Virtual Bronchoscopic Pathfinder (VBP): An Open-Source Web-Based System for Airway Segmentation, Cost-Field Path Planning, and Cross-Device 3D Navigation

Young Kim , Sunggyu Choi , Chulmin Park , Woojin Park , Doohee Lee *

Posted Date: 20 March 2026

doi: 10.20944/preprints202603.1617.v1

Keywords: Virtual Bronchoscopic Pathfinder (VBP); Virtual Bronchoscopic Navigation (VBN); pulmonary airway segmentation; topology-preserving skeletonization; bidirectional Dijkstra algorithm; vtk.js; lung cancer

 doi:10.20944/preprints202603.1617.v1

Article

Virtual Bronchoscopic Pathfinder (VBP): An Open-Source Web-Based System for Airway Segmentation, Cost-Field Path Planning, and Cross-Device 3D Navigation

Young Kim [1], Sunggyu Choi [2], Chulmin Park [3], Woojin Park [4,5] and Doohee Lee [4,*]

[1] Department of Computer Science and Engineering, Seoul National University, 1 Gwanak-ro, Gwanak-gu, Seoul 08826, Republic of Korea
[2] Department of Computer Science Engineering, College of IT, Kangwon National University, Chuncheon 24341, Republic of Korea
[3] Growth Strategy Division, DOUZONEBIZON Co., Ltd., Chuncheon 24465, Republic of Korea
[4] Department of Research and Development, ZIOVISION Co., Ltd., Chuncheon 24341, Republic of Korea
[5] Department of Data Science, Graduate School, Kangwon National University, Chuncheon 24341, Republic of Korea
* Correspondence: getback9@gmail.com

Abstract

Virtual Bronchoscopic Navigation (VBN) is a critical tool for guiding bronchoscopes toward peripheral pulmonary lesions (PPLs), yet its widespread clinical adoption has been limited by the high cost of proprietary software and the fragility of segmentation-dependent path planning in distal airways. In this study, we present **Virtual Bronchoscopic Pathfinder (VBP)**, a complete, open-source, web-based VBN system that addresses both barriers. VBP integrates five components: (i) a connectivity-aware deep learning model for pulmonary airway segmentation incorporating Connectivity-Aware Surrogate (CAS) and Local-Sensitive Distance (LSD) modules; (ii) TotalSegmentator for automated tumor localization; (iii) a topology-preserving 3D thinning algorithm implemented in C++ for centerline extraction; (iv) a bidirectional Dijkstra algorithm operating on a three-tier anatomical cost field (centerline cost 1.0, airway lumen cost 10.0, parenchyma cost 100.0) to guarantee continuous path generation even under partial skeleton disconnection; and (v) a zero-footprint browser-based visualization interface built on the vtk.js engine, providing synchronized 2D axial viewing and interactive 3D volume rendering. VBP was validated on 306 thin-section CT series (154 subjects) from the public Lung-PET-CT-Dx dataset, achieving a path-generation success rate of 100% across the anatomically valid cohort. The system is publicly accessible at https://vbn.ziovision.ai and was confirmed to operate without client-side installation across desktop, laptop, and mobile device configurations. These results demonstrate that a reliable, accessible, and scalable VBN system can be constructed entirely from open-source components, offering a practical foundation for imaging informatics research and future intra-procedural bronchoscopic guidance.

Keywords: Virtual Bronchoscopic Pathfinder (VBP); Virtual Bronchoscopic Navigation (VBN); pulmonary airway segmentation; topology-preserving skeletonization; bidirectional Dijkstra algorithm; vtk.js; lung cancer

1. Introduction

The diagnosis and treatment of peripheral pulmonary lesions (PPLs) represent a critical challenge in respiratory medicine, requiring precise navigation through the complex bronchial tree [1,2]. Virtual Bronchoscopic Navigation (VBN) addresses this challenge by reconstructing three-dimensional airway structures from Computed Tomography (CT) images, thereby providing a pre-procedural roadmap for guiding flexible bronchoscopes toward peripheral targets [3]. However, a persistent limitation of

 doi:10.20944/preprints202603.1617.v1

existing deep learning-based segmentation models is the breakage of small distal bronchi, which disrupts topological continuity and severely undermines the reliability of computer-generated navigation paths [4].

To address these connectivity issues, recent research has introduced specialized training objectives such as the Connectivity-Aware Surrogate (CAS) and Local-Sensitive Distance (LSD) modules, which explicitly enforce topological integrity throughout the bronchial tree [4]. Once a topologically complete airway mask is obtained, extracting a robust 1-voxel-thickness centerline via topology-preserving 3D thinning is essential to construct the graph structure required for path planning [5]. Building upon this centerline graph, an optimal trajectory from the proximal trachea to the target lesion can be efficiently computed using Dijkstra's shortest path algorithm [6].

A substantial body of prior work has addressed individual components of the VBN workflow, but no existing system integrates all of them into a single open-source, browser-accessible platform. On the **benchmarking side**, Lo et al. [7] established EXACT'09, the first systematic evaluation framework for airway extraction algorithms across 15 teams and 20 CT scans, demonstrating that leakage suppression and distal branch detection are in fundamental tension for intensity-based methods. The ATM'22 challenge [8] scaled this evaluation to 500 multi-site CT volumes and confirmed that topological continuity failures in peripheral bronchi remain the defining unsolved problem. On the **deep learning segmentation side**, Nadeem et al. [9] proposed a freeze-and-grow propagation strategy combined with a 3D U-Net to balance completeness against leakage in COPD cohorts, yet the system was designed for disease quantification rather than navigation and does not address downstream path planning. Zheng et al. [10] identified class-wise gradient imbalance as the root cause of peripheral bronchi erosion during CNN training and proposed group supervision with a union loss, but skeleton discontinuities persist in distal airways and no path-planning or visualization layer was provided. Qin et al. [11] introduced feature recalibration and attention distillation to improve sensitivity to tenuous bronchioles simultaneously with arterioles, achieving strong performance on EXACT'09, yet the work addresses only the segmentation stage. Wang et al. [12] presented NaviAirway, a bronchiole-sensitive pipeline with a topology-preserving loss and iterative teacher–student training to improve high-generation bronchiole detection; the system focuses exclusively on segmentation and has not been released as a deployable end-to-end platform. On the **open-source VBN systems side**, Nardelli et al. [13] implemented a 3D Slicer extension for virtual bronchoscopy validated in ex vivo animal models, but the system requires local installation and electromagnetic tracking hardware with no deep learning component. Ramírez et al. [14] presented BronchoX, a multiplatform graph-based planning tool that computes physician-readable bronchoscopic path instructions using classical region-growing segmentation, with no mechanism to bridge distal skeleton breakage and no browser interface. Lervik Bakeng et al. [15] developed Fraxinus on the CustusX image-guided therapy platform with intraoperative CT overlay, but it is a desktop application requiring proprietary hardware unsuitable for large-scale pre-procedural batch processing. On the **path planning and clinical side**, Graham et al. [16] demonstrated that graph-theoretic path planning on CT-derived airway trees is clinically feasible, but the approach depends entirely on the completeness of region-growing segmentation and cannot tolerate skeleton breakage. Chen et al. [17] established the clinical value of automated bronchoscopic path planning in a prospective multicentre robotic bronchoscopy trial (96.2% successful lesion localization), yet the system relies on a closed commercial stack with no open-source or web-accessible counterpart. None of these systems unifies connectivity-aware deep learning segmentation, a skeleton-breakage-resilient cost-field planner, multi-target tumor localization, and a zero-footprint cross-device web interface into a single, publicly deployable pipeline validated at scale.

