Eurographics Conference on Visualization (EuroVis) 2026
N. Andrienko, M. Chen, and B. Wang
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 45 (2026), Number 3

Tubes or Ribbons? Comparing Texture-space Visualization for Multivariate Line Data

B. Russig[1] and R. Rzayev[2,3] and R. Dachselt[2,3] and S. Gumhold[1,3]

TUD Dresden University of Technology
[1]Chair of Computer Graphics and Visualization [2]Interactive Media Lab Dresden
[3]Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig

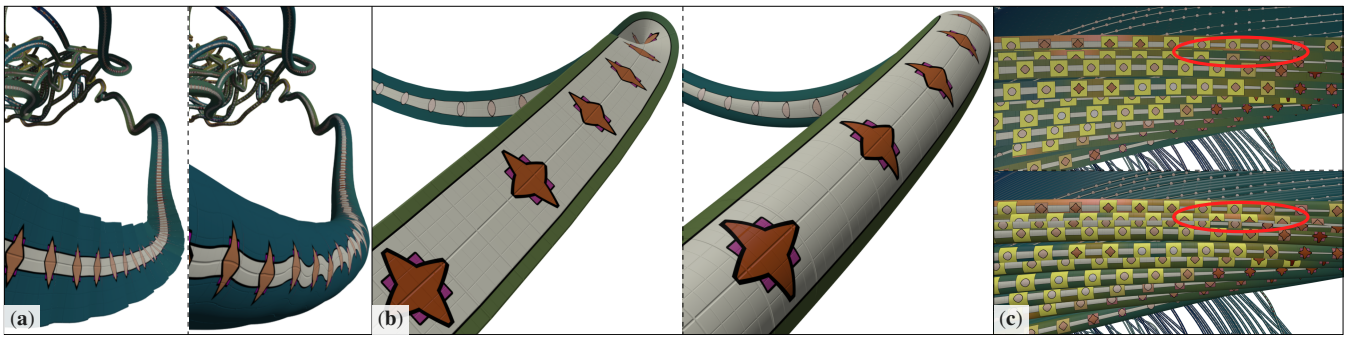

Figure 1: *Notable differences in the display of identical texture-space visualizations on the surfaces of view-aligned Ribbons (left) vs. Tubes (right): (**a**) **Tangent response**: ribbons respond to tangent changes by twisting, on tubes they mostly cause texture distortion. (**b**) **Self-occlusion** when viewing a curving path in-plane is less severe for ribbons. (**c**) **Bundle occlusion** is more likely to hide glyphs (red encirclement) when using ribbons (top), while the curving cross-section of tubes causes some parts to become visible that would otherwise be occluded (bottom). All images use the texture-space visualization of secondary variables that we use throughout this paper (section 5.2).*

Abstract

Multivariate line data is critical for analyzing flow fields, agent systems, and dynamic trajectories. Embedding secondary variables along spatial paths using surface-based primitives such as ribbons and circular tubes introduces challenges related to perspective, scale, and distortion. Perceptual trade-offs due to these challenges remain unclear. We address this gap through a controlled user study with 10 experts in computational fluid dynamics and visualization, performing four identical analysis tasks involving both spatial (requiring location-based relationships) and non-spatial (requiring attribute value comparisons) aspects. The tasks were performed using a prototype with interactivity limited to controlling the camera. While quantitative measures showed no significant performance differences between embedded visualizations on ribbons and tubes, we found a clear subjective preference for tubes among the study participants. Furthermore, their feedback indicates surface-based embeddings are generally helpful and should be utilized more often.

CCS Concepts

• ***Human-centered computing*** → ***Empirical studies in visualization;*** *Visualization techniques; Visualization design and evaluation methods;* • ***Applied computing*** → *Physical sciences and engineering;*

1. Introduction

Multivariate line data poses a unique visualization challenge, as it requires simultaneously conveying both spatial and non-spatial information while relying on a very limited number of visual variables. Typical examples include movement trajectories of cyber-physical systems, such as drones, which have complex internal states associated with each point in time, or streamlines of a flow field with associated secondary fields that all offer important context for interpreting the shape of the flow. Furthermore, when over-utilized, the natural visual variables of a 3D line primitive (most typically color, width/radius, and opacity) can potentially compromise spatial understanding. One way to tackle this challenge is to visualize secondary variables directly onto the primitive, if it has a surface.

Texture patterns have been successfully used to encode additional spatial information like twist around a tube by varying the surface parametrization [SGS05], or to represent non-spatial quantities by varying the pattern's density [RHD*06]. However, when using less abstract representations, such as diagrams or multivariate glyphs, the question of readability arises. Besides the size of

such visualizations relative to the extents of the dataset limiting their effectiveness to the near-field, they are directly affected by the shape of the line primitive through distortion and occlusion. In this paper, we compare the application of this technique to two of the most commonly utilized line primitives with a surface: ribbons – an open surface formed by sweeping an orthogonal line along a curve – and tubes with a circular cross-section.

Ribbons appear to be the obvious choice for hosting secondary visualizations, as their surface has no inherent curvature that causes distortions – an implicit assumption that we found reflected in the preference for ribbons throughout the literature on multivariate line data visualization (detailed in section 3). On the other hand, tubes provide a naturally 3-dimensional line representation that offers certain advantages for spatial perception in combination with lighting [ZSH96] and are thus in common use for dense, not necessarily multivariate line data such as diffusion tensor tractographs [CYA25, KJVC25] and 3D trajectories from diverse sources [BTD16, BLD21, LYR*23]. The open question we address in this paper is whether tubes can be used equally effectively for embedding visualizations of multiple secondary variables.

We tackle this question by first comparing both primitives in an analytical way, and second, conducting a comparative user study with 10 experts from engineering, visualization, and human-computer interaction. To ensure comparability, we use high-quality raycasted rendering, which required tuning an approach from the literature to produce better results in artifact-prone edge cases that frequently arise with view-aligned ribbons. With this, each expert solved 8 realistic multivariate analysis tasks on two streamline datasets, using high-fidelity ribbon- and tube-based line representations, with plot- and glyph-based embedded visualizations of secondary line variables. We found no evidence to support the common assumption that ribbons are better suited to hosting embedded visualizations. Instead, we observed a strong expert preference for embedding such visualizations on tubes.

To summarize, our contributions are:

- A structured, analytical discussion of relevant differences between view-aligned ribbons and tubes for displaying embedded visualizations,
- A tailored ribbon raycasting method with reduced artifacts,
- Qualitative and quantitative insights from an expert user study, indicating no statistical advantage for either primitive, but a strong user preference for tubes.

The remainder of this paper is structured as follows: In section 2 we define our notion of ribbons and tubes to specify our research question. Section 3 gives an overview of related work and its implications for our research question. Section 4 provides a systematic treatment of the differences between ribbons and tubes, and how they are reflected in our visualization. In section 5 we detail the design of the user study, the results of which we discuss in section 6. We conclude and identify possible avenues for further research in section 7.

2. Definitions and Goal

We focus on *line data* as the subject to be visualized. That is, sets of 1-dimensional manifolds that can be parametrized as curves $\underline{\boldsymbol{c}}_i(t) \in \mathbb{R}^3,\ t \in \mathbb{R}$. In this work, we represent lines as Hermite splines, i.e., each $\underline{\boldsymbol{c}}_i(t)$ is a C^1-continuous curve that cubically interpolates the position samples of line i. The required tangents can be heuristically determined (e.g., Catmull-Rom) or taken directly from the data (vector field, measured velocities, etc.).

This notion of line data separates the tubes and ribbons we study from the related concept of *stream tubes* and *stream ribbons*, which are tangent surfaces of a vector field [VS82], meaning their exact shape emerges directly from tracing a set of streamlines rather than resulting from a conscious choice by a visualization practitioner.

The tubes we consider are defined by sweeping a sphere of radius r along the curves $\underline{\boldsymbol{c}}_i(t)$, mirroring the definition introduced by Van Wijk [VW85]. Similarly, ribbons result from sweeping a line segment L of length w (the *width* of the ribbon) centered along $\underline{\boldsymbol{c}}_i(t)$. Unlike the fully symmetric sphere, reference directions $\hat{\boldsymbol{v}}_i(t)$ orthogonal to the tangent directions $\vec{\boldsymbol{c}}_i'(t)$ need to be chosen along which L is aligned at every t. Since we are going to show additional visualizations on the surface, we chose $\hat{\boldsymbol{v}}_i(t)$ such that the normal of the plane $\hat{\boldsymbol{v}}_i(t) \times \vec{\boldsymbol{c}}_i'(t)$ is maximally close to being parallel with the viewing direction; that is, the ribbon will be view-aligned (see section 4.1 for details). This is essentially the *Scalable Self-orienting Surface* introduced by Schussman and Ma [SM02].

