Flow Parameters and their Effect on Extrusion Printing

Master Thesis

Harold Rutten

Delft University of Technology

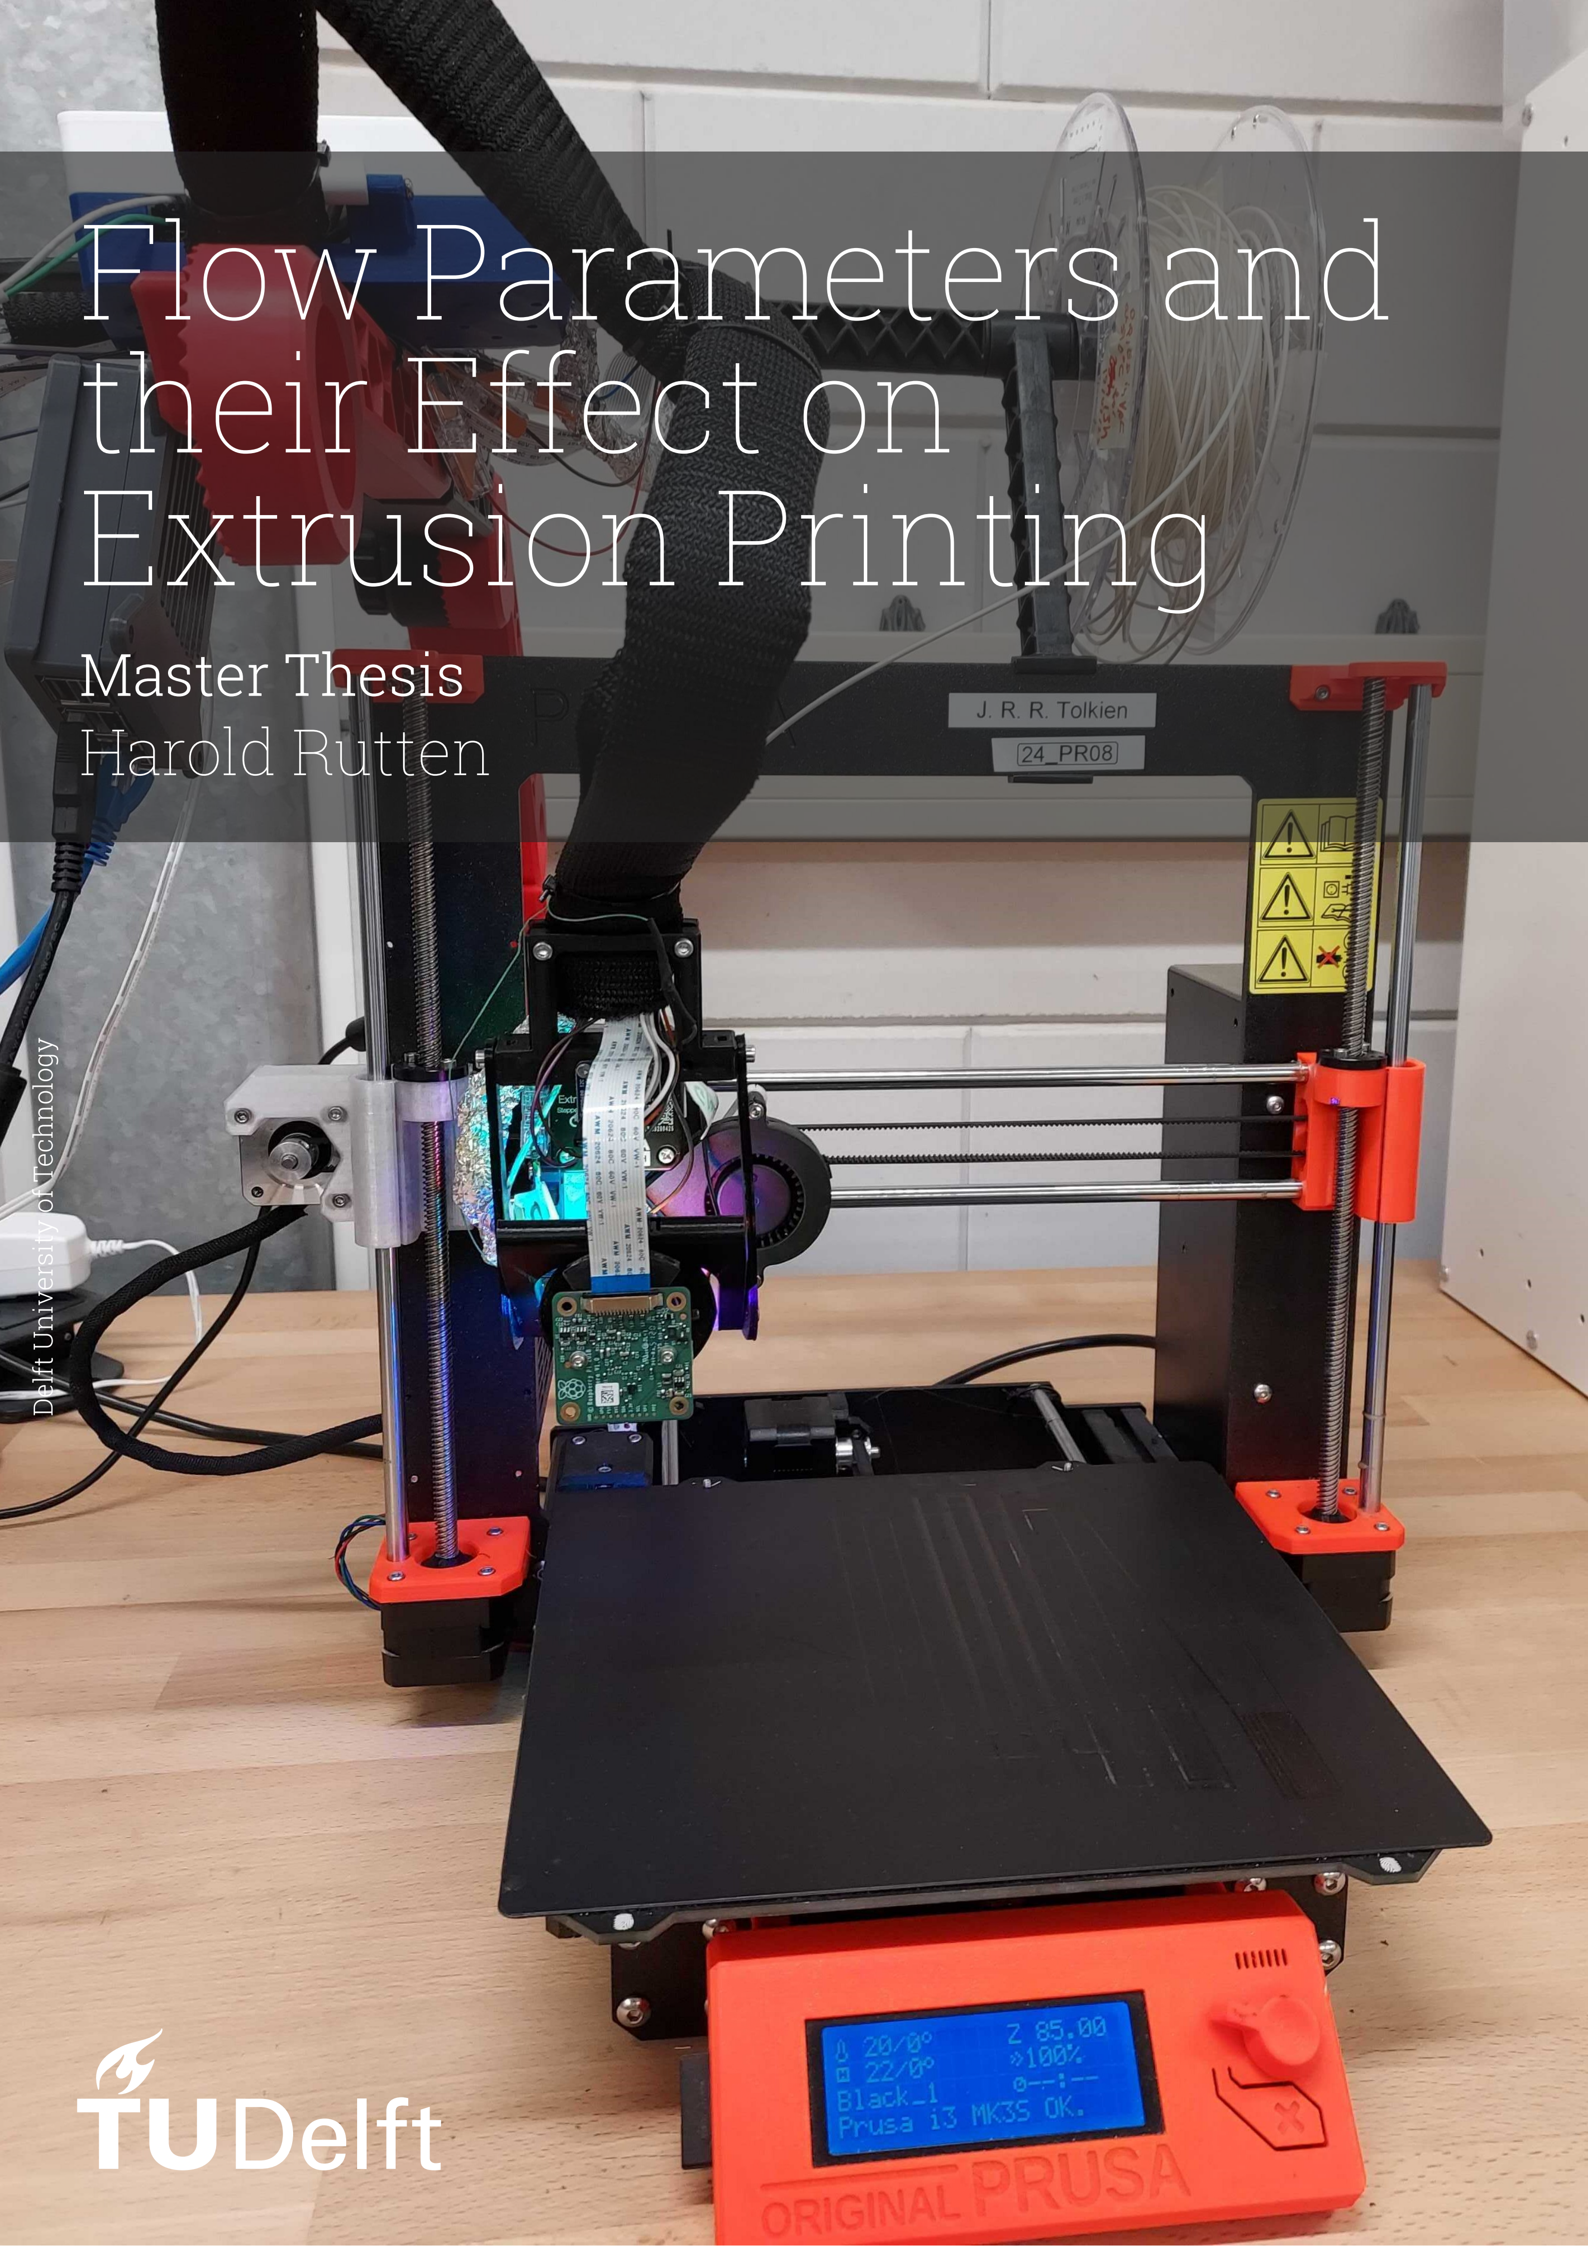

Flow Parameters and their Effect on Extrusion Printing

by

Harold Rutten

Student Name	Student Number
Harold Rutten	4466446

Supervisor: K. Masania
Daily Supervisor: M. Amani
Project Duration: December, 2023 - November, 2024
Faculty: Faculty of Aerospace Engineering, Delft

Cover: Image of the Test Setup that was designed and used for this thesis
Style: TU Delft Report Style, with modifications by Daan Zwaneveld

Preface

This thesis has been an incredible roller-coaster ride of a journey. It has challenged me in unexpected ways, and given me many an unexpected lesson which took a sledgehammer to my expectations and understandings. This thesis at first seemed like a fairly straight road, but has instead been so full of twists and turns as to be almost comical. It seemed at times as though nothing worked as it should, and some things worked as they shouldn't, to my great bewilderment. Of the issues I foresaw, some of the greatest shrunk to naught as I approached them, and some of the smallest reached insurmountable heights. Looking back, I am glad that I still strove forwards, and found ways over or around so many of those issues, but I could not have done it without the many people who kept me company and helped me push on.

One of those is my daily supervisor, Muhamad Amani, who pushed me to keep improving the experimental setup, and always strive for better results. He put up with my many questions and updates at odd hours, and did an admirable job keeping me on track when every rabbit hole enticed me towards information or a goal that was neither relevant nor helpful, but seemed interesting. He showed me the difference between the technical engineering I was accustomed to when starting this thesis, and the world of Academia. He was patient with me when I made the many mistakes on my road to learning how to read and write for academia, and gave me invaluable advice on this journey.

I would also like to thank my thesis supervisor Kunal Masania, who first truly introduced me to additive manufacturing during his course, which had previously only been something I had read about, but never got to try. He has been an invaluable source of information on the many facets of additive manufacturing, and drove me to find its boundaries, and strive to push them even further. He set me down the path of in-situ monitoring, and helped me find a topic that truly interests and motivated me, and a topic I will no-doubt be working on in some shape or form in the future. Last but not least, he has introduced me, through the research group he has assembled, the Shaping Matter Lab, to people whom I will always fondly remember.

Throughout the many long days that made up this thesis, this research group has supported, encouraged, and advised me, helping me grow both as a student, and as a person. I would like to mention in particular Guillermo Presa Magrina, with whom I have spent many a late night in the beginning of my thesis, as we both worked on our respective topics, bouncing ideas off each other and pushing each other ever onwards. During those long sessions in the lab, we could always count on the wonderful humor of Boris Ulyanov, who kept our spirits high with endless memes and anecdotes. I would also like to extend my heartfelt gratitude to Kathrin Weiland, who helped me navigate the ups and downs that come with research, and whose moral support helped me make lemonade out of lemons. I would also like to thank the "late-night-crew", who kept me company through shadows to the edge of night. Thank you to Caroline Houriet and Sourav Patranabish, for your invaluable insights into the weird yet wonderful material that is Vectra. There are too many others to list, but I am grateful to you all.

Finally, I wish to thank the many members of DARE, who have been a source of constant support, and helped me take my mind off of the work when I needed to clear my head. I will fondly remember my time in DARE, the people I met, and the memories made. I wish in particular to thank someone from DARE who came almost out of the blue, not once but twice, at perhaps some of the darkest moments of the thesis when nothing seemed to work. This person reminded me that hope can be found, even in the darkest of times, if one only remembers to turn on the light. You know who you are, thank you for helping me find the light switch.

Harold Rutten
Delft, December 2024

Summary

Additive manufacturing using fused deposition modelling has come a long way in the recent years, with industry beginning to adopt it as a manufacturing technique [8]. The use of filaments as a feedstock to selectively place material exactly where it is necessary has enabled a greater freedom in design for more complicated parts to be manufactured. This further includes using new materials such as liquid crystal polymers, which can be selectively anisotropic, providing greatly improved properties in a controllable direction [10]. This composite-like capability is further improved by being a single material recyclable material, where the actual properties can be tuned and varied during the process [13], and the radius of curvature is much lower compared to conventional composites. These capabilities are still being developed, with the production reliability being poor as both the understanding of the material response to processing and the control of the process parameters are thus far not yet fully understood and controlled [6][16]. This work has combined sensors for process monitoring while the process is ongoing and evaluated those sensors for future process monitoring with the aim of corrective measures. In doing so, it was demonstrated that there is a much greater need for non-intrusive sensors, that do not affect the process. Furthermore, although the fusion of sensor data was achieved by combining the sensor data based on toolhead position, the method by which the toolhead position was found proved to have a lot of drawbacks. This has shown the need to improve position measurement techniques, as well as adding data transfer connections from the printer to a separate system that effectuates the data processing and could eventually act upon the gathered information. As part of the sensor suite, the output geometry was measured to obtain the transition time of the output geometry of the process between set line-widths, and demonstrated how this transition is affected by the process parameters that would be controlled based on sensor readings. It was previously understood that the degree of alignment, cooling time and the linewidth dictated the anisotropy, and that the change in material extrusion was dependent on the pressure in the nozzle. This work has demonstrated that the linewidth will also change as a result of the cooling process warping the deposited geometry, which significantly affects the measured linewidth and lineheight. Further unexpected changes in the transition delay were discovered, which could not be explained by pressure alone. These could be explained by turbulent flow in the nozzle, contradicting previous expectations that flow in the nozzle is linear, that may occur in the nozzle at lower viscosities affecting the rate at which transitions can occur between flow states. This would mean the output geometry is affected in a more convoluted manner by the viscosity of the melt in the nozzle, as well as the nozzle geometry. These discoveries, although not allowing for definite conclusions to be drawn as to the underlying processes, have served as an example of how in-situ monitoring can improve both the understanding of the process as it is occurring, and the need for adaptive control. Separating the effects of this turbulence and the other nozzle parameters could lead to improved final properties, combined with more consistent results, enabling the use of liquid crystal polymers for parts that require stiff lightweight structures. The combination of these properties with the wide operating temperature and relative inertness of the material would make this material and manufacturing system very attractive for future aerospace structures [13].

