

# Making Souls: Methods and a Pipeline for Volumetric Characters

Patrick Coleman  
pcoleman@pixar.com  
Pixar Animation Studios

Laura Murphy  
lmurphy@pixar.com  
Pixar Animation Studios

Markus Kranzler  
bailong@pixar.com  
Pixar Animation Studios

Max Gilbert  
mgilbert@pixar.com  
Pixar Animation Studios




Figure 1: Joe, 22, and three additional souls. ©Disney/Pixar.

## ABSTRACT

The soul characters in Disney/Pixar’s *Soul* have a stylized appearance that sets them into a unique world, which introduced many new challenges. Everyone from the art department and character modelers and shaders to the technical directors and developers in the effects, lighting, and software groups collaborated to bring this new visual style to screen. The soul world is abstract and ethereal; this needed to be balanced with visual clarity and design appeal.

As our character rigging and animation tools use rigged surfaces, a key challenge was presenting a new representation derived from this data that meets our visual goals. To achieve softness of volumetric form and dynamically changing linework, we built a system to procedurally generate this data in Houdini. Significant numerical computation was required to create this data at the fidelity required. We developed an automated system for managing this computation in a configurable way, while keeping data for downstream renders in sync with changes to character performances.

### ACM Reference Format:

Patrick Coleman, Laura Murphy, Markus Kranzler, and Max Gilbert. 2020. Making Souls: Methods and a Pipeline for Volumetric Characters. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH '20 Talks)*, August 17, 2020. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3388767.3407361>

## 1 LOOK DEVELOPMENT

The soul characters are conceptual beings of light that have a prismatic response to lighting (Figure 2, left). A challenge was creating a look that balances between soft/ethereal and enough definition to make shapes and expressions readable. It was important for the characters to not appear as hard surfaces with subsurface scattering.

We wrote a stylized shading node (Iridescence Driver—or IDriver) to feed into the Double Henyey-Greenstein volume shader to achieve

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*SIGGRAPH '20 Talks, August 17, 2020, Virtual Event, USA*

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7971-7/20/08.

<https://doi.org/10.1145/3388767.3407361>




Figure 2: Concept art, rendered, material. ©Disney/Pixar.




Figure 3: Projected normal fields: body gradient, projected gradient, final blend. ©Disney/Pixar.

this non-photorealistic lighting response (Figure 2, right). In essence, the IDriver determines the albedo color by using the normals to look up into a color ramp and faking diffuse and specular lobes. To make the characters look extra soft we use softer normals to filter out some smaller details outside the face. (Figure 3).

For smaller and more complex shapes in faces, we modeled and rigged animatable lines to accentuate contours or wrinkles. (Figure 4, left and center). The rig also outputs modeled detail lines for hair and clothing. We generate lines around the silhouette of the hands to enhance their readability and clarify the pose of the character. The silhouette lines are extracted by tracing the zero iso-contour of the function defined by the dot product between each mesh vertex normals and the vector pointing from that mesh vertex to the camera position. Animators also had control over the visibility, density and size of these lines. (Figure 4, right).

Joe has accessories that manifest in the Soul world. His hat is a major accessory that he feels is a part of him, so it looks similar to his body. His glasses, however, were necessary but lifeless, therefore they are colorless. Their sharp outline allowed us to keep them more transparent, but still readable. Similarly, the background souls have




Figure 4: Rigged lines that respond to face poses (left, center) and silhouette lines used to clarify pose (right) ©Disney/Pixar.

a look similar to Joe’s hat for their hair; major accessories are like a desaturated version of the hat (but still rainbowy); and minor accessories are like Joe’s glasses. Additional portions of the soul look development process are described in [Fong et al. 2020].

## 2 BUILDING VOLUMES

We use Houdini to generate per-frame volumes and volumetric signals that feed the IDriver shading node. To make sure the animation reads well with semi-transparent volumes, most of our workflows are directed towards camera. Each character is made of several individual volumes; most of these are partially or completely culled by self occlusion. This creates a cleaner read with no layered shapes. For example, the teeth and tongue are precisely cut to camera to match the open mouth shape. The eyes are surface geometry, not a volume, so we generate a masking texture instead of cutting them.

