Eurographics Workshop on Visual Computing for Biology and Medicine (2023)
C. Hansen, J. Procter, and R. G. Raidou (Editors)

Bio-Sketch: A new medium for interactive storytelling illustrated by the phenomenon of infection

P. Olivier[1], R. Chabrier[1], P. Memari[1,2], J-L. Coll[3], M-P. Cani[1]

[1]LIX, Ecole Polytechnique, CNRS, IP Paris, Palaiseau, France
[2] INRIA Saclay, Palaiseau, France
[3] Univ. Grenoble Alpes, INSERM, CNRS, Institute for Advanced Biosciences, Grenoble, France.

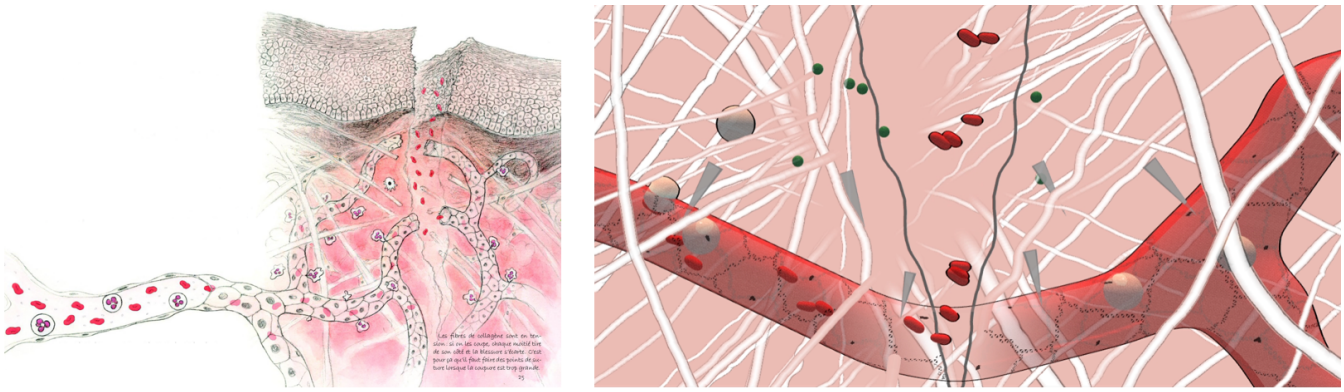

Figure 1: *Immunology illustration from the magazine Globule [Cha18], © Renaud Chabrier (left) served as inspiration for an animated 3D sketch created using our system (right).*

Abstract

In the field of biology, digital illustrations play a crucial role in conveying complex phenomena, allowing for idealized shapes and motion, in contrast to data visualization. In the absence of suitable media, scientists often rely on oversimplified 2D figures or have to call in professional artists to create better illustrations, which can be limiting. We introduce Bio-Sketch, a novel progressive sketching system designed to ease the creation of animated illustrations, as exemplified here in the context of the infection phenomenon. Our solution relies on a new progressive sketching paradigm that seamlessly combines 3D modeling and pattern-based shape distribution to create background volume and temporal animation control. The elements created can be assembled into a complex scenario, enabling narrative design experiments for educational applications in biology. Our results and first feedback from experts in illustration and biology demonstrate the potential of Bio-Sketch to assist communication on the infection phenomenon, helping to bridge the gap between expert and non-expert audiences.

CCS Concepts

• ***Computing methodologies*** → *Computer graphics;* • ***Human-centered computing*** → *Interaction design;*

1. Introduction

By definition, a biological phenomenon is a sequence of events or chemical reactions leading to a transformation. Consequently, biologists study and seek to convey how and why an event occurs and its consequences on the environment. The visualization of biological data has been intensively studied from the molecular level to the mesoscale (interior of cells), with impressive progress in improving the saliency of the objects of interest. However, such 2D or 3D imagery can be complicated to understand and control, as they only capture the motion that occurred in the actual data. These images are thus typically complemented by simpler representations, from simple 2D sketches to 3D animations. However, while mere

DOI: 10.2312/vcbm.20231214

2D sketches may not be sufficient to capture and understand the phenomenon accurately, producing 3D and animated illustrations requires the help of professional artists, leading to lengthy discussions and trials and errors until the scientist's mental vision is conveyed. Biologists are therefore looking for intuitive creation tools to enrich the representation of their objects of study, to visually express their hypotheses, and/or to communicate in educational contexts.

Although existing illustration tools are already helpful for scientific dissemination and education, most of them do not ease the sketching of volumetric environments, where interactions between 3D shapes and changes of state need to be represented. The few methods targeting illustration at the cellular level have focused on modeling vascular networks in 3D using sketching conventions and on blood circulation to represent vessel pathologies. Notably, none of the existing solutions provides a complete system for creating and animating a complex environment inside which cells can navigate freely and evolve dynamically.

In this work, we target a new progressive sketching paradigm for modeling and animation designed to meet biologists' needs to communicate and explore dynamic phenomena. We collaborated with a biologist and a professional artist, both co-authors of this paper, to conduct a pre-study on depiction and narration in biology. This study confirmed that biologists need not only to visualize information from microscopic images but also to depict their understanding of dynamic phenomena in a simplified way. Our solution builds on their expressed needs and focuses on communicating the infection phenomenon.

Our technical contributions are threefold:

- We propose a new sketch-based modeling solution for pattern-based shape distribution in 3D, enabling the modeling of environments such as biological tissues with dense fiber networks in which cells evolve.
- We rely on standard schematic conventions in biology to complement an existing sketch-based modeling system and enable sketch-based specification of motion and deformation of the generated 3D shapes.
- We provide an interface for specifying and controlling temporal narrative scenarios, with a series of events triggered over time.

These three contributions were tested on a complex scenario, depicted in Fig. 1, designed by our expert co-authors. It expresses the human response to infection at the cell level. This provides a proof of concept and an example of the use of our system in a biological application context.

2. Related work

This work is related to sketch-based 3D modeling, example-based shape distributions, sketch-based animation methods, and biological visualization and storytelling. We focus on methods best suited to creating and animating biological phenomena. We refer the reader to [VGH*05, Ise15, LVPI18] for visualization techniques targeting illustrative rendering in biology and to [PS18] for a survey on the virtual representations used for teaching anatomy.

Sketch-based 3D modeling Sketch-based modeling methods have been used in various ways, from taking complete sketches and annotations as input to gradually building models through progressive, interactive sketching. However, only a few of these methods have been applied to represent biological environments. A promising approach to sketching biological shapes is using implicit surfaces, which leverage the organic properties of these shapes. Methods such as Matisse [BPCB08], later improved in Scalis [ZBQC13], proposed a general solution to the progressive sketching of organic-looking shapes. More specific solutions have been explored to represent biological structures from musculoskeletal structures [ABL*21] to vascular systems [PCP10, SBP*16, SSPOJ16]. While following sketching conventions and aimed at supporting anatomy teaching, the method of Pihuit et al. [PCP10] is limited to a single sketch input representing a network of vessels. Closer to our modeling objective is the work by Saalfeld et al. [SBP*16, SSPOJ16] presenting interactive systems aiming at creating and editing vascular systems to explain vascular pathologies. Their first method [SBP*16] is restricted to 2D representation and relies on fluid simulation for the blood flow. Their second model [SSPOJ16] is semi-immersive but requires significant user intervention to define blood flow and is limited to adjusting the weights of the Metaballs composing the vessel.

Example-based distribution synthesis Although extensive research has long been carried out on example-based texture synthesis [WLKT09], only a few methods tackled the synthesis of discrete vector textures, which are the most relevant to creating shape distributions in a 3D environment representing human tissues, from interactive sketches.

