

# BenchLines: Benchmarking Interactive Web Graphics for Parallel Coordinates

Michael Anderson, Jyothish Atheendran, and Filip Ljubotina

706.041 Information Architecture and Web Usability 3VU WS 2025/2026  
Graz University of Technology

03 Feb 2026

## Abstract

This report describes BenchLines, a benchmarking project designed to systematically evaluate and compare four core web rendering technologies: SVG-DOM, Canvas2D, WebGL, and WebGPU, in the context of drawing polylines for a parallel coordinates visualization. To provide deeper insights, BenchLines also incorporates two widely used graphics libraries built on these technologies, namely Pixi.js and Three.js for WebGL and WebGPU.

BenchLines enables standardised performance measurements, allowing for accurate benchmarking of rendering speed as well as common interactive features in a parallel coordinates visualization. This approach facilitates both quantitative and subjective evaluation of performance across multiple interactive features. Based on the results, the report presents recommendations for selecting appropriate rendering approaches depending on performance requirements and workload. These findings aim to guide developers in making informed decisions when building high-performance, visually rich web applications.

© Copyright 2026 by the author(s), except as otherwise noted.

This work is placed under a Creative Commons Attribution 4.0 International (CC BY 4.0) licence.

We have used Grammarly to improve the language in this paper.



# Contents

<b>Contents</b>	<b>ii</b>
<b>List of Figures</b>	<b>iii</b>
<b>List of Tables</b>	<b>v</b>
<b>List of Listings</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Related Work</b>	<b>3</b>
2.1 Parallel Coordinates . . . . .	3
2.2 Core Web Graphics Technologies. . . . .	3
2.2.1 Canvas2D . . . . .	3
2.2.2 SVG-DOM . . . . .	4
2.2.3 WebGL . . . . .	4
2.2.4 WebGL 1.0 . . . . .	4
2.2.5 WebGL 2.0 . . . . .	4
2.2.6 WebGPU . . . . .	5
2.3 Web Graphics Libraries . . . . .	5
2.3.1 Three.js. . . . .	5
2.3.2 Pixi.js . . . . .	5
<b>3 BenchLines System Architecture</b>	<b>7</b>
3.1 Background Layer . . . . .	7
3.2 Rendering Layer . . . . .	11
3.3 Interaction Layer. . . . .	11
3.3.1 Dimension Interactions . . . . .	11
3.3.2 Polyline Hover and Selection . . . . .	12
3.3.2.1 Hover Detection . . . . .	12
3.3.2.2 Selection . . . . .	15
<b>4 Rendering Parallel Coordinates</b>	<b>19</b>
4.1 Core Web Graphics Technologies. . . . .	19
4.1.1 SVG-DOM . . . . .	19
4.1.2 Canvas2D . . . . .	19
4.1.3 WebGL . . . . .	19
4.1.4 WebGPU . . . . .	22

4.2	Using Web Graphics Libraries . . . . .	29
4.2.1	Three.js. . . . .	29
4.2.2	Pixi.js . . . . .	29
<b>5</b>	<b>Benchmarking Results</b>	<b>33</b>
<b>6</b>	<b>Future Work</b>	<b>35</b>
<b>7</b>	<b>Concluding Remarks</b>	<b>37</b>
	<b>Bibliography</b>	<b>39</b>

# List of Figures

1.1	BenchLines WebApp . . . . .	2
3.1	Background Layer . . . . .	8
3.2	Rendering Layer . . . . .	11
3.3	Interaction Layer . . . . .	12
3.4	Hover Detection: Distance from Point to Line . . . . .	13
3.5	Polyline Selection: Line and Box Selection Modes . . . . .	15



# List of Tables

4.1	WebGPU Color Usage . . . . .	22
5.1	BenchLines: Benchmark Results. . . . .	34



# List of Listings

3.1	Background Layer: Canvas Initialization . . . . .	8
3.2	Background Layer: Rasterization . . . . .	9
3.3	Background Layer: Canvas2D Rendering . . . . .	10
3.4	Hover Detection: GPU-Based Point-to-Line Distance Calculation . . . . .	13
3.5	Hover Detection: CPU-Based Dimension Span Detection . . . . .	14
3.6	Hover Detection: CPU-Based Point-to-Line Distance Calculation . . . . .	14
3.7	Line Selection: GPU-Based Line Segment Intersection . . . . .	16
3.8	Line Selection: CPU-Based Line Segment Intersection . . . . .	16
3.9	Box Selection: GPU-Based Box Intersection . . . . .	17
3.10	Box Selection: CPU-Based Box Intersection . . . . .	17
3.11	Line and Box Selection: CPU-Based Multi-Dimensional Span Identification . . . . .	18
4.1	SVG-DOM: Polyline Rendering . . . . .	20
4.2	Canvas2D: Polyline Rendering . . . . .	20
4.3	WebGL: Vertex Shader . . . . .	21
4.4	WebGL: Fragment Shader . . . . .	21
4.5	WebGL: Program . . . . .	22
4.6	WebGL: Polyline Rendering . . . . .	23
4.7	WebGPU: Vertex and Fragment Shaders . . . . .	24
4.8	WebGPU: Render Pipeline . . . . .	25
4.9	WebGPU: Vertex Buffer layout . . . . .	26
4.10	WebGPU: Bind Groups and Uniform Buffers . . . . .	26
4.11	WebGPU: Lines with Custom Thickness . . . . .	27
4.12	WebGPU: Polyline Rendering . . . . .	28
4.13	Three.js: Line Rendering . . . . .	30
4.14	Pixi.js: Polyline Rendering . . . . .	31



# Chapter 1

## Introduction

The landscape of web rendering has undergone significant evolution over the past decade, offering a diverse set of technologies for creating interactive and visually rich web applications. Modern web applications rely on rendering engines not only for displaying graphics, but also for handling complex interactions and visualisations efficiently. The four core web rendering technologies, SVG-DOM, Canvas2D, WebGL, and WebGPU each offer unique capabilities and trade-offs in terms of performance, flexibility, hardware requirements, and the development effort needed to leverage them effectively. Understanding these differences is crucial for developers to optimize user experience, development workflow, and hardware utilization.

This report describes BenchLines, an open-source benchmarking tool designed to evaluate the performance of web rendering technologies in the context of drawing polylines for a parallel coordinates visualization [Anderson et al. 2026]. Parallel coordinates is a widely used visualization technique for exploring multidimensional datasets. The BenchLines project builds upon earlier work called SPCD3 [Andrews et al. 2026].

BenchLines implements parallel coordinates with each of the featured technologies, so that performance can be measured and compared. It provides numerous datasets with differing numbers of dimensions and records. It is possible both to run trials measuring rendering performance and to manually evaluate common interactive features. Figure 1.1 shows the user interface of BenchLines.

In addition to the core technologies, BenchLines also includes support for two widely used libraries built on top of the WebGL and WebGPU platforms: Pixi.js and Three.js. These libraries simplify GPU utilization by providing higher-level abstractions, allowing developers to implement complex graphics, without delving into low-level programming details.

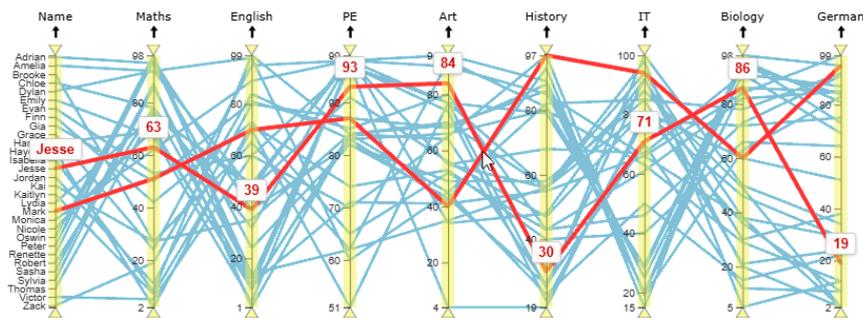
The primary contributions of this work are:

1. Comprehensive benchmarking of web rendering technologies in an interactive visualization scenario.
2. Analysis of the performance and interactivity of each rendering technology.
3. Practical recommendations for selecting an appropriate rendering technology or library based on performance, interactivity, and development requirements.

## BenchLines: Graphics Technologies Benchmark for Parallel Coordinates

Rendering:  Hover:  Dataset:

Avg. rendering time benchmark:



### Past Tests

#	Technology	Dataset	Iterations	Avg Time (ms)
1	SVG-DOM	student_dataset	100	0.392
2	Canvas2D	student_dataset	100	0.759
3	WebGL	student_dataset	100	0.715
4	WebGPU	student_dataset	100	0.711
5	Pixi WebGL	student_dataset	100	0.416
6	Pixi WebGPU	student_dataset	100	0.644
7	Three WebGL	student_dataset	100	0.680

Figure 1.1: BenchLines Website [Screenshot taken by Michael Anderson.]

## Chapter 2

# Related Work

This chapter describes some of the techniques and technologies involved in the BenchLines project.

### 2.1 Parallel Coordinates

Parallel coordinates visualizations are widely used for visualising and exploring multidimensional data. Unlike traditional plots, where a dimension is typically represented by an axis in two or three-dimensional space, parallel coordinates represent each dimension as a vertical axis arranged in parallel. Individual data points are represented as a polyline that intersects each axis at the corresponding value. The technique was popularised by Inselberg [2009].

These visualisations are particularly useful for exploratory data analysis [Weitz 2020]. They allow patterns, correlations, and outliers to be identified across multiple dimensions simultaneously [Gruber and Andrews 2026]. However, the complexity of rendering increases with the number of dimensions and data points, which can lead to performance challenges, especially for interactive applications.

### 2.2 Core Web Graphics Technologies

Modern web-based visualizations rely on a range of graphics technologies that differ in rendering models, performance characteristics, and interaction support. Understanding the historical development and core principles of these technologies provides essential context for the design decisions behind BenchLines.

#### 2.2.1 Canvas2D

The Canvas2D API can be accessed via a `<canvas>` element with a "2d" rendering context. It provides an interface for rendering two-dimensional graphics directly in the browser. Unlike retained-mode graphics systems such as SVG, Canvas2D operates using immediate-mode rendering, meaning the canvas does not “remember” objects, but every frame must be explicitly redrawn [Taivalsaari et al. 2017]. This makes Canvas2D ideal for dynamic graphics, games, and animations where content changes frequently [Aircada 2024]. The API provides methods for direct pixel manipulation, such as `getImageData()` and `putImageData()`, which allow control over individual pixels for custom rendering.