In this study, we present **Virtual Bronchoscopic Pathfinder (VBP)**, a complete, open-source, web-based VBN system that addresses all of these gaps in a single deployable platform, accessible at https://vbn.ziovision.ai. Specifically, the contributions of VBP are as follows: (i) automated airway segmentation using a connectivity-aware deep learning model [4]; (ii) automated tumor localization via TotalSegmentator [18]; (iii) high-performance C++ implementation of a topology-preserving

 doi:10.20944/preprints202603.1617.v1

skeletonization algorithm [5,19]; (iv) multi-target bidirectional Dijkstra-based path planning with an anatomically hierarchical cost field [6]; and (v) a zero-footprint web-based visualization interface built on the vtk.js engine [20], enabling interactive 3D review directly within a standard browser [21]. The entire VBP pipeline was validated on 306 thin-section CT series from 154 subjects in the large-scale public Lung-PET-CT-Dx dataset [22]. This work is primarily a system paper: the objective is to demonstrate the completeness, feasibility, and open accessibility of the integrated pipeline rather than to advance the state of the art in any individual component. Accordingly, evaluation focuses on end-to-end path generation success and system-level operability across a large, heterogeneous CT cohort.

2. Materials and Methods

An overview of the VBP pipeline is shown in Figure 1. The system comprises four sequential stages executed on a shared file system: (i) automated airway and tumor segmentation from raw CT data; (ii) topology-preserving 3D skeletonization of the airway lumen mask; (iii) cost-field-based multi-target path planning via bidirectional Dijkstra search; and (iv) zero-footprint web-based 3D visualization. The right column of Figure 1 illustrates the representative output at each stage, from the raw axial CT stack through progressively refined anatomical structures to the final navigation path rendered in the browser interface. Each stage is described in detail in Sections 2.1–2.5.

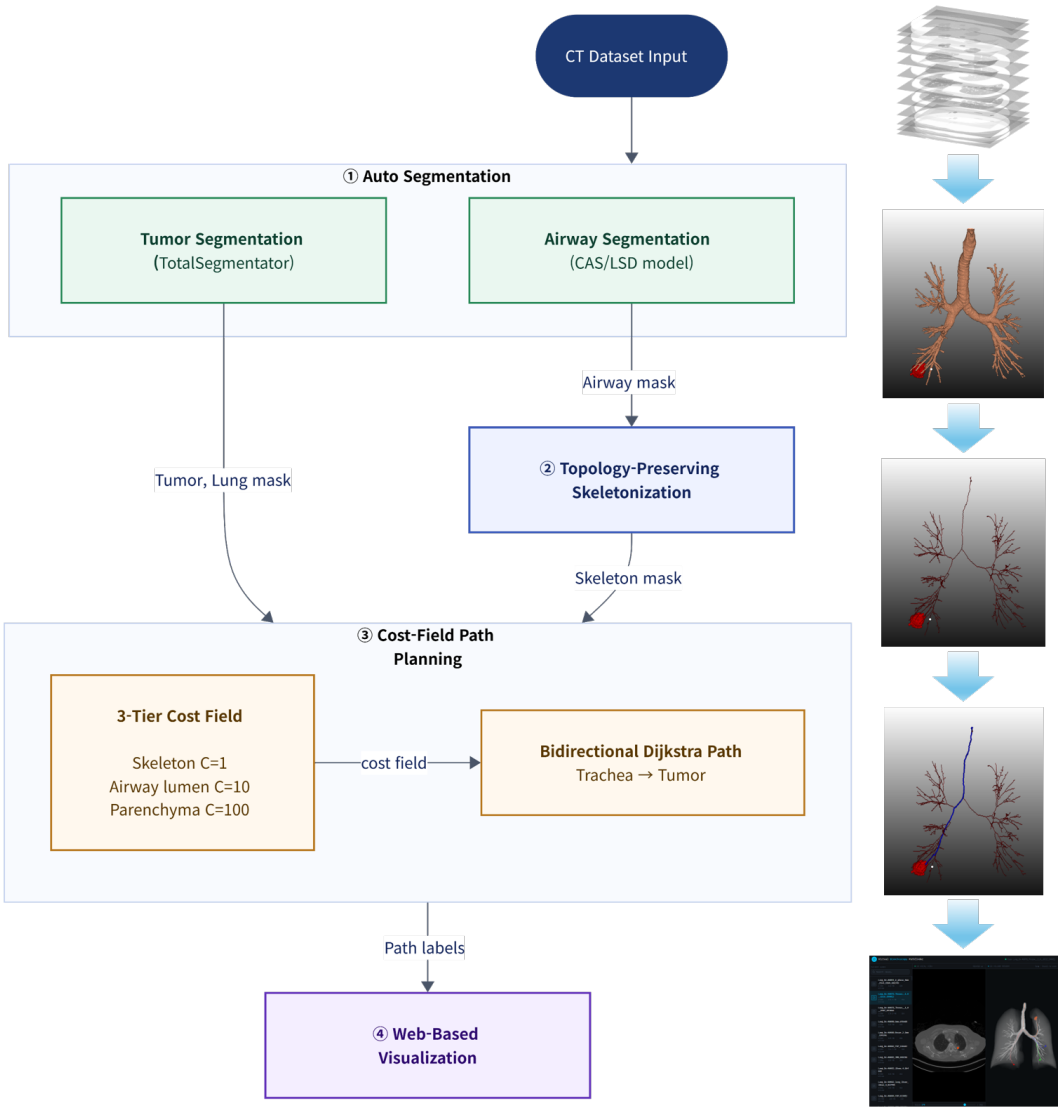

Figure 1. Overview of the Virtual Bronchoscopic Pathfinder (VBP) pipeline. The left column shows the four processing stages and their data flow; the right column shows representative outputs at each stage.

 doi:10.20944/preprints202603.1617.v1

2.1. Dataset

We utilized the **Lung-PET-CT-Dx** dataset, a large-scale public repository of pathologically verified lung cancer cases provided by The Cancer Imaging Archive (TCIA) [22]. Although no publicly available dataset provides bronchoscopy-specific ground truth (e.g., per-generation airway labels or procedural navigability annotations), Lung-PET-CT-Dx offers the largest available cohort of real lung cancer CT data with confirmed pathological diagnoses, making it the most appropriate resource for large-scale feasibility validation of a VBN pipeline. The dataset encompasses three histological subtypes coded in the subject identifier prefix: adenocarcinoma (*Lung_Dx-A*, $n = 116$), large-cell carcinoma (*Lung_Dx-B*, $n = 9$), and squamous cell carcinoma (*Lung_Dx-G*, $n = 29$), yielding 154 unique subjects in total. CT acquisitions were performed between 2009 and 2011 using scanners from three manufacturers: Siemens (82.0%), Philips (12.1%), and GE Medical Systems (5.9%).

To ensure the spatial resolution required for peripheral airway reconstruction, we applied the following inclusion criteria. Only CT series containing **200 or more axial slices** were retained, as this threshold reliably identifies sub-millimetre acquisitions (typically 1.0 mm slice thickness) that preserve the structural continuity of distal bronchi [4]. In addition, only structural CT data stored in 16-bit grayscale format were included; multimodal fused series containing PET information were excluded. These criteria yielded a final cohort of **306 series from 154 subjects** (totalling approximately 54.6 GB and 105,931 axial images). The key characteristics of the retained series are summarised in Table 1.

Table 1. Summary of the CT series retained from the Lung-PET-CT-Dx dataset after quality filtering.

Characteristic	Value
Source repository	Lung-PET-CT-Dx (TCIA) [22]
Total series (after filtering)	306
Unique subjects	154
Adenocarcinoma (*Lung_Dx-A*)	116
Large-cell carcinoma (*Lung_Dx-B*)	9
Squamous cell carcinoma (*Lung_Dx-G*)	29
CT scanner manufacturer	Siemens 251 (82.0%), Philips 37 (12.1%), GE 18 (5.9%)
Acquisition period	2009–2011
Nominal slice thickness	1.0 mm (majority); 3–5 mm (minority)
Axial slices per series	200–852 (mean 346, median 350)
Reconstruction method	FBP 278 (90.8%), Iterative (iDose/IMR) 28 (9.2%)
Total axial images	105,931
Total data volume	≈54.6 GB

Regarding tumor localization, we did not use the XML bounding-box annotations supplied with the dataset [22]. Instead, we performed independent voxel-wise tumor segmentation using TotalSegmentator [18], as described in Section 2.2, to obtain the precise 3D masks required for center-of-mass (COM) computation in the path-planning stage. Of the 306 retained series processed through VBP, 33 (10.8%) produced structurally implausible airway masks attributable to scanner-specific acquisition artifacts; these cases are analysed in detail in Section 4.4.