Among vector texture synthesis methods [BBT*06, IMIM08, HLT*09, MWT11], Landes et al. [LGH13] addressed shape distributions by simplifying the input shapes into proxy geometries and relying on spatial relationships such as distances and relative orientations between them. This method, which is not real-time, is also limited to bounded shapes, which would make the creation of fiber networks difficult. In contrast, Roveri et al. [ROM*15] present an example-based distribution synthesis method for both bounded and unbounded shapes. Shapes are decomposed into point samples encoded in a functional representation. Synthesis is achieved through neighborhood matching and energy optimization. Like other neighborhood-based texture synthesis methods, their approach requires input patterns with a sufficient amount of repetitions. In contrast, Olivier et al. [OMC22] take as input a sketch from which the method constructs a hierarchy of Support Structures used to capture and reproduce multi-scale alignments in real time. Lastly, while they showed promising results for texture synthesis, deep learning methods [SCO17, TLH19, RGF*20] are primarily image-based and have not yet been generalized to handle vector shape distributions effectively.

Rather than aiming for a statistically accurate distribution synthesis, we are looking for an efficient method suitable for an interactive 3D interface. Our solution emphasizes real-time processing and avoids the need for extensive pre-computation, unlike deep learning techniques. While we build on [OMC22], we tackle a new problem, the generation of 3D distributions from a 2D input sketch.

Sketch-based animation in scientific illustrations The use of sketches as input for animation has already been investigated, mostly in the context of fluid simulation: Zhu et al.[ZIH*11] present an interactive method for illustrating and animating fluid systems using sketch inputs and fluid simulation. Direct manipulation interaction enables this system to serve as a tool for visualizing and conveying information about cardiac abnormalities or surgical interventions. Draco [KCG*14] is an interactive system for animated illustrations that combines data samples and motion properties using kinetic textures. Kitty [KCGF14] extends Draco by incorporating interactivity and functional relationships between entities. The Hierarchical Motion Brushes system [MNB*14] allows users to brush animated content onto a 3D scene, supporting coarse-to-fine animation and the reuse of content at different levels of detail. However, it relies on offline processes for initial animation definition and lacks data processing for sample distribution structures, which the authors later addressed using hierarchical spatiotemporal clustering [MGC*16]. Closer to our concerns, the Energy Brushes system [XKG*16] concentrates on passive and secondary animations, where users sketch strokes and define flow particles to generate local energy patterns. Energy brushes represent the direction of energy and forces, continuously emitting particles along the sketched trajectory.

Biological visualization and storytelling While biological visualization has been extensively studied in the context of illustrative rendering or abstraction [Goo09], more recent work have explored visual preferences depending on the context of use or the application [MVB*17, GMF*21]. However, such data are only static and lack the motions and deformations occurring in the phenomenon of study. Therefore, there has been lately a strong interest in interfaces focused on narrative design, first for more general data [GP01, TRB*18, MBS22] and then targeting medical visualization [MGS*21, KSM*22, MGS*22]. In addition, McGill [McG08] highlights the potential of cinematic tools for narrative illustration of molecular phenomena. More recently, a few software specifically targeting narrative illustrations of molecular phenomena were introduced. CellPaint [GAB*18, GAF*21] is a user-friendly interactive digital tool focused on creating illustrations of the molecular structures of cells and viruses. Molecumentary [KSM*23] enables asynchronous communication of design ideas, controlling mesoscale animations of 3D models. The work presented here is complementary since rather than focusing on the reuse of existing 3D animated content in narration and addressing their differences of scale, it targets the interactive, sketch-based creation of content to support new narration ideas. Our work also complements Drew Berry's molecular animation [WB03]. While this work provides accurate visual representations in movie form, it lacks interactivity and requires all steps to be known beforehand. In contrast, we aim to create an interactive tool for biologists to illustrate and explore narrative options.

3. Biological Context and Design Choices

3.1. Preliminary study with biology researchers

To gain insights into terms and symbols used for conveying knowledge in biology, we conducted a preliminary study with four cell biology researchers specializing in different subtopics: one Research Director with 35 years of experience in cancer targets and experimental therapeutics, one Professor of 20 years of expertise in cell mechanics, one Research Director with 26 years of experience in cell polarity and division, and one Research Director of 38 years of expertise in membrane and cytoskeleton dynamics in the context of breast cancer.

The objective was to understand how they describe their phenomena of study regarding shape types, behaviors, and constraints. Each of them explained to us in simple terms the main steps of their studies, employing analogies or visual vocabulary when needed (e.g., "cutting similarly as scissors," "cells hooking onto collagen fibers to squeeze through"). We noted that all of them presented their phenomenon of interest through a narrative, in which an initial animated environment undergoes evolutionary changes in response to specific triggering events. While they managed to provide clear oral explanations using a few sketches or preexisting 2D figures as visual support (see Fig. 2), they lacked a visual medium to effectively illustrate the dynamic events they were describing. This study confirmed that relying solely on static 2D figures is insufficient for biologists because it does not capture the 3D and dynamic nature of a phenomenon, aligned with their mental vision. This observation motivates the framework we propose.

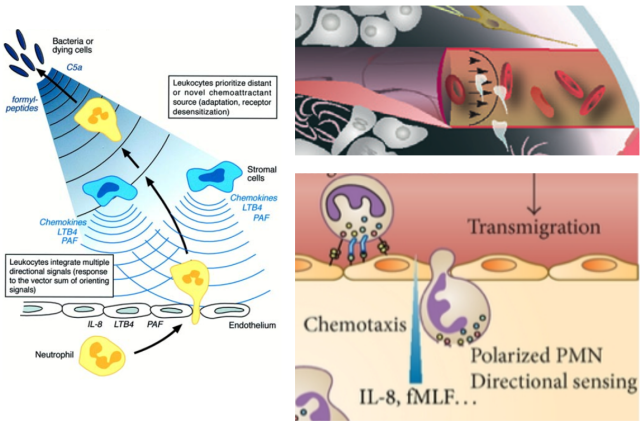

Figure 2: *Examples of 2D illustrations in cell biology: (left) [WSRDL*00], (top right) [SMJ18], (bottom right) [GCD*15].*

3.2. Schematic visual vocabulary in cell-biology

Our biologist co-author supplemented this study by formalizing the schematic vocabulary commonly used in biological illustrations to represent dynamic phenomena (see Fig. 2):

Arrows (left) describe the kinematic trajectory of an element, giving both direction and speed of displacement;

Streamlines (top right) indicate flow within a vessel using three arrows followed by a straight line for homogeneous flows or by a curved line for non-uniform flows (for example, allowing for a lower speed near the vessel membrane due to friction);

Gradient triangles (bottom right) are used to represent the difference of gradient in a medium (e.g., pressure gradient or chemical gradient) denoting attraction of elements.

