MFG322228

# From Concept to Reality in Under a Day with 2.5-Axis Generative Design and CAM

Ben Weiss
Autodesk

Christopher Wade
Autodesk

## Learning Objectives

- Learn how to produce and refine manufacturable generative design outcomes
- Learn how to modify outcomes to incorporate additional design intent and prepare for manufacture
- Learn how to generate efficient CNC toolpaths to machine generative geometries
- Discover the productivity benefits of an end-to-end design-to-manufacturing workflow

## Description

Generative design in Fusion is not just for complex aerospace components that require expensive and time-consuming additive manufacturing or 5-axis CNC. In this session we will demonstrate an end-to-end workflow from 2.5-axis Fusion generative design through CAM for "everyday" CNC parts which can be programmed and machined in a matter of minutes, not hours or days. Starting with generative design, we will explore how to make the best use of the new 2.5-axis machining constraints with one of the Autodesk nerds behind the technology, then discuss editing and refining outcomes. Finally, one of the Autodesk experts in CAM for generative outcomes will share strategies for programming and producing the generative parts. The emphasis will be on improving designer and manufacturing engineer productivity, reducing CAD, CAM, and machining time, and going from problem definition through production in hours, not days.

## Speaker(s)

**Ben Weiss**
Hi! I'm Ben. I joined Autodesk in 2018 after graduating with a doctorate in Mechanical Engineering from the University of Washington, Seattle, where I did nerdy research in design optimization with additive manufacturing constraints. At Autodesk, I've leveraged those experiences to help build new manufacturing constraints for Generative Design in Fusion, including 2D cutting and now 2.5-axis machining. I'm excited to get to come share what I've been working on and get to meet all of you.

**Chris Wade**
Hello! I'm Chris, a keen badminton and squash player. But outside of that I'm a Technical Consultant in the Manufacturing Industry Futures team of Autodesk Research. I've worked at the Birmingham, UK branch intermittently since 2012 – initially alongside studying for my Masters Degree in Integrated Mechanical and Electrical Engineering, but full-time since I graduated in 2015. My current role necessitates a wide-ranging skill-set, such as knowledge of programming and manufacturing best practices. These are used to concoct bespoke solutions for the problems presented by our various customers or to push the boundaries of what manufacturing can achieve.

## Contents

## Table of Step-by-Step Procedures

## Introduction

We all want to design and manufacture faster. We all want to design better products. With generative design, engineers can quickly explore the design space and generate performant, parametric designs for direct use in the design process.

The Generative Design workspace in Fusion 360 enables this process: A series of interface geometries are selected, loads and constraints are applied, and hundreds of potential designs are generated. These can be sorted, compared, and discussed until one or more are selected to proceed forward by saving as a new Fusion document.

Constraints can be applied to the generative design process in Fusion 360 to ensure the final design is manufacturable. This class focuses on the 2.5-axis machining constraint for generative design. We show by walking through an example how to generate parts with 2.5-axis constraints, edit them in the downstream design process, and use Fusion 360’s CAM workspace to generate toolpaths to machine the designs on a CNC mill.

### 2.5-Axis Machining

Without a manufacturing constraint, generative design in Fusion 360 typically produces smooth, organic shapes which are expensive to produce, even though they may be the best solutions to the provided setup. Manufacturing constraints restrict the shapes the generative design engine synthesizes to conform to the design rules for a manufacturing process, producing a less optimal result, but one which is easier to physically produce.

Fusion 360 has three machining constraints, which produce designs suitable for manufacture on CNC milling machines (see figure below for an example). 5-axis milling constraints require that every point in the body is accessible by a tool of the specified size, which can approach the part from any direction. 3-axis milling further restricts the tool to only approach from one of several specified axes. In both of these cases, the machine tool is a ball-nosed end mill, and the geometry created can still consist of compound curved surfaces which require many passes to machine.

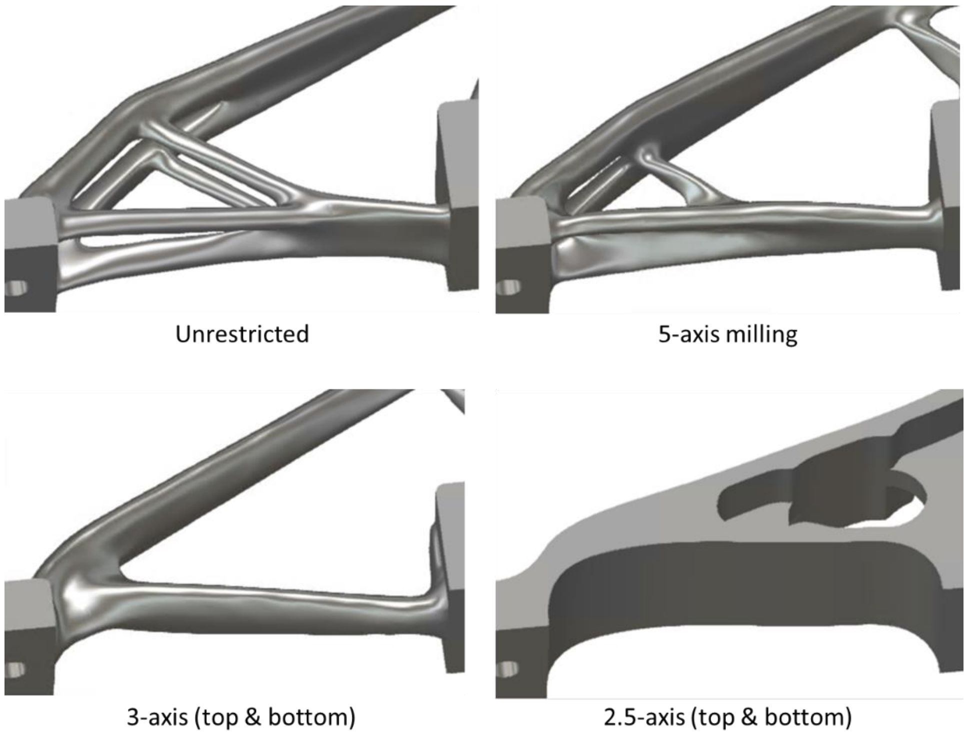


Milling constraints in generative design

The 2.5-axis machining constraint produces layered geometries which can be machined by a flat end mill in a series of passes at different set heights. Not only are these shapes much more efficient to manufacture, even on a 3-axis CNC machine, but the outcomes can be represented using familiar sketches and extrusions instead of T-spline bodies. However, because of these restrictions on the design space, the geometry produced does not perform as well mechanically as the 3- or 5- axis machinable designs. Generative design's Explore feature allows these comparisons to be made dynamically, early in the design process.

## Walkthrough

First, let's dive into the workflow. In this section, we will go step by step through the generative design thru CAM process for a bicycle suspension component. Our goal is to show you the process first, then dig into variations, techniques, and more detailed explanation of each step in the second half of this document.

This walkthrough starts near the beginning of the generative design process and runs through toolpath generation in the Manufacture workspace. We will assume that the generative problem definition is already present (preserves, obstacles, loads, and constraints are already defined). In the first section, we will adapt the generative setup to use the 2.5-axis machining constraint, generate and explore outcomes, and select one to export to a new parametric design. Next, we will discuss editing this outcome design, including basic steps to tune the design for performance and aesthetic considerations. Finally, the design is passed to the Manufacture workspace, where we show how to generate toolpaths to create the final design on a 3-axis CNC machine.

We recognize that the real design process is iterative. For clarity, we have chosen to simplify this walkthrough by describing each step only once, even though we recognize this is unrealistic. Where appropriate, the iterations we went through in developing the part used for this guide are noted.

This guide uses a swingarm design for a Starling Cycles bicycle to explore 2.5-axis machining constraints in generative design and the associated downstream workflows. Starling Cycles are an independent bicycle frame manufacturer based in Bristol, England.  They create bespoke frames for high-performance mountain bikes and so were very interested in the flexibility the new 2.5-axis Generative Design constraint brings to their design processes.  The ease of editability and manufacture of the components they can obtain from the software allows them dramatically improve the speed at which they can produce their high-quality custom parts. Starling Cycles was kind enough to share the 2.5-axis model with all of you for this walkthrough.

Now, buckle your seatbelts; here we go!

## Generative Design with 2.5-Axis Machining Constraints

The Starling Cycles swingarm began as a welded component. Through several engineering discussions, it was selected for redesign. The interfaces to the remaining assembly were extracted and adapted for generative design, and the loads and constraints for this part in the context of the assembly were estimated and applied. The result is the generative study shown in the following two figures and will serve as a starting point for this walkthrough.

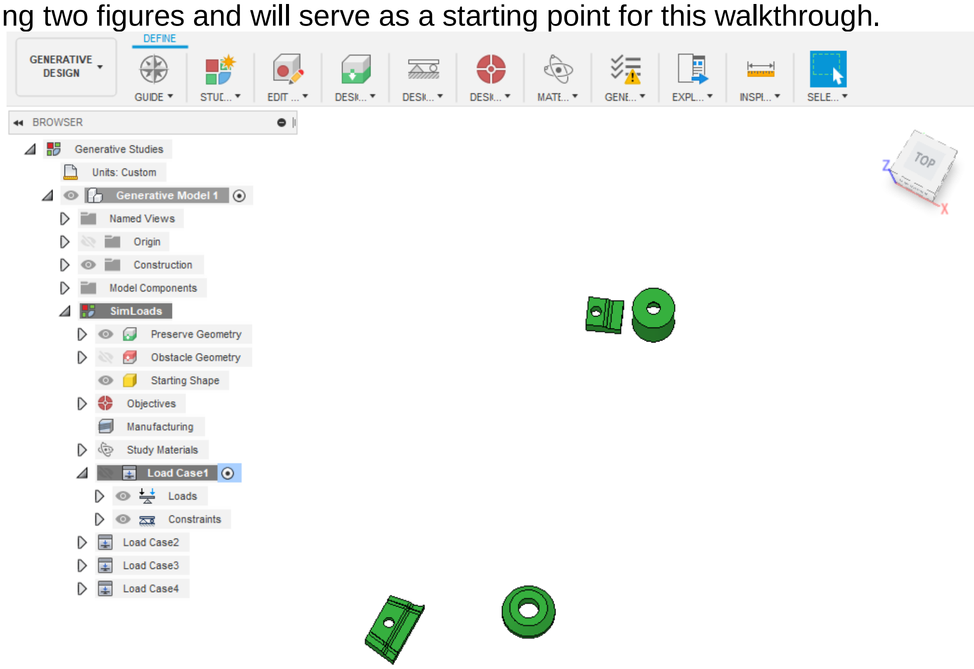


Generative design setup used as a starting point for this walkthrough, showing preserves

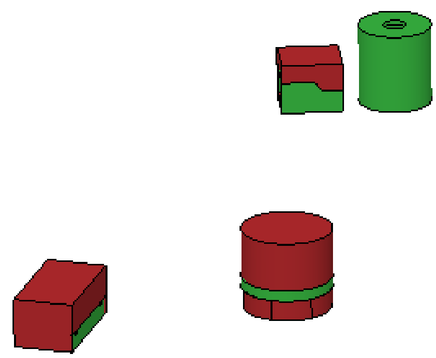

Generative design setup used, showing preserves and obstacles

In this section of the walkthrough, we will describe how this generative setup is used iteratively to refine the design, produce viable outcomes, and select one to develop into the final part. The following subsections describe:

- Determining whether a part is a good candidate for 2.5-axis machining, and how we should orient it on the mill
- Setting up to get a good 2.5-axis machining outcome
- Running and exploring the results of the initial generative run, in which we explore the broad space of possible designs
- Making changes to the generative setup to improve the suitability of results, narrowing in on the desired design
- Exploring outcomes generated
- Exporting a chosen outcome for subsequent editing

### Does this Part Make Sense for 2.5-Axis Machining?

2.5-axis machining can provide time and cost savings over other machining processes. These advantages are driven by the ability to cut most of the geometry in only one or two setups on the machine tool, the speed of cutting using the side of the tool in a smaller number of passes, and the lower cost of 3-axis milling machines.[1]

It's important to think critically about the geometry being generated in deciding whether and how to include 2.5-axis milling manufacturing options in order to take advantage of the benefits mentioned above.

First, think about how would you would orient the design space on the milling machine? Preserves which form a design space which is thinner in one direction are typically good candidates for 2.5-axis because the geometry can be easily accessed from the large top face. Cubic regions can require very long or very large tools to manufacture, increasing expense or reducing the geometric complexity of the results. We usually use the shape of the design space to select a setup direction (more on that in the next section).

[1] Note that "2.5-axis mills" don't actually exist – 2.5-axis geometry is typically cut on a 3-axis CNC milling machine.

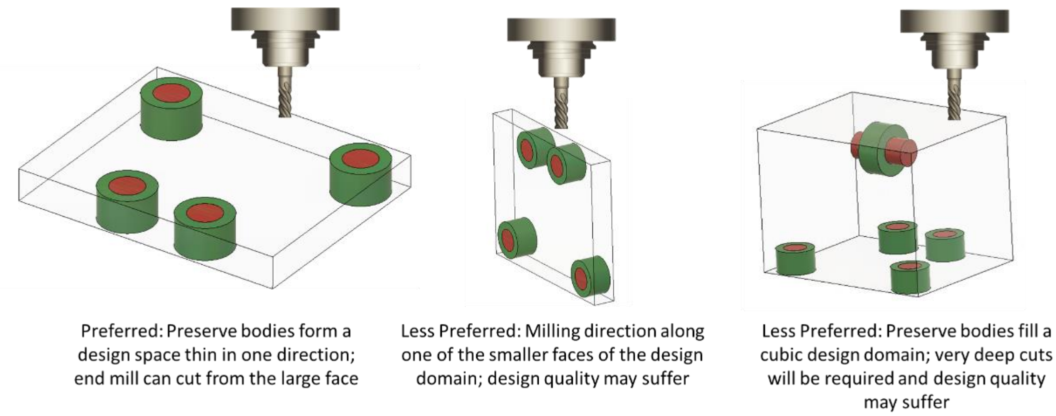


Suitability of a generative setup for 2.5-axis machining based on the shape of a tight box around the preserves

Second, are the preserve geometries themselves machinable using 2.5-axis techniques? Holes or other features in a preserve which are not aligned to the milling direction will require an additional setup to manufacture (this doesn't inhibit their use in 2.5-axis generative setups but will complicate the CAM process down the line). Thinking about how you will machine your preserve bodies before starting a generative solve can often prevent subsequent design iterations later in the process.

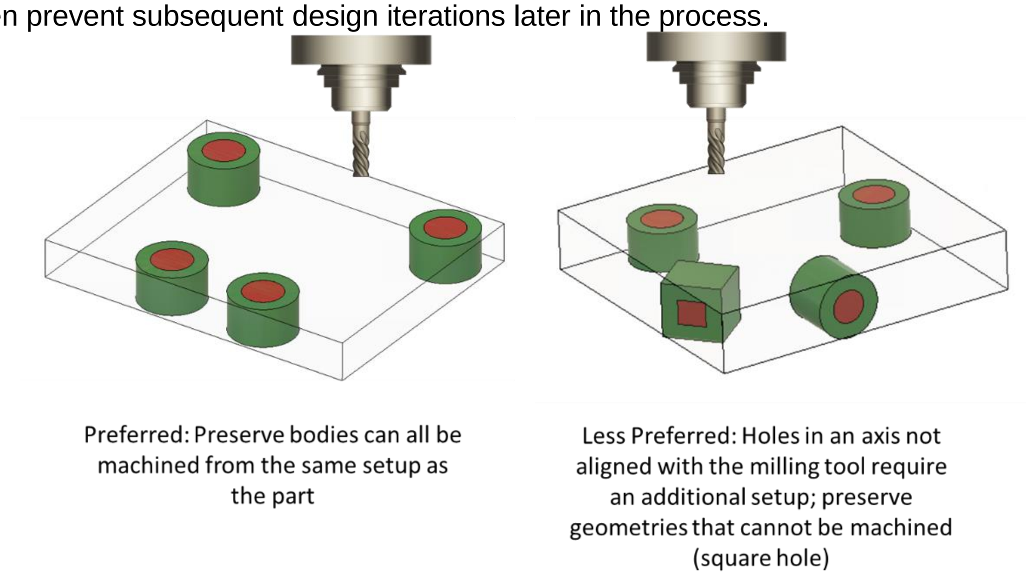


Difficult-to-machine preserve shapes may make 2.5-axis machining less appropriate

## Creating a Good Generative Study for 2.5-Axis Machining

After configuring the rest of the generative design study according to the problem to be solved (preserve & obstacle geometry, load cases, objectives, and materials), enter the Manufacturing dialog to include 2.5-axis machining as a manufacturing option.

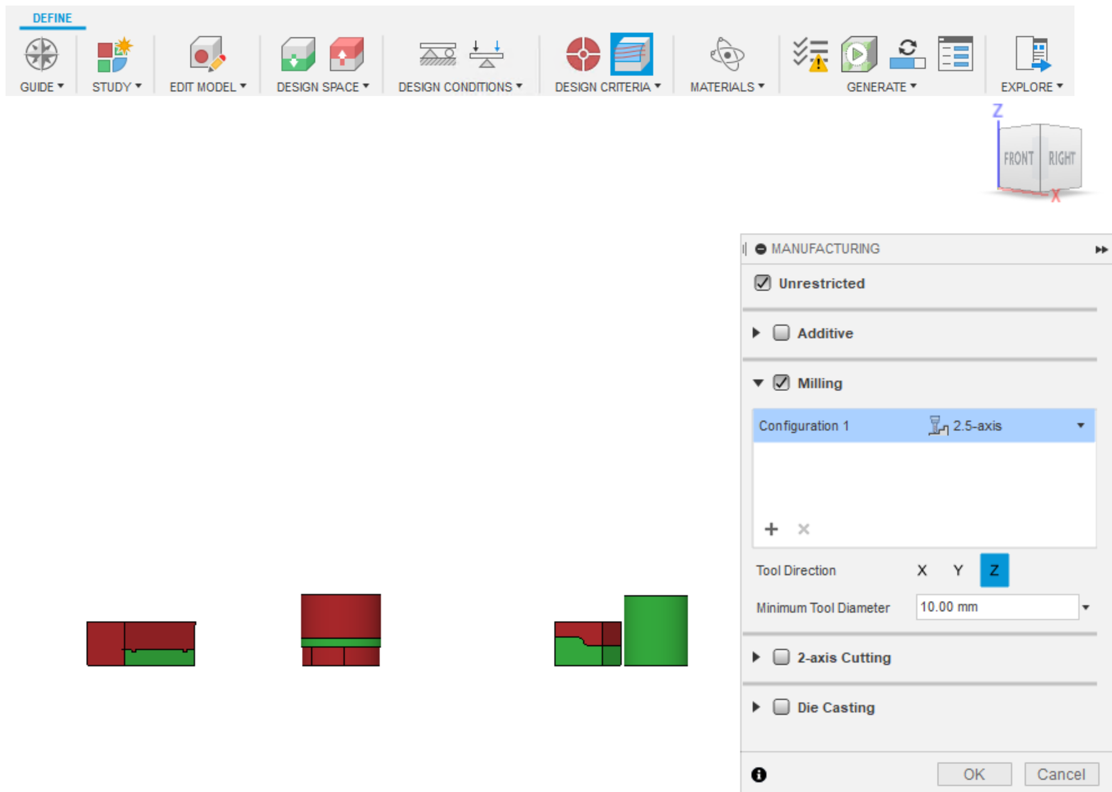


Manufacturing dialog with 2.5-axis milling enabled.

First check the Milling box to expand the relevant portion of the panel, then choose “2.5-axis” from the dropdown by Configuration 1. Two additional controls below the configuration list provide input to the 2.5-axis machining constraint.

- **Tool Direction** specifies an axis of the parallel to the machining direction; results will be generated with the tool approaching from both directions along this axis. The tool direction will typically be the thinnest cross-section of your model. If the preferred milling direction is not one of the coordinate axes, see “Using an Arbitrary Tool Direction” in the Tips, Tricks & Techniques section.
- **Minimum Tool Diameter** selects the smallest endmill you wish to use when cutting the part. Note that the actual radii generated may be increased above this value if it would create an unrealistically long and slender tool. See “Tool Accessibility” in the Tips, Tricks & Techniques section for details.

Unlike 3- and 5-axis milling, information on tool length and head diameter are not required for 2.5-axis machining. Instead, it is assumed that the tool holder remains outside the part throughout the machining process.

Click OK and notice that we now have a pre-check warning for our 2.5-axis study:

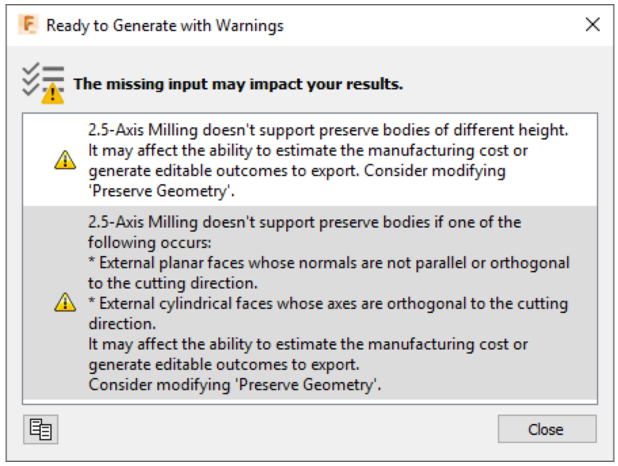


Common 2.5-axis milling precheck warning

What's wrong? The generative design workspace puts up this warning unless all the preserve geometries are straight extrusions along the milling direction which all share the same planes for the top and bottom faces. So only simple setups like this one will *not* generate this warning:

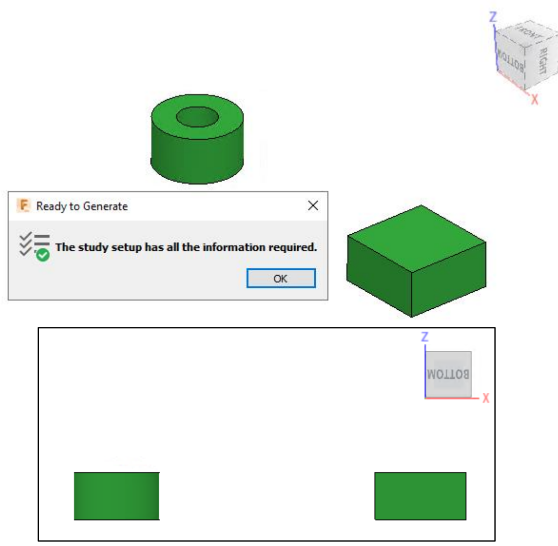


A 2.5-axis milling setup which does not generate a precheck warning

The precheck warning is intended to avoid situations where 2.5-axis generative outcomes may fail or produce unexpected results. But pre-check warnings are warnings, not errors, and it's entirely possible to use more complex geometries for obstacles and preserves without running into problems. We'll give a "Safe Rule of Thumb" which will avoid the errors hinted at by the precheck warnings[2].

[2] We're not guaranteeing that you won't hit errors using the Safe Rule of Thumb…we're still working hard to make Generative Design more robust and have plenty of work to do. Following this rule avoids situations where the system is *guaranteed to fail*.

**Safe Rule of Thumb:**
Preserves + Obstacles = Extrusions
Obstacles should not intersect Preserves

Here's what we mean: So long as preserves have corresponding obstacles so that the two together create an extruded shape, it should work fine (though the precheck warning is still present). In the figure below, each of these preserves violates some aspect of the precheck, but together with their obstacles they make a complete extrusion.[3]

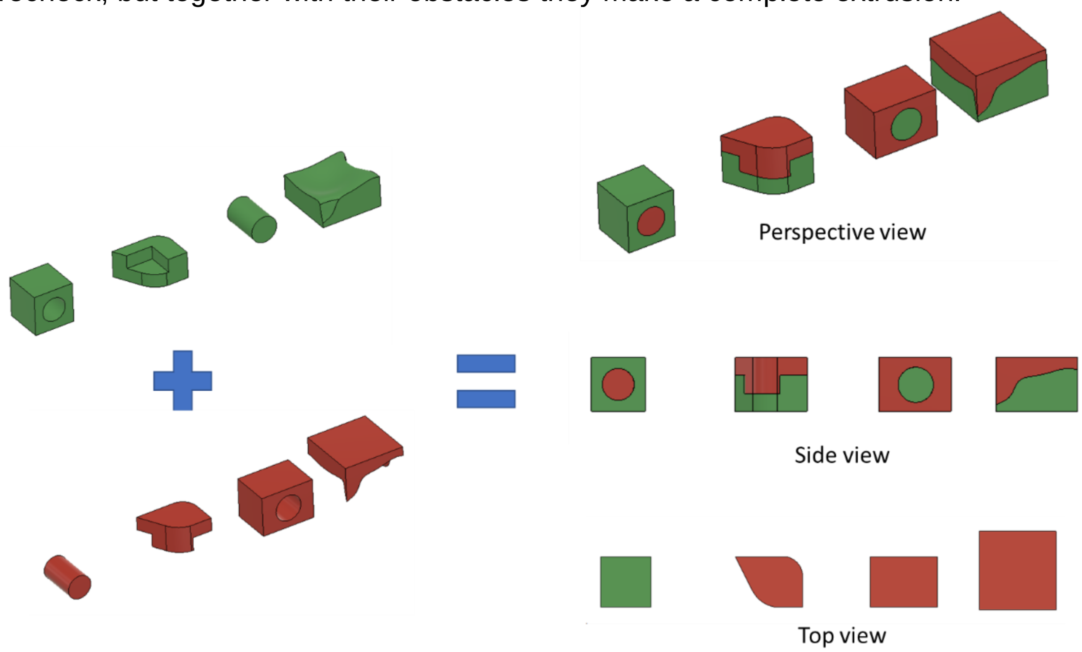


Preserves and obstacles combine to create extrusion-like structures which do not create problems for 2.5-axis generative solves, despite the pre-check warnings

For more details on allowable preserve and obstacle shapes, see the "Preserves" and "Obstacles" discussions in the Tips, Tricks & Techniques section.

Back to our example problem. The initial generative setup already has preserves around three of the interfaces turning them into extrusions, but they do not extend all the way through to the top plane. To fix this, we will extend the top surfaces of each extrusion up to the highest face in the model (along the Z direction, which we selected as the tool direction).

We can make this change in either the Edit Model workspace within Generative Design, or in the original Design workspace. Since we do not need to keep the original obstacles, we will make the change in the Design workspace.

[3] Following this Safe Rule of Thumb doesn't guarantee that generative design doesn't still occasionally have problems for other reasons…

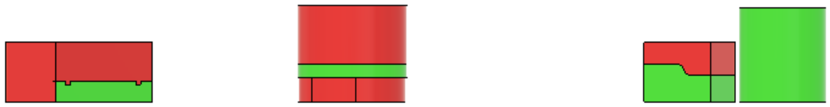

Preserve and obstacle bodies before extending the obstacles

Let’s use two Extrude commands to bring the shorter obstacle bodies up to the height of the taller obstacle.

### Procedure 1: Extend an obstacle up to the top of the part

1. In the Design workspace, click the Extrude button, then select the top face of the left-most preserve as the Profile.

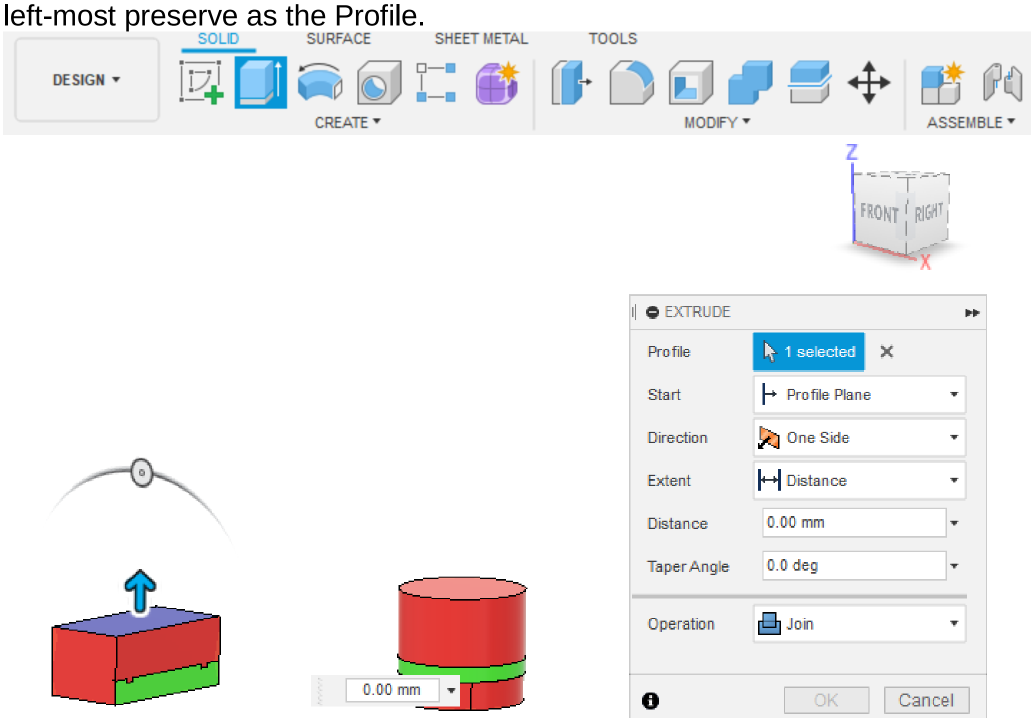

2. Now change the Extent to “To Object” and select the top face of the tallest obstacle as a target Object. This extends the obstacle up to this height.

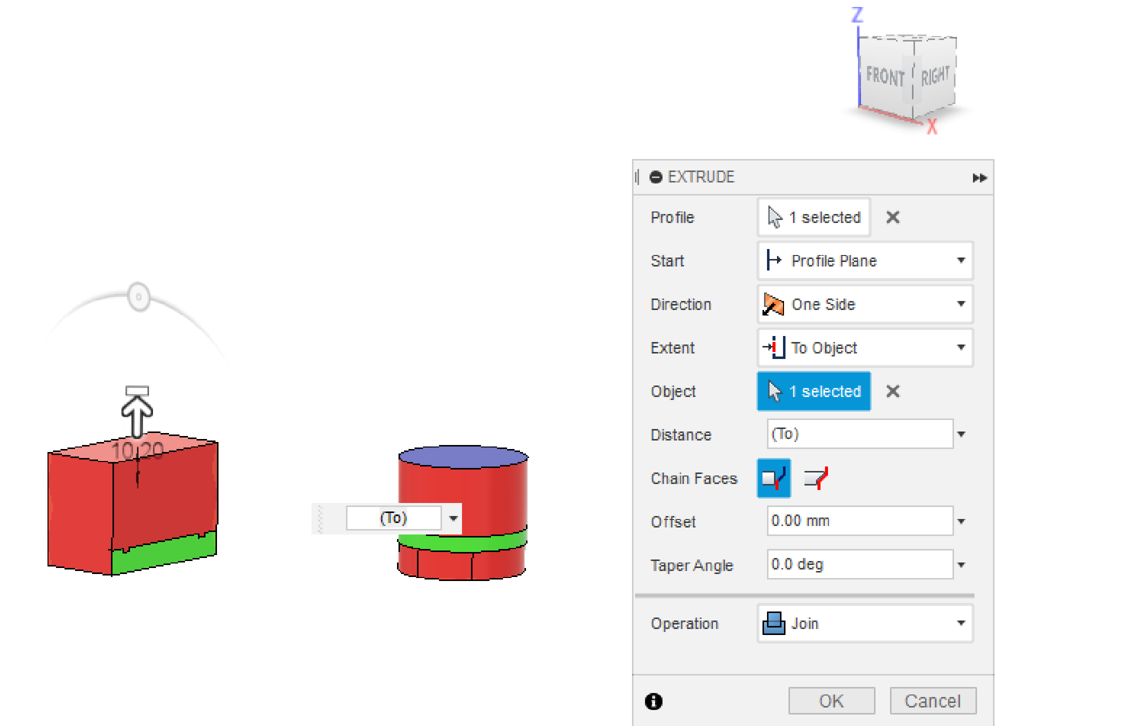


We repeat these steps for the second obstacle to bring all three up to the top plane of the design space.