Contents

Nomenclature

Abbreviations

Abbreviation	Definition
FFF	Fused Filament Fabrication
FDM	Fused Deposition Modelling
LCP	Liquid Crystal Polymer
SBC	Single Board Computer
ADC	Analog to Digital Converter
PID	Proportional–Integral–Derivative
API	Application Programming Interface
EMI	Electro Magnetic Interference
PEI	Polyethylenimine
B&W	Black and White
ROI	Region of Interest
IQR	Inter Quartile Range
CS area	Cross-Sectional area

1 Introduction

1.1. Extrusion based manufacturing

Fused deposition modelling (FDM) as a process is relatively young, but has benefited from a major surge in popularity in the maker world starting with the open-source [2] reprap movement [15], which popularised the technology and made it more accessible, while also developing major improvements and standardizations that made FDM more reliable as a technology. This was however done with the technology available at the time, which consisted of arduino shields powered by the 8-bit microcontrollers of the arduino mega. As a result, despite improvements in speed and quality achieved by improved stepper drivers and firmware, the majority of printers still operate in an open loop manner, relying on stepper motors to accurately position the placement of material. The process of depositing material via the extruder has another open loop stepper, with PID controlled heaters on the nozzle and bed being some of the only closed loops in the average FDM machine. This open loop heritage has led to improvements focusing on controlling the inputs, such as ensuring the filament diameter is consistent and painstakingly keeping the filament moisture content as low as possible. Improved quality and tolerances are achieved by manually tuning parameters for each different type of filament, or in some cases even colour of filament as the additives affect the properties. In industry, where there is a greater interest in precise tolerancing, high volume production and reduced costs, the strategy for FDM has become extruding oversized parts using pellet extruders, and then swapping the toolhead from a FDM extruder to a mill to remove all excess material. This process comes at the cost of reduced possible complexity of parts, and increased waste of materials.

1.2. The need for in-situ monitoring

To maintain the advantages of FDM, and resolve the outstanding issues of tolerances and part quality, one of the most promising solutions is to use what is commonly referred to as in-situ monitoring. In-situ monitoring means monitoring of the part, process or machine using sensors while the part is being printed. This is in contrast to ex-situ monitoring, which is used to denote inspection of the part after it has been manufactured. Others have at times defined in-situ monitoring as strictly monitoring the part being printed, however in the context of this work we will be referring to it in its broader definition which includes the machine and even the inputs to the machine which are a part of the process. Manufacturing a part using FDM lends itself to in-situ monitoring as the part is built up in layers, which therefore means that a sensor with limited depth observation capabilities can still monitor the full depth of a final part, by simply combining multiple readings to build a three-dimensional database of information based off of two-dimensional readings. This is particularly key for the complex and potentially large parts that can be printed using this process, that might otherwise be very difficult to inspect, and therefore certify. Furthermore, by the very nature of the process, it is possible to obtain the data on any deviations from the nominal while they are occurring, allowing for corrective action to be taken during production. These potential benefits are the reason for a lot of effort in recent years towards developing these capabilities.

1.3. The state of in-situ monitoring in Additive Manufacturing

A number of works have focused on detecting errors in the machine, such as the use of acoustic emissions sensors by Wu et. al.[19] to detect filament runout and nozzle clogs. Many of these are aimed at detection with the aim of human intervention if a fault is detected. Other promising techniques include the work of Greeff et. al.[12], who used a camera to detect the actual filament speed exiting the extruder towards the nozzle. This could be used as in input parameter to tune the flow rate. This is where a distinction needs can be made between the capabilities of in-situ monitoring techniques; those that can be used to detect failure post occurrence, those that can predict an incoming failure, and those that by their nature can be used to tune the parameters and prevent failure. An example of a technique that is capable of tuning parameters and preventing failure is the work by Pattinson et. al. [6] [5]. Using an optical imaging camera mounted close to the nozzle, they demonstrated the ability to detect issues such as underextrusion and warping, and were generally able to either output new parameters with which the print would succeed upon being restarted, or change parameters on the fly to improve the output and prevent further issues. These impressive results were obtained in the crucial early layers of a print where the contrast between deposited material and the printbed is high, but others are actively working towards improving the image recognition by utilising digital twins [16], such that it would also work at later stages in the print. This could also open a new capability of in-situ monitoring that has so far rarely been touched upon; that of correcting a print after a detected problem has occurred, by detecting what damage the problem has caused, and how it can be rectified, thus ”repairing” a failure. Such a feat would require the failure detection previously mentioned, but also failure mapping, and then a means of generating a correcting method. The choice of sensors presented in these works are however often geared towards detecting a limited set of problems, which is why multiple sensors would need to be utilised together. This work will try to combine the data from multiple sensors to work towards this goal, utilising in-situ monitoring to improve our understanding of and ability to print a material which is of particular interest for solving the previously mentioned issues.

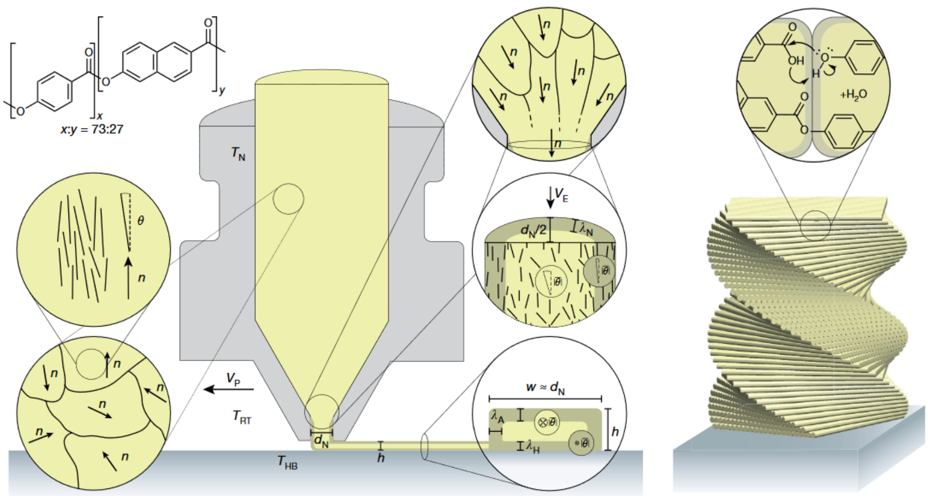

Figure 1.1: A visual represention of the nematic regions of LCP as they pass through the nozzle and are aligned during the extrusion process.[10]

1.4. Liquid Crystal Polymers as a material for fused deposition modelling

Liquid crystal polymers (LCP) are a commonly used in display technologies, but have recently begun to receive more attention for their mechanical properties. These polymers consist of stiff rods that form regions where the rods tend on average towards a particular alignment [10][9]. The ”crystal” refers to these regions, which are free to move and reorient themselves above a certain temperature, but can be fixed in place when cooled. By extruding the material above the melt temperature, these regions can be given a greater degree of alignment, which causes the bulk material to get a high degree of anisotropy, which significantly improves the properties of the material in the direction of alignment, as illustrated in figure 1.1. This enables this material behave similar to a fibre-matrix composite, with the high specific strength and anisotropy commonly associated with those materials, while being made of a single material. It is also uncommonly unreactive for a polymer, which, combined with the low hydrogen porosity, makes it an ideal candidate for example, for making hydrogen tanks, a key technological hurdle in the energy transition [1][17][3]. When additively manufactured, the shear applied to the material in the nozzle aligns the regions, and allows complicated parts such as hydrogen tanks to be made with a much smaller radius of curvature compared to conventional fibre-matrix composites, and at the end of the parts service life, it can be recycled into new filament to be printed again. These are very compelling reasons for the development of its use as an additively manufactured material, however it still requires some development to be used in this manner.

1.5. The requirements of printing liquid crystal polymers

The material, when extruded, is unlike other more commonly printer materials. The outer edge of the material, once deposited, cools the fastest, and keeps most if not all the alignment it was given within the nozzle. This is beneficial for the properties, however it is detrimental to the layer adhesion of the material. The highly aligned regions will not easily mix with and therefore bond to another highly aligned regions deposited on it, as the stiff rods do not rapidly lose their alignment and mix in the same manner as polymer chains do when a more standard thermoplastic polymer will when two pieces are joined at sufficiently high interfacial temperatures to produce a relatively homogeneous mass. This behaviour is in stark contrast to the core of the deposted material, which stays above its melt temperature long enough to lose the alignment gained in the nozzle, and hence become less anisotropic, forming a core-shell structure of a highly anisotropic outer layer that will not easily bond to other deposited material, and a weaker more isotropic internal core. Ideally, this would be reversed, with a highly anisotropic core, surrounded by a more isotropic outer shell ready to bond to any other material deposited alongside it. To improve the bonding, a high enough temperature to allow the stiff rods to interlock and enough contact pressure to encourage motion of the stiff rods is required, ideally without heating the entire outer shell above the melting point and losing the anisotropy and the associated beneficial properties. Therefore, it is imperative that the deposited lines are stacked closely and pressed together during the print process to obtain the best conditions for the bonding of two lines, while controlling the parameters to maintain a high degree of anisotropy in the majority of the printed part.

1.6. State of the art of liquid crystal polymer printing

The parameters that govern this anisotropy were investigated by Houriet et. al.[13]. In this work, it was discovered that reducing the nozzle diameter to a very small diameter, along with printing at a sub-diameter line width was critical to obtained the higher Young's modulus and ultimate tensile strengths of the material. This would also create the greatest amount of inter-layer area within a given sized part, meaning that the layer adhesion issue becomes far more prevalent, and the material is more likely to delaminate, and for cracks to propagate along these boundaries between lines and layers. To this end, their work also explored how to use the example of trees to tailor the properties of the local material to stop crack propagation, and distribute loads at critical junctions to avoid stress concentrations. This ground-breaking research found that by utilising the link between the linewidth and the degree of anisotropy, it was possible to tailor the material properties locally for optimum load distribution within the final part. This does however require actively controlling and changing the linewidth, in a manner that does not leave voids or gaps between lines and maintains as high an interlaminar bond as possible, all the while making complex motion patterns to achieve the geometry recommended by their work in wood-inspired topology optimisation. The goal of this work will be to try to leverage combining in-situ monitoring sensors, and with the help of ex-situ sensors, attempt to gain a greater understanding of the parameters that define the linewidth, and generate a database of in-situ collected data, verified using ex-situ sensors, that could be used to create a control loop to actively monitor the linewidth and predict the parameter changes required to obtain the desired linewidth and linewidth changes required.

1.7. Research Questions

The primary aim of this work is to combine data from sensors mounted on the printer itself and standalone sensors, to evaluate how it works, the problems that may arise, and how best to approach it. Hence, the primary research questions is;
What are the current limitations preventing the use of multiple sensors for direct real time monitoring of the printing process?
In order to further demonstrate the usefulness of in-situ monitoring, this work will vary the linewidth during a print to try and characterise the transition between the two linewidths. This is done with different print parameters to see how they affect this transition and to lay a foundation for what parameters can be changed to achieve a specific result;
What is the effect of print parameters on the extrusion geometry on a time basis?

2

Experimental Setup & Methodology

2.1. Prusa Mk3S/+

A Prusa Mk3S/+ model of 3D printer was chosen for the experiments. This model of printer is well suited for the task as the hardware and firmware are both open source, making it fairly easy to modify when required. It uses the colloquially termed "bed slinger" motion system means that the extruder and nozzle assembly only move in the X-Z plane, and are easily accessible for modifications and observation, and leave a lot of space in the Y direction for modifications. This ability to make modifications was crucial for the experiments. In order to enable it to print at the temperatures required for LCP, the thermistor was replaced by a PT1000 thermistor, and the firmware was adjusted accordingly. To ensure the structure mounting the extruder was capable of resisting the temperatures, it was reprinted out of polycarbonate.

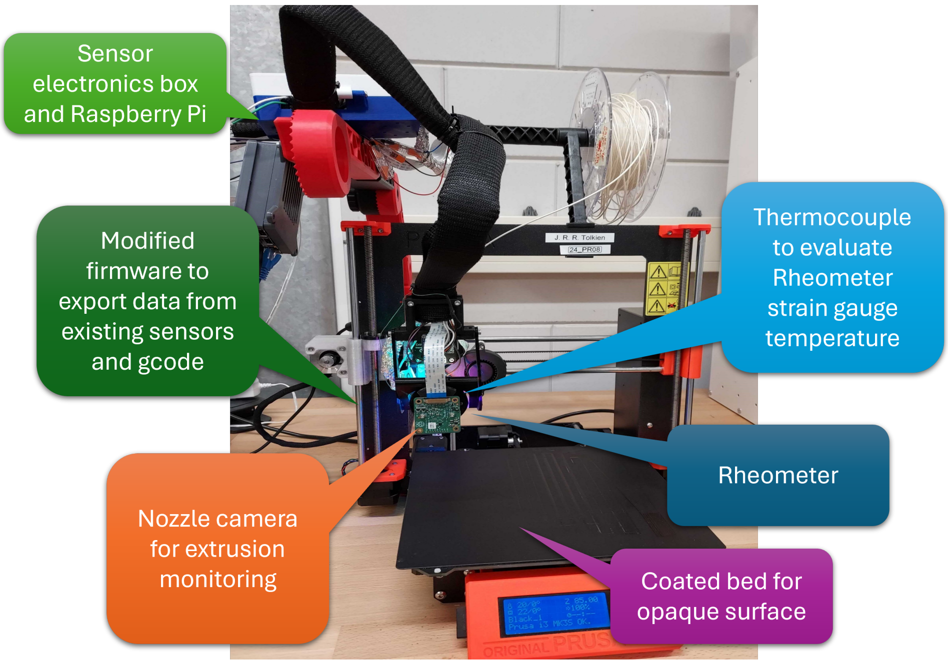

Figure 2.1: An overview of a Prusa Mk3S/+ printer showing

In order to try and sync the datapoints across the in-situ and ex-situ measurements, it was critical to have position data for the in-situ monitoring sensors. As the printer does not have this information by default, and therefore cannot communicate it, an alternative system had to be devised. This took shape in the form of using the G-code M118 command, which serial prints a command. Receiving these serial commands was Octoprint. Octoprint was installed on a raspberry Pi 3B+, which was used as both an interface for the Prusa, and a data gathering tool. It is was also used as a controller for the lighting of the optical setup, by using the general purpose input/output (GPIO) pins of the Raspberry Pi. Octoprint is able to store G-code files, and start and stop prints, communicating directly with the firmware of the Prusa. It does this by hosting a local web interface, as well as having an Application Programming Interface (API). This API allowed a laptop running a python script to receive data from the Prusa firmware, which itself was pre-programmed M118 G-code commands. Octoprint acts as a relay, drip feeding it G-code from its own internal storage, and transmitting serial communication to via the API to the python script gathering all the data. The raspberry Pi was mounted in a grey box, as seen in the top left of figure 2.1.

The M118 g-code command is a standard command to print text to serial, which in this case was preprogramemd to send the position according to the g-code of the printhead at that time via g-code. This position would be timestamped and matched with the sensor data. This serial print was transferred to the Octoprint API from where it was accessed and recorded. To ensure this timestamp was accurate, the firmware of the printer was modified to process the M118 command sequentially, as opposed to it being processed as soon as it entered the buffer of the printer. This meant that the printer would be at the location of the previously defined target position which was reported in the M118 command.

2.2. Rheometer

One of the primary sensors used in this work was a rheometer, based on the work by Coogan et. al.[4]. It consisted of a thermocouple temperature probe, mounted at the tip of a pushrod, which entered the nozzle through a slipfit hole. This would allow the pressure inside the nozzle of the melt to be transferred by the rod onto a strain gauge assembly at the opposite end. This combination of the temperature and pressure of the melt inside the nozzle was used by coogan to calculate the rheological conditions inside the nozzle.

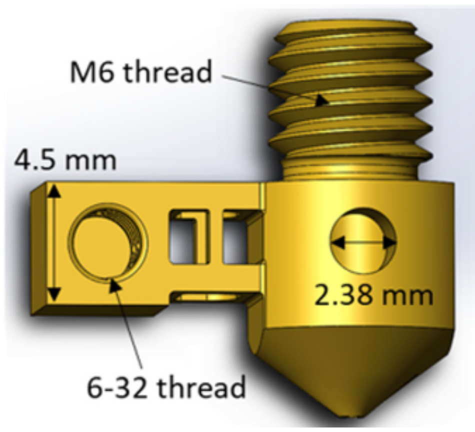

Figure 2.2: External view of the nozzle showing the hole for the pushrod

For use on our printer, the nozzle was slightly modified, such as using metric standard bolts, and changing to a nozzle designed for 1.75mm diameter filament as is standard on the type of printer in use. Furthermore, a more common 0.4mm nozzle diameter was used, as it better represented the standard nozzle size in use by this type of printer. To ensure the strain gauge would not catch on the bed, a small piece of teflon sheet was mounted to the bottom of the strain gauge to enable it to slide.