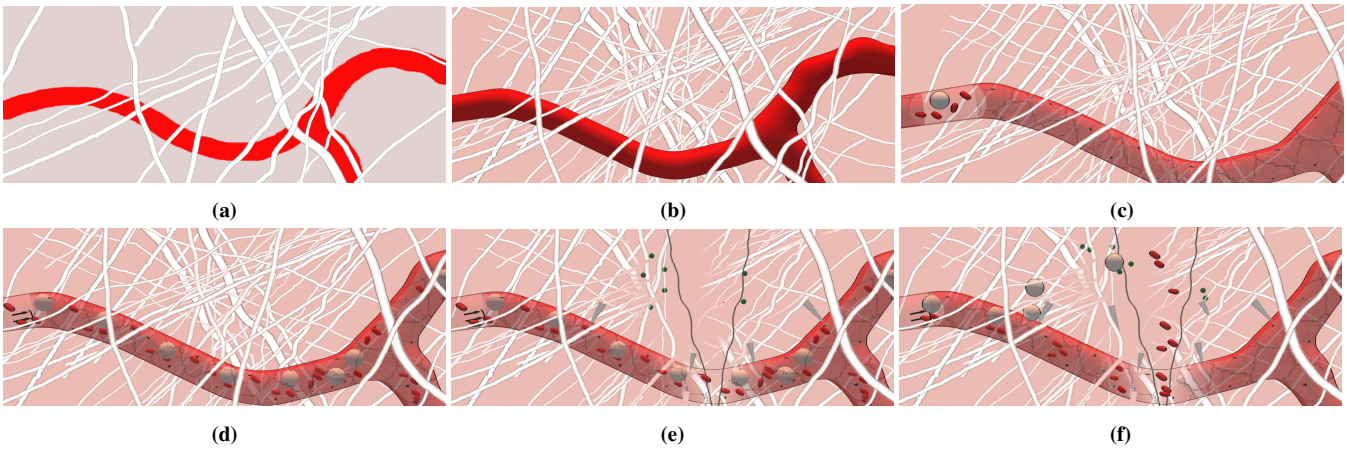

Figure 3: *Overview of Bio-Sketch: from creating the 3D biological environment (a-c) to designing a dynamic scenario (d-f).*

3.3. Findings and design goals for Bio-Sketch

Thanks to this in-depth preliminary study and the illustrations provided as support for narration, we classified the categories of 3D modeling tools, animation mechanisms, and rendering techniques to be targeted in our interactive content-creation tool.

- We noted the need for three main types of shapes: volumetric, organic shapes (e.g., vessel, cells, and their interiors), larger surfaces or volumes, unbounded at the observation scale (e.g., skin), and linear shapes (fibers). We also observed that both organs and larger surfaces often needed to be complemented by textures such as Voronoi-like tessellations standing for cell borders.
- We also noticed the presence of shape distributions, both for unbounded shapes (field of fibers) or nested bounded ones (cells within a vessel). This implies the need for models to extend a user-sketched input into a 3D distribution.
- In animation, we identified the need for sketching both global flows and individual trajectories, as well as controlling shape deformations. When depicting biological phenomena, controlling the timing of animations is also mandatory for building narrative scenarios, which can be specified as sequences of triggering events.
- Lastly, we identified the need for an expressive rendering mechanism using transparency and local shape cuts to convey nested shapes and motions.

3.4. Methodology and technical contributions

Inspired by visual representations and narration in biology, we designed Bio-Sketch, a new creative medium combining the progressive modeling of a biological environment in 3D with the creation of animated narratives. In this interactive sketch-based modeling framework, biologists can gradually convey their understanding of a phenomenon. This approach also serves as an entry point for non-expert users to explore the world of biological phenomena. In either case, the user designs an initial, stable environment and then triggers its temporal transformation. As depicted in Fig. 3, in this work, we focus on the infection phenomenon.

In our prototype implementation, we focused on the most challenging of the identified design goals, trying to offer a good balance between simplicity and richness of representation. We developed solutions to allow users not only to model individual 3D shapes for a sketch but also both bounded and unbounded 3D shape distributions (Sec. 4). In a similar spirit, the sketch-based animation mechanisms are not restricted to rigid motion but enable as well the specification of deformation. This is complemented by a timeline allowing users to both navigate and trigger narrative events (Sec. 5). At all stages of this creative process, an expressive visualization is provided to the user through a combination of 3D shading, contour drawing, transparency, and textures (Sec. 6). All the above methods were designed to offer real-time solutions, which is mandatory to be integrated into an iterative creative media.

Case study: To validate the ability to create an animated biological story with our system, we selected an infection phenomenon, showing the role of white blood cells after wounds (see Fig. 1, left): when the skin, together with the underlying fibers and small vessels, is cut, red blood cells located near the incision area escape through the opening. In contrast, white blood cells, which play a crucial role in the immune response, are attracted to the vessel membrane where they reside. They roll along this inner membrane before finally escaping by undergoing shape deformation to pass through small openings between the membrane cells of the vessel. The white blood cells that successfully escape from the vessel navigate through the collagen fibers to the cut area, to aid in the wound healing process. We have chosen this example to illustrate our technical contributions throughout this article since it comes with an illustration by a professional artist, validated by biologists, and provides an inspiring visual reference for our system. Fig. 3 shows a) the input 2D sketch, b) the resulting 3D environment, c) the example showing the cell distribution inside a vessel, as well as the rendering of its surface, d) the resulting distribution of cells and the use of the flow tool to specify their movement, e) the cutting of tissues triggering an infection modeled using gradient triangles and bacteria (in green) and f) the deformation of the white cells crossing the vascular membrane and attracted by the bacteria, while the red cells escape through the cut.

4. Creation of the 3D biological environment

Given the organic characteristics inherent in the shapes we want to create (e.g., vessels, cells, and the inner structures within cells), we chose implicit surfaces as the most appropriate shape representation in our sketching system. This is built on Matisse [BPCB08], an interactive sketch-based system for 3D implicit modeling. At each stage, brushes are used to paint a 2D region. The medial axis of the region is extracted, simplified into a graph of segments complemented with radii information, and used as a skeleton to generate a scalar field standing for the density of matter around the skeleton, and decreasing with the distance. The surface is defined as an isosurface of this field. While flat skeletons create shapes with flat silhouettes, more complex shapes can be created by sketching parts from different viewpoints and blending them by summing their fields. In this work, we use the SCALIS model [ZBQC13] to improve the fidelity to the specified radii and avoid the blurring of small details. We also complement this system with the ability to sketch both unbounded and bounded shape distributions as input.

4.1. Visually unbounded 3D distributions from a sketch

Our goal is to generate, from a 2D input sketch, a 3D distribution of elements inside which the user can navigate (e.g., the volume of collagen fibers surrounding the vessel in our case study). The insight is to let users sketch a small sample of what they would like to see from the front. We then expand this exemplar in 2D of a given ratio k before immersing the resulting 2D distribution in 3D, using an extension of depth h. The resulting 3D volume of material can then be repeated with adapted rendering to fake an unbounded environment.

The 2D extension of the user-sketched vector texture is computed using the method of Olivier et al. [OMC22]. Based on multi-resolution alignment analysis, this method efficiently generates a larger 2D distribution of strokes that seamlessly extends the 2D sketched sample. More precisely, vector sketched strokes are clustered into a hierarchy of linear or curved support structures. Each structure comprises a main direction and a ribbon that contains the clustered sub-structures. At the synthesis stage, from top to bottom of the hierarchy, the existing supports in the exemplar are extended and duplicated with some randomization to match the larger output space. Clones of the input strokes are finally distributed, with adequate jittering, along the lowest-level structures.

In the second stage, we infer depth parameters for each stroke to extend the 2D vector texture into a 3D distribution. Our solution builds on the usual general-view perceptual assumption: any stroke alignment in the 2D exemplar intentionally depicts a 3D alignment rather than being due to the projection from a very specific viewpoint. Therefore, we immerse clustered strokes using quite similar depths through the following two steps. For each cluster, we first define a midpoint's depth d_e and a slope angle δ with respect to the sketching plane XY. These parameters are randomly chosen in $[0,h]$ for d_e (with h the depth of the output domain) and in $[-\pi/4,\pi/4]$ for the slope (a value chosen from perceptual studies [WBCG09]). Each stroke in the cluster is then immersed in 3D as a texture projected onto an oblique billboard, as follows. While its XY projection is maintained, the depth of its extremities is computed by adding (respectively removing) from d_e the value $\Delta_z = \frac{lg}{2} * tan(\delta)$, with lg, the length of the stroke's support 2D segment.

