Autonomous Vehicles Use New AI Algorithm to Learn from Changes in the Environment

Article

What does data have to do with vehicles?

Every moment of our lives we are internally collecting and processing data points of the world around us, but this process has exceeded being unique to only the biological existence and grown to include the silicon version simultaneously. Like the biological brain in its developmental years, AI is learning how to process enormous amounts of visual data. One of the most fascinating and popular areas of this is within autonomous vehicles.

According to Gartner (gartner.com/en/information-technology/glossary/autonomous-vehicles):

> "An autonomous vehicle is one that can drive itself from a starting point to a predetermined destination in "autopilot" mode using various in-vehicle technologies and sensors, including adaptive cruise control, active steering (steer by wire), anti-lock braking systems (brake by wire), GPS navigation technology, lasers and radar."

The technological feat is powered by the consumption of data at complex and varying degrees, while evolving with changes in human nature. Remember watching shows and movies like Batman, Knight Rider, and Total Recall, sitting with fasciation of the idea that a vehicle could drive you around? That awe-inspiring notion of what the future could hold, well that future is truly here now.

Benchmarks = Standards = Measurements of Growth

Working together with Intel® and the Barcelona Supercomputing Centre (BSC), Lenovo researched the various ways in which autonomous driving data, known as SODA 10M, are consumed, and processed. SODA 10M is “a large-scale object detection benchmark for standardizing the evaluation of different self-supervised and semi-supervised approaches by learning from raw data,” as shown in the image below. (from paperswithcode.com)

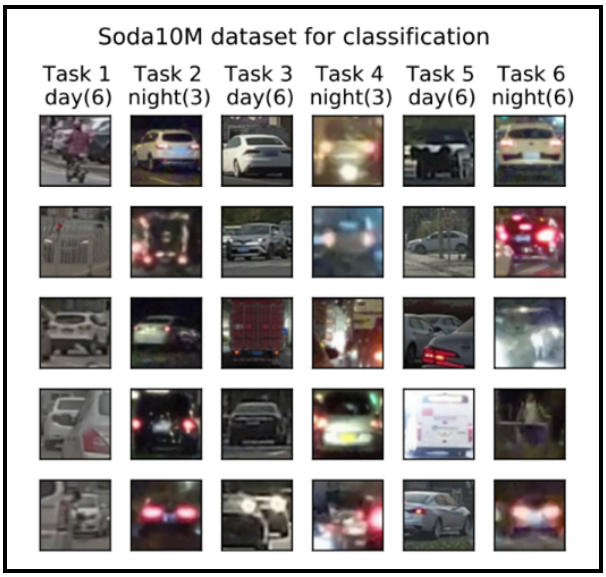

Figure 1. SODA 10M dataset for classification

Machine Learning Continuously from New Data

These processes are powered by Machine Learning (ML), which assumes that all data consumed is the same within the process, meaning that data is assumed to be independent and identically distributed (iid), but that is not really the case. One example of this is when data is assumed to be iid but is correlated in different contexts. This is called a domain shift, which is extremely common in ML. To provide a bit more context around how to understand a domain shift, consider the drive from New York City into rural Pennsylvania. Our brains notice and interpret the context as the scenario changes around us. How we drive and interact with that new environment is seamless and fluid, allowing us to continue onward without needing to reset or reprogram ourselves. Even though MLs have the ability to shift, which is all too common in autonomous driving data, how can it do so to continually advance its capabilities?

Continual Learning (CL) is the study to keep an ML system learning post-deployment, e.g., within an autonomous vehicle. Traditionally, this is done by retaining all learned data and then retraining the system frequently. However, due to various guard rails, this can pose problems around data privacy, storage, or compute restrictions. Knowing these limitations provide the problem known as “catastrophic forgetting” in which a neural network is unable to make a correlation between what it previously learned and the new data it was provided. Again, we can bring this back to our human experience.

Consider an activity that you enjoy, whether it be sports, reading, language, art, music, etc. In the moment, we focus on learning the new skills within that task. Time progresses and we go on to learn other skills, but one day come back to that original skill and simply forget the old skills or techniques after learning new ones. This is catastrophic failure—forgetting old rules after learning new rules—is something that machine learning is susceptible to, just like a human. Through the creation of algorithms, these challenges have been addressed, but often assume rigid task boundaries and known new tasks.

Putting the right team together is where it all begins.

Gusseppe Bravo-Rocca, MS, Peini Liu, MS, Jordi Guitart, PhD, Ajay Dholakia PhD, and David Ellison PhD wrote a paper titled, “Efficient Domain-Incremental Learning through Task-ID Inference using Transformer Nearest-Centroid Embeddings,” hosted in Lenovo’s Center of Discovery and Excellence. The research was run on the Lenovo ThinkSystem SR650 V2 server, while the testbed used in the experiments is as follows:

- Server: Lenovo ThinkSystem SR650 V2
- Processors: 2x Intel 3rd Gen Xeon Platinum 8360Y CPU @ 2.40GHz, 256 GB RAM
- Operating System: Ubuntu 22.04 (64 bit)
- Software:
    - Docker image intel/oneapi-aikit:develubuntu22.04 1 (Intel AI Analytics Toolkit)
    - avalanche-lib 0.3.1 2 (CL library)
    - torch 1.12.0 and torchvision 0.13.0 3 (DL library)
    - intel-extension-for-pytorch 1.12.100+cpu 4 (Intel acceleration for Pytorch)
    - scikit-learn 1.2.2 5 (ML library)
    - scikit-learn-intelex 2023.0.1 6 (Intel acceleration for Sklearn)

The research will show the steps in which the participants conducted group embeddings, trained classifiers, and created an algorithm that can decide if the program should or should not learn a new class.

A task detector was developed that can operate unsupervised and did not have to be informed that any data is changing or carrying a new label. Explaining such a task detector, we can go back to our example of driving from NYC to a rural area, which are two different classes. Much like our brains, this task detector allows the ML to operate without catastrophic failure and sense change while adapting to the new without forgetting the old. Most importantly with this development, the ML within the vehicle did not need to update, but rather was able to learn and run within the context switch it was experiencing first-hand.

Autonomous vehicles are progressing to advance and become a regular occurrence in day-to-day life for society. Lenovo, alongside our partners and fellow scientists will continue to help advance this technology and continue to drive smarter technology for all. This being the first stage of the project we are very excited to share our findings and look forward to bringing you along for the journey.

For more information

For more information, see the following resources:

- Explore Lenovo AI solutions:
  https://www.lenovo.com/us/en/servers-storage/solutions/analytics-ai/
- Engage the Lenovo AI Center of Excellence:
  https://lenovoaicodelab.atlassian.net/servicedesk/customer/portal/3

Related product families

Product families related to this document are the following:

- Artificial Intelligence

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION ”AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2026. All rights reserved.

This document, LP1716, was created or updated on April 14, 2023.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  https://lenovopress.lenovo.com/LP1716
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at https://lenovopress.lenovo.com/LP1716.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both:
Lenovo®
ThinkSystem®

The following terms are trademarks of other companies:

Intel®, the Intel logo and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Other company, product, or service names may be trademarks or service marks of others.