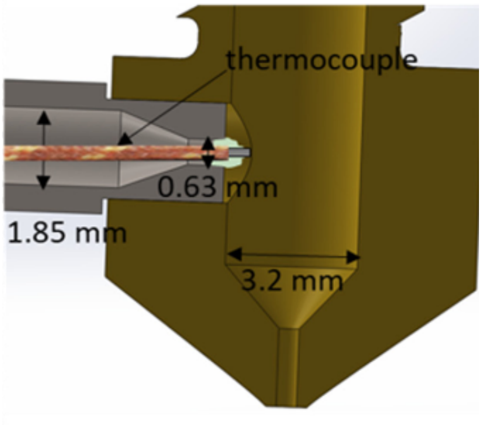

Figure 2.3: Internal view of the nozzle showing the pushrod and the thermocouple

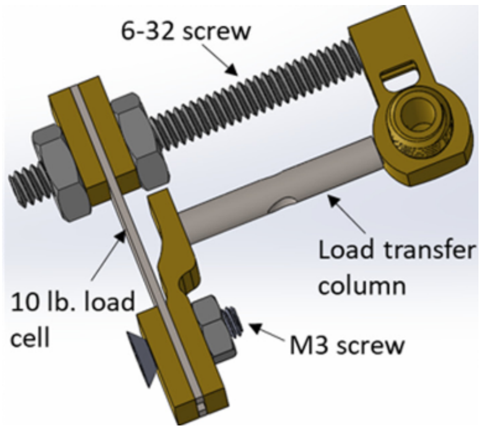

Figure 2.4: Assembled view of the rheometer

Due to the sensitivity of the strain gauge analog signal to electromagnetic interference (EMI), the wires were wrapped in aluminium foil, and this foil was in turn grounded to the same ground as the analog-to-digital converter (ADC). This reduced the effect of EMI on the readings, and resulted in an appreciable decrease in noise in the pressure data.

2.3. Imaging Camera

To be able to observe the material coming out of the nozzle, and generate a database of extrusion images that could be matches to sensor readings and print parameters, a camera was mounted on the extruder pointed at and focused on the material as it was leaving the nozzle. The camera chosen was a raspberry Pi HQ camera with a 25mm telephoto lens as seen in figure 2.6. This provided manual focus and aperture adjustment, specifically on objects that were relatively close to the camera.

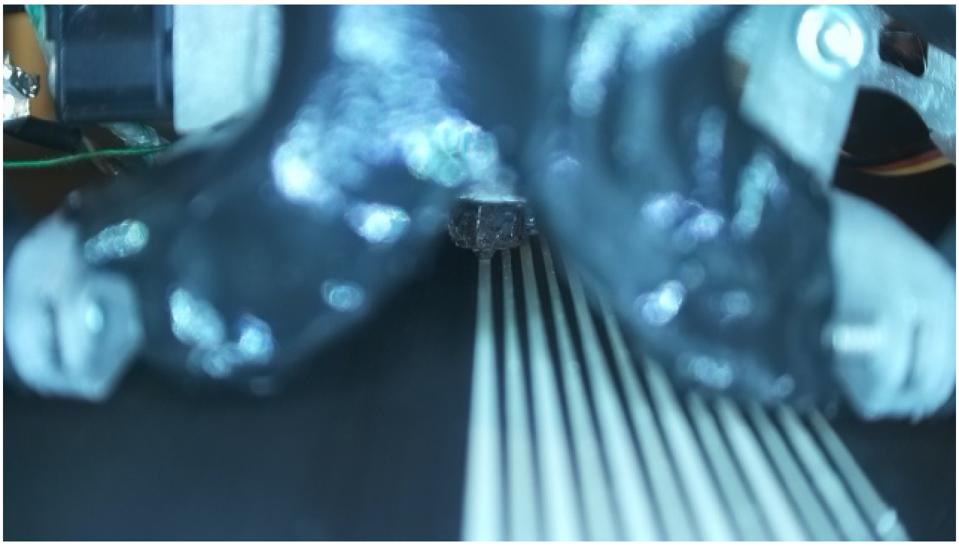

Figure 2.5: The view of the nozzle camera, focused on the material leaving the nozzle.

The mount on the extruder enabled it to move with the nozzle, and prevented it from colliding with the other sensors or having its vision obscured. It also enabled the mounting of a ringlight around the lens of the nozzle to get a consistent level of lighting on the nozzle and the material being deposited.

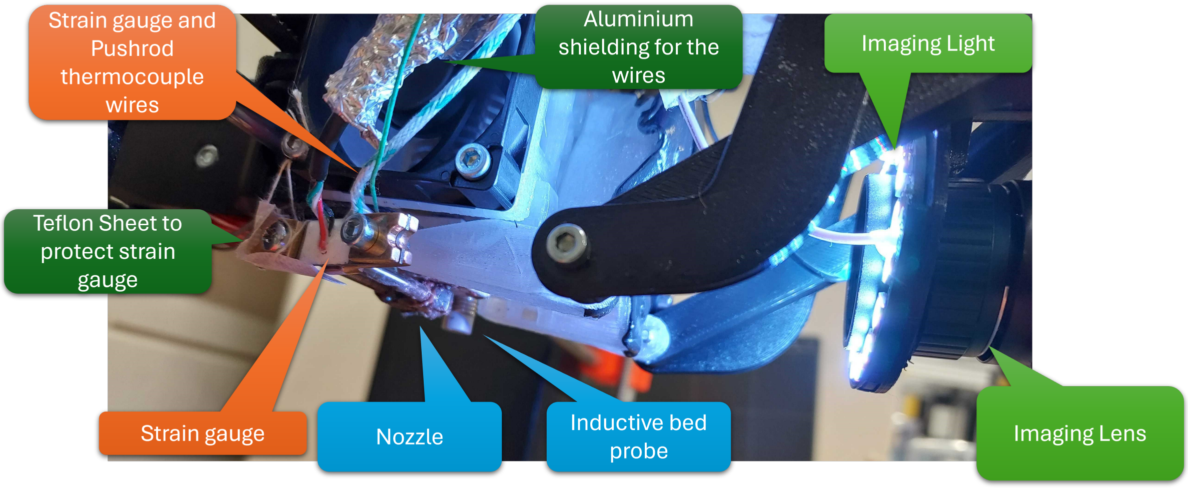

Figure 2.6: Image of the extruder setup with the rheometer and camera sensors.

2.4. Keyence VR-5000 Microscope

The only ex-situ monitoring sensor was the Keyence VR-5000 Microscope. This is a microscope that can take full colour images and generate heightmaps, by scanning and stitching together multiple smaller scan area. This was used as a reliable reference to precisely measure the correct lineheight and linewidth of the samples being printed. It works using a fixed head and a moving stage that will automatically move to map out and then scan the areas of interest designated for scanning by the user.

Figure 2.7: Image of the Keyence microscope. Note the black moving stage, upon which the printbed was placed to scan the sample. This microscope was used to scan the sample with 12x magnification, and was able to automatically move the stage to scan a large area that it would then stitch together into a single image.

Although it is an ex-situ sensor, the in-situ sensors such as the camera were difficult to calibrate, and the image recognition aspect of any optical in-situ measurement technique would require an external sensor to provide the data to train it on. This is the function provided by the keyence microscope, with the goal of training the model and then confirming the model, such that the in-situ sensors can actively provide similar information in a real time manner and replace this ex-situ sensor.

To perform the measurements, the spring steel PEI bed of the Prusa would simply be removed from the Prusa machine and placed on the scanning table of the microscope. The entire area containing the samples were scanned in two sets of 8 images. The bed would cool well below the 90 degree celsius printing temperature in the time it took to set up and scan the samples with the microscope, which could lead to the samples detaching from the bed, so a glue spray was added to the bed to ensure adhesion for a sufficient amount of time for the bed to be scanned under the microscope.

2.5. Sample Lines

To obtain the desired information, 18 experiments were set up. To ensure the repeatability of results and have a basis for comparison, each of these experiments followed the exact same print path on the exact same location on the bed. The geometry consisted of a zig-zag pattern, with 8 lines going one way, and 8 lines going the other way. One set of those 8 lines were the sample lines, each of which had a length of 155 mm. If this was a sample with varying linewidth, then 20mm along this line, the linewidth would suddenly change, for a length of 120mm, before returning back to the initial linewidth.

This zig-zag pattern was printed from right to left on the print bed in the X direction, with the back and forth motion of the zig-zags occuring in the Y direction, such that the samples lines were printed in towards the positie Y direction. The sample lines consisted of the lines where the bed was moving past the nozzle towards the camera, ensuring that the camera had a clear unobstructed view of the material as it left the nozzle. The nozzle would then continue extruding as it made a sharp corner, and moved back in the opposite direction, printing at the minimum line width for the full length, such that

Values in cells: Feedrate (mm/min)		Temperatures (degrees C) 295	310	325	340
Linewidths (mm)	0.6-1.2	200	200	200 360 540 720	200
	0.6-0.6	200	200	200	200
	1.2-1.2	200	200	200	200
	0.6-0.9			200	
	0.9-1.2			200	
	0.9-0.9			200	

Table 2.1: An overview of each of the experiments that were carried out and the variables that were changed. The values inside the cells each represent a feedrate for which an experiment with that temperature and those linewidths were carried out.

the pressure in the nozzle would have the time to return to a steady state. The line would then make a smooth 180 degree curve towards the start of the sample line again. This is because the right angled changes caused pressure spikes due to the sudden overlap areas, compared to the minimal and very gradual overlap of a curve.

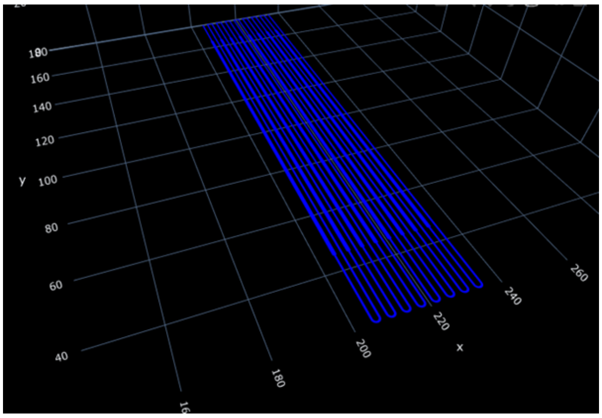

Figure 2.8: 3D render of the g-code showing the print path on the bed, with the rounded corners at the start of the sample lines and the thicker sections where the linewidth changed.

To generate this g-code, a python script called fullcontrol[11] was utilised. This python script simplified the process of generating g-code compared to manually programming it, while allowing the greater control over the geometry and the parameters compared to conventional slicers that was required for the generation of the g-code. A good example of this is the inbuilt rendering feature to visually confirm the planned path as seen in figure 2.8. It furthermore allowed the g-code to be written as short motion steps of half a millimetre, separated by the M118 g-code command, which ensured that at any point in time the location of the nozzle was known to within half a millimetre accuracy. This was a key requirement to be able to combine the data from the previously mentioned sensors.

In order to generate the g-code, the python script was programmed to accept as input parameters the variables, such that all the experiment g-codes could be generated with the exact same parameters except for those that were deliberately changed. An overview of these parameters is presented in table 2.1. These experiments were chosen such that the effects of specific parameters could be isolated, and therefore, in theory, the effect of the combination of these parameters could be found by extrapolation.

3

Data Processing

3.1. Data from the Printer and the Rheometer

The data from the pressure sensor of the rheometer was found to be affected by the temperature, due to the proximity of the strain gauge to the bed and nozzle. Although it was still within the operating temperature, it was outside of the compensated temperature. This caused a gradual shift of the readings over time as the sensor heated up. As a result, a baseline correction had to be applied to the pressure data which the strain gauge measured. This was done using the scipy [18] and pybaselines [7] modules of python. Specifically, by first smoothing the data using a uniform 1D filter, before using the imodpoly function with a polynomial order of 3 and a standard deviation of 0.7 to generated a smoothed baseline. The imodpoly function uses the polynomial order and standard deviation to find a line that minimises the least squares to find a baseline that best represents the baseline of a data set. This baselines was subtracted from the actual data in order to correct for the strain-gauge-temperature-induced drift of the pressure values.

The data of each experiment was separated into the individual sample lines. This was done by separating the data by the X position and storing each the data from line as it went along the Y axis. The messages from octoprint were limited to a 2 Hz frequency, which is a lower rate than both the sampling frequency of the other sensors, and the rate at which M118 g-code commands were sent. This meant that all this data matches to an error of approximately half a second. All the data between these readings was averaged to a single reading to match the octoprint data for the purposes of combining this data with later data from other sensors, which requires the position data from octoprint. Any future work which does not require this position data, can still use the generated dataset and the corresponding images, which were captured at a frequency of approximately 12 Hz. Each line was saved in a separate .csv file, with the corrections applied, to be combined with the microscopy data. This process is shown in 3.1.

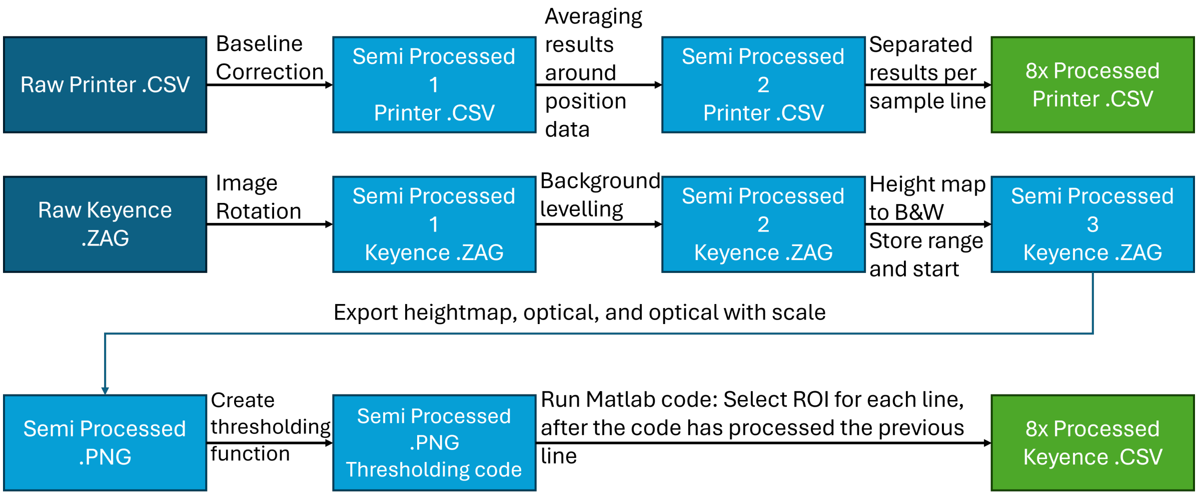

Figure 3.1: High level overview of the steps taken to process the raw data.

3.2. Data from the Keyence VR-5000 Microscope

The data from the Keyence Microscope consisted of images, that first needed to be manipulated and processed in the Keyence software, before being exported and further processed in Matlab. Due to the temperature differential of the cooling bed as it was taken off of the printer and moved onto the microscope base, the printbed tended to warp. Therefore, as seen in figure 3.1, after rotating the image such that the printed lines appeared perfectly horizontal in the image, the software was used to correct for the curvature by using the profile of the print bed between the lines to "flatten" the base.

Following this, a heightmap image was generated, in grayscale, where the intensity of the white pixel was equivalent to a physical height, with the translation values stored for each image. This was done to every image individually, as each image had a different 0 point and height spread, hence the intensity, measured from 0 to 255, would translate to a different height depending on the image. The Keyence software was used to output 3 images in .png format; two optical images, one of which included a scale, and a single grayscale heightmap, which were then used for further processing.

These images were imported into MATLAB [14], with the scale bar being identified and used to find the location of each pixel in the image in mm, while the optical image was used to generate a mask that was used to calculate the width of each line. To generate this mask, each image threshold had to be manually calibrated in MATLAB using the MATLAB app Colour Thresholder [14], to ensure that the slightly translucent white sample lines could be clearly differentiated from the darker black printbed. This mask was also used to differentiate between the bed and the sample lines in the grayscale heightmap, such that the height of the bed could be subtracted from the height of the lines, to find the actual height of each pixel of each line. To find the location of each pixel, the sharp corners at the end of each line was used, which represented the Y position of 180. The exactly location should coincide with the centre of the sample line and the line leading out. This could then be used with the extracted scale to find the location of each pixel along the length of the line.

The script would run from left to right; from the end of the sample line towards the start of the sample line. Using the mask to differentiate between line and background, each pixel of this slice of the sample line which contained the sample would have its pixel intensities converted into a corresponding height. This height would then be averaged for the vertical slice of the sample line, giving the average height of that cross section of the line. This average height was then used with the linewidth to calculate the cross-sectional area.

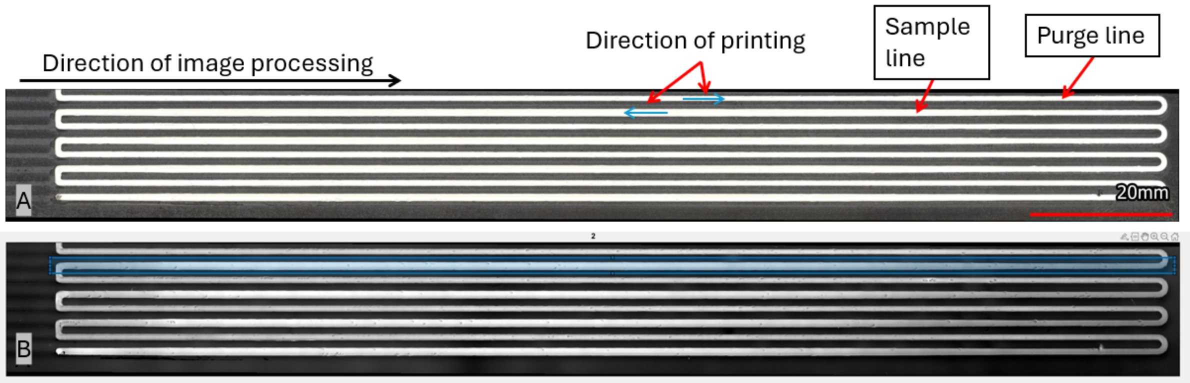

Figure 3.2: A. Optical image showing the scale bar.
B. Selection of the region of interest of a single sample line of the heightmap.