2.2. Automated Segmentation

Segmentation of the pulmonary airway and tumor was performed using two independent deep learning frameworks.

- **Airway Segmentation:** We employed a connectivity-aware model incorporating the **Connectivity-Aware Surrogate (CAS)** and **Local-Sensitive Distance (LSD)** training modules [4]. These modules explicitly penalize topological breakage during training, ensuring that the resulting binary mask maintains the structural continuity of distal bronchi—a prerequisite for generating uninterrupted navigation paths.

 doi:10.20944/preprints202603.1617.v1

- **Tumor Segmentation:** Automated voxel-wise segmentation was performed using **TotalSegmentator** [18], an open-source framework trained to delineate 104 anatomical structures, including pulmonary masses and nodules. The resulting 3D tumor mask was used to compute the target center of mass (COM), which serves as the endpoint for the path-planning algorithm.

2.3. Airway Skeletonization

The segmented airway mask was reduced to a 1-voxel-thickness medial centerline using a **topology-preserving 3D thinning algorithm** [5], implemented in C++ and ported from the Java-based Skeletonize3D plugin of the ImageJ/Fiji framework [19]. The algorithm iteratively removes border voxels from the binary mask while strictly preserving three topological properties: it retains **endpoint voxels** (voxels with exactly one foreground neighbor) to prevent premature termination of branch tips; it tests **Euler invariance** via a pre-computed 256-entry look-up table (LUT) to ensure no topological handles are created or destroyed; and it verifies **simple-point** status through an octree-based connectivity analysis to confirm that removal of a candidate voxel does not disconnect any foreground component.

The outer loop cycles through the six orthogonal border directions (North, South, East, West, Up, Bottom) in sequence. In each direction pass, all voxels that (a) face an empty voxel in the current border direction, (b) are not endpoint voxels, (c) satisfy Euler invariance, and (d) are simple points are collected into a candidate list and then removed in batch. This cycle repeats until all six passes produce no change, at which point the output is a topologically equivalent 1-voxel-thick centerline skeleton. The complete procedure is formalized in Algorithm 1.

Algorithm 1 Topology-Preserving 3D Thinning (`computeThinImage`)

Require: Binary airway mask I of size $W \times H \times D$
Ensure: Skeleton image O (1-voxel-thick medial centerline)

$O \leftarrow$ copy of I (binarized to $\{0,1\}$)
Pre-compute Euler LUT $\mathcal{E}[256]$ ▷ fillEulerLUT
unchangedBorders $\leftarrow 0$
while *unchangedBorders* < 6 **do**
  unchangedBorders $\leftarrow 0$
  for $d \in \{\mathrm{N}, \mathrm{S}, \mathrm{E}, \mathrm{W}, \mathrm{U}, \mathrm{B}\}$ **do** ▷ 6 border directions
    candidates $\leftarrow \emptyset$; *noChange* $\leftarrow$ **true**
    for each foreground voxel $v = (x, y, z)$ in O **do**
      if v is not a border voxel in direction d **then continue**
      end if
      if IsEndPoint(O, v) **then continue**
      end if ▷ ≤ 1 foreground neighbor
      $\mathcal{N} \leftarrow$ 27-neighborhood of v in O
      if not IsEulerInvariant$(\mathcal{N}, \mathcal{E})$ **then continue**
      end if
      if not IsSimplePoint$(\mathcal{N})$ **then continue**
      end if ▷ octree connectivity
      candidates.push(v)
    end for
    for each $v \in$ *candidates* **do**
      Re-fetch $\mathcal{N} \leftarrow$ 27-neighborhood of v in O ▷ concurrent safety
      if IsSimplePoint$(\mathcal{N})$ **then**
        $O[v] \leftarrow 0$; *noChange* $\leftarrow$ **false**
      end if
    end for
    if *noChange* **then** *unchangedBorders* $+= 1$
    end if
  end for
end while
return O

 doi:10.20944/preprints202603.1617.v1

2.4. Path Planning

The path-planning stage takes the skeleton S produced by Section 2.3, the airway mask M_{air}, and up to 25 tumor masks $\{M_k\}_{k=1}^{K}$ as input, and outputs a single integrated NIfTI volume encoding all navigation paths. The stage is organized into three sub-steps: cost field construction, multi-target path finding, and output labeling.

2.4.1. Cost Field Construction

Rather than treating path planning as a pure graph traversal on the skeleton alone, we construct a continuous **3D cost field** $C : V \to \mathbb{R}_{>0}$ over the entire CT volume V so that the planner can bridge local skeleton discontinuities arising from segmentation imperfections. Each voxel v is assigned a traversal cost according to its anatomical classification:

- **Skeleton voxel** ($v \in S$): $C(v) = 1.0$—the planner preferentially follows the bronchial medial axis.
- **Airway lumen** ($v \in M_{\text{air}} \setminus S$): $C(v) = 10.0$—a low-penalty fallback that keeps the path within the airway when skeleton gaps exist.
- **Lung parenchyma / tissue** (otherwise): $C(v) = 100.0$—a high-penalty region traversed only when the target lies beyond the segmentable distal bronchi.

This three-tier weighting guarantees that any reachable tumor COM always has a valid path from the trachea, while enforcing a strong preference for anatomically meaningful routes. The construction procedure is formalized in Algorithm 2.

Algorithm 2 Cost Field Construction for VBP Path Planning

Require: Skeleton S, Airway mask M_{air}, Volume V of size $W \times H \times D$
Ensure: Cost field $C[0 \ldots |V| - 1]$ of type `float32`

```
for each voxel index i ∈ [0, |V|) do
    if S[i] = 1 then
        C[i] ← 1.0                          ▷ Skeleton centerline
    else if M_air[i] = 1 then
        C[i] ← 10.0                         ▷ Airway lumen fallback
    else
        C[i] ← 100.0                        ▷ Lung parenchyma / tissue
    end if
end for
return C
```

2.4.2. Multi-Target Bidirectional Dijkstra Path Finding

The core path-finding engine is a **bidirectional Dijkstra algorithm** [6] adapted for 3D voxel grids with 26-neighbor connectivity, derived from the open-source `dijkstra3d` library [23]. Bidirectional search simultaneously expands two priority queues—one forward from the source and one reverse from the target—and terminates when the two frontiers meet, reducing the typical search radius from $O(|V|)$ to $O(\sqrt{|V|})$ and yielding substantial speedups for large CT volumes.

The **source** v_{src} is automatically identified as the skeleton voxel with the maximum Z-coordinate, corresponding to the superior margin of the trachea. For a patient case containing K tumors ($K \leq 25$), the same cost field C is reused for all K path-finding calls, eliminating redundant computation. Each target $v_k = \text{COM}(M_k)$ is the center of mass of the k-th tumor mask. The complete multi-target path-finding procedure is formalized in Algorithm 3.

 doi:10.20944/preprints202603.1617.v1

Algorithm 3 Multi-Target Bidirectional Dijkstra Path Finding

Require: Cost field C, Skeleton S, Tumor masks $\{M_k\}_{k=1}^{K}$, $K \leq 25$
Ensure: Output label volume L (UINT8), initialized to M_{air} label 2

