# Cogniac

*Observe and Understand the Visible World*

## Purpose

Cogniac offers a superior enterprise platform for Artificial Intelligence (AI)-grade visual observation. In this document, we have provided information about our company and technical capabilities. The Advantages of our solution which we will demonstrate in this paper are:

- **Cogniac offers a flexible & Adaptable System**
- **We embrace continuous improvement, with human in the loop to minimize errors**
- **A superior technical approach**

## About Us

Cogniac is an enterprise computer vision company based in the San Jose, CA with big plans to improve the world by making it easy to automate object recognition. Our innovative platform enables anyone to easily extract actionable information from any source of visual data. Common visual observation tasks such as classifying, detecting, counting and measuring items and conditions can be automated with a high degree of reliability. The Cogniac system uniquely combines the latest artificial intelligence research with human-computer interaction tools and large-scale data management in order to rapidly and automatically learn to identify what matters most to your business so you can take informed, timely, and even automated action. Cogniac is led by William Kish, a serial entrepreneur and engineer who holds over 44 patents and who was previously the co-founder and CTO of Ruckus Wireless. Ruckus, which IPO'd on the NYSE in 2012, grew to be the 3rd largest enterprise WiFi equipment vendor in the world and maintained an international presence with 7 R&D sites worldwide. Early investors include Gradient Ventures, part of Google's AI investment fund.

## Technology Background

Research in a particular branch of AI computer vision, deep convolutional neural

networks (DCNs), has rapidly accelerated since 2012 when a DCN first won the ILSVRC[1] "Imagenet Challenge". Since then DCNs have surpassed human-level accuracy on large scale image classification tasks. These achievements stem from the confluence of decades of research in neural networks and machine learning, the availability of larger training data sets, and the inexorable march of raw compute system performance, in particular GPU performance which is particularly well suited to the training of deep neural networks on large data sets.

For all the promise and excitement, there is a formidable barrier to applying this very powerful but also very raw technology to solve real-world problems. Successful research workflows do not translate well to commercial settings due to a number of problems including tedious data curation and preprocessing requirements, low-level tensor frameworks requiring rare expertise and extensive configuration, and the need to continuously refine and deploy updated prediction models as both requirements and the world-at-large invariably evolve.

## The Cogniac System

Cogniac has developed an elegant end-to-end architecture to enable the commercial deployment of deep convolutional neural networks in enterprise video and image analysis applications. We address the key challenges of data acquisition and management, model training, hyperparameter optimization, model deployment, and ongoing performance assessment. The system is designed with a 'human in the loop' to assist with difficult images, ensuring a high degree of accuracy and also enabling the system to improve over time as it is designed to actively query the relevant human experts within your organization and incorporate their quick visual feedback via a streamlined application interface.

The following diagram provides a high-level overview of the Cogniac system for a single application. Although not shown here the system is multi-tenant so each customer's data is securely partitioned[2]. Each customer can have any number of "applications", which capture a related set of subjects that the customer is interested in detecting, classifying, counting or measuring.