Our part after modifying the obstacle heights

One preserve from the original generative setup does not have an obstacle. We will create one by making an extrusion from the preserve’s “shadow”, then subtracting out the preserve geometry, as described in the following steps:

Procedure 2. Create an extruded obstacle around an arbitrary preserve

1. Create a sketch, selecting the bottom face of one of the preserves or obstacles.

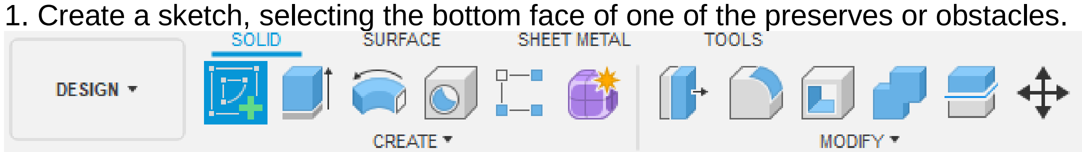

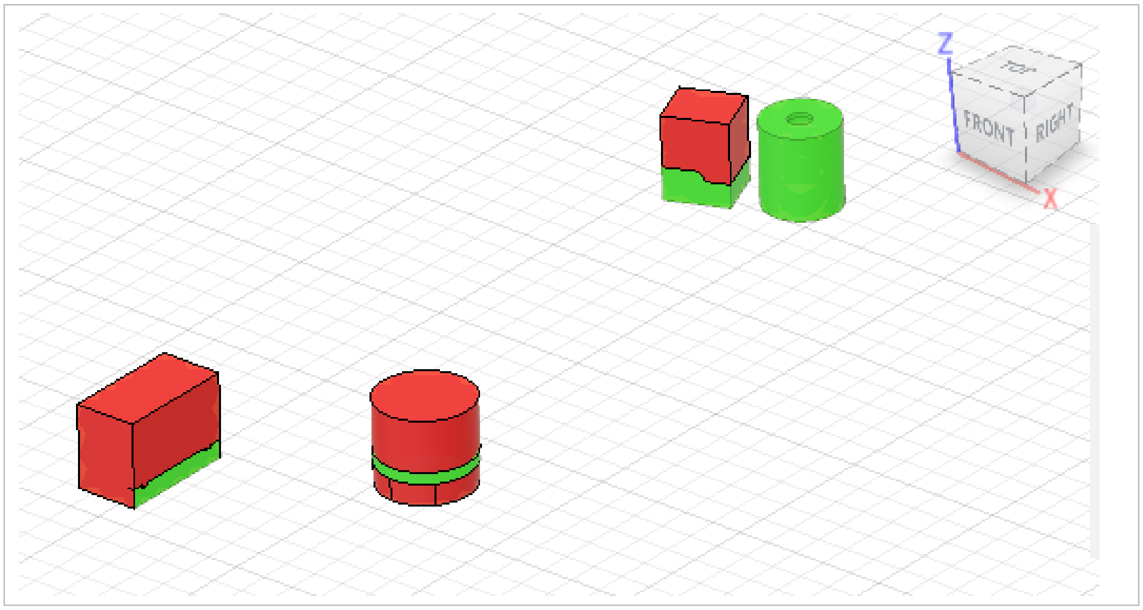
Z
TOP
FRONT
RIGHT
X

2. Select the Project command. In the dialog, change Selection Filter to “Body” and select the preserve body. This creates one or more sketch profiles from the faces of the preserve body.

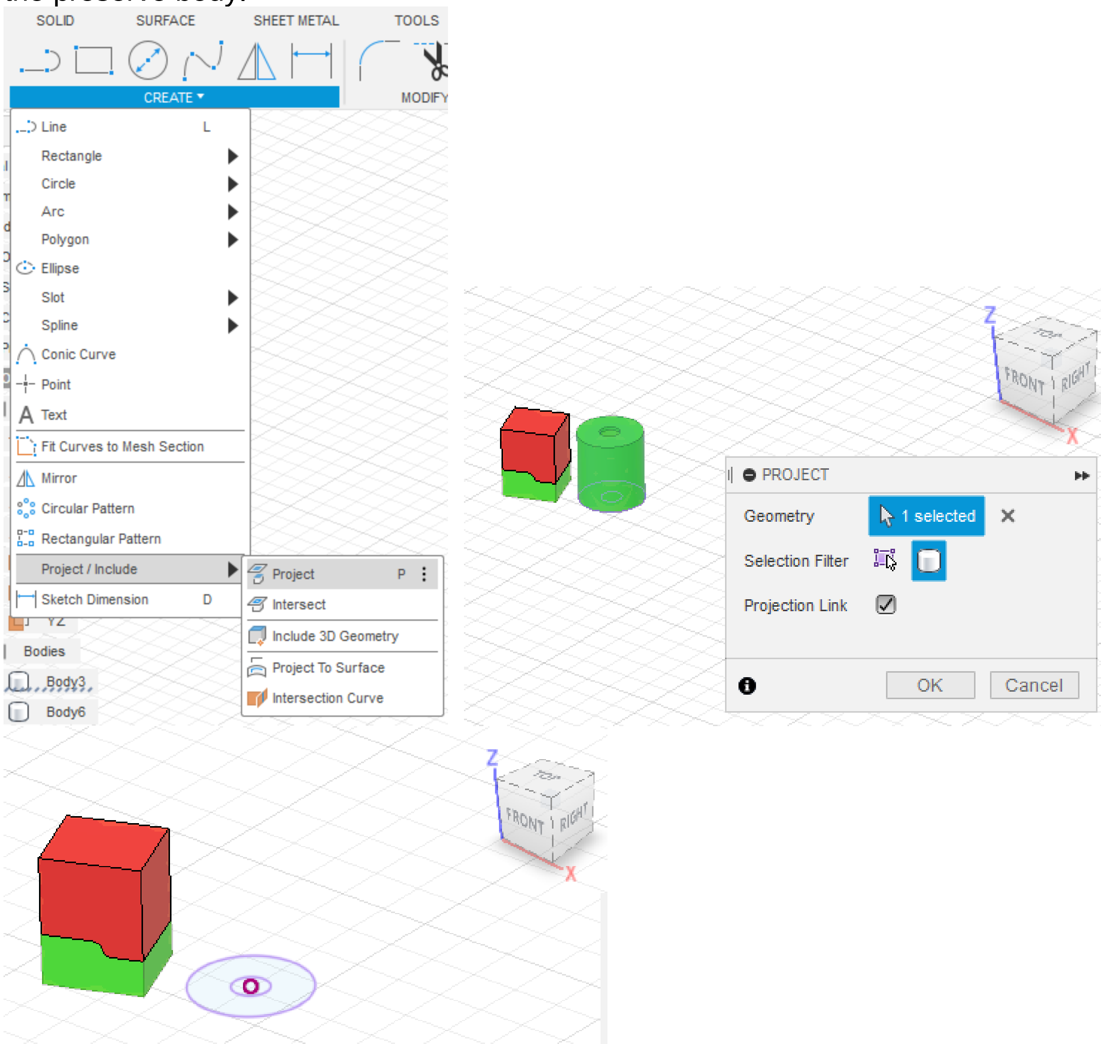

3. Close the sketch and enter the Extrude command. Select all of the profiles from the newly created sketch.

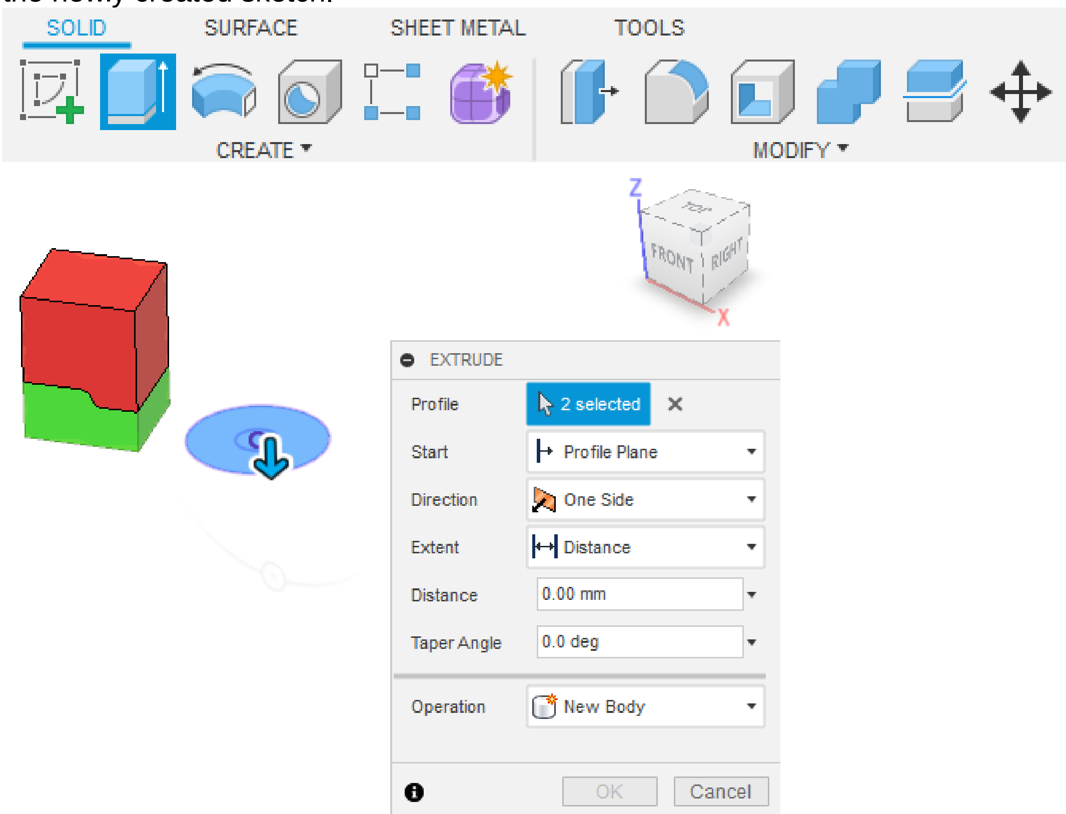

4. Change Extent to “To Object” and select the top face of one of the other obstacles as the Object. Ensure “Operation” is set to “New Body”. Click OK to finish the command.

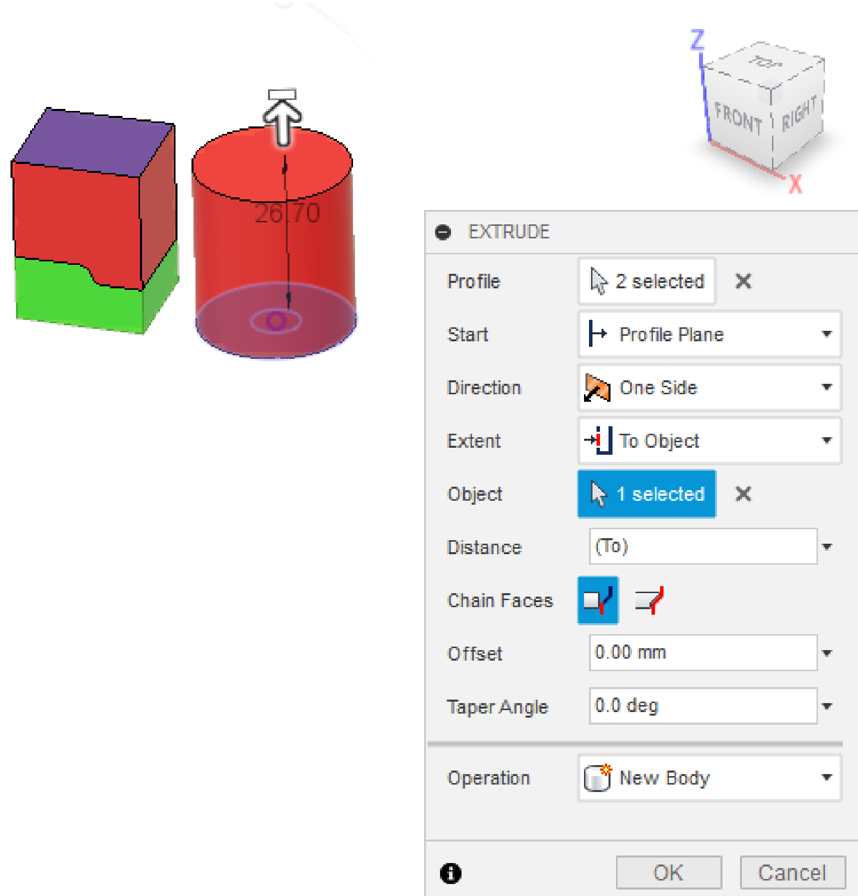

5. Subtract the preserve body from the obstacle body using the Combine command. Select the new body created in the previous step as the Target Body and the preserve as the Tool Body. Choose “Cut” as the Operation, and check “Keep Tools”.

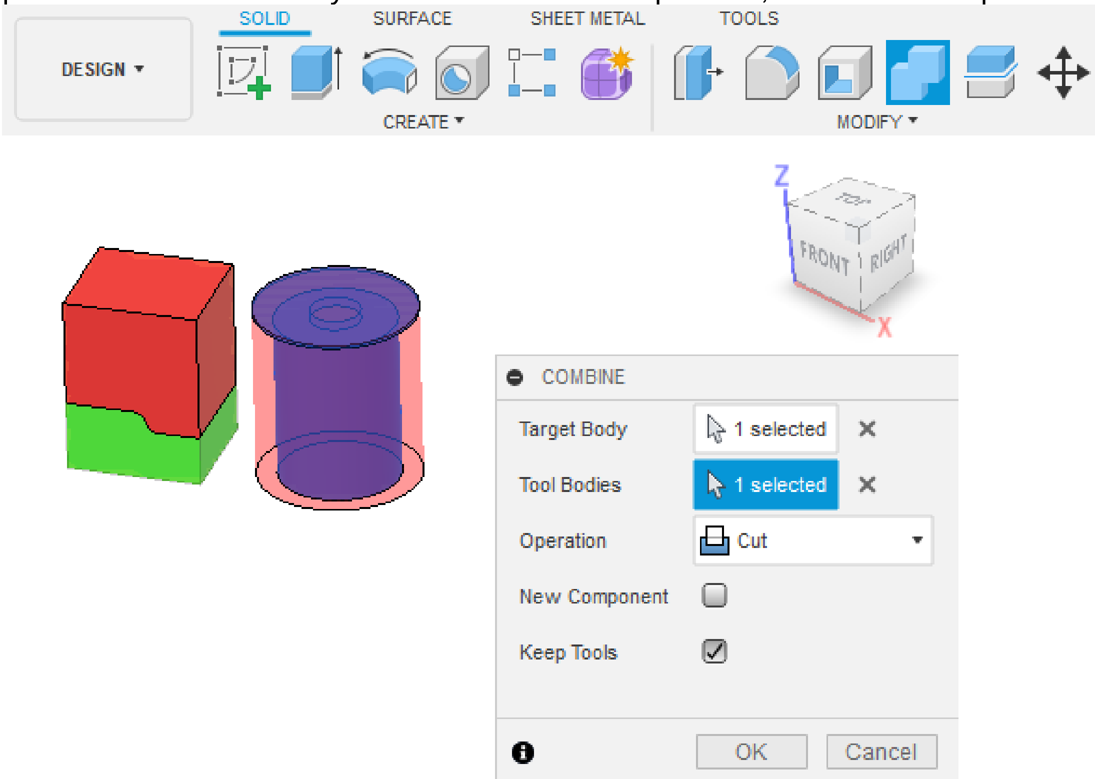


Now the geometry needed for a successful 2.5-axis milling generative solve is in place. Go back to the Generative Design workspace and add the new obstacle to the Obstacle Geometry of the current setup.

## Initial Generative Run

The initial generative solve is intended to explore the design space as broadly as we can. We supply the original engineering loads and the preserve and constraint geometry, updated for compatibility with 2.5-axis milling as described in the previous section. We will use the milling settings and materials to generate a wide variety of solutions to this design problem.

At this point in the design process, the intent is to make this part with utilizing 2.5-axis milling geometry. We don’t know what tooling we want, so we select two 2.5-axis milling configurations, one with a 10 mm tool and another with a 2 mm tool. We also add a 3-axis configuration (Z- tool direction) with a 5 mm tool diameter and check the box to get an unrestricted result. Even though these are not being considered, it is helpful to compare the 2.5-axis results against the unrestricted and 3-axis results to understand the performance implication of the less expensive manufacturing method.

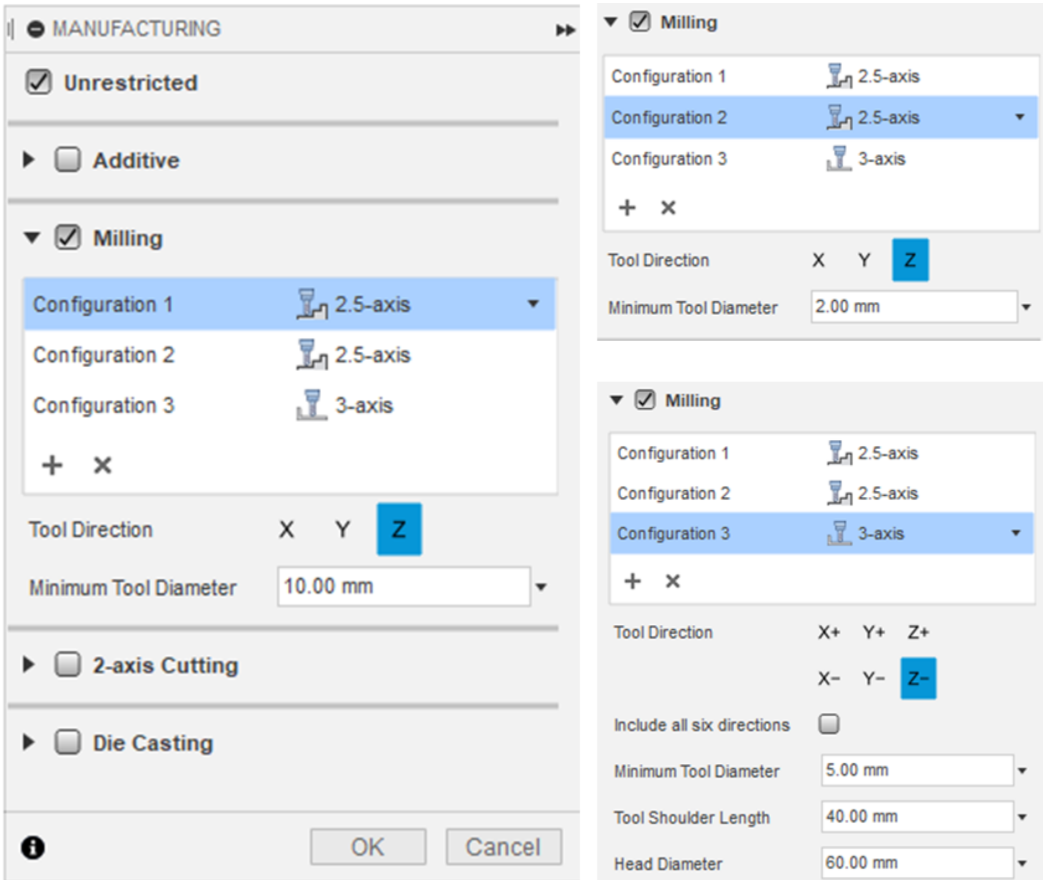

Manufacturing settings used for the first generative run

The intent is to make the final part out of aluminum. Since the purpose of this first generative run is to explore the design space, a selection of aluminum alloys with a nice distribution of yield strength is chosen (yield strength is the most important mechanical property for affecting the results of a generative study).

Table 1. Materials used for first generative study

| Material | Yield Strength (MPa) |
|---|---|
| Aluminum AlSi10Mg | 240 |
| Aluminum 6061 | 275 |
| Aluminum 2014-T4 | 290 |
| Aluminum 6061 T6 0 Cold Formed | 369 |
| Aluminum 2014-T6 | 400 |

These options will be combined by the generative design engine to produce many outcomes, each of which solves the design problem using a unique combination of manufacturing method and material. For the 2.5-axis, each material produces three outcomes with different setup directions and layer counts, as shown in the following figure.

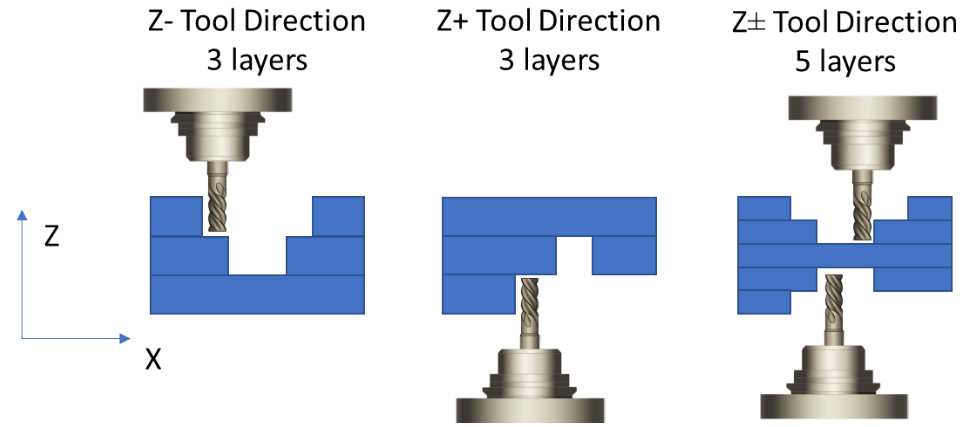


Three outcomes generated for each material when 2.5-axis milling is enabled (Tool Direction Z shown)

The top and bottom of the generated geometry corresponds to the highest and lowest points on the preserves along the tool direction (unless a starting shape is larger; see "Starting Shapes" in the Tips, Tricks & Techniques section). These top and bottom planes are evenly divided into 3 (or 5) layers. The number of layers and layer locations are fixed and cannot be changed.

All told the first generative run produces 8 outcomes for each material (one unrestricted, one 3-axis, and three for each 2.5-axis manufacturing method), for a total of 8x5=40 outcomes. We let this study run until at least a good handful of outcomes have finished generating, but there is no need to wait for the entire process to complete.

Enter the Explore environment to view the 40 outcomes. Looking through the results, it is often helpful to use the Properties view to see which material and manufacturing method were chosen for each outcome.

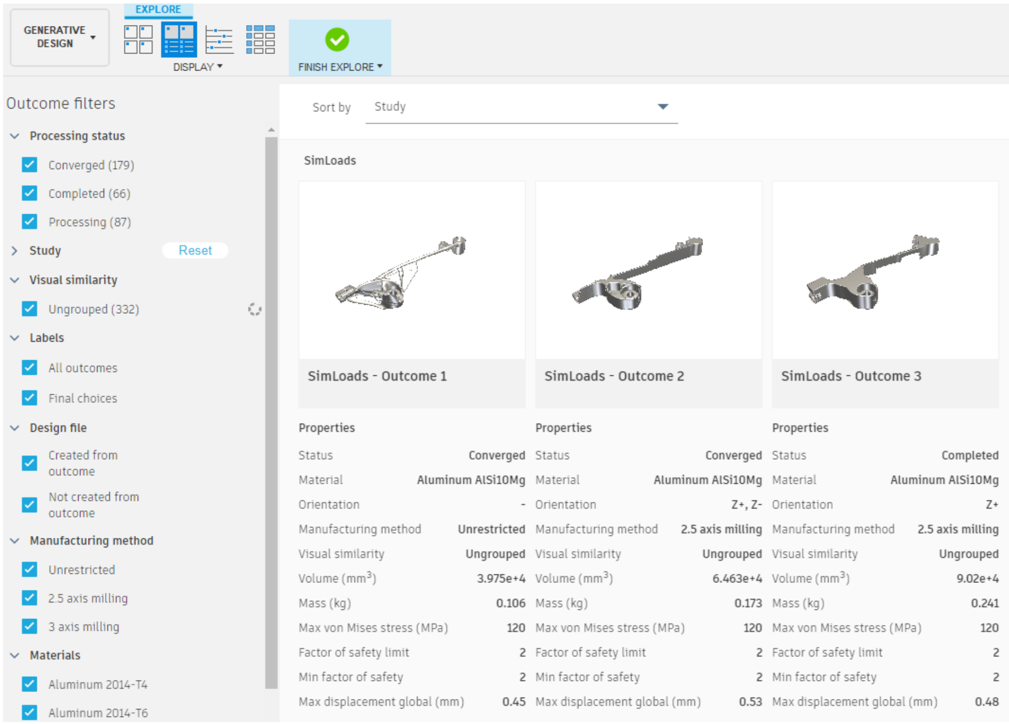

Properties View in the Explore environment

Note that it is not possible to see which 2.5-axis solve a given outcome belongs to (minimum tool diameter is not shown in the list of properties). As a workaround, sort by Study. The first three 2.5-axis outcomes correspond to the first entry in the list of configurations, as shown in the following figure:

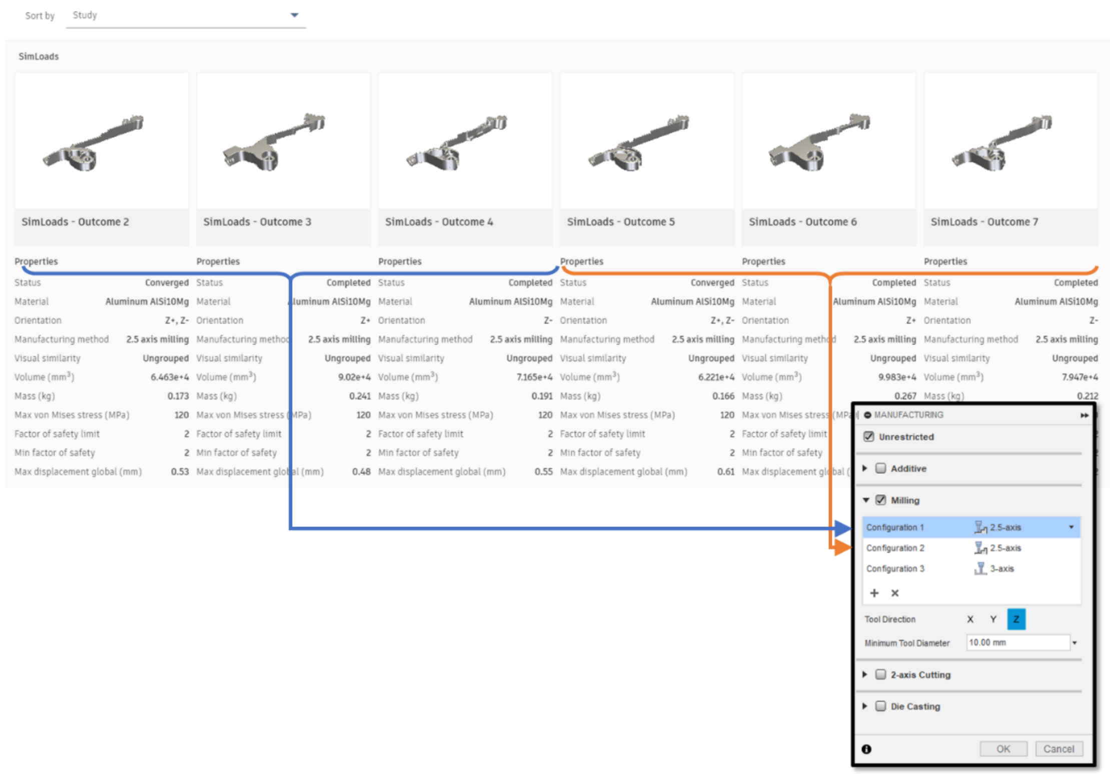


Each group of three 2.5-axis milling results corresponds to the next entry in the list of configurations (when sorted by Study)

Looking at the results, we make two observations.

First, we compare the 2.5-axis results to the unconstrained and 3-axis milling solutions. Using max displacement and outcome volume as proxies for performance[4] and weight[5], we use the Scatter Plot View to explore the overall design space, as shown below.

---

[4] Since the solve process is constrained to limit stress to an acceptable value, max stress is relatively constant across all of the designs produced. Max displacement is not constrained, and is therefore a better rough indicator of "how well will this thing work" at this stage of the process.
[5] We use volume as a proxy for weight because we're interested in the overall size of the final component – we haven't chosen a material for manufacturing yet, and don't want density variations to shade the results (though we don't expect them to be a very big factor in this study, since we only selected aluminum alloys).

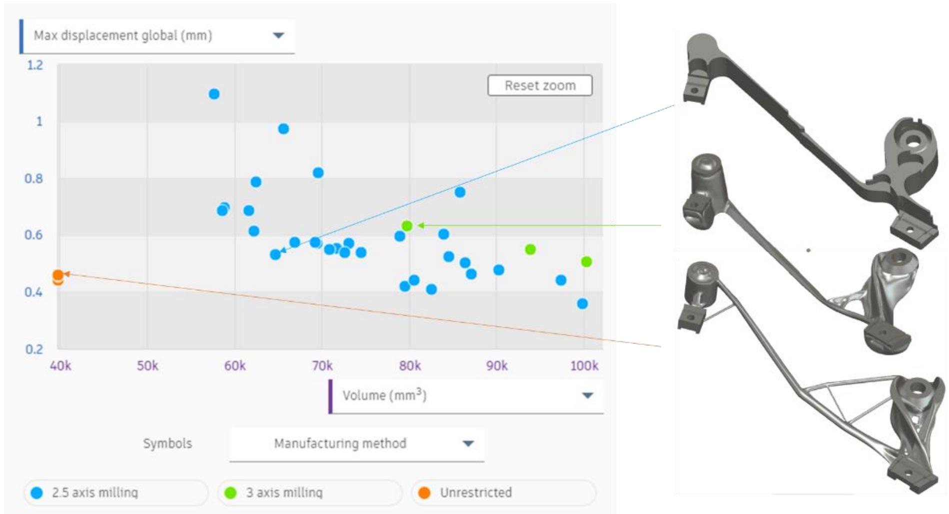

Scatter Plot view of the initial generative study with three outcomes highlighted

Examining these results, we see that 2.5-axis and 3-axis outcomes are substantially larger than the unconstrained solve. That makes sense, considering how the unconstrained solutions all have a diagonal bar from the top pivot down to the bottom bracket and both milling constraints don't allow "floating" bars. Since the 2.5-axis outcomes are performing comparably to the 3-axis outcomes, we conclude that this part is suitable for 2.5-axis milling.