Duplicating the resulting volume along the depth axis is essential to fake an infinite field of 3D strokes. These instances can be set to randomly overlap in depth, allowing for complex arrangements. To achieve this, we duplicate the initial volume and shift each instance in the depth direction by a factor $l * h$, with l an integer in $[1;\infty]$. This process generates a virtually infinite field of anisotropic elements. To visualize them, we rely on two perceptual assumptions. First, distant objects fade away, implying that strokes in deeper layers appear less prominent. Second, we consider the 3D space to have a toroidal topology in depth, meaning that objects that go beyond a certain depth "wrap around" and reappear at the front. By incorporating these perceptual assumptions, we can effectively display the multiple layers and create a visually coherent volumetric representation (see Fig. 4).

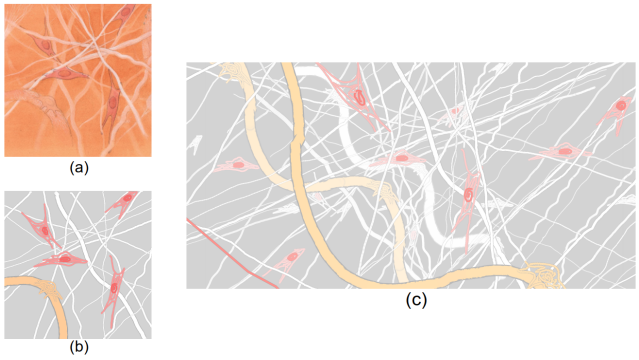

Figure 4: *(a) Biological illustration with cells navigating within a 3D distribution of fibers; (b) 2D input sketch inspired from a); (c) Result of our 3D distribution synthesis method.*

4.2. Volumetric distributions in a container shape

In addition to creating unbounded distributions, the user should be allowed to fill any bounded, volumetric shape with a distribution of smaller shapes (e.g. filling a vessel with blood cells in our case study). Our solution is a second extension of [OMC22], as follows.

The user can select any volumetric shape created by sketching as a container. The latter is thus an iso-surface of a scalar field generated by one or several flat skeleton graphs. Using a camera that faces the local skeleton plane, we allow the user to select a small part of the container (which is then made fully transparent) and create an exemplar of the desired distribution, either by directly drawing sub-shapes or by instantiating a few existing ones. Indeed, this input distribution only consists of bounded shapes.

We then use a 2.5D solution for analysis and synthesis. The distribution exemplar is processed in the skeleton plane, considering the 2D projections of the sub-shapes. The challenge here is to extend the original 2D distribution analysis and synthesis method to free-form regions. Indeed, since synthesis is to be constrained to the inside of the free-form container shape, we cannot directly retrieve shape alignments and their lead directions through successive clustering of straight lines. Fortunately, the scalar field defining the

implicit container offers an excellent base for clustering. Indeed, the notion of shapes being quasi-aligned just needs to be translated into shapes positioned around the same iso-value. Their support structure is then defined by the corresponding iso-line curve in the scalar field rather than by a straight line. As in the original method, support structures are recursively clustered (with ribbons expressing the variance of the elements they represent). To achieve this, their difference of iso-value is used as clustering distance. These simple changes enable us to apply the existing algorithms to analyze the free-form input sample and then extend the distribution to the full 2D projection of the container shape.

Note that the iso-lines do not need to be explicitly computed at the synthesis stage. We iteratively instantiate the next shape to be generated along them through a series of small displacement steps from the previous position. Each step consists of a translation along the tangent direction followed by a projection along the field's gradient to come back to the targeted iso-value. We define the tangent direction as the cross-product between the field's gradient and the depth direction, and the step size from the local curvature. We then add a slight perturbation to this new position to avoid salient repetitions at the synthesis stage.

Once the planar distribution is extended to the full container, the clones created for each of the sub-shapes recover their initial depth coordinates, with some jittering proportional to the deviation in their initial depth distribution. This enables them to spread within the full 3D container. Fig. 5 shows the robustness of the method to shapes with branching, thanks to the definition of support structures as iso-lines.

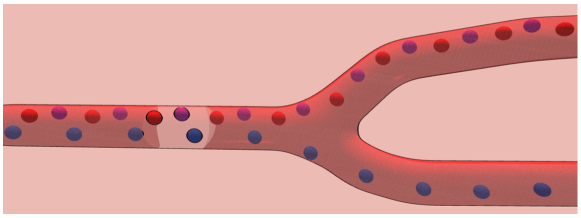

Figure 5: *From an exemplar to a distribution of cells in a vessel.*

5. Design of dynamic scenarios

5.1. Sketching motion

To integrate animation design into our sketch-based creative system, we rely on the visual vocabulary extracted from the analysis of biological illustrations in Section 3. This allows us to establish a clear correspondence between the biologist's familiar schematic representations and the desired motion or deformation. Although always specified in the plane facing the camera and then applied to shapes at constant depth, the motions and deformations described below could be added to the full 3D animation by using the motion design tools from different viewpoints and combining their effects.

Arrow tool: The user draws a trajectory starting from the target object. We sample this trajectory according to the drawing time so that the speed of the gesture controls the speed of movement. We then insert the displacement between two successive points along the timeline to facilitate the synchronization of different motions.

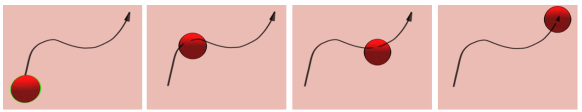

Figure 6: *The cell is following the user's arrow.*

Flow tool: It enables the user to specify a flow strength inside a container shape. The user sketches three arrows followed by a straight or curved line (Fig. 7). Defined in the 2D skeleton plane, the flow defines the velocity of the medium present in the container.

The main issue is to extend the flow speed to a 2D vector field that spans the whole container. From the user's sketch, we retrieve the main direction, maximum speed, and lateral damping of the flow. Even if the sketched flow is small compared to the size of the container, the flow is considered to span its entire cross-section, so that the possible speed reduction (specified through a curved line) occurs near the container's surface. The flow is then propagated to the entire container by considering that its speed remains constant along each of the iso-lines of the field. The flow direction is always considered tangent to these iso-lines.

During the animation loop, the flow applies a viscous friction force to any sub-shape present in the container. This force is applied in the 2D sketching plane, resulting in a 2.5D animation where the animated sub-shapes keep their depth position with respect to the local skeleton plane. Lastly, for implicit shapes such as the container being closed surfaces, we cut the most distant extremities of the container in the flow direction and give it a toroidal topology by having sub-shapes disappearing at one end and reappearing at the other end. Collisions are detected and processed to avoid any unwanted escape of the shapes through the container's surface.

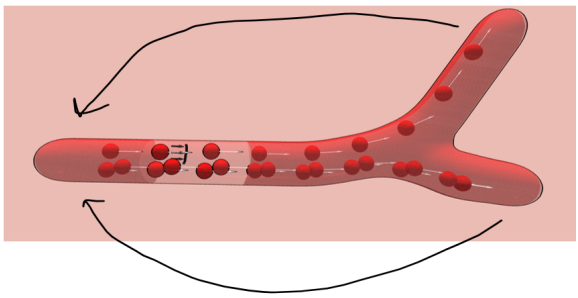

Figure 7: *Use of the flow tool to create a velocity field within a container. The latter being a closed, implicit surface, we give it a toroidal topology (black arrows) by bringing back downstream all the objects that are advected beyond a predefined upstream limit.*

Gradient triangles: They can be positioned by the user in the 2D plane facing the camera to specify directional force fields due to chemical gradient (see Fig. 1, right). These forces are applied to nearby shapes of a given category, with strength decreasing with the distance. In our case study, this can be used to model the chemical procedure that attracts the white blood cells to the site of infection.