Performance-wise, Canvas2D is mostly CPU-based, and performance degrades with complex or high-volume drawing operations [PixelFree 2024]. Since Canvas2D utilizes the CPU to render content directly to the DOM, browsers have optimized Canvas2D to be very performative and even offload some rendering tasks to the GPU in the background [Larabel 2023]. In this sense, in certain scenarios, like plotting lines, Canvas2D remains competitive in terms of rendering performance.

### 2.2.2 SVG-DOM

The SVG Document Object Model (SVG-DOM) refers to the use of JavaScript to programmatically create, manipulate, and interact with Scalable Vector Graphics (SVG) elements in the browser's Document Object Model (DOM). Since SVG describes vector graphics objects, there is no loss of quality when the graphics are scaled up or down. However, since SVG elements are stored in the DOM, graphics containing many hundreds or thousands of elements can significantly strain the CPU.

According to Wikipedia [2025], the World Wide Web Consortium (W3C) began developing the SVG standard in 1998 after six competing submissions for vector graphics standards were submitted. SVG 1.0 became a W3C Recommendation on 04 Sep 2001. The specification continued to evolve, with SVG 1.1 Second Edition released on 16 Aug 2011. Most current web browsers support SVG 1.1. A W3C Candidate Recommendation for SVG 2 was published on 15 Sep 2016, with the latest draft on 14 Sep 2025, but browser support for SVG 2 is feature-based rather than version-based, and no modern browser fully implements the entire SVG 2 spec.

### 2.2.3 WebGL

The Web Graphics Library (WebGL) is a JavaScript API that for rendering interactive 2d and 3d graphics directly in the browser [MDN 2025] using the HTML `<canvas>` element. The first version of WebGL was released at the Game Developers Conference in San Francisco in 2011 [Khronos 2011]. The technology is based on the OpenGL ES (Embedded Systems) standard, which was designed for efficient graphics rendering on mobile or other embedded devices [Gilbert 2025]. As such, WebGL works natively in all major browsers [Deveria 2025], meaning users do not need to enable anything to be able to utilize WebGL.

### 2.2.4 WebGL 1.0

WebGL 1.0 is built on the OpenGL ES 2.0 standard and uses the `<canvas>` element with the "webgl" rendering context. Using the GPU means graphics can be rendered quicker and more efficiently than on the CPU, without the need for additional plugins or adjusting browser settings. WebGL 1.0 supports programmable shaders and 3d geometry, which allows developers to make complex visualizations, games, or simulations entirely within the browser.

WebGL 1.0 also has its drawbacks. Since it is based on OpenGL ES 2.0, some of the advanced features available in later versions of OpenGL ES are not supported, such as multiple render targets, 3d textures, and some instancing. This can limit the ability to create more complex visualizations or simulations as compared to WebGL 2.0. Also, WebGL exposes low-level GPU operations, meaning developers need a solid understanding or background in graphics programming, including the use of shader programming and buffer management, to fully utilize WebGL 1.0 to its full capabilities. To make this easier, some web graphics libraries, such as Three.js and Pixi.js, have been developed. These frameworks simplify working with WebGL, providing easier-to-understand and utilize methods for accessing the GPU powered graphics rendering while obfuscating some of the more complicated details.

### 2.2.5 WebGL 2.0

WebGL 2.0 was released in 2017 [Mo 2017] and is based on the OpenGL ES 3.0 standard. Like WebGL 1.0, it provides a JavaScript API for hardware-accelerated rendering of 2d and 3d graphics directly in the browser. WebGL 2.0 is initialized via the `<canvas>` element with the "webgl2" rendering context. WebGL 2.0 also allows developers access to the GPU which leads to improved rendering performance.

WebGL 2.0 extends the features available in WebGL 1.0 by offering several features made possible by OpenGL ES 3.0. These include support for 3d textures, instanced rendering, multiple render targets, and the ability to query objects directly. These additional features make it possible for developers to create more complex graphics and visualizations while improving the performance. Developers can

now implement complex particle systems and large-scale data visualizations more efficiently than with WebGL 1.0.

Despite this, WebGL 2.0 is not without limitations. While it is supported by most modern browsers, it is slightly less supported than WebGL 1.0, which can affect compatibility with older devices. WebGL 2.0 is also a low-level API, meaning developers must have a strong understanding of graphics programming to take full advantage of the additional features. Libraries such as Three.js and Pixi.js continue to be useful to simplify this process, allowing developers a simpler method to utilize these advanced features without directly interfacing with the GPU.

Overall, WebGL 2.0 allows for better performance and additional features not available in WebGL 1.0. This makes it one of the more common choices for modern, high-performance web graphics applications.

### 2.2.6 WebGPU

The WebGPU API started as an initial idea by Google called “Explicit Web Graphics API” presented by Corentin Wallez [Wallez 2016] on 08 Jun 2016. The main aim was to propose a “WebGL Next” that could replace WebGL. Later, with subsequent evolutions, W3C published their first public working draft of WebGPU on 18 May 2021 [Malyszau and Ninomiya 2021]. The WebGPU API can be accessed by web browsers using JavaScript. This provides access to the Vulkan [Vulkan 2026], Metal [Apple 2026], or Direct3D 12 [Microsoft 2026] technologies of the Linux, Apple or Windows system, respectively.

WebGPU is available in all major desktop browsers [Deveria 2026]. However, in certain Linux systems, it has to be enabled by setting a browser flag. There has been an effort to develop a translation layer that can convert WebGL calls to WebGPU calls [Han et al. 2025]. Although there are some limitations (such as a lack of one-to-one calls from WebGL to WebGPU), and even with the translation overhead, converting code from WebGL to WebGPU was shown to yield better performance.

## 2.3 Web Graphics Libraries

Working directly with low-level web graphics technologies such as WebGL and WebGPU can be challenging. They require deep knowledge of the graphics pipeline and involve large amounts of code. This is why higher-level graphics libraries have become essential in modern web development. They sit on top of WebGL or WebGPU and provide clean, intuitive APIs that simplify the process.

### 2.3.1 Three.js

Three.js [Danchilla 2012] is a popular open-source JavaScript library that simplifies rendering graphics using WebGL and WebGPU. It abstracts away the lower-level programming, thus making it easy and faster for programmers to work on the underlying technologies. Three.js was created by Ricardo Cabello (also known as Mr.doob) and was first released in Apr 2010 on GitHub [Cabello 2026]. When launched, it supported WebGL 1.0. However, through generous community contributions, it was able to support WebGL 2.0 in the r95 version update on 31 Jul 2018. Since version r118 (released on 27 Jun 2020), Three.js uses WebGL 2.0 by default, with WebGL 1.0 as a fallback. Since version r171 published on 29 Nov 2024, WebGPU was made available production-ready with zero-config imports. As of 2026, Three.js is used for several projects such as web graphics rendering, game development, WebXR Expansion and much more [Lecamus 2026].

### 2.3.2 Pixi.js

Pixi.js is an open-source 2D web graphics rendering library [Pixi 2026]. It was developed by Mathew Groves, where the initial commit of the project was made on 21 Jan 2013. It is a fast HTML5 Creation Engine that by default leverages WebGL for hardware-accelerated 2D rendering and uses Canvas as a fallback if certain older platforms cannot handle WebGL. PixiJS v8 introduced WebGPU support along

with many other features and improvements. This is an actively maintained project and is used in gaming, data visualization, interactive web apps, and so on.

## Chapter 3

# BenchLines System Architecture

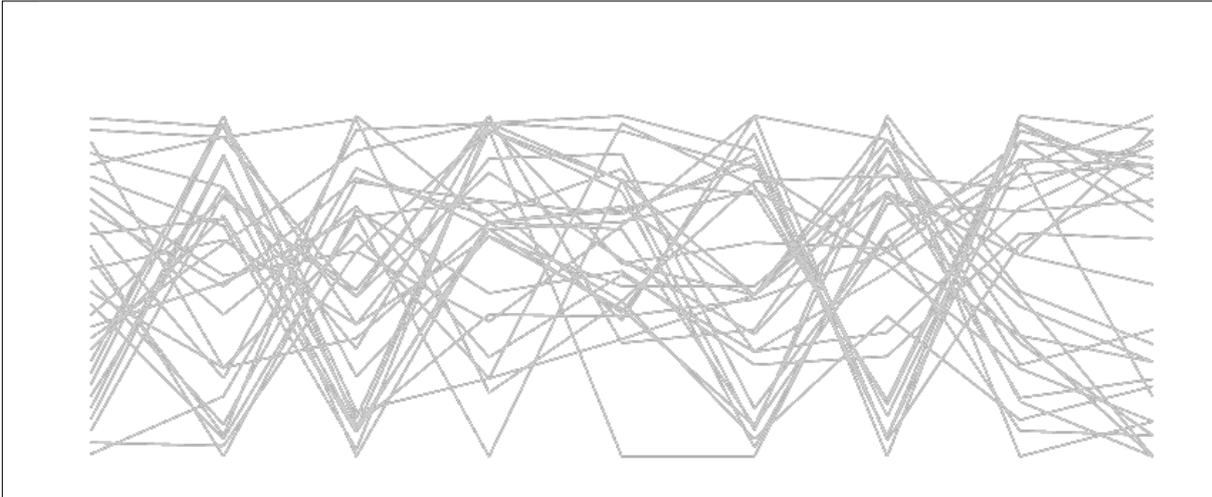
BenchLines utilises three canvases (layers) for drawing and interaction. The lowermost canvas, the background layer, contains a background raster image, showing all polyline rendered in the inactive color (gray). The middle canvas serves as the rendering layer, where active polylines are drawn. The third, foreground, layer is an overlay interaction layer that handles mouse input, such as hovering over polylines and manipulating a filter control.

### 3.1 Background Layer

The background layer consists of a raster image placed beneath all other layers. The idea is to draw all possible polylines in the inactive color (gray) onto a canvas, and then only draw the currently active polylines in blue on top of them. An example of a background canvas is shown in Figure 3.1.

Parallel coordinate visualisations can contain thousands of polylines, and rendering all of them in real time can be computationally expensive. By creating a rasterised background image during scene initialisation, the inactive lines do not need to be re-rendered during normal operations. Some interactive operations, such as reordering or inverting dimensions, require that the background canvas be re-rendered. This process allows for inactive background lines to be displayed when filtering the dataset without having to continuously re-render inactive lines while the user filters the data. During a benchmarking trial, only the rendering of the full set of active polyline is considered.