Each line was separated out by selecting a region of interest, as shown in figure 3.2, where the code would analyse the line present, and store it. To correctly obtain the data, only one line could be selected, and each line had to be selected in the correct sequence to be attributed the correct number. With the understanding of which pixels were part of the line, the pixel to mm scale, and their intensity value which was converted to a height value, all the data for the analysis was obtained and .csv file would then be written for each line, containing the data for each line, as well as identifying which line it was.

3.3. Fusing the Data

The data from the printer and the rheometer are stored with fixed timestamps and images corresponding to each datapoint. However, the data from the microscope does not have any of these. The only way to combine this data is to do so as a basis of the Y position on the bed, as that is a datapoint that all the datasets have in common.

3.4. Extracting the Slope

As the goal is to be able to vary the linewidths predictably and continuously, it is necessary to accurately examine the cause of the delay in the changing of the linewidth. To this end, a database of all the lines was created, with all the sample lines that featured clear visible anomalies being removed from this database. The data from the rest of the lines was averaged per experiment, to be used for further processing.

To analyse the response to a change in linewidth, we are primarily interested in the duration of the delay in the response, in this case the length of the slope of the actual linewidth following the set change. This length is inherently related to the magnitude of the change in linewidth, and so the gradient of this slope is also extracted to serve as a comparison basis for experiments where the magnitude differed. This was done for the cross sectional area data, as well as the linewidth data, as the lineheight was found to unexpectedly vary.

To calculate the slope gradient and length, the region which comprised the slope was first separated out. This was done by finding the regions where no change occurred; the plateaus corresponding to the region where no linewidth change had occured yet, and the region where the linewidth change had finished and stabilised. The slope consists of the region between the end of the first, and the start of the second. A simple straight line was plotted on an empirical basis of it seeming to best fit the data, and the gradient of this line, as well as the length of this transition region as measured in distance along the sample line was calculated and stored.

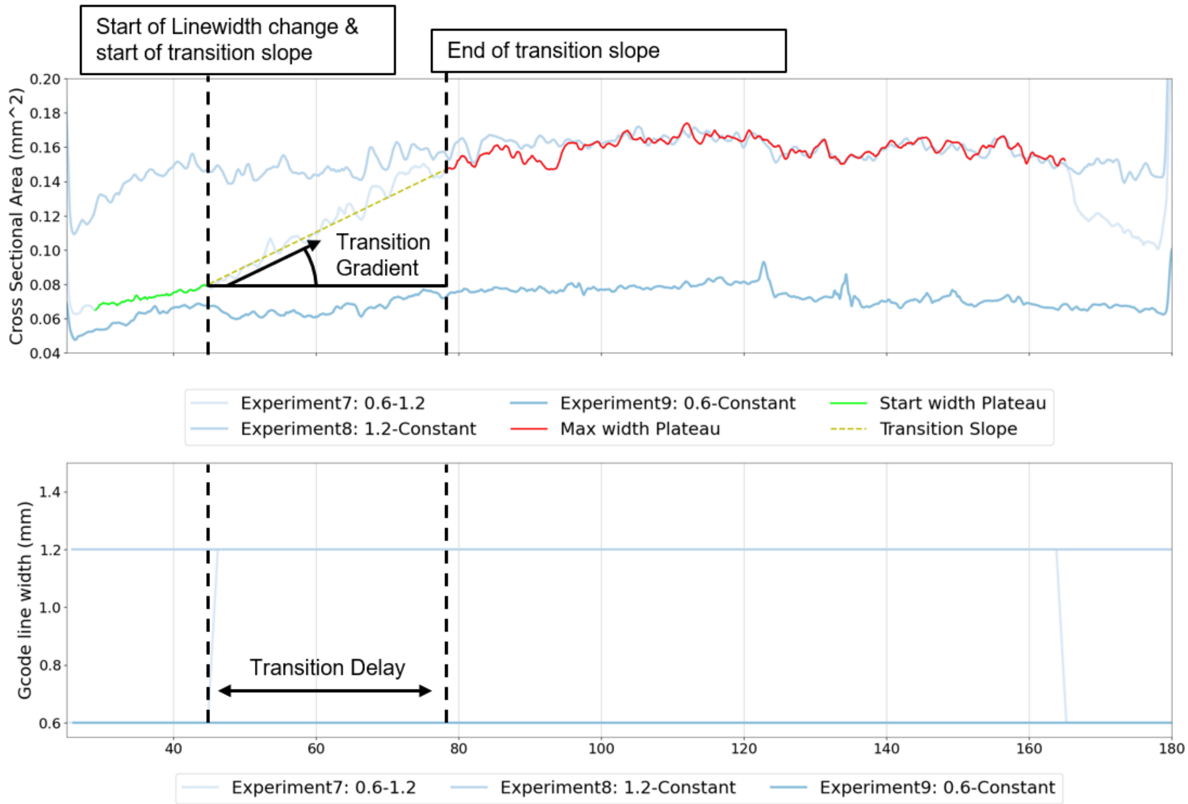

Figure 3.3: Highlighted representation showing the extracted transition delay and gradient.

4 Results and Discussions

4.1. The Pressure Data

One of the key aims of this sensor setup, beyond combining sensors, was to use the combined data to learn more about how the print parameters affect the process, and the final output in the form of the geometry of the output. The combination of this data was successful, with the data from the various sensors being matched to a specific location on the printbed. However not all sensors gave us the data we hoped for.

The rate at which material is extruded from the nozzle can be assumed to be at a steady state and based on the rate at which filament is being fed to the nozzle by the extruder. However, this is only true when continuously extruding and pushing the same amount of material into the nozzle. As mentioned in chapter 1, the work by Houriet et.al. [13] has demonstrated that the linewidth is an important parameter to actively change. This means that the steady state assumption is no longer valid. To then obtain the desired linewidth, in a controlled variation, we need to look at the pressure in the nozzle. This pressure is expected to be a good indicator or predictor of the volume of material that is extruded, and hence it was hoped to be able to combine its data with both the nozzle camera imagery to correctly measure this in situ, with the microscope data used for validation.

Unfortunately, it seems that the rheometer's pressure probe did not function as intended. Except for experiments carried out at the lowest temperature from the experiments, 295 degrees Celsius, and the experiments carried out at the higher feedrates, such as seen in figure 4.1, the pressure data shows a mostly flat, constant line. There are a number of potential reasons for these results. To explain these reasons. we need to first discuss the pressure probe setup and the issues encountered.

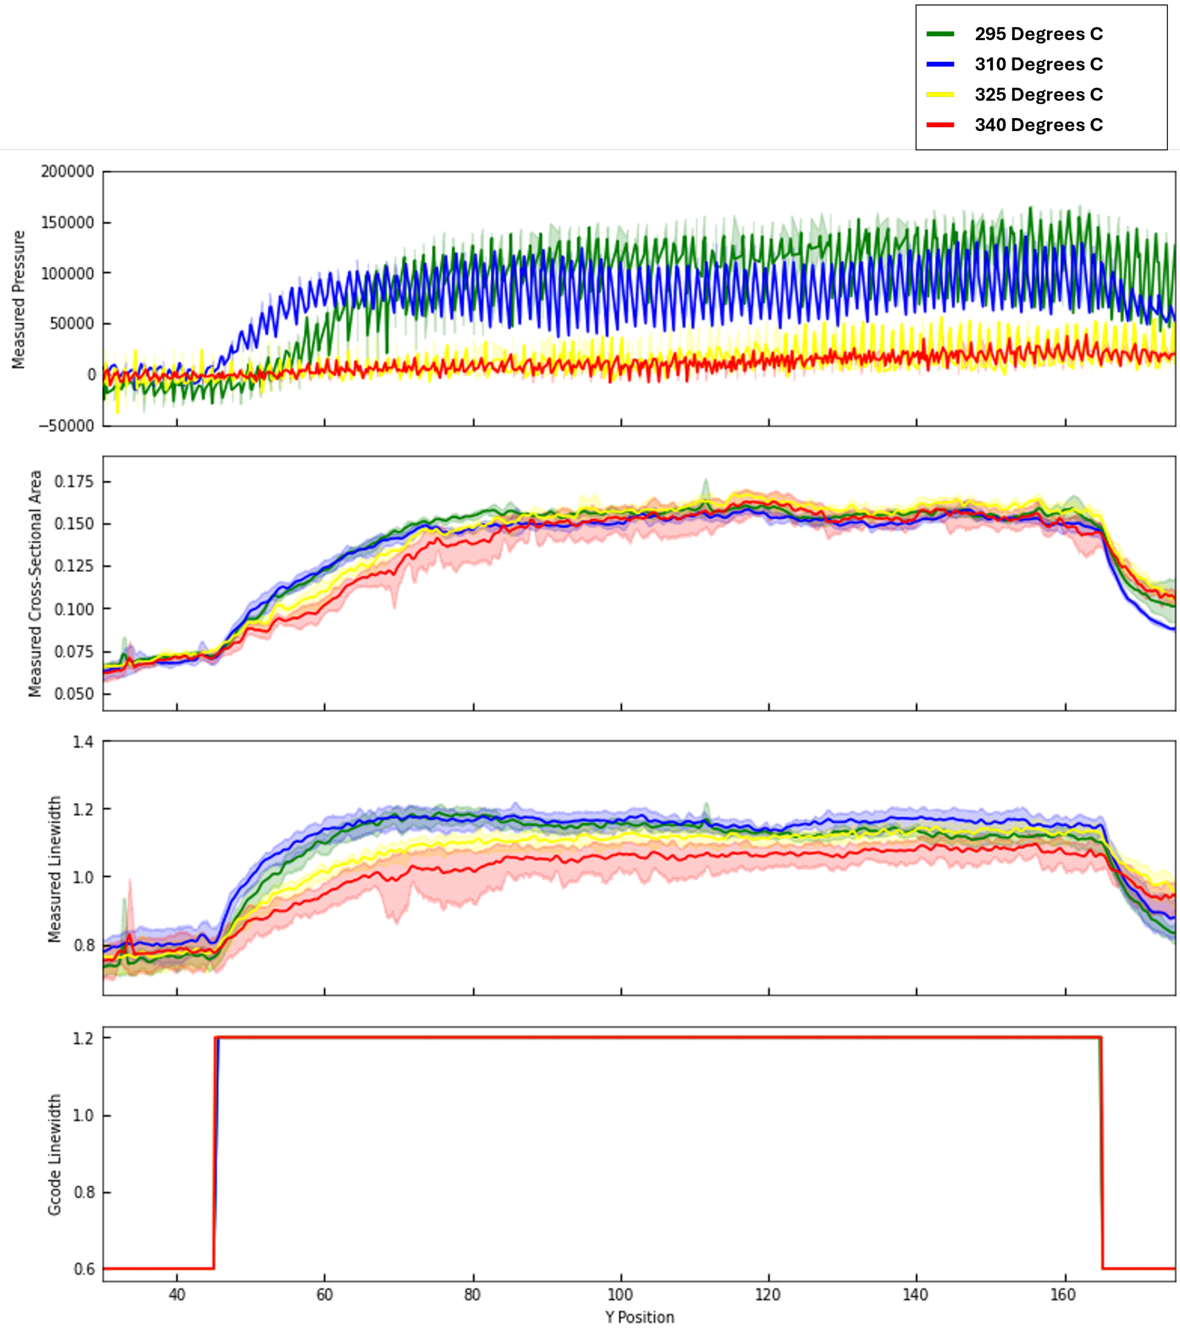

Figure 4.1: Example of the data from multiple separate experiments with different temperatures, showing how the pressure was not always properly recorded at higher temperatures. The spikes in the pressure data is due to an measurement issue which is further discussed in 6. The recorded pressure data that does not stay at 0 shows a similar increase and decrease as the the linewidth and lineheight.

4.1.1. Positioning of the strain gauge

During assembly, it was found that the pressure readings were affected by the position of the strain gauge along the rod. Despite using an M3 bolt, a quarter turn of the nut fixing the position of the strain gauge relative to the nozzle and pushrod would cause a large change in the recorded pressure. Keeping it too loose would create a gap between the pushrod and the strain gauge, which would adversely affect the flow due to allowing the pushrod to create a hole in the inner walls of the nozzle, and was observed to prevent the recording of any readings.

Conversely, applying a degree of pre-tension to the pushrod was seen to enable measurements to be recorded. Low pressure measurements however returned a flat line, or flat sections, similar to a low pass filter. This can be explained by the pressure needing to exceed the pre-tension that was applied by the location of the strain gauge to be measured. Using minute adjustments, the pre-tension was adjusted using the nuts fixing the location of the strain gauge. The final position applied the highest amount of pre-tension that still gave readings and allowed for the lower pressures of the linewidth variations to be seen.

It must be noted that the tests where the position was adjusted was done at the lower temperature of 295 degrees Celsius. The adjustments were carried out with the material was loaded in the nozzle, and involved keeping the nozzle at the temperature without extruding for long durations. However at higher temperatures the material was seen to change over these longer durations. These changes consisted of the material taking on a darker brown hue, and was accompanied by the material becoming more more viscous and sticky.

The result of this was that an unknown level of pre-tension was applied to the strain gauge. Unfortunately, it was not possible to precisely measure the pressure inside the nozzle for calibration, so the pressure readings are not absolute, and the exact pressure is unknown. It was demonstrated however, that the slope of the voltage of the strain gauge varied consistently linearly with applied load. For this reason, although the exact value of the pressure could not be calibrated, the change in pressure could still accurately be measured. This still provides a lot of information from an in-situ monitoring perspective.

4.1.2. Material Leakage

The pushrod probe hole in the nozzle was found to leak material when printing, which forms a ring around the pushrod and on the nozzle. This build-up of material would undergo several heating and cooling cycles when carrying out multiple prints. Therefore, it is suspected that this build-up of material would increase the friction forces acting against the movement of the pushrod. This would restrict or at least delay movement of the pushrod, exacerbating the issues previously mentioned, especially if a gap was left between the pushrod and the strain gauge.

To minimise the effect of the leaking material, prior to every experiment, the accumulated build-up of material that had leaked past the nozzle was removed. Nevertheless, some material was stuck inside the probe hole, between the nozzle and the pushrod. This material could not be removed, and would be pushed out to build up around the pushrod again fairly quickly when printing. This means that the resulting motion restriction potentially still affected the pressure readings by the end of each experiment.

4.1.3. Validity and Cause of the readings

The experiments where a pressure response was obtained show a similar pattern between the extruded geometry and the pressure data as seen in figure 4.1. The extruded geometry still follows this pattern even when no pressure response was recorded. From this, we can conclude that the pressure variation in the nozzle was not correctly measured.

No similar issues were mentioned in the work of Coogan et. al. [4]. However they used a differently shaped nozzle, designed for ”3mm” filament, which is how the 2.85mm filament standard was commonly referred to, and they used different materials printed a lower temperatures and higher feedrates. This meant that their nozzle was larger, reducing any interference from the pressure probe, and their extrusion pressure was likely higher.

The cause of the lack of pressure measurements is likely a combination of the material leakage and the positioning of the strain gauge, which increased the minimum pressure that the sensor could measure. The initial observations where data was recorded were carried out at lower temperature, where the viscosity of the material is lower, and therefore a higher pressure is obtained during printing. This higher pressure is also the case for the experiments carried out at higher flow rates. The low pressure during the process, combined with the higher minimum pressure for recording data, meant that the sensor was effectively used just outside its measurement range.

The result of these issues is that the pressure data from the pressure sensor was not useable for the majority of the experiments. Instead, moving forward on the assumption that the pressure inside the nozzle directly relates to the volume of material being extruded, the cross-sectional area, lineheight and linewidth, as measured based on the microscopy data, was used as an analogue to the pressure, and used for the analysis and the plots presented in the rest of this chapter.

4.2. Reading the data

An example of a set of graphs of the processed line data of a single experiment is shown in figure 4.2. The data plotted in the first, second and third graphs represents the cross-sectional area, the linewidth, and the lineheight data, as measured by the Keyence optical profiler microscope. The dark line represents the median, while the shaded region around it represents the inter-quartile range (IQR). This is the data used for the analysis of the slopes instead of the pressure data.