Our second observation comes from looking through several of the 2.5-axis solutions. For example, look at Outcome 7. It is 2.5-axis, with a -Z tool direction and uses a 2 mm tool diameter. Several iterations of the design evolution are shown in the figure below.

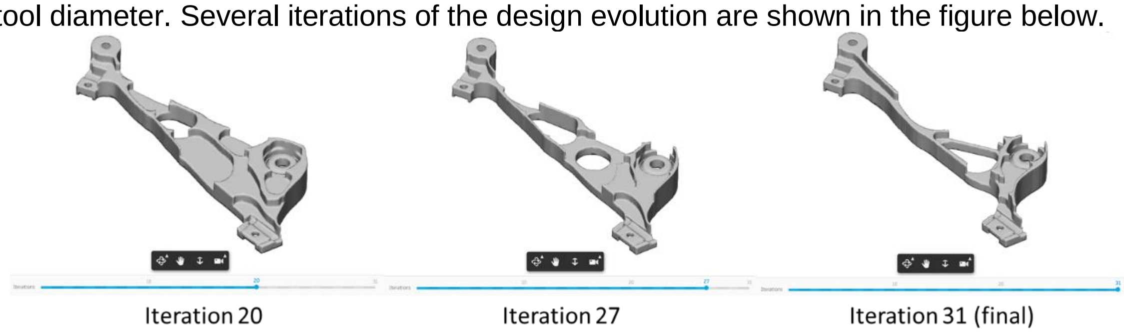

Outcome 7 shown at several iterations of the shape evolution

In reviewing the design, we decided that we liked the triangular of the intermediate design but did not like the tall thin walls which appeared in this and other design outcomes. We also decided that we wanted to make the part from a single setup.

## Setup Revisions

Based on the observations from the first round of generative results, decided to make two changes to our problem setup.

First, we selected four additional load cases which capture forces we hadn't considered originally. These will also encourage the "triangle" shape we liked in some of the original design's intermediate results.

Second, we decided to include a starting shape to discourage high, thin walls by concentrating material towards the "bottom" of the design space. In the next few pages, we will walk through creating a starting shape and making a new generative study in which we can add the load cases without invalidating the original results. For more information on starting shapes, see "Starting Shapes" in the Tips, Tricks & Techniques section.

To create the starting shape, we could go back to the Design workspace, create a starting shape, and add it to a new generative setup, but this would cause all the rest of our generative results to be shown as "out of date" because the part geometry changed. Instead we will create a new Generative Model, as detailed in the following procedure, and add the starting shape in the Edit Model workspace.

A Generative Model is similar to the Simulation Model in the Simulation workspace. It is a parametric "shadow" of the main design in which we can make additional simplifications or changes without modifying the parent part outright. Additionally, we can have multiple Generative Models to allow us to explore various scenarios in generative design without creating multiple separate files. If we change the main design, all existing generative results will be flagged as "out of date". Changes made to a new Generative Model will not cause this, however.

### Procedure 3. Create a new Generative Model with a starting shape

1. Enter the Edit Model workspace by clicking the button in the Generative Design toolbar.


The browser now shifts to show just the geometry from the current Generative Model.

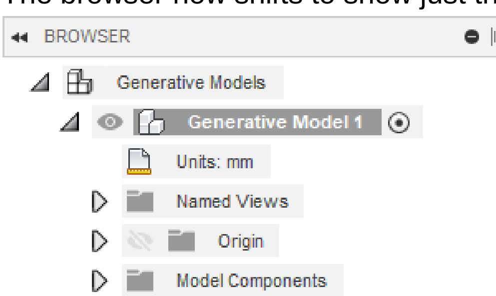

2. Right-click on Generative Model 1 and select “Clone Generative Model”

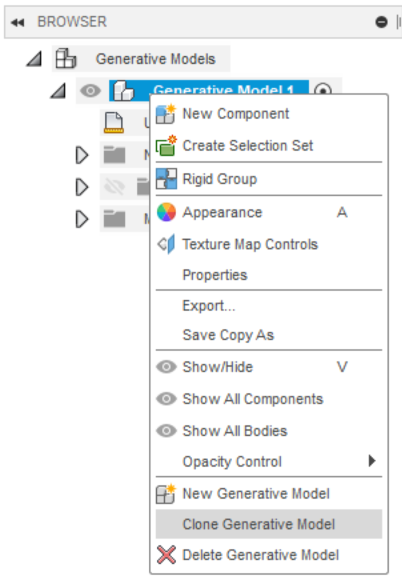


3. A dialog appears asking which studies to include when cloning the generative model. We only have one study; leaving it checked will allow all the loads and constraints to be carried forward.

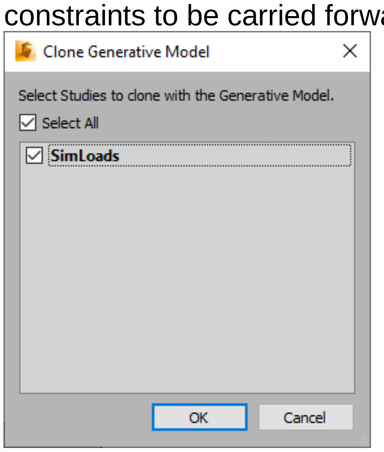

4. Create a sketch on the bottom surface of one of the preserves

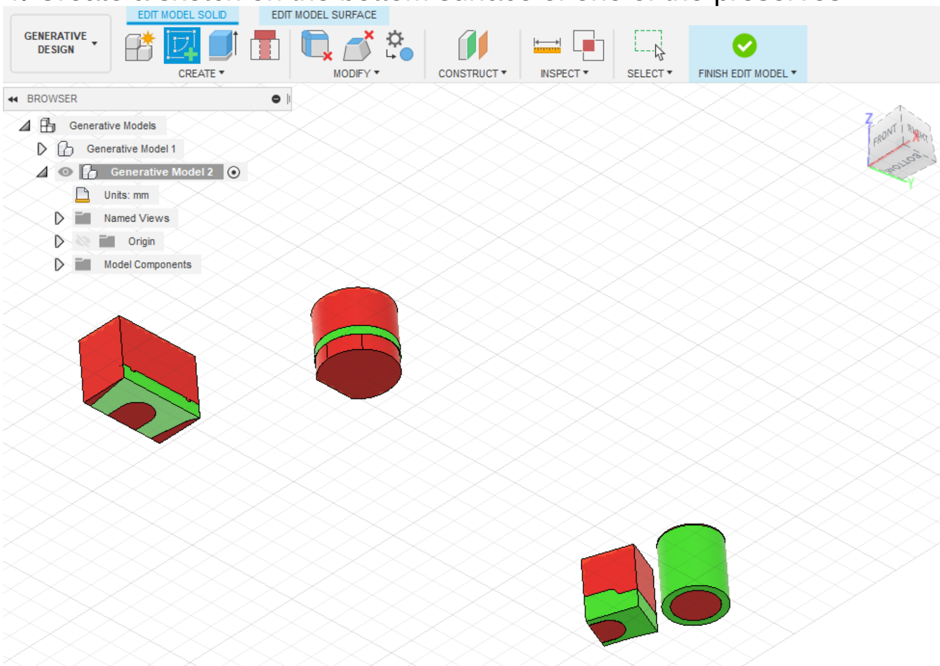


5. Draw a rough outline around the preserves. This does not need to be a high quality sketch; it will just serve as a starting point for the generative design process. We also drew two lines across the sketch to cut the design into profiles we can extrude at different heights.

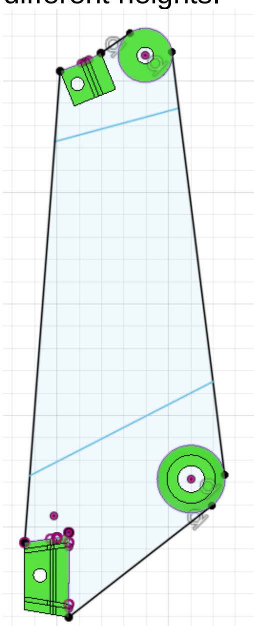

Note that I've flipped the view over to show the top of the preserve bodies.

6. Extrude the sections with three Extrude operations as shown below. Our intent is to be able to use either a 3-layer or 5-layer outcome, so we'll set the height of the central section to a bit less than one fifth of the total height (26.2 mm / 5 ≈ 5mm), and the lower section we will set the lower extrude to about three fifths of the total height.

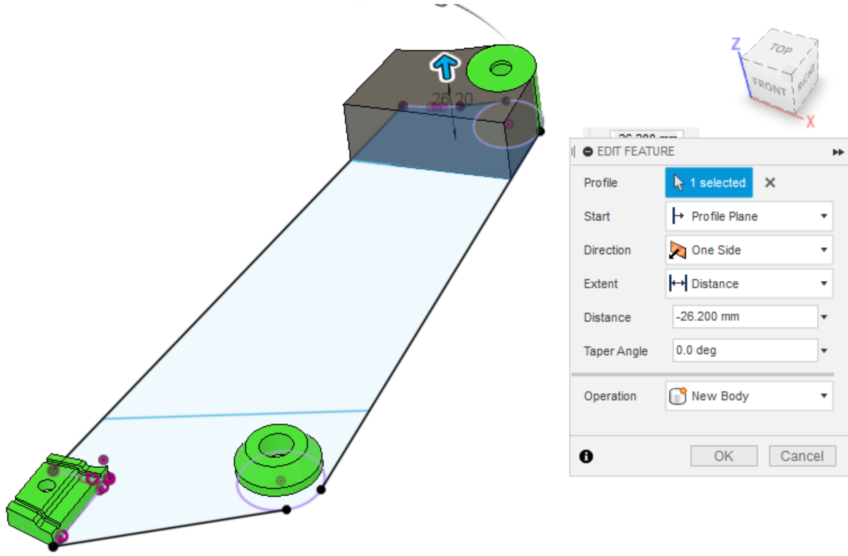


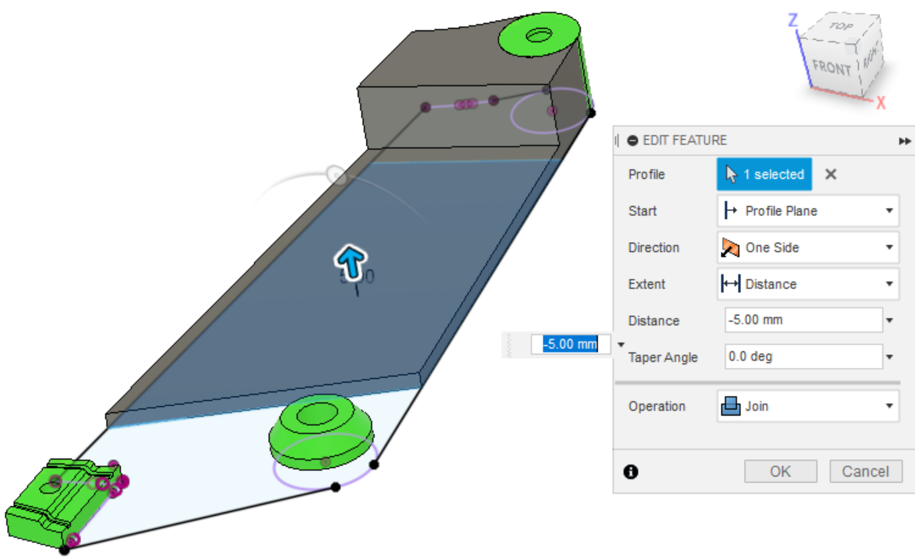

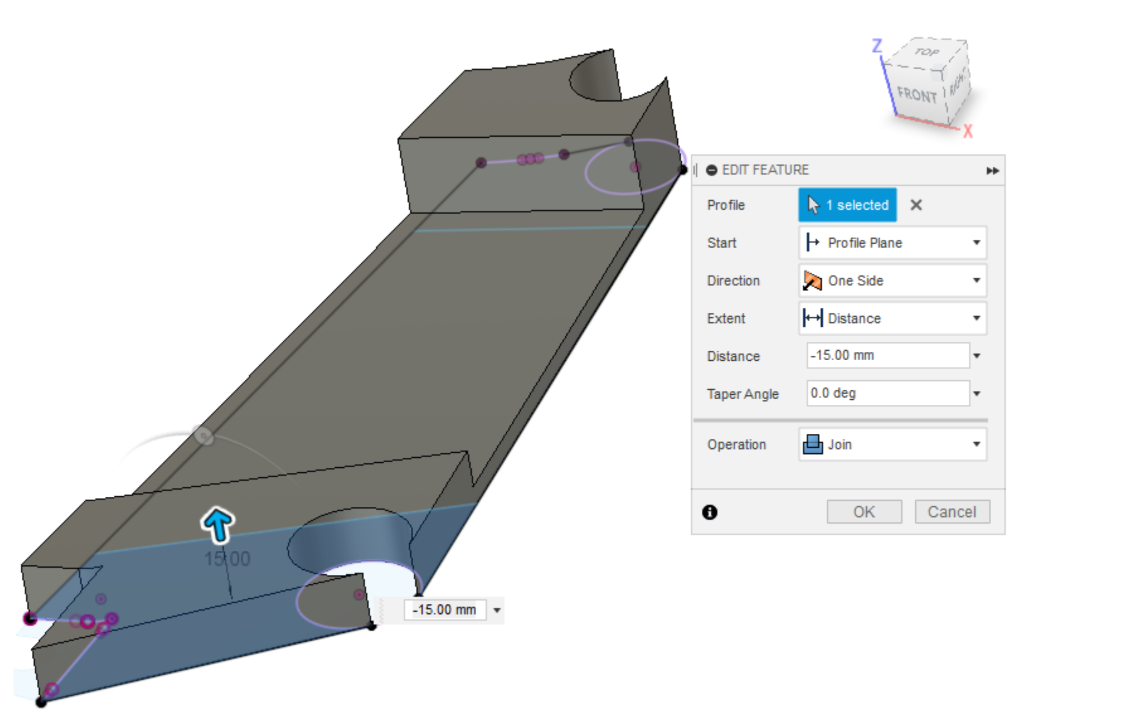


7. Click the Finish Edit Model button to leave the Edit Model workspace.


8. Back in the Generative Design workspace, the browser shows the new generative model activated, along with the copy of the SimLoads setup we included when we created it. Let's rename the SimLoads generative study so we don't confuse it with the one in Generative Model 1. I named it "GenLoads Bottom Fifth".

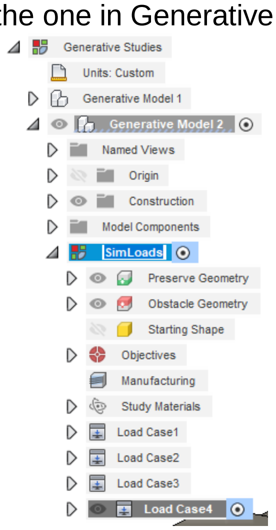

9. Add the starting shape to the generative setup using the Startup Shape item in the Design Space drop-down.

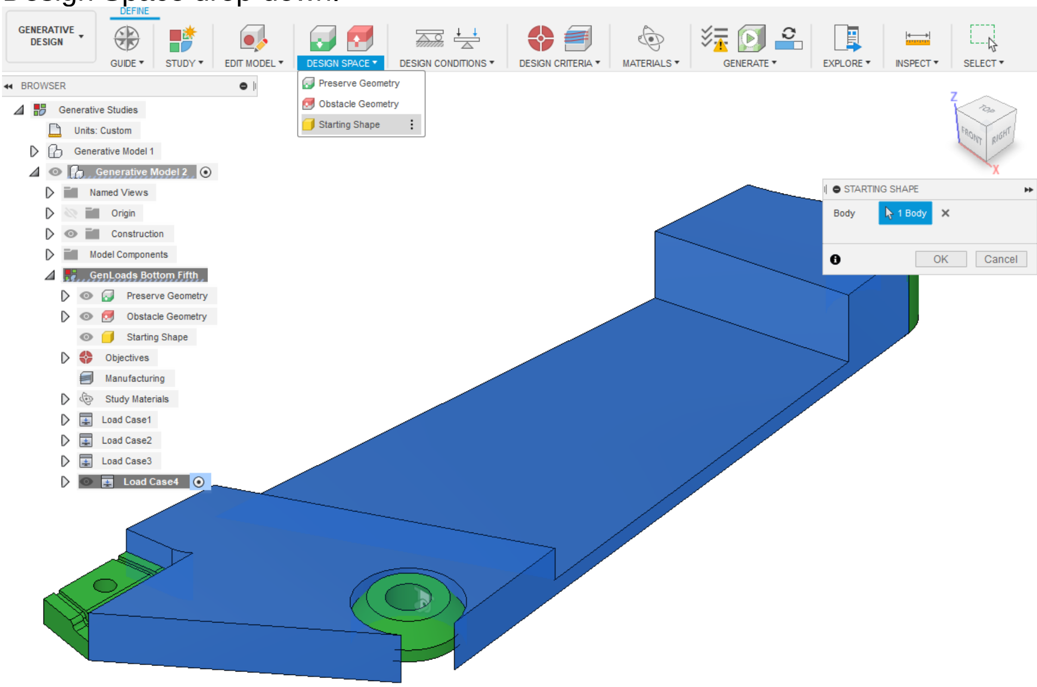


The starting shape will show yellow after finishing the command.

**Note**: Starting shapes affect the initial distribution of material at the start of the generative design process. They do not control the locations of the layers[6]. In the starting shape we created above, the heights of the different extrusions do *not control* the heights of the intermediate layers; they just select which intermediate layers have material at the start of the optimization process. The resulting starting shape in a 3-layer and 5-layer solve is shown schematically in the figure below.

[6] Starting shape can affect the layer locations indirectly when using an advanced technique to change the design domain. See “Starting Shapes” in the Tips, Tricks & Techniques section.

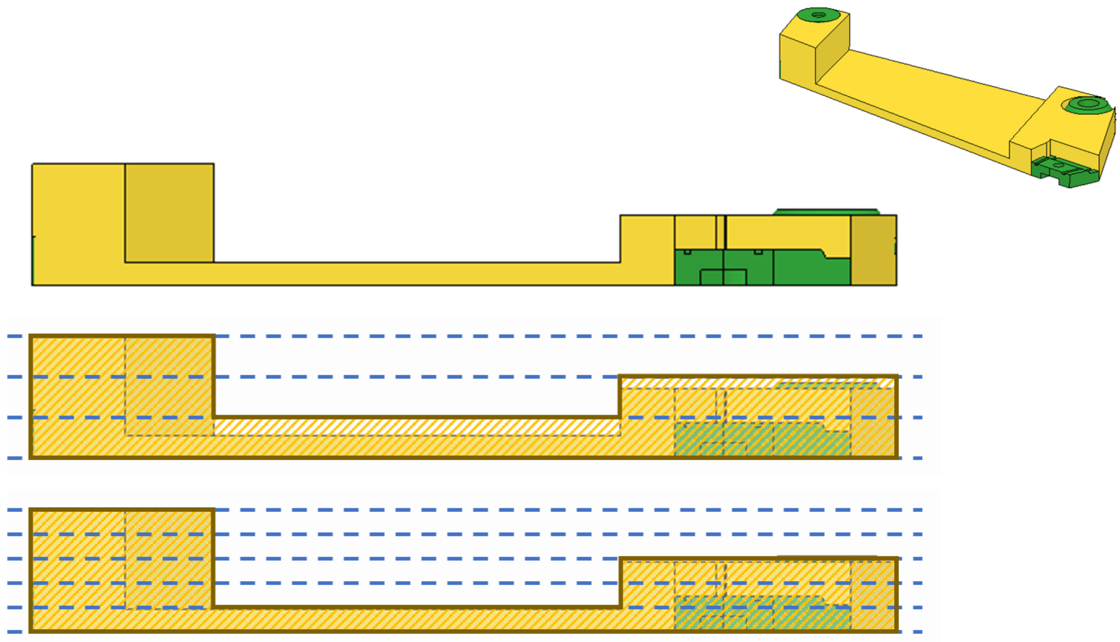

The starting shape we created, and the automatically-corrected starting shape which is used for the 3-layer (-Z) and 5-layer (+/-Z) designs. Note that the layer locations are not selected based on the faces of the starting shape.

Now we can add the new load cases to the generative setup inside Generative Model 2.

### Exploring Results

With the two changes noted above, we re-run the generative solve and explore the results. The Explore environment provides many helpful tools for sorting and organizing the results. A few things we frequently rely upon

<table>
<tr>
<td>We want to see the results of the new study only, so filter by Study in the Outcome filters pane to remove the original SimLoads study.<br><br>We also use this pane to filter for materials and manufacturing methods.</td>
<td>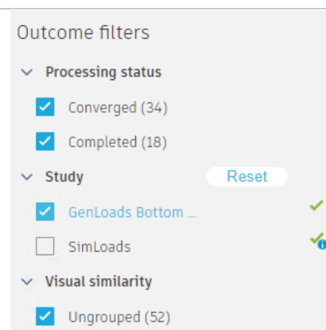
</td>
</tr>
<tr>
<td>As mentioned in the “Initial Generative Run” discussion above, we use Properties View and Sort by Study to determine which 2.5-axis result corresponds to which Configuration (tool size) in the manufacturing dialog</td>
<td>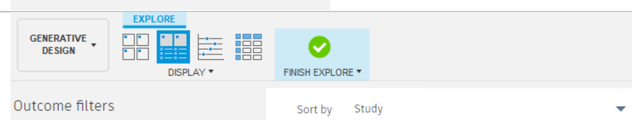
</td>
</tr>
</table>

We like Max Displacement and Volume as axes for the Scatter Plot view, especially for comparing materials or manufacturing methods.

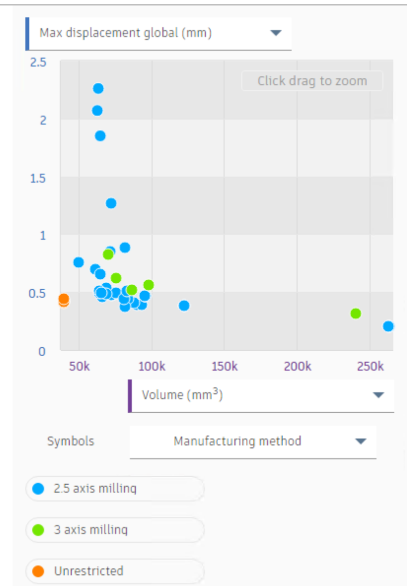


As we find outcomes we like the looks of, we create a few labels so we can find them again. Enter an outcome view (which appears whenever you double-click the image of an outcome or a dot on the scatter plot). Then, use the icon to apply a label to this outcome. Labels appear in the Outcome Filters list for easy access.

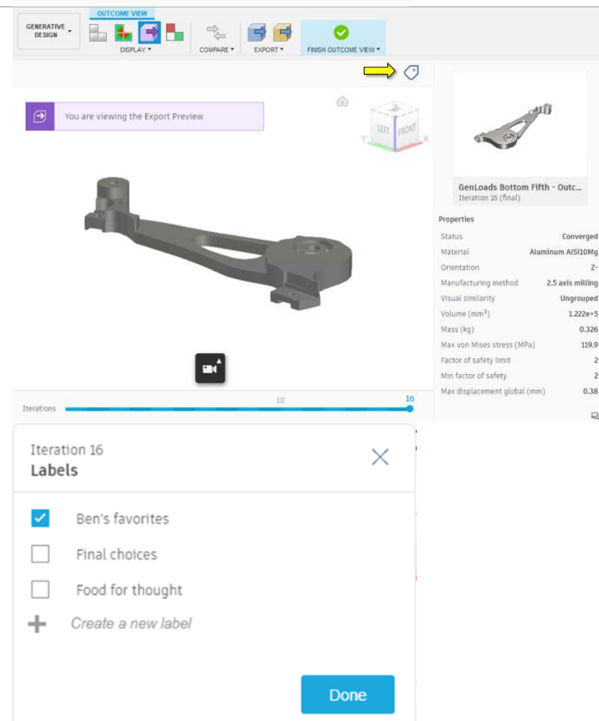

| | |
|---|---|
| After labelling some favorite outcomes, we filter for our label in the Thumbnail View, then use the check boxes to select up to four to compare. Enter compare with the  icon. | 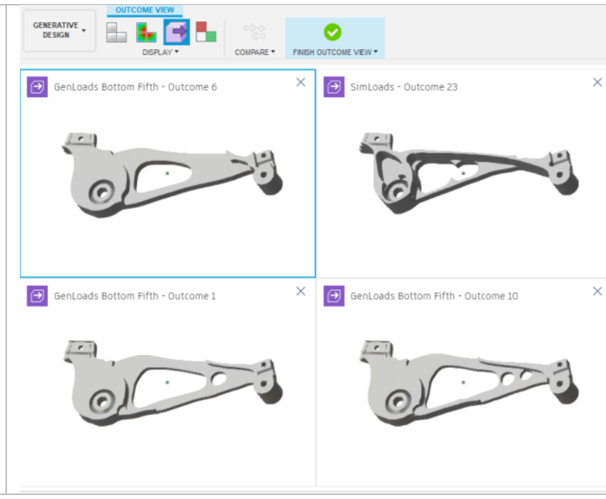  |

## Exporting an Outcome

Ultimately, we selected a 5-layer outcome made from Aluminum AlSi10Mg. This has a somewhat lower yield strength than the intended 6061 alloy, giving us some extra flexibility in making design changes without impacting performance.

To export the outcome, we open outcome view, select the iteration we want, and click Create New Design from Outcome. After a few minutes, the new design is available, either from the


Create New Design from Outcome button in the Outcome View of Explore

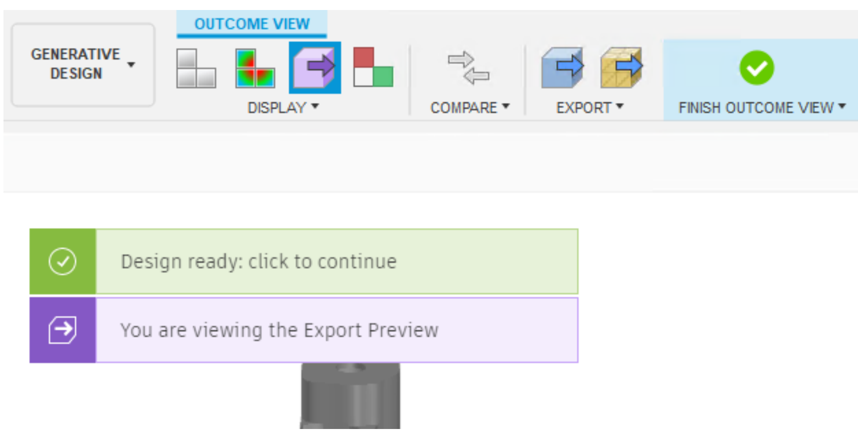


Once the design is ready, click the green icon to open the exported design.

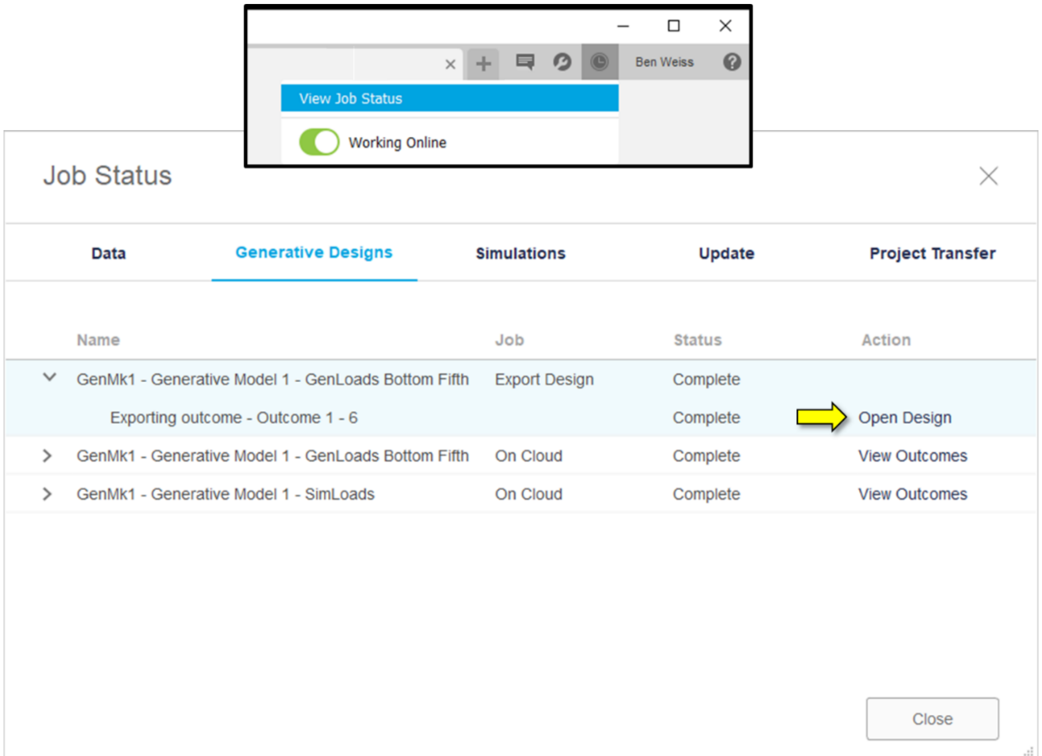


Outcomes can also be accessed in the Job Status page

In addition to exporting the design outcome, look back at the previous few iterations to get a sense of how the design was evolving (see figure below). In this case we notice that a taller wall is slowly encircling the central webbing in the design. In the next section, we'll extend and complete this wall for aesthetic and functional reasons, and it's nice to know that the solver was moving in this direction anyway.

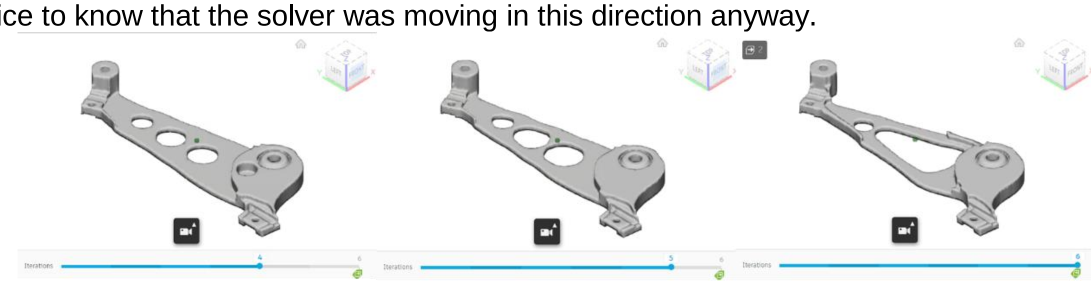

Iteration 4 Iteration 5 Iteration 6 (final)

Last three iterations of the selected design outcome

### Editing 2.5-Axis Outcomes

After exporting a 2.5-axis outcome, a new design is created. While this design could be used as-is, it is typically desirable to make some adjustments and validate the design before manufacturing. Because the resulting design is parametric, aesthetic and functional changes can be easily made. Even if no changes are made, a validation simulation that uses a fine mesh and considers buckling and fatigue (if appropriate) is strongly encouraged, as the generative solver works with relatively low fidelity.