To implement the background image rasterisation, BenchLines uses an offscreen WebGL canvas that matches the height and width of the main background canvas. This canvas is not visible to users. The code used to initialize it is shown in Listing 3.1. Using WebGL, the GPU is employed to render the inactive lines in the same manner as the active lines. Once all polylines are rendered to the offscreen canvas, the content is rasterised into an image and placed onto the designated background canvas, as shown in Listing 3.2. If WebGL is unavailable, Canvas2D is used to render the background lines directly onto the background canvas, as shown in Listing 3.3.



**Figure 3.1:** Background Layer: Every polyline is drawn in the inactive color (gray) into the background canvas. [Screenshot taken by Michael Anderson.]

```
1 //create background lines image
2 inactiveLinesCanvas = document.createElement("canvas");
3 inactiveLinesCanvas.width = canvasEl.width;
4 inactiveLinesCanvas.height = canvasEl.height;
5 inactiveLinesCanvas.style.position = "absolute";
6 inactiveLinesCanvas.style.top = canvasEl.style.top;
7 inactiveLinesCanvas.style.left = canvasEl.style.left;
8 inactiveLinesCanvas.style.pointerEvents = "none";
9 inactiveLinesCanvas.style.width = canvasEl.style.width;
10 inactiveLinesCanvas.style.height = canvasEl.style.height;
11
12 // Insert behind the main canvas
13 canvasEl.parentNode?.insertBefore(inactiveLinesCanvas, canvasEl);
14
15 redrawWebGLBackgroundLines(dataset, parcoords);
```

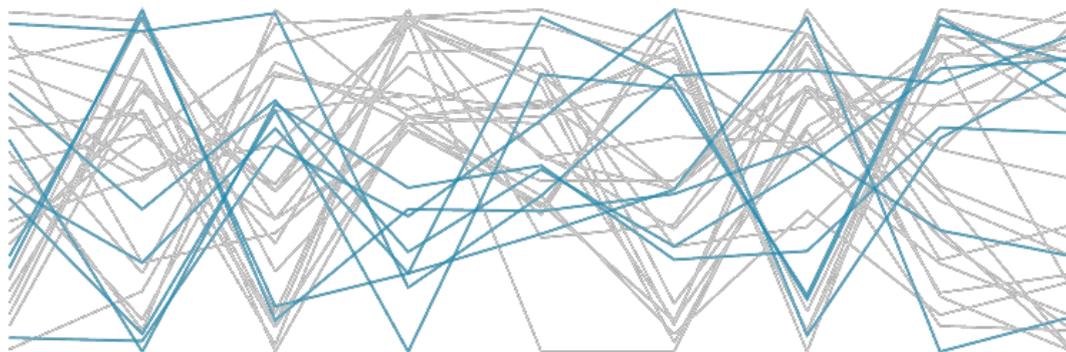
**Listing 3.1:** Background Layer: Typescript code to initialize a background canvas for background lines to be rendered with WebGL.

```
1 export function rasterizeInactiveLinesToCanvas(targetCanvas: HTMLCanvasElement) {
2   if (!gl) {
3     return;
4   }
5
6   const sourceCanvas = gl.canvas as HTMLCanvasElement;
7   const w = sourceCanvas.width;
8   const h = sourceCanvas.height;
9
10  const pixels = new Uint8Array(w * h * 4);
11  gl.readPixels(0, 0, w, h, gl.RGBA, gl.UNSIGNED_BYTE, pixels);
12
13  const ctx = targetCanvas.getContext("2d");
14  if (!ctx) return;
15
16  const imageData = ctx.createImageData(w, h);
17
18  // Flip Y (WebGL bottom-left -> Canvas top-left)
19  for (let y = 0; y < h; y++) {
20    for (let x = 0; x < w; x++) {
21      const src = ((h - y - 1) * w + x) * 4;
22      const dst = (y * w + x) * 4;
23
24      imageData.data[dst] = pixels[src];
25      imageData.data[dst + 1] = pixels[src + 1];
26      imageData.data[dst + 2] = pixels[src + 2];
27      imageData.data[dst + 3] = pixels[src + 3];
28    }
29  }
30
31  ctx.putImageData(imageData, 0, 0);
32 }
```

**Listing 3.2:** Background Layer: Typescript code to rasterize (draw) every polyline in the inactive color (gray) to the background canvas.

```
1 function drawInactiveLinesCanvas2D(dataset: any[], parcoords: any,
2   targetCanvas: HTMLCanvasElement) {
3   const ctx = targetCanvas.getContext("2d");
4   if (!ctx) return;
5
6   const dpr = window.devicePixelRatio || 1;
7
8   // reset transforms
9   ctx.setTransform(dpr, 0, 0, dpr, 0, 0);
10  ctx.clearRect(0, 0, targetCanvas.width, targetCanvas.height);
11
12  ctx.strokeStyle = "rgba(200,200,200,1)";
13  ctx.lineWidth = 1;
14
15  for (const d of dataset) {
16    const id = getLineNameCanvas(d);
17    if (lineState[id]?.active) continue;
18
19    ctx.beginPath();
20    let first = true;
21    parcoords.newFeatures.forEach((name: string) => {
22      const x = (parcoords.dragging[name] ?? parcoords.xScales(name)) * dpr;
23      const y = parcoords.yScales[name](d[name]) * dpr;
24
25      if (first) {
26        ctx.moveTo(x, y);
27        first = false;
28      } else {
29        ctx.lineTo(x, y);
30      }
31    });
32    ctx.stroke();
33  }
34 }
```

**Listing 3.3:** Background Layer: Typescript code to render the background lines directly to the background canvas (in case WebGL is not available).



**Figure 3.2:** Rendering Layer: Active polylines are rendered in the active color (blue) atop the background layer. [Screenshot taken by Michael Anderson.]

## 3.2 Rendering Layer

BenchLines supports polyline rendering through multiple web graphics technologies, allowing users to select the rendering backend best suited to their performance and compatibility requirements. Supported backends include SVG-DOM, Canvas2D, WebGL, and WebGPU. BenchLines additionally provides implementations built on top of two established rendering libraries, Three.js and Pixi.js, for both WebGL and WebGPU. The specific implementation details and design considerations for each rendering technology are discussed in Chapter 4.

Regardless of the selected rendering backend, BenchLines employs a layered canvas architecture to separate rendering from interaction. Rendered polylines are drawn on top of a static background canvas that contains pre-rasterised visual elements, reducing the need for repeated redraws of inactive content, as illustrated in Figure 3.2.

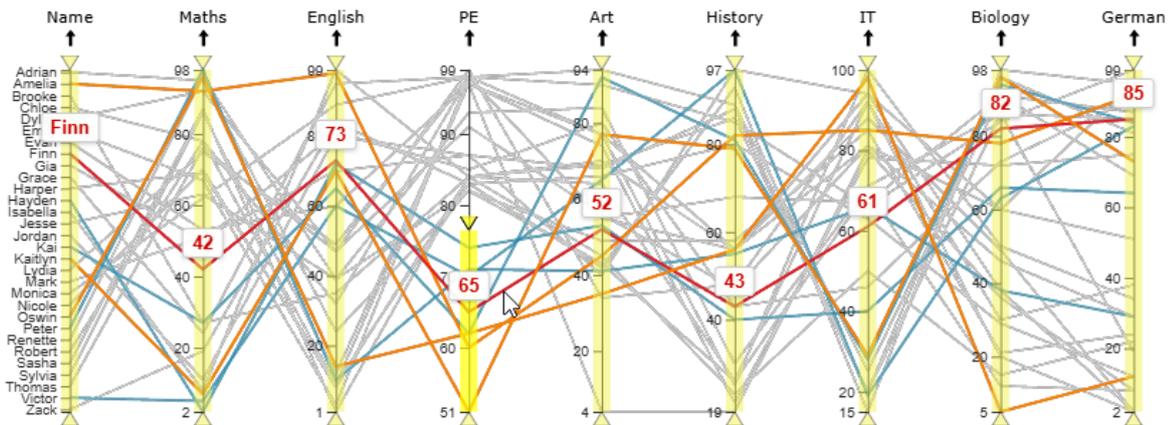
## 3.3 Interaction Layer

The interaction layer is an overlay canvas placed above the rendering layer. It contains interactive controls and is responsible for capturing mouse input, including pointer movement and click events. Interaction events detected on this overlay layer drive updates to the rendering layer, enabling responsive operations such as hovering, highlighting, and filtering without unnecessary recomputation of the background. The interaction layer is shown in Figure 3.3, in front of the other two layers. BenchLines offers three main categories of interactivity: dimension-based interactions, polyline hover effects, and polyline selection.

### 3.3.1 Dimension Interactions

BenchLines provides three primary methods for interacting with dimensions. Users can reorder dimensions by clicking and dragging dimension names horizontally, allowing users to analyze relationships between dimensions. Additionally, clicking the arrow icon beneath a dimension name inverts the sorting order of its values. Finally, dimensions support range-based filtering through draggable yellow arrows positioned at the top and bottom of each axis. By adjusting these arrows, users can define a value range. Any polylines falling outside this range are deactivated and displayed in grey to indicate their filtered status.

It is important to note that all dimension interactions trigger a cascade effect throughout the visualisation. When a user modifies a dimension, all lower layers are automatically notified and updated to maintain consistency.



**Figure 3.3:** Hovered Interactive Lines Atop Render Layer and Background Image. [Screenshot taken by Michael Anderson.]

### 3.3.2 Polyline Hover and Selection

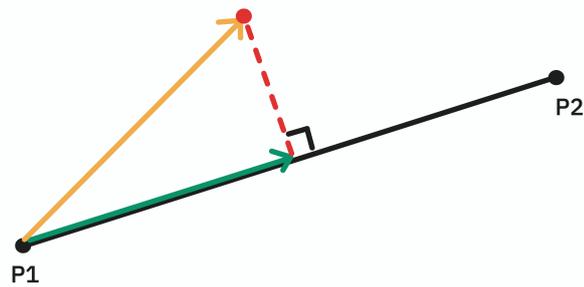
To implement hovering and selection, a hover layer is placed immediately above the plot area containing the chart. It attaches three event listeners (mousemove, mousedown, and mouseup) directly to the plot area. This enables continuous tracking of mouse position and gestures relative to the rendered polylines. Since the hover layer operates on the same dataset used for rendering, it maintains a consistent spatial reference frame, allowing real-time correlation between cursor position and polyline geometry without requiring additional data transformation or synchronisation overhead. The layer supports three distinct interaction modes: point-based hover detection, line-based selection, and box selection.

The core interaction detection mechanism employs raycasting techniques implemented from scratch, providing functionality equivalent to existing library solutions, while maintaining full control over the implementation. The raycasting logic has been deliberately decoupled from any specific rendering technology, enabling it to operate independently of whether the visualisation uses Canvas2D, WebGL, or other rendering backends. This architectural decision facilitates two parallel implementations: a CPU-based JavaScript solution for broad compatibility and a GPU-accelerated version leveraging compute shaders for enhanced performance.

