Making the *Stream of Consciousness* in Pixar's *Inside Out 2*

TOLGA GOKTEKIN, KRZYSZTOF ROST, RAVINDRA DWIVEDI, and FERNANDO DE GOES, Pixar Animation Studios, USA

In Pixar's feature animation *Inside Out 2* (2024), the core Emotions led by Joy find themselves navigating through the Stream of Consciousness in an attempt to reach Riley's Back of the Mind and then retrieve her Sense of Self. Our team was cast to develop the look for the Stream of Consciousness and support its deployment across nearly twenty shots. In particular, it was clear during production that we needed ways to work concurrently with the Animation team, at times even independently, while still able to address notes and changes rapidly. This work discusses our technical solutions for the implementation of the Stream of Consciousness, including in-house procedural tools that facilitated the authoring and stylization of velocity fields interacting with 3D obstacles.

1 Motivation

Finding the look for the Stream of Consciousness in Pixar's *Inside Out 2* was a long journey, mainly because the part it played in the story changed drastically over time. Originally, the stream was portrayed as a body of water filled with glowing bulbs, and the Emotion characters used a boat to travel across it. However, when experimenting with physically based fluid simulation techniques, the recurrent note we received was to make the stream behave less like water and more like a calm substance moving the objects. Ultimately, the look of the Stream of Consciousness settled into three stages presented over the course of the film. First, the stream is a calm substance carrying Riley's favorite food (see upper-left image). Later, a choppier flow is introduced, carrying food objects that Riley hates. Finally, the stream develops into dangerous rapids that form a waterfall at the edge of the "Sarchasm" (see upper-right image). As story development continued, we also needed to take a more flexible approach, allowing the individual objects, their distributions, and input animations to be finalized later in the process. To hit specific look guidelines while reacting to changes quickly, we ended up implementing the Stream of Consciousness fully procedurally by mixing in-house and native Houdini tools, as described next.

2 Designing a Procedural Stream

Our team structured the Stream of Consciousness using five representative layers flowing downstream. At the bottom, there was a volume element that acts as an overall murk. A set of aeration bubbles displayed as flat discs were also added siting on the murk near the shoreline. The main substance was composed by oval shapes arranged into three sub-layers with different depths floating above the bubbles. The rate of the flow and more nuanced details were then depicted by streaming ribbons that laid on top of the substance sub-layers. At last, each floating object and the shore

Authors' Contact Information: Tolga Goktekin, goktekin@pixar.com; Krzysztof Rost, krost@pixar.com; Ravindra Dwivedi, ravindra@pixar.com; Fernando de Goes, fernando@pixar.com, Pixar Animation Studios, USA.

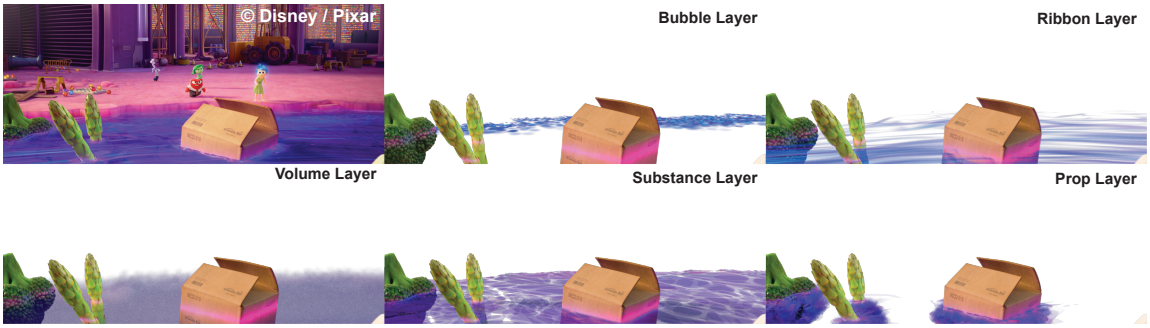

itself were attached to high resolution versions of the substance shapes, resembling their interaction with the main stream. The top figure includes the breakdown of a rendered frame into these different layers.

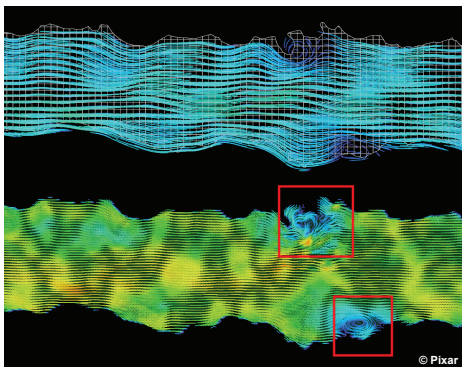

Another important component was controlling the dynamics of each stream layer per shot. We started by generating a baseline for the bulk motion using a two-step procedure. We first computed a 2D version of the bulk velocity by interpolating the tangents to the shoreline via an in-house mesh-based vector field solver [de Goes et al. 2020]. We further refined the velocity interpolation by setting soft constraints along curves drawn inside the river. We then extrapolated this 2D flow field into a base 3D flow field analytically by superposing grab-like Kelvinlets [de Goes and James 2017] instantiated at sample points scattered over the 2D river with a scale that provided full coverage with minimal overlap (top inset). Finally, we modified this base 3D field into distinct flows customized for the advection of each individual layer. For instance, we advected bubble discs along a blend of the base flow, curl noise, and a push field moving away from the shore, which we generated by extracting the gradient part of the shoreline's signed distance function. In the case of the ribbon layer, we manipulated the 3D velocity using a low-frequency noise combined with push fields moving away from the shore and all colliding objects. In general, we found that incorporating these 3D push fields to flow around colliders was more stable compared to constraining the baseline flow with the boundary of the moving objects. We also added vortices in the flow for extra detail as we approached the rapids part. To do so, we flipped the tangents to the shoreline in areas of interest in order to reverse the interpolated 2D flow, and then inserted twist Kelvinlets for point-wise turbulence (bottom inset). To avoid bunching up as the stream density grew, we performed a pressure projection pass on each modified flow to force it to be incompressible. For closeup shots where we needed to see the ribbon layer interacting with the colliders more aesthetically, we replaced the trailing computation with a collision-aware particle simulation. We point the reader to the supplemental video for final renders of the Stream of Consciousness produced by our procedural approach.

References

Fernando de Goes, Andrew Butts, and Mathieu Desbrun. 2020. Discrete differential operators on polygonal meshes. *ACM Trans. Graph.* 39, 4, Article 110 (2020), 14 pages. doi:10.1145/3386569.3392389

Fernando de Goes and Doug L. James. 2017. Regularized kelvinlets: sculpting brushes based on fundamental solutions of elasticity. *ACM Trans. Graph.* 36, 4, Article 40 (2017), 11 pages. doi:10.1145/3072959.3073595