Other typical uses of ribbons align $\hat{\boldsymbol{v}}_i(t)$ according to some spatial attribute, like the orientation of a mover or for showing twist (integrated angular velocity) along a path. A similar capability can be added to circular tubes by parameterizing their surface according to some varying reference frame and showing the orientation via texture. We do, however, not consider such uses in order to answer the following research question:

*When embedding visualizations of additional attributes on the surface of the line primitive, will these visualizations be more effective on **ribbons** than on **tubes**?*

Not keeping the visualization pointing towards the viewer at all times can reasonably be expected to reduce its effectiveness, so we ensure optimal view alignment. We consider studying the exact impact of keeping orientation as a free visual variable out-of-scope for this study. Similarly, allowing r or w to vary along lines opens up additional questions on how to embed the visualizations. Should they scale with the surface, or maintain a constant apparent size? We argue that eliminating radius/width as a free variable enables us to answer our research question more directly.

3. Related Work

In the following, we give an overview of line data visualization research with a special focus on ribbon and tube-based techniques, multivariate data and related studies.

Tubes and ribbons in line data visualization. A number of application domain-specific surveys [AA13, ZFH*22, BSD*24] provide an extensive overview of line data visualization techniques including but not limited to tubes and ribbons. We mention here notable examples. Their use dates back as far as the late 1980s [HH89, Vol89, BMP*90] in the context of visualizing vector field topology. Even earlier, seminal work in computer graphics formalized them as parametric primitives and derived rendering algorithms from there [BK85, VW85].

As graphics workstations and personal computing hardware became more powerful, the possibility of employing these prim-

itives in real-time interactive visualization systems renewed research interest, with significant work being done in improving perception through plausible lighting [ZSH96, EHS13]) and additional cues [MMYK06, EBRI09], as well as on increasingly efficient rendering techniques [SM02, SGS05, RBE*06, HGVV15] including transparency [KRW19, GG21, KNM*21, KJVC25]. Not least for their obvious uses in many areas of computer graphics, there is now a comprehensive set of real-time algorithms available that enable direct raycasting of large amounts of high-fidelity parametric tubes [RL18, RSG20, Yuk22] and ribbons [Res22, PJH23].

Embedded visualization of secondary line variables. Methods that visualize additional data dimensions directly on the surface of the line primitive exist, but they are rare. Ware et al. [WAPW06] represent whale movement with an oriented ribbon that displays additional information on its surface texture and extrude a sawtooth pattern from the center line encoding strength and direction of fluke stroke events. Tominski et al. [TSAA12] plot multiple quantities on trajectory ribbons; however, the trajectories are 2D, and the ribbons are formed by stacking the associated plots. This has proven to be a rather popular technique and was adopted in diverse trajectory-related fields [CTBH13, HMRH15, WESL*23]. Somewhat similarly, Staib et al. [SGG17] plot derived attributes of particle clusters onto flow ribbons [JR05]. Here, the ribbons follow an actual 3D trajectory through space. This limits the number of plots (they display up to 6), as flow ribbons should not be wider than the object they're attached to, and even then they're still quite wide (perhaps more so than is desirable for some applications). Sterzik et al. [SMCL23] studied the perceptual space of varying-density illustrative texture patterns, and show examples using their technique to encode scalar attributes on surfaces (including tubes representing line data). They explicitly mention opening additional channels for encoding information next to color (up to 2 with a cross hatching pattern), but their study does not treat the impact of curvature or perspective-induced distortions on the perceived densities. Russig et al. [RGDG23] draw layers of plots and glyphs directly onto the surface of a 3D tube, allowing for a large number of available visual attributes. They provide some mitigation for distortions introduced by the curved tube surface but cannot eliminate them completely. Neuhauser et al. [NHK*23] utilize this technique to show associated simulation attributes on convection lines, but limit the embedded visualization to multiple color bands, thereby avoiding most of the distortion issues.

Most often however, multivariate line data is visualized in a somewhat more abstract or aggregated way, or the non-spatial data is separated from the spatial view, sometimes linked back via suitable means. He et al. [HCC*19] provide the most recent survey of such methods. In general, there is a plethora of techniques to chose from when designing separate views for the non-spatial attributes [AMST11, FXJ20, ODNN25] as well as their linkage back to the spatial view [JE12].

Perceptual insights into multivariate line data visualization. A number of studies have looked into the effectiveness of conveying additional information embedded with the line primitive. Boukhelifa et al. [BBIF12] studied the perception of dashing/dotting, brightness, blurriness, and sketchiness (a local deformation emulating manual line art sketching) for representing uncertainty along lines. Sterzik et al. [SLK*22] study the use of these techniques (using width instead of blurriness) for showing uncertainty on molecular surfaces via (view-dependent) silhouette lines, as well as other stylized line representations [SMCL23, SLW*23]. More related to classical line data, Bae and Dodge [BD23] study the cognition of movement represented by 2D trajectory lines when using color and width to convey derived movement variables like speed.

All these studies have in common that they utilize the direct visual variables of lines only (color/illumination, shape, width/radius etc.), so we cannot derive insights about the comparative effectiveness of embedded visualizations between different line primitives from their results. They do, however, provide a good selection of task types, some of which we took some inspiration from (see section 5.3). Furthermore, Weigle and Banks [WB08] studied the spatial perception of complex 3D line sets under varying conditions in the context of illuminated stream tube tractograms. As our focus is on embedded visualizations, we opted not to include purely spatial tests. We do, however, observe their findings regarding the benefits of global illumination (see section 5.2).

Summary. When a detailed or complex visualization of additional attributes directly on the line primitive was attempted, ribbons have so far been the preferred choice, as is evident from the body of methods we looked at (e.g., [WAPW06, TSAA12, SGG17]). Only one method attempted to do the same specifically with tubes [RGDG23] which was adopted just once for a specific use-case [NHK*23]. We noticed that this tends to correlate with a heightened emphasis on the non-spatial data attributes, whereas other primitives (lines, e.g., [SMRB22], tubes, e.g., [BTD16] or tube-like shaded ribbons, e.g., [MMYK06]) are employed when the spatial aspects are most important. Then, typically just one or two additional attributes, which are often also space-related (velocity, acceleration etc.), are mapped to color, texture patterns and/or radius/width. As a side note, for conveying spatial orientation, textured tubes (e.g., [SGS05, RBE*06]) and oriented ribbons (e.g., [Web09, WNW*22]) appear to be equally popular.

So far, there is no direct comparison of the suitability of either primitive for embedding secondary visualizations. Intuition might suggest that ribbons are more suitable as they exhibit considerably less distortion on their surface, which could explain the bias towards ribbons in the body of methods that take this approach. This is however not the only relevant difference for embedded secondary visualizations as we will show in section 4, so closer study is warranted.

4. Ribbons versus Tubes

We identified notable effects caused by the geometric differences between view-aligned ribbons and tubes that influence the appearance of embedded visualizations. In this section, we describe each one in detail. First however, we need to establish how the surface of each primitive is parametrized.

4.1. Surface Parametrization and Rendering

The parametrization along the curve $\underline{\boldsymbol{c}}(t)$, for which we use the variable u, is the same for both primitives: it is simply the arc length along the curve using multiples of the half-width (ribbons) or radius (tubes) as unit, respectively. This allows for predictable placement of visualizations and eliminates parameterization-induced stretch-

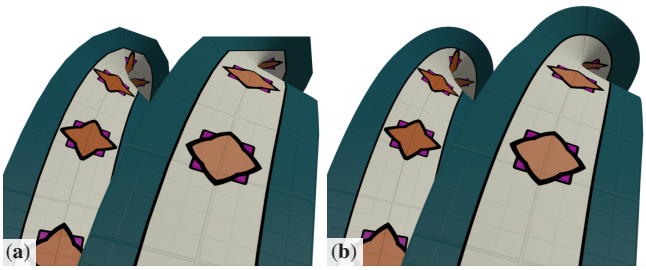

Figure 2: *Adaptation of the ray intersection method from (**a**) Pharr et al. [PJH23] with (**b**) our improvements for view-aligned ribbons. The ribbon shape not breaking down when viewed close to tangentially is important for this study, as this regime gives rise to most of the notable differences between ribbons and tubes.*

ing [RGDG23]. The orthogonal coordinate $v \in [0,1]$ depends on the viewing direction and works slightly differently for each primitive:

Ribbons. As the surface of a ribbon is flat at any given t, the v coordinate is a linear parametrization of the line L (as introduced in section 2) centered around $\underline{\boldsymbol{c}}(t)$ along the reference direction $\hat{\boldsymbol{v}}(t)$:

$$L_t(v) = \underline{\boldsymbol{c}}(t) + w\left(v - \frac{1}{2}\right)\hat{\boldsymbol{v}}(t) \tag{1}$$

where w is the ribbon width. Assuming all coordinates are in eye space, view alignment of $\hat{\boldsymbol{v}}(t)$ is accomplished via

$$\hat{\boldsymbol{v}}(t) \propto \vec{\boldsymbol{c}}'(t) \times \underline{\boldsymbol{c}}(t) \tag{2}$$

Tubes. Care is required when choosing a parametrization around the tube axis. The trivial choice of parametrizing according to angle around the axis (equivalent to using circular arc length) results in a perspective distortion that squishes surface contents close to the silhouette. To avoid this, we use the strategy by Russig et al. [RGDG23] that computes the v coordinate of a surface point $\underline{\boldsymbol{p}}_t$ on the tube with respect to the reference direction $\hat{\boldsymbol{v}}(t) \perp \vec{\boldsymbol{c}}'(t)$ using a simple dot product:

$$v = \langle \underline{\boldsymbol{p}}_t - \underline{\boldsymbol{c}}(t), \hat{\boldsymbol{v}}(t) \rangle \tag{3}$$

Rendering. The other important choice we have to make that will affect the primitives' appearance is how they are rendered. We chose raycasting for both to rule out effects of suboptimal tessellation. For tubes, we again use the implementation by Russig et al. [RGDG23] as it also provides all functionality for drawing visualizations onto the surfaces. For ribbons, we use the method proposed by Pharr et al. [PJH23] and adapt it to better support view alignment. The adaptations were necessary since the intended use case of the original method (hair, blades of grass, and similar) allowed for some performance trade-offs that result in a surface fidelity we found unsuitable for a visualization where curves would be viewed up-close and often tangentially. Also, the algorithm had to be ported to a fragment shader for real-time rendering.

