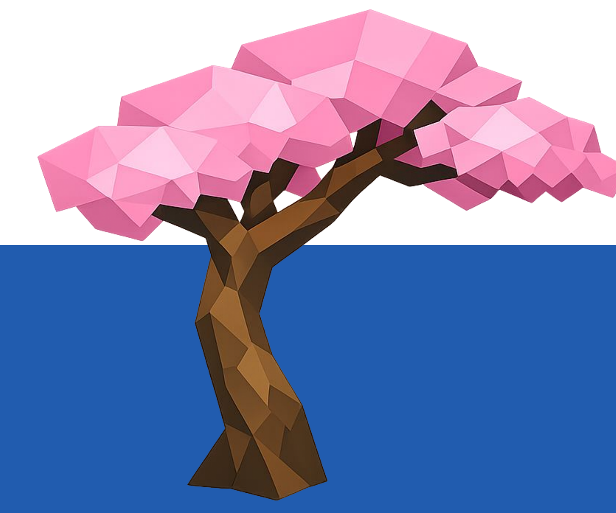

Autoregressive Frontier Expansion: Growing Trees with Graph Machine Learning

Umer Gupta, Saku Peltonen, Martin Ritzert

speltonen@ethz.ch

Abstract

Tree-like branching structures are common in nature, from botanical trees to neurons and respiratory trees. Their **branching shape often reflects function**, making structural modeling central to understanding how these systems work. Acquiring real-world 3D data via imaging can be expensive or infeasible, so **realistic generative models are valuable for simulation and augmentation**. Existing approaches either rely on hand-tuned procedures, or do not jointly generate both tree topology and 3D geometry. We propose a graph neural network that generates trees through an iterative expansion process, **simulating biological growth by expanding step by step**. Experiments on botanical trees show that our method can learn the 3D branching structure across multiple trees.