Depending on your preferences, changes can be made using three main approaches. Each approach will find different aspects of this handout helpful, as explained in the following table.

Table 2. Approaches to editing 2.5-axis generative outcomes. Click a section name to go there.

| Approach | Description | Relevant Sections |
|---|---|---|
| *Start-from-Scratch Approach* | Keep the outcome as a suggestion and fully recreate a parametric model that is similar but uses your own primitives (this can even be done in a separate CAD tool using AnyCAD). | "Validation Simulation"<br>*(No other special techniques required; not covered in this guide)* |
| *Edit-the-Outcome Approach* | Directly edit the parametric features of the generated outcome, modifying or adding primitives in the sketches created by generative design. | "Structure of a 2.5-Axis Design Outcome"<br>"Making a Copy for Reference",<br>"Reworking Profile Sketches",<br>"Validation Simulation" |
| *Model-on-Top Approach* | Keep the generated geometry as-is, and add new features to the timeline to change the final shape | "Validation Simulation"<br><br>*Also see,*<br>"Merging Faces"<br>"Horizontal Sliver Faces"<br>"Vertical Sliver Faces"<br>*In the Tips, Tricks & Techniques section* |

There is no "wrong approach" or "best approach" to editing 2.5-axis outcomes; each of the elements in the table above is equally valuable.

The Edit-the-Outcome Approach was used to create the example part we are following in this walkthrough; it also highlights the unique capabilities of Autodesk's generative design solution, which creates a fully parametric solid body rather than a "dumb solid".

Fusion designs created from 2.5-axis outcomes produce a fully parametric solid consisting of sketches and extrusions (unlike other manufacturing methods, which create T-Spline bodies). In this section, first we will walk you through the structure of the outcome geometry, then describe a selection of best practices for editing the results to more fully meet your needs as a designer. Finally, we will discuss the validation simulation.

### Structure of a 2.5-Axis Design Outcome

The design we exported in the previous section is opened in a new tab of the Fusion 360 interface. Its timeline comes pre-populated with some features that create the body shown.

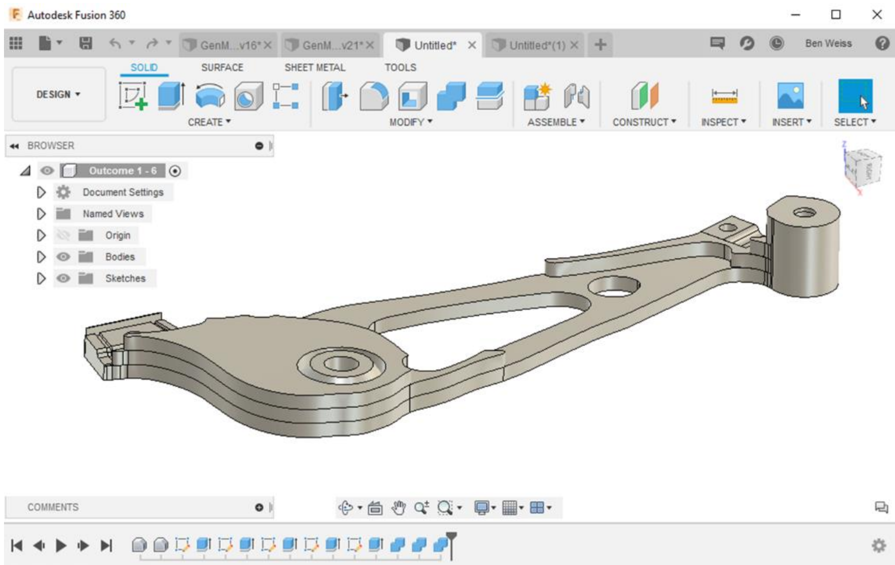

Initial design after export

By rolling back the timeline, we can see what each of these features is doing to the design. The following table explains each group of features.

Table 3. Timeline features of a 2.5-axis generative outcome design

| | |
|---|---|
| <br>The first Base Feature imports the preserve geometry.<br><br>I have colored the preserve features green to make them easier to see going forward. | 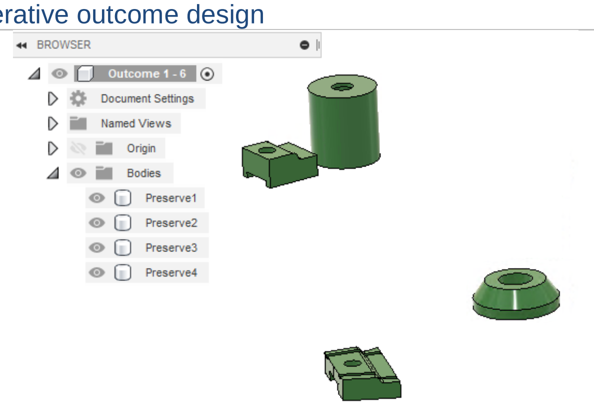<br> |
| <br>The second Base Feature imports the obstacle geometry.<br><br>I have colored the obstacle features red to make them easier to see going forward. | 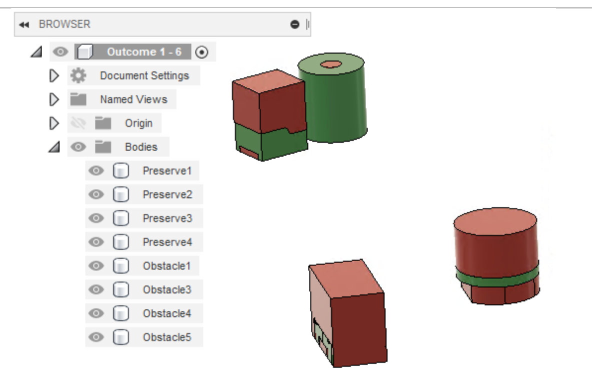<br> |

The next feature creates the sketch for the first layer of the outcome.

The sketch contains spline curves, arcs, and geometry projected from the preserves.

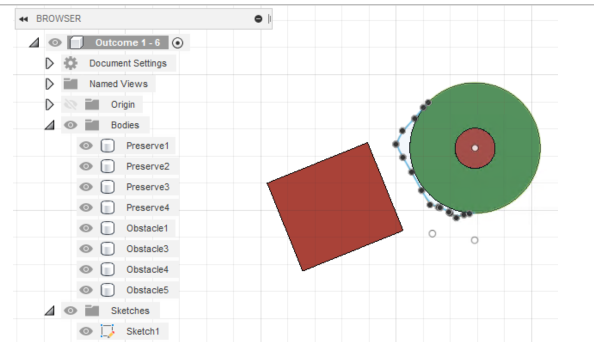


This sketch is extruded to create the first layer body

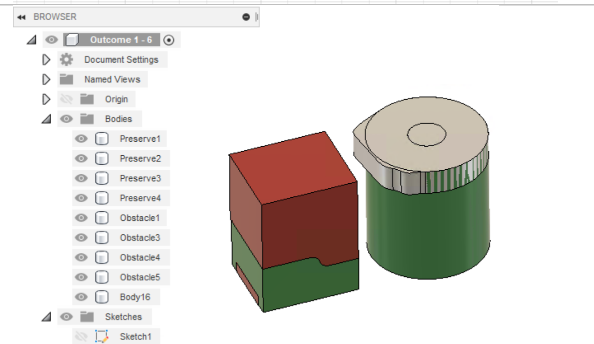


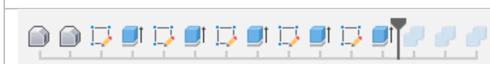

This is repeated for each layer, with each extrude creating one or more new bodies. This outcome has 5 layers, so 5 sketches and 5 extrudes are generated.

The sketch plane for each layer is the bottom face of the previous layer.

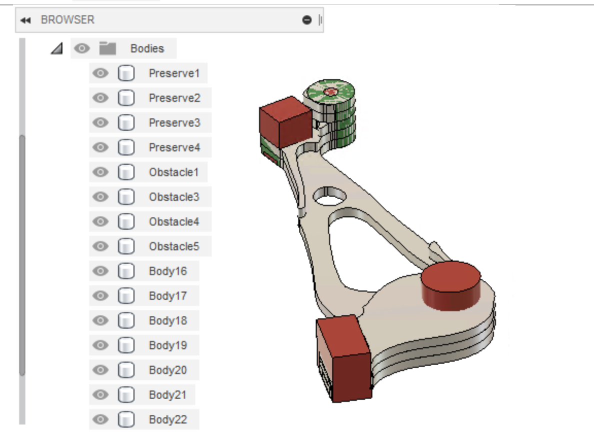


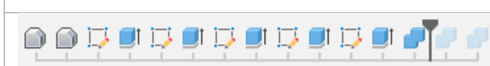

Next, the layer bodies are combined into a single body

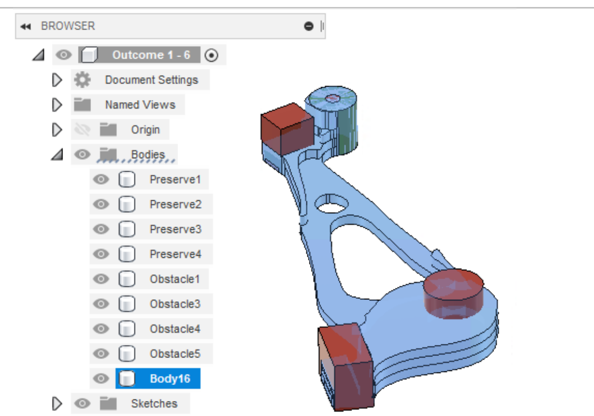

| | |
|---|---|
|  The obstacles are subtracted from this new body using a Combine (cut) feature | 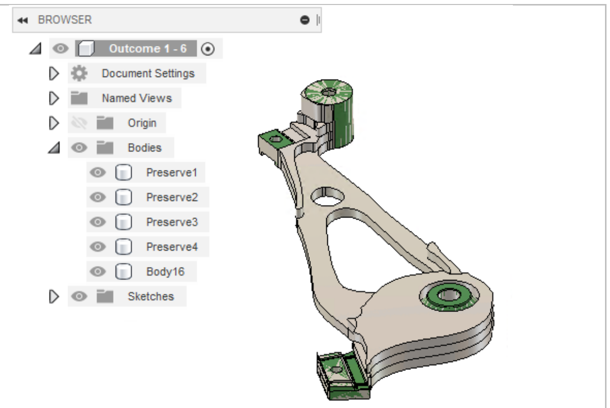  |
|  Finally, the preserve bodies are Combined (Join) with the main body, producing the final shape. | 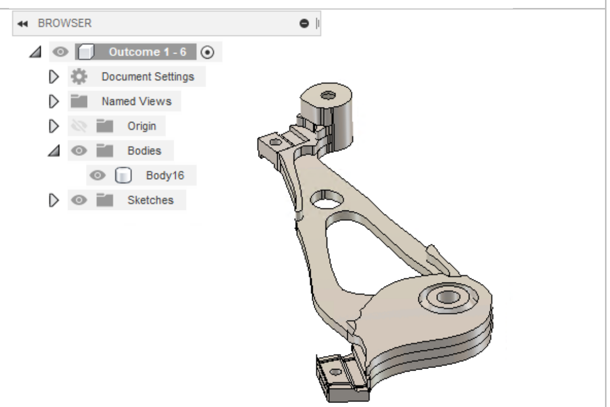  |

Most of the curves inside the sketches are Degree 3 Control Point Splines. For a nice introduction to control point splines in Fusion 360, check out this guide: https://www.autodesk.com/products/fusion-360/blog/sketch-control-point-splines-faq/

### Making a Copy for Reference

As we make changes to the design, it is helpful to be able to refer back to the original shape so we can stay close to the generatively designed geometry. We find it very helpful to create a copy of the original body "frozen" at the initial shape and made available for reference at any step in the timeline. The following procedure describes how to make such a copy.

## Procedure 4. Make a reference copy of a body

1. Right-click the body in the browser (there should only be one), and select Copy.

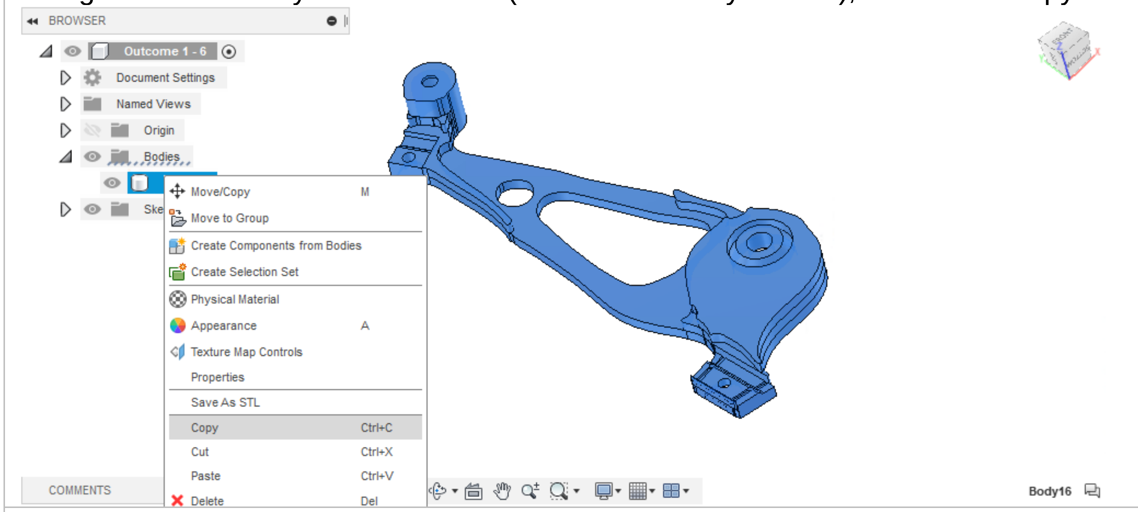


2. From the Create dropdown, select Base Feature. A Base Feature allows direct modelling (history-free geometry) to be incorporated into the timeline.

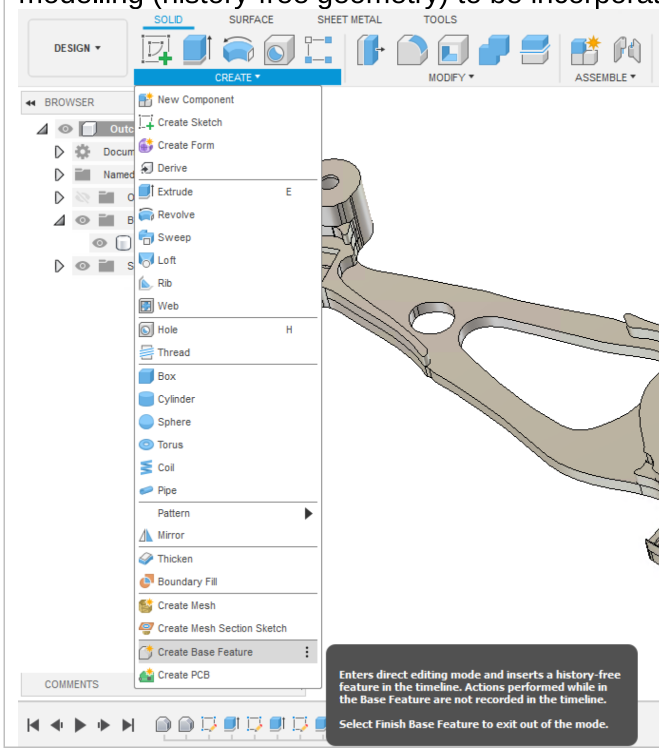

This will enter the Base Feature workspace. This workspace looks almost identical to Design, but note that there are no bodies in the browser.

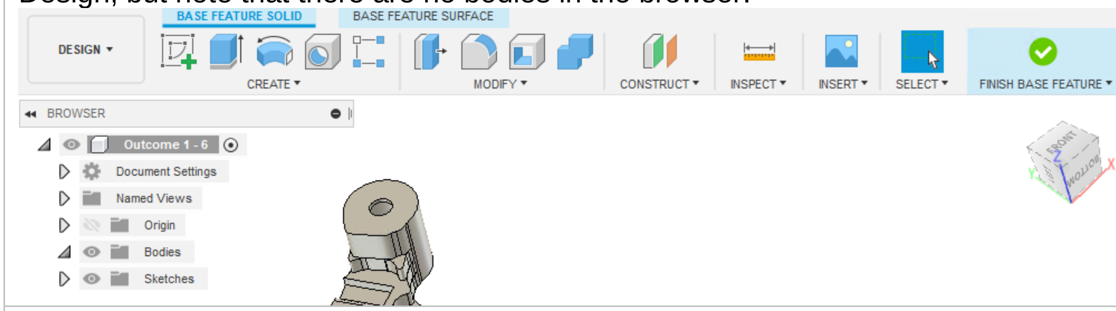


3. Paste the body from the clipboard by right-clicking on the root Component and selecting Paste. Click OK accept the Move/Copy command without moving the part.

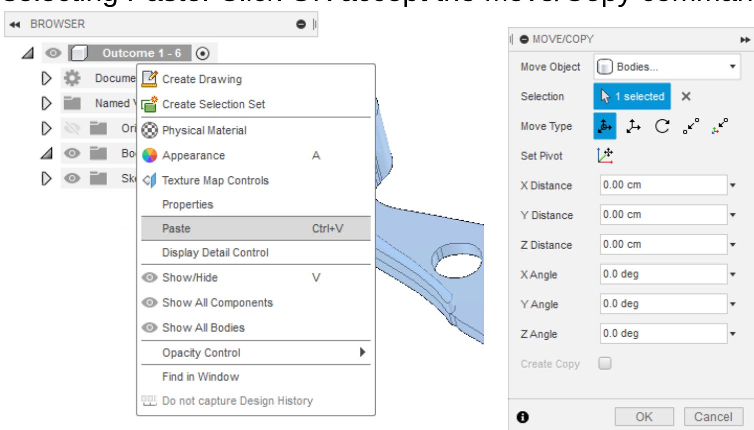


Nothing appears to change in the 3D view, but the body appears in the browser.

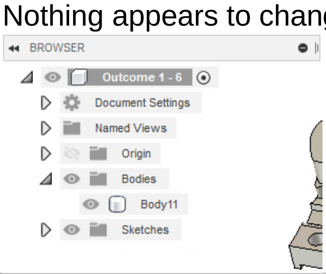

4. Click the green check mark to exit the Base Feature workspace and return to the main model.


Now the browser shows two bodies. One which is history-driven from the timeline and another which is "frozen in time" from the base feature. Note the Base Feature icon at the end of the timeline. We'll rename this body "Original Shape"

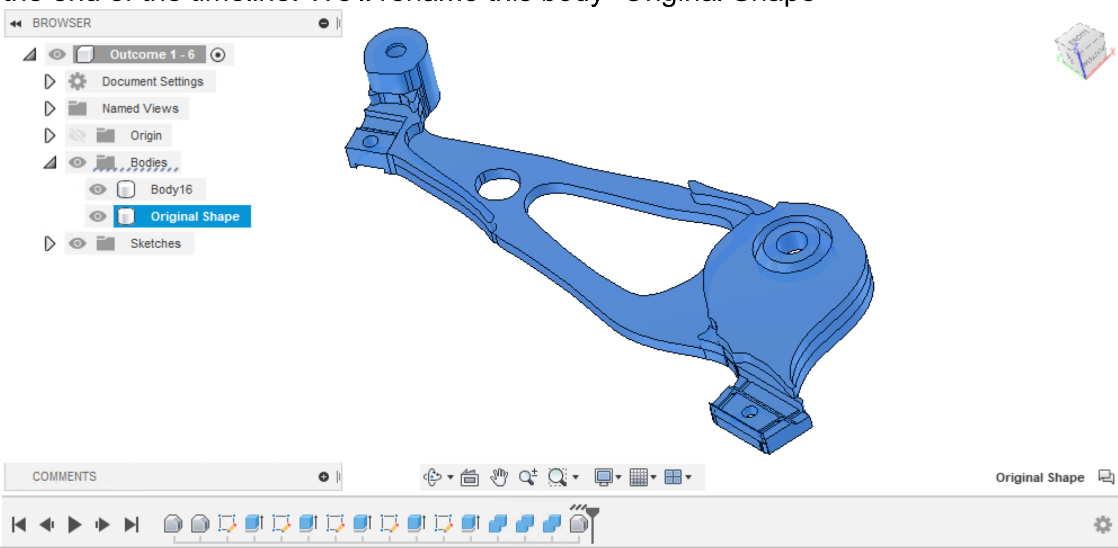


5. Since the base feature doesn't depend on any other elements in the timeline, we can move it back to the very beginning. This way we can see it even when editing features early in the timeline. To do this, drag the Base Feature until it sits at the very start of the timeline.

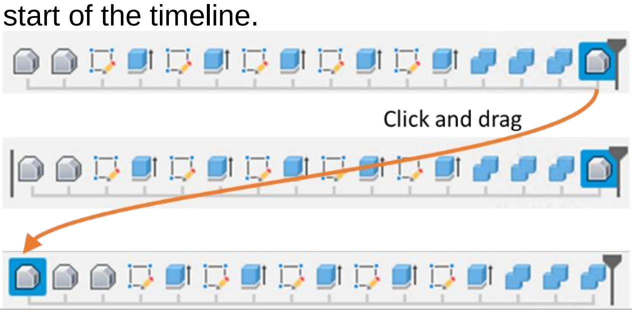


## Reworking Profile Sketches

Instead of walking through every step in the lengthy process of reworking the sketches in this design for aesthetic and performance considerations, in this section we will describe several key techniques and best practices. Using these practices and about an hour of work, we made significant improvements, transforming this from a rough, ugly part to a functional, beautiful machined geometry (see figure).

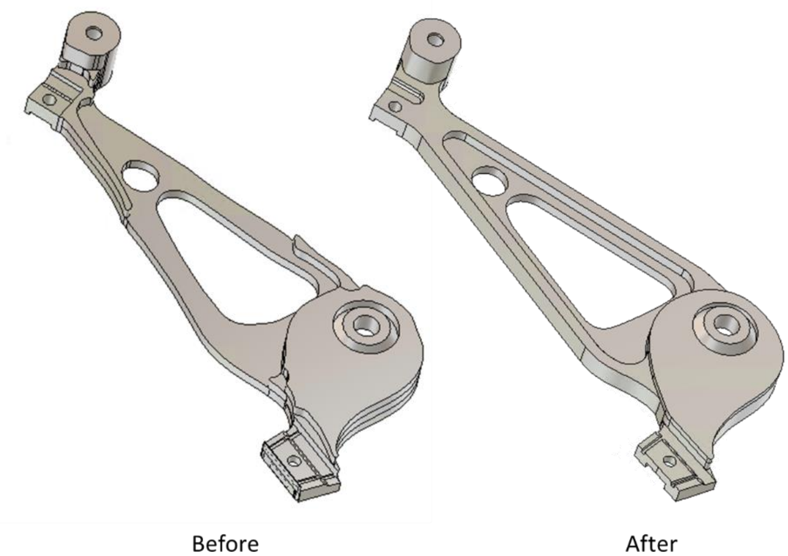


Before and after reworking the design using the tools described in this section

*Cleaning up splines*

The generative outcome produces a sketch for each layer consisting of splines, arcs, and preserve edges. The bulk of the generated geometry is represented as a spline, with arcs inserted where the curvature becomes smaller than the minimum tool diameter calculated for each layer. The figure below shows the raw sketch profile for part of Layer 3 of the example design, which contains each of these features.

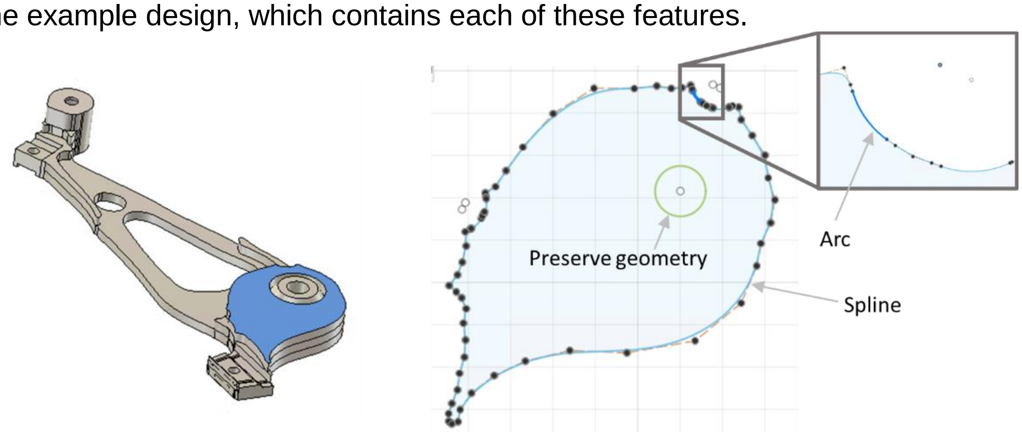


An example sketch from a layer of the design, containing arcs, splines, and preserve geometry

We will make several aesthetic changes to this sketch to smooth the spline and extend it slightly to touch the neighboring preserve in the lower left corner.

Starting on the top, we will remove the two small bumps (these occur whenever the profile transitions from being aligned to the layer below).

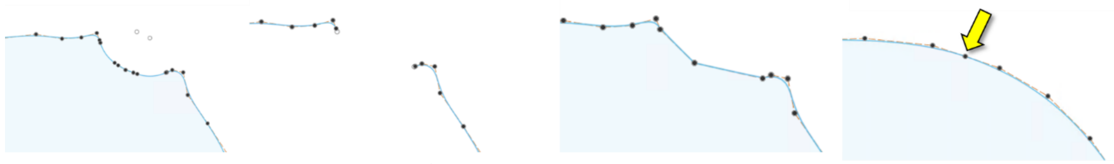

(a) Original (b) Removed arcs/spline (c) Connect spline endpoints (d) Adjust spline points

First edit – bumps in the sketch profile

The above figure shows the steps in this process. Starting with the original profile (a), we delete the two arcs and the spline which connects them (b) and patch the two larger splines together over the gap by moving their endpoints (c). Then, we adjust the spline control points to get a smooth curve (d). Note in (d) the arrow which points to the end of one spline and the start of the next. Use a tangency constraint between the two splines to keep them from producing a cusp later.

At this point, it is advisable to check the modifications against the original design to make sure we haven't drifted too far from the design intent. To do this, show the Original Shape we created in "Making a Copy for Reference" above. I have colored the original shape yellow to make it clearer.

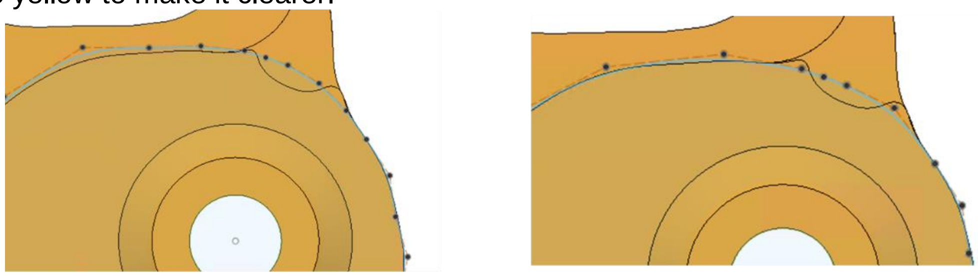

(e) Original Body and tweaked sketch visible (f) Updated sketch

Checking the sketch changes against the original shape contour

Showing the Original Shape (e) reveals that the sketch edits have drifted somewhat from the top profile we did not want to change. Adjusting the spline control points brings it back to lining up on either side of the bumps while cutting an even, smooth line across them.

Continuing clockwise around the contour, we will remove the wobble from the sketch contour, which is an artifact of the solver. As shown in the figure below, we remove about every other sketch point, adjusting the remaining ones to produce a smooth, sweeping contour (h). Then, we again check the curve against the original shape (i) and update to still capture the same overall shape (j).

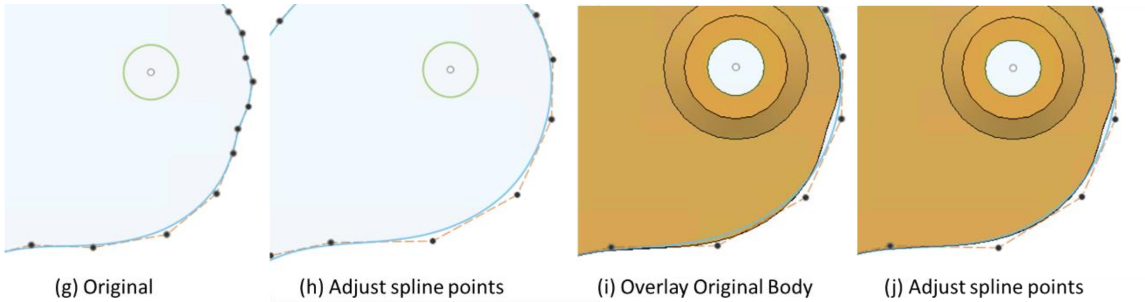


Adjusting the spline to remove “wobbles”

In the lower left corner, we decide to bring the sketch all the way to the edge of the preserve region to avoid a narrow ledge which is not visually appealing. We will effectively replace a section of the spline sketch with a preserve boundary.

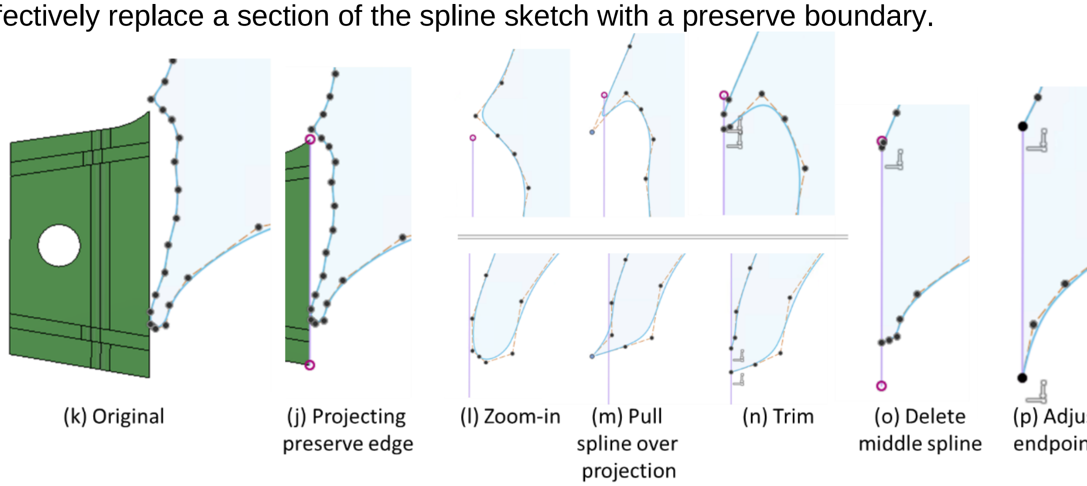


Replacing a spline segment with a preserve projection

To do this, we first project the edge of the preserve body up to the sketch plane using the Project command (Create dropdown -> Project / Include -> Project). We check “Projection Link” to keep the line associated with the preserve geometry in case we make changes down the road (j). Then, at both the top and bottom of the line, we move the spline points so the spline curve intersects the projected line (m). The Trim command (Modify dropdown -> Trim) is used to cut off the spline sections on the left side of the projected line, leaving each spline terminating on the projected edge (n). Deleting the center “shelf” spline (o) and adjusting the spline endpoints (p) complete the task.

