Surface Informed Print Region and Nozzle Angle Refinement for Additive Manufacturing on Concave Parts

Samuel Elliott and Matthew Campbell *

Posted Date: 18 May 2026

doi: 10.20944/preprints202605.1142.v1

Keywords: multi-axis additive; computational geometry; additive for repair; multi-material additive

 doi:10.20944/preprints202605.1142.v1

Article

Surface Informed Print Region and Nozzle Angle Refinement for Additive Manufacturing on Concave Parts

Samuel Elliott and Matthew Campbell *

Oregon State University, 204 Rogers Hall, Corvallis, OR 97331, USA

* Correspondence: mc1@mail.utexas.edu

Abstract

Additive manufacturing is moving towards the use of machines with five or more axes but is limited by the inability to easily generate print paths. This usually requires the creation of custom G-code in order to utilize five-axis printing. However, as additive manufacturing begins to utilize five-axis printing for scenarios such as repair, modification, or composite printing, the existing print surfaces become potential obstacles that need to be accounted for in path planning. This paper shows a novel way of expanding the capabilities for printing on complex, concave parts, in order to prevent collisions between the printer and the part. This is achieved through two main steps: Area Refinement, and Angle Determination. During Area refinement, the proposed print area is altered based on the geometry of the print surface, along with printer parameters. This results in a print region that is feasible with the given machine configuration and geometry, that will not attempt to print too close to any existing surface in the presence of concavities. During Angle Determination, the finalized print paths are adjusted to set a nozzle orientation that prevents the machine from colliding with the part. In this paper, we present the algorithmic details behind these approaches and show computational results for two complex concave geometries.

Keywords: multi-axis additive; computational geometry; additive for repair; multi-material additive

1. Introduction

Multi-axis 3D printing is an area of manufacturing that is largely limited by the computational complexity in defining print paths. Compared to traditional three-axis printing where many slicers exist to allow users to easily generate G-code for printers, the added complexity of five-axis path planning means that there are no widely available methods of creating G-code. Although the kinematics and physical creation of five-axis machines are relatively straightforward given current knowledge and technology, the field is severely limited by the lack of accessible tool planning. This limits multiple areas of manufacturing, including printing parts completely from scratch using five-axis printing where the entire part is printed onto a blank. Additionally, the lack of multi-axis path planning prevents the use of multi-axis machines from printing onto existing, complex geometries, as is the focus and motivation of this research. Given the ability to generate feasible toolpaths on complex geometries, five-axis printing could be used to enable modification and repair of existing structures. 3D printing, especially multi-axis deposition-based approaches, can allow for updating or repairing structures by printing directly onto the existing material. While this would likely be unnecessary for plastic deposition as the cost and time to fabricate new parts are low, this could be incredibly beneficial when using metals deposition, where the costs are higher. Additionally, multi-axis metal 3D printing could be used to repair structures that were not initially created using additive manufacturing. Printing onto existing geometries provides additional path creation complexity over traditional approaches, as the toolpaths need to account for the surface geometry in order to prevent collisions with the print surface. Throughout this paper we will refer to this printing as five-axis

 doi:10.20944/preprints202605.1142.v1

printing because that is the minimum number of axes needed in order to effectively control both the position and orientation of the print head, but many platforms have more the five axes.

One of the motivations for this research is our previous research, where we developed a 5-axis printer capable of printing on objects adjacent to itself [1]. This method, combined with collision detection for complex existing geometries, would allow for the possibility of using additive manufacturing for repair scenarios, as well as complex composite scenarios by printing a new material onto a complex existing surface.

Five-axis machining is a field with a substantial amount of research into generating feasible toolpaths that do not collide with, or "gouge", the surface. In subtractive manufacturing, because material is removed from the existing part, any region that is impossible to reach with a given toolhead is unmachinable. However, in additive manufacturing, an unreachable area may not cause the entire part to be unprintable. Rather, in some cases voids may be able to be left in the construction of the part while still maintaining the desired outer geometry. This is common in many parts in additive manufacturing due to the fact that most 3D printed parts are not completely solid, and instead have some percentage of infill. Thus, the rest of the model may still be able to be built up, while avoiding the unreachable region, as shown in Figure 1. While the entire added region is unable to be printed due to the severity of the concave region, a hollow area can be left in the model while still printing material to maintain the outer geometry of the part.

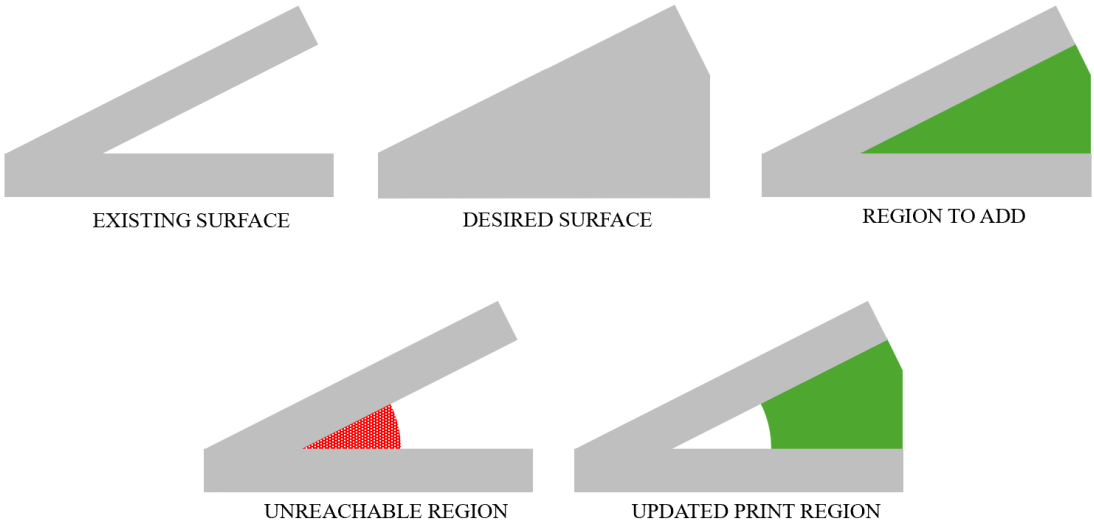

Figure 1. Example Area Refinement based on unreachable area of print. While the inner region is unreachable, the outer geometry of the additional area can still be printed, maintaining the desired outer geometry.

In this paper, we explore a process of creating toolpaths for printing directly onto existing surfaces. This is achieved through an Area Refinement phase that alters an existing print region to not include material that is unprintable based on the surface geometry. Additionally, an Angle Determination phase finds which angle to set the nozzle at for each point in the print path in order to avoid colliding with the existing part. Both of these adjustments help to allow for printing onto existing surfaces, which greatly expands the capabilities and potential uses of five-axis additive manufacturing. This is done while trying to closely adhere to the standard of automation that additive manufacturing has set. While five-axis path planning exists in many CAM software solutions for machining and subtractive manufacturing, these implementations often still require a knowledge of machining and substantial manual effort in order to effectively implement feasible paths. This paper demonstrates a proof-of-concept process for generating five-axis toolpaths for additive manufacturing, and limits the amount of hands-on decisions that the user needs to make.

2. Literature Review

 doi:10.20944/preprints202605.1142.v1

One of the major problems from moving to five-axis printing for additive manufacturing, is that it increases the possibility for collisions. Traditional 3D printing is often referred to as 2.5D printing, because each layer is planar and built from the build plate up. Because of this, as long as the nozzle only moves above the current print layer, then there is no risk of collisions. By moving towards non-planar slicing and five-axis printing, this assumption is no longer true. Papers by Wulle et al. [2] and Tang et al. [3] discuss the solutions of using an extended nozzle to enhance the feasibility of printing in concave regions and by incorporating the ability to angle the toolhead in different orientations to the surface to reduce collisions in their papers respectively. However, this restricts the possibility of printing onto complex parts to only using extended nozzles and does not provide the methods to create these toolpaths. Gunipar and Cam present a method to print onto free-form 3D curves in a way that keeps the nozzle close to the surface normal [4], but this does not account for printing onto existing surfaces with deep, concave regions.

One common method of producing valid toolpaths for complex parts is to segment the print into smaller subsections in order to increase printability and eliminate the possibility for collisions [5-7]. This method works well for parts that are printed onto planar build plates where the model is completely fabricated from scratch, but it does not account for printing onto existing surfaces. Additionally, work by Jayakody et al. uses a more continuous approach that still accounts for the printability of concave regions but keeps the model as a continuous region rather than smaller subsections [8]. However, this method is also intended for entirely new fabrications and does not account for the geometry of the print surface.

Further work by Shembekar et al. and Nishat et al. address printing onto existing surfaces by updating the print angle to avoid collisions with the existing surfaces [9-10]. Other work by Lou et al. employs the use of voxels to discretize the model and perform solid boolean operations to check for printability [11]. Our work differs by also determining areas of the desired print that are unreachable and updating the print regions to account for this.

While our work also uses a discretization approach, we use many 2D slices through the existing print surface to determine both infeasible regions and the necessary angle for the print head. A method of using orthogonal 2D slices to determine collisions is shown in research by Tang and Bohez [12], but their research focuses on subtractive manufacturing and therefore does not account for updates to the model in response to unreachable areas.

Finally, research by Plakhotnik et al. shows the use of Directed Energy Deposition (DED) to print metals onto existing curved surfaces while accounting for the surface geometry [13], but work is limited to prints within a single plane and do not account for any potential changes needed to the print area.

Our research contributes to this field by both presenting a new method of collision detection using 2D slicing and a simplified representation of the print head and updating the desired print region in response to print surface information.

3. Methods

This research focuses on demonstrating two main processes used to increase the feasibility of printing onto complex surfaces using 5-axis 3D printers. The first process is referred to as Area Refinement, where existing slices are altered based on the geometry of the print surface, as well as the geometry of the nozzle. If a user is printing onto a print surface that is not a flat bed, portions of the to-be-printed part will likely be too close to the print surface. This step updates those unprintable regions to only include areas that the nozzle can reach. The second process is Angle Determination, where the geometry of the printed part is used to angle the nozzle to prevent collisions from occurring between the printed part and the print surface. This step also can result in a reduction in print area, as the end of the nozzle still may not be able to reach the desired print point due to a combination of the geometry of the print surface and the printer.

This section outlines these two processes, while also describing the prerequisites and intermediate steps needed in order to complete them. The overall process, as well as the sections

 doi:10.20944/preprints202605.1142.v1

highlighted in this paper, are shown in Figure 2. As shown at the top of the figure, the gray boxes indicate required inputs for this work. Next, the "Slice Initial Print Area" box is completed using conventional slicing routines. The first bolded box "Create Offset Slices" is the first contribution of this research and is an important precursor to the Area Refinement and Angle Determination steps, both described in detail later in this section.

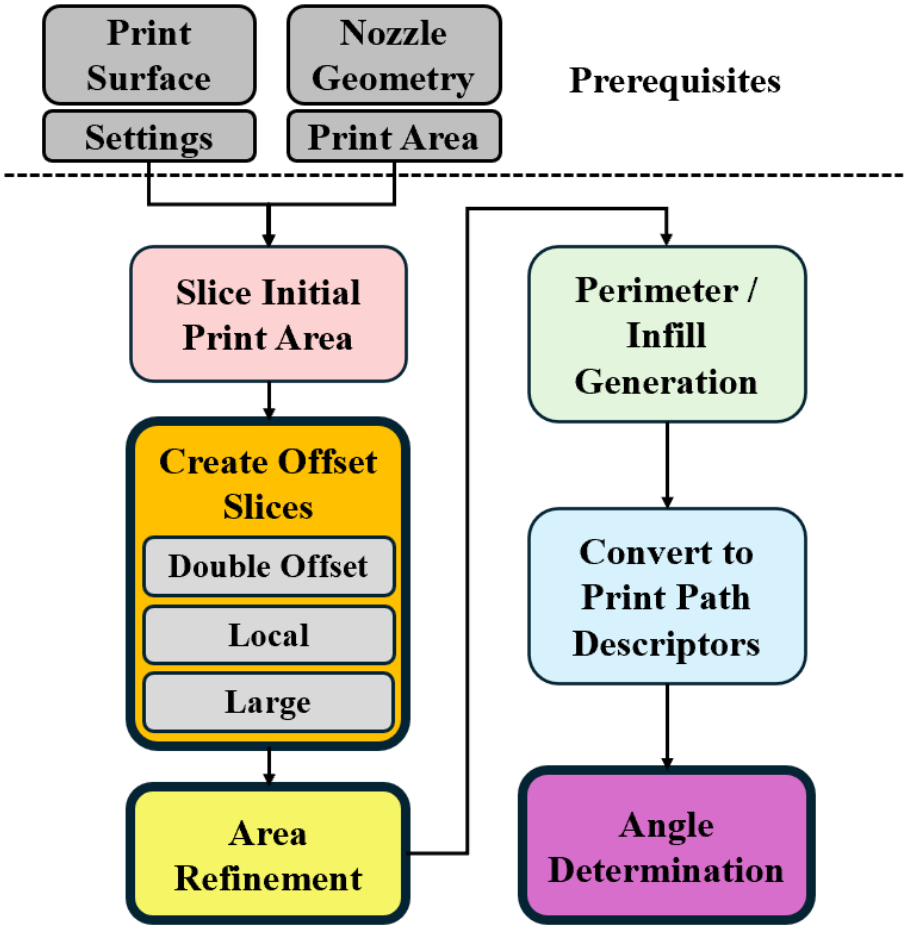

Figure 2. General flowchart of the process depicted in this paper, with the areas of focus highlighted.