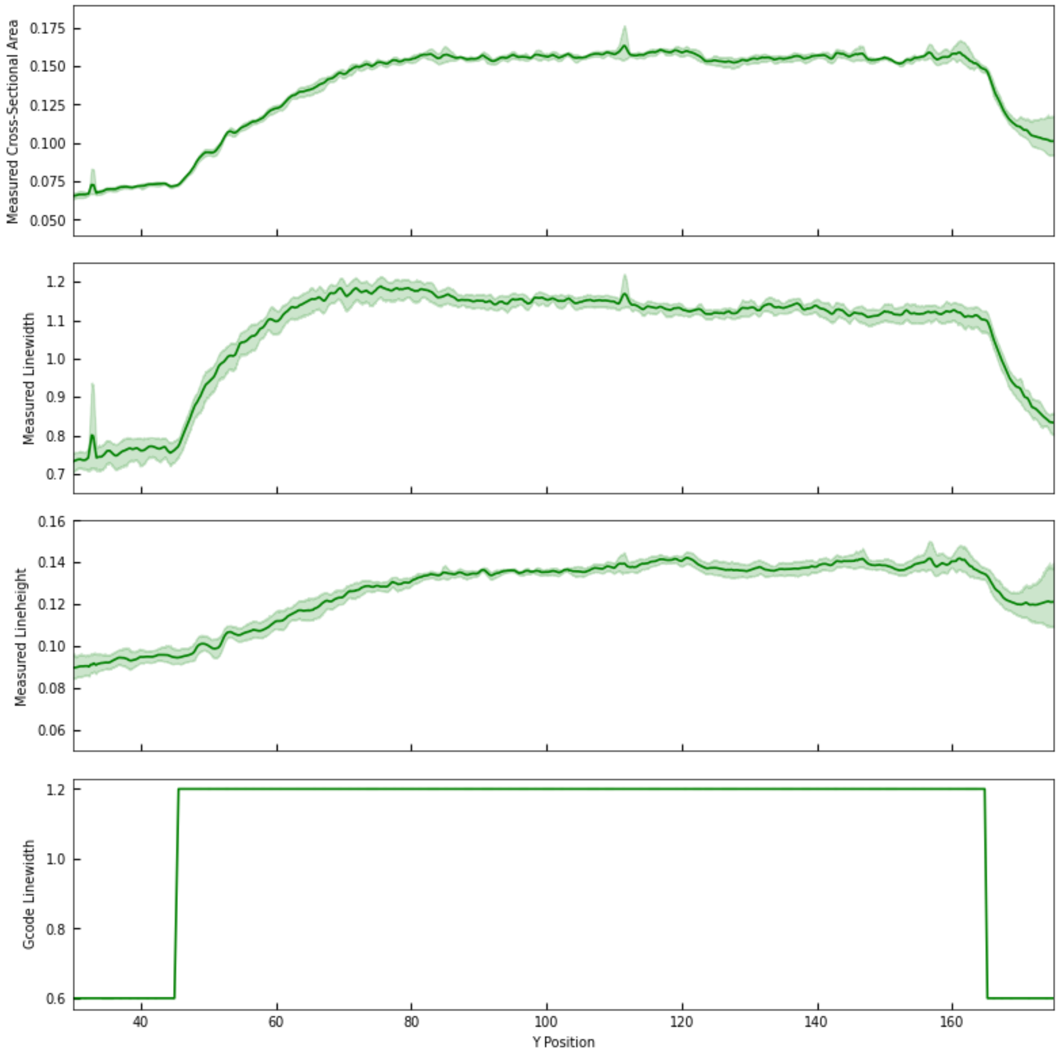

Figure 4.2: Example graph of the data set for a single experiment. The top graph is the cross sectional area, the second is the linewidth, the third is the lineheight, while the fourth is the linewidth as set by the g-code command. The dark line is the mean, while the shaded area represents the inter-quartile range. Note that when plotting multiple experiments together, the lines will be colour-coded to match.

The last plot shows the linewidth as set out by the G-code M118 commands. Due to the 2 Hz sampling rate of this data, it may not perfectly match up time wise, hence some of the plots will show a slight discrepancy between experiments or sample lines. Furthermore, the change in linewidth shows as a slope, this is also due to the sample rate. The command to change the linewidth was issued once, at the Y location of 45, which corresponds to a location between the start end of the slope seen in the linewidth plots.

4.3. Consistency of Data

4.3.1. Selection of Data

Prior to selecting the data for further analysis, the images of the lines were visually inspected. The lines that had clear issues, such as for example where material that was stuck to the nozzle had gotten caught on the line, were removed. These could be removed as their extrusion issues were due to external factors, such as material sticking to the nozzle, and not the print parameters themselves. The measured linewidths, lineheights, and cross-sectional area are quite consistent once we remove the obvious outliers. This is a positive sign, as it would suggest that flow conditions in the nozzle are fairly steady and therefore predictable. This would also mean that to obtain a correct and consistent extrusion geometry, no rapid changes in parameters are likely to be needed.

4.3.2. Variations in the Lineheight

One of the most unexpected finds of this work was that the lineheight of the printed lines was not constant as had originally been assumed and planned for. This parameter is generally expected to be constant due to the ironing effect of the nozzle normally limiting the height of the line. However it is possible that the pressure of the extruded material is high enough to have it push up behind the nozzle similar to a wake wave. This demonstration of its variability, as well as the loss of the pressure data, means that the microscopy data has become much more important than simply verifying the data from the other sensors. Incidentally this also perfectly illustrates the benefit of having redundancy in the sensor suite.

The magnitude of the variation in line height along the length of the sample line is small at up to 70 microns. Relative to the set layer height of 150 microns however, it is significant, and therefore cannot be ignored. Upon initially observing what was thought to be a deviation in lineheight, attempts were made to correct it. It was first assumed that the z-offset, which has to be manually calibrated, was incorrect, but adjusting it seems to have inconsistent results, which were confirmed using a manual micrometer. Upon inspecting the lineheight profile in the Keyence software, it was then seen to vary along the line length. Initial theories were that the pressure of the melt was pushing the nozzle upwards, as the entire extruder assembly is cantilevered off of two smooth rods which are only fixed from rotating at their extremities. It was therefore concluded that it was not controllable, but should instead simply be measured as a parameter, and future printers would need to be stiffened.

4.3.3. Consequences of the lineheight variations

One of the major consequences of the lineheight is that it makes linking this work with that of Houriet et. al.'s [13] work more difficult. The material that is higher than the lineheight is material that is taken away from the linewidth. This means the two cannot be considered separately, as was done in their work. Furthermore, to evaluate the effect of the parameters on the extrusion and its delay, it is now necessary to look at cross-sectional area to investigate the volume of material extruded. However, to evaluate the results in context with other work, the linewidth still bears investigating. By analysing both the linewidth and cross-sectional (CS) area, we can compare their relative changes, which may shed light on the cause of the changing lineheight.

The next sections will discuss the transition rate and delay of both the cross-sectional area and the linewidth, while also plotting the graph of these variables together with the lineheight and the G-code-set linewidth.

4.4. The impact of nozzle temperature

4.4.1. The response of the transition delay to temperature

Contrary to the expectation that the lower viscosity associated with higher temperatures would decrease response time, it seems like the delay in linewidth changes actually increases with temperature. This is visible both in the data of the cross-sectional area, as seen in figure 4.3, and the data of the linewidth, shown in figure 4.4, which can be taken as confirmation of the observed trend. This is potentially due to the shear thinning behaviour of the material, as the lower temperature applies a greater shear to the material. The effect on the viscosity of the shear thinning could potentially be greater than that of the increased temperature.

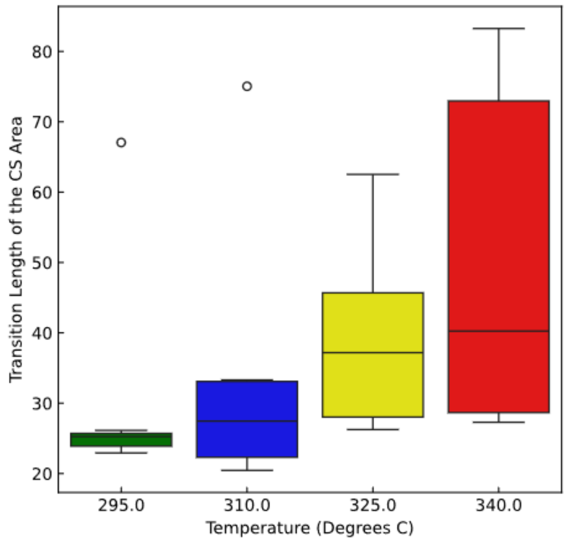

Figure 4.3: The transition delay of the cross-sectional area following changes in linewidth at different temperatures. The variation in the delay between sample lines increases with temperature, making it difficult to determine if the response is linear or exponential.

4.4.2. The different transition delays of the cross-sectional area and the linewidth

A key observation from these plots is that the transition delay is different between the cross-sectional area (figure 4.3)and the linewidth(figure 4.4). The longer length of the cross-sectional area slope suggests that although material is still being extruded, this material is not contributing towards increasing the width, but rather going towards increasing the height. This pattern can also be observed in figure 4.7, which plots the responses of the variables for the different temperatures, further supporting the pattern.

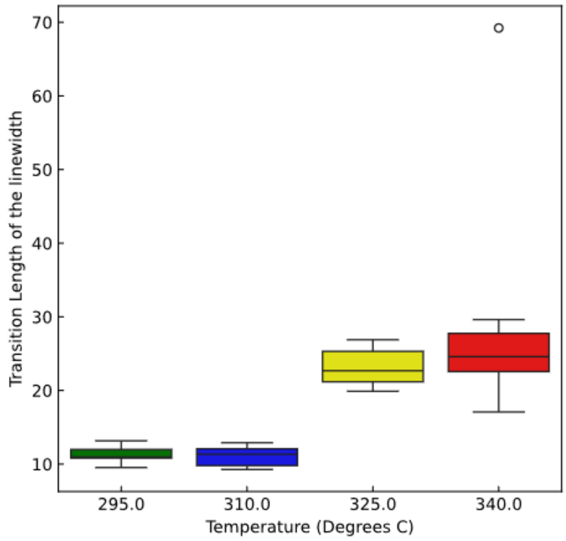

Figure 4.4: The transition delay of the measured linewidth following changes in linewidth at different temperatures. The increase in variability is again visible, while the delay itself appears to change change between two discrete levels.

The way that the linewidth reaches a plateau before the cross-sectional area stops increasing, which appears for all the experiments, suggests that there is a limit on the linewidth for a given set of parameters. Beyond this limit, extruding more material will only increase the lineheight. This was originally suspected to be due to the printing infrastructure; the pressure of the melt could be lifting the nozzle, as it is not mounted very rigidly. This explanation is nevertheless deemed unlikely. However there are many alternative explanations such as material properties like surface tension. Other explanations relate to the printing process, such as the flow behaviour upon leaving the nozzle, or the cooling of the material due to contact with the bed solidifying it and preventing it from spreading out any further. No supporting evidence was found for any of these potential explanations, which means this phenomenon will require future research.

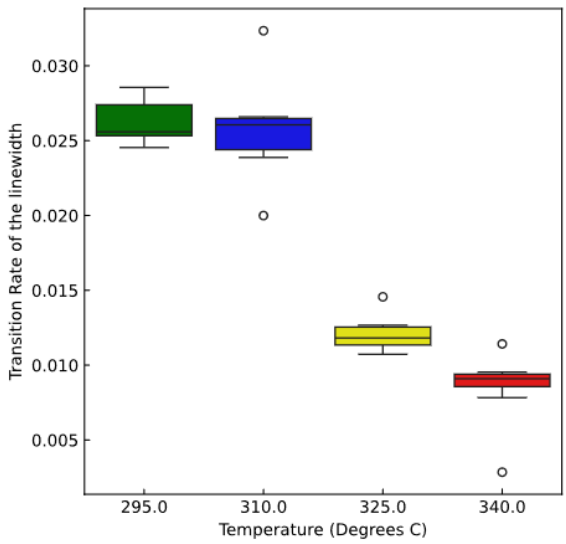

Figure 4.5: Transition rate of the measured linewidth following changes in linewidth at different printing temperatures.

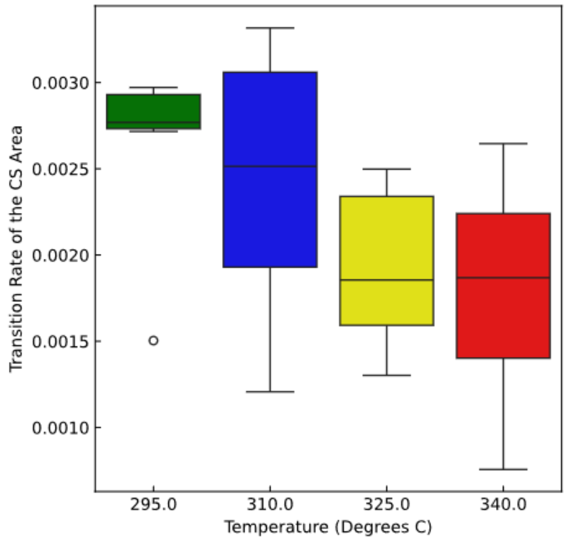

Figure 4.6: Transition rate of the measured cross-sectional area following changes in linewidth at different temperatures.

4.4.3. The variation in transition rates

The transition rate response is not as clear cut as expected. The transition rates do broadly follow the logically similar trends of the transition delays. An increasing delay is associated with decreased transition rate, as seen in figures 4.5 and 4.6. The linewidth appears to undergo a transition at a temperature between 310 and 325 degrees. This is clearly seen in both the transition delay and rate.

The cross-sectional area measurement however seems far less reliable, with the inter-quartile range (IQR) being much larger for the cross-sectional area. This increased spread in the data means that the nature of the cross-sectional area relation to temperature is unclear. It could be either linear, based on the mean values, or exponential, based on the distribution of the IQR.

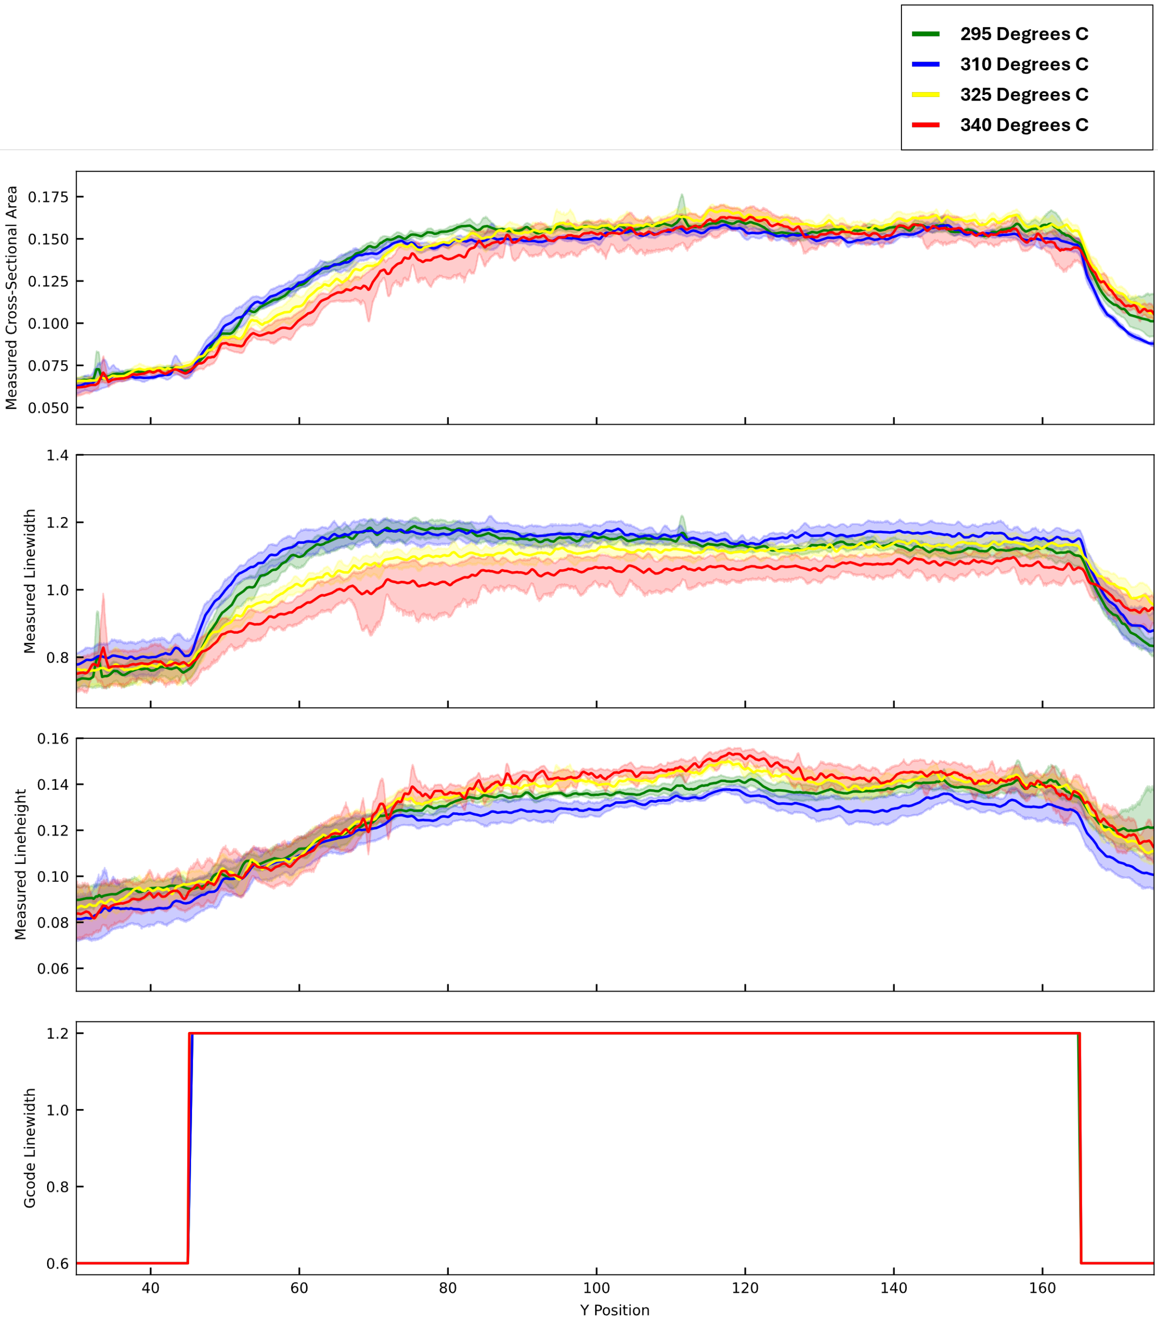

Figure 4.7: A comparison of the results of the experiments investigating the effect of temperature. The dark line is the median, while the shaded region represents the 95% confidence interval.

