GENERAL MOTORS

Deep Learning For Machine Vision

Student Team:
Mabel Chan – EGL (BSE & MSE in Computer Science Engineering)
Konstantinos Chiotinis – Master of Business Administration
Soo Yeon (Sean) Lee – EGL (BSE & MSE in Computer Science Engineering)

Project Sponsors:
Amar Amad – Engineering Manager
William Keller – Sr. Manufacturing Project Engineer

Faculty Advisors:
Jeff Alden – College of Engineering
Sanjeev Kumar – Ross School of Business

General Motors Company, a $79B company, is one of the world's largest automotive manufacturers, producing close to 7 million vehicles a year. The company uses more than 2,000 machine vision cameras across all facilities to inspect various aspects of production. So far, the accuracy of several vision cameras has not been satisfactory to General Motors, since inaccuracy results in unnecessary downtime. General Motors wants to decrease downtime and increase throughput. A potential solution to the problem of camera inaccuracy is the application of deep learning (a form of artificial intelligence learning based on artificial neural networks in which multiple layers of processing are used to extract progressively higher level features from data) on machine vision cameras.

Asked to analyze the performance of deep learning versus traditional methods for machine vision applications within GM, the team implemented deep learning on 13 machine vision cameras that were either performing poorly with traditional vision methods or were performing especially complex inspections. Using deep learning, the team was able to improve camera performance, achieving an accuracy of 99.7% or higher, a number set as the target by GM, an improvement of more than 15 percentage points for some cameras. With deep learning, every camera performed as well as or better than traditional vision methods, indicating an opportunity for large-scale implementation.

Three process improvements were also made, with a fourth proposed possibility: (1) a decision-making process to determine whether a camera is a suitable deep learning candidate; (2) a protocol for training deep learning models, which decreases implementation time by 80%; and (3) a method to reuse deep learning models on different cameras. Fourth, the team identified several bottlenecks, which, if addressed, can decrease deep learning implementation time by up to 50% (more than 3 hours per camera) and the needed space for storing historical machine vision data by 66%, saving GM several terabytes of storage space.

The performance improvements made by the team resulted in over 30 hours of decreased projected downtime per month. The potential economic benefits for implementing deep learning across the assembly line are estimated to be roughly $9 million for the first year and $40 million within 3 years. Also, the increase in machine vision accuracy will positively affect quality control.