On a high level, the ribbon intersection routine recursively subdivides parts of a cubic curve segment until either their bounding boxes are no longer intersected by the ray, or a sub-segment is sufficiently linear for an analytic intersection test. Our fragment shader implementation accomplishes this by emulating a stack in local memory and flattening the recursion into a loop. Beyond that, our adaptation comprises two changes designed to meet the required visual fidelity: (**a**) we keep subdividing until the angle between subsequent reference directions as per eq. (2) is below a threshold – this fixes breakdown of the ribbon when a curving segment is viewed tangentially. And (**b**), we use bilinear patches to compute the actual ray intersections instead of a simple linear approximation, causing the ribbon surface to twist more smoothly between even large subsequent angles. Figure 2 shows the visual impact of these changes.

Note that we cannot eliminate the singularity that forms when ribbons are viewed perfectly tangentially, but this rarely happens in practice, so these measures proved effective enough for our study. A detailed description of the rendering pipeline can be found in the appendix.

4.2. Differences and Effects

We identified 5 effects that impact embedded visualizations differently on ribbons and tubes. We are unaware of any prior systematic treatment of the geometric differences between these two line primitives. It is entirely possible that additional effects are at play, but after extensive examination, we are reasonably confident that these effects would be the major sources of observable differences in user performance. We characterize them as follows:

Base distortion. The most obvious difference is that circular tubes have an inherently curved surface, which causes a base distortion that is always present and which becomes increasingly visible as the viewing direction gets closer to the curve tangent direction $\vec{\boldsymbol{c}}'(t)$, since the tube cross section arc then stretches across more screen space. Ribbons, with their perfectly flat surface (when $\underline{\boldsymbol{c}}(t)$ is itself straight) do not have this problem. Note that this is not related to perspective distortions, which all 3D content is subject to even when using orthographic projection.

Tangent response. Changes in the tangent $\vec{\boldsymbol{c}}'(t)$ transfer to the view-aligned reference direction $\hat{\boldsymbol{v}}(t)$, and the strength of this influence is amplified as the view direction becomes more parallel. This manifests strongly in the presence of high-frequency geometric features in the curve, most typically caused by noise in the position samples. For ribbons, the resulting jiggle of the reference direction results in a noticeably wavy surface but induces very little distortion on the visualizations. For tubes, the effect on the geometry is much less noticeable, but the squiggly trace of the reference direction $\hat{\boldsymbol{v}}(t)$ causes visible distortions due to the view-aligned parametrization (cf. section 4.1). This effect is depicted in fig. 1a.

Self-occlusion. When $\underline{\boldsymbol{c}}(t)$ curves in some interval $[t_0, t_1]$, and the camera resides on the plane $\vec{\boldsymbol{c}}'(t_0) \times \vec{\boldsymbol{c}}'(t_1)$, then portions of the line primitive behind the apex of the curve will get occluded by the part before it. Here, the view alignment of the ribbon causes it to twist in a way that actually un-obscures part of the surface behind the apex. A tube will just twist its parametrization, leaving the geometry unaffected and the parts behind fully occluded (fig. 1b). Additionally, the flip of the reference direction $\hat{\boldsymbol{v}}(t)$ causes a massive distortion for tubes, but this happens close to the apex and is thus half-occluded and not very prominent.

Bundle Occlusion. Line data often contains structures with tight bundles. Unless the width/radius of the line primitive is adjusted to prevent it (which we explicitly rule out due to the reasons motivated

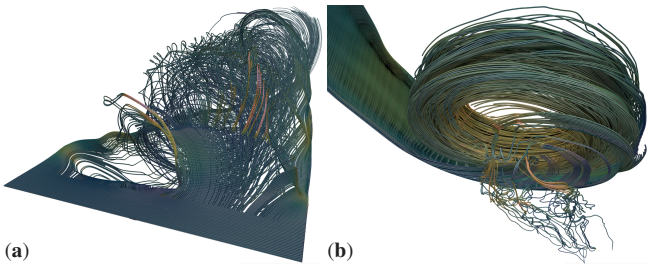

(**a**) (**b**)

Figure 3: *Overview rendering of both datasets with the same colormap (section 5.2) and embedded visualization we used in the study: (**a**) dataset 1 (propeller), (**b**) dataset 2 (fish-friendly turbine).*

in section 2), these regions will typically contain many overlapping lines. Here, their circular cross section can cause tubes situated behind others (from the perspective of the viewer) to partially end up in front anyway. Since ribbons are locally flat and in our case view-aligned, they are always either behind or in front of other ribbons at any given curve point t (and if close to the same distance, cause z-fighting and related effects). This means that visualizations on otherwise occluded tubes can still be (partially) readable even in such a situation, while on ribbons, they cannot (fig. 1c).

Radial Displacement. Any visualization on a tube surface will be displaced from the point on the line the underlying data attributes were sampled at, since tubes have non-zero radius. The surface of a ribbon on the other hand always lies on the curve it represents by construction, so there is no displacement at all. This effect is again most noticeable when viewing tangentially (cf. fig. 1a-b).

It is obvious that when studied in isolation, each of these effects is advantageous/disadvantageous for one primitive over the other. While there is value in quantifying each effect, it would not be possible to infer from that if one primitive works better than the other in practice. Instead, we believe the research question of this paper is best answered by an approach that remains agnostic to individual aspects and instead studies the net effect under a realistic workload.

5. Method

This section presents our user study, outlining the overall strategy and the major aspects of its design and procedure.

5.1. Overall Study Design and Strategy

Since our primary objective was to assess performance under realistic workload conditions, we decided to use real data rather than fully synthetic data that could be more tightly controlled but would lack ecological validity. Working with real data allows us to engage domain experts who can interpret the visualizations meaningfully and avoid the risk of biased feature selection. Consequently, the study was designed as an expert-focused evaluation, which limits the feasible sample size and reduces the statistical power to detect differences. Thus, we designed a within-subjects mixed-design study with the visual primitive as the independent variable, with levels ribbon and tube.

Participants performed a series of analysis tasks (section 5.3) while we measured the time and camera interactions needed to complete each task. The tasks were identical for both conditions, differing only in the visual primitive (ribbon or tube). By keeping every other aspect of the visualization constant, including the embedded texture-space representation, we isolated the independent variable. Because the same visual content is shown in both conditions, the study is confined to a single application domain. This trade-off avoids the need to redesign the visualization for multiple domains. To minimize learning effects, we employed two distinct datasets and randomized task order in each trial. We counterbalanced the order of the visual primitives and the assigned datasets using a Latin square. Consequently, each participant completed two trials (one ribbon and one tube), each with a different dataset, and across participants, all four possible primitive+dataset pairings were represented equally often. After each task, we asked a short set of qualitative questions, and after completing all tasks, each participant took part in a semi-structured interview. The interview probed perceived effectiveness of the embedded visualization, overall preference for ribbons versus tubes, and any domain-specific suggestions for improvement.

Data domain and individual datasets. We chose 3D *computational fluid dynamics* (referred to as CFD in the remainder of this paper) as the source of the line data. The reasons are threefold:

1. Fluid simulations in engineering typically produce vector fields with complex topology, which in turn results in line sets that form structures within which lines become correlated.
2. While there typically is a canonical "up" direction, the highly 3-dimensional nature of the structures does not favor a single view orientation and requires considerable camera movement by the user to fully grasp. We expect this to amplify the difference between tubes and ribbons, most of which are view-dependent.
3. The flow field is typically accompanied by several additional fields, such as temperature, pressure, and vorticity, among others, each contributing to the dynamic effects of interest to engineers (e.g., flow separation increases vorticity and involves decreasing velocity and increasing pressure). This makes flow field analysis an inherently multivariate problem.

In keeping with the goal of emulating realistic workloads, we used streamlines generated using *VTK* [SML06] from output of the CFD software *OpenFOAM* [Jas09] as test data. The simulation results were provided to us by domain experts and involved two distinct scenarios: **dataset 1** simulates the flow of water around a ship's propeller, and **dataset 2** simulates the flow of water through a fish-friendly turbine. A rendering of both datasets is provided in fig. 3.