$v_{\text{src}} \leftarrow \arg\max_{v \in S} Z(v)$ ▷ Superior tracheal voxel
$L \leftarrow$ encode S as label 3, M_{air} as label 2, background as 1
for $k = 1$ **to** K **do**
  $v_k \leftarrow \text{round}(\text{COM}(M_k))$ ▷ Target: tumor center of mass
  //—Bidirectional Dijkstra —
  $dist_{\text{fwd}}[v] \leftarrow +\infty \; \forall v$; $dist_{\text{fwd}}[v_{\text{src}}] \leftarrow 0$
  $dist_{\text{rev}}[v] \leftarrow +\infty \; \forall v$; $dist_{\text{rev}}[v_k] \leftarrow 0$
  $Q_{\text{fwd}}.\text{push}(0, v_{\text{src}})$; $Q_{\text{rev}}.\text{push}(0, v_k)$
  $cost^* \leftarrow +\infty$; $v_{\text{meet}} \leftarrow$ null
  $forward \leftarrow$ **false**
  while $Q_{\text{fwd}} \neq \emptyset$ **and** $Q_{\text{rev}} \neq \emptyset$ **do**
    $forward \leftarrow \neg forward$ ▷ Alternate between queues
    if $forward$ **then**
      $u \leftarrow Q_{\text{fwd}}.\text{pop\_min}()$
      if $dist_{\text{rev}}[u] < +\infty$ **then**
        $c \leftarrow |dist_{\text{fwd}}[u]| + |dist_{\text{rev}}[u]| + C[v_k] - C[u]$
        if $c < cost^*$ **then** $cost^* \leftarrow c$; $v_{\text{meet}} \leftarrow u$
        else break
        end if
      end if
      Relax 26-neighbors of u in $dist_{\text{fwd}}$, Q_{fwd}, $parents_{\text{fwd}}$
    else
      $u \leftarrow Q_{\text{rev}}.\text{pop\_min}()$
      if $dist_{\text{fwd}}[u] < +\infty$ **then**
        $c \leftarrow |dist_{\text{fwd}}[u]| + |dist_{\text{rev}}[u]| + C[v_k] - C[u]$
        if $c < cost^*$ **then** $cost^* \leftarrow c$; $v_{\text{meet}} \leftarrow u$
        else break
        end if
      end if
      Relax 26-neighbors of u in $dist_{\text{rev}}$, Q_{rev}, $parents_{\text{rev}}$
    end if
  end while
  //—Path Reconstruction & Labeling —
  if $v_{\text{meet}} \neq$ null **then**
    $P_k \leftarrow$ backtrack from v_{meet} via $parents_{\text{fwd}}$ and $parents_{\text{rev}}$
    for each $v \in P_k$ **do**
      $L[v] \leftarrow 10k + 2$ ▷ Path#k label
    end for
    $L[v_k] \leftarrow 10k + 1$ ▷ Tumor#k label
  end if
end for
return L

2.4.3. Output Label Schema

The final output is stored as a NIfTI file using the following label encoding, which is shared between the C++ backend and the web visualization system:

- Label **1**: Lung mask (background context)
- Label **2**: Airway lumen
- Label **3**: Skeletonized centerline tree
- Label $\mathbf{10k + 1}$ ($k = 1, \ldots, 25$): Tumor k
- Label $\mathbf{10k + 2}$ ($k = 1, \ldots, 25$): Navigation path to Tumor k

 doi:10.20944/preprints202603.1617.v1

2.5. Web-Based Visualization System

The complete VBP pipeline executes five sequential stages as independent modules: (1) CT filtering and NIfTI conversion; (2) airway and tumor segmentation; (3) C++ topology-preserving skeletonization; (4) multi-target Dijkstra path planning; and (5) web-based visualization. Intermediate outputs are written to a shared file system, allowing individual stages to be re-run or replaced without reprocessing the entire pipeline. To maximize clinical accessibility, the final stage delivers results through a **zero-footprint** web-based interface that allows clinicians to review VBP outputs on any device with a modern browser—including desktop workstations, laptops, and mobile devices—without installing dedicated software. The VBP visualization system is publicly accessible at https://vbn.ziovision.ai. The system architecture separates concerns into a server-side processing backend and a client-side rendering frontend, following design principles established in prior web-based medical imaging platforms [20,21].

2.5.1. Backend Architecture

The backend is implemented as a **Flask** (Python) application exposing a RESTful API. It handles two distinct data pathways.

For **2D axial viewing**, the server asynchronously processes CT and mask NIfTI volumes upon case selection. The mask is simultaneously colorized using a pre-computed look-up table (LUT) that maps the label schema of Section 2.4.3 to RGBA values with label-specific opacity. All slices are Base64-encoded as PNG images and delivered to the client in a single payload. A dedicated polling endpoint exposes preprocessing progress, enabling a real-time progress indicator in the frontend.

For **3D volume rendering**, a separate endpoint returns the complete mask volume as a raw `uint8` byte stream encoded in Base64, allowing the client to reconstruct the full `vtkImageData` object in one operation.

2.5.2. Frontend Interface

The frontend is a responsive single-page application (SPA) built with adaptive CSS layouts that reconfigure automatically for the available viewport. On **desktop and laptop** displays, the interface presents a **dual-panel layout**: the left panel hosts an interactive 2D axial viewer with a synchronized slice slider, and the right panel hosts a 3D volume renderer powered by **vtk.js**. A collapsible sidebar lists all available patient cases with metadata and provides a search filter. On **mobile devices**, the layout switches to a vertically stacked single-column arrangement, with touch-optimised controls for slice navigation and 3D interaction, enabling point-of-care review directly from a smartphone or tablet. Both views are illustrated in Figure 2.

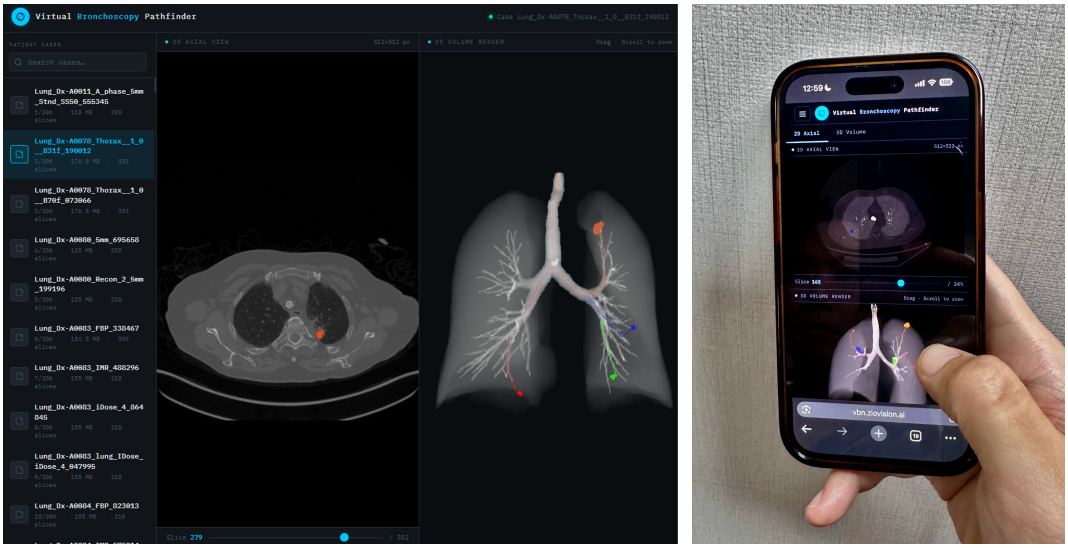

Figure 2. The VBP visualization system (https://vbn.ziovision.ai) across device form factors. (**Left, desktop/laptop**)

 doi:10.20944/preprints202603.1617.v1

The 3D volume transfer functions are configured to reflect the anatomical label hierarchy: the lung parenchyma ($\alpha = 0.01$) serves as a subtle spatial reference; the airway lumen ($\alpha = 0.10$) provides structural context; the skeletonized tree ($\alpha = 0.50$) delineates the navigable centerline; and all tumor and path labels are fully opaque ($\alpha = 1.0$). The camera is initialized in an anterior-to-posterior orientation with the superior axis aligned to Z, replicating the standard bronchoscopic planning perspective.

3. Results

3.1. Quantitative Evaluation

VBP was validated on all 306 thin-section CT series (154 subjects) from the Lung-PET-CT-Dx dataset [22]. Pipeline-level outcomes are summarised in Table 2.

