



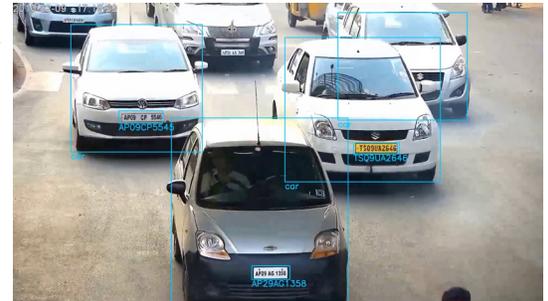
## SMARTCOW LAUNCHES SENTINAL ALPR



### DEEP LEARNING: FROM R&D TO REALITY

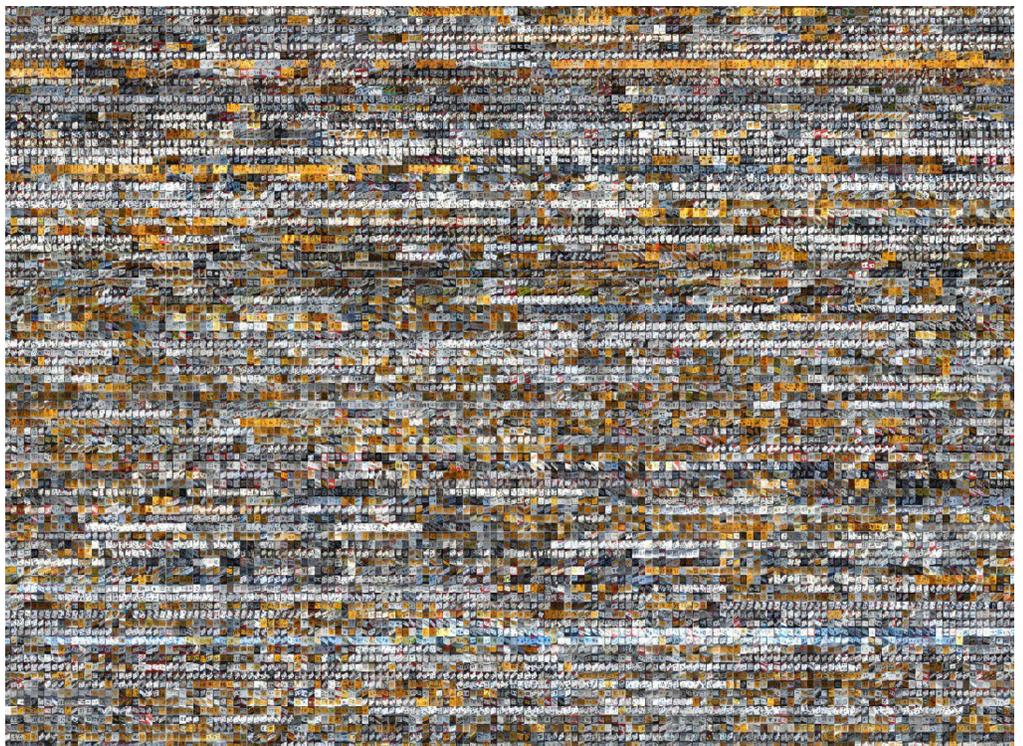
The following describes how one AI engineering company leveraged the NVIDIA® Jetson™ ecosystem to rapidly deploy accurate, cost-effective, commercial-ready, and industrial-grade automated license plate recognition (ALPR) solutions on a potentially massive scale.

SmartCow is a multi-national AI engineering firm focused on deep learning-based computer vision technologies. The company develops proprietary inferencing algorithms for ALPR, which are supported by a large team whose sole responsibility is to label large data sets of images that are fed into deep learning models.



From a technical perspective, the most important requirement of an ALPR system is that it processes video streams and runs neural network algorithms locally. By performing character recognition on or near the device, eliminates the latency of sending video streams back to the cloud for analysis, minimize the network transmission costs associated with wirelessly streaming video, and reduce any security concerns related to sending information over the network.

The challenge this presents is that massive computational and memory resources are required to run neural networks. This is why they are typically run in a data center, not at the edge. For some perspective, Figure 1 below shows a snapshot of the more than 350,000-image dataset leveraged for the ALPR inferencing algorithm.



(Figure 1) SmartCow's ALPR inferencing algorithm is based on a dataset of more than 3 million images. Pictured here is a 350,000-image snapshot of that 3-million image data set.

### Challenge

The inability to compress neural network models into a footprint suitable for deployment on embedded processors at the edge.

The lack of suitable embedded compute platforms with the performance required for tasks like image recognition and object classification, as well as the performance per watt to satisfy resource-constrained edge environments.