[1] image-net.org/challenges/LSVRC/

[2] Cogniac also offers hybrid-cloud and pure on-premises versions of the system

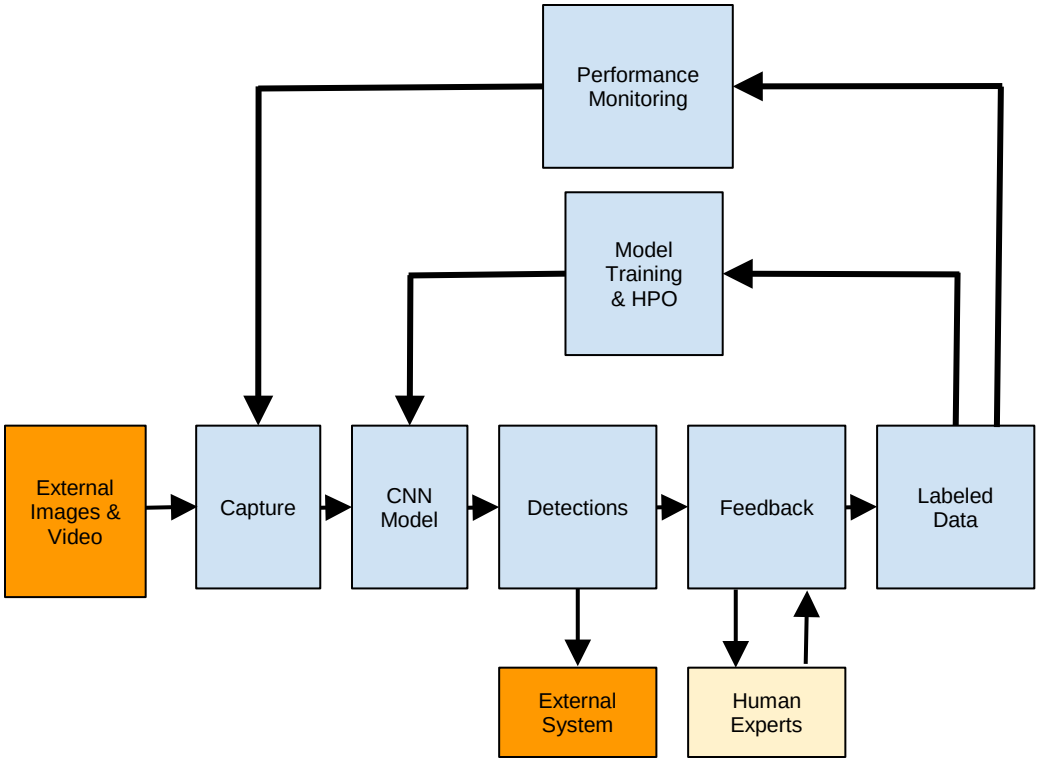


Figure 1 - Cogniac Overview -- The system automatically generates best-in-class deep convolutional neural network (DCN)

External image data or video is input into the system via the capture interface which performs basic indexing of the input data and assigns unique image identifiers which will persist throughout the system. The system supports most common image and video formats such as JPEG, TIFF, PNG, MJPEG, and H.264 and can process images extremely large images (over 10,000 x 10,000 pixels) as well as large videos. The image data is passed to an image processing pipeline which runs the per-tenant, per-application specific DCN classification and detection models. The output of the models produces a set of zero or more detection events, which are a software indication that one of the configurable object classes has been detected in the input image or video stream. These detection events are sent to the customer's external systems subject to configurable filters that control which detection events should be surfaced based on the type and confidence level of the detections.

Detection events are forwarded to a feedback management subsystem which determines the detection events that should be surfaced to the relevant human domain experts for confirmation of the application detections. The rate at which feedback is requested is dependent on the quality of the model. More feedback is required in the early model training

phases. Over time the rate of requested feedback will naturally decrease as the model accuracy increases. The feedback subsystem also prioritizes feedback requests based on the amount of new information that the prospective feedback item will provide to the system. This serves to minimize disruption to the human experts by avoiding similar or redundant feedback queries. While the feedback can be supplied from any computing system via our public API, Cogniac provides a reference iOS application for iPhone and iPad to enable efficient workflow integration with minimum custom development.

Application prediction, model training, and optimization is an ongoing process, and typically models will be retrained after a certain quantity of new labeled data is created. The key to good application model performance is the selection of the best hyperparameters. Hyperparameters are the dozens or hundreds of variables that define a deep learning model, the associated configuration of the model, and the configuration of all possible image preprocessing options. In traditional workflows, hyperparameter selection has been somewhat of an art form practiced by experienced data scientists. The Cogniac system fully automates this tedious process through an automated hyperparameter optimization (HPO) subsystem that explores values of all relevant model dimensions in order to find the best model configurations and continuously adapts model configurations as dictated by growing dataset sizes.

The final Cogniac component is a performance monitoring subsystem that provides real-time and historic assessment of model performance and accuracy. This is especially important in an online system since the test and validation sets are never ‘done’. In addition to providing users with system performance analytics, the results of this assessment are used to adapt certain internal configuration items such as feedback requirements so that as the accuracy of the system improves the system will automatically request less feedback from the users.

## Demonstration Applications

The Cogniac system is designed to facilitate the easy creation of deep convolutional neural network models via simple user interfaces. There are dedicated operator screens for the common tasks as follows:

- Application Overview

The application overview provides and administrative dashboard for configuring an application and monitoring its performance.

- Application Feedback

The application feedback view provides the primary interface for confirming or correcting application model predictions and generation of new training data

- Subject Ground Truth

The subject ground truth view provides an easy mechanism to review or audit the 'ground truth' data used for application model training.

- Application Detections

The application detection screen provides a real-time view into each individual application output prediction and associated confidence.