The remainder of the contour is cleaned up using the same techniques described above, removing the two remaining bumps and smoothing out the profile to be more visually appealing. Finally, the hole in the center of the layer is removed (it is redundant because a later Combine (Cut) feature will remove the obstacle geometry it represents.

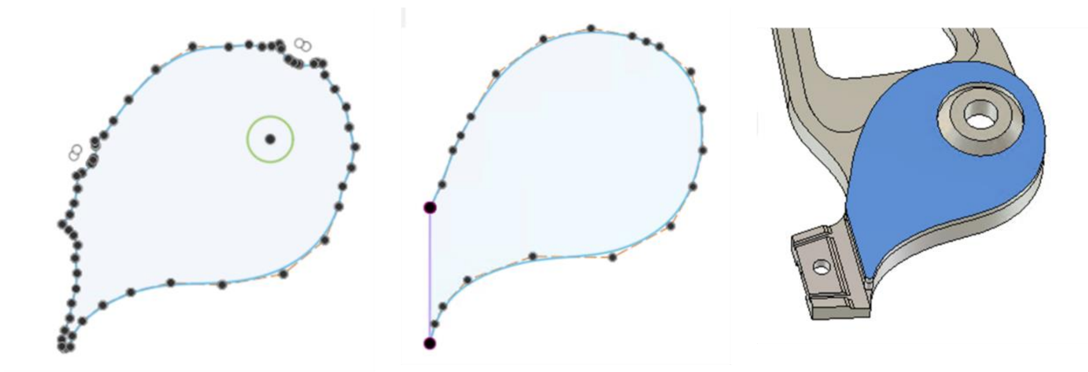

(q) Original (r) Fully Edited (s) Layer in final part

Before and after the cleaning up the spline for this section of the part

Before closing the sketch environment, ensure that the profile is still closed (indicated by the blue shading on the interior and a darker blue highlight when selected). If it is not, check for small gaps between entities or see other tips in “Fixing Open Contours” in the “Tips, Tricks & Techniques” section.

### *Linking layer profiles*

The next layer in the timeline is the layer below the one used above.[7] This layer includes a large amount of geometry shared with the layer above, as pictured in the figure below.

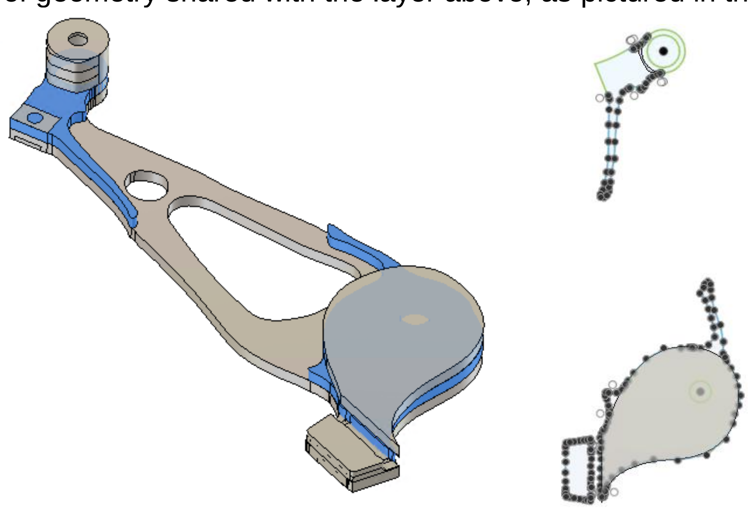

The fourth layer in the timeline produces the blue-shaded region in the part.

We want the curves which are nearly coincident with the layer above to line up exactly, for visual appeal, to avoid stress concentrations by avoiding sharp corners, and to make the CAM process more efficient. In the original design, these two profiles were aligned;

[7] This file happens to arrange layers from the logical “top” to the logical “bottom” of the part, but about half the time it is the opposite, with the first layer in the timeline being the “bottom” one in the geometry. The order in the timeline is the important aspect for this section.

the splines had control points at the same places and the face was smooth. Now, however, the edits made to the layer above have broken the alignment.

When using the Edit-the-Outcome approach, our best practice is to project the geometry from the previous layer onto the current sketch, creating a parametric link between the two faces and ensuring that any further changes in the layer above are automatically captured here. Then, only the “new bits” of this layer need to be adjusted for aesthetics and performance using the techniques shared previously.

First, we will delete the curves in this sketch which overlap with the projected geometry as shown in (b) of the figure below. Then, using the Project command (Create drop down -> Project / Include -> Project), we select the top face of the previous layer (layer above), check the Projection Link box, and click OK (c). The result is a purple curve which will automatically track the face we selected if changes are made to its sketch (d).

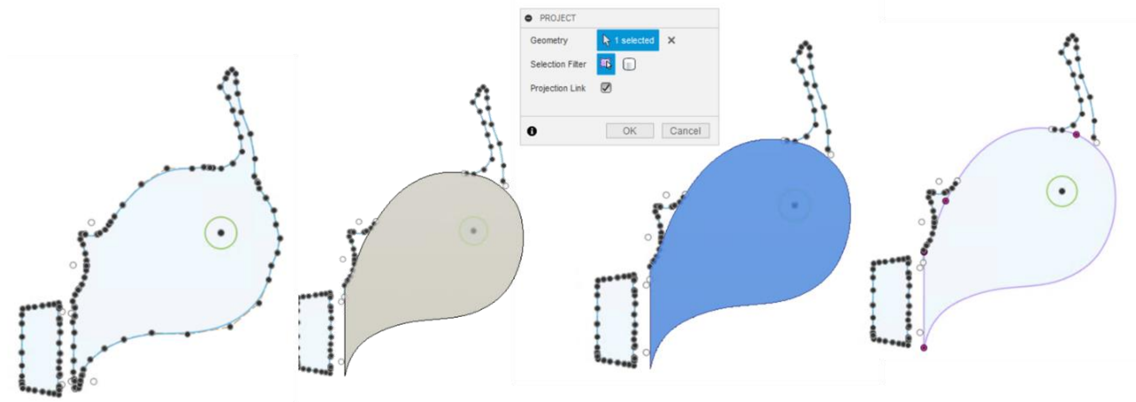


(a) Original (b) Redundant curves deleted (c) Projection command (d) Complete

Creating a linked projection to the top face of the layer above

Now we can adjust the remaining curves to intersect the projected geometry, and select both the projected profile and the newly created profiles in the extrusion command.

Note that there is no way to apply a tangency constraint in Fusion between a primitive and a point on a spline (we can’t apply a tangency constraint on the blue spline where it is coincident with the green one in the figure below). As a result, our approach is to create a sharp corner here, then add a fillet feature to the end of the timeline later (see “Adding fillets and the minimum radius tool” later in this section).[8]

[8] Even the sketch fillet tool does not work when one of the splines is a linked projection.

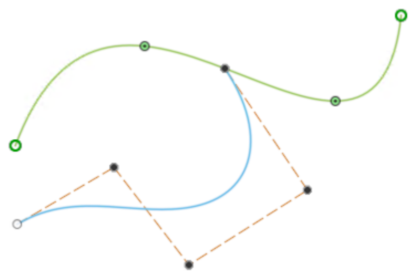

Spline-spline (not endpoint) tangency constraint isn't available in Fusion

As an example, the interface between the protrusion from this layer on the left side and the projection of the layer above is cleaned up following the steps in the figure below. Parts of the curves left from the original culling of the profiles are removed to simplify the sketch (f). On the bottom side, the spline endpoint is moved to make it coincident with the projected profile, while on the top a line is added between the arc and the profile to complete the closed contour. Neither line nor spline can be made tangent to the projected geometry, so the corners are left sharp and will be cleaned up later with a fillet feature.

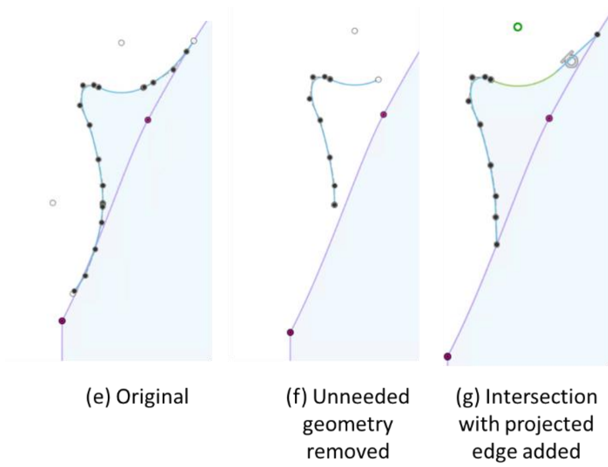


Cleaning up the interface between existing geometry and projected geometry

So far, we have associatively projected geometry from the previous layer onto the current one. For this layer, we desire to extend the partial walls at the top and bottom of the central webbing all the way around the body in order to improve buckling performance (see orange sketch in the figure below). To do this, we wish to project the outer contour of the layer below onto this sketch. Unfortunately, that sketch is inaccessible (it is forward on the timeline, see yellow arrow).

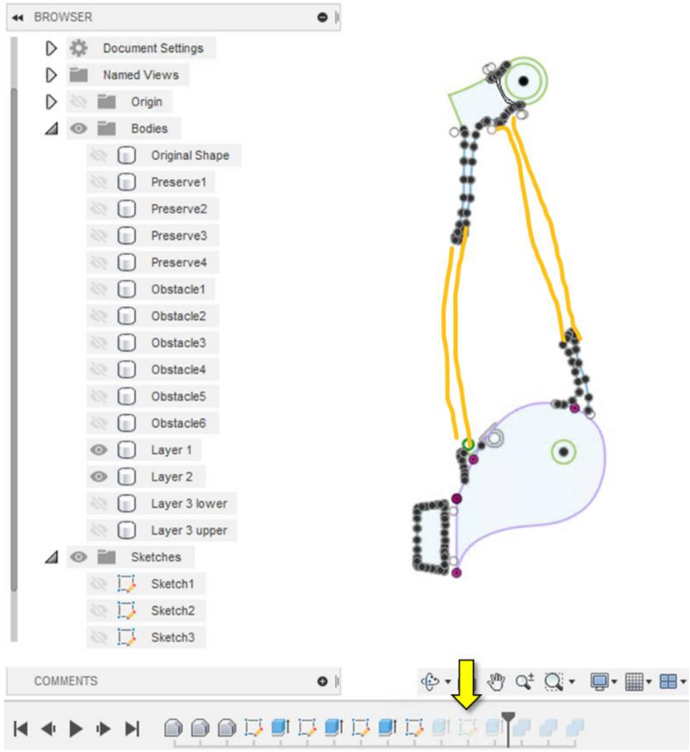


We want to project geometry from a layer that is forward in the timeline

We cannot reorder the sketches in the timeline. Instead, we can project the curves we want from the Original Shape body, as shown in the following figure. We show the body (i), then use the Project command to grab the top face of the next layer down (j). This time, we uncheck the Projection Link option so that we get editable geometry which we can clean up in this sketch (k). Removing the unwanted pieces of the projected profile completes the operation (l). When we move forward to the next layer in the timeline, we will project these edges back from the layer above using the same procedure as before.

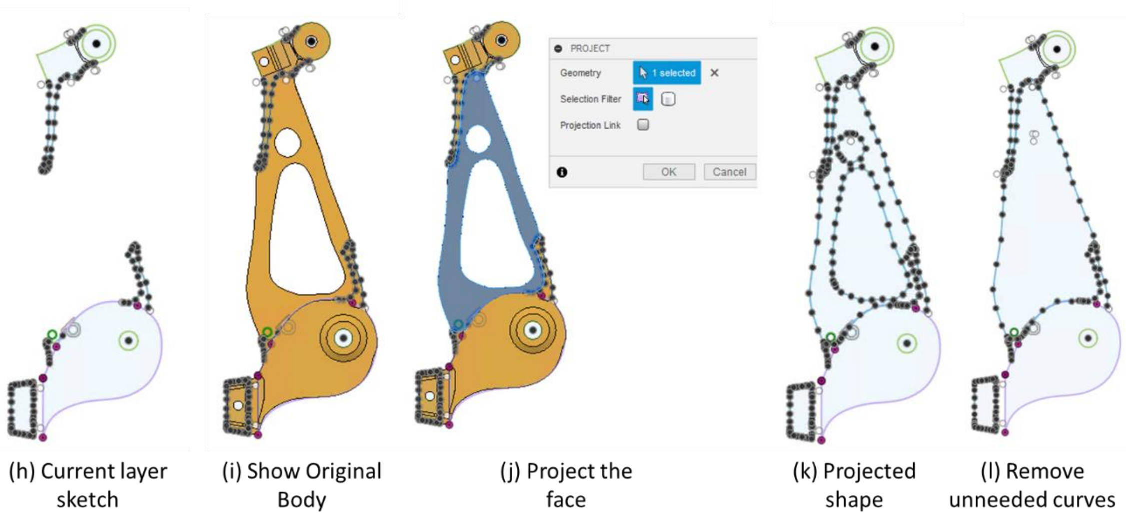


Projecting geometry from a layer forward in the timeline

*Replacing a spline with a line*

The long edges on the profile we just created should probably be represented as a line instead of a spline curve. While there are various ways to achieve this, here's one we find very helpful in editing 2.5-axis outcomes.

First, we'll draw a polyline, as shown in (b) in the figure below. It should be three sides of a rough box, placed as follows. Starting at the top, draw a line segment across the spline, ending a little to one side (*don't* let the polyline snap to the spline). The next segment should stay parallel to the spline's path and follow it as long as it is straight (this segment will become the new edge of the profile). Finally, draw a segment back across the spline and finish the polyline. These two short segments will serve as "tool lines" to make the process easier.

Next, use the Trim tool to remove the spline between the two tool lines (c), and drag the long line segment over until it overlaps the original edge (checking the Original Shape if needed, see (d)).

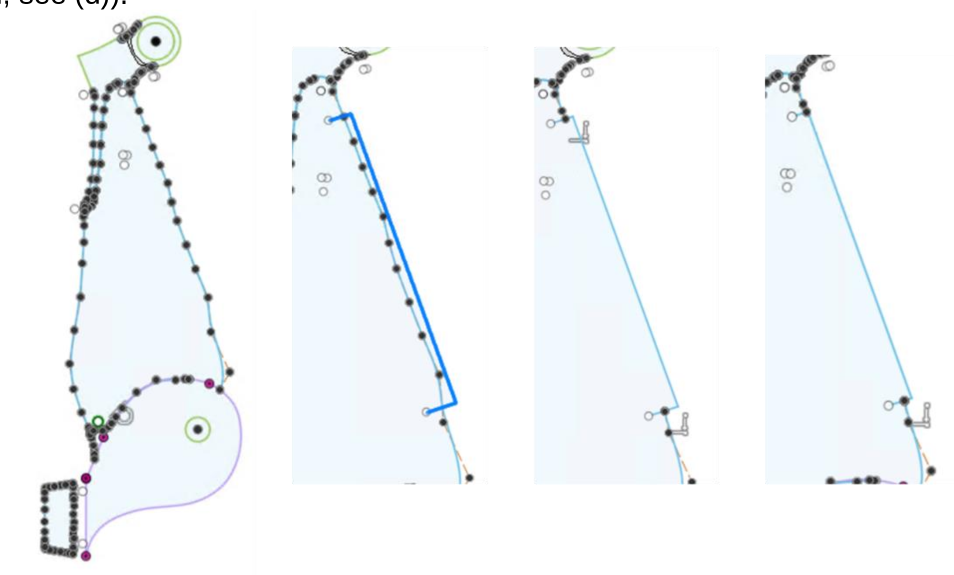

(a) Current layer sketch (b) Line + tools (c) Trim spline (d) Move line into position

Replace a spline section with a straight line, part 1

We want to keep the line still while we move the splines on each end to match it. To do this, use the Fix constraint on the end points of the line (figure below, (g)). Once the line is fixed, apply a coincident constraint between each spline end point and the line end points, then add a tangency constraint between the splines and the lines (h). The line replacement is now complete, and the tool lines can be deleted (i).

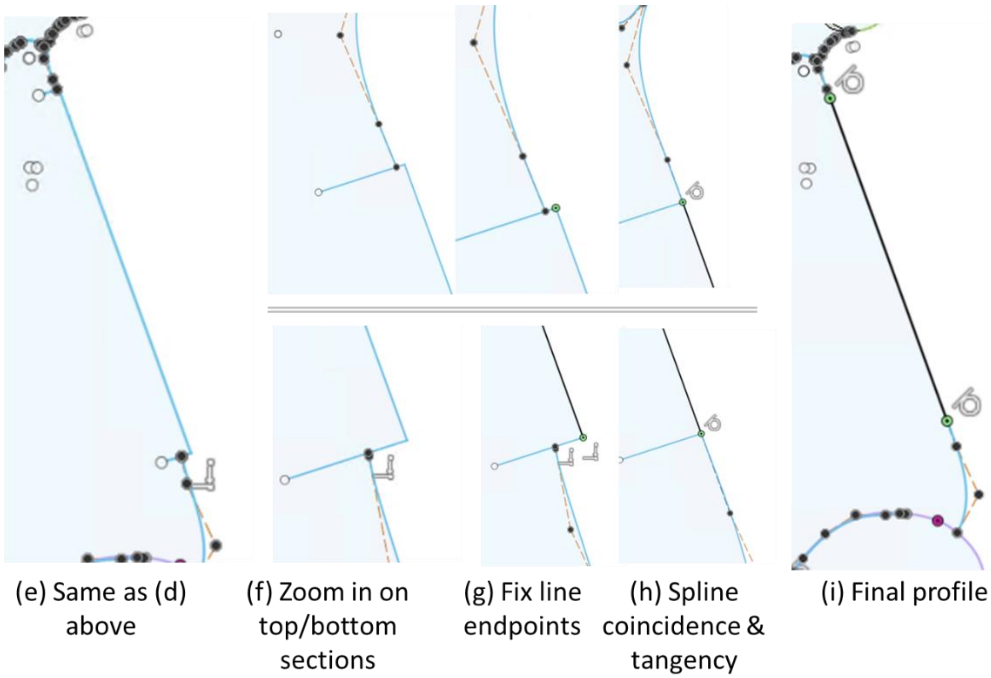


Replace a spline section with a straight line, part 2

*Replacing a spline with an arc*

We may wish to replace a curve in a sketch profile with an exact arc. Again, there are many ways to do this; here is one we like.

Consider the figure below. In (a), the original sketch has a curve which we want to replace with an arc. First, we draw a 3-point circle over the spline[9], adjusting it after drawing until we have the shape we want (c). Next, we'll draw a pair of tool lines that pass through the circle's center and the points where we want to splice the arc onto the curve (c). Next, trim the circle and spline as shown (d). We sometimes find it helpful to shrink the circle while performing the trim to avoid multiple intersections. After trimming, we fix the center and end points of the arc using a Fix constraint to keep them from moving in the next few steps.

On the left side, we can snap the two points in (e) together and apply a tangency constraint (f), but when we attempt to do so on the right side (g), the constraint does not work. Instead, we add a construction line tangent to the end of the arc (h), and make the first control point of the spline coincident with the construction line (i). This ensures tangency between the spline and the arc. The final sketch is obtained by removing the tool lines.

[9] We use a 3-point circle instead of a 3-point arc to avoid automatic coincident constraints, which make it tricky to adjust after drawing.

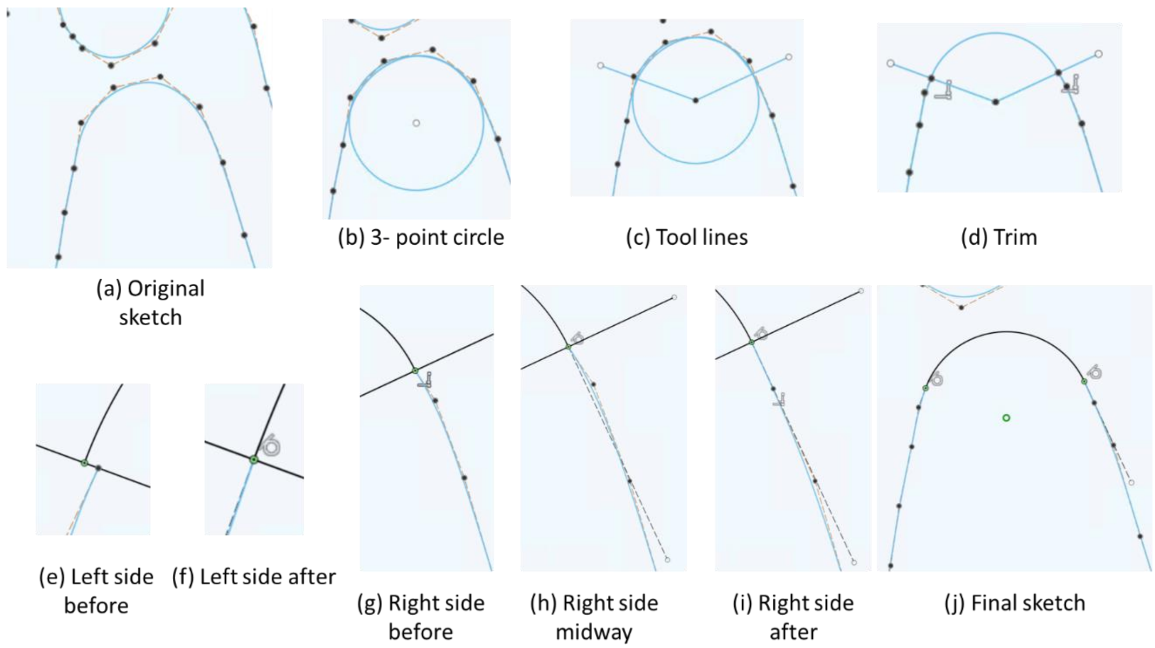


Replacing a spline section with an arc

*Fixing the combine features*

After making edits like those described above, we roll the history marker forwards and almost invariably one of the three combine features in the timeline is unhappy. Enter each combine and re-select the appropriate bodies, as described below.

Each combine has a different purpose and uses different bodies:

1. The first combine unions together all of the layers. Its Target Body should be any one of the layers, and the tool bodies should consist of every body created by the layer extrusions (but none of the preserves or obstacles).
2. The second combine subtracts out obstacles. The Target Body is the geometry produced by Combine 1, and the tool bodies are all of the obstacles.
3. The third combine joins the generated geometry with the preserves. The Target Body is the result of Combine 2, and the tool bodies are all of the preserves.

After the combine operations, there should only be two bodies in the timeline: the Original Shape copy we made at the start of our edits, and the new part.

*Adding fillets and the minimum radius tool*

After all the sketch edits are complete and any combine feature errors are resolved, we want to go over the part and ensure that all points can be cut with the endmill(s) we intend to use.

To support this, Fusion 360 includes a Minimum Radius Analysis tool in the Inspect drop down.

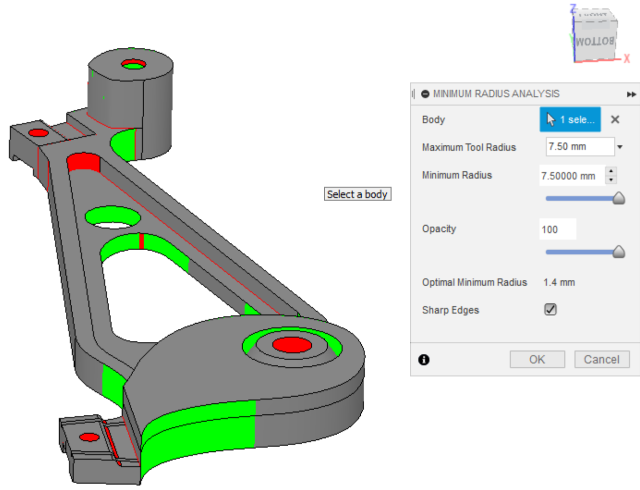


Minimum Radius Analysis tool

Activate the Minimum Radius Analysis from the Inspect dropdown and select the main part body. For this part, we want to be able to cut the exterior profile with a 15mm diameter end mill, so we specify a Maximum Tool Radius of 15mm and slide the Minimum Radius all the way to the right. This will cause any corner with a radius less than 15mm to be flagged. In addition, if we enable the Sharp Edges mode, concave sharp edges are also highlighted. Click OK to finish the dialog.

The command stays visible and updates as we make changes to the design. Note that this visualization has an entry in the Analysis folder in the browser.

For each red highlighted region, we consider our plan for machining and either adjust the sketch, insert a fillet, or revise our plan accordingly. Note that although the sharp corner between each pair of layers is highlighted red, we do not want to add fillets between these layers (it makes the CAM task more difficult).

We can also use the Accessibility Analysis tool (also in the Inspect dropdown) to ensure that all of the features in our design can be machined from one of the two milling setups. As shown in the figure below, we select our part as the Body, the Z axis as the Direction, and the part is color-coded based on its accessibility from the -Z setup (Change Direction can be used to toggle to the +Z setup or both +Z and -Z).

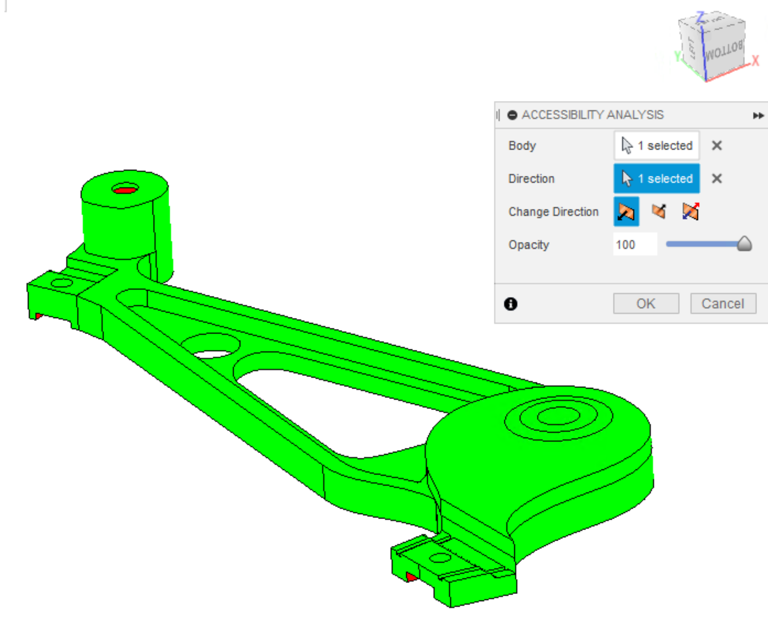


Accessibility Analysis

### Validation Simulation

It is important to validate any design against the mechanical requirements. Because this design began as a generative design outcome, it carries with it the generative design setup's load and constraint information. Entering the Simulation workspace, we are greeted with a message offering to set up a static stress simulation using the generative loads and constraints.

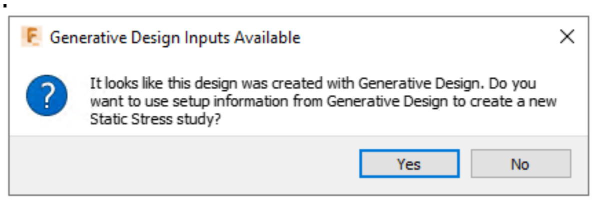


Message on entering the Simulation workspace

In the browser, we see the imported load cases; each with a warning indicating that some of the geometry that the load case references has vanished or changed shape in this new model.

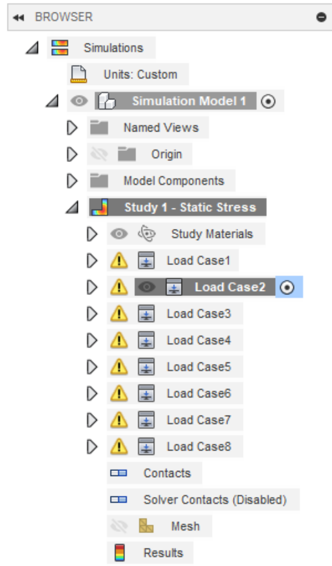


Simulation browser

To resolve the warnings, enter each load case and re-select the faces for the loads and constraints having warnings.

You can use the Simplify workspace inside the Simulation environment to include fillets on the layer-to-layer interfaces if desired for stress or fatigue simulation without modifying the original model, since these fillets make the CAM process more difficult.

## Manufacturing Workflow

### Overview

The Manufacture workspace in Fusion is how a user might transition from the digital design stage of the workflow to a final part. There are many ways to make parts created through Generative Design, with new manufacturing constraints for various manufacturing types being added year on year. Although the part discussed in this documentation is perfectly creatable through additive manufacturing methods, it was created for the subtractive process of milling. Milling is a very fast and reliable process used to produce the majority of industrial parts in modern society. The 2.5-axis constraint for Generative Design produces prismatic parts which are perfect for a milling process, as vertical walls and flat faces are very easy to machine.

In the Manufacture workspace you can access the tools necessary to set up a CAD model for milling. However, before starting this process, it is important to think about how a part might be machined. In this tutorial the part will be created from a block of aluminium.

### The Milling Process

Because milling is a subtractive process, it starts with a block of material larger than the part. These can be cuboidal, cylindrical or potentially even shaped from a forging[10]. This part has detail in the Z+ and the Z- directions[11]. Therefore, it must be machined in two stages, with a series of processes in each direction. This needs to be taken into account when planning the machining process. It is usually best to plan out which features should be machined from which direction before starting the manufacturing setup within Fusion.

There are two routes to go down when machining a part like this – machining from the smallest block possible (which may require a custom fixture to also be made) or an oversized block, which is easier to fix and set up but requires more material. In this section we will explore both options, but the main walkthrough uses an oversized block. A smaller starting block means less material needs to be removed and thus the manufacturing process is faster. The dimensions of a bounding box on the final part can be found by right-clicking on the body in the browser and selecting Properties. The result is shown in the figure below.

