imec

An imec.icon research project | project results

SMART INDUSTRIES

RADIANCE

Real-time and Context-aware Anomaly Detection for Connected Devices

Anomaly detection (AD) identifies outlying data points that potentially indicate some kind of issue or problem. Traditionally, AD is performed by defining thresholds on measured parameters and derived features. Such techniques are easy to understand and perform very efficiently in resource-constrained environments but require deep understanding of the domain where they are deployed. Moreover, they can lead to many false positives if the data characteristics or context of deployment change unless there is frequent human involvement to tune the thresholds, which is difficult and expensive.

Consequently, many businesses are turning to AD based on machine learning (ML) to avoid the need for human intervention. However, such techniques are more resource intensive, especially during the training phase and so training is often carried out only once. But as anomalies manifest differently in different contexts, deploying the solution in other situations and configurations often leads to false positives that flood the user with irrelevant notifications.

The RADIANCE project designed AD algorithms that automatically adapt to changes in context, deployment environment and configurations, data stream parameters, and available resources. This enables accurate, real-time detection of anomalies while reducing false positives. Specifically, RADIANCE aimed to:

1. Support full, end-to-end real-time AD
2. Create algorithms that tune themselves to the available resources
3. Reduce the effort, resources and data required to handle changes in the context or deployment environment and configuration
4. Minimize human involvement to tune the machine learning
5. Enable non-technical users to explore the gained insights through dynamic dashboards

THE OUTCOMES

1. Cross-context learning and visualization

The project developed a cross-context learning methodology that allows AD models trained in one context to easily work correctly in a second context. The contexts in which the model is trained and deployed are accurately represented by knowledge graphs, and a contextual difference algorithm was designed to compare these graphs and pinpoint the three most important differences.

When the AD is adopted in a new context, the difference algorithm finds the most similar context for which the AD model was already trained. In addition, the project designed optimized incremental and transfer learning algorithms that allow AD models to be transferred and quickly optimized to the new context. The context and model from which the AD was transferred plus the 3 key differences can be visualized through a dynamic dashboard, which uses semantic reasoning without any hard-coded configuration. This methodology and visualization reduce the training data required in the new context by 60%.

2. Optimizing AD model resource usage

We developed a systematic approach for selecting the most accurate AD models while also considering their resource usage. Our multi-objective optimization solution extends existing model selection and hyperparameter tuning frameworks to deliver trade-offs between accuracy metrics and resource usage characteristics. These trade-offs take the form of Pareto curves that allow AD models which meet resource constraints to be selected. The solution can speed up model exploration for similar cases by leveraging the hyperparameters of the AD models on the Pareto curve. Moreover, when comparing centralized with federated learning for the given use cases, we were able to speed up training through distributed parallelization and reduce network overhead by up to two orders of magnitude.

3. Three demonstrators

The impact of our solution was showcased in three demonstrators.

Barco's Nexxis video-over-IP platform delivers uncompressed video streaming services to digital operating rooms. The RADIANCE solution allowed Barco to perform real-time AD within its NexxisCare software and monitor the health of a fleet of Nexxis devices and more proactively steer maintenance operations.

The second demonstrator integrated context-aware AD algorithms into Skyline's DataMiner software, enabling better differentiation between genuine network-related anomalies and false positives in various contexts. The algorithms were efficiently deployed to accurately detect anomalies in a very resource-constrained equipment environment.

In the third demonstrator, ML6 used the context-aware AD technique to detect anomalies in vehicle behavior in a city, based on automated number plate recognition (ANPR).

NEXT STEPS

Barco is adding automated, adaptive anomaly detection and visualization to its NexxisCare cloud-based software system for remote Nexxis operating room management. This will significantly reduce false positives, increase OR uptime and enable efficient, cost-effective troubleshooting and maintenance. The RADIANCE solution could also be beneficial in other, similar business cases.

Skyline plans to incrementally augment its DataMiner software through multiple release cycles, based on the results of the RADIANCE project. This will enable it to redefine network fault detection with increased accuracy, reduced set-up effort, easier troubleshooting and a broader scope that includes largescale, resource-intensive environments.

Further, the RADIANCE project has enabled ML6 to develop a new anomaly detection technique, known as bagged frequency sampling, that can handle very large datasets while maintaining state-of-the-art accuracy. This algorithm is currently being implemented at customers and will allow ML6 to become a reference in high-performance anomaly detection.

The RADIANCE project was co-funded by imec, with project support from Agentschap Innoveren & Ondernemen.

AGENTSCHAP INNOVEREN & ONDERNEMEN

Vlaanderen is ondernemen

FACTS

NAME	**RADIANCE**
OBJECTIVE	Real-time, context-aware anomaly detection for connected devices
TECHNOLOGIES USED	knowledge graphs, federated learning, transfer learning, incremental learning, anomaly detection, microservices, hyperparameter tuning, communication networks, automated number plate recognition (ANPR)
TYPE	imec.icon project
DURATION	01/09/2018 – 30/11/2020
PROJECT LEAD	Geert Gheysen, Barco
RESEARCH LEAD	Femke Ongenae, imec - IDLab – UGent
BUDGET	2,402,306 euro
PROJECT PARTNERS	Skyline communications, Barco, ML6 - Skyhaus
RESEARCH GROUPS	imec - IDLab IBCN – UGent; imec - IDLab Data Science Lab – UGent; imec - DistriNet - KULeuven

RADIANCE project partners:

BARCO ML6 skyline communications

KU LEUVEN GHENT UNIVERSITY

WHAT IS AN

IMEC.ICON PROJECT?

The imec.icon research program equals demand-driven, cooperative research. The driving force behind imec.icon projects are multidisciplinary teams of imec researchers, industry partners and/or social-profit organizations. Together, they lay the foundation of digital solutions which find their way into the product portfolios of the participating partners.

CONTACT US

WWW.CONTACTIMEC.COM