- Image Metadata

This view includes all historic metadata for each individual image or video including predictions and any user corrections.

To illustrate these views and demonstrate the applicability of the Cogniac system applications the following screenshot shows the web browser based view of the 'freighthopper' (illegal rider) application overview screen in the Cogniac System. From this view, the application owner can review current model and historical application performance statistics.

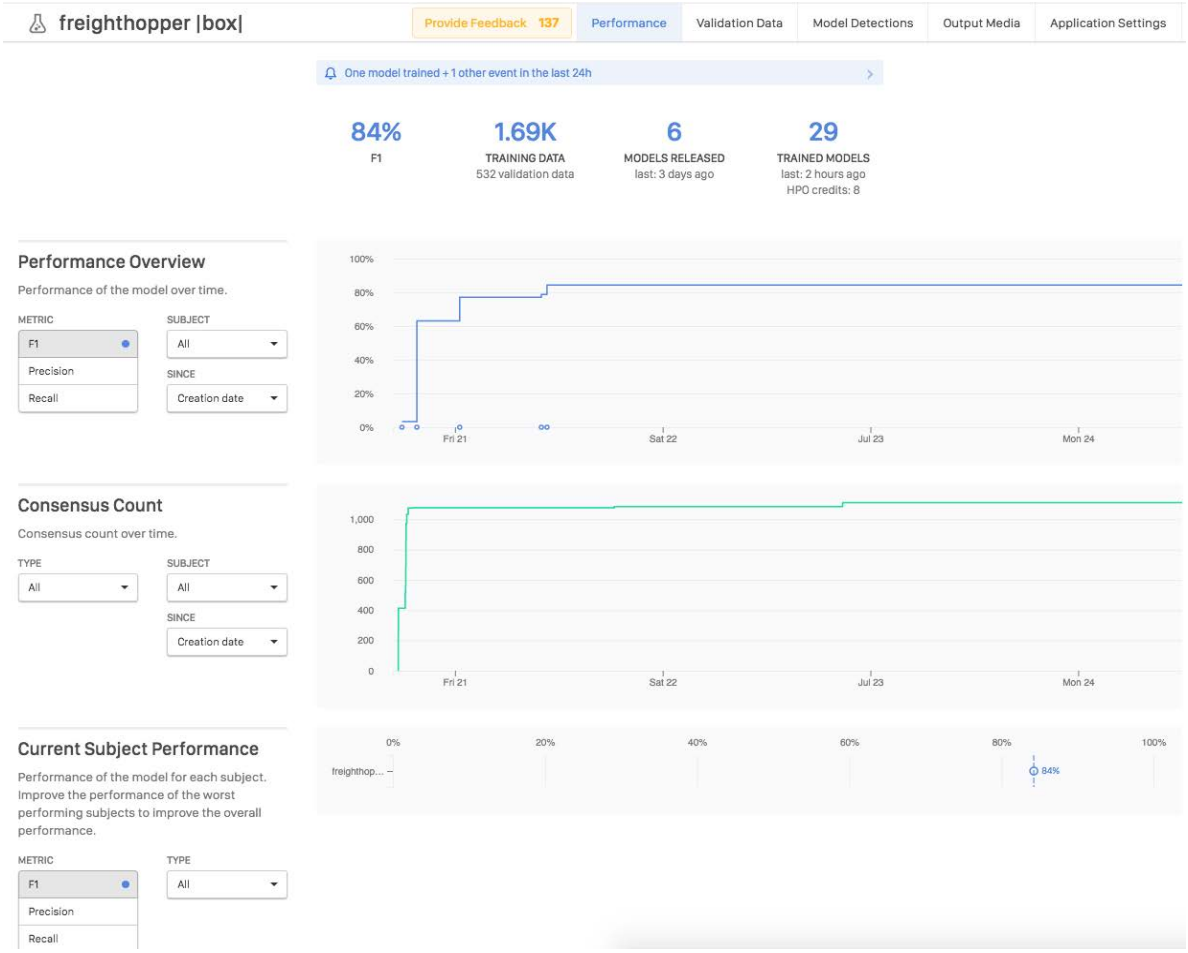


*Figure 2 - Application Overview -- (Web). Provides a dashboard view of application statistics and performance over time*

## Security and Surveillance Applications

The Cogniac Platform has been utilized in several applications that deal with real time security, perimeter security, and aerial/drone based surveillance. The system is capable of ingesting real time video streams (RTSP/HTTP) and provide detections. Some example applications are:

**Border Protection** - trail cameras for detecting illegal border crossings

**Crowd Control -** crowd density detection and crowd counting for key locations

In each of these uses, Cogniac is performing real time analysis on live video and returning detection probabilities to the integrated solutions. We have an on-premise gateway that prevents internet bandwidth bottlenecks and allows for immediate detections with the graphical processing units (GPUs).