4.4.4. The linearity of the responses

Further investigation by plotting the results in figure 4.7 shows that the response of the cross-sectional area, which is naturally a combination of the linewdith and lineheight, is different due to the contribution of the lineheight. While the lineheight varies almost linearly following a change in the set linewidth, the linewidth itself follows an asymptotic curve. This asymptotic curve is clearly different for temperatures of 310 and below, compared to the temperatures of 325 and above, a difference only seen in the linewidth data.

When extracting the slopes from this data, the assumption was made that the transition changes could be approximated as linear slopes. Initial evaluation of the cross-sectional area plot corroborated this. However the plot of the linewidth demonstrated a clearly non-linear response. The increase in lineheight does appear to follow a linear trend, but for both of these datasets, until the underlying cause is known, assuming any shape in the response will be prone to potential errors. For the purpose of comparing the responses however, the linear approximation does allow us to compare the duration of the transition, and provides a simplified, albeit potentially oversimplified way of evaluating the transition rate response. From a purely process control perspective, it provides a way of inferring what parameters to change to achieve certain results. This is sufficient for making a very simple loop that should eventually reach the correct parameter value. Ideally, using a predictive capability, we would obtain the exact parameter value without the need for iteration, however this requires a further fundamental understanding of the causes of the transition response.

4.4.5. Potential causes for the discrepancy between lineheight and linewidth

Surface tension in the deposited melt would cause the material to minimise its surface area, which would cause it to change shape from a stadium (shape commonly described in literature and used to calculated extrusion values), to more of a circular cross section. While the existence of surface tension in the material has been observed, no information was found as to its significance. This rounding of the cross section would theoretically be counteracted by the adhesion to the bed, however this adhesion has been found to be very low, potentially being overcome by the surface tension. This surface tension could also be one of the causes for the LCP often separating from the bed, as it reduces the contact area. Increasing the temperature would lower the viscosity, and increase the period of time that the material is above its melt temperature after being extruded. This combination means that increasing temperature would give the surface tension more time to further deform the deposited line, and at a higher rate. However, the evidence is not definitive, in particular when looking at the cross sections of the lines. These were cut and then imaged under the same Keyence microscope, with little evidence of the surface tension causing a significant change in the cross section.

As the flow leaves the nozzle, it might form vortices. This is another potential cause for the difference in layer height and width. As the material leaves the nozzle and is pushed into the bed, it could form symmetric vortices in the plane formed by the nozzle axis and perpendicular to the direction of travel, that cause the lineheight to increase as material is pushed upwards by the vortices [13]. As the temperature increases, the vortices are able to move more before being frozen in place, hence increasing the temperature allows for the material to move more and curl on itself more, giving it a greater lineheight. This is difficult to predict, as it will depend on many material and process parameters. Furthermore, it will occur at the opening of the nozzle, where the geometry of the nozzle tends to obstruct this viewpoint, making in-situ observations more difficult.

It is very likely that the actual cause is a combination of these potential phenomenons. The impact of these phenomenon should decrease as the linewidth approaches the nozzle diameter. The presence, or lack thereof, of these phenomenons when extruding at widths below the nozzle diameter, such as proposed in Houriet et. al.'s [13] work, would require more experiments to be carried out. These phenomenons fail to offer an explanation for the varying delay in the changes to the cross-sectional area. The increased delay in the correct cross section would suggest some mechanism is restricting the additional material from leaving the nozzle, which is counter-intuitive to the increased viscosity associated with increasing temperature.

4.4.6. Potential cause of the transition delay

One potential explanation which could cause an increased delay with increased temperature, is that the lower viscosity enables turbulent vortices to occur in the nozzle. These turbulent vortices would likely reach a relatively steady state when consistenly extruding with the same parameters. This steady state would need to change when changing the linewidth. The "inertia" of this steady state could be what is causing the observed transition delay. An increase in the linewidth would mean more material needs to move downwards to the nozzle opening. Depending on the amount and degree of vortices, this could mean that material will require a greater change in its motion, hence why it could be described as inertia.

This turbulent motion within the nozzle would likely be quite random, which would result in the actual delay also exhibiting a large degree of randomness. This matches the data, as both the linewidth and cross sectional area have an increasing variability in the results as temperature increases. Furthermore airflows generally undergoes a transition point, where the flow transitions from laminar to turbulent, with the turbulent flow often growing from the turbulent boundary layer present at the edge of the flow. Based on this, it is also possible that the changes in the transition rate and length of the linewidth, which seems to show two distinct levels, is due to the onset of turbulent flow somewhere in between the temperatures of 310 and 325 for the geometry of this nozzle. These potential explanations should be treated as hypotheses, which will be further examined, but not conclusively supported or discarded by this work in the upcoming sections and chapters, and warrant further investigations.

4.5. The impact of varying speed

Print speed, otherwise referred to as feed rate, is the speed of the nozzle when printing. This is one of the base parameters, usually defined in mm/min, as that is how it is coded in G-code. It must be noted that the printer will automatically adapt the extrusion rate to match the feed rate, which is to say that increasing or decreasing the feedrate will not affect the amount of material the extruder pushes in for a given distance travelled. This means that when increasing the feedrate, the amount of filament being pushed into the nozzle per unit time is increasing, hence the shear rate of the material is increasing, as more material is being pushed through the nozzle per unit time. This increase in shear rate should increase the degree of alignment of the material. Furthermore as LCP is a shear thinning material, it should lover the viscosity of the melt.

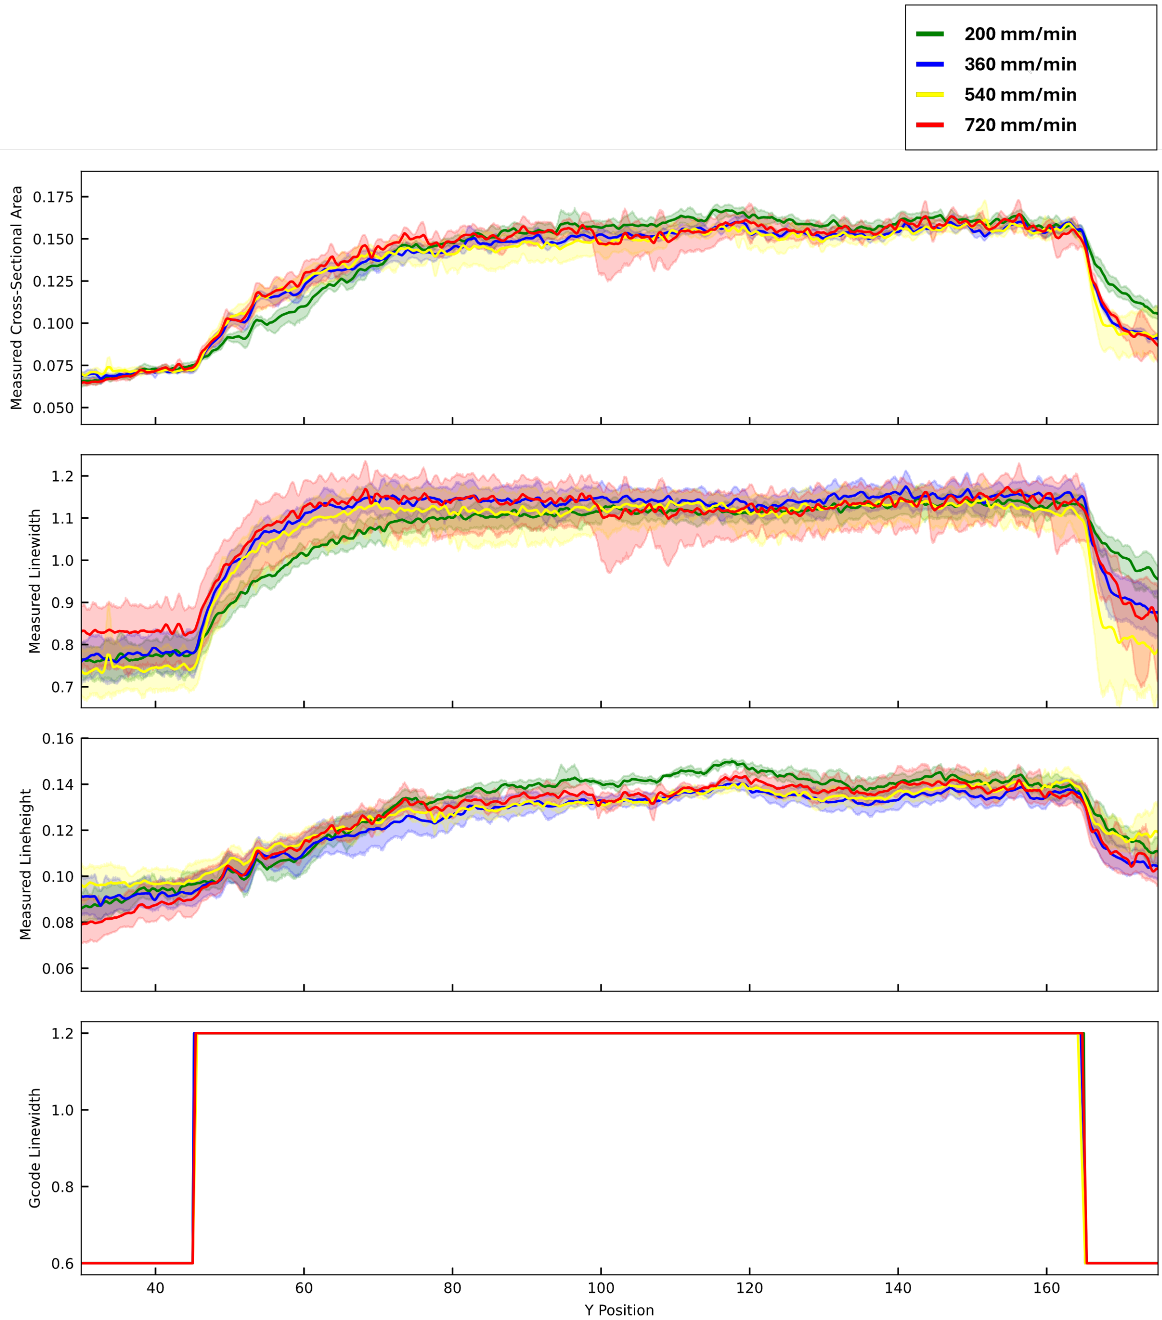

Figure 4.8: A comparison of the results of the experiments investigating the effect of varying feedrates. The dark line is the median, while the shaded region represents the 95% confidence interval. The measured lineheight shows little difference between the different feedrates. The measured linewidth however shows a clearly different slope for the slowest feedrate of 200mm/min.

Initial planning aimed to have a broader range of printing speeds that more closely resembled commonly used print speeds. The aim for the slow speeds is to try to get a very high image and data rate per distance travelled, and the speeds would increase from that point. However, when reaching a print speed of 720 mm/min, it was found that the printer emitted a loud noise, and the printbed was observed to visibly stutter. Investigations into the issue revealed that the printer itself was stuttering. This was due to the sheer amount of G-code the processor was dealing with. To obtain the position data, the motion commands were split into small commands, each followed by an M118 command where the printer would report its position. Under normal operations, the next line of code would be processed as the move was being made, which is sufficient for smooth motion. In the case of these samples, the rate at which commands were given was much higher, and more commands were given such as reporting the position to serial. This did not allow sufficient time for the 8-bit processor to process the next set of commands, hence the printer suddenly stopped moving as the stepper drivers waited for the next motion commands to be received. The impact of this stuttering meant that the data from that experiment needs to be discarded. This can also be seen from the data itself, such as in figure 4.9, as the clear trend set by the lower speeds is gone in the results from this experiment.

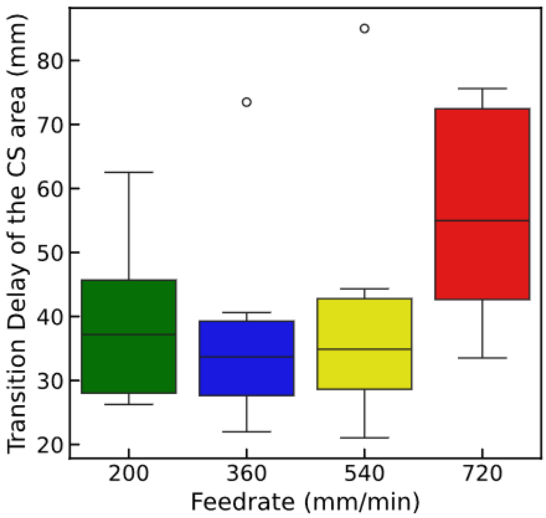

Figure 4.9: Transition delay of the cross-sectional area following changes in linewidth at different feedrates, defined in mm/min, as it is defined in G-code. Feedrate is the speed at which the nozzle moves, and not the rate at which lengths of filament are fed into the extruder.

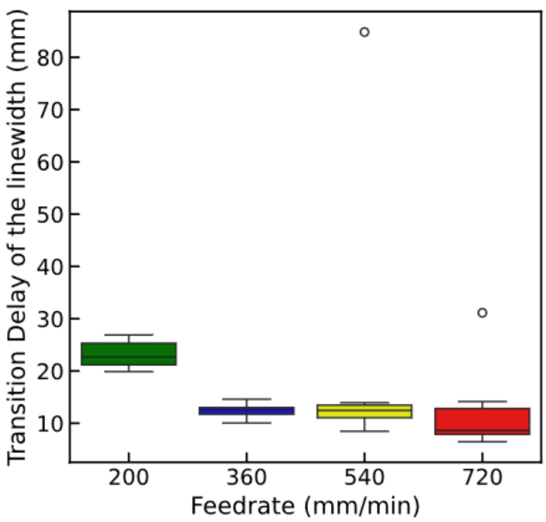

Figure 4.10: Transition delay of the linweidth area following changes in linewidth at different feedrates

Looking at the delays in linewidth and cross-sectional area, plotted in figures 4.9 and 4.10, increasing the speed of the printhead causes a reduction in the delay of the material being extruded, which is likely due to the shear thinning effect reducing the viscosity and enabling more rapid changes in the flow. Furthermore, the the linewidth gradient rapidly increases with feed rate, and the length of the slope decreases, suggesting that the linewidth responds much more rapidly to changes, and is more sensitive to the speed than the cross-sectional area. It is possible that the shear thinning effect results in the material being locked in place earlier. This is because the shear thinning also causes the viscosity to decrease more rapidly once the material has left the nozzle as no more shear is being applied. This would affect the potential causes of the discrepancy, such as the surface tension, from deforming the extruder material as much.

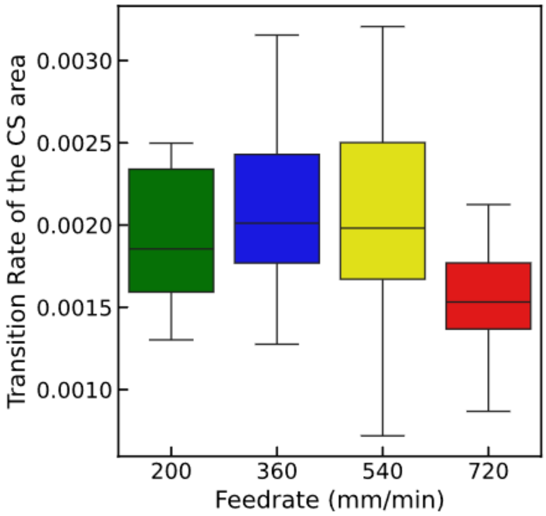

Figure 4.11: Transition rate of the cross-sectional area following changes in linewidth at different feedrates

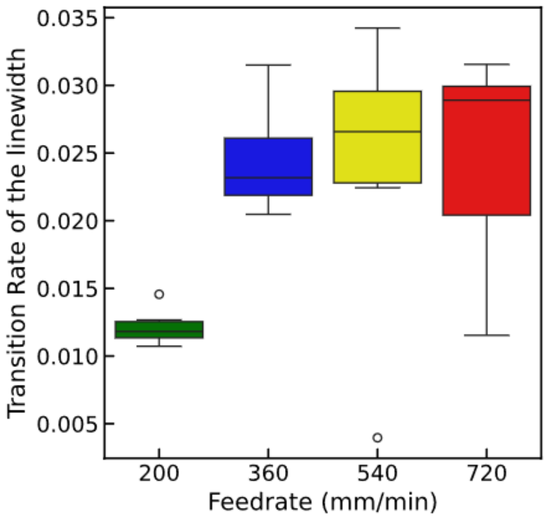

Figure 4.12: Transition rate of the linewidth area following changes in linewidth at different feedrates

As discussed previously, when looking at the lineplot in figure 4.13, we see that at speeds above 200 mm/min, the pressure sensor is again able to give readings, corroborating the idea that the issues with the sensors are due to its use at pressure outside of its measurement range. Unfortunately, the lack of pressure data coverage across all the experiments makes it difficult to draw conclusions. However it is interesting to see that while the pressure also has a sloped region of mounting pressure, it is distinctly shorter and steeper than that of the cross-sectional area. In this regard, it's behaviour is more akin to that of the measured linewidth.

