Exploiting Spatial and Temporal Consistency for Fault Detection in Semantic Segmentation NN for Indoor and Outdoor Exploration

In safety-critical applications such as autonomous driving and robotics, the reliability of Deep Neural Networks (DNNs) performing video semantic segmentation is paramount. However, detecting hardware-induced faults in these complex tasks remains a challenge. State-of-the-art methods often rely on metrics like mean Intersection over Union or Pixel Accuracy, which are inapplicable during real-world inference due to the unavailability of ground truth masks, or fail to capture transient faults by focusing exclusively on single-frame analysis. To address these limitations, this work proposes a novel, unsupervised Spatio-Temporal Fault Detection framework. The proposed methodology leverages a dual-layer strategy: a Single Frame Consistency (SFC) Analysis, utilizing geometric features (Area, Position, Symmetry, Shape) to identify structural degradations within single frames, and a Temporal Consistency Analysis, which monitors frame-to-frame dynamics to detect transient glitches that violate physical continuity. To validate this approach, we present an extension of the Faulty Output Dataset (FOD), now covering both outdoor automotive scenarios (Fast-SCNN on Cityscapes) and indoor robotic environments (ESANet on NYU Depth V2). Experimental results for both permanent and transient faults demonstrate that the synergy between spatial and temporal monitoring significantly enhances reliability, achieving a detection performance increase of up to 5.5% compared to single-frame reference, by operating in a fully black-box way.

Lorenzo Fezza, Antonio Porsia, Annachiara Ruospo, Ernesto Sanchez, Matteo Sonza Reorda, Vittorio Turco

CAD & Reliability Group
Dipartimento di Automatica ed Informatica, Politecnico di Torino
Corso Duca Degli Abruzzi, 24 • +39 011 090 102 • www.polito.it