For the hover layer, polylines are represented internally using a flattened, fixed-size buffer structure. Each polyline is allocated a predetermined segment of a Float32Array buffer capable of storing up to `maxPointsPerLine` coordinate pairs. Points are encoded sequentially in the form  $[x_1, y_1, x_2, y_2, x_3, y_3, \dots]$ , where consecutive pairs define individual line segments. This representation decomposes each polyline into a series of line segments, enabling efficient segment-by-segment distance calculations during raycasting operations.

#### 3.3.2.1 Hover Detection

The hover detection mechanism determines whether a polyline is within proximity of the cursor by computing the perpendicular distance from the mouse position to each line segment. As illustrated in Figure 3.4, for a given line segment defined by points  $P_1$  and  $P_2$ , and a cursor position, the projection of the cursor onto the segment vector is calculated. This is achieved through vector mathematics: the dot product of the cursor-relative vector with the segment direction vector, normalized by the segment's squared length, yields a clamping parameter that determines the closest point on the segment. The Euclidean distance from the cursor to this projected point represents the minimum distance to the line segment. If this distance falls below a configurable threshold, the polyline is considered hovered and subsequently rendered in red. The configurable distance threshold enables fuzzy selection capabilities,



**Figure 3.4:** Hover Detection: Calculating the distance from a point to a line using the vector product.  
[Diagram drawn by Filip Ljubotina.]

```

1 // Hover mode: point-to-line distance
2 for (var i = 0u; i < maxPts - 1u; i++) {
3     let idx = lineIdx * maxPts + i;
4     let p1 = lineData[idx];
5     let p2 = lineData[idx + 1u];
6
7     let pa = mousePos - p1;
8     let ba = p2 - p1;
9     let h = clamp(dot(pa, ba) / dot(ba, ba), 0.0, 1.0);
10    let d = length(pa - ba * h);
11
12    if (d < hoverDist) {
13        hit = 1u;
14        break;
15    }
16 }

```

**Listing 3.4:** Hover Detection: WGSL code for GPU-based point-to-line distance calculation.

allowing adjustments to the hover sensitivity.

The distance calculation follows a standard point-to-line algorithm. Given a cursor point and line segment endpoints, we first compute the vector from the segment start to the cursor, then project this onto the segment direction. The projection is clamped to the  $[0,1]$  interval to ensure the closest point lies within the segment bounds rather than on its infinite extension. The final distance is simply the magnitude of the vector from the cursor to this clamped projection point.

Two implementations are available for computing the point-to-line distance: one GPU-based for WebGPU, and the second CPU-based written in standard JavaScript. Both implementations employ identical mathematical logic, yet differ significantly in their execution strategy. The GPU-based implementation is shown in Listing 3.4. It is written in WGSL for WebGPU compute shaders and assigns one thread per polyline. Each thread iterates through all segments of its assigned polyline sequentially, computing the distance to each segment until either a segment within the hover threshold is detected or all segments in the polyline have been examined. This parallelization strategy processes all polylines simultaneously, with the workgroup size of 256 threads enabling efficient utilisation of GPU compute units.

The CPU-based implementation adopts a different optimization strategy suited to sequential processing. Rather than examining all segments of every polyline, it first determines which two adjacent dimensions

```

1 private findSegmentIndex(x: number): number {
2   const positions = this.dimensionXPositions;
3   if (positions.length < 2) return -1;
4
5   for (let i = 0; i < positions.length - 1; i++) {
6     const leftX = positions[i];
7     const rightX = positions[i + 1];
8
9     const minX = Math.min(leftX, rightX);
10    const maxX = Math.max(leftX, rightX);
11
12    if (x >= minX && x <= maxX) {
13      return i;
14    }
15  }
16
17  return -1;
18 }

```

**Listing 3.5:** Hover Detection: TypeScript code for CPU-based detection of the dimension span.

```

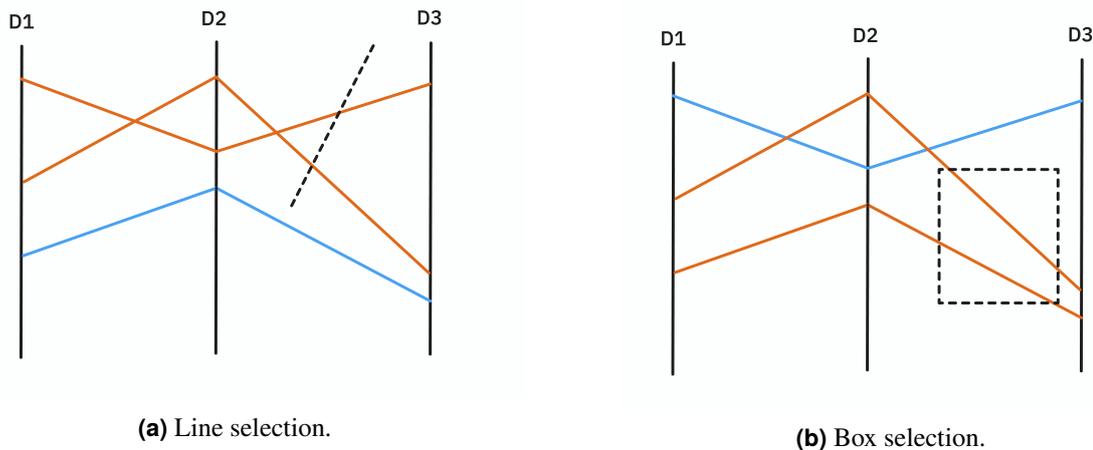
1 private pointToSegmentDistance(
2   px: number, py: number,
3   x1: number, y1: number,
4   x2: number, y2: number
5 ): number {
6   const pax = px - x1;
7   const pay = py - y1;
8   const bax = x2 - x1;
9   const bay = y2 - y1;
10
11   const dotBA = bax * bax + bay * bay;
12   if (dotBA === 0) {
13     return Math.hypot(pax, pay);
14   }
15
16   const h = Math.max(0, Math.min(1, (pax * bax + pay * bay) / dotBA));
17   const projX = pax - bax * h;
18   const projY = pay - bay * h;
19
20   return Math.hypot(projX, projY);
21 }

```

**Listing 3.6:** Hover Detection: TypeScript code for CPU-based point-to-line distance calculation.

the cursor is positioned between. As shown in Listing 3.5, this is achieved by iterating through the stored x-coordinates of each dimension axis and identifying the dimensional span containing the cursor's horizontal position. The method accounts for potential dimension reordering by comparing against both the minimum and maximum x-values of each span. Once the relevant dimensional span is identified, only the polyline segment within that span undergoes the distance computation for each polyline. The code for distance computation is shown in Listing 3.6. This spatial filtering significantly reduces the computational workload, as each polyline typically contains many segments corresponding to the number of dimensions, yet only one falls within the cursor's horizontal position at any given time.

One might question why the same dimensional filtering optimization strategy was not applied to the GPU-based implementation. While technically feasible, this approach was not pursued, because the



**Figure 3.5:** Polyline Selection: Line and box selection modes. [Both diagrams created by Filip Ljubotina.]

GPU’s parallel architecture diminishes its relative benefit. Since all polylines are processed simultaneously and each thread terminates immediately upon detecting a hover, the additional segments checked represent minimal overhead. Furthermore, computing the line segment index on the CPU and passing it to the shader via a uniform buffer on every mouse movement would introduce synchronization latency that could offset any computational savings.

### 3.3.2.2 Selection

BenchLines supports two selection modes that enable users to select multiple polylines simultaneously: line selection and box selection. Line selection, illustrated in Figure 3.5a, detects all polylines whose segments intersect a user-drawn line. Box selection, shown in Figure 3.5b, identifies polylines that either pass through or have segments contained within a rectangular region defined by the user. Both selection modes are initiated by clicking and dragging within the plot area, with the selection type determined by a configurable setting.

The line selection algorithm determines whether two line segments intersect using a parametric approach. For a user-drawn line from point  $P_1$  to  $P_2$  and a polyline segment from  $P_3$  to  $P_4$ , the algorithm computes the denominator of the intersection equation to check for parallel lines. If the segments are not parallel, two parameters ( $u_a$  and  $u_b$ ) are calculated, representing the relative position of the intersection point along each segment. If both parameters fall within the  $[0,1]$  interval, the segments intersect within their defined bounds. Listing 3.7 and Listing 3.8 show the GPU-based and CPU-based implementations, respectively, demonstrating identical mathematical logic.

The box selection algorithm extends the intersection logic to rectangular regions. A polyline is considered selected if any of its segments either has an endpoint within the box boundaries or intersects with any of the four box edges. Listing 3.9 shows the GPU-based implementation, which first performs a quick containment check for segment endpoints before testing against each edge. The CPU-based implementation in Listing 3.10 follows the same two-phase approach: an initial endpoint containment test followed by iteration through all four box edges, utilising the line segment intersection function from Listing 3.8.