Autoregressive Frontier Expansion

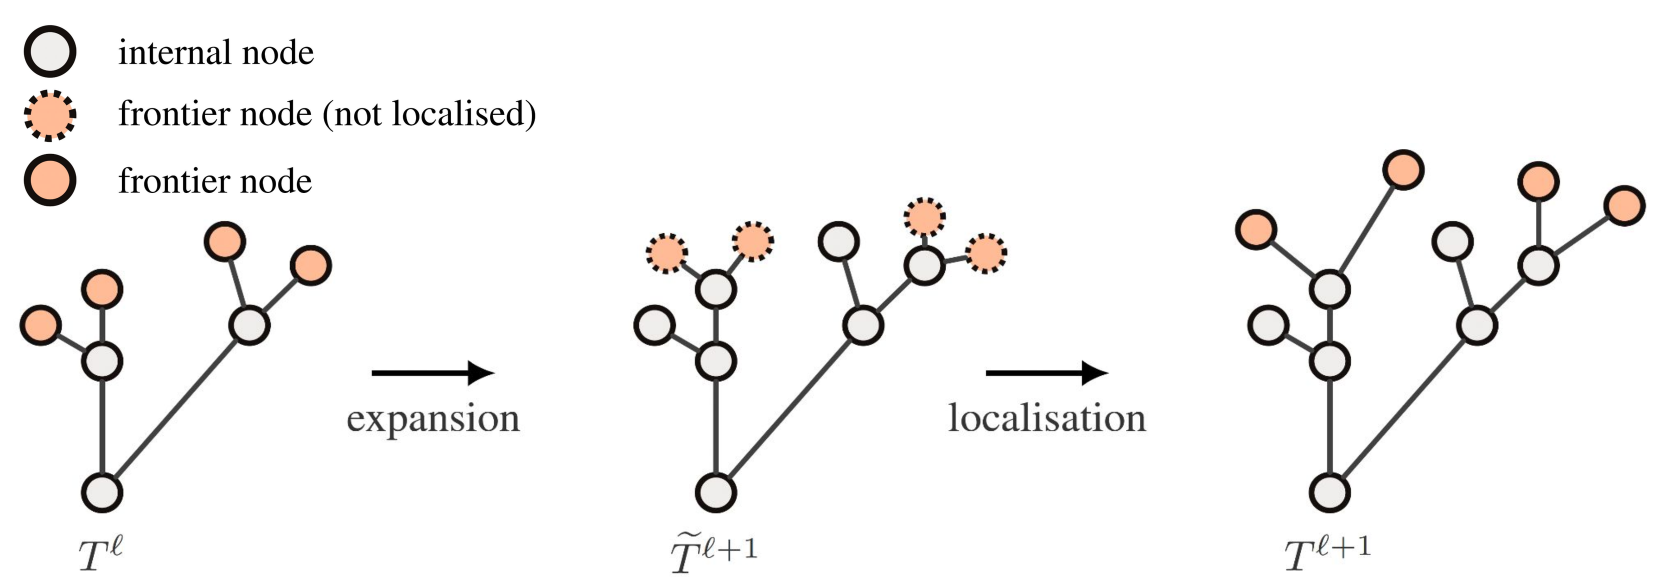

Figure 1: Overview of a generation step. Starting from a partial tree T^{ℓ}, we (i) **expand** the current frontier (a subset of leaves) by deciding for each frontier node whether it branches or terminates, yielding a tree $\widetilde{T}^{\ell+1}$; then (ii) **localise** by predicting 3D parent-relative offsets for the previous frontier nodes, producing $T^{\ell+1}$. The process repeats until all frontier nodes terminate.

Modeling and Training

Diffusion

We model each generation step as a **denoising problem on the current frontier.** Starting from noisy 3D offsets and expansion states, the model iteratively refines both using the current partial tree as context.

SO(2)-equivariant GNN

Our model is designed to be equivariant to **rotations around the vertical axis**, reflecting that trees have a preferred growth direction.

Topological Conditioning

To guide generation in early stages, we condition on a summary of the target tree's topology. We compute persistence-based descriptors from **filtrations** such as tree distance and radial distance (middle), and then embed the resulting **persistence diagrams** (right) as graph-level conditioning signals.

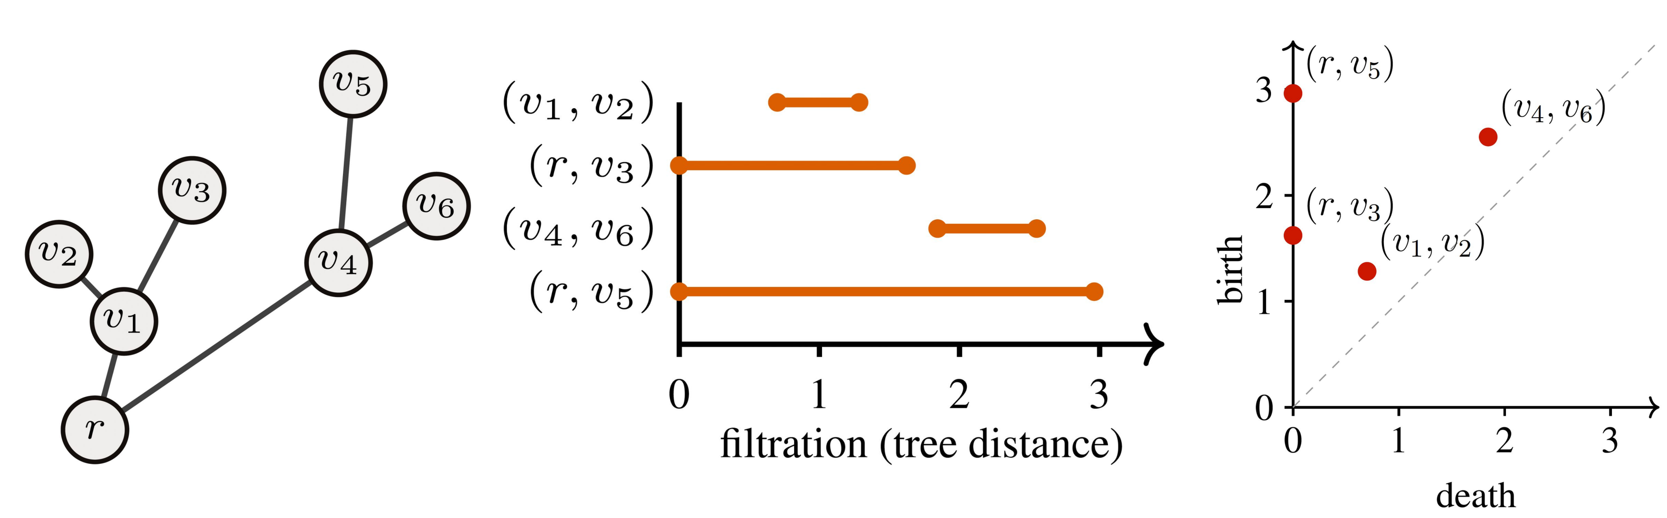

Results

The model reconstructs **botanical trees** with closely matching global statistics and low tree edit distance.