3.1. Process Prerequisites

In order to complete the Area Refinement and the Angle Determination, multiple prerequisites processes need to be completed. The first is that the print surface needs to be defined as a solid, which should be in a known position relative to the printer. For conventional printers, this requires a simple model of the thin planar build plates, but the focus of this paper is for printing on more complex surfaces. For the example shown in this section, the geometry that we want to alter is used as the print surface, seen in Figure 3. The initial surface is shown in tan, representing the existing surface before any 3D printing occurs. The image on the right represents the desired final surface with the area to print shown in red.

Additionally, the material to be added needs to be included as a solid model, located at the correct coordinates in the global system. While the current system requires that only the additional material be modeled, it could be changed to a system where the initial condition and final conditions are provided, and the added geometry is found through a Boolean solid modeling operation. The Boolean method would allow for the modification of a scanned print surface, similar to our previous work using scanning and additive manufacturing to fill defects on a surface [14].

 doi:10.20944/preprints202605.1142.v1

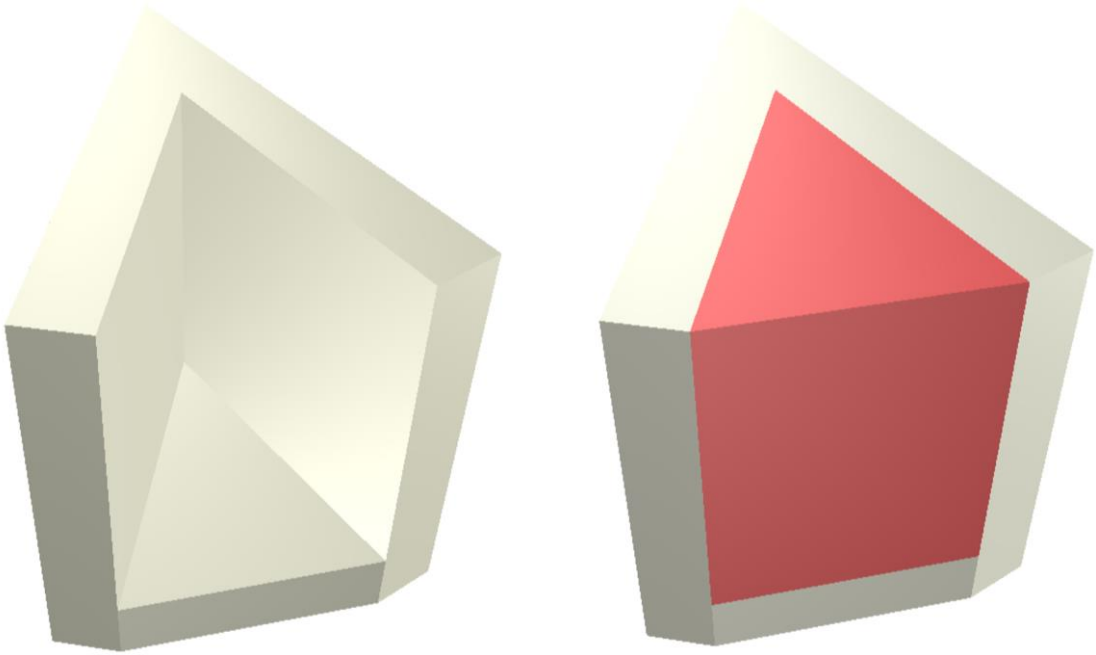

Figure 3. Print surface (left) and print surface with filled region highlighted (right).

In order to effectively and accurately determine where the printer can print onto the existing model, information about the dimensions of the nozzle are needed. The current system requires information about the height and aperture of the nozzle cone, as well as the distance to the next joint. Because we are using a modified 3-axis printer (shown in Figure 4), anything above the top of the extended nozzle is considered to be infeasible space, as the 3D printing gantry poses interference issues (see the gantry interference zone in the enlarged image). Note the red cone in our current model showing the representation of the nozzle, and the blue arrow showing the clearance to the rest of the print head assembly. Therefore, we define the top of the nozzle as the maximum plunge height of the nozzle. However, this could be modified in future work by representing more aspects of the kinematic chain as links, which would allow for better planning on systems such as robotic arms.

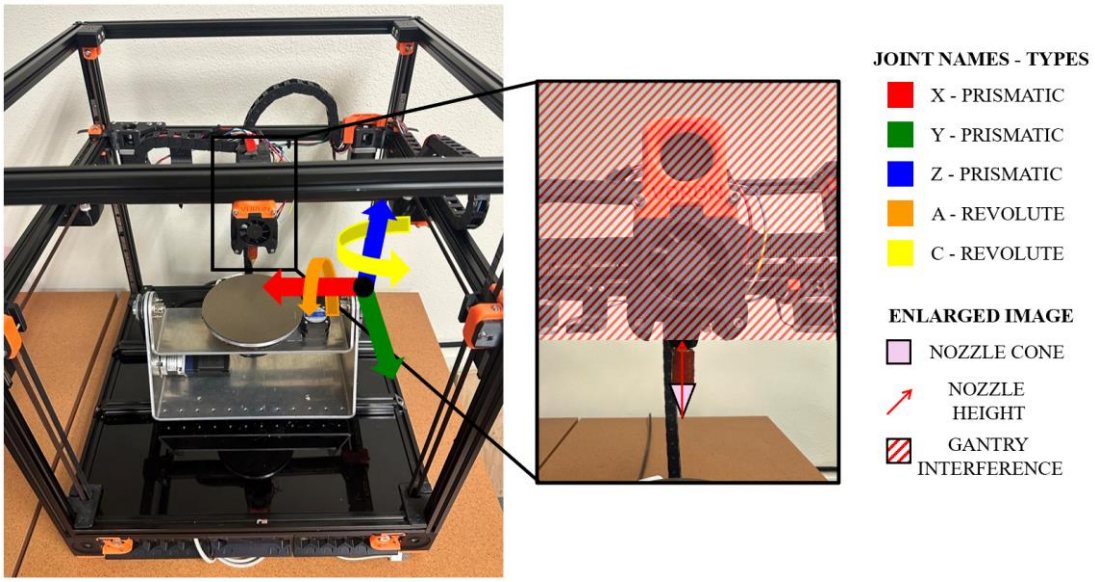

Figure 4. Parameters for the code in the methodology section based on the above printer, a 3-axis printer modified to include two rotational axes.

 doi:10.20944/preprints202605.1142.v1

Finally, there are adjustment parameters to give additional clearance for various measurements within the system, in order to give more or less conservative estimates of the feasible print area. While these are intended to not need modification, the user could adjust them based on the accuracy of their machine, and how close they want to print to the existing part.

3.2. Function Prerequisites

Prior to the Area Refinement step, the area to be printed needs to be sliced. This is currently achieved by creating planer slices through the area based on a provided layer height from the user, akin to how traditional 3-axis machines slice a 3D model. The initial slice direction through the area to print is aligned with the XY plane, which sets the default print angle as aligned with gravity. This also means the default angle of the nozzle is perpendicular to the build plate. However, this direction could be changed depending on the printer orientation and configuration, in order to change the orientation of the slice planes. The initial slices are shown in Figure 5.

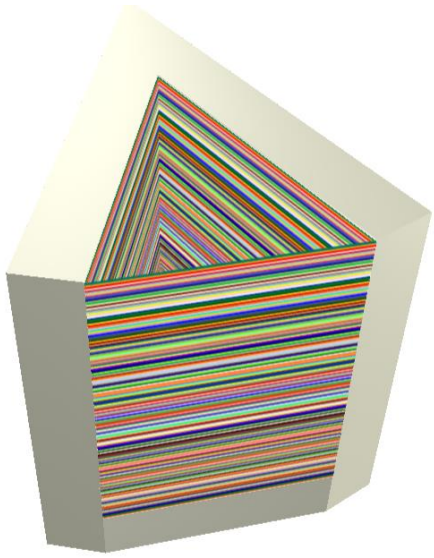

Figure 5. Initial slices through the print region, aligned with the Z axis.

The overall process in this paper largely relies on the creation of additional 2D slices through the print surface to determine concavity information of the overall model. The first set of slices are directly aligned with the initial slice planes through the print area, which simplifies the process of altering the regions in response to concavities on the print surface. These slices through the print surface are altered by applying a double offset, which is a method that modifies the original polygon to have rounded corners inside concave regions, as shown in Figure 6. This works because the adjacent lines being offset are joined with a radius if they no longer meet after the offset. For concave regions, the process of offsetting the shape out and back causes the concavities to become rounded, as demonstrated in Figure 7. The specific radius value that is used is based on the width of the nozzle cone. This step simplifies the initial area refinement step, as the in-plane concave regions can be altered using polygon subtraction.

The next step is to slice through the surface model along the orthogonal axes of the printing plane, generating further information about the concavities of the overall model. While the individual slices are in two dimensions, the combinations of these slices along different planes are used in to gain information about the geometry of the overall model, while existing as 2D regions that are easier to use for manipulating the existing print paths. This discretized version of the model with offsets is ultimately used in both the Area Refinement and the Angle Determination processes. The discretization helps to simplify later steps, as the analysis is conducted on 2D slices, rather than the entire 3D model. Each of the slicing orientations for the double offsets are shown in Figure 8. The current method uses slicing in three orientations but could be expanded to include slicing along off-

 doi:10.20944/preprints202605.1142.v1

axis planes in the future to get more information about concavities in the model. This potential update is expanded upon further in the Future Work section. The current slicing algorithm gives a conservative estimate of where material can be deposited in order to not print in infeasible regions.

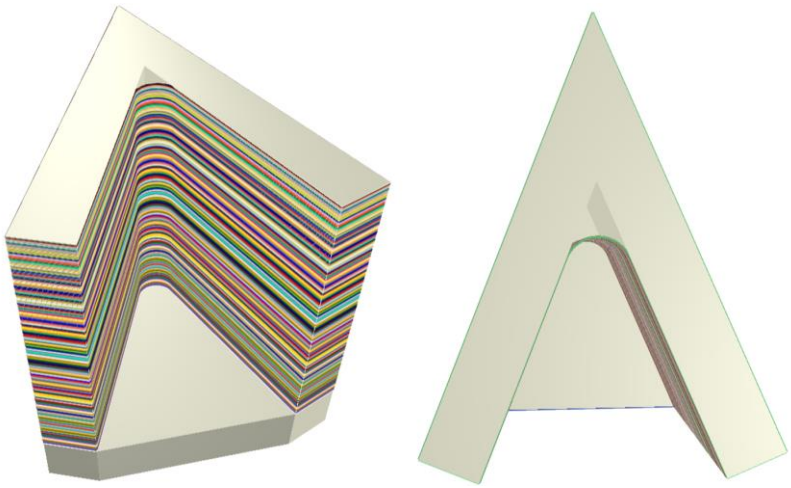

Figure 6. Print surface shown with rounded corner along demonstrated concavity.

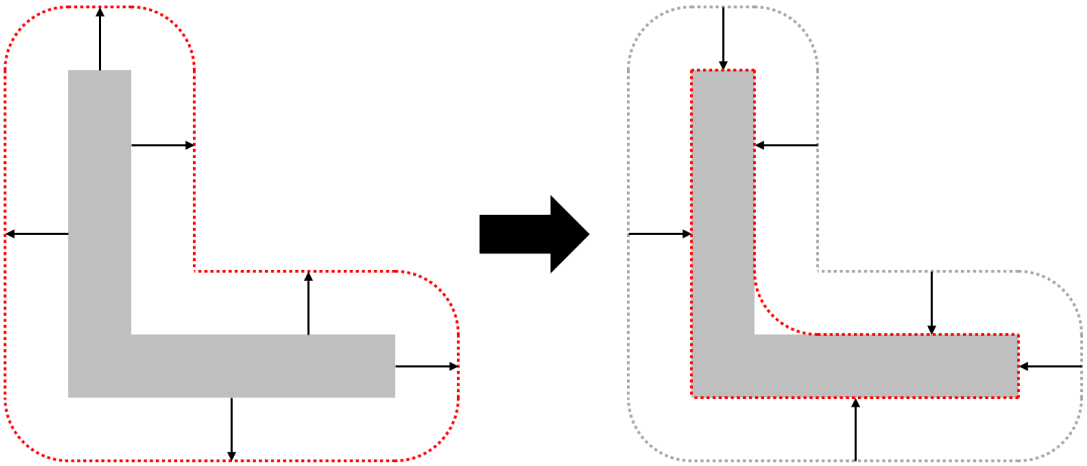

Figure 7. The double-offset process resulting in reduced concavity.

The process of creating discretized slicing planes in three orthogonal orientations is completed at three different levels. The first is the double offset, as was discussed previously. While useful for detecting concavities, the double offset also is useful to determine regions of the model that are too close together for the nozzle to print into. The double offset gives the region that the nozzle cannot reach, while preserving other areas of the region that are feasible, as shown in Figure 9. This set of slices is used for the Area Refinement step, as it shows portions of the print surface where the geometry is too small for the print nozzle.