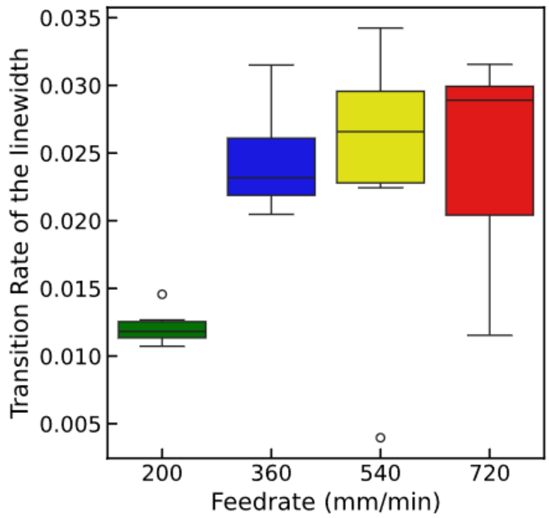

Figure 4.13: Transition rate of the linewidth area following changes in linewidth at different feedrates

4.6. The impact of varying linewidths

The effect of changing the linewidths is similar to what was expected. Increasing the step change results in an increase in the gradient of the slope, as the pressure in the nozzle should ramp up to a higher pressure, which would increase the shear rate. However, the effect of starting at a higher initial linewidth, which would mean a higher initial pressure, does not appear to create a reduction in the slope length that was expected. For the same step change in linewidth, the slope length is only a little longer for both the cross section and the linewidth. The results suggest that the slope length is governed primarily by the magnitude of the change in linewidth, with a potentially smaller impact of the start and end pressure, as seen in figures 4.14 and 4.15. This is in contrast to the gradients.

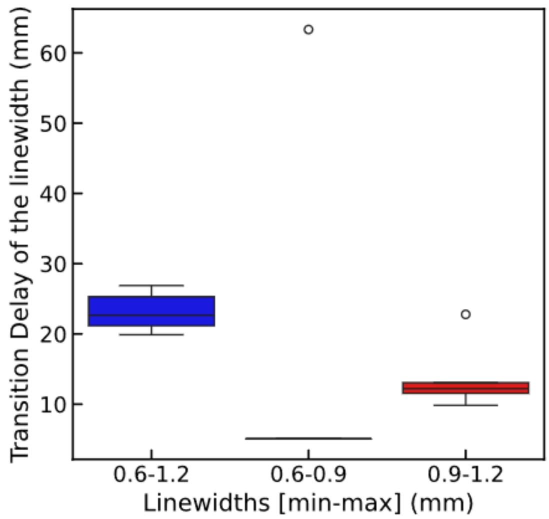

Figure 4.14: Transition rate of the measured cross-sectional area following changes in linewidth at different feedrates

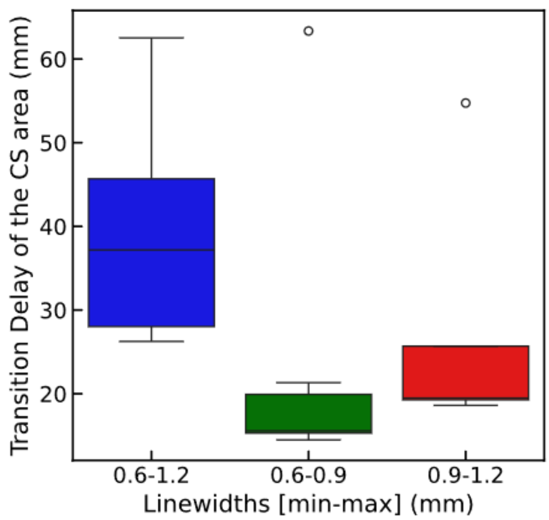

Figure 4.15: Transition rate of the measured cross-sectional area following changes in linewidth at different feedrates

The gradients, as showin in figures 4.16 and 4.17, appear to be more affected by the final linewidth, rather than the step change. Interestingly, the slope gradient of both the linewidth and the cross-sectional area are similar for experiments 7 and 14, suggesting that the gradient is affected primarily by the final pressure, as opposed to either the starting pressure or magnitude of the change in linewidth. This can be explained by the gradient being the direct response of the pressure in the nozzle, and suggests that the pressure required to extrude the final linewidth is applied rapidly, which corresponds to the pressure readings seen during the speed comparison. The gradient should have a large impact on the final slope length, together with the step change that was discussed. Therefore applying an excessively large pressure rapidly, before reducing to the final extrusion pressure, should be a valid way of reducing the delay in changes.

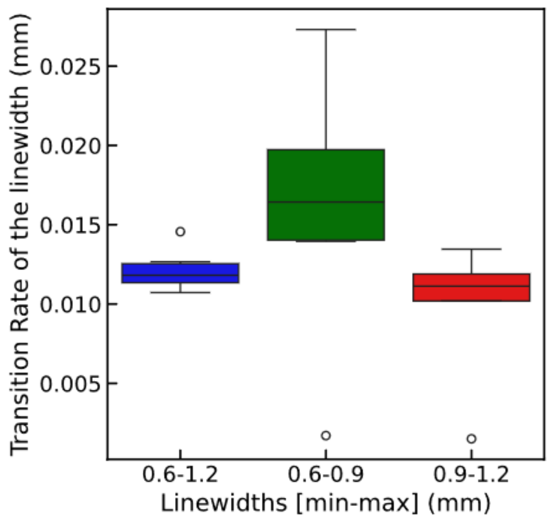

Figure 4.16: Transition rate of the measured cross-sectional area following changes in linewidth at different feedrates

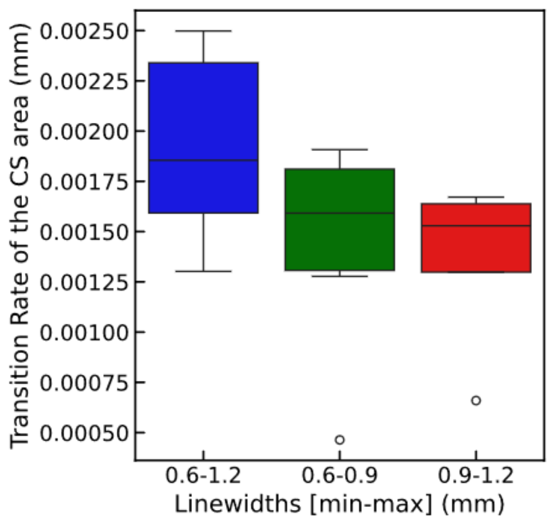

Figure 4.17: Transition rate of the measured cross-sectional area following changes in linewidth at different feedrates

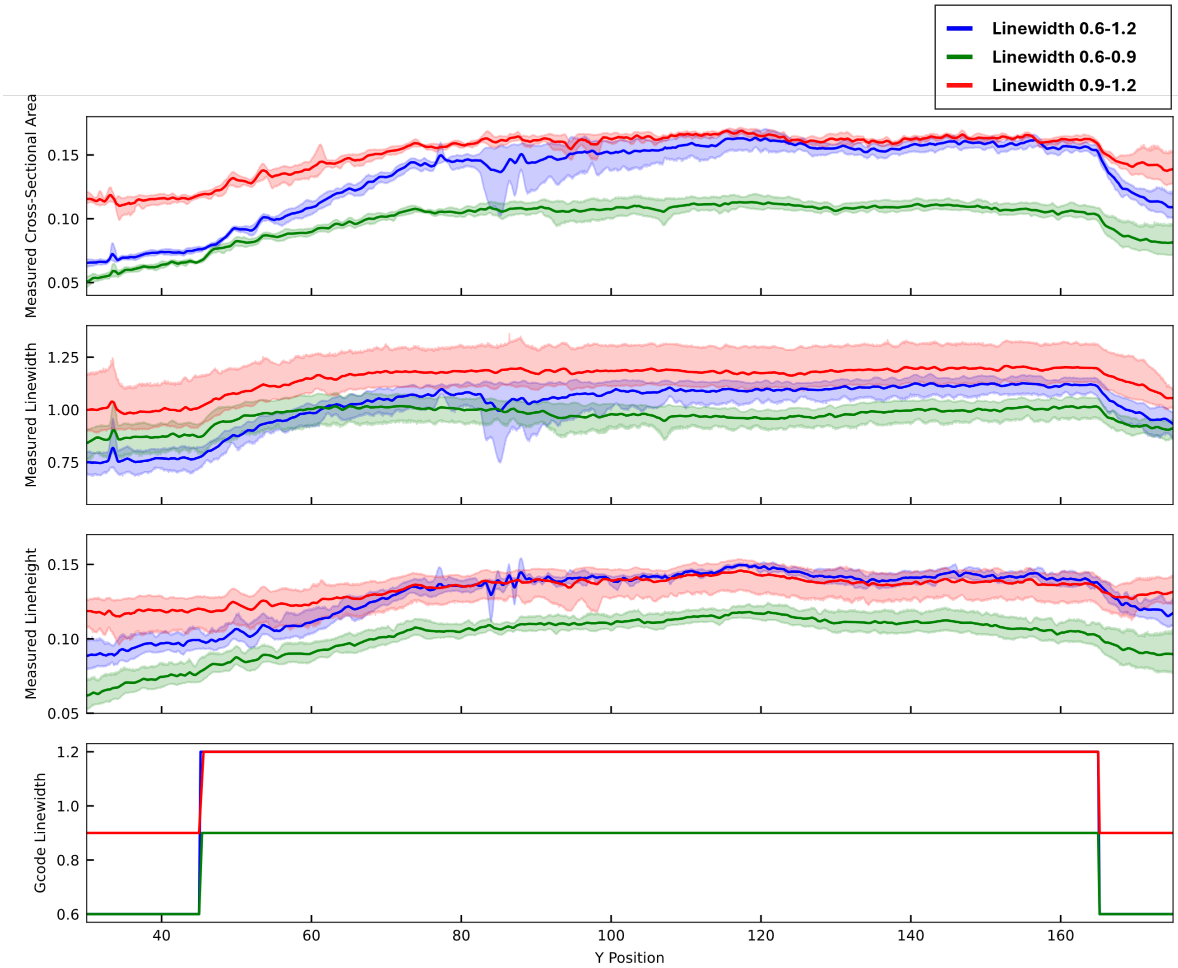

Figure 4.18: A comparison of the results of the experiments investigating the effect of varying the changes in linewidth. The dark line is the median, while the shaded region represents the 95% confidence interval.

4.7. Results in the context of current work with LCP

One of the major aims of this work is to work towards achieving the level of control that is required for the topology optimisation and material properties set out in the work of Houriet et. al.[13]. It is necessary to first address some differences in how this work was carried out.

4.7.1. The experimental differences in line geometry

Their work strictly used a layer height of 0.1, however this was not continuously measured, hence variations might be possible, as was discovered in this work. However, for the sake of the comparison, we will assume they obtained a constant layer height. When attempting to print with a layer height of 0.1, it was found that the strain gauge of the rheometer would scrape the bed, and therefore all the prints and the extrusion quantity was set with the assumption of a layer height of 0.15, with the layer widths defined by the ratio of lineheight to linewidth as set out in their work. The difference in layer height means that direct comparisons are difficult, but we should still be able infer patterns from the data that are applicable to both since the ratio is kept constant. This ratio was considered the more important, because this ratio is believed to define the flow path as it left the nozzle. For this reason, we can still reach conclusions for how to best approach the different requirements set by their work.

Their work established that the alignment had a major impact on the properties obtained, and this alignment was affected by how much the flow spread out from the nozzle. Their comparison is based on a ratio of the linewidth to the nozzle diameter. This ratio does not incorporate the lineheight, which could cause reductions in the properties if too high, similar to the linewidth. Ideally a way to incorporate this parameter will to be added to the equation, or a second separate relation will need to be made for it. Depending on the cause of the observed limit of the linewidth, it may be possible to effectively find a minimum linewidth, below which the lineheight may indeed be assumed constant, and its effect be considered negligible on the properties.

4.7.2. The effect of temperature in context

In their work, they found that printing with lower temperature was beneficial for achieving higher anisotropy and higher temperatures. This bodes well for printing precise geometries, as this work suggest that printing at lower temperatures is beneficial for reducing the delay in changing exrusion parameters, hence fine control of both the geometry and the strength is possible at lower temperatures.

4.7.3. The effect of linewidth and nozzle diameter in context

As part of their work, they suggested that for optimal strength, a narrow nozzle with a small linewidth that is at or below nozzle diameter should be used, as this ”draws” the filament out of the nozzle aligning it. The results of this work suggest that having a smaller nozzle and the associated increased shear rate should be beneficial to reducing the delay in parameter changes, and therefore improving the predictability of flow geometry for a given linewidth. However, there are two barriers to making more conclusive statements. The first is that due to the rheometer, this work could only use one nozzle diameter, and more work needs to be done to see if changing the nozzle diameter affects the results. The second is that to obtain a similar geometry ratio, this work used very wide linewidths, hence this work is unlikely to accurately predict the effect of the printing parameters when printing linewidths below the nozzle diameter. The sub-nozzle diameter linewidth recommend for the absolute best properties mean that the material is no longer being sheared due to pressure in the nozzle; it is more akin to drawing, which is governed by different flow phenomena, and the research carried out here is unlikely to accurately represent this. However, it is possible that the kind of high alignment generated by drawing out the filament could potentially be achieved by printing faster and increasing the shear rate.

4.7.4. The effect of feedrate in context

Their experiments were carried out with a motion speed of 60mm/s, which was not possible with our setup; the highest speed achieved (with stuttering) was 12 mm/s, hence the shear rates can be expected to be much lower than for the lines printed as parts of their experiments. This also means that the properties of the sample lines in this work can be expected to be far more isotropic, with lower properties. Nevertheless, even at these lower shear rates, it was demonstrated that increasing the shear rate is beneficial for achieving a faster response in parameter changes. It is worth investigating if higher speeds can be achieved that would increase the shear rate, perhaps in combination with nozzle geometry changes, that could result in a similar level of alignment.

Nevertheless, if we assume that we can apply the equation derived in their work to the results obtained in this work, then it is possible to calculate the theoretically achievable rate of change of Young’s Modulus based on the measured rate of change of the linewidth. It should be noted that the true highest rate of change is likely higher, as the calculated transition rate assumed a linear rate of change. While this seemed correct for the cross-sectional area and lineheight, the measured linewidth does not seem to follow a linear trend, but instead an asymptotic one. Hence the true achievable rate of change of the modulus could be even higher.

Temperatures (degrees C)	Feed Rate (mm/min)	Min Linewidth	Max Linewidth	Measured Transition Rate	Calculated Achievable Rate of Change of Young's Modulus
295	200	0.6	1.2	0.026	1.64E+07
310	200	0.6	1.2	0.025	1.76E+07
325	200	0.6	1.2	0.012	7.30E+07
340	200	0.6	1.2	0.009	1.27E+08
325	200	0.6	0.9	0.014	5.43E+07
325	200	0.9	1.2	0.011	8.63E+07
325	360	0.6	1.2	0.022	2.26E+07
325	540	0.6	1.2	0.028	1.42E+07
325	720	0.6	1.2	0.028	1.42E+07

Table 4.1: A table showing the input parameters and the resultant measured transition rate, which is then converted according to the equation for the Young's modulus as proposed by Houriet et. al. [13]

5 Conclusion

This work set out to utilise a combination of in-situ and ex-situ sensors to investigate the effect of print parameters on the extruded geometry of fused deposition modeled LCP. To this end, a test setup was built which was able to successfully combine the data from these sensors for the experiments. In doing so, a dataset was created that contains a large amount of data from each of the sensors, that can be used at a later date to verify print parameters, as well as train an algorithm to predict what changes in parameters would be beneficial to obtaining a specific geometry.

This ability to predict the effect of the parameters on the geometry was demonstrated to be necessary, as a clear delay in the extrusion change was observed, the magnitude of which changed as parameters were changed. This will need to become predictable in the future if precise complicated geometry will need to be printed that incorporates continuous linewidth changes in the pursuit of specific material properties. It was further demonstrated that the linewidth was not the sole parameter that needs monitoring with regards to the geometry of the output, as the lineheight demonstrably changed and was also affected by the print parameters.

Based on the data, it would seem that print geometry is very consistent and reliable across the tested parameters, with transitions being quite smooth and gradual, suggesting that actively modifying parameters to change the final geometry is a valid approach that will not result in the flow conditions entering an unstable or turbulent state that requires high speed monitoring and compensating. This also opens up the use of more sensors for in-situ monitoring that might otherwise be disregarded due to their sampling speed.

This sampling speed need not be an issue, as this work has further demonstrated the viability of mixing sensor datasets sampled at different times. However it is difficult to do, as not all sensors have the same datarate, and it is critical to have some measured parameter that is shared between them and sampled at a high frequency, such that their datasets can be combined. For in-situ measurements, time is the easy and obvious choice, however when using an ex-situ sensor after the print, this is no longer an option. When using position as the reference. and if the sampling frequency is low, then it is probably that the data cannot be combined with the same level of positional accuracy, unless a separate high precision sensor can be used.

In-situ monitoring techniques are likely to be required in the future, as the changes in extrusion width are difficult to predict, and likely too computationally expensive to predict fast enough for an open loop system, especially when considering the vast variety of materials and their potential additives. This only becomes more critical when dealing with LCP's, where the properties of the material are intrinsically linked to the parameters of the process. For this reason, the chosen approach of combining multiple sensors was considered a promising avenue of research, and this work has demonstrated that the information that can be gleaned from combining datasets can be useful to understanding the behaviour of a material or process.

