

The faculty of Engineering of the Vrije Universiteit Brussel invites you to attend the public defense leading to the degree of

DOCTOR OF ENGINEERING SCIENCES

of **Remco Royen**

The public defense will take place on **Thursday 13th June 2024 at 4:00 pm** in the room **I.0.01** (Building I, VUB Main Campus)

To join the digital defense, please click [here](#)

Meeting ID: 373 158 396 644

Passcode: Wjqwb2

**ADDRESSING LABELLING, COMPLEXITY, LATENCY, AND SCALABILITY
IN DEEP LEARNING-BASED PROCESSING OF POINT CLOUDS**

BOARD OF EXAMINERS

Prof. dr. ir. Bram Vanderborght

Prof. dr. ir. Wendy Meulebroeck

Prof. dr. ir. Bruno Da Silva Gomes

Dr. ir. Nastaran Nourbakhsh Kaashki

Prof. dr. ir. Aleksandra Pizurica

Dr. ir. Ward van der Tempel

PROMOTORS

Prof. dr. ir. Adrian Munteanu

Dr. ir. Leon Denis

Abstract of the PhD research

The field of 3D technology is attracting considerable academic and industrial interest due to its expanding range of potential applications. Noteworthy domains include, but extend beyond, automotive, gaming, extended and augmented reality, drone inspection, robotics, medical imaging, and 3D modeling and design. An essential aspect for these applications is 3D scene understanding. Point clouds, consisting of a collection of points, play a crucial role in capturing the spatial information of physical environments. This results in a lightweight and precise 3D representation, preserving fine details and enabling efficient integration with real-world data.

In recent years, deep learning has gained widespread use, demonstrating its significance across various domains. Its ability to automatically learn intricate patterns from vast datasets has resulted in a transformative impact, driving advancements in technology, and reshaping the landscape of artificial intelligence applications. The ongoing development of increasingly sophisticated neural network architectures continues to push the boundaries of what is achievable across diverse sectors.

As a result, deep learning has become ubiquitous. However, certain limitations hinder its broad applicability. This thesis delves into four crucial challenges associated with deep learning-based point cloud processing: (i) the precise labeling of extensive datasets, (ii) the model complexity requirements, (iii) the latency introduced during inference, and (iv) the concept of scalability. The initial challenge stems from the necessity for extensive datasets with highly accurate annotations. Particularly in the 3D domain, obtaining such high-quality annotations proves challenging and, consequently, expensive. The second challenge arises from the development of more intricate and memory-intensive, facilitated by advancements in high-power-consuming graphics cards. While these methods achieve higher performance levels, they impose constraints on deployment, particularly for embedded devices. Furthermore, the escalating complexity of these networks is accompanied by an increased inference time, impeding real-time applications. Lastly, deep learning-based solutions lack the concept of scalability which have proven vital in traditional methods.

In this thesis, we tackle these challenges and propose diverse solutions within the deep learning paradigm. The thesis commences with the introduction of a rapid 3D LiDAR simulator, designed to mitigate the labeling problem by learning from perfectly annotated synthetic data. We demonstrate its applications in 3D denoising and semantic segmentation. A second contribution can be found within the domain of point cloud instance segmentation. Through the joint learning of prototypes and coefficients, we present an efficient and rapid method that demands relatively low GPU memory. To further improve our method, we introduce an enhanced block merging algorithm. As a third main contribution, we achieve deep learning-based quality scalability by learning embedded latent representations, demonstrating compelling results in applications such as image reconstruction, point cloud compression, and image semantic hashing. The final contribution introduces resolution-scalable 3D semantic segmentation of point clouds. When applied to resolution-scalable 3D sensors, it enables joint point cloud acquisition and processing.

Our proposed methods consistently outperform established benchmarks across diverse datasets, as demonstrated through comprehensive experimentation. The research findings have been disseminated in various reputable journals and conferences, and have led to a patent submission, highlighting their impact in both academic and industrial contexts.