The second set of slices created are the local offset. These are based on the geometry of the nozzle cone and are used to determine the angle that the nozzle can tilt before colliding with the existing surface. Figure 10a shows the two measurements that define the cone, the angle, θ, and the height, h. Figure 10b then shows how those two measurements can be used to determine the maximum angle of tilt, denoted as the local offset, LO. This is the value that is used for the local offset. Next, Figure10c shows the resulting triangle used to determine the local offset value, shown in the following equation:

$$Local\ Offset = Nozzle\ Height * sin\ (\theta) \quad (1)$$

 doi:10.20944/preprints202605.1142.v1

Finally, Figure10d shows how applying the offset h results in both a feasible zone and an infeasible zone. The latter of which is determined if the end of the normal vector of the cone is inside the infeasible zone; the nozzle has tilted too much. These do not use the double offset method, and instead just offset outwards to create an exclusionary zone around the model that is used for Angle Determination, as shown in Figure 11.

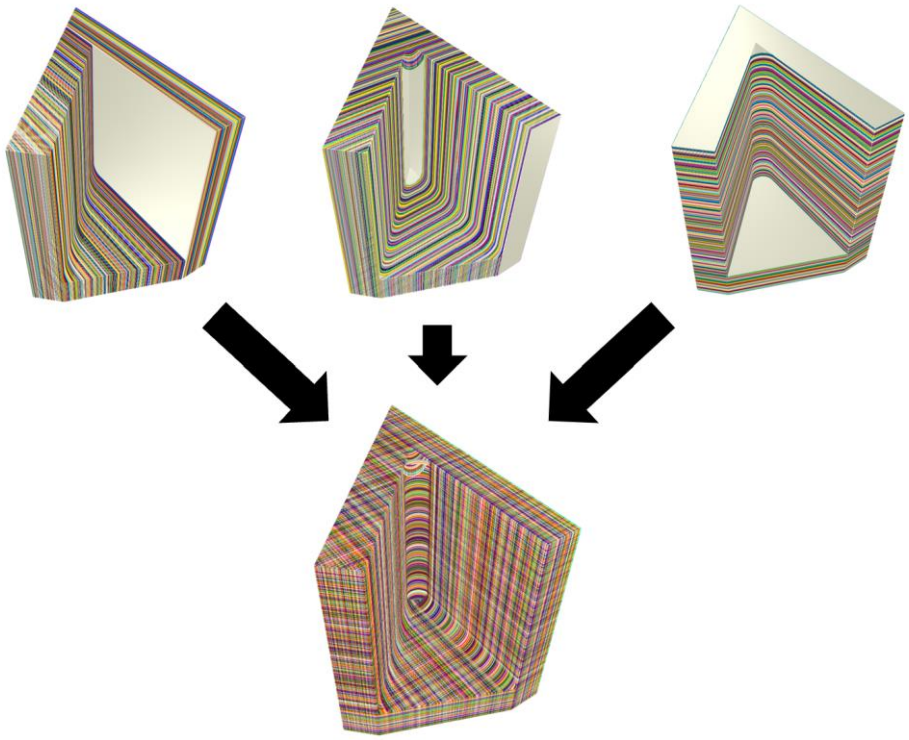

Figure 8. Each set of slice planes (in X, Y, and Z) and the combination of slices used to approximate the 3D object.

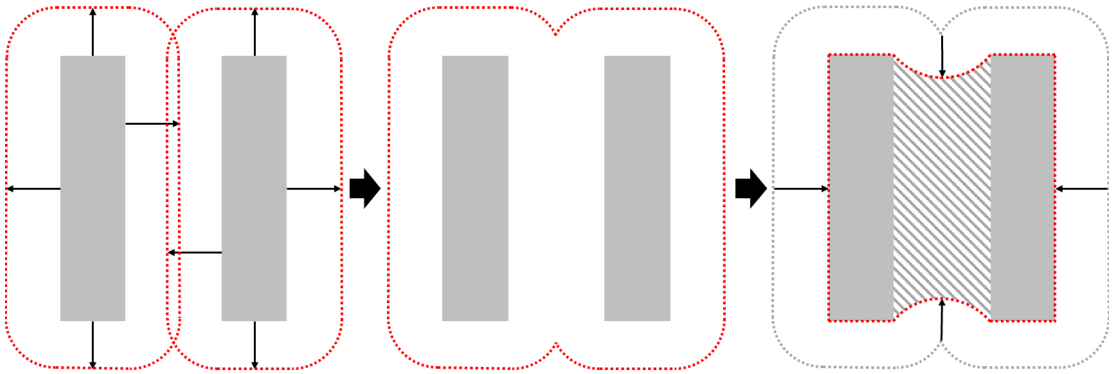

Figure 9. Process showing how a double-offset can be used to determine infeasible regions with a model.

The final set of slices is the large offset, representing a no-go plane for the printer. In our prototype printer (shown earlier in Figure 4), this represents the region that the gantry inhabits and may not be required for all configurations of printers. This is also a single offset operation and is based on the distance from the tip of the nozzle to the top of the nozzle, shown in Figure 12. Now that the various polygon offsets have been defined (as shown in the orange box in Figure 2), we move to the Area Refinement operation.

 doi:10.20944/preprints202605.1142.v1

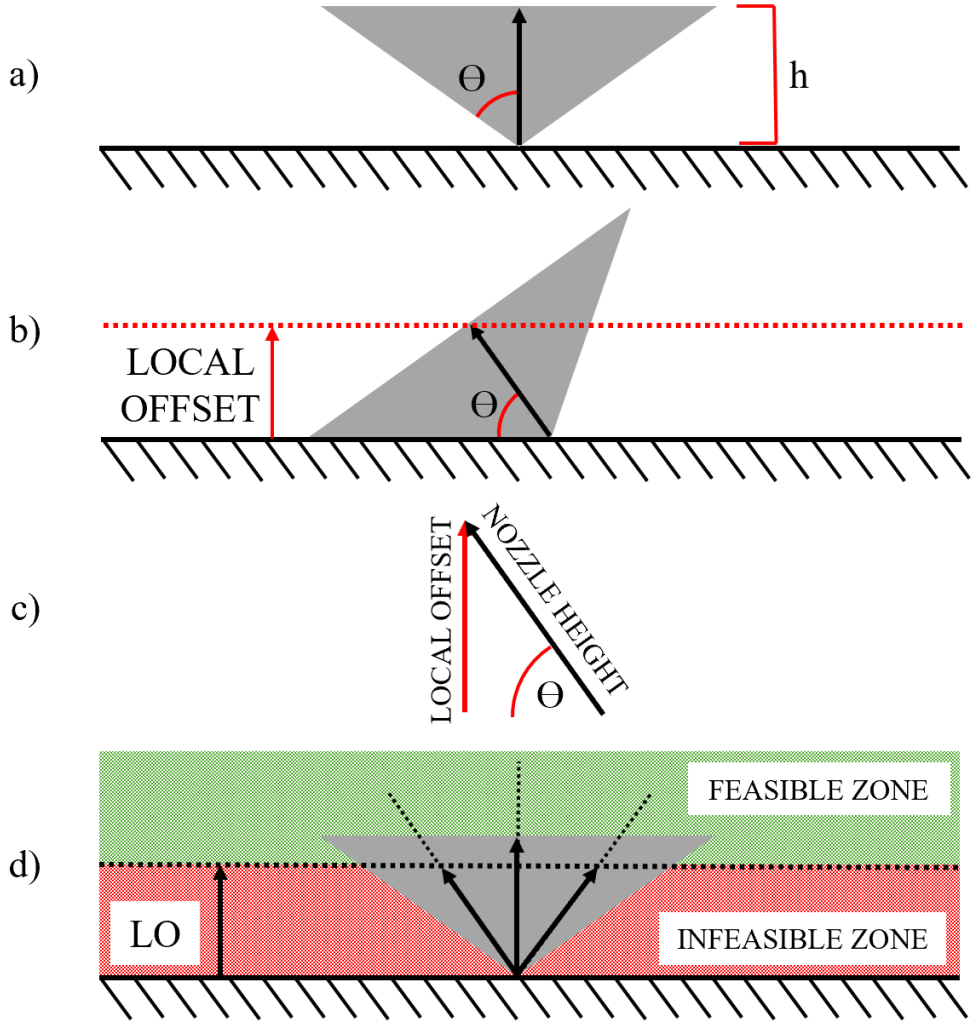

Figure 10. Maximum nozzle tilt demonstrated, used to determine the local offsets.

Figure 11. Surface model with local offset values applied.

 doi:10.20944/preprints202605.1142.v1

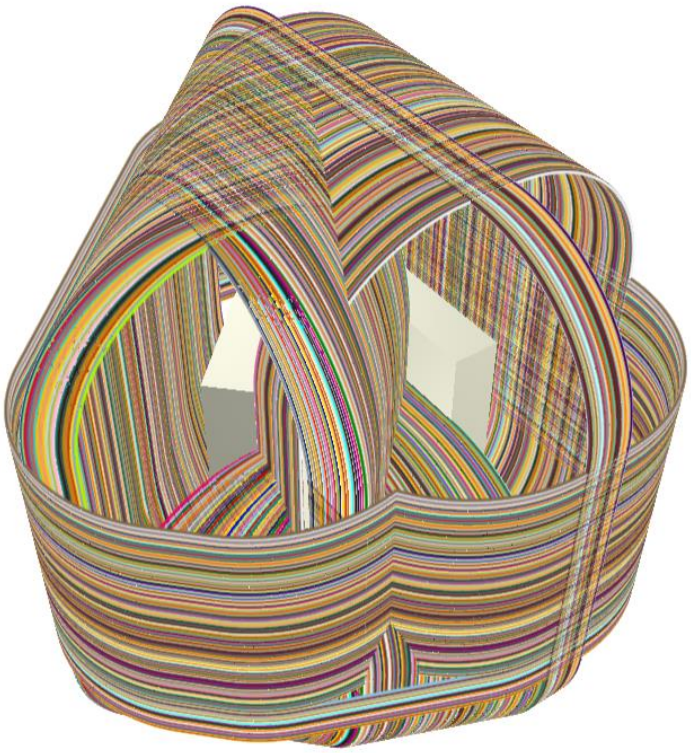

Figure 12. Surface model with large offset values applied.

3.3. Area Refinement

The Area Refinement process starts with the list of current print layers, as well as the list of double offset slices along each plane. Each print layer is composed of regions, showing the individual polygons that need to be printed within each single plane. However, if these regions are within a concave region of the print surface (especially in the cases of filling a hole or a void), then the nozzle will likely not be able to reach everywhere on the print region, as demonstrated in Figure 13. The Area Refinement step alters the print regions within each print layer based on information from the surface of the model (generated through the double offset slices) and the geometry of the nozzle. This alteration affects the polygon outlines of the print regions, which will later be converted into perimeters and infill.

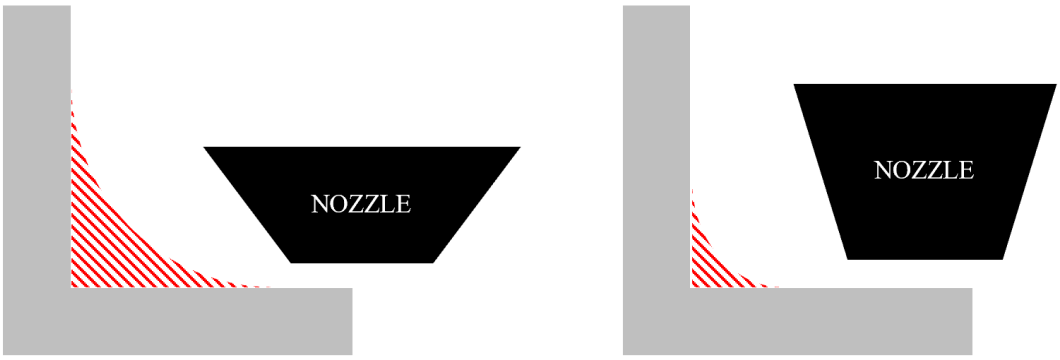

Figure 13. The angle of the nozzle cone largely affects the feasible area of printing in concavities on the surface model.

 doi:10.20944/preprints202605.1142.v1

The main Area Refinement function operates on each individual print layer. While the process applies to each point individually, the movements for points that need to move are informed by other points within the layer. The overall flowchart for the Area Refinement process is detailed in Figure 14 and will be elaborated upon throughout the rest of this section.