Validation. To inform our attribute selection (section 5.2) and task design (section 5.3), we interviewed two domain experts before the main study. We elicited their use of flow field visualization tools and the types of insights they aim to extract in their current work. Expert 1 (henceforth referred to as *E1*) mostly deals with large-scale flood simulations which are typically modeled in 2 dimensions. Expert 2 (*E2*), a civil engineer, specifically pointed out identifying and delimiting instances of flow separation, for which they often combine stream or path lines around the object under study with (sometimes animated) scalar field visualizations to display flow velocity, pressure, and vorticity simultaneously. Both experts highlighted identifying instances and patterns as a frequent task in their visualization workflows. Based on these insights, we identified **finding instances of certain attribute combinations**

and gradients as a task well-suited to texture-space visualizations, considering our datasets. Besides, *finding instances* is a task that has been successfully used for studying line data visualizations [BD23]. Thus, we created visualizations for several attribute combinations that included flow velocity, pressure, and vorticity, and asked two experts for their preferences and feedback to inform our final design.

5.2. Visualization Design

Our visualization encodes four scalar secondary variables along streamlines: flow velocity, pressure, magnitude of vorticity, and angular velocity. All color scales we use were selected from the collection of *Scientific Color Maps* by Crameri and Shepard [Cra21]. We linearly map flow velocity $\|\vec{c}'(t)\|$ to surface color using the *batlow* color scale. The other variables are linearly mapped to shapes on top using the following encoding:

1. **Pressure** is represented by the *thickness* of a light gray continuous **line** in the center of the view-aligned texture space. Line thickness goes up to a width of 75% of the view-facing surface, ensuring the surface color is never fully obscured.
2. Magnitude of flow **vorticity** is double-coded by *side length* and *color* of a **square**. E2 mentioned that they often need the components of vorticity, but showing vectors in a spatially intuitive way embedded on a 2D surface of unpredictable orientation is not straightforward. They stated, however, that the magnitude is still useful, especially when combined with angular velocity and the streamline geometry. For color, we use the *buda* scale.
3. **Angular velocity** is double-coded by *shape* and *color* (using the diverging *vik* scale) of a **sign glyph** [RGDG23] that morphs between a minus (counter-clockwise), circle (neutral or 0), and plus (clockwise) shape. In CFD, angular velocity is defined as the scalar component of the vorticity vector in flow direction divided by two. It thus combines length and flow parallelism of the vorticity vector into a signed scalar. Notably, this means wherever flow and vorticity are close to orthogonal in a dataset with many long vorticity vectors (like our two datasets), angular velocity can be "high" (mapped to large shape/color values) where vorticity is mapped to "small" squares with data windowing tuned to show the fluctuation of high vorticity levels.

All shapes utilize a black outline to help visually separate them from their surroundings. This is required to mitigate occurrences of similar colors, which cannot be ruled out completely, although we carefully selected color scales with limited hue overlap.

For the discrete mappings of **vorticity** and **angular velocity**, we chose shapes that maximize visibility of their defining features: Plus, neutral, and minus forms of the sign glyph ensure that the corners of the underlying square remain visible when both glyphs are placed at the same location. The minimum side length of the square is chosen such that it can only be completely occluded when vorticity is close to 0. Since we always sample them at identical locations, they form a compound glyph displaying vorticity and angular velocity simultaneously. Both experts preferred this integrated design over separate placement. Figure 4 shows legends for each encoding. Finally, we superimpose a subtle reference grid via normal mapping to help users mentally rectify the embedded visualizations even when distorted by a curving surface (cf. figs. 1 and 2).

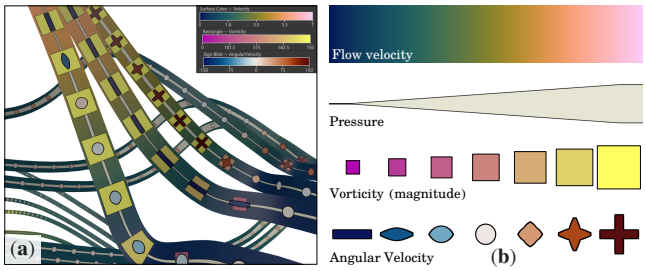

Figure 4: *The embedded visualization we designed with domain experts. (**a**) A closeup of the visualization applied to dataset 1. (**b**) The full range of the visual variables for each attribute.*

Lighting. Lighting enhances the structural perception of dense line sets [ZSH96], with global illumination contributing valuable depth and other structural cues [WB08, EHS13]. We attach a directional light to the camera so that view-aligned surfaces are directly illuminated, while most illumination is provided from ambient light, approximated using voxel-cone tracing for real-time ambient occlusion [CNS*11]. We include these effects not only for their perceptual benefits but also because they are widely used in modern line data visualization and represent realistic rendering conditions.

5.3. User Task Design

We measure participants performing 8 tasks in total, 4 distinct tasks per condition. Each task requires finding an instance of a particular attribute constellation (henceforth referred to as *instance type*) as motivated in section 5.1. We use 4 instance types in total (named *A*, *B*, *C*, and *D*), so each must be found once for each dataset. The attribute constellations we used for each instance type were chosen with semantic relevance as a first priority, but some compromises were required to get feasible tasks.

Firstly, we did not consider large-scale features, as the embedded visualizations become unreadable at some point when zooming out. Secondly, we restricted ourselves to constellations with at least 4 clearly recognizable instances in both datasets. The task descriptions for the 4 instance types are as follows:

A. "Find an instance on a line where *angular velocity* quickly changes its sign, i.e. from ***+positive*** (clockwise) to ***-negative*** (counter-clockwise) or vice-versa."
We verbally informed the participants that the flip of sign should occur within 3 samples, i.e. there should be no more than 1 neutral value in between.

B. "Find an instance where *pressure* is ***low*** and *angular velocity* is maximally ***+positive*** (clockwise)."

C. "Find an instance where *pressure* is ***low***, *vorticity* is ***high***, and *angular velocity* is ***neutral o***."

D. "Find an instance where *angular velocity* on neighboring, mostly parallel lines has **opposite *sign +/-***."

Types *B* and *C* involve 2 and 3 variables, respectively, and test only the readability of individual glyphs and the pressure plot. *A* and *D* ask for just a single variable, but involve spatial aspects as well; a change along a single line for *A*, and between close-by lines for *D*. We did not include flow velocity and thus surface color in any of the tasks, as the color of lines can be seen very clearly even from

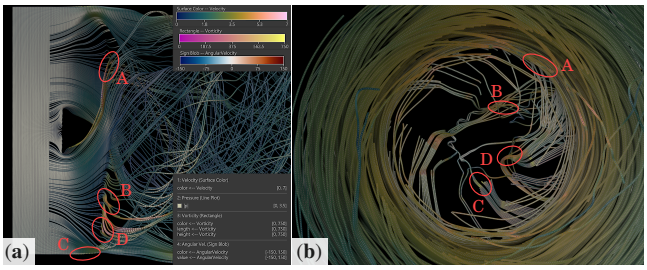

Figure 5: *The common initial viewpoints for all tasks on (**a**) dataset 1 (including the legends we show to participants) and (**b**) dataset 2, respectively. Possible instances for each task type (A, B, C and D) are always already within the initial view; the locations of one example each are marked by a red circle. Uncropped screenshots as seen by participants can be found in the supplementary material.*

far away for both line primitives. We still included it in the visualization since line color is one visual variable that is very likely to be used one way or another in practice, so we want to include any influence it might have on other embedded visualizations. We also created 4 similar tasks for training that all participants perform, two per dataset, using 4 instance types not used outside training. Finally, we used just one starting viewpoint per dataset that is shared across tasks to eliminate this as a confounding variable. These viewpoints were chosen such that instances of all types are already inside the initial view frustum; a fact all participants were informed of. Both initial viewpoints are shown in fig. 5.

The readability of the visualizations afforded by each primitive will come into play primarily when participants decide whether a variable exhibits the desired behavior or value range, typically when zoomed in to some extent. However, it will also be important how well the primitives preserve distinctive visual features of the embedded visualization at a distance.

5.4. Measurements

During the study, we recorded objective performance metrics and gathered subjective data via between-task questionnaires and a final semi-structured interview. Task performance was measured by completion time and the number of camera interactions, including orbiting (2 DoF/degrees of freedom), panning (2 DoF), and zooming. Interactions are performed by dragging the mouse with left (*orbiting*) or right (*panning*) buttons pressed, and scrolling the mouse wheel (*zooming*).

We measure the cumulative amount for each interaction type; that is, we sum up the magnitudes of all individual interactions. Due to their nature, all types use different units of measurement:

orbit	degrees around rotational axis
zooming	% distance to orbital focus
panning	% screen width

Zooming sensitivity was tuned such that one wheel tick zooms by approximately 22%. All raw data we recorded is available in the supplementary material.