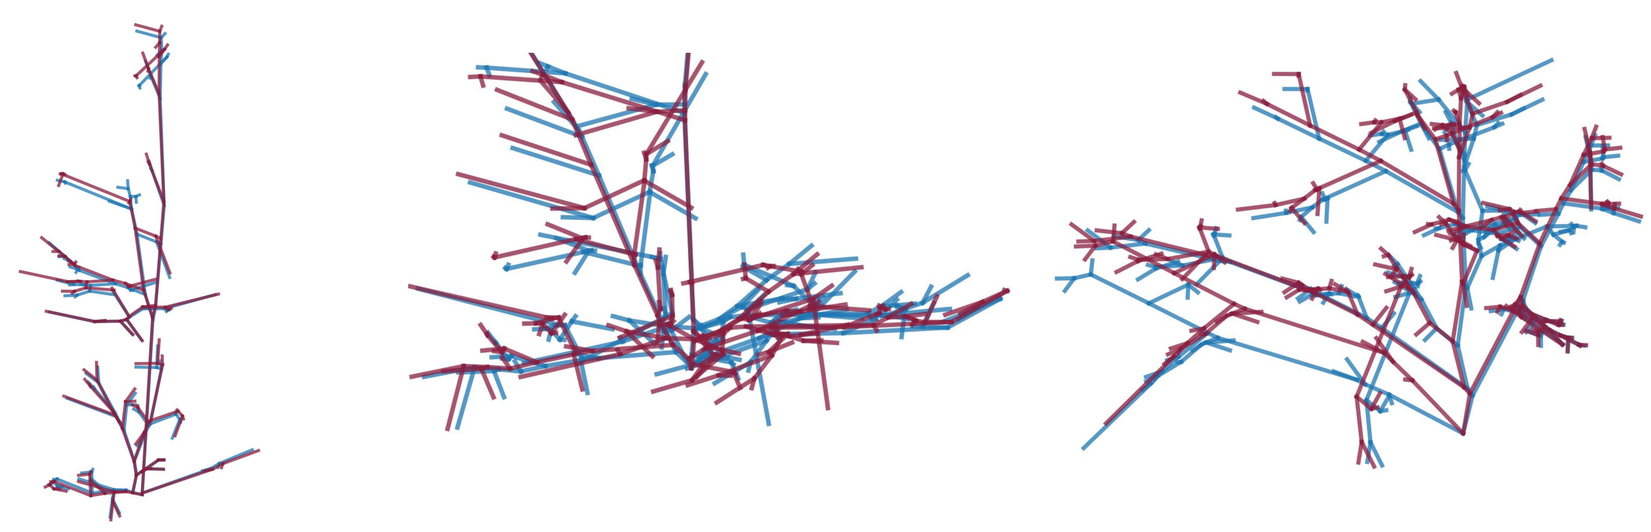

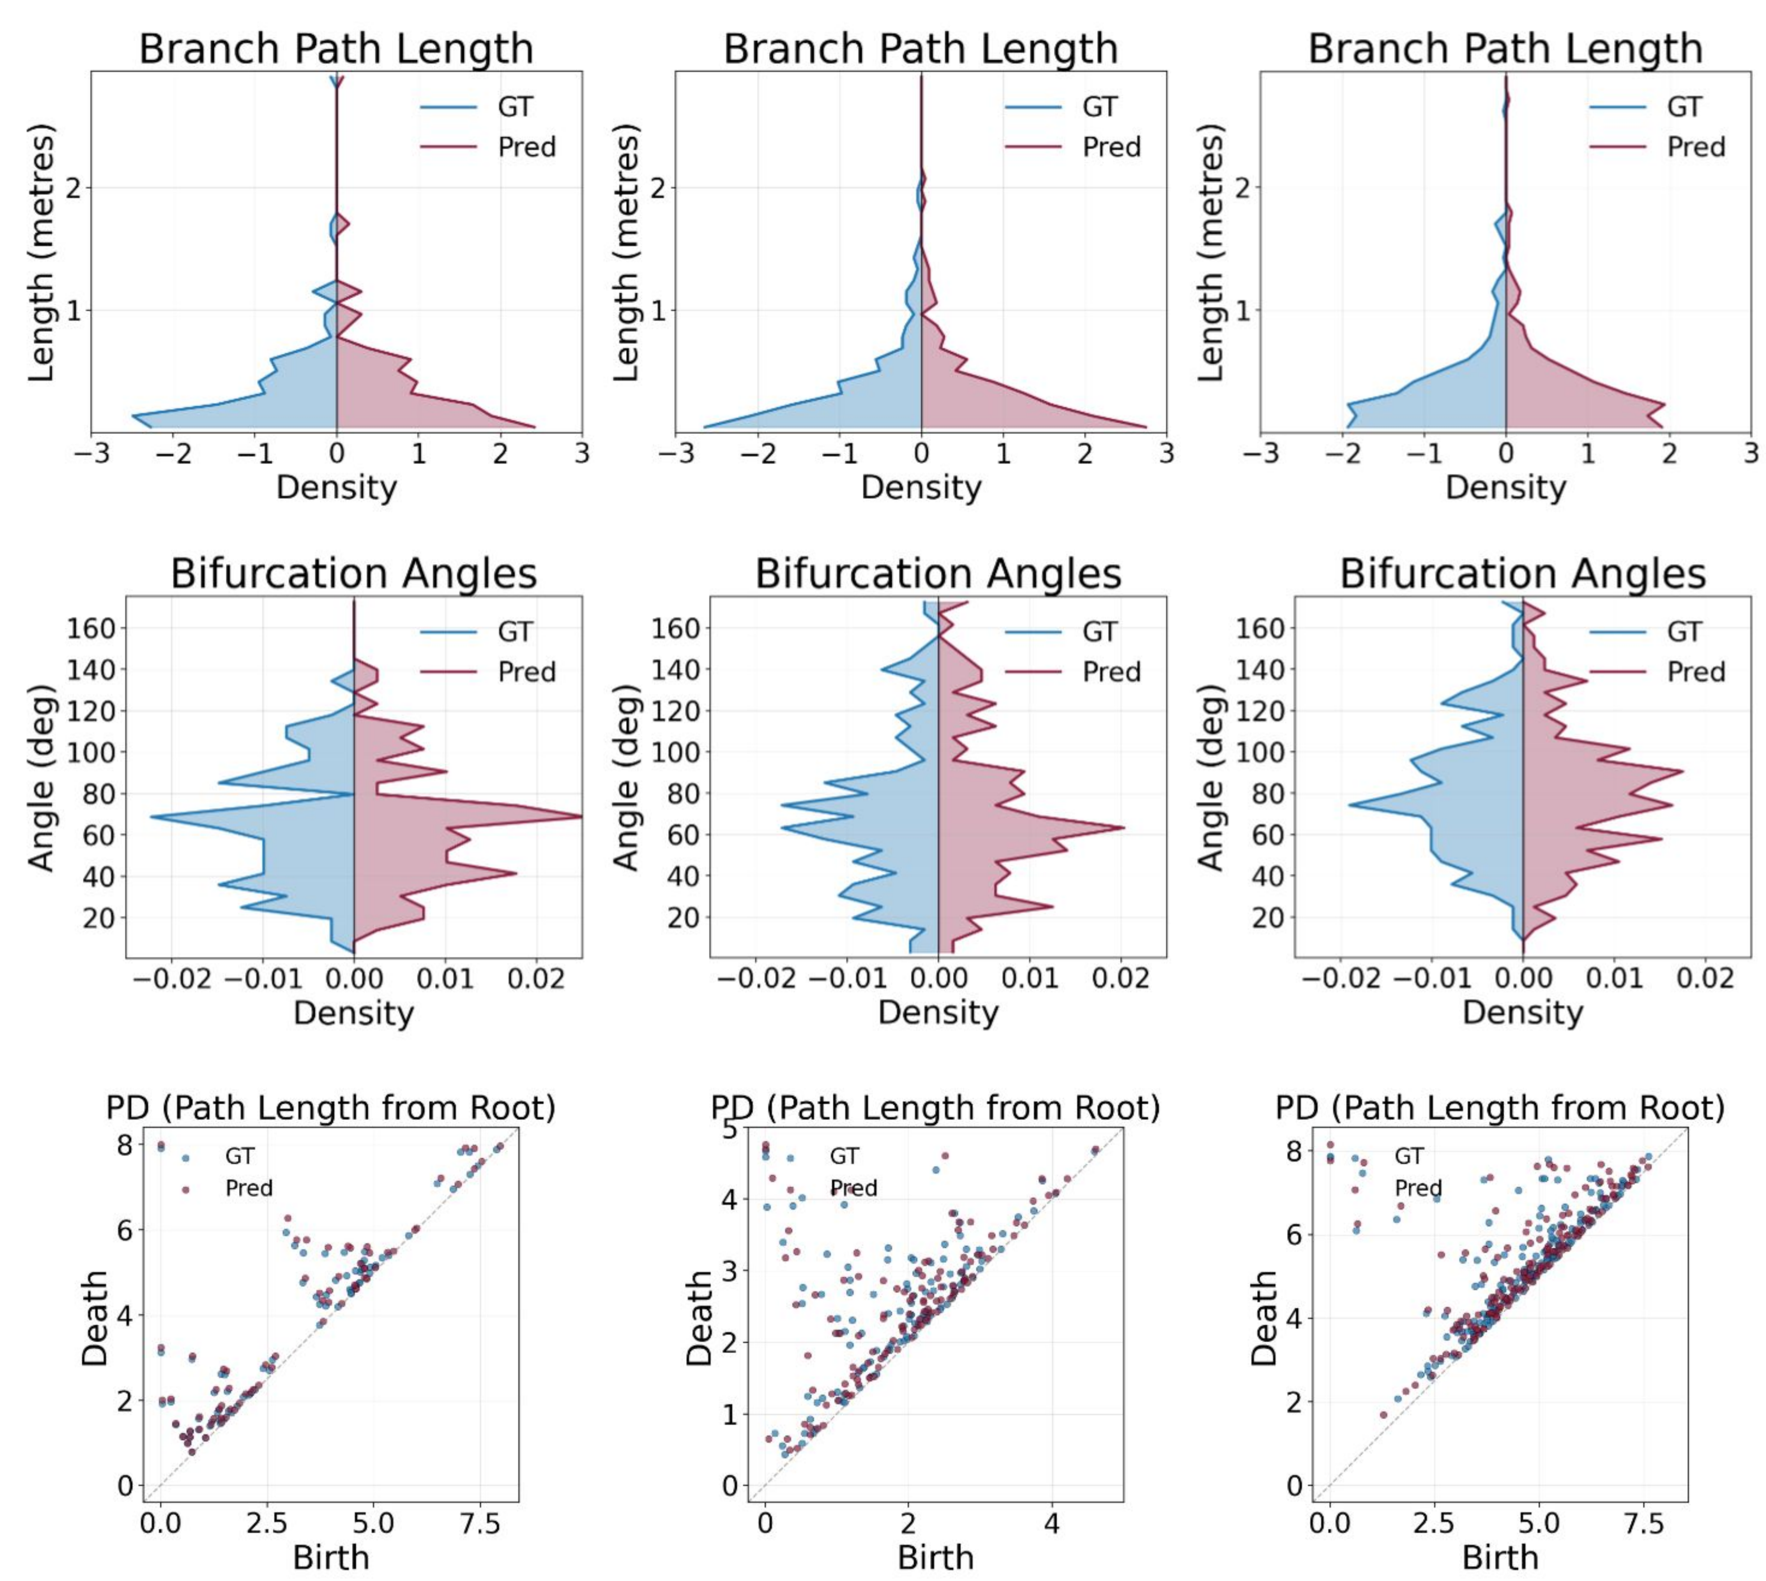

Tree		$\lvert V\rvert$	MBPL	MASB	Height	Radial Span	Box Diagonal	Chamfer Distance	TED
Tree 1	*GT*	149	0.34	64.42	7.65	4.32	9.33	0.08	0.01
	Pred	145	0.36	64.41	7.77	4.46	9.55		
Tree 2	*GT*	235	0.34	69.31	3.55	6.12	8.38	0.23	0.02
	Pred	235	0.32	72.15	3.59	6.37	8.43		
Tree 3	*GT*	327	0.36	81.50	6.37	9.46	12.76	0.39	0.03
	Pred	315	0.36	80.82	6.35	9.04	12.50		

Because generation is iterative, small early deviations can propagate to later branches. As a result, Chamfer distance may penalize reconstructions that are structurally accurate but slightly misaligned.

Takeaways and Next Steps

The method learns both branching topology and 3D geometry of tree structures. Next steps include

- scaling to larger and more diverse datasets
- extending beyond botanical trees to different branching structures

Check out our paper!