The GPU-based implementation, as with hover detection, assigns one thread per polyline and iterates through all polyline segments until an intersection is found. The CPU-based implementation employs an extended version of the dimensional filtering optimization. As shown in Listing 3.11, rather than identifying a single-dimensional span between two adjacent dimensions as in hover detection (Listing 3.5), the algorithm identifies all dimensional spans that fall within the horizontal extent of the drawn selection

```

1 fn lineSegmentsIntersect(
2   p1: vec2<f32>, p2: vec2<f32>,
3   p3: vec2<f32>, p4: vec2<f32>
4 ) -> bool {
5   let denom = (p4.y - p3.y) * (p2.x - p1.x) -
6               (p4.x - p3.x) * (p2.y - p1.y);
7   if (abs(denom) < 0.0001) {
8     return false;
9   }
10
11  let ua = ((p4.x - p3.x) * (p1.y - p3.y) -
12            (p4.y - p3.y) * (p1.x - p3.x)) / denom;
13  let ub = ((p2.x - p1.x) * (p1.y - p3.y) -
14            (p2.y - p1.y) * (p1.x - p3.x)) / denom;
15
16  return ua >= 0.0 && ua <= 1.0 &&
17         ub >= 0.0 && ub <= 1.0;
18 }

```

**Listing 3.7:** Line Selection: WGSL code for GPU-based line segment intersection.

```

1 private lineSegmentsIntersect(
2   p1x: number, p1y: number, p2x: number, p2y: number,
3   p3x: number, p3y: number, p4x: number, p4y: number
4 ): boolean {
5   const denom =
6     (p4y - p3y) * (p2x - p1x) -
7     (p4x - p3x) * (p2y - p1y);
8   if (Math.abs(denom) < 0.0001) {
9     return false;
10  }
11
12  const ua =
13    ((p4x - p3x) * (p1y - p3y) -
14     (p4y - p3y) * (p1x - p3x)) / denom;
15  const ub =
16    ((p2x - p1x) * (p1y - p3y) -
17     (p2y - p1y) * (p1x - p3x)) / denom;
18
19  return ua >= 0 && ua <= 1 &&
20         ub >= 0 && ub <= 1;
21 }

```

**Listing 3.8:** Line Selection: TypeScript code for CPU-based line segment intersection.

```

1 fn lineSegmentIntersectsBox(
2   p1: vec2<f32>, p2: vec2<f32>,
3   minX: f32, maxX: f32, minY: f32, maxY: f32
4 ) -> bool {
5   // Check if either endpoint is inside the box
6   if ((p1.x >= minX && p1.x <= maxX &&
7       p1.y >= minY && p1.y <= maxY) ||
8       (p2.x >= minX && p2.x <= maxX &&
9       p2.y >= minY && p2.y <= maxY)) {
10    return true;
11  }
12
13  // Check intersection with each of the four box edges
14  let topLeft      = vec2<f32>(minX, minY);
15  let topRight     = vec2<f32>(maxX, minY);
16  let bottomLeft  = vec2<f32>(minX, maxY);
17  let bottomRight = vec2<f32>(maxX, maxY);
18
19  if (lineSegmentsIntersect(p1, p2, topLeft, topRight) ||
20      lineSegmentsIntersect(p1, p2, topRight, bottomRight) ||
21      lineSegmentsIntersect(p1, p2, bottomRight, bottomLeft) ||
22      lineSegmentsIntersect(p1, p2, bottomLeft, topLeft)) {
23    return true;
24  }
25
26  return false;
27 }

```

**Listing 3.9:** Box Selection: WGSL code for GPU-based box intersection.

```

1 private lineSegmentIntersectsBox(
2   plx: number, ply: number, p2x: number, p2y: number,
3   minX: number, maxX: number, minY: number, maxY: number
4 ): boolean {
5   if ((plx >= minX && plx <= maxX &&
6       ply >= minY && ply <= maxY) ||
7       (p2x >= minX && p2x <= maxX &&
8       p2y >= minY && p2y <= maxY)) {
9     return true;
10  }
11
12  const edges: [number, number, number, number][] = [
13    [minX, minY, maxX, minY], // top
14    [maxX, minY, maxX, maxY], // right
15    [maxX, maxY, minX, maxY], // bottom
16    [minX, maxY, minX, minY], // left
17  ];
18
19  for (const [ex1, ey1, ex2, ey2] of edges) {
20    if (this.lineSegmentsIntersect(
21      plx, ply, p2x, p2y, ex1, ey1, ex2, ey2
22    )) {
23      return true;
24    }
25  }
26
27  return false;
28 }

```

**Listing 3.10:** Box Selection: TypeScript code for CPU-based box intersection.

```
1 private findSegmentRange(  
2   minX: number, maxX: number  
3 ): [number, number] {  
4   const positions = this.dimensionXPositions;  
5   if (positions.length < 2) {  
6     return [-1, -1];  
7   }  
8   let startIdx = -1;  
9   let endIdx = -1;  
10  for (let i = 0; i < positions.length - 1; i++) {  
11    const segMinX = Math.min(positions[i], positions[i + 1]);  
12    const segMaxX = Math.max(positions[i], positions[i + 1]);  
13    if (segMaxX >= minX && segMinX <= maxX) {  
14      if (startIdx === -1) {  
15        startIdx = i;  
16      }  
17      endIdx = i;  
18    }  
19  }  
20  return [startIdx, endIdx];  
21 }
```

**Listing 3.11:** Line and Box Selection: TypeScript code for CPU-based multi-dimensional span identification. It returns the start and end indices for dimensional spans overlapped by the selection object.

object. For line selection, this encompasses all spans from the leftmost to the rightmost x-coordinate of the drawn line; for box selection, it includes all spans overlapped by the rectangle's width. Only polyline segments crossing these identified dimensional spans are then tested using the intersection functions.

## Chapter 4

# Rendering Parallel Coordinates

BenchLines supports multiple rendering methodologies for parallel coordinates, each with differing rendering models and performance characteristics.

### 4.1 Core Web Graphics Technologies

Implementations are provided for each of the four core web graphics technologies: SVG-DOM, Canvas2D, WebGL, and WebGPU.

#### 4.1.1 SVG-DOM

BenchLines builds upon the SPCD3 project [Andrews et al. 2026], which used SVG-DOM as its primary rendering method. Consequently, the original SVG-DOM rendering pipeline was preserved during the development of BenchLines and adopted as a visual baseline against which other rendering approaches were compared. Visual properties such as color, line thickness, and interaction behavior were defined to mirror the SVG-DOM implementation and subsequently replicated across the remaining rendering technologies to ensure visual and functional consistency. Listing 4.1 shows an example of SVG-DOM-based rendering. This code is not used within the BenchLines project and is included solely as a reference to illustrate how such an approach may be implemented by interested developers in future work.

#### 4.1.2 Canvas2D

The Canvas2D implementation uses the HTML5 `<canvas>` element and its "2d" rendering context, providing a CPU-based approach to drawing graphics. While not as performant as GPU-accelerated methods for large datasets, Canvas2D offers simplicity and broad browser compatibility. The 2d context is obtained from the canvas element, with DPR scaling applied via `setTransform()`. For the background canvas containing inactive lines, BenchLines leverages WebGL rendering, which is then rasterised onto a 2D canvas, combining simple compositing with GPU-accelerated performance. Active polylines are rendered by iterating through the dataset and constructing paths using `beginPath()`, `moveTo()`, and `lineTo()` calls, then stroking with the appropriate color and width. Listing 4.2 demonstrates this approach.

#### 4.1.3 WebGL

The WebGL implementation in BenchLines is designed to leverage GPU acceleration for rendering large parallel coordinate datasets efficiently. The process begins with the initialisation of the WebGL context. As part of this initialisation, BenchLines first calculates the device pixel ratio (DPR), which ensures that the visualisation scales appropriately across a variety of screen resolutions and high-DPI displays.

Next, BenchLines creates the core GPU programs required for rendering. A vertex shader, shown in Listing 4.3, is used to process the position of each vertex in the dataset. A fragment shader, shown in

```

1  const svg = document.createElementNS("http://www.w3.org/2000/svg", "svg");
2  svg.setAttribute("width", canvas.width);
3  svg.setAttribute("height", canvas.height);
4  document.body.appendChild(svg);
5
6  for (const d of dataset) {
7    const polyline =
8      document.createElementNS("http://www.w3.org/2000/svg", "polyline");
9    polyline.setAttribute("points", getPolylinePoints(d));
10   polyline.setAttribute("stroke", "steelblue");
11   polyline.setAttribute("stroke-width", "2");
12   polyline.setAttribute("fill", "none");
13   svg.appendChild(polyline);
14 }

```

**Listing 4.1:** SVG-DOM: JavaScript code for drawing polylines.

```

1  export function redrawCanvasLines(newDataset: any, parcoords: any) {
2    if (!ctx || !canvasEl || !newDataset) return;
3    dataset = newDataset;
4    ctx.clearRect(0, 0, canvasEl.width, canvasEl.height);
5
6    for (const d of newDataset) {
7      const id = getLineNameCanvas(d);
8      const active = lineState[id]?.active ?? true;
9      if (!active) continue;
10
11     const pts = getPolylinePoints(d, parcoords);
12     if (!pts.length) continue;
13
14     ctx.beginPath();
15     ctx.moveTo(pts[0][0], pts[0][1]);
16     for (let i = 1; i < pts.length; i++) {
17       ctx.lineTo(pts[i][0], pts[i][1]);
18     }
19     ctx.lineWidth = 2;
20     ctx.strokeStyle = "rgba(0,129,175,0.5)";
21     ctx.stroke();
22   }
23   redrawHoverOverlay();
24 }

```

**Listing 4.2:** Canvas2D: TypeScript code for Canvas2D rendering of active polylines.

```
1 const vertexShaderSrc = '  
2 attribute vec2 position;  
3 attribute vec4 a_color;  
4 uniform vec2 resolution;  
5 varying vec4 v_color;  
6  
7 void main() {  
8     vec2 zeroToOne = position / resolution;  
9     vec2 clipSpace = zeroToOne * 2.0 - 1.0;  
10    gl_Position = vec4(clipSpace * vec2(1, -1), 0, 1);  
11    v_color = a_color;  
12 }  
13 ';
```

**Listing 4.3:** [WebGL: TypeScript code for a vertex shader.

```
1 const fragmentShaderSrc = '  
2 precision mediump float;  
3 varying vec4 v_color;  
4 void main() {  
5     gl_FragColor = v_color;  
6 }  
7 ';
```

**Listing 4.4:** WebGL: TypeScript code for a fragment shader.

Listing 4.4, handles the coloring of pixels. These shaders are compiled and then linked into a WebGL program, shown in Listing 4.5, which enables the GPU to perform the rendering operations efficiently.

Once the program is set up, BenchLines creates buffers to store the vertex positions, colors, and other attributes required for rendering each line. At this stage, BenchLines also initializes its three-layer architecture, consisting of the background canvas, the rendering canvas, and the interaction overlay canvas. Each layer serves a specific purpose: the background canvas stores rasterized inactive polylines, the rendering canvas handles displaying the active polylines, and the interaction canvas is used for hovering, selection, and other interactive controls.

Rendering the active polylines involves iterating through the dataset of polylines. Each polyline is split into individual line segments, and each segment is represented as two triangles. This conversion allows precise control over the line thickness, at a slight cost in terms of rendering performance compared to rendering as WebGL lines. The three vertices for each triangle are loaded into the vertex buffer and drawn using the WebGL program (Listing 4.5). By performing these operations on the GPU, BenchLines can efficiently render thousands of lines while maintaining interactivity, even for high-dimensional datasets. The code for drawing polylines with WebGL is shown in Listing 4.6.

```

1 function createProgram(
2   gl: WebGLRenderingContext,
3   vShader: WebGLShader,
4   fShader: WebGLShader
5 ) {
6   const program = gl.createProgram();
7   if (!program) throw new Error("createProgram failed");
8   gl.attachShader(program, vShader);
9   gl.attachShader(program, fShader);
10  gl.linkProgram(program);
11  if (!gl.getProgramParameter(program, gl.LINK_STATUS)) {
12    console.error(gl.getProgramInfoLog(program));
13    throw new Error("Program link failed");
14  }
15  return program;
16 }

```

**Listing 4.5:** WebGL: TypeScript code for a program.

State	Color	RGBA Values
Clear	Transparent	(0,0,0,0)
Active	Light Blue	(0.502, 0.749, 0.839, 1.0)
Inactive	Light Gray	(0.922, 0.922, 0.922, 1.0)
Hover	Red	(1.0, 0.2, 0.2, 1.0)
Selected	Orange	(1.0, 0.502, 0.0, 1.0)

**Table 4.1:** WebGPU: Color usage.

#### 4.1.4 WebGPU

Similar to the WebGL implementation, WebGPU also takes advantage of the GPU to render graphics. The program initially declares the global state variables that help maintain the rendering context and interaction state. Afterwards, constants that specify the active and inactive line widths are declared.