After completing each task, participants answered four statements using a 7-point Likert scale, where 1 represented "strongly disagree" and 7 represented "strongly agree." The statements were: (1) "The task was difficult to complete." (2) "The information presented in the visualization was complex." (3) "It was easy to interpret the visualization for this task." (4) "I am satisfied with the visualization for this task." Participants could also indicate any issues or challenges they faced during the task. Finally, after finishing all tasks, we conducted a semi-structured interview to gather qualitative feedback. During the interview, participants were free to explore both datasets and switch the line display between ribbons and tubes at their discretion. The interview questions focused on comparative preferences for ribbons and tubes, the effectiveness of each in solving tasks, strategies used with both visual primitives for tasks involving line/trajectory or vector/tensor field data, typical work goals, background influences, and suggestions for improvement.

5.5. Experimental Setup

Apparatus. We run the trials on a graphics workstation with enough power (Intel Xeon E3-1270v3 CPU with 64GB RAM, NVIDIA Quadro RTX 5000), providing sustained rendering performance above 60 frames per second at full HD (1920×1080) using the ray-casted primitives introduced in section 4.1. The machine was connected to two 25” monitors: one displayed the visualization in full screen for the participant, while the other showed trial controls for the experimenter. During each task, participants freely manipulated the camera on their screen using the mouse.

Procedure. After welcoming the participants, we explained the goal of the study and asked them to sign an informed consent form. We then collected demographic data and asked participants about their personal experiences with line data. Next, participants were introduced to the visualization tools and camera controls. The researcher provided a detailed explanation of the visualization interface, including how to navigate using the camera controls. Participants were then given time to explore and practice interacting with the visualization freely. To further acclimate participants to the visualization and interaction techniques, they completed four training tasks. These tasks were divided evenly across two datasets, with two tasks per dataset. Once the participants felt comfortable with the visualization and interaction options, we proceeded to the main part of the study.

Participants solved tasks by starting with either a ribbon (4 tasks in random order) as a visual primitive, then switching to a tube (4 tasks in random order) with different datasets, or vice versa. Before beginning each task, the camera was reset to the starting location. Each task was displayed on a second screen, allowing participants to read the instructions carefully and ask the experimenter for clarification on any points they were unsure about. When a participant indicated they were ready, the experimenter started the task, which initiated the timing and recording of camera interactions. Participants were encouraged to think aloud at all times, especially when they believed they had found a relevant instance. If the instance met the criteria, the experimenter would stop the task, thereby saving the elapsed time and all captured data. If the instance did not satisfy the criteria, the participant had to continue searching until reaching the maximum time limit of 5 minutes, at which point the task would be counted as failed. However, none of the partic-

ipants experienced this; all found a valid instance within the time limit, typically in under a minute (see section 6.1). After each task, participants answered subjective qualitative questions. Finally, we conducted the semi-structured interview. Throughout the tasks and during the interviews, we used a smartphone to record audio for later analysis.

5.6. Participants

The study participants were 10 experts (9 male, 1 female) from engineering and visualization research groups at our university, all experienced with line/trajectory data. Their ages varied from 28 to 38 years (M = 32.2, SD = 3.3), with eight holding master's degrees and two Ph.D.s. Their expertise included mechanical and civil engineering, computer graphics, and human-computer interaction. None reported color vision deficiencies.

Prior experience. Participants' professional experience with line data included vector field analysis with *ParaView* (4 participants) and *PyVista* (2), 3D particle trajectory analysis using *sciview* (1) or custom tooling (1), trajectory analysis of biological cells using *Mastodon* (1) or custom tooling (1), and VR/AR user pose trajectory analysis using *Unity*-based (1) or custom tooling (1). Participants often obtained line and trajectory data from CFD simulations, vector fields, human and VR movement tracking, medical imaging (e.g., blood vessels), and cell tracking. They used visualizations to derive insights into movement patterns, flow behavior (e.g., turbulence and streamlines), high-velocity regions, trajectory inconsistencies, brain connectivity, and tissue development. The most frequently reported challenges were occlusions and data overload (4 mentions), software and technical limitations (4), data management and integration issues (3), and concerns about visualization clarity and techniques (2).

6. Results and Discussion

This section presents the results of our study, including statistical analyses, participants' qualitative feedback, and a discussion of the findings in the context of our research questions.

6.1. Statistical Analysis

All ten participants completed both experimental conditions, each comprising four tasks. We performed quantitative analyses on both objective and subjective data. For the subjective measures, scores for the two negatively worded items were reversed so that higher scores consistently indicated better outcomes (i.e., "The task was difficult to complete" and "The information presented in the visualization was complex"). The study primarily aimed to compare ribbons and tubes as line primitives for embedding visualizations of additional surface attributes. The tasks were designed to be diverse, allowing participants to engage with and gain experience using both primitives. Accordingly, our analysis focuses on the main effects of the line primitive and its interactions with task type. Complete descriptive statistics are provided in table 1, while detailed statistical analyses, including possible main effects of tasks, are presented in the appendix.

For the analysis, we examined whether the tasks produced any interaction effects with the line primitives. A repeated-measures analysis of variance (ANOVA) was conducted on all subjective and objective measures across combinations of line primitives and tasks. We found a statistically significant interaction between tasks and line primitives only for the statement, "The information presented in the visualization was complex," $F(3, 27) = 5.695$, $p = 0.004$, $\eta_p^2 = 0.388$. For tasks A and B, participants perceived the visualizations as equally complex. In task C, ribbons were rated as more complex, while in task D, tubes were easier (see table 1). However, post-hoc tests revealed no statistically significant differences. Furthermore, we found no other statistically significant interaction effects between tasks and line primitives, nor any main effects of line primitives on the subjective quantitative measures (see table 1 in the appendix for full statistics). Likewise, the objective measurements showed no significant interaction effects or main effects of the line primitives (see table 2 in the appendix for full statistics).

Given that tasks were spatial or non-spatial, we examined whether interactions existed between the spatial/non-spatial nature of tasks and the type of line primitive. We averaged scores for each line primitive and participant across spatial (A and D) and non-spatial (B and C) tasks, then conducted repeated-measures ANOVAs. For subjective measures, a statistically significant interaction effect between task spatiality and line primitive was found for ratings of task difficulty ($F(1, 9) = 5.12$, $p = 0.05$, $\eta_p^2 = 0.363$) and for perceived visualization complexity ($F(1, 9) = 5.192$, $p = 0.049$, $\eta_p^2 = 0.366$). In non-spatial tasks, tasks with ribbons were perceived as slightly less difficult to complete (M = 5.35, SD = 1.56) than the ones with tubes (M = 4.90, SD = 1.56), whereas for spatial tasks, the opposite pattern was observed (M = 4.05, SD = 2.13 and M = 5.15, SD = 0.94 for ribbons and tubes, respectively). For complexity, ribbons were seen as slightly less complex than tubes in non-spatial tasks (M = 4.05, SD = 1.83 for ribbons; M = 4.3, SD = 1.81 for tubes), but in spatial tasks, the pattern was reversed (M = 4.25, SD = 1.92 for ribbons; M = 3.65, SD = 1.80 for tubes). However, for both outcomes, post hoc pairwise comparisons using paired-samples t-tests with Bonferroni correction did not yield statistically significant differences (all $p > 0.25$ for task difficulty and all $p > 0.3$ for complexity). We found no other statistically significant interaction effects between tasks and line primitives, nor any main effects of line primitives (see table 3 in the appendix for full statistics).

6.2. Qualitative Feedback

The semi-structured interviews yielded a number of relevant insights, which we organized into four categories. Note that no participant was informed of our analysis regarding the differences between view-aligned ribbons and tubes (cf. section 4.2). Whenever they mentioned one of them, they noticed the effect independently. When referring to individual participants, we use a capital **P** followed by the sequential number we assigned to each participant.

Ribbon/tube preference. Feedback collected during the interviews reveals a clear subjective preference for tubes over view-aligned ribbons: 6 participants stated preferring tubes, only 1 preferred the view-aligned ribbons – the rest were undecided. Participants provided the following reasons **in favor of tubes**:

- *Structural perception* – **P3**: "I just get a better sense of depth from the tubes." **P8**: "To me, the sense of depth from the shading

Table 1: *Subjective and objective measurements reported as Mean (SD) for Ribbon and Tube across tasks A, B, C, and D. The scores for difficult and complex are inverted (i). The subjective measures are on a Likert scale from 1 to 7.*

		Subjective (Mean (SD))				Objective (Mean (SD))			
		difficult (i)	complex (i)	easy	satisfying	orbiting, °	panning, %	zooming, %	task time, s
Ribbon	A	4.5 (2.321)	3.9 (2.079)	5.1 (1.524)	5.4 (1.713)	337.40 (559.00)	179.78 (275.81)	604.9 (344.3)	44.10 (28.587)
	B	5.6 (1.506)	4.5 (2.014)	5.9 (1.197)	5.9 (1.287)	30.90 (48.49)	46.33 (60.46)	379.1 (428.0)	18.98 (11.158)
	C	5.1 (2.025)	3.6 (2.066)	5.7 (1.337)	5.6 (1.506)	92.20 (210.51)	147.58 (242.25)	650.7 (470.5)	28.10 (15.851)
	D	3.6 (2.221)	4.6 (2.066)	5.6 (1.174)	5.7 (1.252)	467.67 (637.44)	442.73 (521.54)	1206.2 (966.8)	74.54 (58.026)
Tube	A	6.1 (0.876)	4.0 (1.944)	6.4 (0.516)	6.0 (1.155)	149.68 (226.14)	111.93 (137.84)	713.3 (978.3)	30.14 (17.585)
	B	5.2 (2.044)	4.4 (2.011)	6.2 (0.632)	5.9 (0.738)	20.36 (36.54)	76.58 (78.36)	433.9 (318.7)	19.40 (7.612)
	C	4.6 (1.713)	4.2 (1.989)	5.7 (1.160)	5.8 (1.229)	187.39 (290.78)	109.67 (168.91)	703.9 (614.0)	34.39 (22.253)
	D	4.2 (1.687)	3.3 (1.829)	5.5 (1.354)	5.9 (0.994)	423.60 (520.85)	249.12 (184.26)	1401.7 (1785.2)	65.25 (48.235)

that the tubes have is just really helpful." **P9**: "The ambient occlusion effect is more visible on them, they convey space better."