### SmartCow Solution

SmartCow's Sentinel is an automatic license plate recognition (ALPR) system based on the NVIDIA Jetson Nano processor module. It runs four deep learning algorithms based on a 3 million image data set. The result is **near-human** accuracy in license plate recognition tasks.

### Connect Tech Solution

Connect Tech Inc. developed a custom, Jetson Nano carrier board solution for the SmartCow Sentinel platform with a unique mix of I/O, extra NVMe storage, custom BSP and maintained a compact 103 mm by 72 mm footprint.

Connect Tech Inc. has 35 years of experience in embedded hardware, software, and systems design, and is presently NVIDIA's largest embedded ecosystem partner.

Ready to deploy your Deep Learning Application? Contact Connect Tech at [sales@connecttech.com](mailto:sales@connecttech.com)

A variety of computer vision-based chipsets are available in the market but what ultimately sets Jetson products apart is a vibrant software ecosystem. For SmartCow AI, this meant access to TensorRT, a programmable software platform that accelerates inferencing workloads running on NVIDIA GPUs.

TensorRT is a C++ library based on the CUDA parallel programming model that optimizes trained neural networks for deployment in production-ready embedded systems. It achieves this by compressing the neural net into a condensed runtime engine and adjusting the precision of floating-point and integer operations for the target GPU platform. This correlates directly to improved latency, power efficiency, and memory consumption in deep learning-enabled systems, all of which are of course essential in ALPR-type applications.

TensorRT is part of the NVIDIA JetPack SDK, a software development suite that includes a Ubuntu Linux OS image, Linux kernel, bootloaders, libraries, APIs, an IDE, and other tools. The SDK is continuously updated by NVIDIA to help accelerate the AI development lifecycle.

Take ALPR in India, for example. The technology is being deployed to serve a number of purposes, from checking the registration status of vehicles and electronic toll collection to enforcing traffic laws and general security surveillance. The issue on the subcontinent is that each state uses a different style of license plate, and these styles even vary between different classes of vehicles from the same state. As shown in Figure 2 below, the variations include different sizes, layouts, colors, script types, and placement on a vehicle.

Today, low-cost ALPR solutions are available in the Indian market for \$1,500 plus licensing fees. However, these are packaged with one 2 MP camera with a 15 fps frame rate, which limits their ability to capture high-resolution shots of the various license plates positioned at various locations on vehicles.

But there are also more technical challenges presented by the wide variance in Indian license plates. For one, the sheer number of ways that plates could be represented in an image means that an ALPR system must be trained against a much larger data set compared to more controlled use cases. And, because a fixed ALPR system will inevitably perform inferencing on vehicles that are parked or idling at traffic signals, the devices will re-register the same license plate numbers over and over. Both of these facts mean that storage and processing requirements increase.



(Figure 2) Indian license plates vary by size, color, layout, and script.

At the same time, of course, the design requirements of a low-cost, low-power, small-form-factor ALPR system still apply. The ability to perform inferencing locally is also a must.

## SCALING DOWN TO SCALE OUT



(Figure 3) The NVIDIA® Jetson Nano™ is a small form factor embedded processor module designed for computer vision and deep learning applications.

The Jetson Nano module is a compact 70 mm x 45 mm embedded processor module based on a Tegra processor you'd expect to find in the data center. The system on chip at the heart of the board contains a Maxwell architecture GPU with 128 CUDA cores alongside a quad-core Arm Cortex-A57. This equates to 472 GFLOPS of total performance at the expense of only 5 to 10W power consumption.

The Nano SoC also integrates 720p, 1080p, and 4K HEVC video encoders and decoders that deliver 250 MP/s and 500 MP/s of performance, respectively. These are, understandably, key components in an ALPR application that has to apply neural networking algorithms to streaming video in real time.

In addition, the Nano supports 4 GB of 1600 MHz LPDDR4 memory for the fast, frequent memory accesses required by deep learning applications, as well as 16 GB of additional eMMC storage. Interfaces such as GbE, HDMI 2.0, eDP, and USB 3.0 are brought out through a companion carrier board.

The low power, low cost Jetson Nano provided a solid foundation for SmartCow's deep learning-based ALPR solution, Sentinel.



(Figure 4) SmartCow's Sentinel is an automatic license plate recognition (ALPR) system based on the NVIDIA® Jetson Nano™ processor module.

Sentinel (see Figure 4) is a miniature ALPR system that runs four deep learning algorithms based on a 3 million image data set. The result is near-human accuracy in license plate recognition tasks.