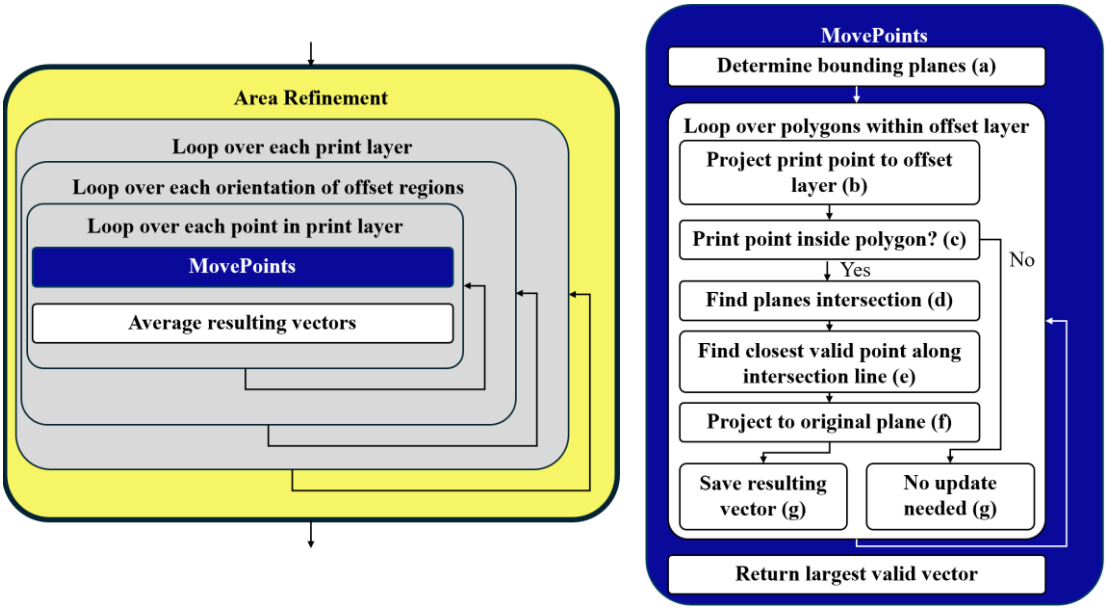

Figure 14. The overall flowchart representing the area refinement process (left), and the movePoints function shown in more detail (right). The various processes depicted in this flowchart will be discussed throughout the remainder of this section. Some subsections of movePoints are labelled to be referenced in the test of this paper.

The specific per-point operation takes every single point that is set to be printed and determines if it needs to be moved in order to be a valid point. It does this by checking if the current print point is within the boundaries of the double offset polygons. If so, then that point is too close to a concave or restricted region and needs to be updated.

The original process for updating print points, as shown in Figure 15, was to find the closest point on the polygon, and move the print point to that point. Then, the updated print point was projected back to the original print plane. However, this often resulted in new points that remained within the infeasible region, due to the projection back to the original print plane. Because of this, we use the expanded MovePoints function (highlighted in Figure 14) to determine how to update the position of print points. This function is applied to all points in the desired print path, including both polygon outlines and hole outlines, as those define the areas of the printed shape. Therefore, all points in these areas can be impacted by the geometry of the print surface.

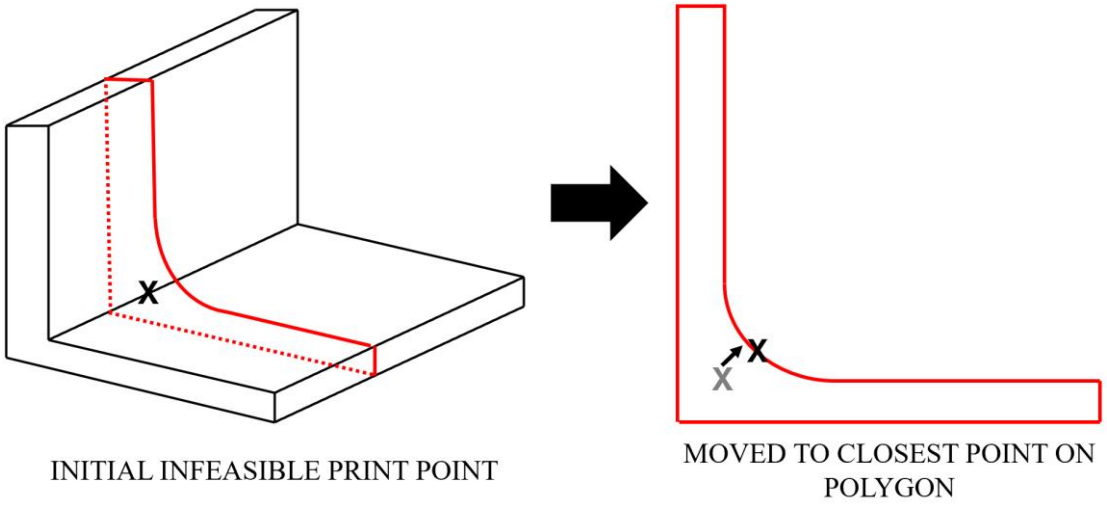

 doi:10.20944/preprints202605.1142.v1

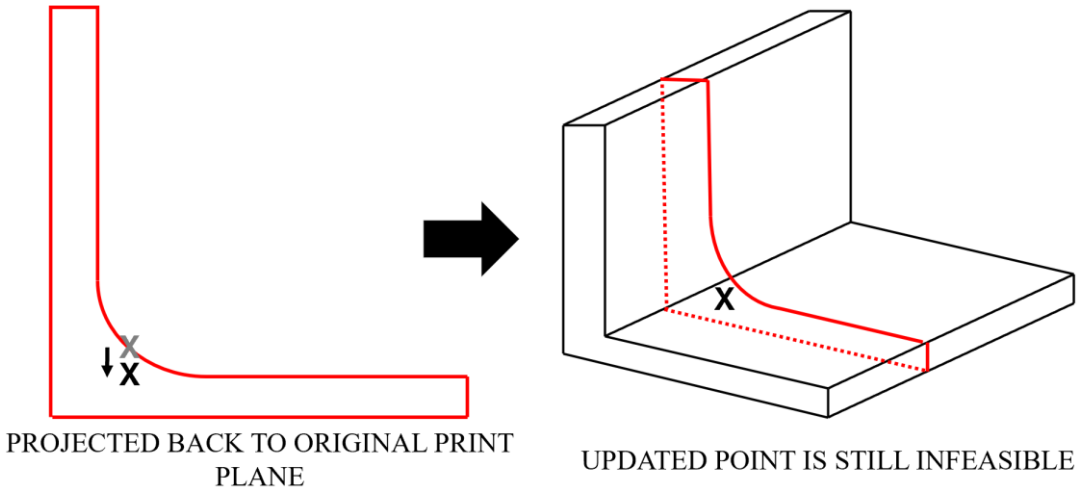

Figure 15. An initial infeasible point is projected to the closest point on the polygon, then back to the slice plane, resulting in an updated point that is still infeasible.

The MovePoints function begins with the specific point being considered and the list of slices that create the exclusion zone along one orientation of slices. The actual slice planes that have any impact on the considered point need to be determined, as shown in Figure 16 (and mentioned in Figure 14a in the flowchart). There can be at most two planes that are impacting the current point (the two planes bounding it on either side). In order to determine these planes from the list of all possible planes, we find which planes have the smallest distance to the point being considered. Only the closest plane with a positive distance to the print point and the closest plane with a negative distance are analyzed.

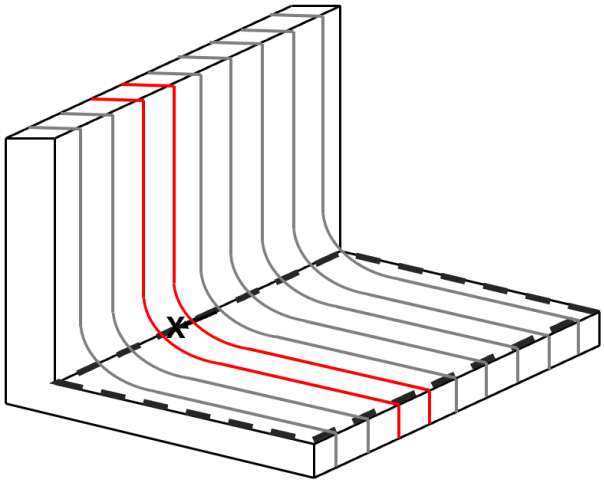

Figure 16. The two closest (one on each side) offset planes are found that bound the print point.

The following process for moving an invalid point outside of the infeasible region is shown in Figure 17. For both of the slices bounding the print point, we loop over each polygon within the offset slice to see if the print is inside any of the offset polygons. While the slices are close together because their distance is based on the provided print layer height, it is still important to project the print point to the slice planes in order to get the best positional information about the point (Figure 14b). This allows the print point to be analyzed to see if the point is internal to any of the polygons on that slice (Figure 14c). If it is, then the point needs to move to be outside of the polygon. Because the polygons represent areas of infeasibility, moving the print point outside of the polygon makes it valid. The movements are restricted to only be within the initial print plane. This is determined by finding the line on which the print plane and the exclusion planes intersect (Figure 14d). Then, the print point is

 doi:10.20944/preprints202605.1142.v1

moved to the best possible point of the bounding polygon that intersects the newly found line (Figure 14e). This ensures that the following two conditions are met: the new point is outside of the exclusion zone, and the new point is still valid within its original print plane. The first condition is verified because the intersection points of the exclusion polygon and the line are defined by the limits of the exclusion polygon and are therefore valid, and the second condition is verified because the line of intersection between the two planes is within the plane of the print plane, so the newly found point does not need to be projected back to the print plane.

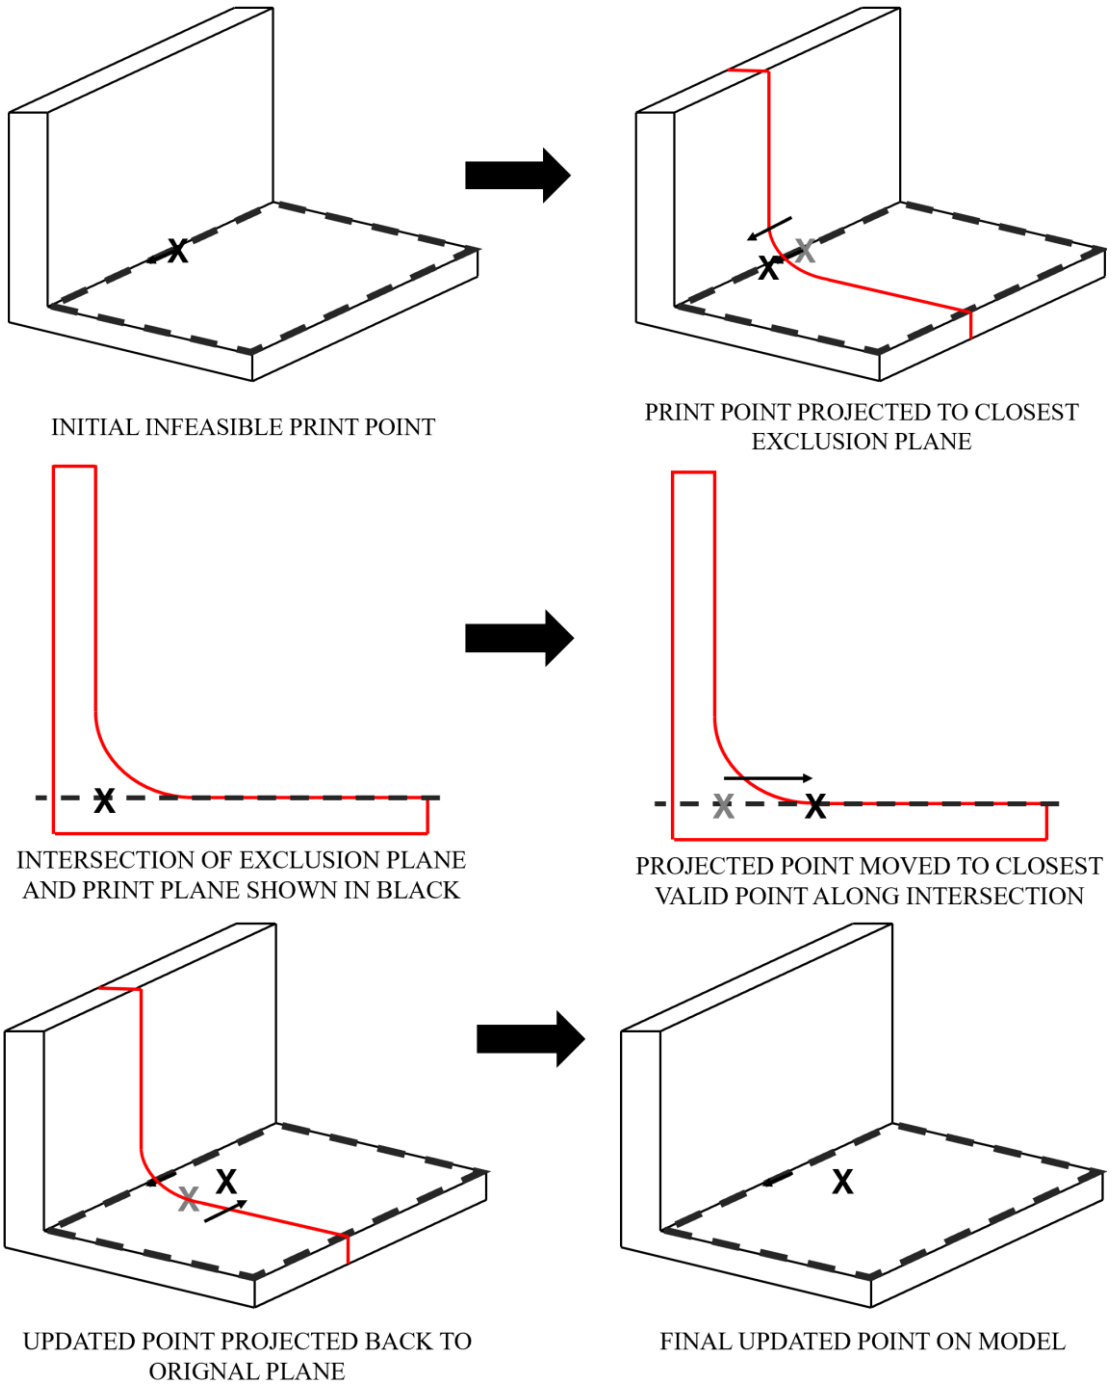

Figure 17. The process of determining how to adjust a print point inside of an exclusion zone. First, the print point is projected to the nearest plane (top). Then, the projected point is moved to along the print plane to the closest valid point (middle). Finally, the point is projected back to the original plane (bottom).