Once a call to render the polylines is made, the navigator interface is initially used to check if a GPU is available. If yes, the process continues; else, the rendering attempt stops after throwing an error indicating the lack of a GPU. Continuing the process, a WebGPU context is acquired from the HTML canvas element. Afterwards, the context is configured to use the WebGPU device acquired before. Also, the preferred canvas format is set up, and alpha mode is declared. In this instance, alpha mode is set as premultiplied, which means that the colors put in the canvas must have their color values already multiplied by the alpha value. Furthermore, a secondary canvas positioned absolutely over the main canvas is dynamically created for rendering the hovered and selected lines. Also, a third canvas is inserted behind the main canvas to display inactive (filtered) polylines.

Listing 4.7 shows the WGSL (WebGPU Shading Language) code for the vertex and fragment shaders. The vertex shader takes 2d vertex positions as input and converts them to 4d clip-space coordinates. Meanwhile, the fragment shader simply returns the uniform color for all pixels.

Listing 4.8 defines how vertices are processed and rendered. And, the Listing 4.9 shows the creation of a vertex buffer that holds the vertex data in memory. Furthermore, the Listing 4.10 shows how the colors are managed for the rendering. Each separate color state is defined with its own uniform buffer and bind group. The Table 4.1 shows the RGBA Values of the different colors used for various line states.

```

1 export function redrawWebGLLines(newDataset: any[], parcoords: any) {
2   if (!gl || !vertexBuffer || !colorBuffer) return;
3   dataset = newDataset;
4
5   gl.useProgram(program);
6   gl.uniform2f(resolutionLoc, canvasEl.width, canvasEl.height);
7   gl.clearColor(0, 0, 0, 0);
8   gl.clear(gl.COLOR_BUFFER_BIT); // clears active lines canvas
9
10  const dpr = window.devicePixelRatio || 1;
11  const vertices: number[] = [];
12  const colors: number[] = [];
13
14  for (const d of newDataset) {
15    const id = getLineNameCanvas(d);
16    const active = lineState[id]?.active ?? true;
17    if (!active) continue; // skip inactive lines
18    const pts = getPolylinePoints(d, parcoords, dpr);
19    if (pts.length < 2) continue;
20
21    const color = [128 / 255, 191 / 255, 214 / 255, 1];
22    const width = ACTIVE_LINE_WIDTH;
23
24    for (let i = 0; i < pts.length - 1; i++) {
25      const x1 = pts[i][0], y1 = pts[i][1];
26      const x2 = pts[i + 1][0], y2 = pts[i + 1][1];
27      const dx = x2 - x1;
28      const dy = y2 - y1;
29      const length = Math.sqrt(dx * dx + dy * dy);
30      if (length === 0) continue;
31      const perpX = -dy / length * (width / 2);
32      const perpY = dx / length * (width / 2);
33
34      vertices.push(x1 + perpX, y1 + perpY);
35      vertices.push(x1 - perpX, y1 - perpY);
36      vertices.push(x2 + perpX, y2 + perpY);
37      vertices.push(x1 - perpX, y1 - perpY);
38      vertices.push(x2 - perpX, y2 - perpY);
39      vertices.push(x2 + perpX, y2 + perpY);
40
41      for (let j = 0; j < 6; j++) {
42        colors.push(...color);
43      }
44    }
45  }
46
47  const vertexData = new Float32Array(vertices);
48  const colorData = new Float32Array(colors);
49
50  gl.bindBuffer(gl.ARRAY_BUFFER, vertexBuffer);
51  gl.bufferData(gl.ARRAY_BUFFER, vertexData, gl.DYNAMIC_DRAW);
52  gl.vertexAttribPointer(posLoc, 2, gl.FLOAT, false, 0, 0);
53
54  gl.bindBuffer(gl.ARRAY_BUFFER, colorBuffer);
55  gl.bufferData(gl.ARRAY_BUFFER, colorData, gl.DYNAMIC_DRAW);
56  gl.vertexAttribPointer(colorLoc, 4, gl.FLOAT, false, 0, 0);
57
58  gl.drawArrays(gl.TRIANGLES, 0, vertexData.length / 2);
59
60  redrawHoverOverlay();
61 }

```

**Listing 4.6:** WebGL: TypeScript code to render the active polylines.

```

1 // Create a new shader module on the GPU device
2 const cellShaderModule = device.createShaderModule({
3 // The label is used for debugging purposes.
4 label: "Shaders",
5 code: `
6   @group(0) @binding(0) var<uniform> color: vec4<f32>;
7   struct VSOut {
8     @builtin(position) position : vec4<f32>,
9   };
10  @vertex
11  fn vs_main(@location(0) pos: vec2<f32>) -> VSOut {
12    var out: VSOut;
13    out.position = vec4<f32>(pos, 0.0, 1.0);
14    return out;
15  }
16  @fragment
17  fn fs_main() -> @location(0) vec4<f32> {
18    return color;
19  }
20 `;
21 });

```

**Listing 4.7:** WebGPU: WGSL code for vertex and fragment shaders.

WebGPU does not natively support rendering lines with custom thickness. To overcome this, the vertices are extrapolated to form rectangles out of two triangles joined together. The Listing 4.11 shows a sample code showcasing this.

Finally, Listing 4.12 shows how polylines are rendered with WebGPU. The function begins with the Command Encoder, which is a recording interface that builds a list of GPU commands to be executed later. Then the render pass begins. This clears the canvas to transparent, prepares the GPU to receive the drawing commands and sets up the rendering pipeline state. The canvas is then set up using the dpr value obtained from the current window, so that the render is also presented well on high-resolution displays without blurriness. Afterwards, the vertices are processed based on their active or inactive state, and the data values are transformed to pixel positions using scale functions.

To display the lines in their correct thickness, 6 vertices are required for each line segment, which represents the edges of the two triangles, essentially forming one thick line. The `addVertex()` function converts the pixel coordinates to the WebGPU clip space. Then, a vertex buffer is created to allocate GPU memory for the drawing. The buffer can store vertex data and can receive data from the CPU. Next, the pipeline is configured to activate the shader programs and rendering settings configured during initialisation. The vertex buffer is then bound to slot 0, making it available to the vertex shader. From here, each polyline is given the appropriate color based on its active or inactive state using the bind groups defined before. The draw command accepts the number of vertices to draw, the number of instances, the starting position in the vertex buffer and the first instance index. Finally, the end of the render pass is marked, and the command encoder converted into an immutable command buffer is sent to the GPU for execution. The GPU processes the commands independently while the CPU continues with its other tasks.

```
1 pipeline = device.createRenderPipeline({
2   layout: device.createPipelineLayout({
3     bindGroupLayouts: [bindGroupLayout],
4   }),
5   vertex: {
6     // GPUShaderModule containing vertex shader
7     module: cellShaderModule,
8     // name of function in shader code called for every vertex invocation
9     entryPoint: "vs_main",
10    // array of GPUVertexBufferLayout objects describing how data is packed
11    buffers: [vertexBufferLayout],
12  },
13  fragment: {
14    module: cellShaderModule,
15    entryPoint: "fs_main",
16    targets: [
17      {
18        format: canvasFormat,
19        blend: {
20          color: {
21            // factor for source color (color being drawn)
22            srcFactor: "src-alpha",
23            // factor for destination color (color already in framebuffer)
24            dstFactor: "one-minus-src-alpha",
25            // blending operation to apply
26            operation: "add",
27          },
28          alpha: {
29            // factor for source alpha (alpha being drawn)
30            srcFactor: "one",
31            // factor for destination alpha (alpha already in framebuffer)
32            dstFactor: "one-minus-src-alpha",
33            // blending operation to apply
34            operation: "add",
35          },
36        },
37      },
38    ],
39  },
40  primitive: {
41    // we are drawing triangles
42    topology: "triangle-list",
43    // For pipelines with strip topologies ("line-strip" or "triangle-strip"),
44    // this determines the index buffer format and primitive restart value
45    stripIndexFormat: undefined,
46  },
47 });
```

**Listing 4.8:** WebGPU: TypeScript code for the render pipeline.

```
1 const vertexBufferLayout: GPUVertexBufferLayout = {
2   // number of bytes tGPU needs to skip forward for next vertex
3   arrayStride: 8,
4   attributes: [
5     {
6       format: "float32x2" as GPUVertexFormat,
7
8       // offset into vertex for this attribute
9       offset: 0,
10
11      // arbitrary number between 0 and 15, unique for every attribute
12      shaderLocation: 0,
13    } as GPUVertexAttribute,
14  ],};
```

**Listing 4.9:** WebGPU: TypeScript code for vertex buffer layout.

```
1 // Create hover color buffer (red)
2 const hoverColorBuffer = device.createBuffer({
3   size: 16,
4   usage: GPUBufferUsage.UNIFORM | GPUBufferUsage.COPY_DST,
5 });
6 device.queue.writeBuffer(
7   hoverColorBuffer,
8   0,
9   new Float32Array([1.0, 51.0 / 255.0, 51.0 / 255.0, 1.0]) // Red with alpha 1.0
10 );
11 hoverBindGroup = device.createBindGroup({
12   layout: bindGroupLayout,
13   entries: [{ binding: 0, resource: { buffer: hoverColorBuffer } }],
14 });
```