Of the 306 series, 33 (10.8%) were excluded from path planning due to structurally implausible airway masks arising from scanner-specific acquisition artifacts (detailed in Section 4.4). For the remaining 273 series (89.2%), the bidirectional Dijkstra algorithm successfully generated a continuous path from the proximal trachea to the center of mass (COM) of every TotalSegmentator-detected tumor, yielding a path generation success rate of 100% within the anatomically valid subset.

Table 2. Pipeline-level outcome summary across all 306 CT series.

Metric	Value
Total series processed	306 (154 subjects)
Segmentation failure (scanner artifact)	33 (10.8%)
Type 1: Bed centre air channel	27 / 33 (81.8%)
Type 2: Bed lateral edge	5 / 33 (15.2%)
Type 3: Leg skin–clothing gap	1 / 33 (3.0%)
Anatomically valid series	273 (89.2%)
Path generation success (valid subset)	273 / 273 (100%)

3.2. Qualitative Assessment

Qualitative review confirmed that generated trajectories were anatomically plausible in representative cases. The computed paths followed the bronchial midline through distal branches reconstructed by the connectivity-aware segmentation model [4], and correctly transitioned to the airway-lumen cost layer where minor skeleton discontinuities were present. In cases where tumors were located beyond the extent of the segmentable bronchi, the planner successfully traversed the parenchyma cost layer to reach the target COM, as designed by the three-tier cost field.

3.3. Web-Based System: Cross-Device Operability

The VBP visualization system described in Section 2.5 was evaluated as a proof-of-concept for zero-footprint deployment across three device classes: desktop workstation, laptop, and mobile smartphone. The objective was to confirm basic cross-device operability rather than to conduct a formal usability study. Server-side NIfTI preprocessing and slice encoding completed without observable latency issues on the test server, after which all 2D slices and the 3D mask volume were resident in the client browser without further server calls. The 3D volume renderer produced smooth interactive rotation and zoom on desktop and laptop hardware. On mobile devices, the responsive layout and touch controls functioned correctly, confirming that point-of-care review is feasible without application installation. Operation was verified in current releases of Chrome, Firefox, and Safari. A formal evaluation of clinical usability, including task completion rates and user experience metrics, is identified as a priority for future work.

 doi:10.20944/preprints202603.1617.v1

4. Discussion

4.1. Significance of the Open-Source Pipeline

The primary contribution of this work is the end-to-end integration of individually established components—connectivity-aware segmentation, topology-preserving skeletonization, cost-field-based path planning, and browser-native 3D visualization—into **VBP**, a single, openly reproducible open-source system. Each component has been validated in isolation in prior literature [4–6,18], but no prior work has assembled them into a deployable, publicly accessible VBN pipeline. The engineering challenge of reliable integration—ensuring consistent coordinate spaces, label conventions, and data handoffs across heterogeneous tools—is non-trivial, and the result is a system that enables large-scale informatics studies that would be practically infeasible with existing commercial or desktop-only alternatives. While conventional VBN systems often depend on proprietary software licenses and dedicated workstations, VBP demonstrates that equivalent end-to-end functionality can be achieved using exclusively open-source tools and a public dataset [18,22], lowering the barrier to adoption for research institutions worldwide.

4.2. Robustness Through Cost Field Design and Multi-Target Support

A technically novel aspect of VBP is the three-tier anatomical cost field, which, to the best of our knowledge, has not been previously described in the VBN literature. Prior path-planning approaches for bronchoscopic navigation operate exclusively on the skeleton graph and therefore fail whenever a skeleton discontinuity interrupts the route [14,16]. By constructing a dense volumetric cost field that assigns graded traversal penalties to skeleton voxels, airway-lumen voxels, and parenchymal voxels, VBP guarantees that a continuous path from the trachea to any reachable tumor COM always exists, regardless of local skeleton breakage. This graceful degradation property—preferring skeleton routes but falling back through the lumen and, if necessary, parenchyma—directly addresses the fundamental fragility of skeleton-only planners in the presence of segmentation imperfections [4]. Furthermore, the cost field is constructed only once per patient case and reused for all K tumor targets ($K \leq 25$), making the multi-target extension computationally efficient. The single integrated output label volume encodes all tumor-path pairs in a compact UINT8 format, enabling seamless handoff to the web visualization layer.

4.3. Limitations and Future Work

Excluding the 33 cases with scanner-artifact segmentation failures analysed in Section 4.4, the remaining 273 anatomically valid series are subject to two additional path-planning failure modes that reflect the boundaries of the current cost-field design, both illustrated in Figure 3 using case *Lung_Dx-A0110_5mm_197561*. Although the precise prevalence of each mode across all 273 cases was not quantified in this study—an acknowledged limitation—representative examples are presented to characterise their mechanisms and to guide future cost-field refinements.

Trachea-proximal path shortcuts (yellow arrow). When a tumor's center of mass (COM) lies in immediate proximity to the trachea or a main bronchus, the bidirectional Dijkstra algorithm selects the geometrically shortest—and cost-minimizing—path, which may be a direct trajectory from the tracheal centerline to the tumor COM. This shortcut bypasses the normal descending bronchial routing entirely and yields a path that is graph-theoretically optimal but clinically non-navigable: a bronchoscope cannot reach the target via such a trajectory. The root cause is a cost-field formulation that does not encode anatomical depth as a constraint—the planner has no incentive to descend to a minimum bronchial generation before reaching the target. Incorporating a minimum-depth criterion, or weighting edge costs by anatomical level, would prevent premature termination.

Excessive parenchymal traversal for peripheral tumors (red arrows). When a tumor is located at or beyond the peripheral boundary of the lung—where the CAS/LSD model [4] produces no segmentable distal bronchi—the planner is forced entirely into the parenchyma cost layer ($C = 100.0$) for the distal segment of the route. The resulting path is structurally valid and reaches the target

 doi:10.20944/preprints202603.1617.v1

COM, but the extended trajectory through lung parenchyma far exceeds the distance achievable by a navigable bronchial route, rendering the plan clinically meaningless. This failure mode is not a segmentation error but a cost-field design limitation: the three-tier weighting guarantees path existence at the expense of anatomical plausibility when distal airway coverage is insufficient. A coverage-aware penalty—for example, scaling parenchymal cost by the distance from the nearest skeleton voxel—would discourage excessively peripheral excursions.