After determining the intersection points between the exclusion plane and the original plane, the optimal intersection point needs to be determined. While this is often the closest intersection point

 doi:10.20944/preprints202605.1142.v1

to the original point, this is not always true. The closest intersection can sometimes result in an infeasible move, as shown in Figure 18. In this example, the closest point would result in a move through the print surface, which is infeasible. Instead, the point needs to actually move to the further point, which is valid. To determine the best point to move to, each intersection point is considered. The area of the printed polygon is then calculated using the updated point, and all other current points in that region. For this reason, information about the entire region is necessary for the Area Refinement step. The best point to move to is ultimately the point that maintains the largest polygon on the original slice plane. This is true because moving the point must make the original polygon area smaller, or the move is not valid.

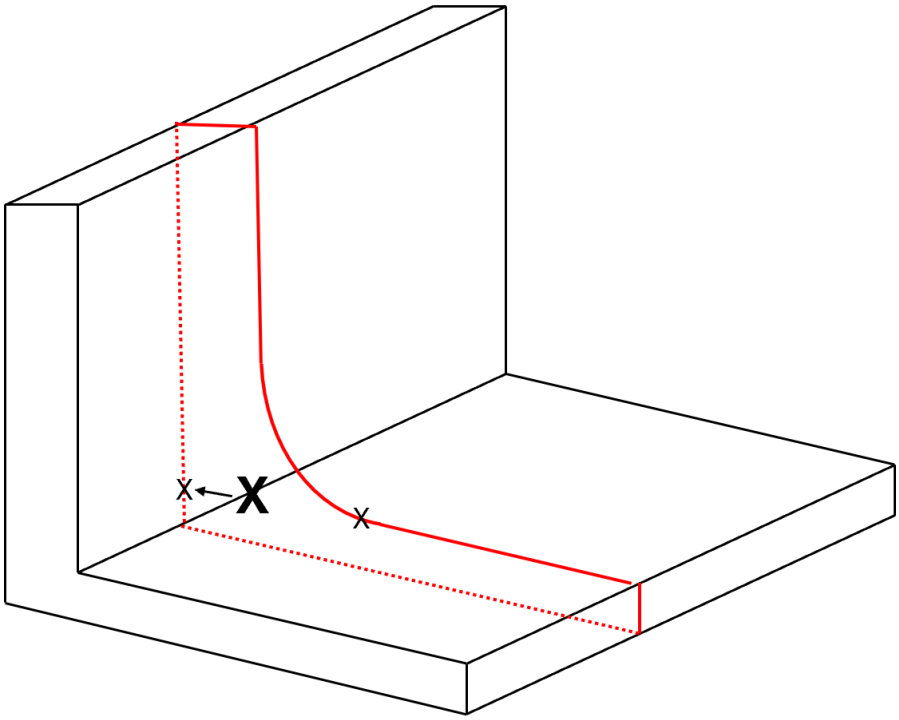

Figure 18. The closest point is not always valid, as it will sometimes result in an attempt in the wrong direction.

This process of determining the intersection points and finding the best point to move to is done for both adjacent exclusion slices. This process is only done if the point of interest is within an exclusion zone on both sides. If it is only constrained on one side, then the point is valid. This allows for printing onto walls of parts, where the print point is internal to the exclusion polygon on one adjacent slice plane, but not the other. If the point falls within the exclusion zone on both sides, then it needs to be updated using the process described previously. Ultimately, this process determines two possible points to move to – the best point for each adjacent plane. The point that is ultimately chosen is the move that is furthest from the original point and is projected back to the original plane (Figure 14f), or the original point is returned if no moves are feasible (Figure 14g).

Following this process, the print point from the original polygon is updated to the new location. Each offset slice only represents a 2D exclusion zone and can therefore only move the selected point along that plane. However, in many cases the point will need to move in multiple directions in order to become valid. To account for this, the possible positions of where the point can feasibly move along each plane are stored, and the final position is determined using an average of the vectors. This is done by using vector addition on the possible moves, and then normalizing the resulting vector based on the vector weights of the valid movements. For example, consider the corner point shown with an "X" in Figure 19. It is initially positioned inside the infeasible region, shown by the dashed red

 doi:10.20944/preprints202605.1142.v1

outline. Possible movements outside of the infeasible region are found along each slice orientation, and then they are combined and scaled by averaging the vectors to find the resulting movement vector.

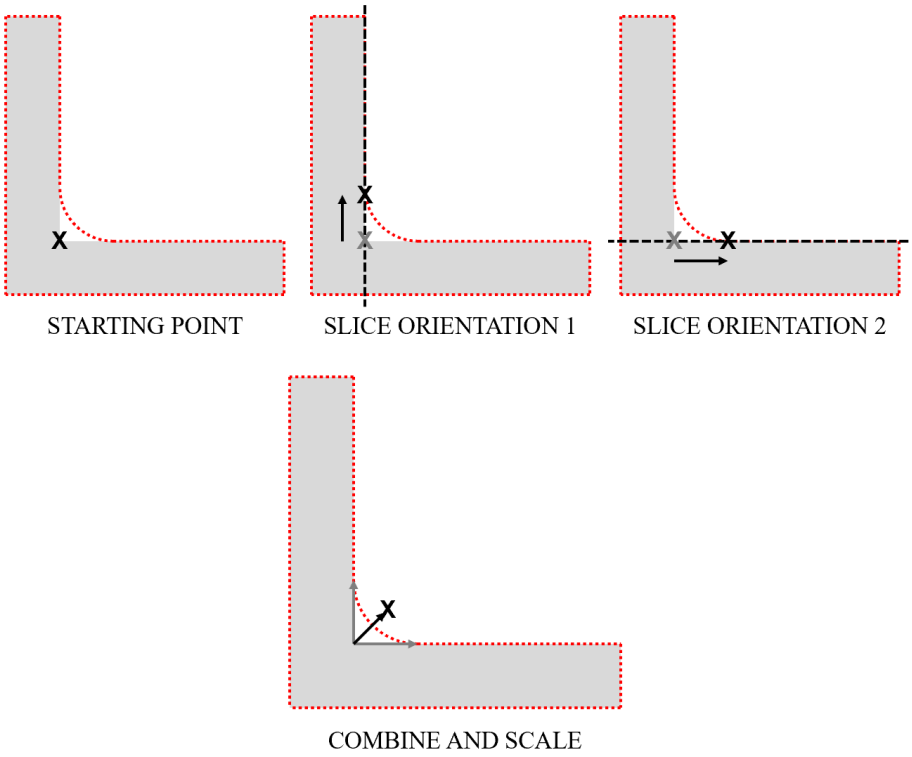

Figure 19. The process used to reconcile the point adjustments from each 2D plane.

This process can be parallelized for each slice or for the points within an individual region, although parallelizing the individual points requires a copy of the initial locations to be made, as the step of using the updated polygon area to determine the best updated point requires the positions of the original points for comparison.

At this point in the method, the print regions are all updated and reflect a feasible solid that can be printed based on the print surface geometry. However, the feasible region still may be updated further in the Angle Determination step, as the Print Area Refinement does not ensure that the given nozzle can actually reach the point. The updated regions from the Print Area Refinement step depend on the geometry of the nozzle used as is shown in Figure 20.

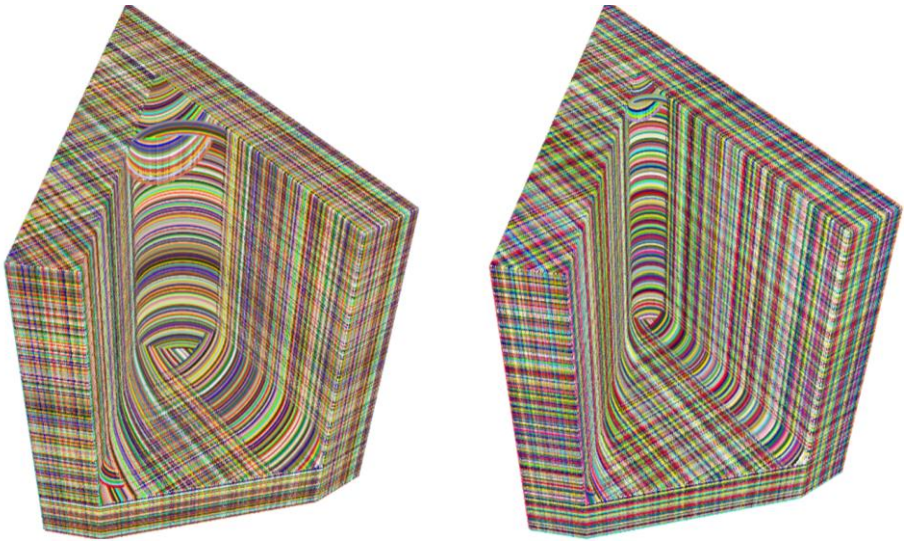

 doi:10.20944/preprints202605.1142.v1

Figure 20. The impact of printing with different nozzle width parameters on the feasibility regions.

3.4. Intermediate Steps

Once the Area Refinement step is complete, the resulting print areas need to be converted into Print Path Descriptors. These are objects in the code which contain the position of the point being printed, the angle of the nozzle at that point (represented as a 3D vector), and the extrusion value. The Print Path Descriptors are generated based on the current print area information, as well as printer and printing parameters provided by the user like the layer height and flow factor. Once these are generated, they can be used in the next step to update the angle of the nozzle to allow for valid printing configurations.

3.5. Angle Determination

Following the generation of the valid print areas, we convert each polygon region into a print path, which is an array of Print Path Descriptors. For most printers, including in our test case, we initially assume that the Print Path Descriptor angle is in the unit Z direction (0, 0, 1). This ensures that for printing on traditional surfaces with a build plate, the default angle points directly away from the build plate. Additionally, this aligns the initial print angle with gravity, even for printers that do not have fixed orientations of the nozzle.

The process behind the Angle Determination is similar to the Area Refinement process. Sliced, 2D exclusion zones are used along with the points making up the print path to see if they are valid, or if they need to be updated to become feasible. The initial Print Path Descriptors are aligned with the Z axis, but this orientation is not feasible for printing into many concave regions.

The first way that this process differs from the Area Refinement process is that it validates the normal of the print path descriptor, not the position. The tested vectors use the Print Path Descriptor angles, with magnitudes equal to the length of the nozzle cone. This helps to determine the amount that the nozzle can be angled before it collides with the existing print surface, as previously depicted in Figure 10. If the end of the vector is outside of the local offset area, then it is valid.

This process loops over every point in the entirety of the print path. This can be parallelized for every single point, as this step does not require information about any other points within the same region, as it is only updating the printing angle for each point.

This process mirrors the previous section, where the two adjacent exclusion planes are found closest to the point, and the point of interest is projected to both planes to see if it is internal to the offset polygons. If so, then the Print Path Descriptor needs to be updated to angle the vector in a feasible direction. This is done by finding the closest point on the exclusion polygon to the end of the Print Path Descriptor. Because the polygons used to determine the exclusion zones for this step are single offsets, the closest point on the polygon will always angle the vector away from the current print surface along the selected plane, as shown in Figure 21. However, in some cases the Print Path Descriptor will need to be angled only along one of the planes, while a different set of exclusion zones might try to produce an angle that is infeasible and still produces collisions. To validate a candidate angle, the length of the vector from the actual print point to the updated point is compared to the original length of the vector, and if it increases then the move is considered invalid. If invalid, the Print Path Descriptor is returned to its original angle from before the move, and the next slice orientation is tested. If invalid and no other slice orientations remain, then the Print Path Descriptor is removed. An example of an updated print orientation is shown in Figure 22. In the left image, the original Print Path Descriptor is shown, aligned with the Z axis. However, due to the height of the walls and concavity of the part, this is not a valid print orientation for the given point. The results of the Angle Determination are shown in the right image, where the orange vector shows the Print Path Descriptor angled towards the opening of the part.

 doi:10.20944/preprints202605.1142.v1

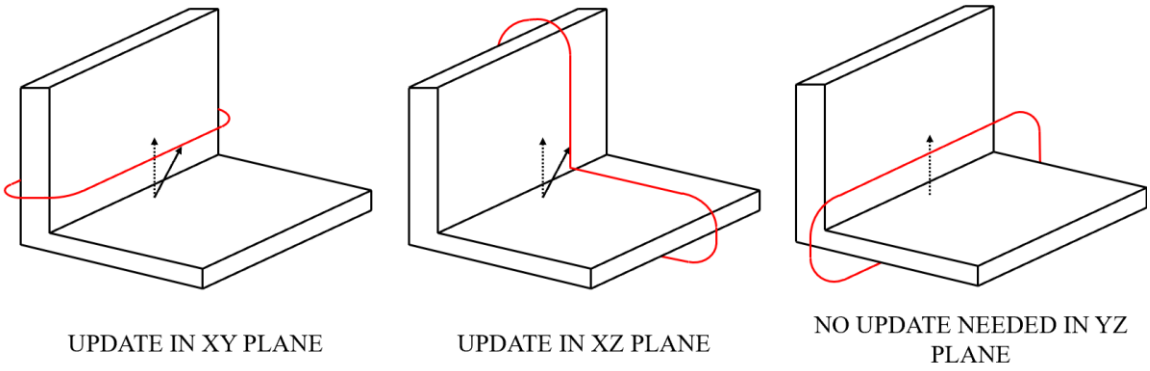

Figure 21. Original vector (dashed) and possible updated vector (solid) as a result of moving to the closest point on each slice plane.