6 Outlook and Recommendations for Future Work

6.1. Rheometer

6.1.1. A critical review

The rheometer had a number of issues in functioning, such as the lack of an ability to calibrate it once mounted. This meant that although it was possible to reliably get correct values for changes in pressure, the absolute values could not be obtained. The temperature sensor in the tip, although located in contact with the melt, was clearly affected by the thermal conductivity of the pushrod. This was mostly clearly visible when switching on the part cooling fan, which would drop the measured temperature by around 60 degrees. Initial differences of over 100 degrees were measured at first, before it was discovered that the thermocouple was improperly in the pushrod, resulting in a new thermocouple being mounted into the pushrod. It nevertheless still measured a temperature that was consistently around 10 degrees lower than the temperature according to the thermistor of the printer. Measurements with a temporary thermocouple inserted into the nozzle without any extrusion taking place gave readings more in accordance with the thermistor, suggesting that even without the part cooling fans blowing over the pushrod, the pushrod was acting as a radiator for the thermocouple measurements.

These issues alone do not preclude the use of the rheometer, however it also imposed severe restriction on the motion of the printer due to the strain gauge being located so close to the bed. It furthermore required a significant amount of cleaning prior to every print to clear up the leaking material. These factors, combined with the particularly low viscosity of LCP, make the design as a whole impractical for further experiments. Despite this, a sensor which measures the pressure of the melt would still be very useful. In this case, a sensor that is also capable of measuring the pressure on the nozzle tip by the melt as it is being deposited is particularly interesting. This design would ideally also be nozzle agnostic. The results of this work suggest that further work looking at varying the geometry of the nozzle could lead to some very interesting results, especially in combination with varying the nozzle diameter.

6.1.2. Potential improvements

This leads to the considering of sensors that instead mounted on the entire cold end assembly, similar to how the newer Prusa MK4 uses a strain gauge to probe the bed. This sensor would ideally be capable of measuring the pressure on the nozzle, regardless of motion direction. For use in spin printing, and to further evaluate conventional printing, it would also be intersting to apply strain gauges in lateral axes parallel to the motion, to measure the sideways forces on the nozzle. This might give further insight into the effect of extruding material out as in conventional printing, and allow for comparison to drawing the material out as in spin printing, which incidentally currently achieves the highest degree of anisotropy.

6.2. Sampling rate and hardware/software limitations

One of the fundamental issues encountered in the development of the test setup used in this work was a legacy of how additive manufacturing machines have been developed since the reprap movement. Initial printers were built using cheap on-hand electronics, such as arduino microcontrollers, simple stepper drivers and stepper motors. These served their purpose well, and although some changes have been made more recently, such as the move to 32 bit microcontrollers, the fundamentals have remained the same. This is also true for the software, which is designed to work almost headless, necessitating the most bare bones display and interface. This made these machines much more accessible and easy to use.

Unfortunately, this legacy means that these machines lack the outputs that would allow rapid data transfer to an external machine. Due to the high processing requirements that come with in-situ monitoring, especially if optical data is used, the small microcontrollers common to older printers are no longer sufficient. Even modern single board computers (SBC) are unable to keep up with such demands. Hence, an external ”command centre” is needed, that is capable of rapidly receiving large quantities of data from the printer and it’s sensor, and process them. It must then also be able to directly send commands that modify parameters and intercede on the base g-code provided. Recent firmwares such as klipper, where a SBC will first do some initial processing to the g-code before feeding it to the printer are a step towards such capabilities, however more steps are needed, as this is still open-loop control. This needs changes to how the firmware works, but also hardware changes, such as motion tracking sensors that enable the printer or the command centre to know what the current position of the printer is, as opposed to sending commands and blindly hoping they are correctly followed. A successfully attempt towards this was made using the existing hardware, and M118 g-code command, however this rapidly overburdened the 8-bit microcontroller and lacked a sufficiently high datarate. Furthermore, existing g-code seemed limited in terms of what commands could be fed to the printer to change print parameters, with the currently supported g-code commands being changing temperatures and flow rate. These commands would be insufficient to make corrective moves and extrusions.

Future software work could focus on separating the firmware into two sections, one which interprets given commands and executes the low level commands to the motor drivers or heaters, and one which actively processes data from the printer and the original g-code and decides what high level commands to send to the first section. This second section would retain full control of motion commands and can be directly interfaced with by the user. Future hardware work could look into suitable motion sensors, such as for example encoders, and how to rapidly sample and transfer this data from the sensors or printer board to this new command centre for processing via high bandwidth data links. The number of potential sensors is very high, as many are suited to very particular tasks, hence to obtain information from the various parts of the process, a multitude of sensors will be required, likely including a number of optical ones, which are notoriously bandwidth and processing hungry.

6.3. Positioning data and its necessity for sensor fusion

The positioning technique used in this work has demonstrated its usefulness, as it enabled all the gathered data to be overlayed based on the position of the toolhead at the time that the measurements were taken. This is a far more intuitive method of overlaying sensor data compared to time itself, as the motion occurring is also important to the process, and in a case where feedrate is not constant, finding the motion that is occuring based upon a timestamp during or after the print is much more difficult.

The positioning technique is however not without it’s flaws. It significantly limited the maximum feedrate achieved by the toolhead. Based on the raw data of the feedrate experiments, it seems that although imperceptible to a user sitting next to the machine, it is likely that micro stutters were occuring even in the slowest prints. This is not an acceptable drawback for such a sensing system, especially for industrial use.

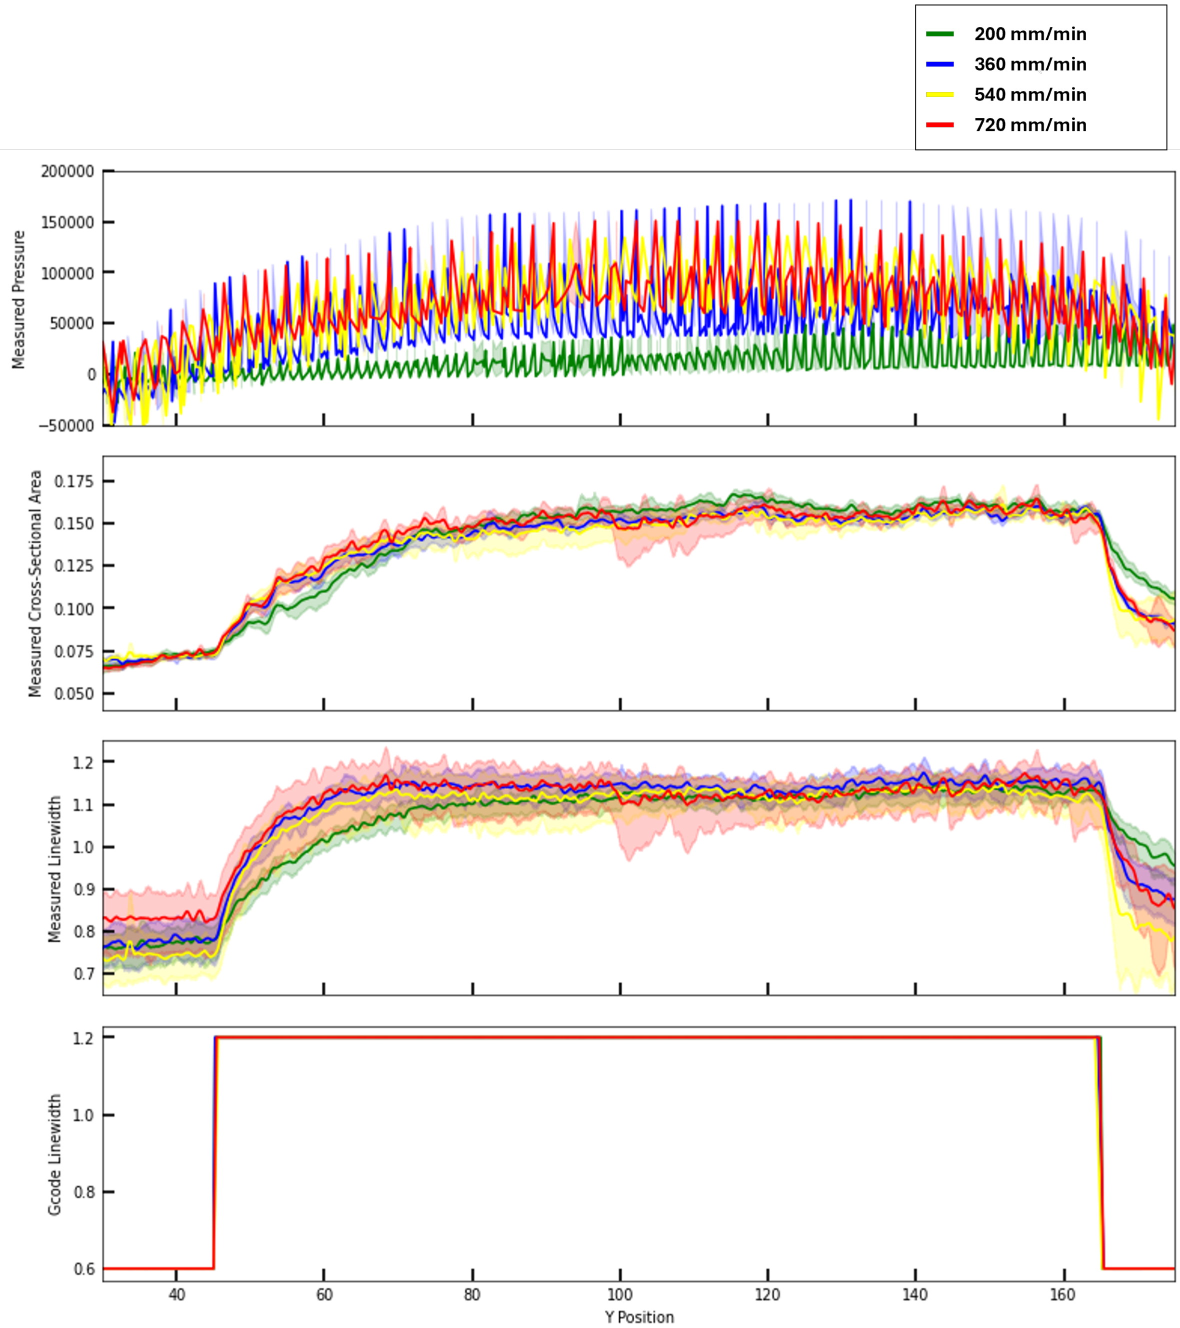

Figure 6.1: Results from the feedrate experiments, with the top graph showing the pressure data. Note the spikes in the data, this is due to the start-stop motion of the toolhead as the microcontroller stutters. This motion translates to a vibration which shakes the strain gauge measuring the pressure, resulting in the vibrations being recorded as well as the pressure. As can be seen by the shaded spikes, the vibrations do not necessarily all occur at the same location along the line. Note that only the red vibrations were of sufficient magnitude due to the higher speed and the pauses sufficiently long that the stuttering was visibly and audibly noticeable. At 540 mm/min, the yellow line, it was suspected to vibrate only after rerunning the print and touching the 90 degree bed to feel the vibrations which were almost imperceptible.

Positioning information can be obtained in a number of ways, such as for example encoders on the stepper motors controlling the motion, or a camera tracking the toolhead. The measurement system would ideally measure the position at a very high frequency, as it will be critical for sensor fusion, hence precision of the combination of data will be dependent on it's sampling rate. Furthermore, with work like that of Petsiuk et. al. [16], the ability to use cameras mounted on the toolhead to corroborate the geometry with that of a digital twin will require precise knowledge of the positions of the cameras, and therefore the toolhead.

6.4. Future work with LCP

The results suggest that some of the assumptions made on the flow in the nozzle, may be incorrect. It has already been demonstrated that utilising narrower nozzles can provide higher degrees of anisotropy, with the understanding that the material was being drawn out more and experience more shear aligning the crystals. Based on the results shown here, and the potential cause for the measured delays, it is distinctly possible that more flow turbulence is occurring in larger diameter nozzles, at least in the melt sections. This turbulence can be reasonably expected to affect the alignment of the material as it leaves the nozzle. This requires more research to confirm, but may need to be addressed, or alternatively controlled and even induced. It is known that high degrees of anisotropy prevent two layers from bonding well, causing weak interlaminar properties. If a relatively isotropic outer shell for the deposited material can be generated by deliberately inducing small controlled turbulence in the outer region, then potentially this interlaminar strength could be improved.

References

[1] *AR6 Synthesis Report: Climate Change 2023*. en. URL: https://www.ipcc.ch/report/ar6/syr/ (visited on 11/16/2024).

[2] Erik de Bruijn. "On the viability of the open source development model for the design of physical objects Lessons learned from the RepRap project". en. In: ().

[3] Yi-Sheng Chen et al. "Hydrogen trapping and embrittlement in metals – A review". en. In: *International Journal of Hydrogen Energy* (Apr. 2024), S0360319924013326. ISSN: 03603199. DOI: 10.1016/j.ijhydene.2024.04.076. URL: https://linkinghub.elsevier.com/retrieve/pii/S0360319924013326 (visited on 11/16/2024).

[4] Timothy J. Coogan and David Kazmer. "In-line rheological monitoring of fused deposition modeling". In: *Journal of Rheology* 63.1 (Jan. 2019), pp. 141–155. DOI: 10.1122/1.5054648.

[5] D.A.J. Brion, M. Shen, and Sebastian W. Pattinson. "Automated recognition and correction of warp deformation in extrusion additive manufacturing". In: *Additive manufacturing* (May 2022), pp. 102838–102838. DOI: 10.1016/j.addma.2022.102838.

[6] D.A.J. Brion and Sebastian W. Pattinson. "Generalisable 3D printing error detection and correction via multi-head neural networks". In: *Nature Communications* 13.1 (Aug. 2022). DOI: 10.1038/s41467-022-31985-y.

[7] Donald Erb. *pybaselines: A Python library of algorithms for the baseline correction of experimental data*. DOI: 10.5281/zenodo.5608581. URL: https://github.com/derb12/pybaselines.

[8] Yanzhou Fu et al. "In situ monitoring for fused filament fabrication process: A review". In: *Additive manufacturing* 38 (2020), p. 101749. DOI: 10.1016/j.addma.2020.101749.

[9] Silvan Gantenbein et al. "Spin-Printing of Liquid Crystal Polymer into Recyclable and Strong All-Fiber Materials". en. In: *Advanced Functional Materials* 31.52 (Dec. 2021), p. 2104574. ISSN: 1616-301X, 1616-3028. DOI: 10.1002/adfm.202104574. URL: https://onlinelibrary.wiley.com/doi/10.1002/adfm.202104574 (visited on 04/02/2024).

[10] Silvan Gantenbein et al. "Three-dimensional printing of hierarchical liquid-crystal-polymer structures". In: *Nature* 561.7722 (Sept. 2018), pp. 226–230. DOI: 10.1038/s41586-018-0474-7.

[11] Andrew Gleadall. "FullControl GCode Designer: Open-source software for unconstrained design in additive manufacturing". In: *Additive Manufacturing* 46 (Oct. 2021), p. 102109. ISSN: 2214-8604. DOI: 10.1016/j.addma.2021.102109. URL: https://www.sciencedirect.com/science/article/pii/S2214860421002748 (visited on 01/09/2024).

[12] Gabriel Pieter Greeff and Meinhard Schilling. "Closed loop control of slippage during filament transport in molten material extrusion". en. In: *Additive Manufacturing* 14 (Mar. 2017), pp. 31–38. ISSN: 2214-8604. DOI: 10.1016/j.addma.2016.12.005. URL: https://www.sciencedirect.com/science/article/pii/S2214860416302366 (visited on 05/16/2023).

[13] Caroline Houriet et al. "3D Printing of Flow-Inspired Anisotropic Patterns with Liquid Crystalline Polymers". en. In: *Advanced Materials* 36.11 (2024). _eprint:https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.202307444, p. 2307444. ISSN: 1521-4095. DOI: 10.1002/adma.202307444. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/adma.202307444 (visited on 06/24/2024).

[14] The MathWorks Inc. *MATLAB version: 9.13.0 (R2022b)*. Natick, Massachusetts, United States, 2024. URL: https://www.mathworks.com.

[15] Rhys Jones et al. "RepRap – the replicating rapid prototyper". en. In: *Robotica* 29.1 (Jan. 2011), pp. 177–191. ISSN: 0263-5747, 1469-8668. DOI: 10.1017/S026357471000069X. URL: https://www.cambridge.org/core/product/identifier/S026357471000069X/type/journal_article (visited on 11/16/2024).

[16] Aliaksei L. Petsiuk and Joshua M. Pearce. “Towards Smart Monitored AM: Open Source in-Situ Layer-wise 3D Printing Image Anomaly Detection Using Histograms of Oriented Gradients and a Physics-Based Rendering Engine”. In: *arXiv: Computer Vision and Pattern Recognition* (Nov. 2021). MAG ID: 3209914635.

[17] *The Future of Hydrogen – Analysis - IEA*. URL: `https://www.iea.org/reports/the-future-of-hydrogen` (visited on 11/16/2024).

[18] Pauli Virtanen et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17 (2020), pp. 261–272. DOI: `10.1038/s41592-019-0686-2`.

[19] Haixi Wu, Yan Wang, and Zhonghua Yu. “In situ monitoring of FDM machine condition via acoustic emission”. en. In: *The International Journal of Advanced Manufacturing Technology* (Sept. 2015). ISSN: 0268-3768, 1433-3015. DOI: `10.1007/s00170-015-7809-4`. URL: `http://link.springer.com/10.1007/s00170-015-7809-4` (visited on 08/15/2023).