5.2. Specifying shape deformation

Being able to deform shapes and possibly change their topology through merging and fusion is a major challenge in biological animations. For example, in our case study, white blood cells are deformed to escape through small gaps between the vessel's membrane cells and then navigate among fibers. We design two mechanisms to allow the specification of deformation.

Shape morphing: Sketch-based modeling is used to specify initial and final shapes, which are then placed at specific instants along the timeline. A morphing algorithm based on Minkowski sums [GA96] is used to generate the intermediate shapes (see Fig. 8).

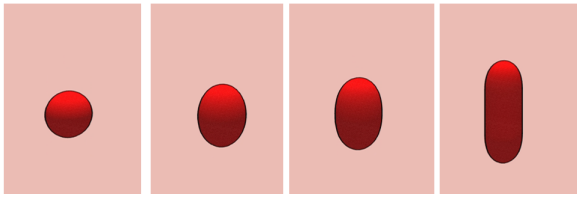

Figure 8: *Morphing between the user-specified initial (left) and final (right) states.*

Shape deformation through a hole: Let T be the target shape, which is to traverse a hole H through a surface S. We define R_t and R_h, respectively as the radius of the bounding sphere of the target and hole. We want to model the following behavior:

- if $R_t < R_h$, the target will escape without undergoing any deformation,
- if $R_t > 2R_h$, the compression of the target will be excessive, making it impossible to escape,
- Finally, if neither of the above conditions applies, the target will deform while maintaining a quasi-constant volume.

As the target T and the surface S (i.e., a vessel) are modeled with our system, they are represented using implicit surfaces. The hole itself is modeled as an implicit primitive of a negative field generated by a skeleton point. To simulate the effect of passing through the hole on the target shape (which should squeeze and extend in the other direction before retrieving its shape on the other side of the surface), we add a new deformation field to the one defining T, as follows.

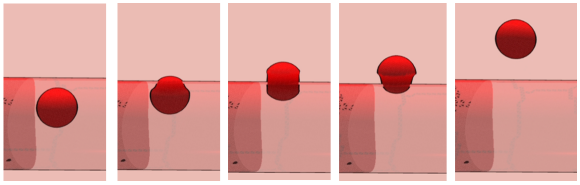

Figure 9: *Cell deforming to cross a membrane through a hole.*

We center a deformation primitive on the hole, with a region of influence divided into three distinct parts: an outer ring, a middle ring, and the remaining spherical region around the center. In the outer interior ring, we apply a negative field to locally compress target shapes ensuring that they fit within the available hole space. In the middle ring (of the hole's size), we use a positive field to help preserve the volume of the target during deformation: it makes the approaching target shape stretch along its trajectory when approaching the hole, as well as when it leaves it on the other side. The deformation field is set to zero in the spherical inner region.

The effect of this deformation primitive, applied to all the targets passing by, is depicted in Fig. 9. In our case study, this is used to model white cells passing through the vessel's membrane.

5.3. Narrative scenarios

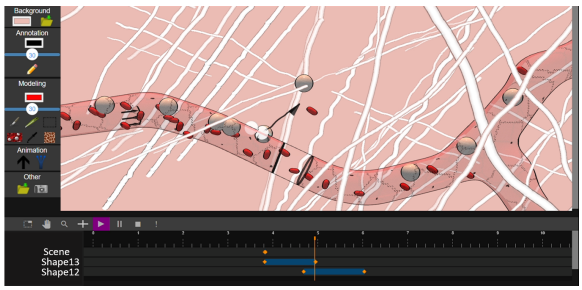

Figure 10: *Bio-Sketch interface. Left: Toolbox; Bottom: Timeline.*

As depicted in Fig. 10, the interface of our creative system consists of a toolbox and a timeline. When an animation is triggered, the timeline becomes active. We handle two types of animations: periodic animations such as the flow in a container or the random motion of bacteria, which last for the whole duration of a session, and non-periodic ones, which are triggered by a specific event (e.g., a chemical gradient appearing due to a cut). Only these events and the subsequent animations are shown along the timeline.

We propose two modes for our system: creation mode, where the user navigates to create or edit shapes or animations and to position their triggering events along the timeline, and play mode, where the user visualizes the current, animated biological sketch from a given point in time. When a new event is instantiated, such as the one at time 4.8 in Fig. 10 (using the arrow tool to attract a white blood cell out of a vessel through a hole), we establish an event state in which we record the event time as well as the current state of all the elements in our environment. When the user returns to this event, we simply reuse these recorded states.

Switching between these modes enables users, particularly biologists, to gradually illustrate the desired phenomenon, as well as design and record an associated narration if needed.

6. Expressive rendering

6.1. Illustration-inspired rendering

Throughout the creative process, the user is provided with an expressive, non-photo-realistic rendering, designed to allow a good perception of the 3D scene and the animated events, while making the whole scene look like a shaded 3D sketch inspired by the artistic illustration in Fig. 1 (left).

Each shape created using our sketch-based modeling system is

rendered using a default material that combines Phong shading with a procedural value noise (to give a paper-like effect). We also add a black outline effect using a pen-and-ink shader (see the vessel and fibers in Fig. 1). This choice of rendering is flexible enough to be further adapted according to the needs of the animation scenario.

Organic tesselation of surfaces: To give an organic membrane look to surfaces such as the vessels, we offer the option to complement a 3D surface with Voronoi-like regions standing for the membrane cells, visually delimited by black dots. We call *kernels* the centers of these regions.

To achieve this, the user clicks on the desired locations to position kernels over the chosen surface. At each click, we use the raycasting from the camera viewpoint to find $P1$ the first intersection with the implicit surface that we select as a kernel point. Rather than using the second ray-surface intersection point $P2$ as the kernel on the other side of the surface, we jitter the ray direction before computing the next intersection point, starting from $P1$. This randomization avoids unrealistic cell alignments.

Once all kernels are specified, we compute an approximate Voronoi diagram over the free-form surface, directly encoded in the surface geometry attributes to enable 3D navigation. We first classify the mesh vertices relative to their closest kernel, using the Euclidean distance as a coarse approximation of the geodesic distance. We also retrieve the vertex closest to each kernel, which we display in black. We then scan the mesh faces to identify triangles whose vertices are associated with three different kernels. These faces are the location of the corners of the Voronoi regions we are looking for, which we approximate as their barycenter. Therefore, we add the barycenter together with its interpolated normal as a new mesh vertex and use it to tessellate the face into three sub-triangles connecting it to the original vertices. Similarly, we add a new vertex at the barycenter of each edge between two vertices associated with two different kernels and insert it within the mesh by cutting the triangle into two parts. This results in salient Voronoi-like regions, delimited by dotted lines, with a shape error of the order of the length of triangle edges in the input mesh (see Fig 11).

Figure 11: *Organic-looking tesselation of surfaces. a) Shape to texture; b) User-selected kernels; c) Result.*

Visualization of nested shapes: To make it easier to design nested shapes such as cells inside a vessel, we drew inspiration from the artistic illustration in Fig. 1 (left), and, in particular, the way the vessel was made locally transparent to show the interior at the far left. Therefore, while all surfaces are by default only partially transparent, we allow users to visually "erase" a portion of their membrane using a brush or the intersection with a box. This makes their interior fully visible, which eases the sketching of sub-shapes. In addition, as shown in the supplemental video, we simplified the phagocytosis phenomenon by progressively reducing the transparency of the bacteria and then the white cell one.