Sentinel currently supports four 30 fps RTSP cameras and can integrate with existing solutions. However, SmartCow is implementing direct memory access (DMA) features that would bypass the CPU when video streams are being decoded. By freeing up these extra resources, one Sentinel box will be able to support six to eight cameras simultaneously in the very near future. The device retails for around \$1,000 with no additional licensing fees, making them less than 1/8th the overall price of today's lowest-cost solutions with the potential to be less than 1/16th the cost once eight-camera support is implemented.

To offset the problem of re-registering images of the same license plate, SmartCow developed a feature called similarity search. This registers the first time a license plate is detected and discards all duplicate images until the vehicle leaves the frame, saving valuable memory space.

In addition to enough storage capacity to save more than 2.5 million images directly on the device, Sentinel also supports low-cost 4G connectivity modules that allow metadata to be streamed back to transit authorities.

All of this at only 10W of power.

Sentinel boxes are being deployed on Marina Beach in Chennai, India, where the integrated CPU can act on inferences to open and close gates, garage doors, or cycle traffic signals.

## MOVING AI INTO THE REAL WORLD

Although the core functionality of the Jetson Nano power platforms like SmartCow's Sentinel, the off-the-shelf NVIDIA modules cannot be deployed in production-grade systems. After all, the module was originally designed as a development kit, and this application required custom design.

In real-world deployment settings, Jetson platforms need a carrier board. The carrier board routes input and output signals through standard connectors, and also provides access to additional functionality. For example, the SmartCow deep learning platforms needed GPIO outputs for control tasks like driving gates and triggering relays, which were delivered over a terminal block in this design.

In the case of this ALPR solution, additional memory was a prerequisite. In order to support the 3 million region of interest (ROI) deep learning algorithms and store more than 2.5 million images, much more was needed than the 4 GB RAM and 16 GB eMMC natively available on the Jetson Nano. This meant adding extra non-volatile memory to the design.

And, of course, all of the systems' hardware components had to be spec'd out for operation in potentially harsh outdoor environments. Once that was complete, both platforms had to be designed into compact, rugged enclosures that provided ample protection against the elements as well as sufficient heat dissipation.

But perhaps the most critical and most often overlooked element of taking these systems to market was the software stack. As part of the JetPack SDK, Jetson development kit modules come pre-loaded with a stock Linux for Tegra operating system (OS) image from NVIDIA. This provides a great starting point for software prototyping, but these OS images are not

suitable for use in commercial environments.

However, in order to maintain software compatibility with JetPack tools and features that are constantly being updated by NVIDIA, a commercial-grade board support package (BSP) is required. This means creating a modified BSP with a lifecycle that aligns with JetPack releases.

As an AI, deep learning, and video analytics firm, these engineering disciplines fell outside of SmartCow's core competency. After evaluating several potential design partners, SmartCow selected Connect Tech Inc. (CTI) to help transition their ALPR technologies into the real world.

## FROM PROTOTYPE TO PRODUCTION GRADE

CTI has 35 years of experience in embedded hardware, software, and systems design, and is presently NVIDIA's largest embedded hardware ecosystem partner. To date the company has completed more than 100 custom Jetson-based product designs that are currently deployed in the field, spanning the AGX Xavier, TX1, TX2, and Nano product offerings.

In the case of Sentinel, CTI designed the custom, Nano-compatible carrier board that integrated the aforementioned GPIO pins for driving relays and other control functions and additional NVMe storage in support of the platform's significant memory demands. They also removed JTAG headers for additional system security.



*(Figure 5) Connect Tech Inc. developed a custom, Jetson Nano™ - compatible carrier board solution for the SmartCow Sentinel platform that integrated GPIO pins and extra NVMe storage while maintaining a compact 103 mm by 72 mm footprint.*

Another substantial feature added to the Sentinel platform by CTI was support for 4G cellular modules over the M.2 interface. This allows customers to simply drop in a 4G dongle to gain access to wireless backhaul networks, allowing metadata to be transmitted back to operations centers and software updates to be deployed over-the-air to the device without the need for any cabling. This is a particularly economical option in India, where 30 GB of LTE service can be had for a modest \$3.

Perhaps the greatest benefit of the SmartCow/CTI partnership was the ability to complete a turnkey design quickly. CTI was able to modify the BSP from the stock Tegra for Linux OS to a custom, small-footprint Ubuntu Linux distribution from SmartCow that required much less memory.

The full system – including the hardware design, BSPs, and supporting software, and firmware – was actually complete before production-ready Jetson Nano modules were even available. CTI also provided “Flashing-as-a-Service”, helping SmartCow flash Sentinel platforms en masse to further accelerate time to market.