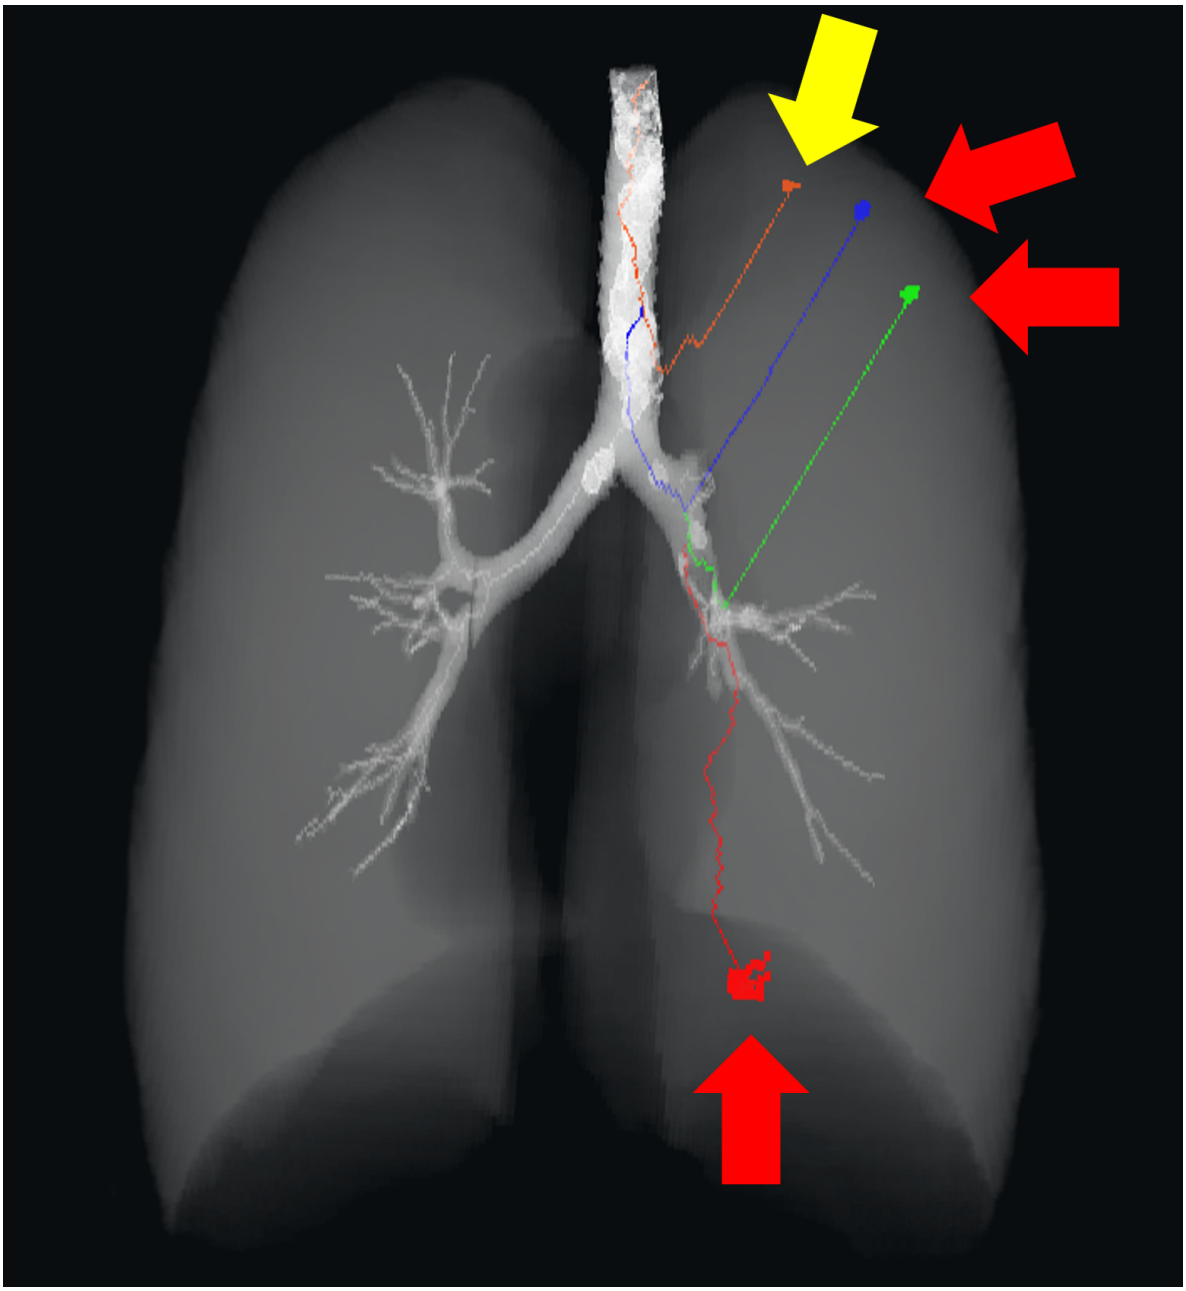

Figure 3. Representative path-planning failure cases of VBP (case *Lung_Dx-A0110_5mm_197561*). (**Yellow arrow**) A tumor whose center of mass lies in immediate proximity to the trachea causes the bidirectional Dijkstra algorithm to select a direct path from the tracheal centerline to the target, bypassing normal bronchial routing and yielding a clinically non-navigable trajectory. (**Red arrows**) Tumors located at or beyond the peripheral boundary of the lung, where no segmentable distal bronchi are present, force the planner to traverse an extended parenchymal cost layer. The resulting paths are geometrically valid but anatomically implausible, spanning an excessive distance through lung tissue rather than following a bronchoscopic route.

It should be noted that the explicit characterisation of these failure modes is itself a contribution of this work. Existing commercial VBN systems do not report failure taxonomies in the

 doi:10.20944/preprints202603.1617.v1

open literature, making systematic comparison difficult. By documenting failure mechanisms and their DICOM-metadata signatures, VBP provides a reproducible basis for future benchmarking and algorithmic improvement.

Future work will pursue four directions to address these limitations. First, a CT-anatomy-aware pre-filter will be introduced to exclude non-airway structures from the segmentation mask before skeletonization. Second, a minimum bronchial generation depth constraint will be added to the path-planning cost function to prevent anatomically implausible shortcuts in proximal cases. Third, a prospective study with independent thoracic radiology review will be conducted to formally quantify the prevalence of each path-planning failure mode and to assess the clinical navigability of generated routes. Fourth, VBP will be extended to support real-time intra-procedural guidance by integrating live bronchoscopic image registration with the pre-computed virtual path, and the web backend will be optimised for low-latency multi-user collaboration to support multidisciplinary case review.

4.4. Airway Segmentation Failure Analysis

Manual inspection of all 306 series processed through VBP after airway segmentation revealed that 33 cases (10.8%) produced structurally implausible masks that could not be used for skeletonization or path planning. These failures arise from scanner-specific acquisition artifacts and weaknesses in the segmentation model itself. They cluster into three geometrically distinct types illustrated in Figure 4.

Type 1—Bed center air channel (27 cases, 81.8%): The segmentation model incorrectly labeled a narrow longitudinal air channel located beneath the patient table as the airway lumen (yellow circle, Figure 4, top row). This artifact is predominantly associated with **Philips** scanners: 26 of the 27 affected series record the manufacturer as *Philips* and the study description as *Chest 3D IMR* or *Chest 3D*, corresponding to the Philips Ingenuity TF PET/CT platform whose carbon-fibre couch contains a hollow longitudinal channel. The remaining case originates from a **GE Medical Systems** scanner (series: *Recon 2 5mm*, study: *chest.3d*), indicating that a geometrically similar bed structure is present in at least one GE platform as well. Because this structure presents as a continuous, low-attenuation (≈ -1000 HU) tubular region at the inferior boundary of the CT field of view, it satisfies the same Hounsfield criteria as tracheal air.

Type 2—Bed lateral edge (5 cases, 15.2%): The tapered lateral edge of the scanner bed produced a narrow wedge-shaped air pocket that was misclassified as a bronchial segment (red circle, Figure 4, middle row). This type spans two scanner platforms: three cases arise from **Siemens** whole-body acquisitions (series: *CT WB 3.0 B30f*; studies: *PET02/PET08WholebodyOnly Adult*; subjects A0164, A0216, A0221), and two cases arise from **Philips** chest acquisitions (series: *IMR* and *lung IDose iDose 4*; study: *Chest 3D IMR*; subject A0084). The mechanism is related to Type 1 but geometrically distinct: rather than the central channel, it is the sloped outer rim of the bed that creates the artifact.

Type 3—Leg skin–clothing air gap (1 case, 3.0%): In one Siemens whole-body case (*Lung_Dx-A0220*, series: *CT WB 1.0 B30f*, study: *PET08WholebodyOnly Adult*, 677 slices), the imaging field extends from the head to the lower extremities. When this volume is passed to a thorax-trained segmentation model, the low-attenuation gap between the patient's thigh skin and clothing is misidentified as airway lumen (blue circle, Figure 4, bottom row), producing a skeleton that diverges entirely from the bronchial tree.

The root cause common to all three failure types is the limited generalizability of the airway segmentation model under real-world acquisition diversity. Because the entire downstream pipeline—skeletonization, path planning, and visualization—depends entirely on the quality of the segmentation mask, a failure at this stage propagates unconditionally to all subsequent stages, rendering the pipeline output clinically meaningless regardless of the robustness of later components. This upstream dependency represents the most critical vulnerability of any segmentation-driven VBN system. While heuristic countermeasures such as DICOM-metadata-based triage or Hounsfield-distribution pre-filtering could reduce the incidence of specific artefact types, such approaches introduce their own selection bias and do not address the underlying problem. The fundamental requirement is a generalizable, high-performance airway segmentation model that is robust to the full spectrum of clinical

 doi:10.20944/preprints202603.1617.v1

acquisition variability—including heterogeneous scanner platforms, reconstruction protocols, and patient populations—without relying on post-hoc filtering. Development and prospective validation of such a model is identified as the highest-priority direction for future work.