**Listing 4.10:** WebGPU: TypeScript code for bind groups and uniform buffers.

```
1 for (const line of allLines) {
2   const width = line.active ? ACTIVE_LINE_WIDTH : INACTIVE_LINE_WIDTH;
3   for (let i = 0; i < line.pts.length - 1; i++) {
4     const p1 = line.pts[i];
5     const p2 = line.pts[i + 1];
6     const x1 = p1[0], y1 = p1[1];
7     const x2 = p2[0], y2 = p2[1];
8     const dx = x2 - x1;
9     const dy = y2 - y1;
10    const length = Math.sqrt(dx * dx + dy * dy);
11    if (length === 0) continue;
12    const perpX = -dy / length * (width / 2);
13    const perpY = dx / length * (width / 2);
14
15    // Triangle 1
16    addVertex(x1 + perpX, y1 + perpY);
17    addVertex(x1 - perpX, y1 - perpY);
18    addVertex(x2 + perpX, y2 + perpY);
19
20    // Triangle 2
21    addVertex(x1 - perpX, y1 - perpY);
22    addVertex(x2 - perpX, y2 - perpY);
23    addVertex(x2 + perpX, y2 + perpY);
24  }
25 }
```

**Listing 4.11:** WebGPU: TypeScript code for rendering lines with custom thickness.

```

1 export function redrawWebGPULines(newDataset: any[], parcoords: any) {
2   dataset = newDataset;
3   currentParcoords = parcoords;
4   // Create command encoder to encode GPU commands
5   encoder = device.createCommandEncoder();
6   pass = encoder.beginRenderPass({
7     colorAttachments: [
8       {
9         view: context.getCurrentTexture().createView(),
10        loadOp: "clear",
11        clearColor: { r: 0, g: 0, b: 0, a: 0 },
12        storeOp: "store",},],});
13   const dpr = window.devicePixelRatio || 1;
14   const canvasWidth = canvasEl.width;
15   const canvasHeight = canvasEl.height;
16   const allLines: { pts: [number, number][]; active: boolean }[] = [];
17   let totalVertexCount = 0;
18   for (const d of newDataset) {
19     const id = getLineNameCanvas(d);
20     const active = lineState[id]?.active ?? true;
21     if (!active) continue;
22     const pts = getPolylinePoints(d, parcoords, dpr);
23     if (pts.length >= 2) {
24       allLines.push({ pts, active });
25       // Each segment has 6 vertices (2 triangles)
26       totalVertexCount += (pts.length - 1) * 6;
27     }
28   }
29   if (totalVertexCount === 0) {
30     // If no lines to draw, finish the pass and submit empty command buffer
31     pass.end();
32     device.queue.submit([encoder.finish()]);
33     return;
34   }
35   const totalBufferSize = totalVertexCount * 2 * 4;
36   const allVerts = new Float32Array(totalVertexCount * 2);
37   let currentOffset = 0;
38   function addVertex(x: number, y: number) {
39     const xClip = (x / canvasWidth) * 2 - 1;
40     const yClip = 1 - (y / canvasHeight) * 2;
41     allVerts[currentOffset++] = xClip;
42     allVerts[currentOffset++] = yClip;
43   }
44   const vertexBuffer = device.createBuffer({
45     label: "all-polyline-vertices",
46     size: totalBufferSize,
47     usage: GPUBufferUsage.VERTEX | GPUBufferUsage.COPY_DST,});
48   device.queue.writeBuffer(vertexBuffer, 0, allVerts);
49   pass.setPipeline(pipeline);
50   pass.setVertexBuffer(0, vertexBuffer);
51   let vertexOffset = 0;
52   for (const line of allLines) {
53     const lineVertexCount = (line.pts.length - 1) * 6;
54     pass.setBindGroup(0, line.active ? activeBindGroup : inactiveBindGroup);
55     pass.draw(lineVertexCount, 1, vertexOffset, 0);
56     vertexOffset += lineVertexCount;
57   }
58   pass.end();
59   device.queue.submit([encoder.finish()]);
60 }

```

**Listing 4.12:** WebGPU: TypeScript code for rendering polylines.

## 4.2 Using Web Graphics Libraries

Implementations are provided for two web graphics libraries: Three.js and Pixi.js.

### 4.2.1 Three.js

Three.js was used to implement both WebGL (using Three.js's `WebGLRenderer`) and WebGPU (using Three.js's experimental `WebGPURenderer`) versions of polyline rendering. Listing 4.13 shows some sample Three.js code to render a line. Depending on the underlying rendering technology needed, the renderer can be swapped. In Three.js, the scene allows the user to define what is to be rendered and where. It acts as a container for all visual elements (objects, lights, cameras) in the 3d or 2d world. A scene is only created once, and objects are added dynamically to it later. The camera defines the viewpoint and projection type. An `OrthographicCamera` is used in the example provided since, in this case, the size of the object rendered stays constant regardless of its distance from the camera. This prevents any perspective distortion. The camera decides how the scene gets projected onto the canvas. The points in the example are a flat array of coordinates `[x1, y1, z1, x2, y2, z2, ...]` defining the line's vertices. Finally, the `LineGeometry` specifies the vertices that form a polyline and the `Line2` is used in the given example, so that the width or thickness of the line can also be configured.

### 4.2.2 Pixi.js

The Listing 4.14 demonstrates example line rendering using Pixi.js. The example code uses the `WebGPU` renderer if it is set up (check using the navigator object), and if not, it falls back to the `WebGL` renderer. After the renderer is set up, the code creates a stage and a `linesContainer`, which is added to the stage. The `graphics` object helps to draw shapes such as rectangles, circles, or any polygons. In this case, it is used to render a polyline. The `pts` array defines the vertices of the polyline. Once all the segments are defined, line properties such as width, color and alpha are configured. Finally, the renderer renders the stage to the canvas.

```
1 // Shared imports and setup
2 import * as THREE from "three"; // For WebGL
3 import { Line2 } from "three/examples/jsm/lines/Line2.js"; // Shared
4 import { LineGeometry } from "three/examples/jsm/lines/LineGeometry.js"; // Shared
5 import { LineMaterial } from "three/examples/jsm/lines/LineMaterial.js";
6
7 // Renderer setup (swap based on backend)
8 let renderer;
9 if (backend === "webgl") {
10   renderer = new THREE.WebGLRenderer({ canvas: canvasEl, alpha: true });
11 } else if (backend === "webgpu") {
12   renderer = new THREE.WebGPURenderer({ canvas: canvasEl, alpha: true });
13   await renderer.init();
14 }
15 renderer.setSize(width, height);
16 renderer.setPixelRatio(window.devicePixelRatio || 1);
17
18 // Material (swap based on backend)
19 let lineMaterial;
20 if (backend === "webgl") {
21   let lineMaterial: LineMaterial | null = null;
22   lineMaterial = new LineMaterial({
23     color: 0x80bfd6,
24     linewidth: 3,
25   });
26   lineMaterial.resolution.set(width, height);
27 } else if (backend === "webgpu") {
28   let lineMaterial: THREE.Line2NodeMaterial | null = null;
29   lineMaterial = new Line2NodeMaterial({
30     color: 0x80bfd6,
31     linewidth: 3,
32     worldUnits: false,
33     alphaToCoverage: true,
34   });
35 }
36
37 // Scene and camera (shared)
38 const scene = new THREE.Scene();
39 const camera = new THREE.OrthographicCamera(0, width, height, 0, -1, 1);
40
41 // Draw a simple thick line
42 const geometry = new LineGeometry();
43 // Horizontal line across canvas
44 const points = [0, height / 2, 0, width, height / 2, 0];
45 geometry.setPositions(points);
46 const line = new Line2(geometry, lineMaterial);
47 line.computeLineDistances();
48 scene.add(line);
49
50 // Render
51 renderer.render(scene, camera);
```

**Listing 4.13:** Three.js: TypeScript code for rendering a line.

```
1 import * as PIXI from "pixi.js";
2 async function initPixiLine(canvas: HTMLCanvasElement) {
3   const dpr = window.devicePixelRatio || 1;
4   const width = canvas.clientWidth;
5   const height = canvas.clientHeight;
6   // Choose WebGPU if available, otherwise use WebGL
7   let renderer: PIXI.Renderer | PIXI.WebGPURenderer;
8   if ("gpu" in navigator) {
9     const webgpuRenderer = new PIXI.WebGPURenderer();
10    await webgpuRenderer.init({
11      canvas,
12      width,
13      height,
14      resolution: dpr,
15      backgroundAlpha: 0,
16      antialias: true,
17      autoDensity: true,
18      clearBeforeRender: true,
19    });
20    renderer = webgpuRenderer;
21  } else {
22    renderer = new PIXI.Renderer({
23      view: canvas,
24      width: width / dpr,
25      height: height / dpr,
26      resolution: dpr,
27      backgroundAlpha: 0,
28      autoDensity: true,
29      clearBeforeRender: true,
30    });
31  }
32  const stage = new PIXI.Container();
33  const linesContainer = new PIXI.Container();
34  stage.addChild(linesContainer);
35  const graphics = new PIXI.Graphics();
36  linesContainer.addChild(graphics);
37  const pts: [number, number][] = [
38    [50, 50],
39    [150, 80],
40    [250, 40],
41    [350, 120],
42  ];
43  graphics.clear();
44  if (pts.length > 0) {
45    graphics.moveTo(pts[0][0], pts[0][1]);
46    for (let i = 1; i < pts.length; i++) {
47      graphics.lineTo(pts[i][0], pts[i][1]);
48      // Use WebGPU-style stroke API if available, else fall back to lineStyle API.
49      if ("stroke" in graphics) {
50        graphics.stroke({width: 3, color: 0x0081af, alpha: 1,});
51      } else {
52        (graphics as any).lineStyle(3, 0x0081af, 1);
53      }
54    }
55    renderer.render(stage);
56  }
```

**Listing 4.14:** Pixi.js: TypeScript code for rendering polylines.



## Chapter 5

# Benchmarking Results

Using BenchLines, rendering time was analysed across three datasets of 10 dimensions and each of 100, 1000, and 10,000 records for 10 iterations. BenchLines subsequently reports the average rendering time. The tests were performed on a MacBook Pro with an Apple M3 Pro and 18 GB of RAM. The results can be seen in Table 5.1.

WebGPU consistently delivers the fastest overall rendering performance. This advantage stems from its modern, low-level design, which enables more efficient and direct use of GPU resources. By leveraging contemporary GPU rendering techniques, WebGPU achieves superior performance compared to the other approaches. This performance advantage comes at the cost of increased implementation complexity, as fully leveraging WebGPU requires explicit GPU programming and a strong understanding of modern graphics pipelines when compared to approaches such as SVG-DOM. Despite this trade-off, WebGPU proved to be the most efficient rendering technology for parallel coordinates visualisations.

Canvas2D and WebGL are both competitive across varying dataset sizes. This result was somewhat surprising, as WebGL was initially expected to outperform Canvas2D. However, modern browsers are highly optimized for Canvas2D in these scenarios and can effectively leverage GPU acceleration. WebGL also remains competitive in terms of performance, although the BenchLines WebGL implementation does not scale as efficiently as the library-based implementation provided by Three.js. The reduced overhead from avoiding an external library is offset at larger scales by enabling more direct and efficient WebGL rendering.