- *Bundle occlusion* – **P2**: "If they're super close, it's much less distracting with tubes, less occlusion", **P8**: "Especially here, where lines are close by, it just looks clearer."
- *Color perception* – **P4**: "From far away, I think it's easier to observe the color of the tube."
- *Geometric consistency* – **P9**: "I find it visually easier to digest when the tubes are static, and only the visualization moves around the surface.", **P10**: "I find it easier if you don't have to mentally undo their orientation, which here has no actual meaning."

As for **preferring ribbons**, **P1** again gave *color perception* as the reason: "I like that for the ribbons, there is just one level of brightness, it looks flatter." When queried if they thought that the type of primitive affected their performance, only **P10** felt a negative impact from the view-aligned ribbon, while 3 participants (**P2**, **P4**, **P9**) felt no difference. The others were unsure.

Task strategies. The common strategy mentioned by all participants was to first look for characteristic color patterns, e.g., **P4**: "I found color is the most important thing [...] then the grey line, it was easy to find." – **P5**: "I mainly went after colors, like blue or red for angular velocity, or high vorticity, which is neither blue nor red." – **P6**: "I first checked whether something catches my eye, the yellow and red from the boxes and pluses popped out from afar." Then they looked at candidate locations in sequence, sometimes reassuming an overview position in between to regain bearings, e.g., **P6**: "I went back out since you said all instances are somewhere in view from here." **P1** additionally reported using domain knowledge to prioritize likely areas with the help of attributes not currently asked for: "Since I knew there is a pressure difference there, I looked there first." When queried, no participant reported changing their approach based on the line primitive. Consistent with this, only 3 participants even noticed during the study that ribbons and tubes were switched with the dataset.

General feedback on the visualization. All participants expressed liking the visualization overall, but raised notable points which we summarize as follows (we forego quotes for brevity; more relevant quotes are compiled in table 5 in the appendix):

- *Glyph readability*: 9 of the 10 participants reported being able to read the embedded visualizations well, particularly praising that it was easy to discern areas of high pressure and high vorticity, even with surface color prominently taken up by flow velocity. **P1** and **P3** mentioned that low vorticity (small pink square) with negative angular velocity (blue minus bar) tends to look like a red plus sign from afar, requiring more zooming, and **P5** found the minus form more difficult to recognize than the plus form. **P10**, the only participant to report some difficulty reading the visualization, would have preferred alternating separate glyphs.
- *Filtering*: **P2**, **P4**, **P6**, and **P7** mentioned they would rather change visualization settings to hide variables they deem unimportant at any given moment. All participants also highlighted the need for data filtering in a real application.
- *Focus+context*: **P7** found it hard to keep a mental map of where they were when zoomed in and would have liked a context view.
- *View alignment*: All participants liked the view alignment of embedded visualizations for tubes. For ribbons, however, participants **P7** and **P10** expressed a preference for the ribbon orientation to be derived from the data instead.

Applicability. We asked participants whether embedding secondary visualizations of this kind could be useful in their own work, and eight answered positively. Five participants specified the number of additional variables they would consider: **P2** and **P6** mentioned three, **P5** and **P8** up to four, and **P7** up to five. Among those who responded positively, five participants (**P1**, **P2**, **P5**, **P6**, **P7**) indicated they would use embedded secondary visualizations if supported in their existing tools, while three (**P3**, **P8**, **P9**) would consider integrating a similar approach into their custom tooling.

6.3. Discussion

Our quantitative results did not indicate a clear advantage of either primitive for the effectiveness of the embedded visualizations. We observed some statistically significant interaction effects for the subjective measures. However, post-hoc tests did not reveal any significant pairwise differences. Given the relatively small sample size of ten participants, these findings should be interpreted with caution, as the statistical power to detect subtle effects is limited. In interpreting the results, we focus primarily on patterns that align with qualitative feedback gathered from the semi-structured interviews and our observations, which provided additional insights and guidance for future research. In the following, we discuss the main takeaways from these findings.

Tubes preference. While several participants provided subjective reasons for preferring tubes, only **P10** explicitly noted greater difficulty when solving tasks with view-aligned ribbons, directly attributing this to the view alignment itself (see below). Furthermore, participant comments regarding bundle occlusion (cf. section 4.2, 6.2) may relate to the increased subjective difficulty observed for spatial tasks involving ribbons (cf. section 6.1), but more data would be required to reduce uncertainty surrounding this effect. The only preference expressed for ribbons (by **P1**) was justified by the flatter appearance of lighting. Interestingly, this contrasts with **P4**'s observation that colors were easier to distinguish on tubes, suggesting that the perceptual role of lighting warrants further investigation.

Finally, **no participant mentioned texture-space distortions** induced by the inherent curvature of tubes and their tangent response (cf. section 4.2). This is surprising and leads us to hypothesize that in general, users tolerate them quite well – at least if the tasks do not require reading exact values.

View alignment. Two participants (**P9**, **P10**) mentioned the geometric consistency of tubes as a reason they preferred them. **P10**, in particular, noted they had to mentally undo the orientation of the ribbon as it carried no meaning. Although this was a minority opinion, we note that it constitutes a valid concern, especially when producing still images, where it will not be obvious that the ribbon orientation is the result of view alignment. For tubes on the other hand, no participant reported being distracted by the visualization moving around the tube surface. The higher number of orbit and panning interactions observed for ribbons (c.f. table 1) may be related to this effect, though it did not reach statistical significance for our sample size (see table 2 in the appendix).

Practical implications. Our qualitative results suggest that when choosing a primitive for visualizing line data with embedded visualizations, one should go with tubes. However, our quantitative findings indicate that, given the minimal impact on user performance, this choice should primarily depend on other practical considerations. We think the main takeaway for practitioners is that any tube-based line data visualization can be extended with embedded visualizations, and there is no need to switch out the line primitive for this purpose. This is especially relevant for domains like CFD and tractography, where tubes have become the de facto standard.

Furthermore, the positive reactions to the embedded visualization demonstrate that encoding four additional variables this way is comprehensible to users with very little training. Finally, self-reported task strategies suggest that double-coding with color is generally beneficial for embedded visualizations, enhancing their effectiveness at larger distances.

Limitations. The small sample size of 10 experts limits the statistical power of our study, which is compounded by random error introduced through search and navigation behavior, as reflected in the large standard deviations of the objective measures in table 1. Increasing the certainty of our measurements would require a much larger study, which may not be feasible when relying solely on expert participants. Furthermore, we cannot determine from the data we collected how much of our experts' subjective preference for tubes is related to their background and experience (although accommodating expectations from domain experts is certainly advisable either way).

Lastly, our study does not quantify relative visualization readability directly. There likely are nuanced effects depending on a variety of factors that amplify or reduce the impact of the differences between the two primitives, like local line geometry (flat vs. curved), primitive parameters like width/radius, or prevalence of tangential viewing angles in the application. These effects appear to be offsetting each other in our end-to-end treatment, but there is value in studying their exact impact on readability for both line primitives, thus better informing how to factor in data characteristics and expected usage of the visualization into the choice of primitive.

7. Conclusion and Future Work

We compared embedded visualizations for multivariate line data, where secondary attributes are visualized on the surface of view-aligned ribbons or tubes. Besides contributing a theoretical analysis of their differences relevant for embedded visualizations, we conducted a study with expert participants. We found that for the task of visually identifying regions with specific high/low combinations of secondary variable values, tubes work at least as well as ribbons, contrary to our expectations. Our participants' subjective preferences, on the other hand, even strongly favor tubes. In sum, our work indicates that tube-embedded visualizations are a viable and interesting alternative for multivariate line data.

Further research is required to test if these findings hold and generalize to other kinds of multivariate line data and also to different analysis tasks. Furthermore, participants raised interesting questions worth exploring further: Lighting remains indispensable for spatial perception, but opposing views expressed by participants suggest that fine-tuning of lighting parameters and their influence on the perception of embedded visualizations is required. This in turn needs to be informed by a focused perceptual study. Similarly, a large-scale perceptual study could attempt to directly measure the readability of embedded visualizations under varying conditions, helping practitioners tune their designs for specific situations.

Finally, we received positive feedback about the usefulness of embedding visualizations of secondary variables: 8 out of 10 participants were interested in using this technique in their own work with 3D line data. Therefore, we encourage visualization practitioners to look at the surfaces of their 3D line primitives as a canvas that can and should be put to use for rich and expressive visualizations.