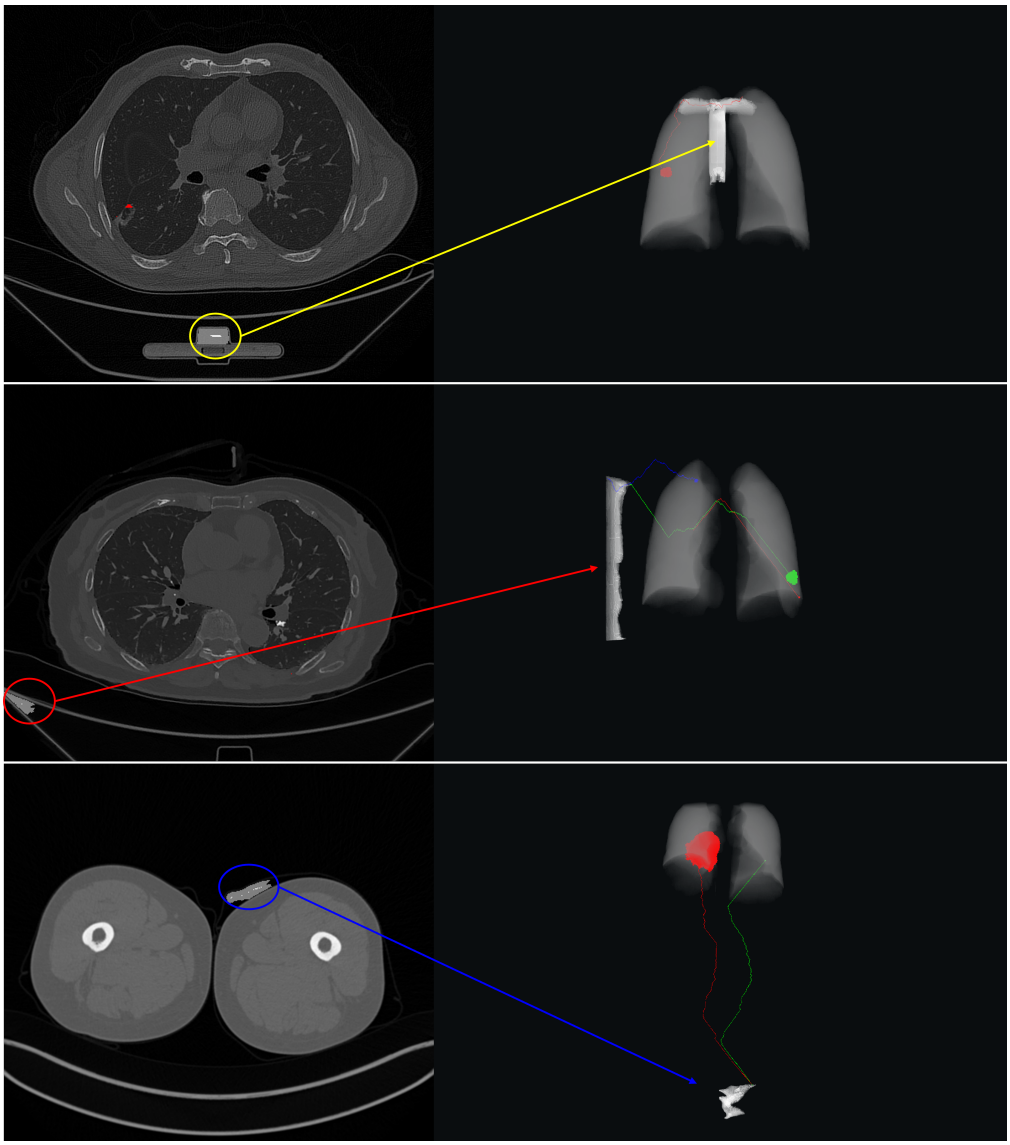

Figure 4. Three types of airway segmentation failure observed in the Lung-PET-CT-Dx dataset. For each row, the axial CT slice (left) and the resulting 3D path-planning output (right) are shown. (**Top, yellow**) Type 1: the central longitudinal air channel of a Philips scanner bed is misclassified as the trachea (case *Dx-A0083_FBP_338467*; Philips, *Chest 3D IMR*). (**Middle, red**) Type 2: the tapered lateral edge of the scanner bed is misclassified as a bronchial segment (case *Dx-A0084_IMR_575314*; Philips, *Chest 3D IMR*). (**Bottom, blue**) Type 3: in a Siemens whole-body acquisition (*CT WB 1.0 B30f*, 677 slices), the air gap between the patient's thigh skin and clothing is misidentified as airway lumen (case *Dx-A0220_CT_WB_1_0_B30f_160816*).

5. Conclusions

We developed and validated **Virtual Bronchoscopic Pathfinder (VBP)**, a complete, open-source, web-based VBN system that integrates automated airway and tumor segmentation, topology-preserving skeletonization, multi-target bidirectional Dijkstra path planning with a three-tier anatomical cost field, and zero-footprint browser-based 3D visualization. Validated across 306 heterogeneous CT series from 154 lung cancer subjects, VBP achieved a path generation success rate of 100% in the 273 anatomically valid cases. The three-tier cost field—a design not previously described in the VBN literature—provides a principled mechanism for tolerating skeleton discontinuities without pipeline failure. By building exclusively on open-source components and a public dataset, VBP demonstrates

 doi:10.20944/preprints202603.1617.v1

that reliable, end-to-end bronchoscopic navigation guidance can be delivered without proprietary hardware or software. The system is openly accessible at https://vbn.ziovision.ai and serves as a scalable, reproducible foundation for imaging informatics research in bronchoscopic navigation, with a modular architecture that supports further extension toward intra-procedural guidance and multi-user clinical workflows.

Author Contributions: Conceptualization, D.L.; methodology, Y.K.; software, W.P. and D.L.; validation, C.P.; formal analysis, Y.K.; investigation, Y.K.; resources, S.C.; data curation, C.P.; writing—original draft preparation, Y.K.; writing—review and editing, W.P., D.L., S.C. and C.P.; visualization, Y.K.; supervision, D.L.; project administration, D.L. All authors have read and agreed to the published version of the manuscript.

Funding: This study was supported by the Global Innovation Special Zone Innovation Project(RS-2024-00488398) of the Ministry of SMEs and Startups.

Institutional Review Board Statement: Not applicable (publicly available dataset).

Informed Consent Statement: Not applicable.

Data Availability Statement: The Lung-PET-CT-Dx dataset is publicly available via The Cancer Imaging Archive [22]. The VBP system is publicly accessible at https://vbn.ziovision.ai. Dataset available on request from the authors.

