A&A*, *603*, Article A3, A3. https://doi.org/10.1051/0004-6361/201630133

Vilalta, R., Dhar Gupta, K., Boumber, D., & Meskhi, M. M. (2019). A General Approach to Domain Adaptation with Applications in Astronomy. *PASP*, *131*(1004), 108008. https://doi.org/10.1088/1538-3873/aaf1fc

Voulodimos, A., Doulamis, N., Doulamis, A., Protopapadakis, E., & Andina, D. (2018). Deep learning for computer vision: A brief review. intell. *Neuroscience*, *7068349*. https://doi.org/10.1155/2018/7068349

Weidenschilling, S. J. (1977). Aerodynamics of solid bodies in the solar nebula. *MNRAS*, *180*, 57–70. https://doi.org/10.1093/mnras/180.2.57

Williams, J. P., & Best, W. M. J. (2014). A Parametric Modeling Approach to Measuring the Gas Masses of Circumstellar Disks. *ApJ*, *788*(1), Article 59, 59. https://doi.org/10.1088/0004-637X/788/1/59

Wölfer, L., Facchini, S., Kurtovic, N. T., Teague, R., van Dishoeck, E. F., Benisty, M., Ercolano, B., Lodato, G., Miotello, A., Rosotti, G., Testi, L., &

Ubeira Gabellini, M. G. (2021). A highly non-Keplerian protoplanetary disc. Spiral structure in the gas disc of CQ Tau. *A&A*, *648*, Article A19, A19. https://doi.org/10.1051/0004-6361/202039469

Wurster, J. (2021). Do we need non-ideal magnetohydrodynamic to model protostellar discs? *MNRAS*, *501*(4), 5873–5891. https://doi.org/10.1093/mnras/staa3943

Wurster, J., Bate, M. R., & Bonnell, I. A. (2021). The impact of non-ideal magnetohydrodynamic processes on discs, outflows, counter-rotation, and magnetic walls during the early stages of star formation. *MNRAS*, *507*(2), 2354–2372. https://doi.org/10.1093/mnras/stab2296

Wurster, J., Bate, M. R., & Price, D. J. (2018). The effect of extreme ionization rates during the initial collapse of a molecular cloud core. *MNRAS*, *476*(2), 2063–2074. https://doi.org/10.1093/mnras/sty392

Wurster, J., Price, D. J., & Bate, M. R. (2016). Can non-ideal magnetohydrodynamics solve the magnetic braking catastrophe? *MNRAS*, *457*(1), 1037–1061. https://doi.org/10.1093/mnras/stw013

Xu, J., Pan, Y., Pan, X., Hoi, S., Yi, Z., & Xu, Z. (2021). RegNet: Self-Regulated Network for Image Classification. *arXiv e-prints*, Article arXiv:2101.00590, arXiv:2101.00590.

Youdin, A. N., & Goodman, J. (2005). Streaming Instabilities in Protoplanetary Disks. *ApJ*, *620*(1), 459–469. https://doi.org/10.1086/426895

Zhang, K., Bergin, E. A., Schwarz, K., Krijt, S., & Ciesla, F. (2019). Systematic Variations of CO Gas Abundance with Radius in Gas-rich Protoplanetary Disks. *ApJ*, *883*(1), Article 98, 98. https://doi.org/10.3847/1538-4357/ab38b9

Zhang, S., Zhu, Z., & Kang, M. (2022). PGNets: planet mass prediction using convolutional neural networks for radio continuum observations of protoplanetary discs. *MNRAS*, *510*(3), 4473–4484. https://doi.org/10.1093/mnras/stab3502

Zhao, Z.-Q., Zheng, P., Xu, S.-t., & Wu, X. (2018). Object Detection with Deep Learning: A Review. *arXiv e-prints*, Article arXiv:1807.05511, arXiv:1807.05511.

Zhou, W., Yang, Y., Yu, C., Liu, J., Duan, X., Weng, Z., Chen, D., Liang, Q., Fang, Q., Zhou, J., Ju, H., Luo, Z., Guo, W., Ma, X., Xie, X., Wang, R., & Zhou, L. (2021). Ensembled deep learning model outperforms human experts in diagnosing biliary atresia from sonographic gallbladder images. *Nature Communications*, *12*, Article 1259, 1259. https://doi.org/10.1038/s41467-021-21466-z

Zormpas, A., Birnstiel, T., Rosotti, G. P., & Andrews, S. M. (2022). A large population study of protoplanetary disks. Explaining the millimeter size-luminosity relation with or without substructure. *A&A*, *661*, Article A66, A66. https://doi.org/10.1051/0004-6361/202142046