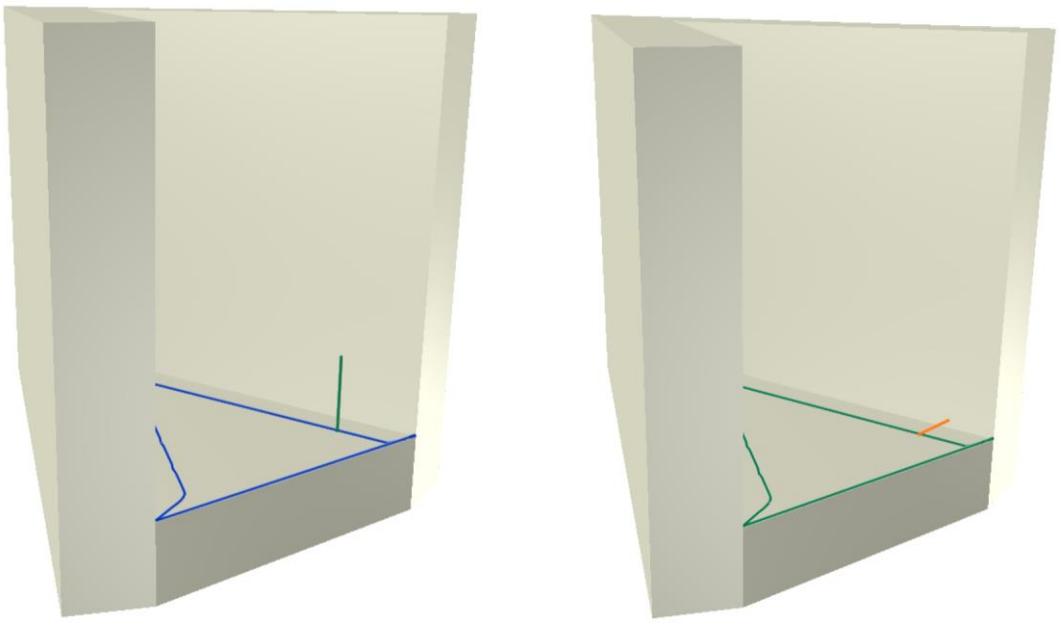

Figure 22. Visualization of Print Path Descriptor angle shown as a green vector before Angle Determination (left), and shown as an orange vector after (right).

Finally, the large offset values are used to determine if the Print Path Descriptor is too deep into the part to print. This step uses the large offset values and the total nozzle height. In this case, if the resulting vector ends outside of the offset regions, then the plunge depth is too large and cannot be printed at that point, causing it to be removed from the print path.

4. Results

The following results were generated based on parameters from a modified Voron 2.4 printer, as shown in the grey input boxes in Figure 2. The modified printer contains two rotational axes, converting it into a 5-axis printer. The computation for this project has no dependencies on existing slicers and was developed on top of previous research in the lab where multiple C# projects build upon a lightweight geometry library.

The results from this method show the overall Area Refinement process being applied to every layer in the test model, as shown in Figure 23. As can be seen, the Area Refinement step resulted in reductions to the printable areas along the most concave regions. Figure 24 shows the results of the Angle Determination for the first layer, and the Angle Determination for the 50th layer. The Print Angle Adjustment step also removed infeasible points if they were too far recessed into the model.

 doi:10.20944/preprints202605.1142.v1

Because the first layer is more recessed, a much larger portion of it is determined to be unprintable with the current printer and nozzle configuration, while the 50th layer requires angling of many of the Print Path Descriptors, but is ultimately printable. Only the perimeters of the model are shown for clarity, but this process can be applied to the infill as well. Although, by their nature infill results will show less reduction in the adjusted angles. The area refinement can work for infill, but it is more efficient to first reduce the perimeters then determine infill on the new results. Applying Area Refinement and Angle Determination to the perimeters for every layer of this model took an average of 15 seconds and 34 seconds respectively.

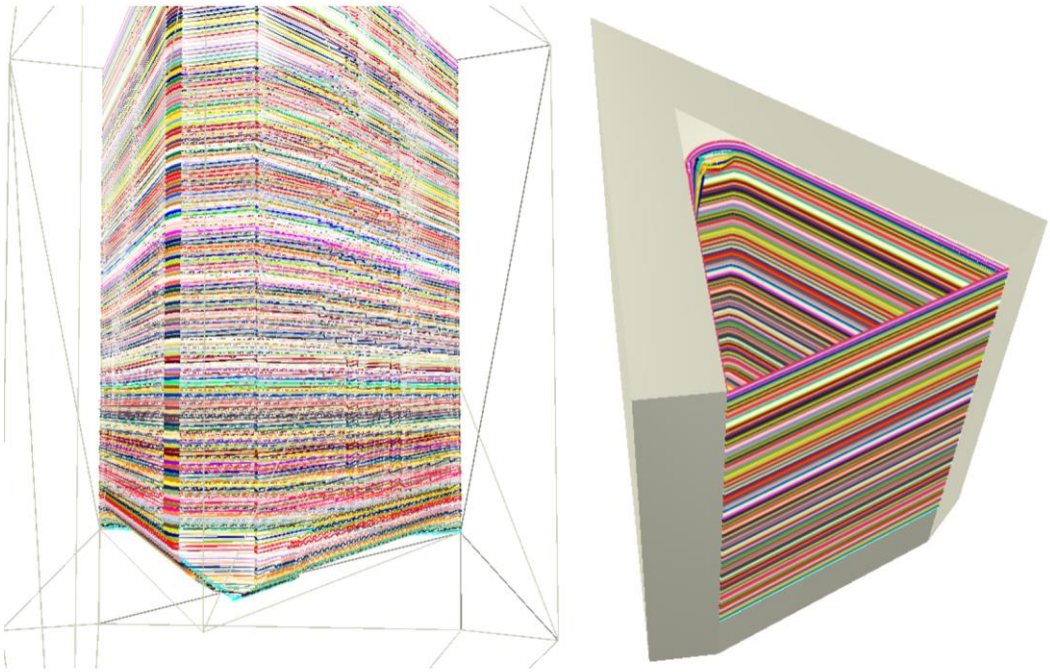

Figure 23. Results of Area Refinement, showing material removed near the corners.

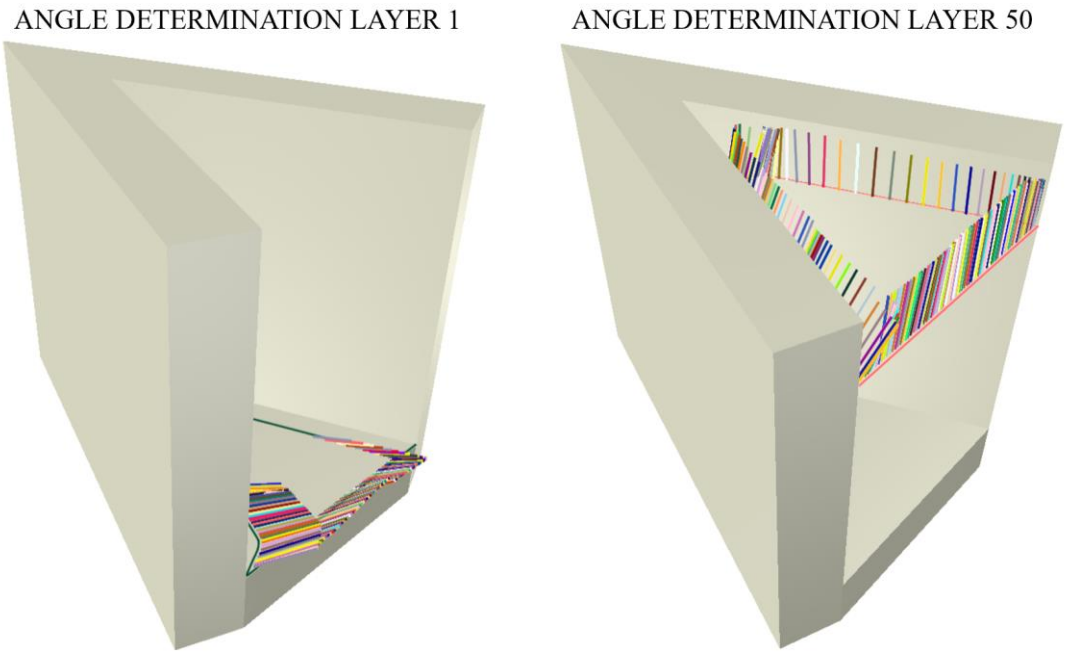

Figure 24. Results of Angle Adjustment showing updated print angles and infeasible points removed (left), and updated angles with no points removed (right).

 doi:10.20944/preprints202605.1142.v1

Another test that was conducted was a shape made up of three distinct regions, shown in Figure 25. The first region is a 90-degree corner to represent printing in a concavity, similar to the first example shown. The second region contains a curved wall, and the third region features a low, 1 mm tall wall that divides the shape into quadrants. This low wall is used to test the ability of the method to recognize that angling the nozzle is not needed for this portion of the model. While the print area covers most of the flat portion of the print surface, the areas of interest are the closest areas to the obstacle in the center of the print surface.

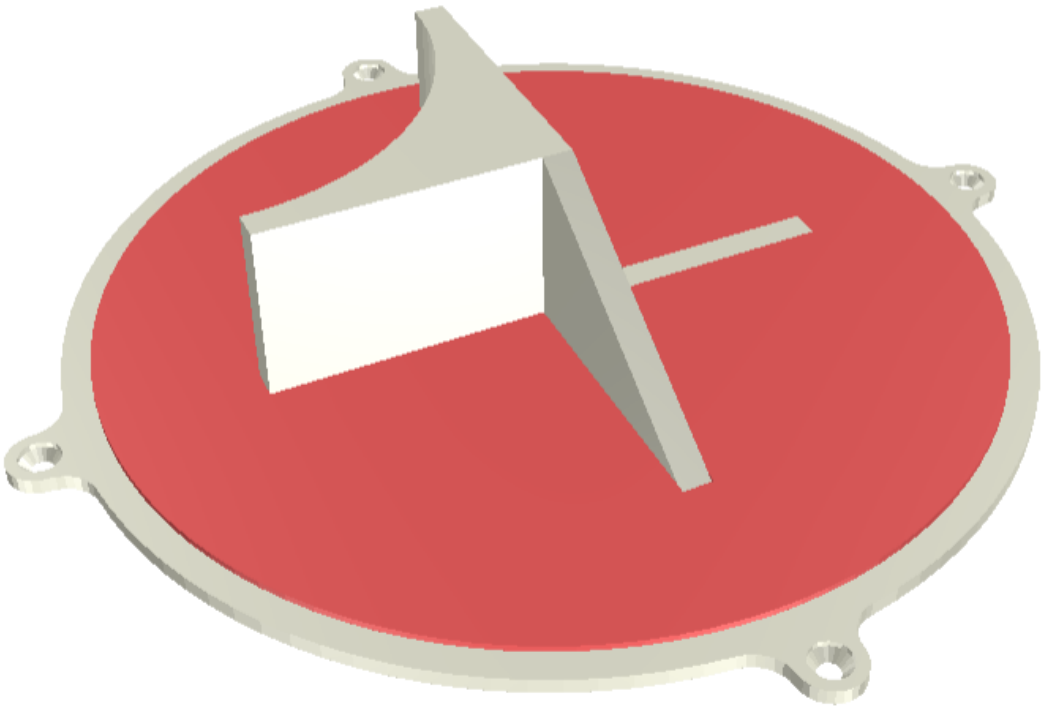

Figure 25. The test surface showing the print surface (tan) and the single-layer area to print (red).