Acknowledgments: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

VBP	Virtual Bronchoscopic Pathfinder
VBN	Virtual Bronchoscopic Navigation
PPL	Peripheral Pulmonary Lesion
CT	Computed Tomography
CAS	Connectivity-Aware Surrogate
LSD	Local-Sensitive Distance
COM	Center of Mass
LUT	Look-Up Table
TCIA	The Cancer Imaging Archive
SPA	Single-Page Application
FBP	Filtered Back Projection
HU	Hounsfield Unit

References

1. Asano, F.; Eberhardt, R.; Herth, F.J.F. Virtual Bronchoscopic Navigation as an Adjunct to Bronchoscopy: A Systematic Review and Meta-Analysis. *J. Thorac. Oncol.* **2014**, *9*, 569–576. https://doi.org/10.1097/JTO.0000000000000100.
2. Ishida, A.; Matsumoto, K.; Suzuki, M.; Akazawa, K.; Seki, R.; Kikuchi, T.; Miyazawa, T. Fluoroscopy-Guided Bronchoscopy versus Radial Endobronchial Ultrasound with a Guide Sheath for Peripheral Pulmonary Lesions: A Randomized Trial. *Chest* **2012**, *142*, 1455–1461. https://doi.org/10.1378/chest.11-2505.
3. Seijo, L.M.; Zulueta, J.J. Tools for Quaternary Prevention: Optimizing the Diagnosis of Peripheral Lung Nodules. *Arch. Bronconeumol.* **2016**, *52*, 577–578. https://doi.org/10.1016/j.arbres.2016.07.002.
4. Zhang, M.; Gu, Y. Towards Connectivity-Aware Pulmonary Airway Segmentation. *IEEE J. Biomed. Health Inform.* **2023**, *27*, 921–932. https://doi.org/10.1109/JBHI.2022.3218478.
5. Lee, T.C.; Kashyap, R.L.; Chu, C.C. Building Skeleton Models via 3-D Medial Surface/Axis Thinning Algorithms. *CVGIP Graph. Models Image Process.* **1994**, *56*, 462–478. https://doi.org/10.1006/cgip.1994.1042.
6. Dijkstra, E.W. A Note on Two Problems in Connexion with Graphs. *Numer. Math.* **1959**, *1*, 269–271. https://doi.org/10.1007/BF01386390.

 doi:10.20944/preprints202603.1617.v1

7. Lo, P.; van Ginneken, B.; Reinhardt, J.M.; Yavarna, T.; de Jong, P.A.; Irving, B.; Fetita, C.I.; Ortner, M.; Pinho, R.; Sijbers, J.; et al. Extraction of Airways from CT (EXACT'09). *IEEE Trans. Med. Imaging* **2012**, *31*, 2093–2107. https://doi.org/10.1109/TMI.2012.2209674.
8. Zhang, M.; Wu, Y.; Zhang, H.; Qin, Y.; Zheng, H.; Tang, W.; Arnold, C.; Pei, C.; Yu, P.; Nan, Y.; et al. Multi-Site, Multi-Domain Airway Tree Modeling. *Med. Image Anal.* **2023**, *90*, 102957. https://doi.org/10.1016/j.media.2023.102957.
9. Nadeem, S.A.; Hoffman, E.A.; Sieren, J.C.; Comellas, A.P.; Bhatt, S.P.; Barjaktarevic, I.Z.; Abtin, F.; Saha, P.K. A CT-Based Automated Algorithm for Airway Segmentation Using Freeze-and-Grow Propagation and Deep Learning. *IEEE Trans. Med. Imaging* **2021**, *40*, 405–418. https://doi.org/10.1109/TMI.2020.3029013.
10. Zheng, H.; Qin, Y.; Gu, Y.; Xie, F.; Yang, J.; Sun, J.; Yang, G.Z. Alleviating Class-Wise Gradient Imbalance for Pulmonary Airway Segmentation. *IEEE Trans. Med. Imaging* **2021**, *40*, 2452–2462. https://doi.org/10.1109/TMI.2021.3078828.
11. Qin, Y.; Zheng, H.; Gu, Y.; Huang, X.; Yang, J.; Wang, L.; Yao, F.; Zhu, Y.M.; Yang, G.Z. Learning Tubule-Sensitive CNNs for Pulmonary Airway and Artery-Vein Segmentation in CT. *IEEE Trans. Med. Imaging* **2021**, *40*, 1603–1617. https://doi.org/10.1109/TMI.2021.3062280.
12. Wang, A.; Tam, T.C.C.; Poon, H.M.; Yu, K.C.; Lee, W.N. NaviAirway: A Bronchiole-Sensitive Deep Learning-Based Airway Segmentation Pipeline for Planning of Navigation Bronchoscopy. arXiv Preprint arXiv:2203.04294, 2022. Available online: https://arxiv.org/abs/2203.04294.
13. Nardelli, P.; Jaeger, H.A.; O'Shea, C.; Khan, K.A.; Kennedy, M.P.; Cantillon-Murphy, P. Pre-Clinical Validation of Virtual Bronchoscopy Using 3D Slicer. *Int. J. Comput. Assist. Radiol. Surg.* **2017**, *12*, 25–38. https://doi.org/10.1007/s11548-016-1447-7.
14. Ramírez, E.; Sánchez, C.; Borràs, A.; Diez-Ferrer, M.; Rosell, A.; Gil, D. BronchoX: Bronchoscopy Exploration Software for Biopsy Intervention Planning. *Healthc. Technol. Lett.* **2018**, *5*, 177–182. https://doi.org/10.1049/htl.2018.5074.
15. Lervik Bakeng, J.B.; Hofstad, E.F.; Solberg, O.V.; Eiesland, J.; Tangen, G.A.; Amundsen, T.; Langø, T.; Reinertsen, I.; Selbekk, T.; Leira, H.O. Using the CustusX Toolkit to Create an Image Guided Bronchoscopy Application: Fraxinus. *PLoS ONE* **2019**, *14*, e0211772. https://doi.org/10.1371/journal.pone.0211772.
16. Graham, M.W.; Gibbs, J.D.; Cornish, D.C.; Higgins, W.E. Robust 3-D Airway Tree Segmentation for Image-Guided Peripheral Bronchoscopy. *IEEE Trans. Med. Imaging* **2010**, *29*, 982–997. https://doi.org/10.1109/TMI.2009.2035813.
17. Chen, A.C.; Pastis, N.J.; Mahajan, A.K.; Khandhar, S.J.; Simoff, M.J.; Machuzak, M.S.; Cicenia, J.; Gildea, T.R.; Silvestri, G.A. Robotic Bronchoscopy for Peripheral Pulmonary Lesions: A Multicenter Pilot and Feasibility Study (BENEFIT). *Chest* **2021**, *159*, 845–852. https://doi.org/10.1016/j.chest.2020.08.2047.
18. Wasserthal, J.; Breit, H.C.; Meyer, M.T.; Pradella, M.; Hinck, D.; Sauter, A.W.; Heye, T.; Boll, D.T.; Cyriac, J.; Yang, S.; et al. TotalSegmentator: Robust Segmentation of 104 Anatomic Structures in CT Images. *Radiol. Artif. Intell.* **2023**, *5*, e230024. https://doi.org/10.1148/ryai.230024.
19. Arganda-Carreras, I. Skeletonize3D (ImageJ Plugin). ImageJ/Fiji Plugin. Available online: https://imagej.net/plugins/skeletonize3d, 2023. Accessed: 7 February 2026.
20. Lee, D.; Lee, K.; Park, D.H.; Moon, G.; Park, I.; Jeong, Y.; Sung, K.Y.; Choi, H.S.; Kim, Y. A Refined Approach to Segmenting and Quantifying Inter-Fracture Spaces in Facial Bone CT Imaging. *Appl. Sci.* **2025**, *15*, 1539. https://doi.org/10.3390/app15031539.
21. Dho, Y.S.; Lee, D.; Ha, T.; Ji, S.Y.; Kim, K.M.; Kang, H.; Kim, M.S.; Kim, J.W.; Cho, W.S.; Kim, Y.H.; et al. Clinical Application of Patient-Specific 3D Printing Brain Tumor Model Production System for Neurosurgery. *Sci. Rep.* **2021**, *11*, 7005. https://doi.org/10.1038/s41598-021-86546-y.
22. Li, P.; Wang, S.; Li, T.; Lu, J.; HuangFu, Y.; Wang, D. A Large-Scale CT and PET/CT Dataset for Lung Cancer Diagnosis (Lung-PET-CT-Dx) [Data set]. The Cancer Imaging Archive, 2020. https://doi.org/10.7937/TCIA.2020.NNC2-0461.
23. Silversmith, W. dijkstra3d: Dijkstra's Shortest Path Algorithm for 3D Volumetric Images. GitHub Repository. Available online: https://github.com/seung-lab/dijkstra3d, 2021. Seung Lab, Princeton University. Accessed: 7 February 2026.