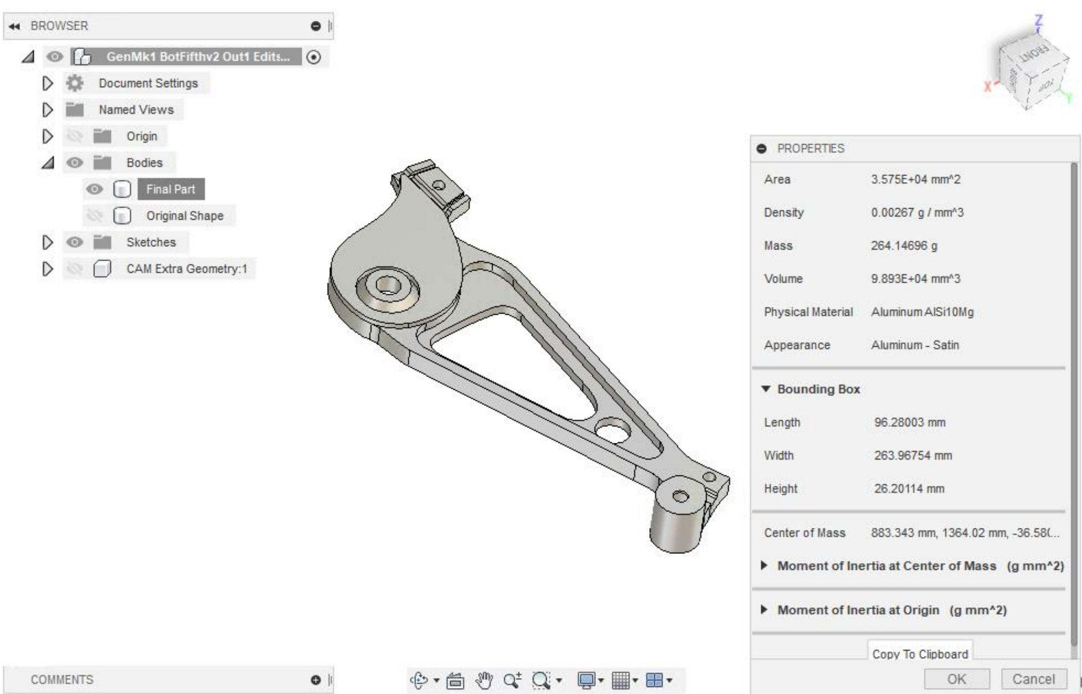


Bounding box dimensions of the final part

Usually, adding a couple of millimetres to each machined face of the part results in a correctly-dimensioned starting block. It is good practice to have a stock to remove on

---

[10] A forging is a part produced from identically-named process that uses compressive forces to shape a part. It is not uncommon that subtractive manufacturing starts with a block formed through this process.

[11] Even though the generated geometry only requires a single setup, one of the preserve bodies has a pocket which must be machined from the back side.

each face, even if the top surface will remain flat. This is so that any potential defects or issues in the surface of the starting block can be removed and the part will be in a bright, machined condition. Once you have material of the correct dimensions, the part can be set up in the Manufacture workspace.

The other option requires slightly more work up front within Fusion. However, this time will be saved when machining the part. This is because a dedicated fixture does not need to be made and the clamps holding the part to the machine will not need to be moved. This method of manufacturing will create a part in the centre of a larger block, attached by tabs that can be later sawed off. This process is quite common for parts which are machined as a one-off. For larger quantities of parts, it can be more prudent to create a dedicated fixture to sit the part on.

### CAD Modifications for Manufacture

In order to use the oversized stock approach for machining this geometry in Fusion CAM, we will explicitly model what we want the part and stock to look like after machining is complete. The image below shows the additional CAD created for manufacturing purposes. The oversized stock allows the material to be clamped to the table without worrying about collisions between the clamps and the cutting tool. The black blocks in the image show where the clamps might be positioned in this example. The aim of this setup is to ensure that the entirety of the part can be machined, while remaining securely attached to the machine throughout the manufacturing procedure. The tabs are machined at the end of the program to allow for easy removal of the part with a bandsaw or similar machine.

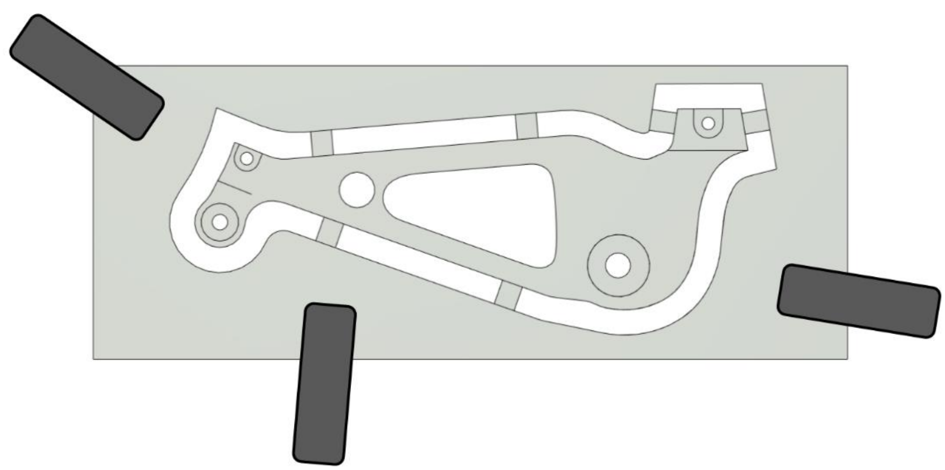

Layout view of the part in the stock (back side; clamps shown)

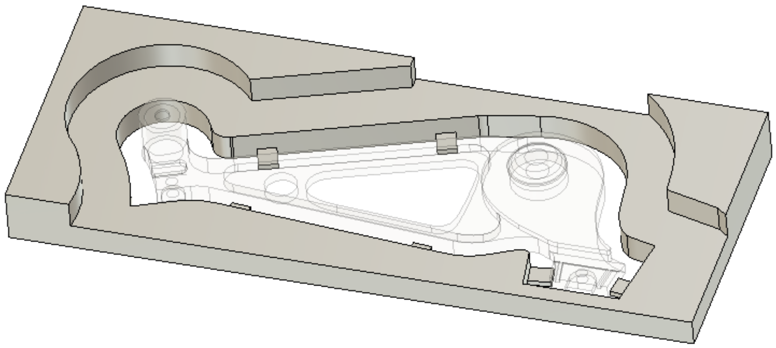

Perspective view of the part in the stock (front side; part greyed)

The first step is to create a model of the oversized block.  Since we want to be sure to remove material from all sides, the stock should have at least a couple of millimetres added on all faces, but in this case, we add additional material in the X and Y directions to allow room for clamping.  We chose to use a stock with dimensions 360mm x 140mm x 30mm (corresponding to an additional 96mm, 43mm, and 3.8mm for each axis, respectively).

Before creating the stock geometry in the Design workspace, we will add a new component to separate the machining-support geometry from the main part body. Use New Component in the Create drop-down, then rename the component in the browser to "CAM Extra Geometry".

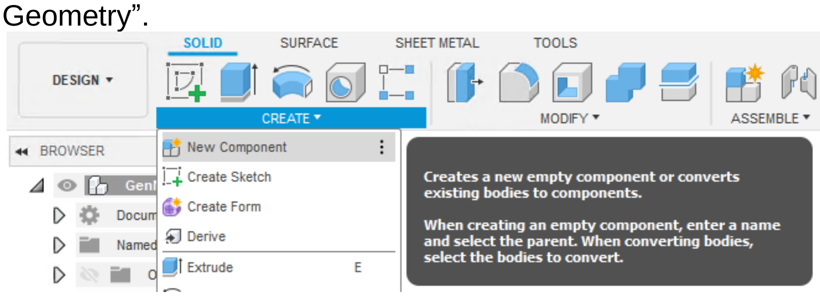


New Component command

The new component is automatically set to Active (note the black radio button next to its entry in the browser in the figure below). This will place new modelling operations inside this component instead of the root component of the design.

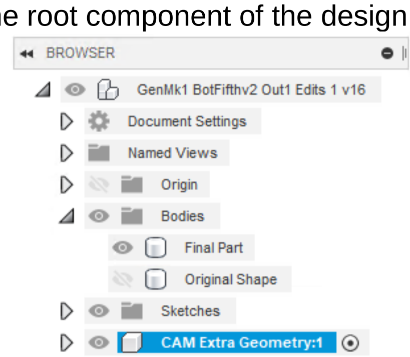


Browser after creating and renaming the "CAM Extra Geometry" component

The stock body is modelled using a Box feature (in the Create dropdown). Select the bottom face of the part, then draw a rectangle around the part using the dimensions of our stock (360mm x 140mm)

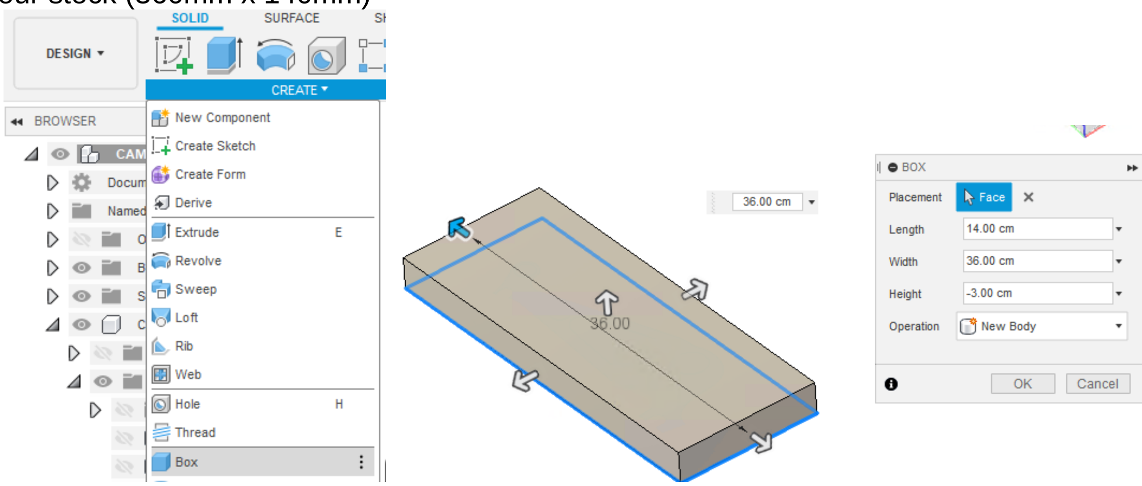


Box command. The box completely covers the part body

After box has been created to the correct dimensions, it should be repositioned, both to ensure there is room to clamp as shown in the figure on the previous page, as well as to move the stock body down 2mm so that there is stock material on all sides of the part. This can be accomplished using the Move/Copy command in the Modify dropdown.

The as-machined stock body needs to have a cut out so there is space for the tool to remove material around the part profile.  Therefore, a pocket must be cut into this block that the parts and the tabs can sit inside.  The easiest way to accomplish this is to project the part body onto a sketch and extrude its "shadow" all the way through the stock body.

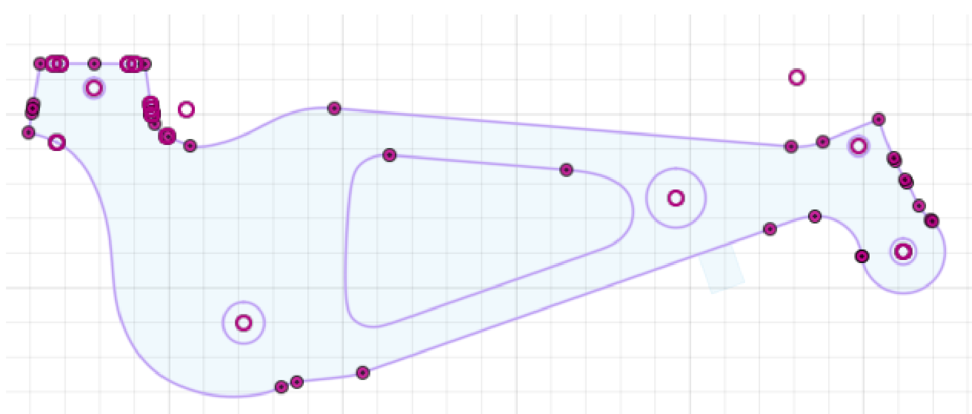

Projection of the part body onto a sketch on the bottom face of the stock.

We extrude this profile to create a new body and make it large enough to pass all the way through the stock body, as shown below.

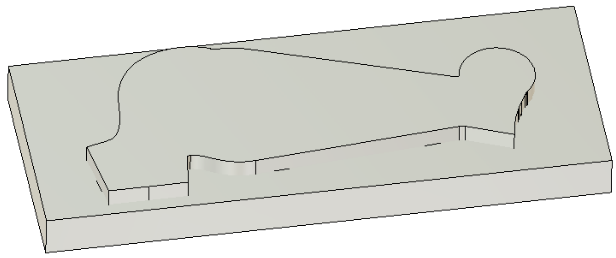

Projection sketch extruded far enough to pass through the stock body.

This extruded model can then have its outer faces offset by a set distance using the Offset Face command (Modify dropdown) by an amount larger than the largest tool diameter we plan to use. In this example the largest tool we plan to use is 10mm and the we offset the faces by 12mm.

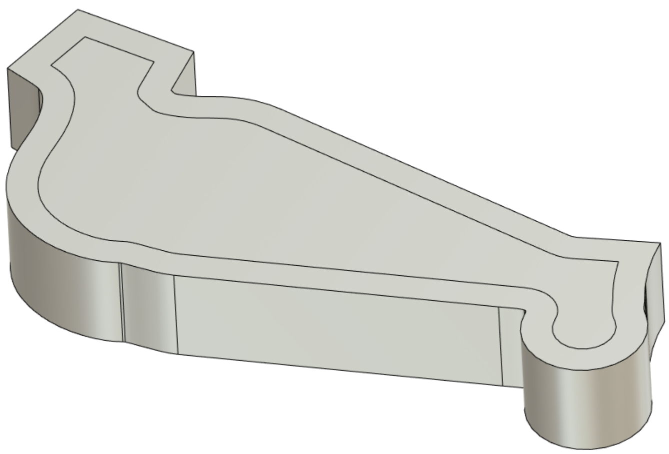

Offset of the profile extrusion faces

Once the solid has been formed it can be cut from the block using a Combine feature (Modify dropdown), which gives the result in the image below.

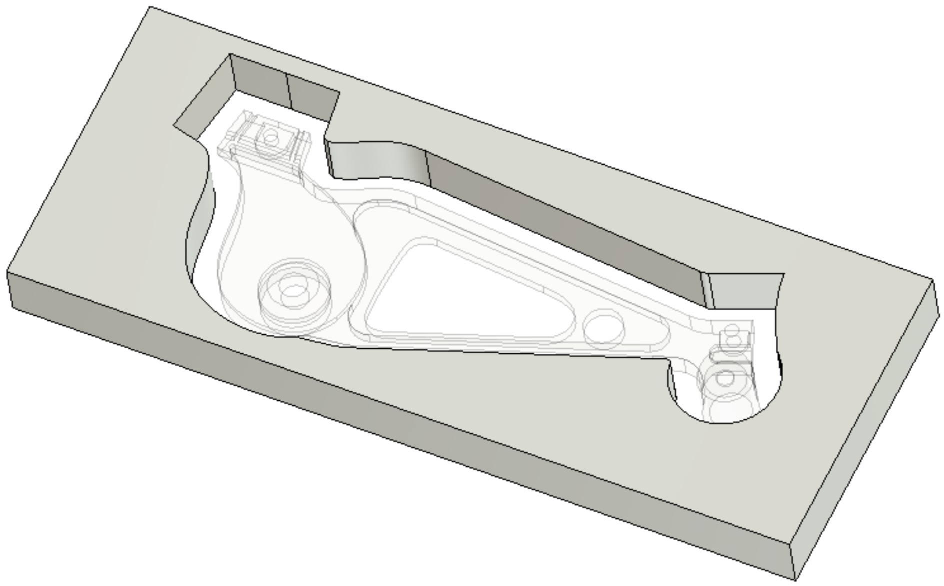

Stock body with offset profile subtracted out

There are two additional steps to modifying the CAD. The first is creating the tabs. This can be done with simple sketches around the periphery of the part. Whenever you create a part in this way, these tabs should be placed evenly around the part so that the part can remain as supported as possible. The tabs should also be placed on areas that are flatter, if feasible, as this makes them easier to remove.

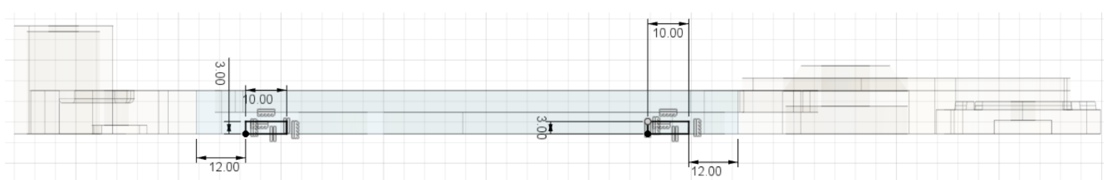


Sketching tabs on a flat face of the bottom layer

The tabs should then be extruded the same distance as the offset for the pocket (12mm in this example).

We also need to consider the shape of the cutting tool and milling machine itself. A cutting tool should only be a certain length, dictated by its diameter, in order to stiffen the tool against machining forces. This means that, depending on the depth of the part, the tool may not be able to reach where it needs to cut without the tool holder and the part colliding, as illustrated below.

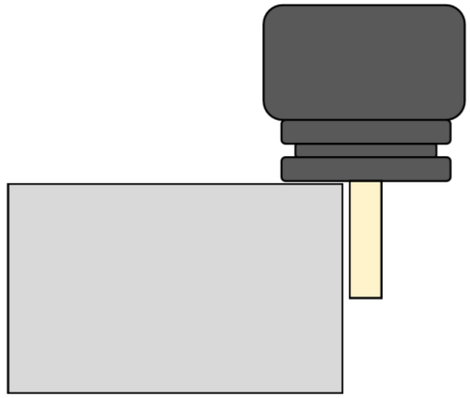

Tool cannot reach all the way into the part without the tool holder colliding

To account for this space required by the tool holder, part of the block can be machined away. This requires a small amount of extra time but means that the tool is able to reach the correct depth safely.

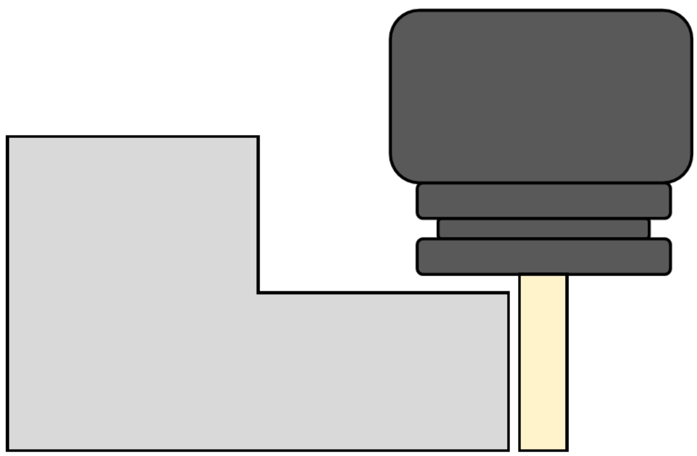

Tool holder relief in the residual stock enables the tool to reach all the way to the bottom of the pocket

The process of modifying the block is completed in a similar way to creating the pocket just prior. Another sketch is created following the profile of the pocket. However, instead of extruding the sketch and then offsetting as we did above, the sketch in the example is offset directly by 25mm.

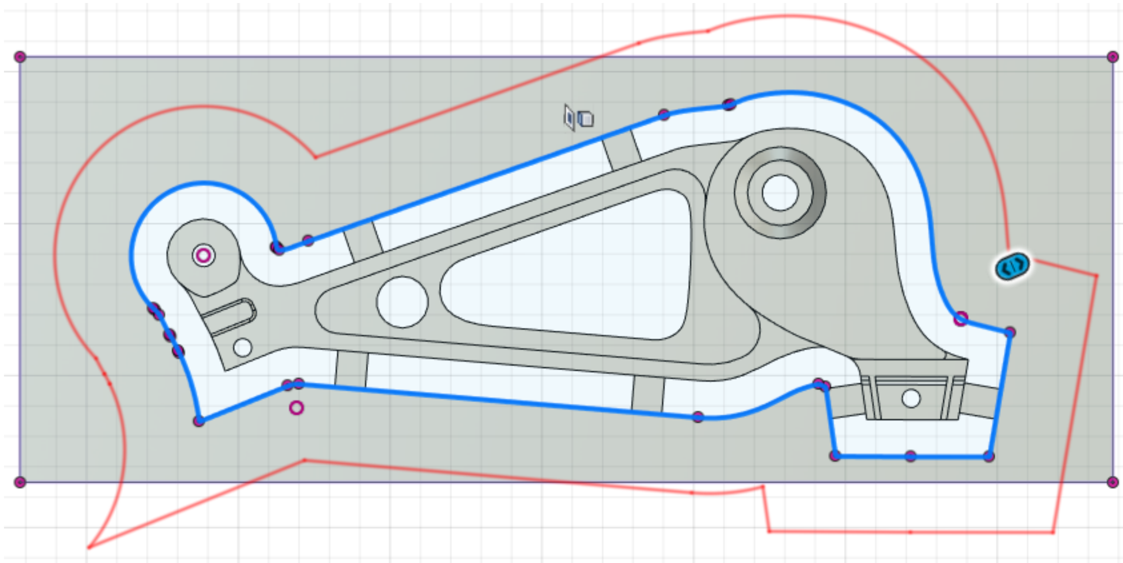

Tool holder relief pocket, initial sketch

This sketch can then be cut into the part and everything would be fine, but some small optimisations can be made. This block is not part of the design, so this sketch can easily be modified without much impact to the quality of the final result. Thin walls are not great to machine and can create sharp and dangerous edges, so the thin sections at the bottom and right sides of the image above are removed by adding to the sketch and trimming away unnecessary curves.

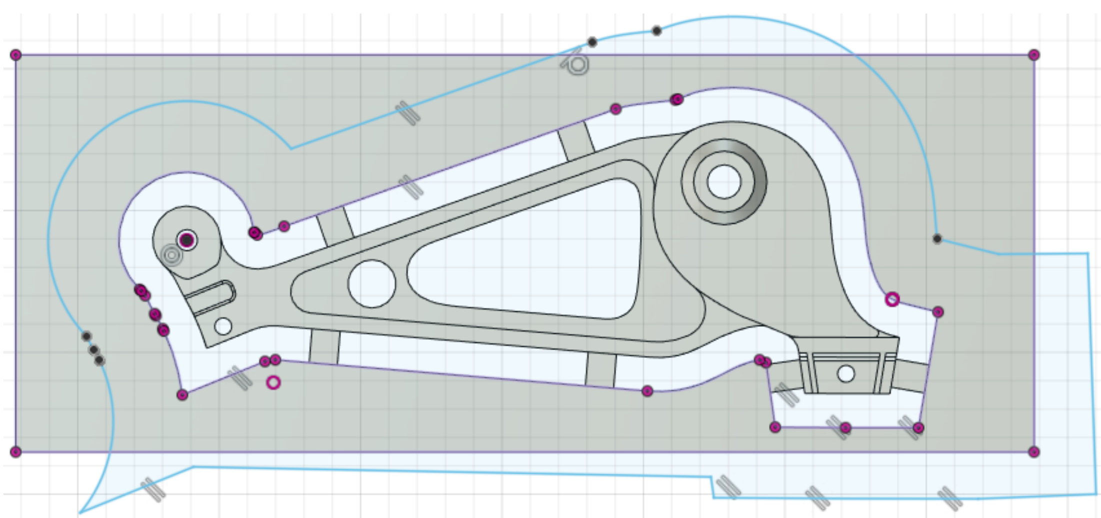

Edited tool holder relief pocket, mid-way through edits

A final modification can be to remove the sharp edge that would be formed at the top of the image. Currently, the offset profile forms an acute angle as it passes over the edge of the block. This is worth removing, for safety purposes, again by editing the sketch curve. Radii can be added to corners in the profile for the same reason but also for ease of manufacture, as flowing curves prevent stresses on the tool and can also mean that parts can be cut faster due to fewer sharp changes in tool direction.

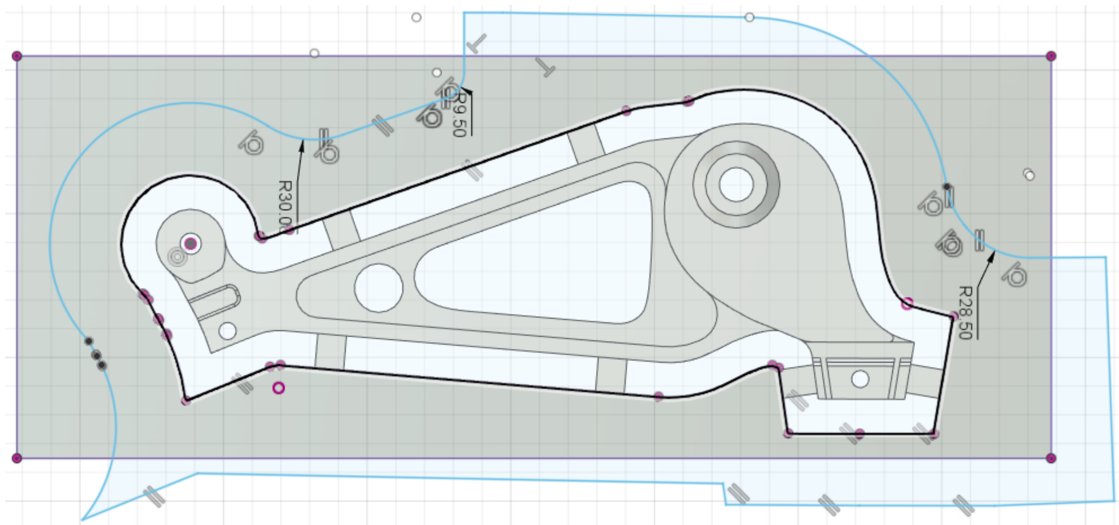

Tool holder relief pocket, final sketch

This sketch is extruded/cut part-way down into the stock body to produce the final shape shown in the figure below. The additional CAD geometry for manufacture is now be complete. Now we switch to the Manufacture workspace and program the part.

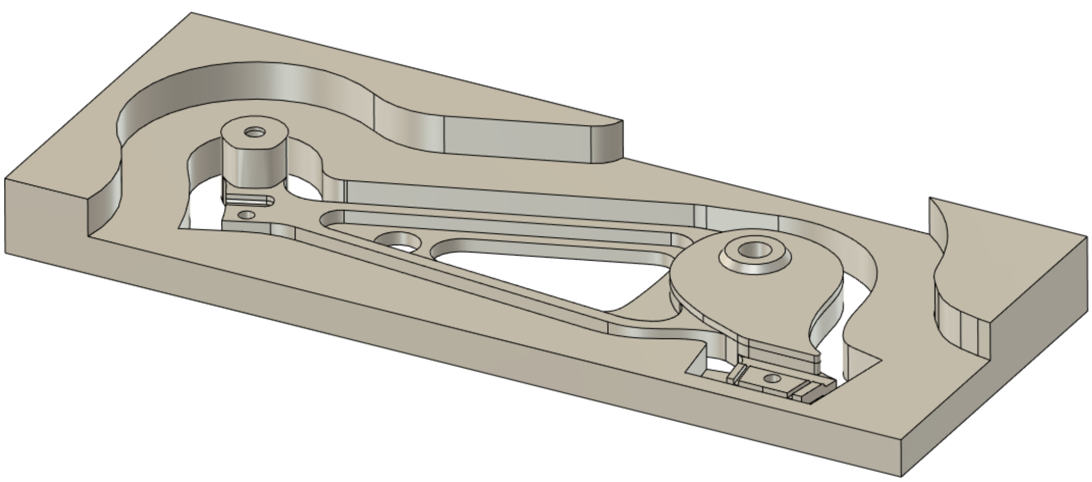
Final CAD geometry (part plus modelled stock)

## Setting Up For CAM

We switch to the Manufacture workspace using the menu in the top left.

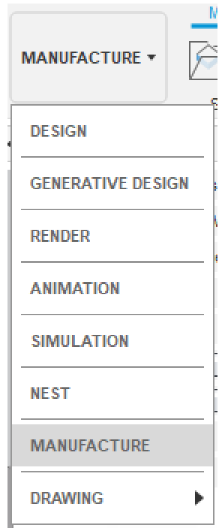

Manufacture workspace in the workspace menu

The Setup function is the first step towards realising the physical part you have within Fusion.

Setup button

The Setup dialog allows you to select or create the block of material you have for your part (stock geometry). As mentioned previously, the dimensions of the block used for cutting this example are 360mm x 140mm x 30mm. This block as already been created as a CAD model, so either a relative or fixed size box will work in this scenario. We select all of the machining models (the original part, the tabs, and the modelled remaining stock) in the Model section of the Setup tab. Select the remaining stock for the Fixture, and choose the coordinate system as shown in the figure below.

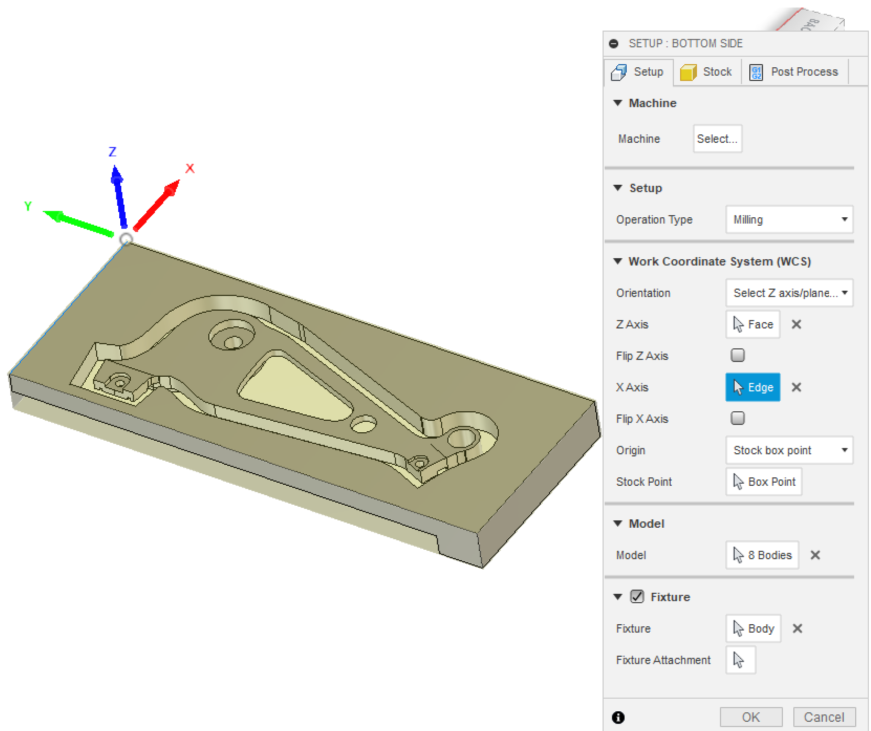


Setup dialog, Setup tab

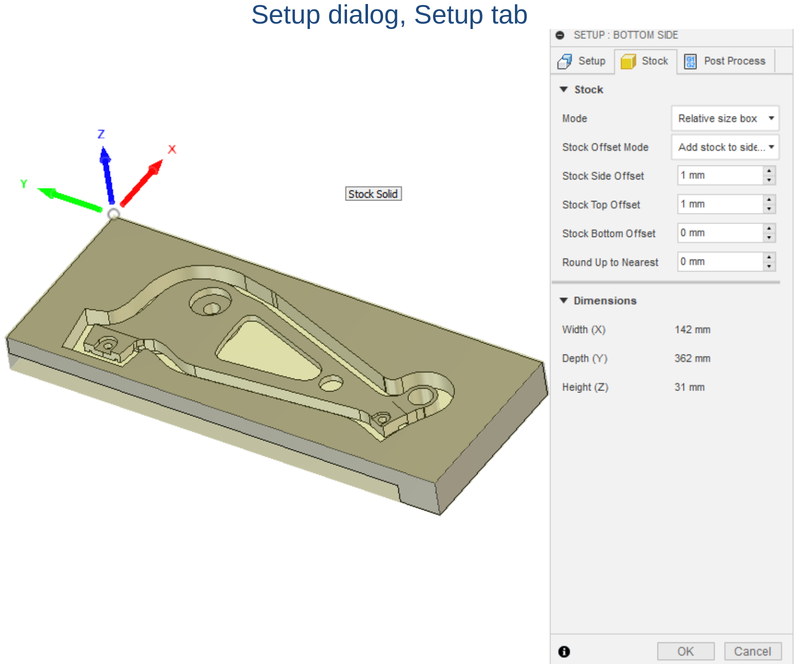


Setup dialog, Stock tab

The part can then be cut as per a usual milling program.