Faking a cut through tissues: While modeling and animating a cut through biological tissues would have been a challenge, in particular, considering their anisotropic behavior due to fibers and vessels, we used a visualization trick to symbolize a cut through tissues in association with a temporal event that triggers changes in the animation modes.

The user is provided a cutting tool inspired by a traditional cutter to draw a cut trajectory in the 2D viewing plane. Assuming that due to the tension of tissues, a real cut would result in a triangular opening, we simply disable the display of all the elements that lie within an isosceles triangle (of angle $\frac{\pi}{12}$ in our implementation) whose apex is at the extremity of the cut stroke. To achieve this, we record all the shapes (implicit surfaces and 3D textures such as fibers) intersected by the user's stroke, given its predefined depth. We complement this list with the other elements in the triangular region, using raycasting at grid-based sampling points over the triangular shape.

We then delimit two areas where we adapt the rendering: inside the triangular cut where only silhouettes are displayed, and in an offset region around the cut where we progressively blend toward standard rendering. For all the intersected shapes in the cut region, we remove the mesh triangles that, when projected onto the plane of the triangle, fall within the boundaries of the cut using the fragment shader. The user's cut stroke is duplicated and displayed twice to delimit the opening (see Fig. 12).

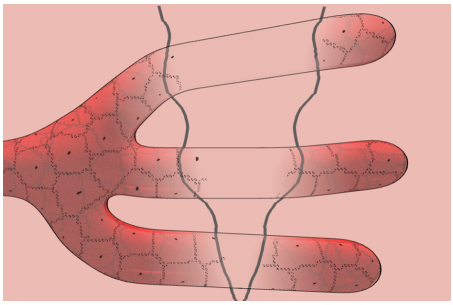

Figure 12: *Rendering effect faking a vertical cut through organs.*

6.2. Visualization and navigation within stroke distributions

As detailed in Section 4, most of the static, 3D environment that serves as a background is made of 3D strokes. To effectively visualize them, we assign different materials to their inner area and their outline. The central inner area is assigned the stroke's original color, while the outline color ranges from light gray to black to achieve a smoother transition with a darker background. To enhance the three-dimensional appearance of elongated strokes such as fibers, we smoothly blend the inner and outline colors, which gives the stroke a cylindrical look.

In our implementation, we use a dolly zoom to smoothly transition from the interactive sketching session, where the user sketches

the input 2D vector texture, to the exploration of the 3D output scene. The volumetric stroke-distribution layers are displayed in each frame with increasing transparency and decreasing stroke thickness with distance to the camera. This operation is performed stroke by stroke, depending on the distance between the stroke center point and the camera. We also transform the color of distant strokes into white by converting its RGB color to HSL format and updating the luminance parameter from its original value to 1. When a stroke is behind the camera, we shift it in depth to the deepest visible volume, giving the illusion of infinite space in depth. This delivers better visual results during exploration than applying the same transformation to volumetric layers one by one while maintaining real-time performance (see Fig. 4).

7. Results

Performances: We implemented our prototype system in WebGL. The table presented below was generated by measuring the runtime performance of Google Chrome on an Intel(R) Core(TM) i7-7920HQ CPU operating at a frequency of 3.10GHz while using our case study as a reference at each stage.

Process	Execution time in seconds
3D field of fibers	0.098
Vessel creation	0.791
Surface texturing	1.13
Cut (on the shader)	Instant
Cell creation (per cell)	0.090
Cell distribution	2.44s
Flow	0.015
Vertical cut event	3
Cell deformation	0.058

Case study and experts feedback: The full test case was created in less than 15 minutes by a non-expert user with a few hours of experience with our system. Fig. 1, 3, and 10 already showed a few key steps of our experiment on the chosen case study, described in Section 3. In addition to these results, we refer the reader to the supplementary video for an interactive presentation of how this example was created using our prototype and the final animated 3D sketch. The video also shows a few extra examples.

Additionally, we engaged two experts, one specializing in biology and the other in biological illustrations, to gather their expert feedback (see below). Our objective was to assess the potential of a specialized a system like Bio-Sketch in the context of evolving into a more universally applicable tool for narrative illustrations in cellular biology.

"As a biologist involved in the development of nanomedicines, I would use Bio-Sketch to explain how mechanical, physical, and chemical forces may guide the diffusion of a (nano) drug through different biological tissues. In the current form, Bio-Sketch proposes very interesting tools for starting my demo and insists on critical parameters such as "crossing the endothelial barrier", "impact of interstitial pressures", "osmolarity", "density of the tumor microenvironment", "density of extracellular matrices" ... This is currently difficult with the graphic tools we are using. In addition to its 3D rendering, Bio-Sketch would also give me access to a 4th dimension (time), a 5th one (pressures), and many others."

"As an expert in scientific transmission through drawing and animation, I believe Bio-Sketch provides an interesting demonstrator for explaining biological phenomena in the context of education, workshops, and general public conferences. Conveying an idea of the richness of the biological environment is usually difficult, even if realism is not the goal: it requires skills in drawing, modeling, and/or animation. Bio-Sketch offers an elegant solution to the evocation of collagen fiber landscape in depth. It also provides a good way to address the distribution and behavior of immune cells in vessels and tissue. Further applications of such a tool remain strongly dependent on the explanation given by the user/speaker to point out the necessary simplifications and omissions (typically, the real timing of phenomenon is only loosely respected in such animations; also, in the present study, coagulation was simply not represented). In general, I don't believe such a tool can be a stand-alone solution, eliminating the need for complex static images. However, it provides a good way to initiate discussions and questions on aspects that can be easily passed over when only static representations are presented."

8. Discussion & Limitations

We proposed a full pipeline for interactively modeling, animating, and rendering a 3D scene, to synthesize an animated 3D sketch in the context of the infection phenomenon. While our design choices were strongly influenced by biological illustrations and our motion design interface incorporates conventions from biological sketching, we designed our user interface to make it accessible to both experts and non-experts. In addition, our system leaves great freedom to the user to model various shapes at different scales thanks to the sketch-based system for 3D implicit modeling and our cutting tool. When it comes to animation, different phenomena might require the inclusion of additional motion or deformation behaviors.

It is worth noting that our intention was not to create exact replicas of our inspirational figure. As confirmed by our scientific expert, our tool complements detailed static images rather than replicating them precisely. In our current prototype, users have the option to add extra information through annotations. However, we recognize that enabling users to insert text and labels would be a valuable feature, enhancing comprehension and usability.

Although not exhaustive, our classification (Sec.3.3) was sufficient to express the cell biology scenario we studied. While future extensions could explore more comprehensive sets of biological notions and phenomena, the limitations of our system so far are mostly due to our implementation of these basic tools. We list the main ones below.

Firstly, while our iterative, sketch-based modeling system enables users to create free-form 3D shapes by adding shape parts from different viewpoints, our solution for animation is only 2.5D. In particular, it restricts the flow in a container to shape parts with a flat skeleton, such as a flat portion of a vessel. Moreover, while the strength of the flow can be efficiently encoded based on the iso-value, the direction of the flow along an iso-line of the implicit field function is currently determined using a scalar product with the viewing direction. This would not allow the modeling of the flow in a curved vessel which would go to the right and then bend

back to the left. To solve this and extend the flow motion to 3D vessels, we would need the introduction of local frames propagating in a consistent way along the vessel.

Secondly, while implicit modeling eased the interactive design and deformation of organic-looking shapes, modeling Voronoi cells on these surfaces was not obvious. As their lack of parameterization makes texture mapping difficult, we relied on explicit mesh tessellation to achieve this, with a method that leaves way for improvement. The user needs to manually define all kernels of the Voronoi cell, while an automatic sampling method would have been desirable. Moreover, the resulting tessellation is only visually acceptable for dense enough sampling with respect to the surface's curvature. The Euclidean distance used in Voronoi cell computation should therefore be replaced by geodesic distance along the surface, as described in Meuschke et al. [MVB*17].