This test shows how – for points far from the obstacles – both the print path and the print angles are not updated, as they are already feasible (shown in Figure 26a and Figure 26d). Along the curved portion of the model the print area is updated uniformly, and the updated angles of the Print Path Descriptors curve along with the surface, shown in Figure 26b and Figure 26e. Finally, the corner section also results in a fairly uniform offset, as shown in Figure 26c and Figure 26d, but the angle of the nozzle becomes more aligned with the z-axis as the print gets closer to the bottom of the ramped portion. This test shows how the process responds to a variety of geometries, such as interior corners, concave curves, sloped walls, and planar regions. Currently, the process produces some artifacts (see Figure 26a and Figure 26b) in interior corners as a result of only using three orientations of slice planes (and only two that are not co-planar with the original) to determine concavity information. This is a topic discussed in the Future Work section. These results show how this process is more useful than just offsetting based on existing geometry, as there are cases where the print path can be closer to the existing surface, such as at the bottom of the ramped section, where the print area moves closer to the print surface obstacle. For the perimeters of this model, Area Refinement took an average of 10 seconds, while Angle Refinement took 6 seconds, although this model only has 3 layers.

The final test that was completed was a bracket piece, shown in Figure 27. While the filled portion of this example is similar to the concave pocket of the previous, this surface model has a lot more complexity. This test shows how only the geometry near to the print area should ultimately affect the final print area and print orientation. This test was used to validate that the process still works even with a more complex part geometry. The results of the Area Refinement and Angle Determination processes are shown in Figure 28 and show how the print area is refined based on the

 doi:10.20944/preprints202605.1142.v1

print surface, and how the Print Path Descriptors angle to give better access to the part. For this model, Area Refinement took 29 seconds and Angle Determination took 1 minute and 45 seconds.

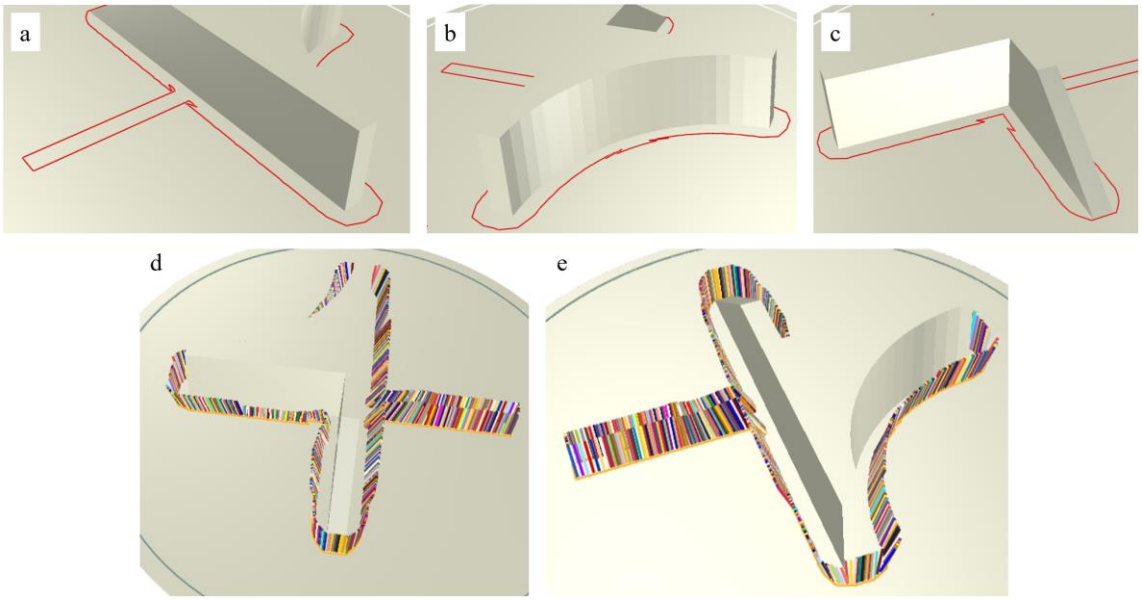

Figure 26. The results of Area Refinement, showing the effects of different geometries on the resulting area (top) and the results of Angle Determination (bottom).

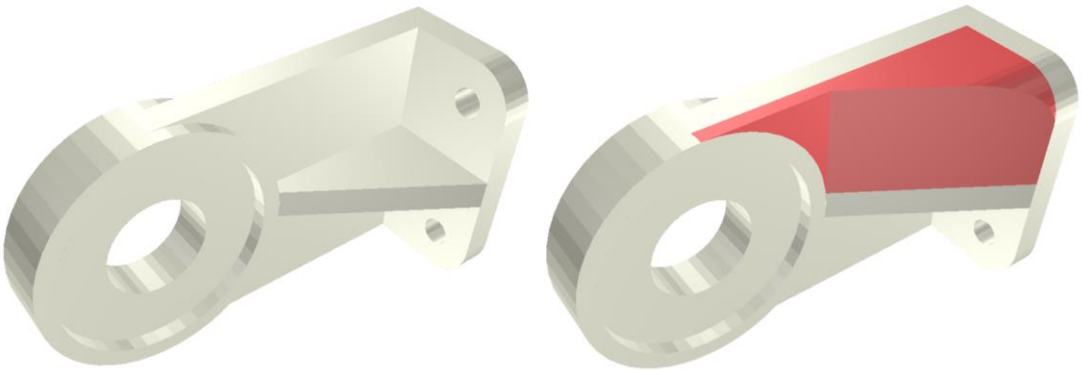

Figure 27. The bracket test part, with the print surface (tan) and the print area (red).

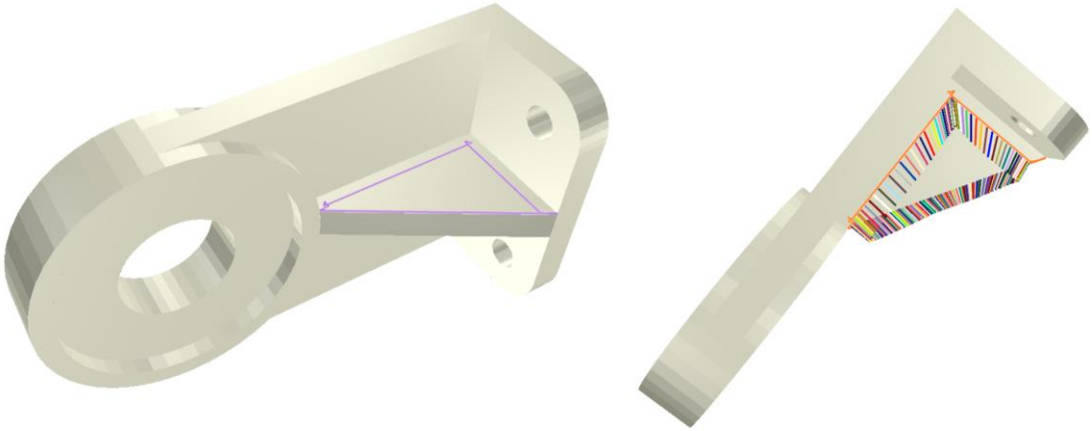

Figure 28. The first layer results from Area Refinement (left) and the results of Angle Determination (right).

 doi:10.20944/preprints202605.1142.v1

5. Future Work

This overall process has a variety of avenues for future work, some focused on increasing the capabilities and generality of the process, and others focused on adapting this for use in different applications.

The first potential area for expanding the capabilities of this process is to allow for a more iterative version to use on robotic arms, where the clearance of consecutive portions of the arm are calculated. Each new link would begin at the endpoint of the previous link, and a similar method to the angle refinement could be used. This would require some changes to what is considered valid during the large offset step of the angle refinement, as the first link could be interior to the print surface and still be valid, if the additional links could be used to reach outside of the box, as seen in Figure 29. This change to the method would allow for printing into deeper and more complex concavities (including potentially better and more feasible results for pockets with overhangs) and would allow for printing into complex concavities using robotic arm printers.

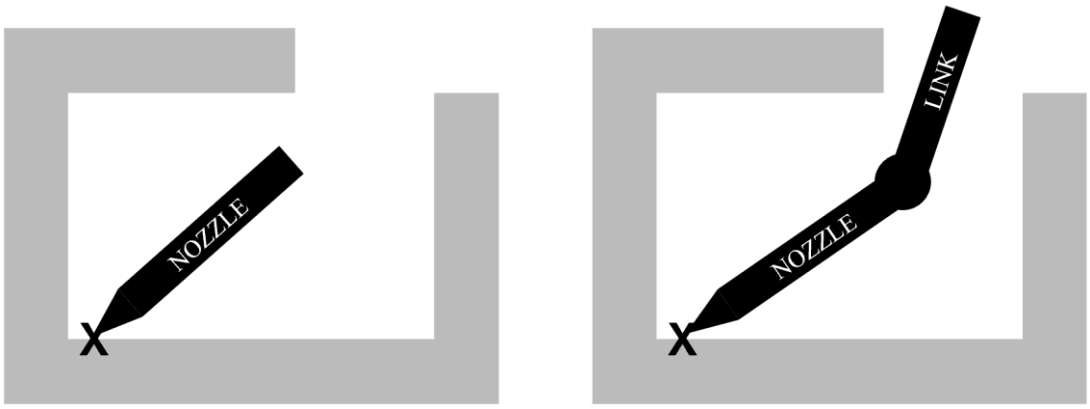

Figure 29. Demonstration of an infeasible setup with the current process (left) and how representing additional links could make it feasible (right).

The next area of expansion is to utilize more slicing orientations for the print surface. Because only three exclusion plane orientations are currently used, the final print area is affected by the orientation of the chosen print planes. While the results still produce feasible print areas, choosing different slicing orientations result in discrepancies in the final results. For example, Figure 30 shows a highlighted region that the existing process returns as infeasible due to the chosen orientation of the slice planes. However, because this is so close to the surface of the model, it could feasibly be printed. One possible way to mitigate this would be to use additional plane orientations. However, this would make the reconciliation of the updated points harder, as it is possible that some slice planes show a region as unprintable, where – in reality – it could be reached This will also help to solve the problem of corner loops, as these are caused by specific orientations of slice planes where the point can only move along one of the orthogonal axes, next to a point that can move in both. The second point moves a slightly different amount because of the averaging of the vectors. In severe cases, this causes small loops in corners, as shown in Figure 31. Because these are small areas they do not currently pose a significant problem to the print quality or feasibility, but they slow down the print as they cause the nozzle to change direction more. For a completely solid print this would cause feasibility issues, but due to the nature of 3D prints being partially hollow, this will not significantly affect the model as there is room for some excess filament in these small regions. However, adding additional planes affects the calculation used to update points in the Area Refinement step. An average could still be used to produce similar results, or this equation could be iterated upon to produce different results depending on the information from that step.

 doi:10.20944/preprints202605.1142.v1

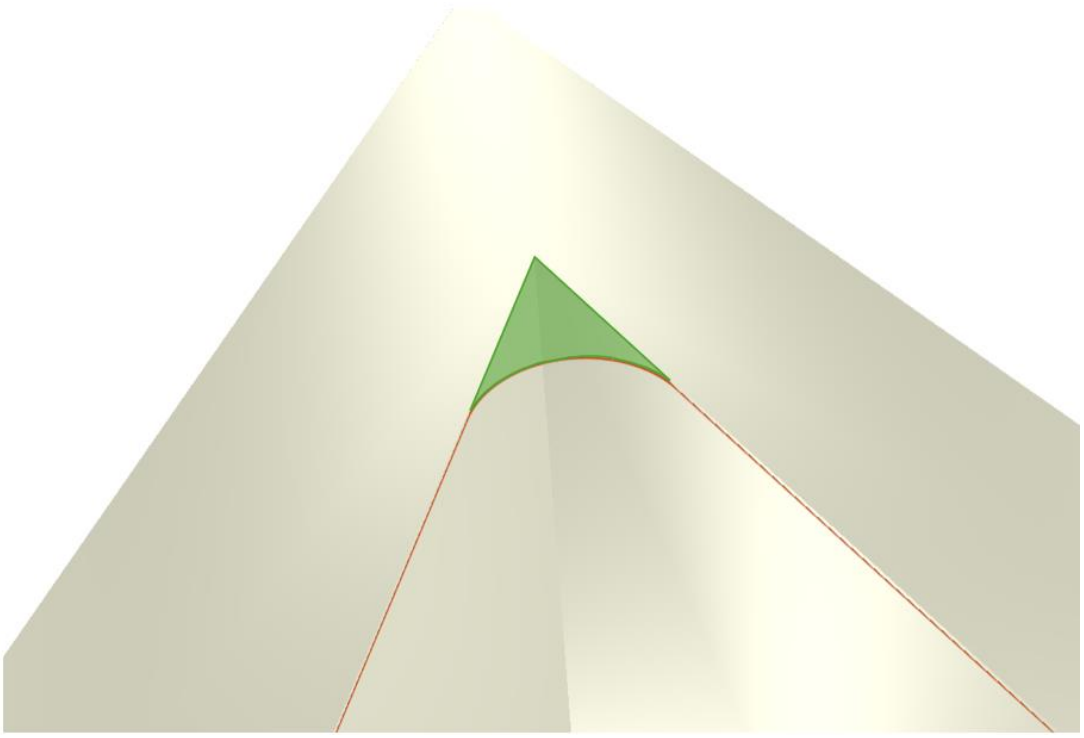

Figure 30. The area highlighted in green is not filled due to the slice orientation.