Acknowledgments

We thank the anonymous reviewers for their valuable input. This work was funded in part by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) via DFG grant 389792660 as part of TRR 248 (see https://perspicuous-computing.science) and as part of Germany's Excellence Strategy – EXC 2050/2 – Project ID 390696704 – Cluster of Excellence "Centre for Tactile Internet with Human-in-the-Loop" (CeTI) of Technische Universität Dresden. The authors acknowledge the financial support by the German Federal Ministry of Research, Technology and Space and by Sächsisches Staatsministerium für Wissenschaft, Kultur und Tourismus in the program Center of Excellence for AI-research "Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig", project ID: ScaDS.AI.

References

[AA13] ANDRIENKO N., ANDRIENKO G.: Visual analytics of movement: An overview of methods, tools and procedures. *Information Visualization 12*, 1 (Jan. 2013), 3–24. doi:10.1177/1473871612457601. 2

[AMST11] AIGNER W., MIKSCH S., SCHUMANN H., TOMINSKI C.: *Visualization of Time-Oriented Data*. Human-Computer Interaction Series. Springer London, London, 2011. doi:10.1007/978-0-85729-079-3. 3

[BBIF12] BOUKHELIFA N., BEZERIANOS A., ISENBERG T., FEKETE J.-D.: Evaluating Sketchiness as a Visual Variable for the Depiction of Qualitative Uncertainty. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (Dec. 2012), 2769–2778. doi:10.1109/TVCG.2012.220. 3

[BD23] BAE C. J., DODGE S.: Assessing the cognition of movement trajectory visualizations: interpreting speed and direction. *Cartography and Geographic Information Science 50*, 2 (Mar. 2023), 143–161. doi:10.1080/15230406.2022.2157879. 3, 6

[BK85] BRONSVOORT W. F., KLOK F.: Ray tracing generalized cylinders. *ACM Transactions on Graphics 4*, 4 (Oct. 1985), 291–303. doi:10.1145/6116.6118. 2

[BLD21] BÜSCHEL W., LEHMANN A., DACHSELT R.: MIRIA: A Mixed Reality Toolkit for the In-Situ Visualization and Analysis of Spatio-Temporal Interaction Data. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama Japan, May 2021), ACM, pp. 1–15. doi:10.1145/3411764.3445651. 2

[BMP*90] BANCROFT G., MERRITT F., PLESSEL T., KELAITA P., MCCABE R., GLOBUS A.: FAST: a multi-processed environment for visualization of computational fluid dynamics. In *Proceedings of the First IEEE Conference on Visualization: Visualization '90* (San Francisco, CA, USA, 1990), IEEE Comput. Soc. Press, pp. 14–27,. doi:10.1109/VISUAL.1990.146360. 2

[BSD*24] BELGHIT H., SPIVAK M., DAUCHEZ M., BAADEN M., JONQUET-PREVOTEAU J.: From complex data to clear insights: visualizing molecular dynamics trajectories. *Frontiers in Bioinformatics 4* (Apr. 2024), 1356659. doi:10.3389/fbinf.2024.1356659. 2

[BTD16] BUSCHMANN S., TRAPP M., DÖLLNER J.: Animated visualization of spatial–temporal trajectory data for air-traffic analysis. *The Visual Computer 32*, 3 (Mar. 2016), 371–381. doi:10.1007/s00371-015-1185-9. 2, 3

[CNS*11] CRASSIN C., NEYRET F., SAINZ M., GREEN S., EISEMANN E.: Interactive indirect illumination using voxel cone tracing: a preview. In *Symposium on Interactive 3D Graphics and Games* (San Francisco California, Feb. 2011), ACM, pp. 207–207. doi:10.1145/1944745.1944787. 6

[Cra21] CRAMERI F.: Scientific colour maps, Sept. 2021. doi:10.5281/ZENODO.5501399. 6

[CTBH13] CHENG T., TANAKSARANOND G., BRUNSDON C., HAWORTH J.: Exploratory visualisation of congestion evolutions on urban transport networks. *Transportation Research Part C: Emerging Technologies 36* (Nov. 2013), 296–306. doi:10.1016/j.trc.2013.09.001. 3

[CYA25] CHAMBERLAND M., YANG J. Y.-M., AYDOGAN D. B.: Real-time tractography: computation and visualization. *Brain Structure and Function 230*, 5 (May 2025), 62. doi:10.1007/s00429-025-02928-2. 2

[EBRI09] EVERTS M., BEKKER H., ROERDINK J., ISENBERG T.: Depth-Dependent Halos: Illustrative Rendering of Dense Line Data. *IEEE Transactions on Visualization and Computer Graphics 15*, 6 (Nov. 2009), 1299–1306. doi:10.1109/TVCG.2009.138. 3

[EHS13] EICHELBAUM S., HLAWITSCHKA M., SCHEUERMANN G.: LineAO—Improved Three-Dimensional Line Rendering. *IEEE Transactions on Visualization and Computer Graphics 19*, 3 (Mar. 2013), 433–445. doi:10.1109/TVCG.2012.142. 3, 6

[FXJ20] FANG Y., XU H., JIANG J.: A Survey of Time Series Data Visualization Research. *IOP Conference Series: Materials Science and Engineering 782*, 2 (Mar. 2020), 022013. doi:10.1088/1757-899X/782/2/022013. 3

[GG21] GROSS D., GUMHOLD S.: Advanced Rendering of Line Data with Ambient Occlusion and Transparency. *IEEE Transactions on Visualization and Computer Graphics 27*, 2 (Feb. 2021), 614–624. doi:10.1109/TVCG.2020.3028954. 3

[HCC*19] HE J., CHEN H., CHEN Y., TANG X., ZOU Y.: Diverse Visualization Techniques and Methods of Moving-Object-Trajectory Data: A Review. *ISPRS International Journal of Geo-Information 8*, 2 (Jan. 2019), 63. doi:10.3390/ijgi8020063. 3

[HGVV15] HERMOSILLA P., GUALLAR V., VINACUA A., VÁZQUEZ P.-P.: Instant Visualization of Secondary Structures of Molecular Models. In *VCBM '15: Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine* (2015), The Eurographics Association. ISBN: 9783905674828, ISSN: 2070-5786. URL: https://diglib.eg.org/handle/10.2312/vcbm20151208, doi:10.2312/VCBM.20151208. 3

[HH89] HELMAN J., HESSELINK L.: Representation and display of vector field topology in fluid flow data sets. *Computer 22*, 8 (Aug. 1989), 27–36. doi:10.1109/2.35197. 2

[HMRH15] HAB K., MIDDEL A., RUDDELL B. L., HAGEN H.: TraVis - A visualization framework for mobile transect data sets in an urban microclimate context. In *2015 IEEE Pacific Visualization Symposium (PacificVis)* (Hangzhou, China, Apr. 2015), IEEE, pp. 167–174. doi:10.1109/PACIFICVIS.2015.7156374. 3

[Jas09] JASAK H.: OpenFOAM: Open source CFD in research and industry. *International Journal of Naval Architecture and Ocean Engineering 1*, 2 (Dec. 2009), 89–94. doi:10.2478/IJNAOE-2013-0011. 5

[JE12] JAVED W., ELMQVIST N.: Exploring the design space of composite visualization. In *2012 IEEE Pacific Visualization Symposium* (Songdo, Korea (South), Feb. 2012), IEEE, pp. 1–8. URL: http://ieeexplore.ieee.org/document/6183556/, doi:10.1109/PacificVis.2012.6183556. 3

[JR05] JOSHI A., RHEINGANS P.: Illustration-inspired techniques for visualizing time-varying data. In *VIS 05. IEEE Visualization, 2005.* (Minneapolis, MN, USA, 2005), IEEE, pp. 679–686. doi:10.1109/VISUAL.2005.1532857. 3

[KJVC25] KRAAIJEVELD B., JALBA A. C., VILANOVA A., CHAMBERLAND M.: Real-Time Rendering of Dynamic Line Sets using Voxel Ray Tracing, 2025. doi:10.48550/ARXIV.2510.09081. 2, 3

[KNM*21] KERN M., NEUHAUSER C., MAACK T., HAN M., USHER W., WESTERMANN R.: A Comparison of Rendering Techniques for 3D Line Sets With Transparency. *IEEE Transactions on Visualization and Computer Graphics 27*, 8 (Aug. 2021), 3361–3376. doi:10.1109/TVCG.2020.2975795. 3

[KRW19] KANZLER M., RAUTENHAUS M., WESTERMANN R.: A Voxel-Based Rendering Pipeline for Large 3D Line Sets. *IEEE Transactions on Visualization and Computer Graphics 25*, 7 (July 2019), 2378–2391. doi:10.1109/TVCG.2018.2834372. 3

[LYR*23] LUO W., YU Z., RZAYEV R., SATKOWSKI M., GUMHOLD S., MCGINITY M., DACHSELT R.: Pearl: Physical Environment based Augmented Reality Lenses for In-Situ Human Movement Analysis. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg Germany, Apr. 2023), ACM, pp. 1–15. doi:10.1145/3544548.3580715. 2