## Design to Part Time & Effort

Because the example we used in this tutorial is a real part for a real customer, we actually walked through a design process fairly similar to the one described above (which we've simplified slightly to keep the presentation concise). Roughly speaking, the time required for each step is shown in the following table.

Table 4. Materials used for first generative study

| Step | Time Required | Comment |
|---|---|---|
| Generative setup | - | This was the assumed starting point of the exercise |
| Generative solve 1 | ~2 hours | Didn't wait for it to finish before drawing conclusions |
| Examine results; update setup | 0.5 hours | |
| Generative solve 2 | ~4 hours | Over night. In practice, we ran several solves in parallel here with different configurations |
| Examine results; select outcome | 0.5 hours | |
| CAD geometry cleanup; validation simulation | 1 hour | |
| Design iteration | 1 hour | Preserve geometry changed when we discovered we couldn't machine it as designed. |
| CAM programming | 2-3 hours | Includes creating stock geometry |
| Machining | 1 hour | |
| **Total** | **5 hours hands-on**<br>**6 hours hands-off (generative solves)** | |

If the process started mid-afternoon so the main generative design solves occurred overnight, a machined part could be ready to hold by the same time the next day. This ability to design a part given its place in an assembly in under a day is a major step in improved engineering productivity.

Depending on industry, additional design iterations or validation steps may be required, so your results may vary. 2.5-axis generative design produces a high-quality initial design, in fully-editable parametric CAD geometry, optimized for an inexpensive manufacturing process. Your design is generated on the cloud while you go home and have dinner with your family, and because the results are parametric, you can still incorporate your aesthetic style and design preferences, along with responding flexibly to requirements changes.

This ends the walkthrough portion of this guide. The rest of the document contains an appendix of sorts, covering a variety of additional tips and tricks for the 2.5-axis generative design process. We hope you've found this content helpful, and if you're ever in Portland, OR, USA or Birmingham, UK, please drop by the Autodesk office and say "hi" to Chris or Ben. We'd love to meet you!

## Tips, Tricks & Techniques

### Generative Design with 2.5-Axis Machining Constraints

This section emphasizes the differences between the 2.5-axis machining constraint and other manufacturing methods in generative design. For some additional background and tips for other manufacturing methods, check out the great videos on Fusion 360's YouTube channel:
https://www.youtube.com/watch?v=LaG9uVj_grE&list=PLmA_xUT-8UlJxhI506AoIAziSq5HFs3Cv&index=4

#### Using an Arbitrary Tool Direction

As mentioned in the walkthrough, for 2.5-axis machining constraints, the tool direction must be one of the coordinate axes. In order to use a different direction instead, the part must be rotated to align the desired tool axis with one of the coordinate planes. The easiest way to do this is using the Edit Model feature of the Generative Design workspace. Follow these steps:

Procedure 5. Rotate a model to use an arbitrary tool direction

1. Let's start with this design, which is 2.5-axis compatible and has a clear milling direction, as shown, but isn't aligned with the coordinate axes.

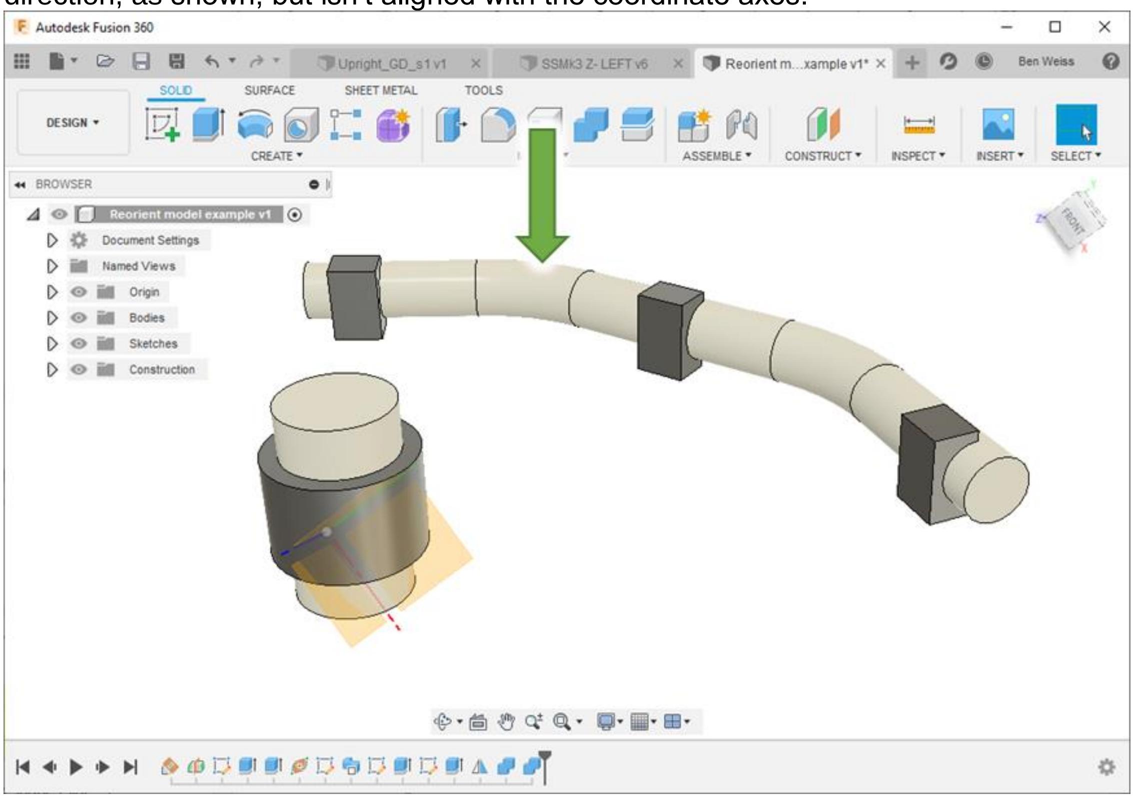

2. In order to make this work, we need all the bodies to be part of a single component. Use the Add Component button to create a new, empty component[12]

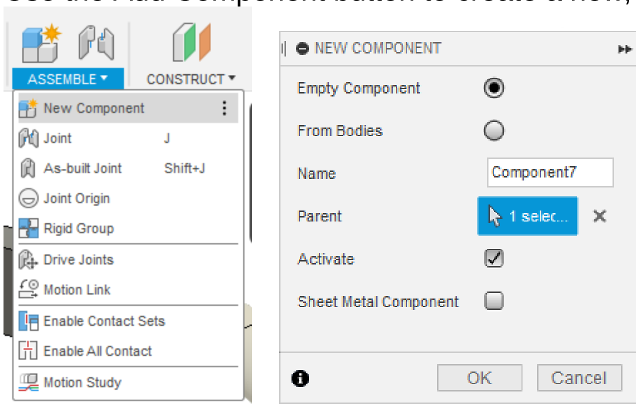


3. Move the bodies to the new component by dragging them into the component in the browser

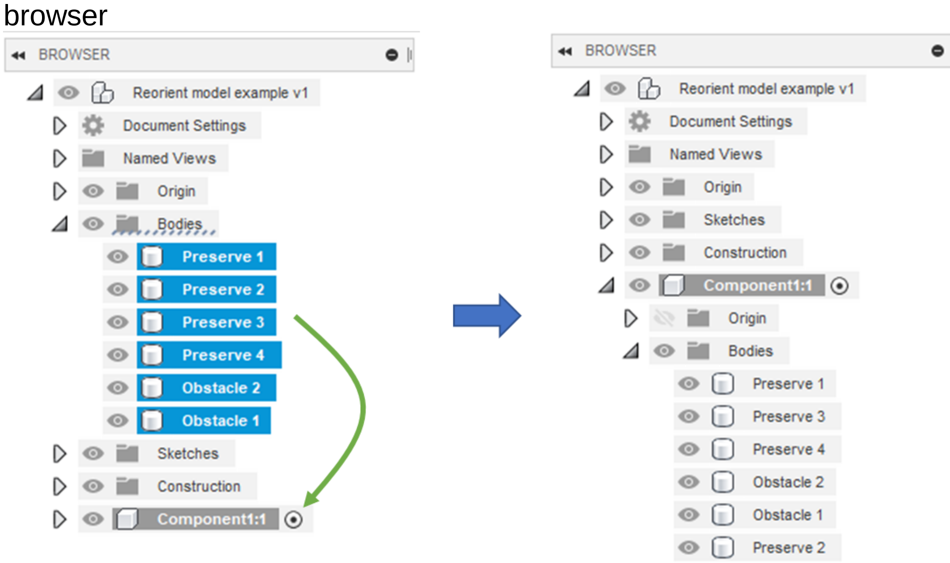


3. Switch to the Generative Design workspace and enter the Edit Model environment. Changes made here are applied only to the current Generative Model, and do not affect the original design.

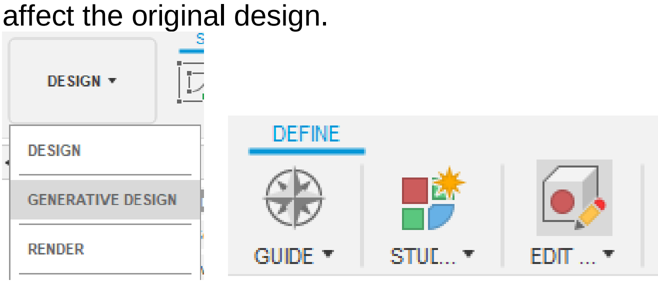

4. In the Modify dropdown, select the Align command

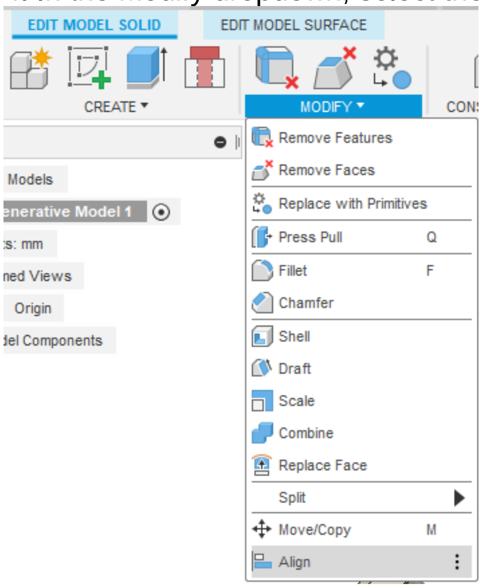


5. In the Align dialog, select Components from the "Object" dropdown, then select a face perpendicular to the milling direction for "From" and the XZ plane in the Origin for "To". The whole component should be re-oriented so that the selected face is coplanar with the selected plane.

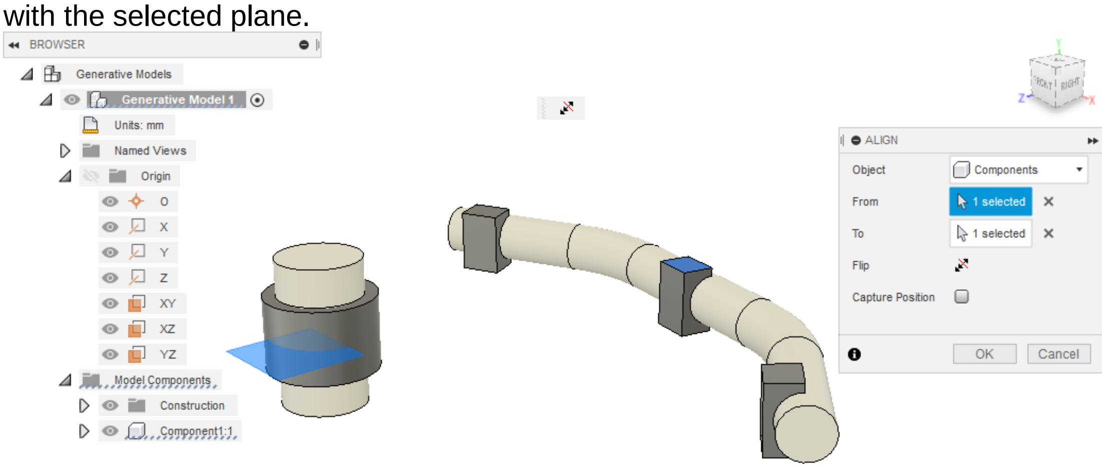


Now when selecting a direction in the Manufacturing dialog, the desired tool axis corresponds with the "Y" direction.

## Preserves: What's with those Precheck Warnings

2.5-axis frequently generates two scary, obscure-sounding pre-check warnings:

[12] We need the spatial relationships between all of our preserve/obstacle/starting shape bodies to be retained during the Align transformation; otherwise we re-align each body separately and lose the distance between them.

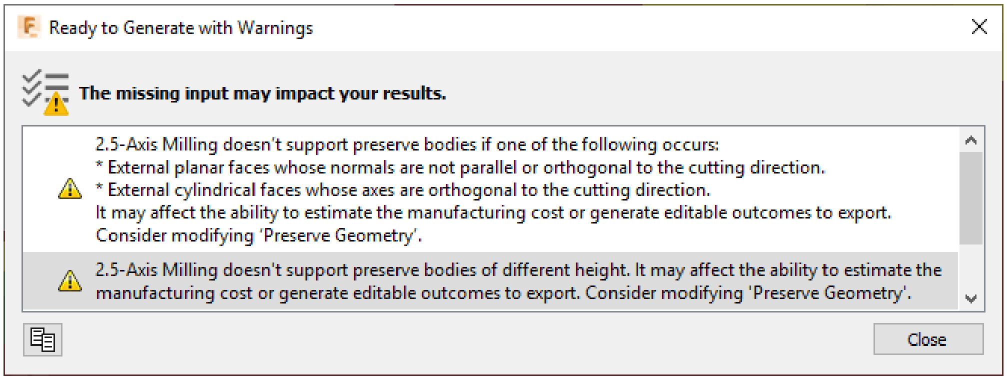


Pre-check warnings from a setup with 2.5-axis machining enabled

The only way to remove these warnings entirely is to provide preserves which could be all created with a single Extrude, like this:

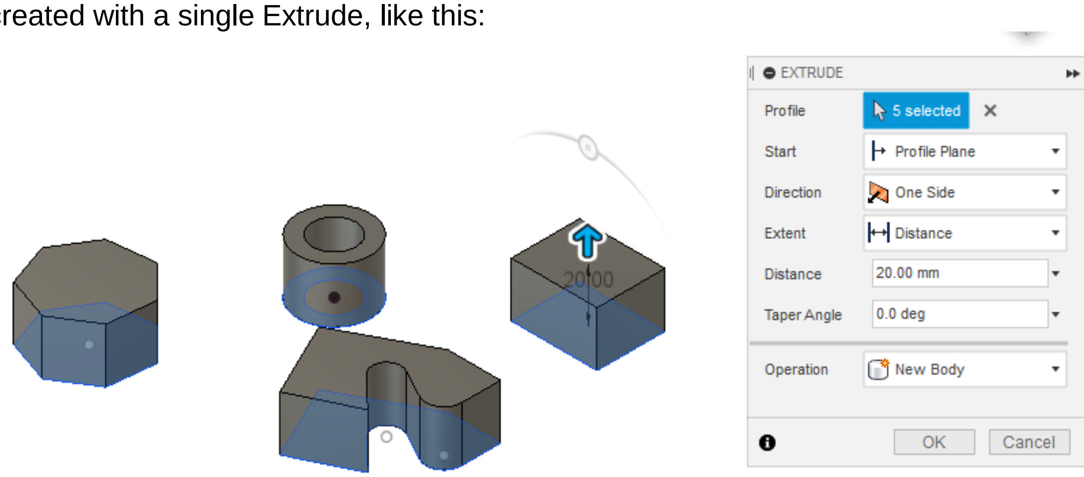


All these bodies can be used together in a 2.5-axis generative setup without creating pre-check warnings

But pre-check warnings are warnings, not errors, and it's entirely possible to use more complex geometries for obstacles and preserves without running into problems. In the walkthrough section we gave a "Safe Rule of Thumb" which will avoid the errors hinted at by the precheck warnings[13]. It bears repeating here:

**Safe Rule of Thumb:**
Preserves + Obstacles = Extrusions
Obstacles should not intersect Preserves

Here's what we mean: So long as preserves have corresponding obstacles so that the two together create an extruded shape, it should work fine (though the precheck warning is still present). A horizontal hole used to bolt your gadget to something? Fine. A

13 We're not guaranteeing that you won't hit errors using the Safe Rule of Thumb…we're still working hard to make Generative Design more robust and have plenty of work to do. Following this rule avoids situations where the system is *guaranteed to fail*.

shallower tab in a low stress region (with an obstacle above it to keep it shallow)? Totally cool. The intricate interface your coworker designed up to connect this part to another piece of the assembly? So long as you have an obstacle keeping material from being added on top of it, everything should work. Note that the obstacles can't intersect the preserves without creating unexpected results.

The figure below shows four example preserve + obstacle combinations which satisfy the safe rule of thumb.

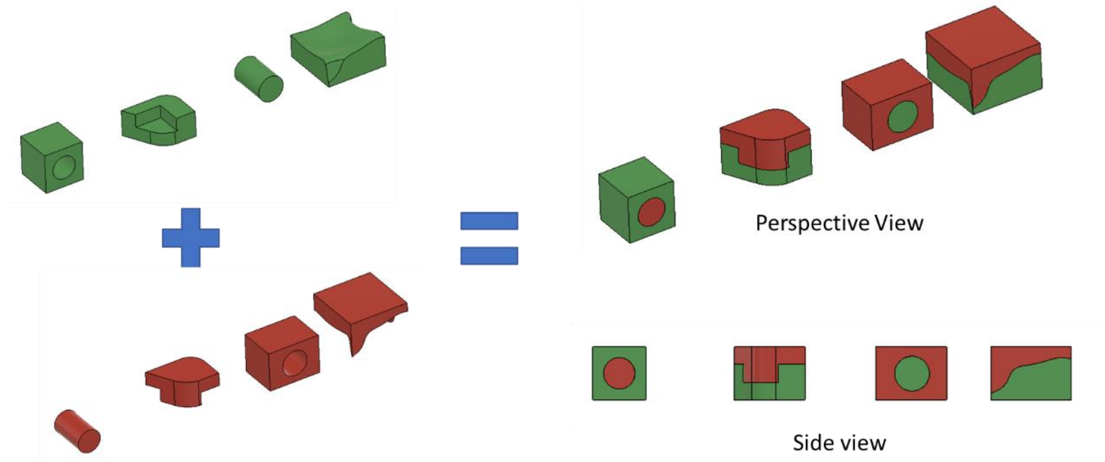


Valid preserves/obstacles for use in 2.5-axis machining constraints

There is also a helpful walkthrough video which explains a simplified form of this rule, available here: https://f360ap.autodesk.com/courses/generative-design-tips/lessons/preserve-geometry-for-2-axis-cutting-and-2-5-axis-milling

**Pushing the boundaries:** The Safe Rule of Thumb can still be broken without disastrous results but solves might fail and results may be unexpected. For the brave souls who dare, here's what happens:

- If a preserve's top and bottom surfaces both land on layer boundaries (and it's otherwise an extrusion), the solve, costing, and export all work as expected. Because each 2.5-axis solve produces three outcomes with different attack directions and layer locations, only one of these will avoid potential problems, however. It is better to create a complementary obstacle to make your intentions about the region above & below the preserve explicit, following the rule of thumb.
- If a preserve has non-extruded faces without a complementary obstacle, the generative solve will complete, but costing results will be inaccurate, validation simulations may not match the values predicted in the Explore environment, and the exported design will leave the interface between the generated geometry and the preserves empty – it will need to be filled in manually. In extreme cases, the design conversion can fail in this case.[14]

[14] We are working hard to improve this workflow and remove the restrictions on preserve geometries.

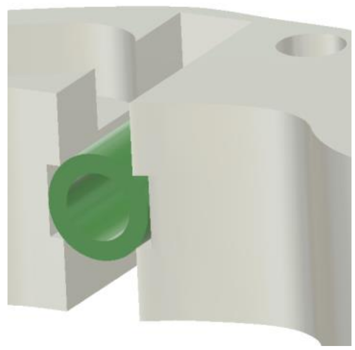

Example design outcome from a design showing the conversion result for a non-suggested preserve

### Obstacles: How to Cheat

In addition to being used in conjunction with obstacles (see previous tip), obstacles can be used to "cheat" the manufacturability rules for a 2.5-axis generative design.

Obstacles have none of the requirements placed on preserves for 2.5-axis machining. They can be any shape, can extend beyond the bounding box of the preserves, and do not generate pre-check warnings.

Obstacles are removed from the design domain after the manufacturing constraint is applied. If the manufacturable design extends into an obstacle region, the final shape is trimmed by the obstacle without further consideration of manufacturability.[15] This process occurs internally inside the generative engine at each iteration.

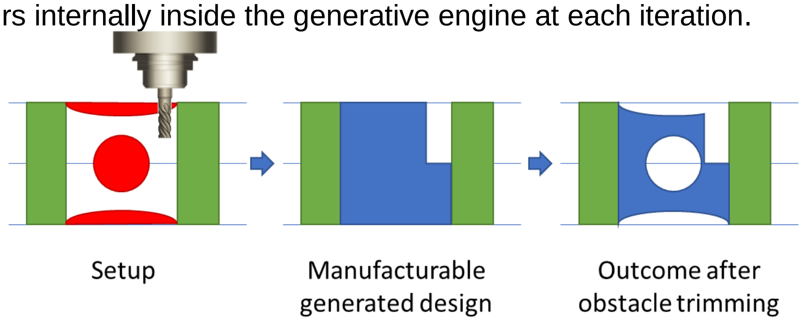


Illustration of non-2.5-axis-compatible obstacles being used to trim a manufacturable design to "cheat" and produce a non-2.5-axis output shape

This can be powerful in situations where selected a 2.5-axis machined part will be supplemented with a few 3- or 5- axis features for aesthetic or functional reasons.

[15] If you wish to ensure that the resulting design is strictly 2.5-axis machinable, please ensure that all obstacles are themselves 2.5-axis machinable.

## Starting Shapes: A Powerful Ally

The starting shape (or lack of starting shape) can have a dramatic impact on the outcomes produced by a generative study. Starting shapes can be used to drive design variation and are also helpful in achieving specific effects.

When a starting shape is provided, material is added to make the shape a layered 2.5-axis design. The shape is projected down to the bottom of the solid, then the intersection of each layer is extruded to fill out the layer, as shown in the figure below.

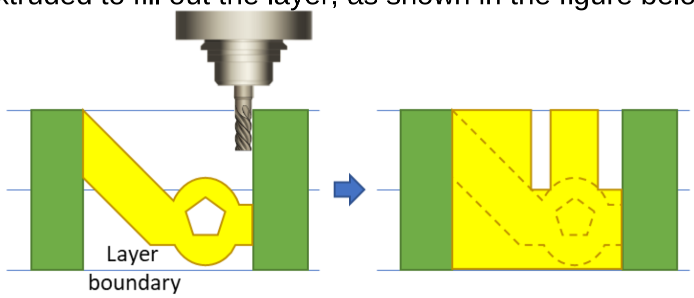


Starting shapes are projected to the top of the layer they intersect, and down to the bottom of the part.

The 2.5-axis generative solver will introduce holes as the solve progresses in order to encourage pocket-like designs and speed up the optimization process. This behavior can be avoided by providing a starting shape already containing holes which does not leave enough flat surface to make new hole generation possible. When providing a starting shape containing holes or pockets, be sure that they can be easily machined by the tool selected – twice the tool diameter is a good rule of thumb. Otherwise the solver may remove these features as unmanufacturable.

The figure below shows a variety of starting shapes and the corresponding design outcomes for a simple generative model.

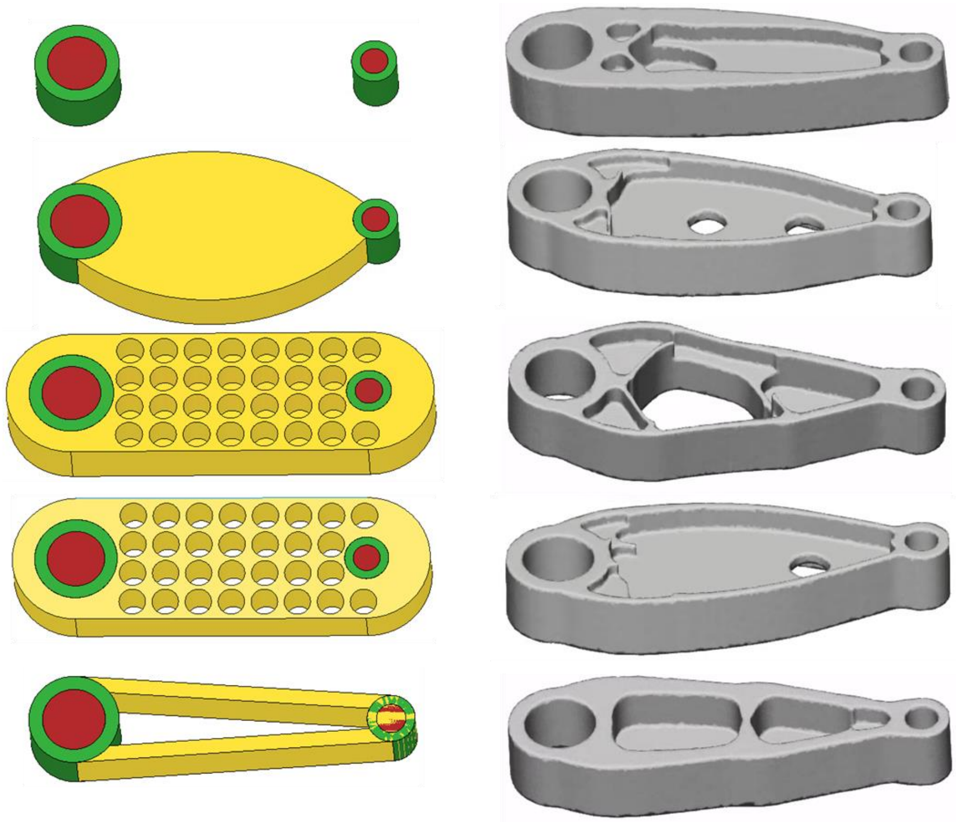

A variety of starting shapes and the corresponding converged outcomes; all for the same generative setup

Note how the 2.5-axis constraint is more sensitive to the starting shape than the unconstrained outcome. This occurs because the 2.5-axis constraint keeps the solver from removing material from the top or bottom faces of each layer, and the evolution of the profiles will follow a different path depending on the starting configuration.

Starting shapes can also be used to add material extending beyond the preserve geometries in the milling direction. Because 2.5-axis enforces all geometry to be contained in a layer, the part cannot “grow” beyond the top and bottom planes of the preserves, unless a starting shape hints to the solver that the layers should start or end at a different height. Note that this is very specific to 2.5-axis generative design because we select the layer locations at the start of the solve. All other manufacturing constraints can grow beyond the specified starting shape.

In the example below, the setup on the left can only generate a thin structure because the shape can’t grow above the top and bottom planes (milling direction is from the top). Adding a starting shape increases the space available for optimization and produces a more desirable solution. Note that this example ignores the Preserve Rule of Thumb because the preserve geometry does not have a complementary obstacle, so it fails the export conversion process.

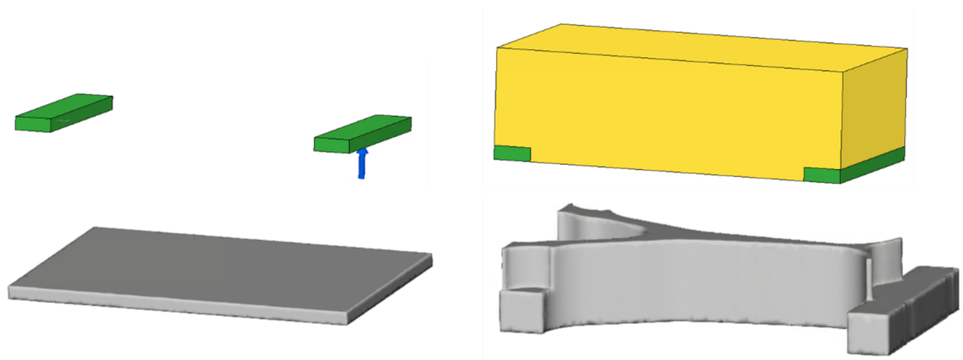

No starting shape
(not converged; FOS 0.17)

Tall starting shape
(converged)

Setups with and without starting shape and the affect on the usable domain

### Tool Accessibility: Making it Machinable

2.5-axis generative designs should be as easy-to-machine as possible. The manufacturing dialog gives the option to input a minimum tool size to be incorporated into the generative solver, ensuring that the part has no profiles which could not be cut by an endmill of this diameter.[16]

The tool accessibility function is implemented in such a way that it *adds material to the design* to ensure accessibility on each iteration (material is never removed to ensure accessibility). As a result, when the evolving design contacts an obstacle which contains radii smaller than the allowed tool size, these features are *not* corrected to make them accessible.

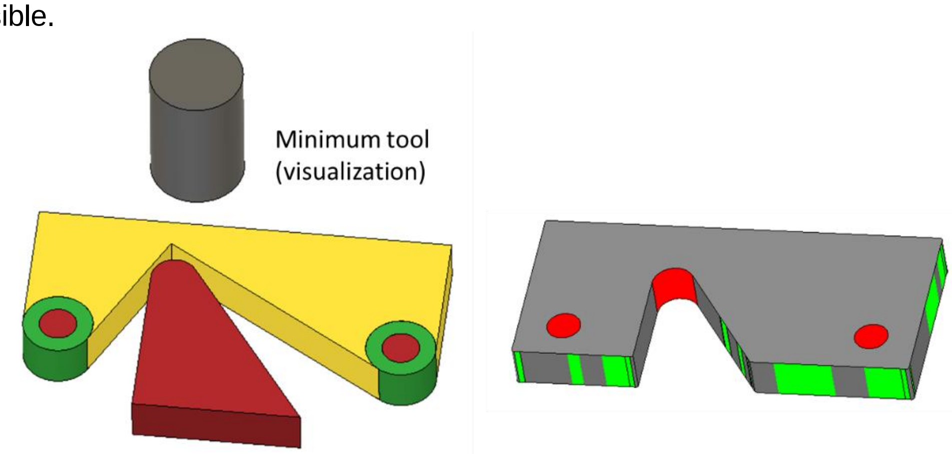


Generative setup

Minimum Radius Analysis on
outcome

Obstacles with radii which are not manufacturable may produce unmanufacturable results

[16] In accordance with machining best practices, the radii produced will typically be slightly larger than the specified tool diameter to reduce the tool contact surface while cutting the corner.