Lastly, to enable real-time cutting through tissues, we chose to preserve the geometry of elements and rely on a rendering trick to selectively remove the unwanted areas at the rendering stage. This does not reflect the complex local deformation, partly irreversible, that happens during a cut.

Animating cutting processes in a more realistic way would require a volumetric physically-based simulation in anisotropic material by encoding the tension that fibers and vessels give to biological tissues and the latter's relaxation when these elements are cut. While animating this in real-time would be a challenge, it could be an interesting direction for future work.

9. Conclusion

We presented Bio-Sketch, an interactive medium designed to help biologists convey their mental image of an infection phenomenon in the form of an animated 3D sketch, but also to help non-specialists discover the immune response to infection. After classifying the shapes and animation tools that would be useful to this end, thanks to a pre-study with experts in biology, we proposed a full pipeline to achieve it and implemented a prototype system. The latter was experimented on a challenging test case: the modeling of the immune reaction of tissues when the skin gets cut. It enabled users to create a bio-sketch, which could serve as support for narration, in a few tenths of minutes.

Although not fully tested yet, we hope that our new system and its future extensions will open up new avenues for visual communication, education, and research in the field of biology and beyond.

In future work, we would like to explore the extension of Bio-Sketch to the design of other phenomena and especially to multi-resolution sketches, letting users both interactively create and zoom into animated scientific illustrations with unlimited levels of detail. As in [WBCG09], some machine learning would be required to allow the duplication of sketch-based details to other parts of a support shape, which is mandatory for navigation and zoom-out.

Acknowledgments

We would like to thank the reviewers for their comments and suggestions that help us improve this paper. We are also grateful to Julien Husson of Ecole Polytechnique, Philippe Chavrier and Mathieu Piel of Institute Curie for their participation to our pre-study. We are thankful to Tara Butler and Pascal Guehl for their help on the project. This paper was partly funded by Marie-Paule Cani's fellowship on Creative AI from Hi!Paris.

References

[ABL*21] Abdrashitov R., Bang S., Levin D., Singh K., Jacobson A.: Interactive Modelling of Volumetric Musculoskeletal Anatomy. *ACM Trans. Graph. 40*, 4 (2021), 1–13. doi:10.1145/3450626.3459769. 2

[BBT*06] Barla P., Breslav S., Thollot J., Sillion F., Markosian L.: Stroke pattern analysis and synthesis. *Computer Graphics Forum 25*, 3 (2006), 663–671. doi:10.1111/j.1467-8659.2006.00986.x. 2

[BPCB08] Bernhardt A., Pihuit A., Cani M. P., Barthe L.: Matisse: Painting 2d regions for modeling free-form shapes. In *Proceedings of the Fifth Eurographics Conference on Sketch-Based Interfaces and Modeling* (2008), p. 57–64. doi:10.2312/SBM/SBM08/057-064. 2, 5

[Cha18] Chabrier R.: *Globule, the magazine of all cells*. 2018. 1

[GA96] Galin E., Akkouche S.: Blob Metamorphosis based on Minkowski Sums. *Computer Graphics Forum 15*, 3 (1996), 143–153. doi:10.1111/1467-8659.1530143. 7

[GAB*18] Gardner A., Autin L., Barbaro B., Olson A. J., Goodsell D. S.: Cellpaint: Interactive illustration of dynamic mesoscale cellular environments. *IEEE Computer Graphics and Applications 38*, 6 (2018), 51–66. doi:10.1109/MCG.2018.2877076. 3

[GAF*21] Gardner A., Autin L., Fuentes D., Maritan M., Barad B. A., Medina M., Olson A. J., Grotjahn D. A., Goodsell D. S.: Cellpaint: Turnkey illustration of molecular cell biology. *Frontiers in bioinformatics 1* (2021), 660936. doi:10.3389/fbinf.2021.660936. 3

[GCD*15] Gamara J., Chouinard F., Davis L., Aoudjit F., Bourgoin S.: Regulators and effectors of arf gtpases in neutrophils. *Journal of Immunology Research 2015* (2015), 1–15. doi:10.1155/2015/235170. 3

[GMF*21] Garrison L., Meuschke M., Fairman J., Smit N. N., Preim B., Bruckner S.: An Exploration of Practice and Preferences for the Visual Communication of Biomedical Processes. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2021), pp. 1–12. doi:10.2312/vcbm.20211339. 3

[Goo09] Goodsell D. S.: The machinery of life. doi:10.1007/978-0-387-84925-6. 3

[GP01] Gershon N., Page W.: What storytelling can do for information visualization. *Commun. ACM 44*, 8 (2001), 31–37. URL: https://doi.org/10.1145/381641.381653, doi:10.1145/381641.381653. 3

[HLT*09] Hurtut T., Landes P.-E., Thollot J., Gousseau Y., Drouillhet R., Coeurjolly J.-F.: Appearance-Guided Synthesis of Element Arrangements by Example. In *Proceedings of the 7th International Symposium on Non-Photorealistic Animation and Rendering* (2009), p. 51–60. doi:10.1145/1572614.1572623. 2

[IMIM08] Ijiri T., Měch R., Igarashi T., Miller G.: An Example-based Procedural System for Element Arrangement. *Computer Graphics Forum 27*, 2 (2008), 429–436. doi:10.1111/j.1467-8659.2008.01140.x. 2

[Ise15] Isenberg T.: A Survey of Illustrative Visualization Techniques for Diffusion-Weighted MRI Tractography. In *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data* (2015), pp. 235–256. doi:10.1007/978-3-319-15090-1_12. 2

[KCG*14] KAZI R. H., CHEVALIER F., GROSSMAN T., ZHAO S., FITZMAURICE G.: Draco: Bringing Life to Illustrations with Kinetic Textures. In *Proceedings of the SIGCHI* (2014), CHI '14, p. 351–360. doi:10.1145/2556288.2556987. 3

[KCGF14] KAZI R. H., CHEVALIER F., GROSSMAN T., FITZMAURICE G.: Kitty: Sketching Dynamic and Interactive Illustrations. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (2014), p. 395–405. doi:10.1145/2642918.2647375. 3

[KSM*22] KLEINAU A., STUPAK E., MÖRTH E., GARRISON L. A., MITTENENTZWEI S., SMIT N. N., LAWONN K., BRUCKNER S., GUTBERLET M., PREIM B., MEUSCHKE M.: Is there a Tornado in Alex's Blood Flow? A Case Study for Narrative Medical Visualization. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2022). doi:10.2312/vcbm.20221183. 3

[KSM*23] KOUŘIL D., STRNAD O., MINDEK P., HALLADJIAN S., ISENBERG T., GRÖLLER M. E., VIOLA I.: Molecumentary: Adaptable narrated documentaries using molecular visualization. *IEEE Transactions on Visualization and Computer Graphics 29*, 3 (2023), 1733–1747. doi:10.1109/TVCG.2021.3130670. 3

[LGH13] LANDES P.-E., GALERNE B., HURTUT T.: A shape-aware model for discrete texture synthesis. *Computer Graphics Forum 32*, 4 (2013), 67–76. doi:10.1111/cgf.12152. 2

[LVPI18] LAWONN K., VIOLA I., PREIM B., ISENBERG T.: A Survey of Surface-Based Illustrative Rendering for Visualization. *Computer Graphics Forum 37*, 6 (2018), 205–234. doi:10.1111/cgf.13322. 2