[MMYK06] MELEK Z., MAYERICH D., YUKSEL C., KEYSER J.: Visualization of Fibrous and Thread-like Data. *IEEE Transactions on Visualization and Computer Graphics 12*, 5 (Sept. 2006), 1165–1172. doi:10.1109/TVCG.2006.197. 3

[NHK*23] NEUHAUSER C., HIERONYMUS M., KERN M., RAUTENHAUS M., OERTEL A., WESTERMANN R.: Visual analysis of model parameter sensitivities along warm conveyor belt trajectories using Met.3D

(1.6.0-multivar1). *Geoscientific Model Development 16*, 16 (Aug. 2023), 4617–4638. doi:10.5194/gmd-16-4617-2023. 3

[ODNN25] ORTIGOSSA E. S., DIAS F. F., NASCIMENTO D. C., NONATO L. G.: Time Series Information Visualization - A Review of Approaches and Tools. *IEEE Access 13* (2025), 161653–161684. doi:10.1109/ACCESS.2025.3609404. 3

[PJH23] PHARR M., JAKOB W., HUMPHREYS G.: *Physically Based Rendering: From Theory To Implementation*, 4 ed. The MIT Press, Mar. 2023. URL: https://www.pbr-book.org/. 3, 4

[RBE*06] REINA G., BIDMON K., ENDERS F., HASTREITER P., ERTL T.: GPU-Based Hyperstreamlines for Diffusion Tensor Imaging. In *EUROVIS '06: Proceedings of the Eurographics/IEEE VGTC Symposium on Visualization* (2006), The Eurographics Association. ISBN: 9783905673319, ISSN: 1727-5296. doi:10.2312/VISSYM/EUROVIS06/035-042. 3

[Res22] RESHETOV A.: Ray/Ribbon Intersections. *Proceedings of the ACM on Computer Graphics and Interactive Techniques 5*, 3 (July 2022), 1–22. doi:10.1145/3543862. 3

[RGDG23] RUSSIG B., GROSS D., DACHSELT R., GUMHOLD S.: On-Tube Attribute Visualization for Multivariate Trajectory Data. *IEEE Transactions on Visualization and Computer Graphics 29*, 1 (Jan. 2023), 1288–1298. doi:10.1109/TVCG.2022.3209400. 3, 4, 6

[RHD*06] RITTER F., HANSEN C., DICKEN V., KONRAD O., PREIM B., PEITGEN H.-O.: Real-Time Illustration of Vascular Structures. *IEEE Transactions on Visualization and Computer Graphics 12*, 5 (Sept. 2006), 877–884. URL: http://ieeexplore.ieee.org/document/4015442/, doi:10.1109/TVCG.2006.172. 1

[RL18] RESHETOV A., LUEBKE D.: Phantom Ray-Hair Intersector. *Proceedings of the ACM on Computer Graphics and Interactive Techniques 1*, 2 (Aug. 2018), 1–22. doi:10.1145/3233307. 3

[RSG20] RUSSIG B., SALM M., GUMHOLD S.: GPU-based Raycasting of Hermite Spline Tubes. In *2020 IEEE Visualization Conference (VIS)* (Salt Lake City, UT, USA, Oct. 2020), IEEE, pp. 26–30. doi:10.1109/VIS47514.2020.00012. 3

[SGG17] STAIB J., GROTTEL S., GUMHOLD S.: Temporal Focus+Context for Clusters in Particle Data. In *Vision, Modeling & Visualization* (Goslar, DEU, 2017), Hullin M., Klein R., Schultz T., Yao A., (Eds.), VMV '17, The Eurographics Association, pp. 85–93. doi:10.2312/vmv.20171263. 3

[SGS05] STOLL C., GUMHOLD S., SEIDEL H.-P.: Visualization with stylized line primitives. In *VIS 05. IEEE Visualization, 2005.* (Minneapolis, MN, USA, 2005), IEEE, pp. 695–702. doi:10.1109/VISUAL.2005.1532859. 1, 3

[SLK*22] STERZIK A., LICHTENBERG N., KRONE M., CUNNINGHAM D., LAWONN K.: Perceptual Evaluation of Common Line Variables for Displaying Uncertainty on Molecular Surfaces. *Eurographics Workshop on Visual Computing for Biology and Medicine* (2022), 41–51. ISBN: 9783038681779. doi:10.2312/VCBM.20221186. 3

[SLW*23] STERZIK A., LICHTENBERG N., WILMS J., KRONE M., CUNNINGHAM D. W., LAWONN K.: Perception of Line Attributes for Visualization. *IEEE Transactions on Visualization and Computer Graphics* (2023), 1–11. doi:10.1109/TVCG.2023.3326523. 3

[SM02] SCHUSSMAN G., MA K.: Scalable self-orienting surfaces: a compact, texture-enhanced representation for interactive visualization of 3D vector fields. In *10th Pacific Conference on Computer Graphics and Applications, 2002. Proceedings.* (Beijing, China, 2002), IEEE Comput. Soc, pp. 356–365. doi:10.1109/PCCGA.2002.1167879. 2, 3

[SMCL23] STERZIK A., MEUSCHKE M., CUNNINGHAM D. W., LAWONN K.: Perceptually Uniform Construction of Illustrative Textures. *IEEE Transactions on Visualization and Computer Graphics* (2023), 1–11. doi:10.1109/TVCG.2023.3326574. 3

[SML06] SCHROEDER W., MARTIN K., LORENSEN B.: *The Visualization Toolkit*, 4 ed. Kitware, 2006. 5

[SMRB22] SCHERTZER J., MERCIER C., ROUSSEAU S., BOUBEKEUR T.: Fiblets for Real-Time Rendering of Massive Brain Tractograms. *Computer Graphics Forum 41*, 2 (May 2022), 447–460. doi:10.1111/cgf.14486. 3

[TSAA12] TOMINSKI C., SCHUMANN H., ANDRIENKO G., ANDRIENKO N.: Stacking-Based Visualization of Trajectory Attribute Data. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (Dec. 2012), 2565–2574. doi:10.1109/TVCG.2012.265. 3

[Vol89] VOLPE G.: Streamlines and streamribbons in aerodynamics. In *27th Aerospace Sciences Meeting* (Reno,NV,U.S.A., Jan. 1989), American Institute of Aeronautics and Astronautics. doi:10.2514/6.1989-140. 2

[VS82] VENNARD J. K., STREET R. L.: *Elementary Fluid Mechanics*, 6th ed., si version ed. Wiley, New York, 1982. 2

[VW85] VAN WIJK J.: Ray tracing objects defined by sweeping a sphere. *Computers & Graphics 9*, 3 (Jan. 1985), 283–290. doi:10.1016/0097-8493(85)90055-X. 2

[WAPW06] WARE C., ARSENAULT R., PLUMLEE M., WILEY D.: Visualizing the underwater behavior of humpback whales. *IEEE Computer Graphics and Applications 26*, 4 (July 2006), 14–18. doi:10.1109/MCG.2006.93. 3

[WB08] WEIGLE C., BANKS D.: A Comparison of the Perceptual Benefits of Linear Perspective and Physically-Based Illumination for Display of Dense 3D Streamtubes. *IEEE Transactions on Visualization and Computer Graphics 14*, 6 (Nov. 2008), 1723–1730. doi:10.1109/TVCG.2008.108. 3, 6

[Web09] WEBER J. R.: ProteinShader: illustrative rendering of macromolecules. *BMC Structural Biology 9*, 1 (2009), 19. doi:10.1186/1472-6807-9-19. 3

[WESL*23] WEIDINGER N., EL SAYED N., LUZHNICA G., SCHRECK T., VEAS E.: Immersive Analysis of Spatiotemporal Racing Data. In *Proceedings of the 2023 ACM Symposium on Spatial User Interaction* (Sydney NSW Australia, Oct. 2023), ACM, pp. 1–12. doi:10.1145/3607822.3614542. 3

[WNW*22] WANG J., NEUHAUSER C., WU J., GAO X., WESTERMANN R.: 3D-TSV: The 3D trajectory-based stress visualizer. *Advances in Engineering Software 170* (Aug. 2022), 103144. doi:10.1016/j.advengsoft.2022.103144. 3

[Yuk22] YUKSEL C.: High-Performance Polynomial Root Finding for Graphics. *Proceedings of the ACM on Computer Graphics and Interactive Techniques 5*, 3 (July 2022), 1–15. doi:10.1145/3543865. 3

[ZFH*22] ZHOU L., FAN M., HANSEN C., JOHNSON C. R., WEISKOPF D.: A Review of Three-Dimensional Medical Image Visualization. *Health Data Science 2022* (Jan. 2022), 9840519. doi:10.34133/2022/9840519. 2

[ZSH96] ZOCKLER M., STALLING D., HEGE H.-C.: Interactive visualization of 3D-vector fields using illuminated stream lines. In *Proceedings of Seventh Annual IEEE Visualization '96* (San Francisco, CA, USA, 1996), ACM, pp. 107–113,. doi:10.1109/VISUAL.1996.567777. 2, 3, 6