This creates the opportunity to "override" portions of the design where specific non-2.5-axis operations are planned or acceptable by providing an appropriate obstacle.

Typically, more than one tool will be used to cut the profile of any 2.5-axis part. As the depth of cut increases, larger endmills are preferred to reduce expense and increase accuracy. In 2.5-axis generative parts, the deeper the tool must cut into the bounding box of the part, the larger a radius is selected. A tool will never be smaller than the specified minimum tool diameter but may become larger if the part is comparatively tall.

The following figure illustrates this approach with a side view of a hypothetical 2.5-axis geometry. Note that the tool size is increased even if the whole flute length is not required to cut the face in question.

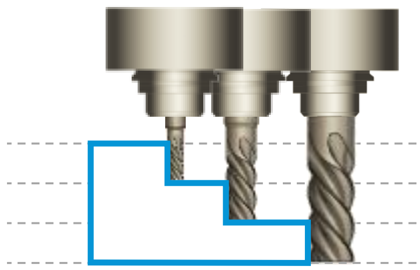

Schematic of automatic tool size increase as the depth into the part increases (side view).

An example of this effect is visible in the following result. Near the top surface, a small tool can be used and fine features are resolved, but as the depth of the cut increases, the profile becomes smoother and accessible to a larger tool.

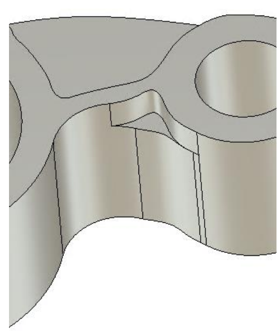

Features near the surface can have smaller radii and be more detailed because a finer tool can be used to machine them. Deeper features require larger radii (tool approaches from above).

### Solve Resolution

The resolution of a generative solve can be adjusted using the context menu of the Study item in the tree, as shown below.

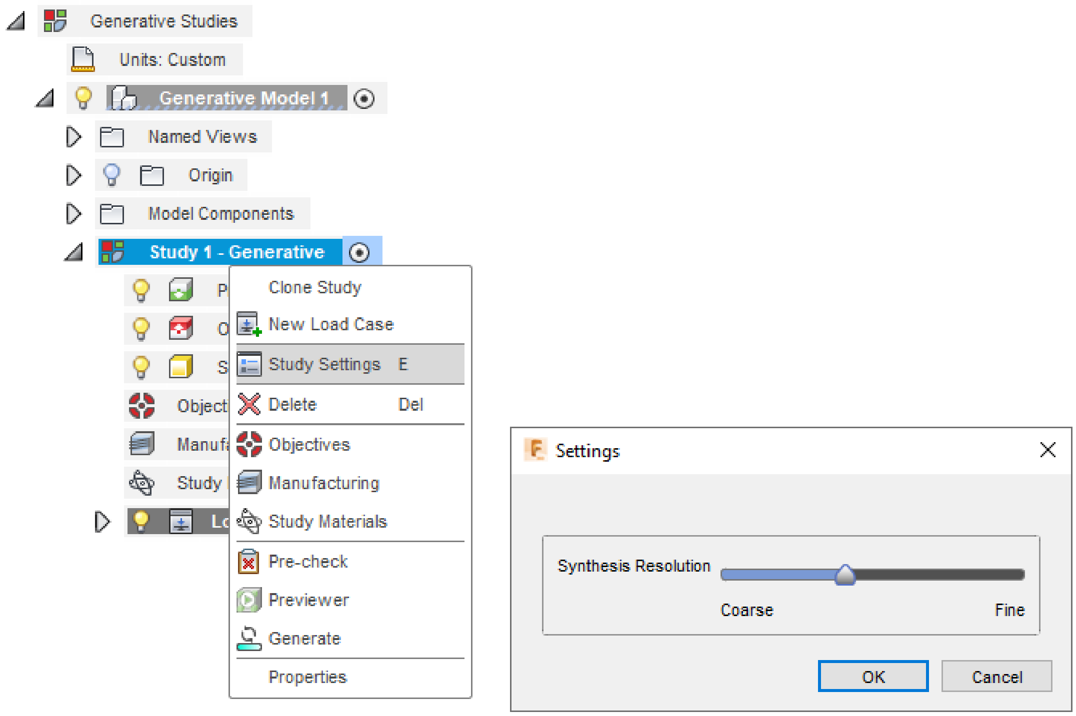


Accessing Study Settings to change Synthesis Resolution

Typically, coarser resolutions result in faster generation of outcomes but produce less detailed geometry, while finer resolutions take longer and produce more fine features in the results. For 2.5-axis synthesis, the effect is compounded – fine resolution takes longer to solve at each iteration and tends to progress more slowly from one iteration to the next. Because the geometries created by 2.5-axis manufacturing methods tend to not involve very fine features, selecting "fine" synthesis resolution is not recommended as it adds little value to the results and requires much more time to generate.

### Selecting Materials

Even if the final material is known ahead of time, getting an aesthetically-pleasing outcome from generative design can be challenging, requiring multiple generations with changes to loads, factor of safety, and mass target.

In our practice, we find it helpful to include several materials (up to 5 can be specified) in each solve in order to increase the chances of producing a viable or thought-provoking outcome. Typically, we select several materials with a range of yield strengths that bracket the material we plan to use. The 2.5-axis constraint does not treat materials differently than any other generative manufacturing constraint.

## Editing 2.5-axis Generative Outcomes

In this section we provide a few additional techniques, especially helpful in the Model-on-Top Approach, where the generated outcome geometry is modified by adding additional features to the timeline (instead of editing the sketches).

## Merging Faces

Generative design outcomes produce sketches which are extruded to form the various layers of the part. Sometimes, two edges which should be perfectly aligned are not quite recognized by Fusion 360, resulting in extra seams between (identical) faces. Fusion lets you delete these artifact faces to produce a smooth, continuous structure.

### Procedure 6. Merge redundant faces

1. Consider the nearest face of the model in the picture below. The three layer tiers are all exactly aligned, but Fusion still recognizes 3 faces.

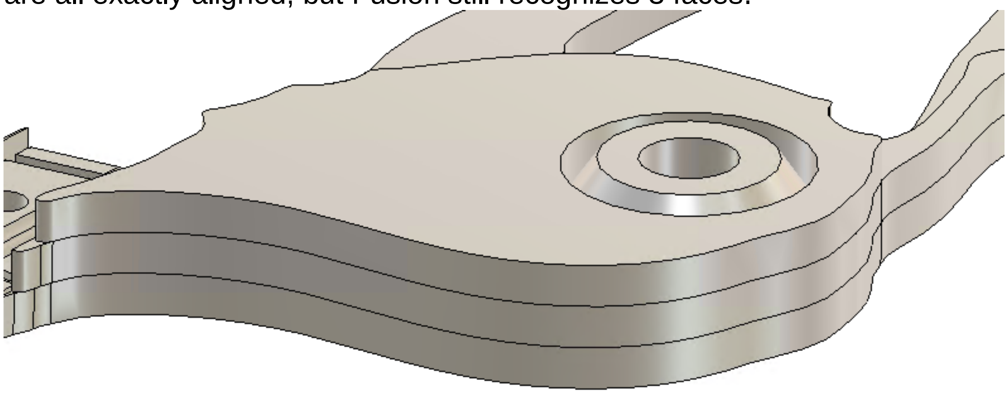

2. Select two of the faces and right-click, selecting the Delete option. Typically, I select the lower faces to delete; deleting the top face can shift sketch geometry near the end points.

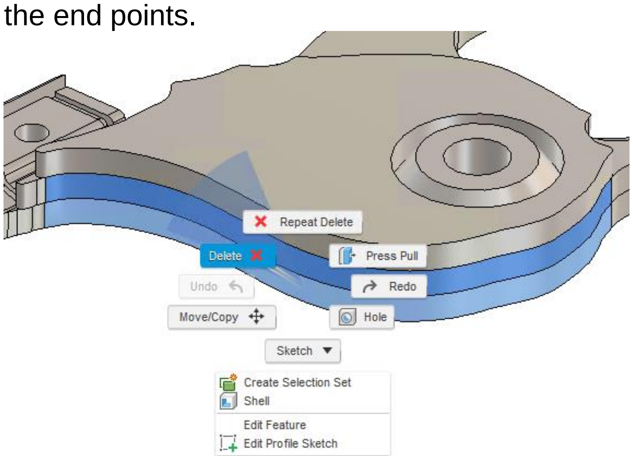


Fusion will remove the redundant edges and produce a single face:

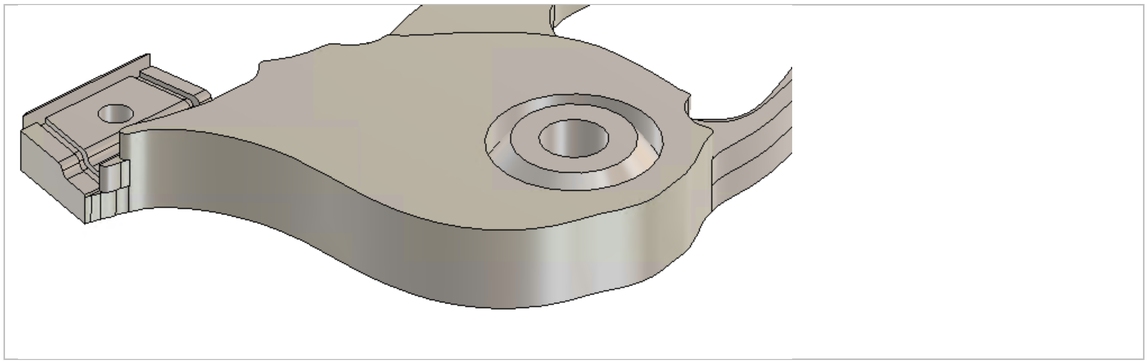

Sometimes this operation fails, either with an error or by producing unexpected and unwanted geometry. In that case, there may be a sliver face between the layers; see the next section.

### Horizontal Sliver Faces

"Sliver faces" occur occasionally in 2.5-axis generative results when two profiles are almost (but not quite) aligned. For example, in the figure below, the three lowest layers are all appear aligned coming up to the cylindrical preserve geometry, but closer examination reveals thin shelves. Sometimes these shelves even make faces inaccessible from the requested machining direction.

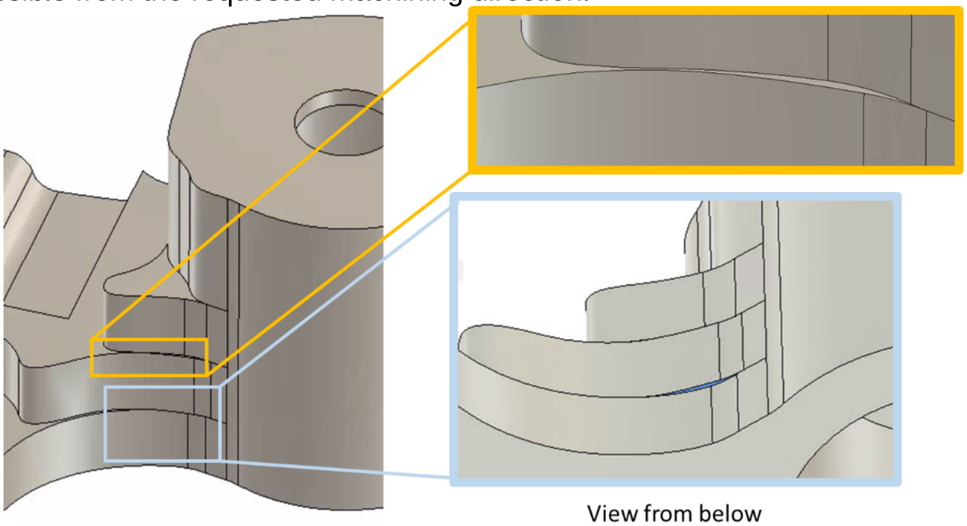


Examples of sliver faces

In this section we will show a few tricks for handling sliver faces.

**Trick #1** is to delete the sliver face itself by selecting it and picking Delete from the right-click menu. In this case, deleting the bottom sliver face removes both problems highlighted above.

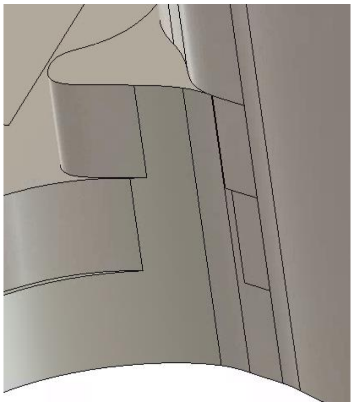

Result after deleting the lower sliver face

This trick fails a fairly high percentage of the time, but when it works, it often creates the cleanest outcome.

**Trick #2** is to extrude the problem face up or down to the next layer, either cutting away the sliver or joining it with the next layer's surface.

Procedure 7. Extruding sliver faces

| |
|---|
| 1. Select the sliver face (zooming in may be required)<br>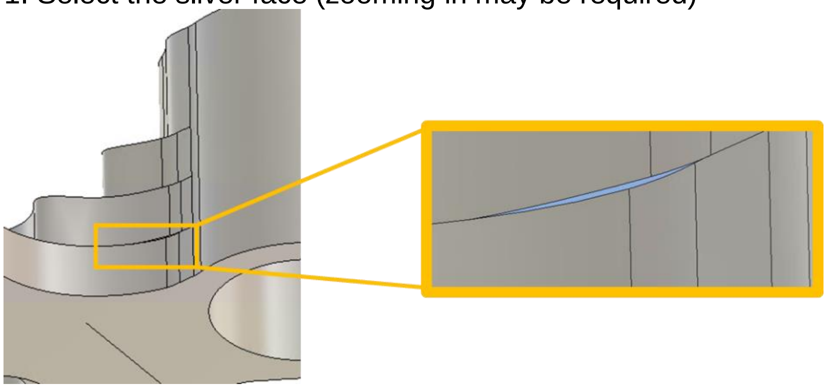 |
| 2. Add an Extrude feature by clicking the Extrude button in the Create drop-down. The profile will be pre-selected |

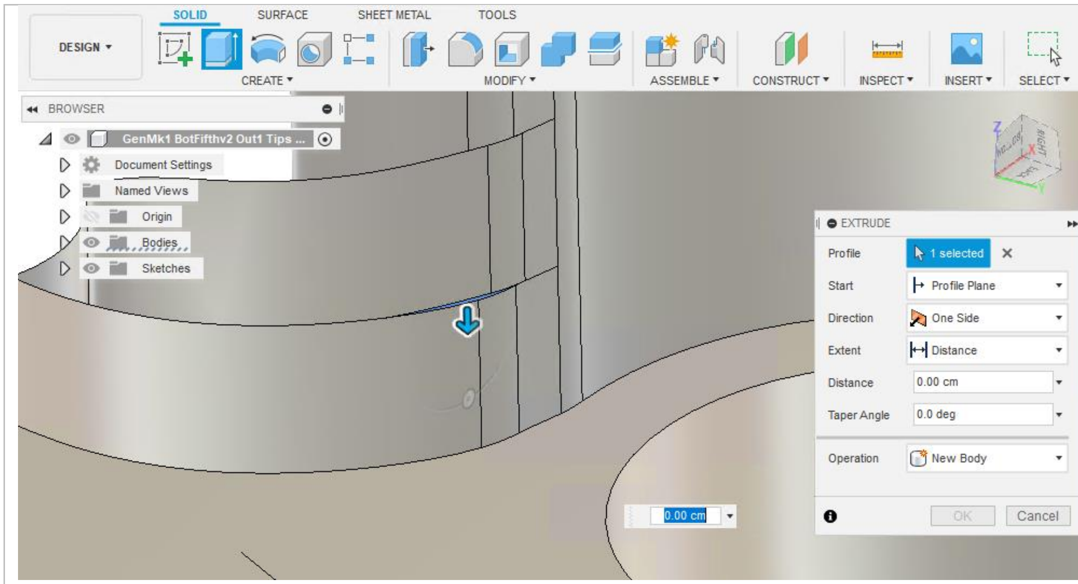


Change Extent to To Object and select the next-highest layer to join the sliver body with the existing part, removing the sliver face by slightly expanding the profile of the modified layer.

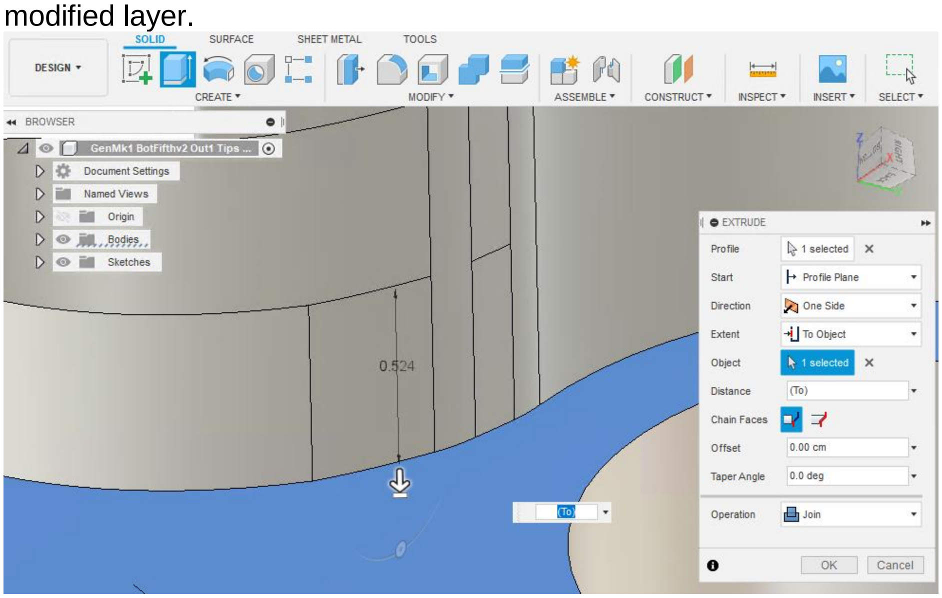


Click OK to finish the command.

Sometimes, the sliver is connected to a larger face profile, and it is not desirable to extrude the entire face. In this case, add a sketch on the layer surface with a line or curve which breaks the layer into two closed profiles, as shown below.

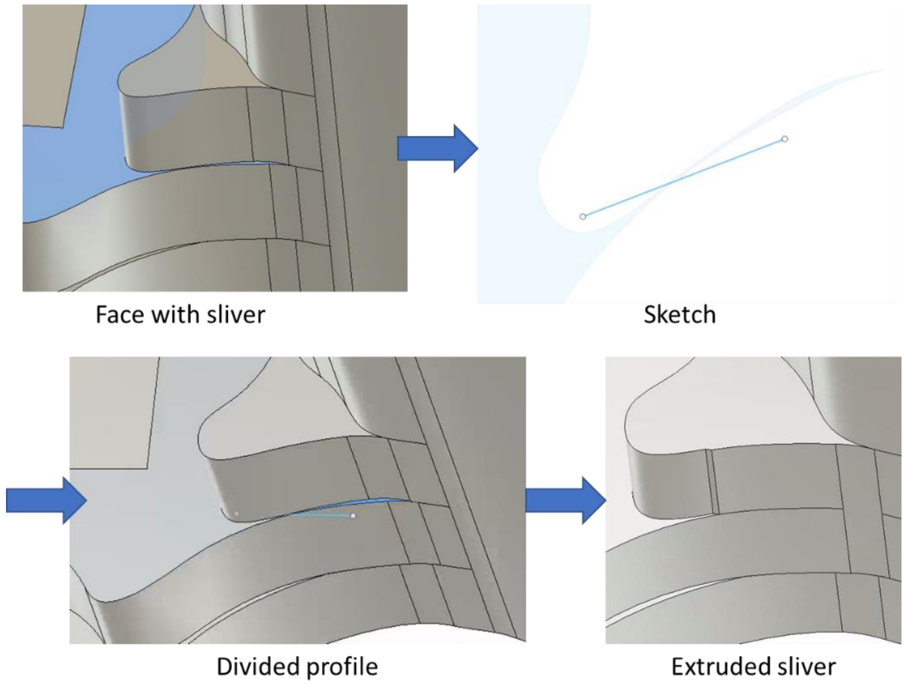


Extruding a sliver when the sliver is not isolated from the larger face surface

An alternative approach is to use an Extrude/Cut instead of an Extrude/Join feature, cutting the sliver face down to the far side of the part. In the example below, an Extrude/Cut from the bottom face resolves both slivers with one feature.

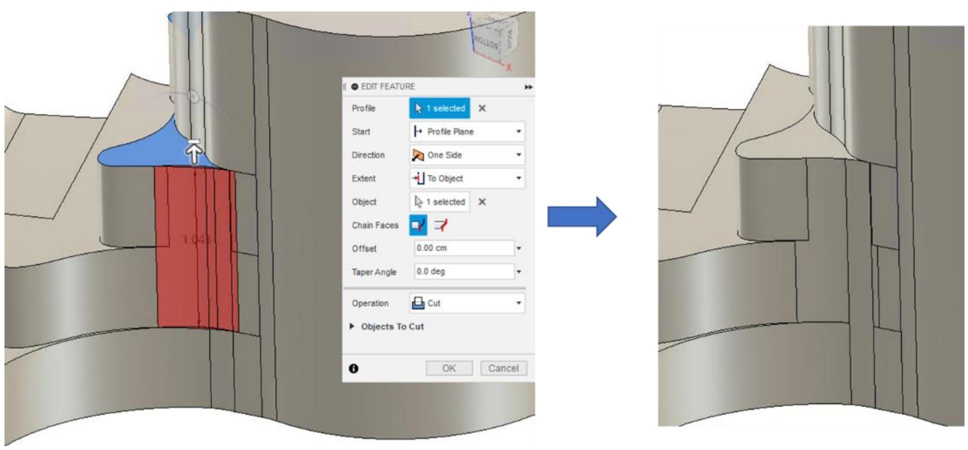


Extrude/Cut feature to remove a sliver face.

Note that it is typically required to add fillet features after using extrusions to correct sliver faces, as the extrusions often leave very slight sharp corners in adjacent layer profiles.

## Vertical Sliver Faces

In other situations, sliver faces appear vertically in the model, frequently near the transitions between preserve geometry and generated geometry.

**Trick #1:** Delete the faces in the same manner as described above. Select the face and choose Delete from the right-click menu. Fusion will make its best attempt to join the adjacent faces cleanly, often in a way that makes intuitive sense.

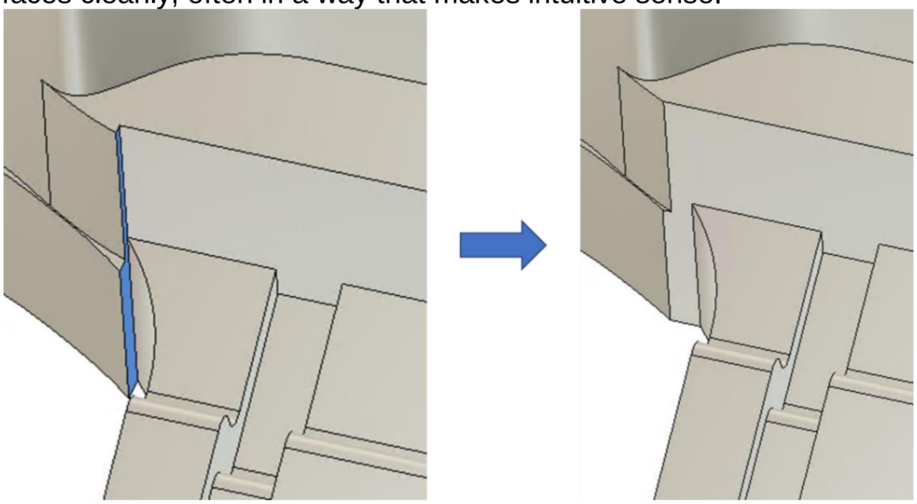

Vertical sliver faces removed by deleting the face

**Trick #2:** Try adding a Fillet feature and selecting the concave vertical edges on adjacent to the face.

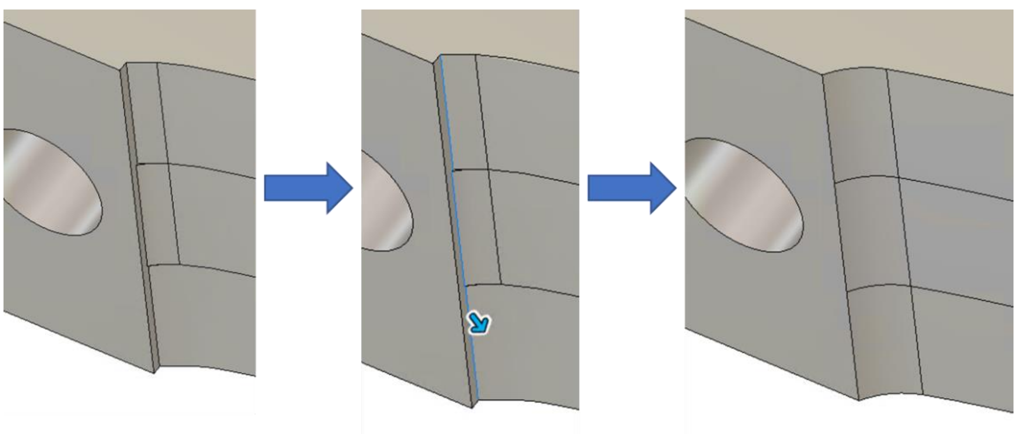

Adding a fillet to remove a vertical sliver face

## Fixing Open Contours

In our experience, the sketch contours of 2.5-axis results sometimes fail to form closed contours which can be extruded. We're working hard on improving the robustness of these sketches and have some great tools for analyzing sketches coming in Fusion. For now, though, here are a few tricks we have found helpful.

**Trick #1**: Look for “free” sketch points which are not part of a closed contour, or “stacked” sketch points which may be hiding a free sketch point (see figure below). These appear as open circles in the sketch view and indicate there may be a tiny gap in the sketch that needs to be closed.

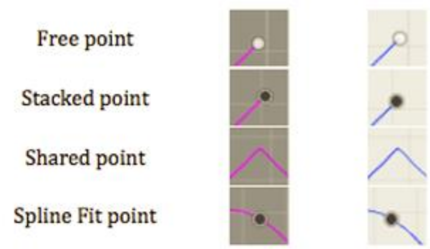


Symbols for points used in sketches

As an example, the stacked sketch point in the figure below, when zoomed in, reveals a free sketch point and a hole in the sketch contour.

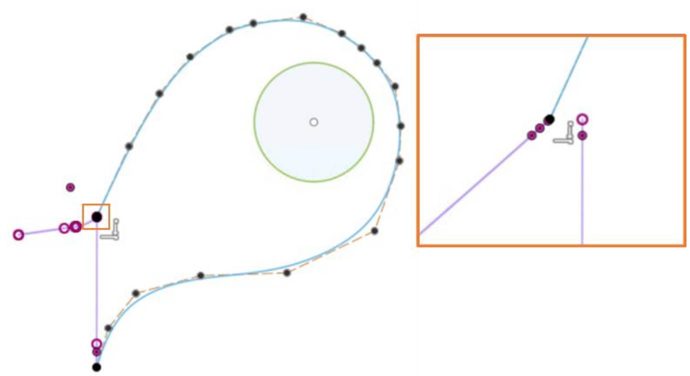

Stacked sketch point reveals an open contour when zoomed in

**Trick #2**: Use a line across your sketch to narrow in on the problem area. When a contour is closed it produces a blue shaded region, and the line can immediately tell you which side of the sketch is closed and which contains the problem. See the figure below, which is the same sketch as shown previously, with an opening on the left side. By moving the line back and forth, the problem vertex is easy to identify.

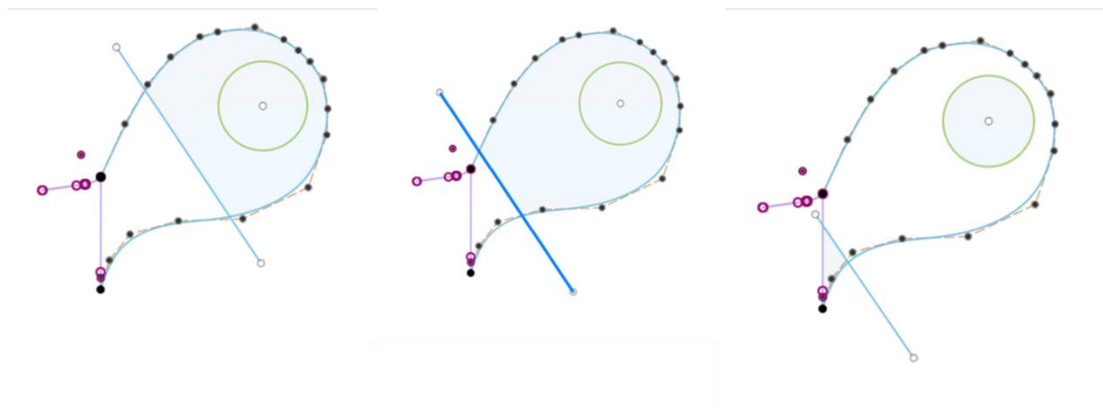

Open sketch debugging by adding a line around the sketch. In three moves we identify the problem vertex and can investigate.

**Trick #3:** Sometimes, it isn't a matter of creating a closed contour, but of the contour not all laying on the same plane. Even though it shouldn't happen, we have occasionally noticed 2.5-axis generative outcomes which, after editing, produce slightly non-planar sketches which will not be recognized as closed contours. We're still working on why this is occurring, but the fix is easy. Select all of your sketch geometry, right-click, and choose "Move to sketch plane". If this item is not visible, then the sketch is already coplanar. Sometimes it takes two tries before the sketch comes back into line.

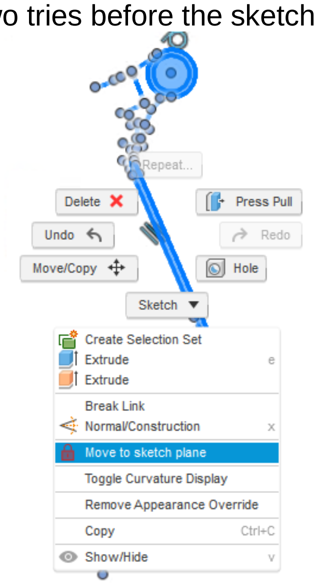


The "Move to sketch plane" command appears only when the sketch is not already coplanar