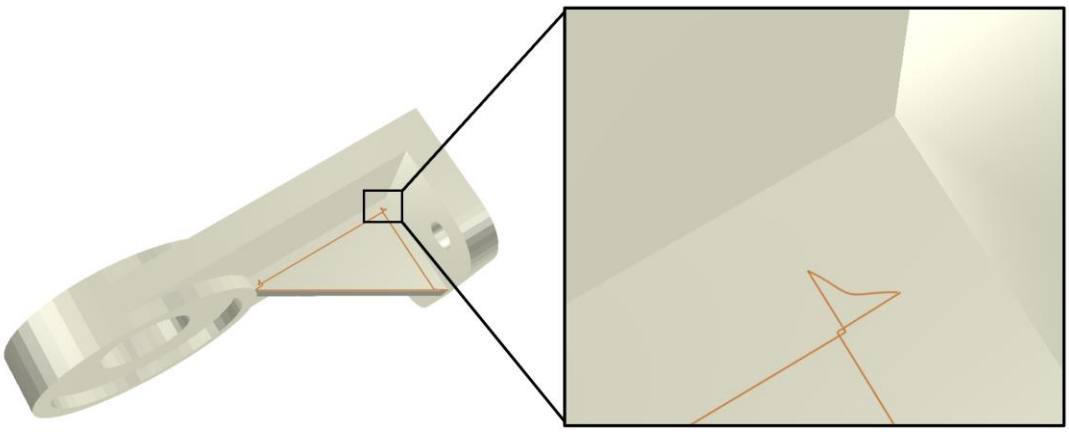

Figure 31. Looping/Overlapping shown as result of current orthogonal planes method.

Finally, the process could be improved through refinement of the large offset value. This might need to be fully redefined for more complicated parts, especially in the case of wanting to extend into large concavities with something like a robot arm. In those cases, it might be valid for the endpoint of the link to still be fully contained within the print surface, as long as the remaining links could make it out of the enclosure. Additionally, a new orthogonal slice plane could be added to the end of the vector representing the nozzle (or arm link), which could then be used to determine extra information about the clearance of that vector, as shown in Figure 32. This could be used to ensure there are no obstacles along the new slice plane within a specified distance of the vector. In other cases, this new slice plane could be used to ensure that the vector extends out of the print surface bounds, and that there are no clearance issues if using a printer with a gantry. In this second case, the new slice should not result in any polygons. If it does, then the attempted angle is invalid.

 doi:10.20944/preprints202605.1142.v1

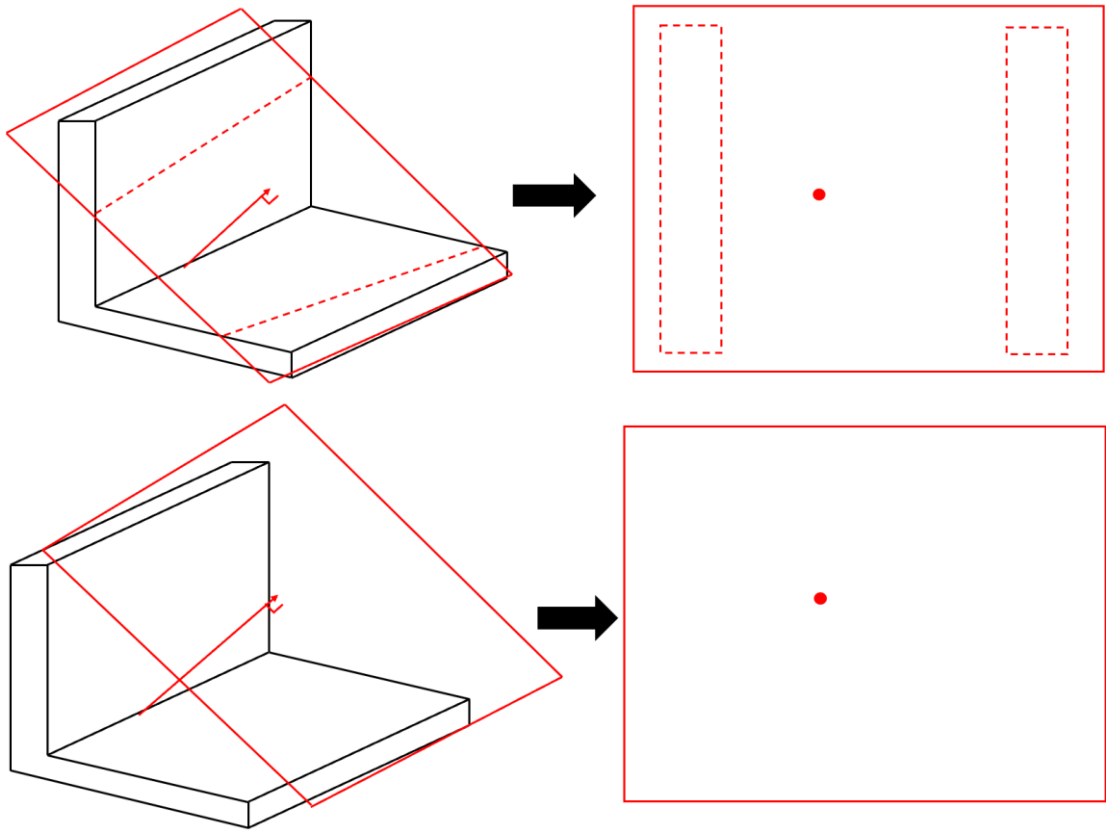

Figure 32. Orthogonal slice plane to the angle vector showing range of clearance (top) and slice plane showing no obstacles (bottom).

6. Conclusions

The methods in this paper demonstrate the ability to generate print paths on complex surfaces, which would increase the capabilities of multi-axis printers. While printers with five or more axes are not limited by hardware capabilities, the inability to easily generate the toolpaths for printing limits the usefulness of this technology. This research demonstrates proof-of-concept methods to alter a desired print area based on the print surface geometry, and to determine the toolhead angle needed to successfully print onto the surface without collision. This work could be extended to research on metals and composites printing, in order to increase the capabilities of additive manufacturing.

In combination with some of the previously mentioned future work, this research could be expanded to test with a variety of other machines. This includes both similar machines with different dimensions and machines with different kinematic chains. As mentioned previously, this could be applied to multi-axis robotic arm printers to explore the use of these methods for other deposition-based approaches such as DED. The use of five-axis capabilities for DED expands the printability of objects, and enables the use of hybrid manufacturing, where material could be deposited to alter an existing surface, and then machined down to achieve a precise surface finish [15-16]. While the potential modification and repair that these methods help enable could be used for plastics printing, the time and cost savings are likely not high enough to warrant use in most cases. However, for metals additive printing, the ability to modify existing geometries – even with complex surfaces – potentially allows for repair of existing parts such as dies and molds. Both of these uses have the potential to reduce cost and time spent on modifying or repairing existing surfaces. Testing these methods with metals printing would help to explore the feasibility of this technology for alteration, rather than just for first-time fabrication.

The final future use is the application of these methods for composite manufacturing. Multi-material printing exists in a couple different forms for deposition approaches. The first, cheaper approach is to use a single nozzle and swap the filament each time the material being printed needs

 doi:10.20944/preprints202605.1142.v1

to change. This has many limitations, as it produces a lot of waste, and it requires the various printed materials to be similar in properties. The second approach is to use multiple different nozzles that trade off, which reduces the amount of waste produced and expands the material capabilities, as each nozzle can be different in order to accommodate for different materials. The use of five-axis printing expands the possibilities of multi-material printing to allow for printing continuous, non-planar paths after an initial surface has been printed. While the Area Refinement method might not be as extendable to this work, the Angle Adjustment could be used to plan for the toolhead orientation in order to allow for separate materials to be printed onto an existing object. This could potentially enable use of continuous fiber to be printed onto complex surfaces as a final layer, while a more traditional material and printing process is used for the infill.

While the research in this paper helps to establish a general way to expand additive manufacturing capabilities for printing onto complex parts, there are still many different aspects to explore. Further refinement of the methodology and code, along with expanded testing on more machines and platforms, could help to increase the potential for five-axis additive manufacturing.

References

1. [1] Rudd, Liam, Matthew I. Campbell, and Ghazi Alonayni. “Multi-Axis 3D Printer Design Challenges for In-Situ Additive Manufacturing.” Accessed May 24, 2025. https://doi.org/10.1115/MSEC2023-100996.
2. [2] Maier. “Workpiece and Machine Design in Additive Manufacturing for Multi-Axis Fused Deposition Modeling.” *Procedia CIRP*, Complex Systems Engineering and Development Proceedings of the 27th CIRP Design Conference Cranfield University, UK 10th – 12th May 2017, 60 (January 1, 2017): 229–34. https://doi.org/10.1016/j.procir.2017.01.046.
3. [3] Tang, Pengfei, Xianfeng Zhao, Hongyan Shi, Bo Hu, Jinghu Ding, Buquan Yang, and Wei Xu. “A Review of Multi-Axis Additive Manufacturing: Potential, Opportunity and Challenge.” *Additive Manufacturing* 83 (March 5, 2024): 104075. https://doi.org/10.1016/j.addma.2024.104075.
4. [4] Gunpinar, Erkan, and Serhat Cam. “4 and 5-Axis Additive Manufacturing of Parts Represented Using Free-Form 3D Curves.” *Graphical Models* 120 (March 1, 2022): 101137. https://doi.org/10.1016/j.gmod.2022.101137.
5. [5] Jayakody, Don Pubudu Vishwana Joseph, Tak Yu Lau, Hyunyoung Kim, Kai Tang, and Lauren E. J. Thomas-Seale. “Topological Awareness towards Collision-Free Multi-Axis Curved Layer Additive Manufacturing.” *Additive Manufacturing* 88 (May 25, 2024): 104247. https://doi.org/10.1016/j.addma.2024.104247.
6. [6] Murtezaoglu, Yavuz, Denys Plakhotnik, Marc Stautner, Tom Vaneker, and Fred J. A. M. van Houten. “Geometry-Based Process Planning for Multi-Axis Support-Free Additive Manufacturing.” *Procedia CIRP*, 6th CIRP Global Web Conference – Envisaging the future manufacturing, design, technologies and systems in innovation era (CIRPe 2018), 78 (January 1, 2018): 73–78. https://doi.org/10.1016/j.procir.2018.08.175.
7. [7] Xiao, Xinyi, and Sanjay Joshi. “Process Planning for Five-Axis Support Free Additive Manufacturing.” *Additive Manufacturing* 36 (December 1, 2020): 101569. https://doi.org/10.1016/j.addma.2020.101569.
8. [8] Tang. “Convexity and Surface Quality Enhanced Curved Slicing for Support-Free Multi-Axis Fabrication.” *Journal of Manufacturing and Materials Processing* 7, no. 1 (February 2023): 9. https://doi.org/10.3390/jmmp7010009.
9. [9] Robot Trajectories for Conformal Three-Dimensional Printing Using Nonplanar Layers.” *Journal of Computing and Information Science in Engineering* 19, no. 031011 (April 1, 2019). https://doi.org/10.1115/1.4043013.
10. [10] Nishat, Rahnuma, Yeganeh Bahoo, Konstantinos Georgiou, Robert Hedrick, and Jill Urbanic. “Collision-Free Multi-Axis Tool-Path for Additive Manufacturing.” *Computer-Aided Design and Applications*, March 6, 2023, 1094–1109. https://doi.org/10.14733/cadaps.2023.1094-1109.
11. [11] Lou, Alexandre, Yeganeh Bahoo, Robert Hedrick, and Austin Vuong. “Voxel-Based Collision Avoidance for 5-Axis Additive Manufacturing.” *Computer-Aided Design and Applications*, January 7, 2025, 845–66. https://doi.org/10.14733/cadaps.2025.845-866.

 doi:10.20944/preprints202605.1142.v1

12. [12] Tang, Tran Duc, and Erik L. J. Bohez. "A New Collision Avoidance Strategy and Its Integration with Collision Detection for Five-Axis NC Machining." *The International Journal of Advanced Manufacturing Technology* 81, no. 5 (November 1, 2015): 1247–58. https://doi.org/10.1007/s00170-015-7293-x.
13. [13] Plakhotnik, Denys, Lothar Glasmacher, Tom Vaneker, Yury Smetanin, Marc Stautner, Yavuz Murtezaoglu, and Fred van Houten. "CAM Planning for Multi-Axis Laser Additive Manufacturing Considering Collisions." *CIRP Annals* 68, no. 1 (January 1, 2019): 447–50. https://doi.org/10.1016/j.cirp.2019.04.007.
14. [14] Elliott, Sam, and Matthew I. Campbell. "Calibration and Inspection of Additive Manufacturing With Laser Profilometry." Accessed May 27, 2025. https://doi.org/10.1115/IMECE2024-144176.
15. [15] Gibson, Brian T., Paritosh Mhatre, Michael C. Borish, Celeste E. Atkins, John T. Potter, Joshua E. Vaughan, and Lonnie J. Love. "Controls and Process Planning Strategies for 5-Axis Laser Directed Energy Deposition of Ti-6Al-4V Using an 8-Axis Industrial Robot and Rotary Motion." *Additive Manufacturing* 58 (October 1, 2022): 103048. https://doi.org/10.1016/j.addma.2022.103048.
16. [16] Ahn, Dong-Gyu. "Directed Energy Deposition (DED) Process: State of the Art." *International Journal of Precision Engineering and Manufacturing-Green Technology* 8, no. 2 (March 1, 2021): 703–42. https://doi.org/10.1007/s40684-020-00302-7.