Every shading signal (e.g. eye shadow, cheeks, lips, etc.) also needs to be defined as a volume; and again, many are projected from the camera. We use these to adjust the color and shader lobe properties in the IDriver and shading network. We build the normals volume by projecting the surface normals through the character from camera. For the soft normals volume, we convert the surface to an SDF, dilate it, apply an iterative mean value blur, and convert it back to a surface to project the normals of this smoothed geometry.

For all the background and secondary souls, we generalize the houdini network to work for all of their designs (Figure 1, two right characters). They only need a handful of parameters for specific control. While each soul is modeled as their own unique shape, most share the same topology, which greatly streamlines the task. The same network also creates rest pose volumes for our soul crowd pipeline that uses volumetric deformation [Ouellet et al. 2020].

Volumes for facial linework (Figure 4, left and center) are generated by extruding the rigged flat meshes away from camera to form the volume. For silhouette lines that are procedurally generated, we create ribbons perpendicular to camera along the curves before we extrude those also away from camera (Figure 4, right).

## 3 PIPELINE

Our Houdini networks generate multiple passes of high fidelity volume data. This is computationally intensive and needs to happen regularly; we designed a system called *hexport* (houdini export) that automatically generates the data by parallelizing it on our renderfarm. The system is highly tunable, giving departments different parallelization and scheduling behavior based on turnaround needs.

Early in production, character artists used our existing effects pipeline. Effects artists also work with computationally expensive volumes, and we have developed tools for parallelization and large




Figure 5: Soul character data pipeline. ©Disney/Pixar.

data storage. Other departments export animated USD assets in a single file with a time sample per frame, but this does not allow for immediate parallelization, as a race condition exists on file writes. Instead, our effects artists export a file per frame and use the utility *usdstitchclips* to assemble per frame data into *USD Value Clips*. On a USD stage, this appears identical to time samples in a single file<sup>1</sup>.

We used this for early look development, but this did not scale to character assets that are shared among multiple users and departments. We developed a versioning system that includes multiple unpublished working directories, which we store in high performance storage for parallel renderfarm use. Publishing a working directory for use in downstream departments involves only a symlink and database change. In addition, we moved from storing houdini network data in shot level hip files to character asset HDA (houdini digital asset) files that are shared among relevant shots. When exporting a shot, the hexport system creates a shot level hip file and references the current character HDAs into it; these HDA subnetworks read shot-specific cached surface animation, convert it to volumes and linework, and export results to disk (Figure 5).

An important part of the hexport pipeline is partitioning and scheduling the Houdini network computation. The effects department parallelizes as much as possible, but the farm scheduler does not scale to a task for every character on every frame. Most computation in our networks is not history dependent; this increases partitioning and scheduling flexibility. When exporting a shot, it is initially split into subsets of characters, and the network for each character is partitioned into three classes of compute tasks: expensive but parallelizable across frames, tasks for aggregating data from prior tasks, and independent single process tasks. Portions of each character’s network computation are assigned to appropriate tasks.

Hexport also schedules time-dependent character simulations. We assign these to the independent single-process task and rely on houdini for thread parallelization. We tune thread allocation and resources for each task type based on turnaround needs and experimentally measured data. As we create tens of terabytes of high fidelity volume data per night that downstream renders read in parallel, we use a cached filesystem designed for heavy write loads and many parallel reads. To avoid copying data, we use symlink labels to indicate whether data is IP or published for rendering.

**Acknowledgements :** We thank George Gnguyen, Junyi Ling, and Ling Tu for character development, Matt Kuruc and Elliott Cattell for pipeline development, and Boris Dalstein, Tolga Goktekin, and Fernando de Goes for silhouette software development.

## REFERENCES

- J. Fong, J. Jarvers, M. Kranzler, and R. Michero. 2020. Finding the Look of Souls. In *SIGGRAPH 2020 Talks (SIGGRAPH '20)*. ACM, New York, NY, USA, 2.
- S. Ouellet, D. Garcia, S. Gustafson, M. Kuruc, J. Ling, M. Lorenzen, G. Nguyen, G. Gilbert, and F. Eder. 2020. Rasterizing Volumes and Surfaces for Crowds in *Soul*. In *SIGGRAPH 2020 Talks (SIGGRAPH '20)*. ACM, New York, NY, USA, 2.

<sup>1</sup>[graphics.pixar.com/usd/docs/USD-Glossary.html#USDGlossary-ValueClips](https://graphics.pixar.com/usd/docs/USD-Glossary.html#USDGlossary-ValueClips)