SVG-DOM, while conceptually straightforward and easy to implement, demonstrated limited scalability. It performs adequately for small visualisations, such as those containing fewer than approximately 100 polylines, but does not scale well to larger datasets. Consequently, SVG-DOM is best suited for small-scale visualisations, where simplicity and direct DOM-based interaction are prioritised over rendering performance.

For developers seeking to avoid low-level GPU programming, high-level libraries such as Pixi.js and Three.js provide abstraction layers over WebGL and WebGPU. While these libraries reduce implementation complexity, they still require familiarity with library-specific features to achieve optimal performance. Additionally, the abstraction introduces a small performance overhead compared to custom WebGPU or WebGL implementations. Overall, these libraries offer a practical compromise by combining GPU acceleration with improved developer accessibility.

#	Technology	Dataset	Iterations	Avg Time (ms)
1	WebGPU	10D_100	10	0.460
2	Canvas2D	10D_100	10	0.480
3	WebGL	10D_100	10	1.840
4	Pixi-WebGPU	10D_100	10	2.350
5	Pixi-Canvas2D	10D_100	10	2.470
6	SVG-DOM	10D_100	10	3.140
7	Three-WebGPU	10D_100	10	3.590
8	Three-WebGL	10D_100	10	3.800
9	Pixi-WebGL	10D_100	10	4.760
1	WebGPU	10D_1000	10	1.170
2	WebGL	10D_1000	10	3.240
3	Canvas2D	10D_1000	10	3.300
4	Three-WebGPU	10D_1000	10	3.630
5	Three-WebGL	10D_1000	10	5.410
6	Pixi-WebGPU	10D_1000	10	6.100
7	Pixi-Canvas2D	10D_1000	10	7.650
8	SVG-DOM	10D_1000	10	11.360
9	Pixi-WebGL	10D_1000	10	12.650
1	WebGPU	10D_10000	10	11.040
2	Canvas2D	10D_10000	10	14.270
3	Three-WebGL	10D_10000	10	21.930
4	Three-WebGPU	10D_10000	10	22.360
5	WebGL	10D_10000	10	26.730
6	Pixi-Canvas2D	10D_10000	10	39.790
7	Pixi-WebGPU	10D_10000	10	64.140
8	SVG-DOM	10D_10000	10	67.740
9	Pixi-WebGL	10D_10000	10	82.010

**Table 5.1:** BenchLines: Benchmark results for three different 10-dimensional datasets, with 100, 1000, and 10,000 records. Tested using a MacBook Pro, with Apple M3 Pro and 18 GB RAM. The results for each dataset are sorted in increasing order of average rendering time for all the polylines in the dataset.

## Chapter 6

### Future Work

BenchLines has significant potential for extension and improvement. Additional rendering libraries, such as Orillusion [Orillusion 2023], could be incorporated into the system. By taking advantage of the existing three-layer canvas architecture, additional features can be integrated with minimal system-level modifications. Benchmarking capabilities could also be expanded beyond rendering performance to include interactive features, such as hover behavior, using the existing JavaScript and WebGPU-based hover algorithms. Further implementation of customizable interaction features, including rendering color, line thickness, and other visual properties, would enhance both usability and flexibility.

Accessibility improvements represent another avenue for future work, including the implementation of color blind friendly palettes and alternative color options. As well as the ability to upload or generate custom-sized datasets for further comparison capabilities. Together, these enhancements would broaden the utility of BenchLines as both a benchmarking tool and a framework for developing interactive web-based visualisations.



## Chapter 7

# Concluding Remarks

BenchLines was developed to evaluate and compare web-based rendering technologies within a parallel coordinates visualisation. The system enables the comparison of common rendering methods as well as widely used graphics libraries, providing both quantitative performance assessment and qualitative evaluation of interaction features. By separating background, rendering, and interaction layers, BenchLines allows consistent benchmarking across different technologies while maintaining the interactive features expected of a parallel coordinate system.

Based on the evaluation, WebGPU offers the highest performance, making it the preferred choice when implementation complexity is not a consideration and the target audience supports modern GPU-enabled browsers. For scenarios where low-level GPU programming is not desired, libraries such as Three.js and Pixi.js provide abstraction over WebGL or WebGPU, offering a balance between performance and ease of use, albeit with a minor overhead and the need for library-specific knowledge. Canvas2D remains a practical alternative, delivering competitive performance while providing a simpler and more accessible programming model.



# Bibliography

- Aircada [2024]. *Picking Sides in WebGL vs Canvas*. 2024. <https://aircada.com/blog/webgl-vs-canvas> (cited on page 3).
- Anderson, Michael, Jyothish Atheendran, and Filip Ljubotina [2026]. *BenchLines: Benchmarking Interactive Web Graphics for Parallel Coordinates*. Graz University of Technology, 03 Feb 2026. <https://github.com/filip-ljubotina/BenchLines> (cited on page 1).
- Andrews, Keith, Romana Gruber, Philipp Drescher, Jeremias Kleinschuster, Sebastian Schreiner, and Burim Vrella [2026]. *Steerable Parallel Coordinates in D3*. Graz University of Technology, 03 Feb 2026. <https://github.com/tugraz-isds/spcd3> (cited on pages 1, 19).
- Apple [2026]. *Metal*. Apple Developer Documentation, 2026. <https://developer.apple.com/documentation/metal> (cited on page 5).
- Cabello, Ricardo [2026]. *three.js: JavaScript 3D Library*. 03 Feb 2026. <https://github.com/mrdoob/three.js> (cited on page 5).
- Danchilla, Brian [2012]. *Three.js Framework*. In: *Beginning WebGL for HTML5*. Berkeley, CA: Apress, 2012, pages 173–203. ISBN 1430239972. doi:10.1007/978-1-4302-3997-0\_7 (cited on page 5).
- Deveria, Alexis [2025]. *WebGL– 3D Canvas graphics*. 11 Oct 2025. <https://caniuse.com/webgl> (cited on page 4).
- Deveria, Alexis [2026]. *Can I use WebGPU?* 2026. <https://caniuse.com/webgpu> (cited on page 5).
- Gilbert, Kelsey [2025]. *WebGL 2.0 Specification*. 07 Feb 2025. <https://registry.khronos.org/webgl/specs/latest/2.0/> (cited on page 4).
- Gruber, Romana and Keith Andrews [2026]. *Parallel Coordinates: An Interactive Tutorial*. 2026. <https://tugraz-isds.github.io/pcee/> (cited on page 3).
- Han, Yudong, Weichen Bi, RuiBo An, Deyu Tian, Qi Yang, and Yun Ma [2025]. *GL2GPU: Accelerating WebGL Applications via Dynamic API Translation to WebGPU*. Proc. The ACM Web Conference (WWW 2025) (Sydney, Australia). 28 Apr 2025, pages 751–762. doi:10.1145/3696410.3714785 (cited on page 5).
- Inselberg, Alfred [2009]. *Parallel Coordinates: Visual Multidimensional Geometry and Its Applications*. Springer, 08 Oct 2009. 554 pages. ISBN 0387215077 (cited on page 3).
- Khronos [2011]. *Khronos Releases Final WebGL 1.0 Specification*. 03 Mar 2011. <https://khr.io/ba> (cited on page 4).
- Larabel, Michael [2023]. *Firefox 110 Released With Better WebGL Performance, GPU-Accelerated 2D Canvas*. 14 Feb 2023. <https://www.phoronix.com/news/Firefox-110-Released> (cited on page 3).
- Lecamus, Jocelyn [2026]. *What's New in Three.js (2026): WebGPU, New Workflows & Beyond*. Utsubo, 10 Jan 2026. <https://utsubo.com/blog/threejs-2026-what-changed> (cited on page 5).

- Malyszau, Dzmitry and Kai Ninomiya [2021]. *WebGPU*. W3C Working Draft. First Public Working Draft. W3C, May 2021. <https://w3.org/TR/2021/WD-webgpu-20210518/> (cited on page 5).
- MDN [2025]. *WebGL: 2D and 3D Graphics for the Web*. Mozilla Developer Network, 15 Jul 2025. [https://developer.mozilla.org/en-US/docs/Web/API/WebGL\\_API](https://developer.mozilla.org/en-US/docs/Web/API/WebGL_API) (cited on page 4).
- Microsoft [2026]. *Direct3D 12 programming guide*. Microsoft Learn, 2026. <https://learn.microsoft.com/en-us/windows/win32/direct3d12/directx-12-programming-guide> (cited on page 5).
- Mo, Zhenyao [2017]. *WebGL 2.0 Arrives*. 27 Feb 2017. <https://khronos.org/blog/webgl-2.0-arrives> (cited on page 4).
- Orillusion [2023]. *Orillusion*. 2023. <https://github.com/Orillusion/orillusion> (cited on page 35).
- PixelFree [2024]. *WebGL vs. Canvas: Which is Better for 3D Web Development?* 2024. <https://blog.pixelfreestudio.com/webgl-vs-canvas-which-is-better-for-3d-web-development/> (cited on page 3).
- Pixi [2026]. *PixiJS Repository*. 03 Feb 2026. <https://github.com/pixijs/pixijs> (cited on page 5).
- Taivalsaari, Antero, Tommi Mikkonen, Cesare Pautasso, and Kari Systä [2017]. *Comparing the Built-In Application Architecture Models in the Web Browser*. Proc. International Conference on Software Architecture (ICSA 2017) (Gothenburg, Sweden). IEEE, 03 Apr 2017, pages 51–54. doi:10.1109/ICSA.2017.23. <https://si.usi.ch/assets/publications/conf/icsa/icsa2017/TaivalsaariMPS17.pdf> (cited on page 3).
- Vulkan [2026]. *Vulkan Documentation*. 03 Feb 2026. <https://docs.vulkan.org/spec/latest/> (cited on page 5).
- Wallez, Corentin [2016]. *Explicit Web Graphics API*. Presentation to the WebGL Working Group. 08 Jun 2016. <https://docs.google.com/presentation/d/1FkfvYgw1Fmdrx8GhqomFgL0B0ppb6LNXBPT9DyQE0qY/> (cited on page 5).
- Weitz, Dario [2020]. *Parallel Coordinates Plots*. 27 Jul 2020. <https://medium.com/data-science/parallel-coordinates-plots-6fcfa066dcb3> (cited on page 3).
- Wikipedia [2025]. *SVG*. 07 Dec 2025. <https://wikipedia.org/wiki/SVG> (cited on page 4).