[MBS22] MÖRTH E., BRUCKNER S., SMIT N. N.: Scrollyvis: Interactive visual authoring of guided dynamic narratives for scientific scrollytelling. *IEEE Transactions on Visualization and Computer Graphics* (2022), 1–12. doi:10.1109/TVCG.2022.3205769. 3

[McG08] MCGILL G.: Molecular movies… coming to a lecture near you. *Cell 133*, 7 (2008), 1127–1132. doi:10.1016/j.cell.2008.06.013. 3

[MGC*16] MILLIEZ A., GUAY M., CANI M.-P., GROSS M., SUMNER R. W.: Programmable Animation Texturing using Motion Stamps. *Computer Graphics Forum 35*, 7 (2016), 67–75. doi:10.1111/cgf.13004. 3

[MGS*21] MEUSCHKE M., GARRISON L., SMIT N., BRUCKNER S., LAWONN K., PREIM B.: Towards narrative medical visualization, 2021. arXiv:2108.05462. 3

[MGS*22] MEUSCHKE M., GARRISON L. A., SMIT N. N., BACH B., MITTENENTZWEI S., WEISS V., BRUCKNER S., LAWONN K., PREIM B.: Narrative medical visualization to communicate disease data. *Comput. Graph. 107*, C (2022), 144–157. doi:10.1016/j.cag.2022.07.017. 3

[MNB*14] MILLIEZ A., NORIS G., BARAN I., COROS S., CANI M.-P., NITTI M., MARRA A., GROSS M., SUMNER R. W.: Hierarchical motion brushes for animation instancing. In *Proceedings of the Workshop on Non-Photorealistic Animation and Rendering* (2014), p. 71–79. doi:10.1145/2630397.2630402. 3

[MVB*17] MEUSCHKE M., VOSS S., BEUING O., PREIM B., LAWONN K.: Glyph-based comparative stress tensor visualization in cerebral aneurysms. *Computer Graphics Forum 36*, 3 (2017), 99–108. doi:10.1111/cgf.13171. 3, 10

[MWT11] MA C., WEI L.-Y., TONG X.: Discrete Element Textures. *ACM Trans. Graph. 30*, 4 (2011), 1–10. doi:10.1145/2010324.1964957. 2

[OMC22] OLIVIER P., MEMARI P., CANI M.-P.: Structured shape-patterns from a sketch: A multi-scale approach. In *Proceedings of Graphics Interface 2022* (2022), pp. 161 – 170. doi:10.20380/GI2022.17. 2, 5

[PCP10] PIHUIT A., CANI M.-P., PALOMBI O.: Sketch-based modeling of vascular systems: A first step towards interactive teaching of anatomy. In *Proceedings of the Seventh Sketch-Based Interfaces and Modeling Symposium* (2010), p. 151–158. doi:10.2312/SBM/SBM10/151-158. 2

[PS18] PREIM B., SAALFELD P.: A survey of virtual human anatomy education systems. *Computers & Graphics 71* (2018), 132–153. doi:10.1016/j.cag.2018.01.005. 2

[RGF*20] REDDY P., GUERRERO P., FISHER M., LI W., MITRA N. J.: Discovering pattern structure using differentiable compositing. *ACM Trans. Graph. 39*, 6 (2020), 1–15. doi:10.1145/3414685.3417830. 2

[ROM*15] ROVERI R., ÖZTIRELI A. C., MARTIN S., SOLENTHALER B., GROSS M.: Example Based Repetitive Structure Synthesis. *Computer Graphics Forum 34*, 5 (2015), 39–52. doi:10.1111/cgf.12695. 2

[SBP*16] SAALFELD P., BAER A., PREIM U., PREIM B., LAWONN K.: A sketch-based interface for 2d illustration of vascular structures, diseases, and treatment options with real-time blood flow. In *Computer Vision, Imaging and Computer Graphics Theory and Applications* (2016), pp. 19–40. doi:10.1007/978-3-319-29971-6_2. 2

[SCO17] SENDIK O., COHEN-OR D.: Deep Correlations for Texture Synthesis. *ACM Trans. Graph. 36*, 5 (2017), 1–15. doi:10.1145/3015461. 2

[SMJ18] STYLIANOPOULOS T., MUNN L. L., JAIN R. K.: Reengineering the Physical Microenvironment of Tumors to Improve Drug Delivery and Efficacy: From Mathematical Modeling to Bench to Bedside. *Trends in Cancer 4*, 4 (2018), 292–319. Special Issue: Physical Sciences in Oncology. doi:10.1016/j.trecan.2018.02.005. 3

[SSPOJ16] SAALFELD P., STOJNIC A., PREIM B., OELTZE-JAFRA S.: Semi-Immersive 3D Sketching of Vascular Structures for Medical Education. In *Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine* (2016), p. 123–132. doi:10.5555/3061507.3061528. 2

[TLH19] TU P., LISCHINSKI D., HUANG H.: Point Pattern Synthesis via Irregular Convolution. *Computer Graphics Forum 38*, 5 (2019), 109–122. doi:10.1111/cgf.13793. 2

[TRB*18] TONG C., ROBERTS R., BORGO R., WALTON S., LARAMEE R. S., WEGBA K., LU A., WANG Y., QU H., LUO Q., MA X.: Storytelling and visualization: An extended survey. *Information 9*, 3 (2018). doi:10.3390/info9030065. 3

[VGH*05] VIOLA I., GRÖLLER M. E., HADWIGER M., BÜHLER K., PREIM B., EBERT D.: Illustrative Visualization. In *Eurographics 2005 - Tutorials* (2005). doi:10.2312/egt.20051052. 2

[WB03] WHITBY M., BERRY D.: Molecular visualizations of dna. In *ACM SIGGRAPH 2003 Video Review on Electronic Theater Program on Electronic Theater Program - Volume 144* (2003), p. 4. doi:10.1145/1006032.1006036. 3

[WBCG09] WITHER J., BOUDON F., CANI M.-P., GODIN C.: Structure from silhouettes: a new paradigm for fast sketch-based design of trees. *Computer Graphics Forum 28*, 2 (2009), 541–550. doi:10.1111/j.1467-8659.2009.01394.x. 5, 10

[WLKT09] WEI L.-Y., LEFEBVRE S., KWATRA V., TURK G.: State of the Art in Example-based Texture Synthesis. In *Eurographics 2009 - State of the Art Reports* (2009). doi:10.2312/egst.20091063. 2

[WSRDL*00] WITKO-SARSAT V., RIEU P., DESCAMPS-LATSCHA B., LESAVRE P., HALBWACHS-MECARELLI L.: Neutrophils: Molecules, functions and pathophysiological aspects. *Laboratory Investigation 80* (2000), 617–653. doi:10.1038/labinvest.3780067. 3

[XKG*16] XING J., KAZI R. H., GROSSMAN T., WEI L.-Y., STAM J., FITZMAURICE G.: Energy-Brushes: Interactive Tools for Illustrating Stylized Elemental Dynamics. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (2016), p. 755–766. doi:10.1145/2984511.2984585. 3

[ZBQC13] ZANNI C., BERNHARDT A., QUIBLIER M., CANI M.-P.: SCALe-invariant Integral Surfaces. *Computer Graphics Forum 32*, 8 (2013), 219–232. doi:10.1111/cgf.12199. 2, 5

[ZIH*11] ZHU B., IWATA M., HARAGUCHI R., ASHIHARA T., UMETANI N., IGARASHI T., NAKAZAWA K.: Sketch-based dynamic illustration of fluid systems. *ACM Trans. Graph. 30*, 6 (dec 2011), 1–8. doi:10.1145/2070781.2024168. 3