Whether it's counting people in an area for crowd control, looking for real perimeter violations vs false alarms, or classifying events (fallen person) we can easily create those applications in a matter of hours. Traditional computer vision systems would take a large team and many man hours to do the same.

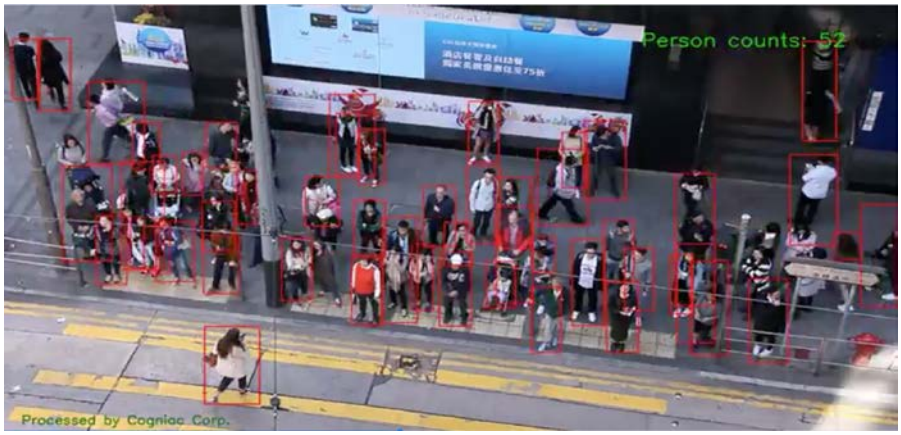


*Figure 3 - Crowd Control. The system can count people very quickly to maintain control of a particular area.*

**Illegal Rider Detection** - freight car video surveillance (illegal riders)

The following screenshot from the mobile application shows the feedback interface through

which a user is able to train the system to detect illegal riders. By confirming or drawing a box around the object of interest (illegal rider seen at the top of the train), the user is training the system in real time.

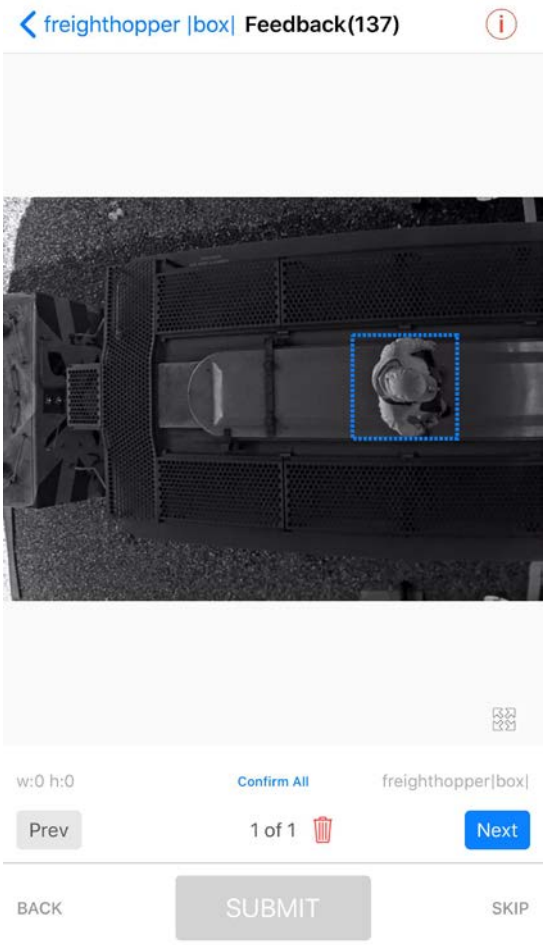


*Figure 4 - Application Feedback -- (Mobile). A domain expert can quickly provide confirmation or corrections of Cogniac application predictions.*

The following screenshots from the web application interface shows the output of the classification of a number of images. Images shown with high probability close to 100% are likely to contain illegal riders and images shown with low probabilities close to 0% are unlikely to contain illegal riders.

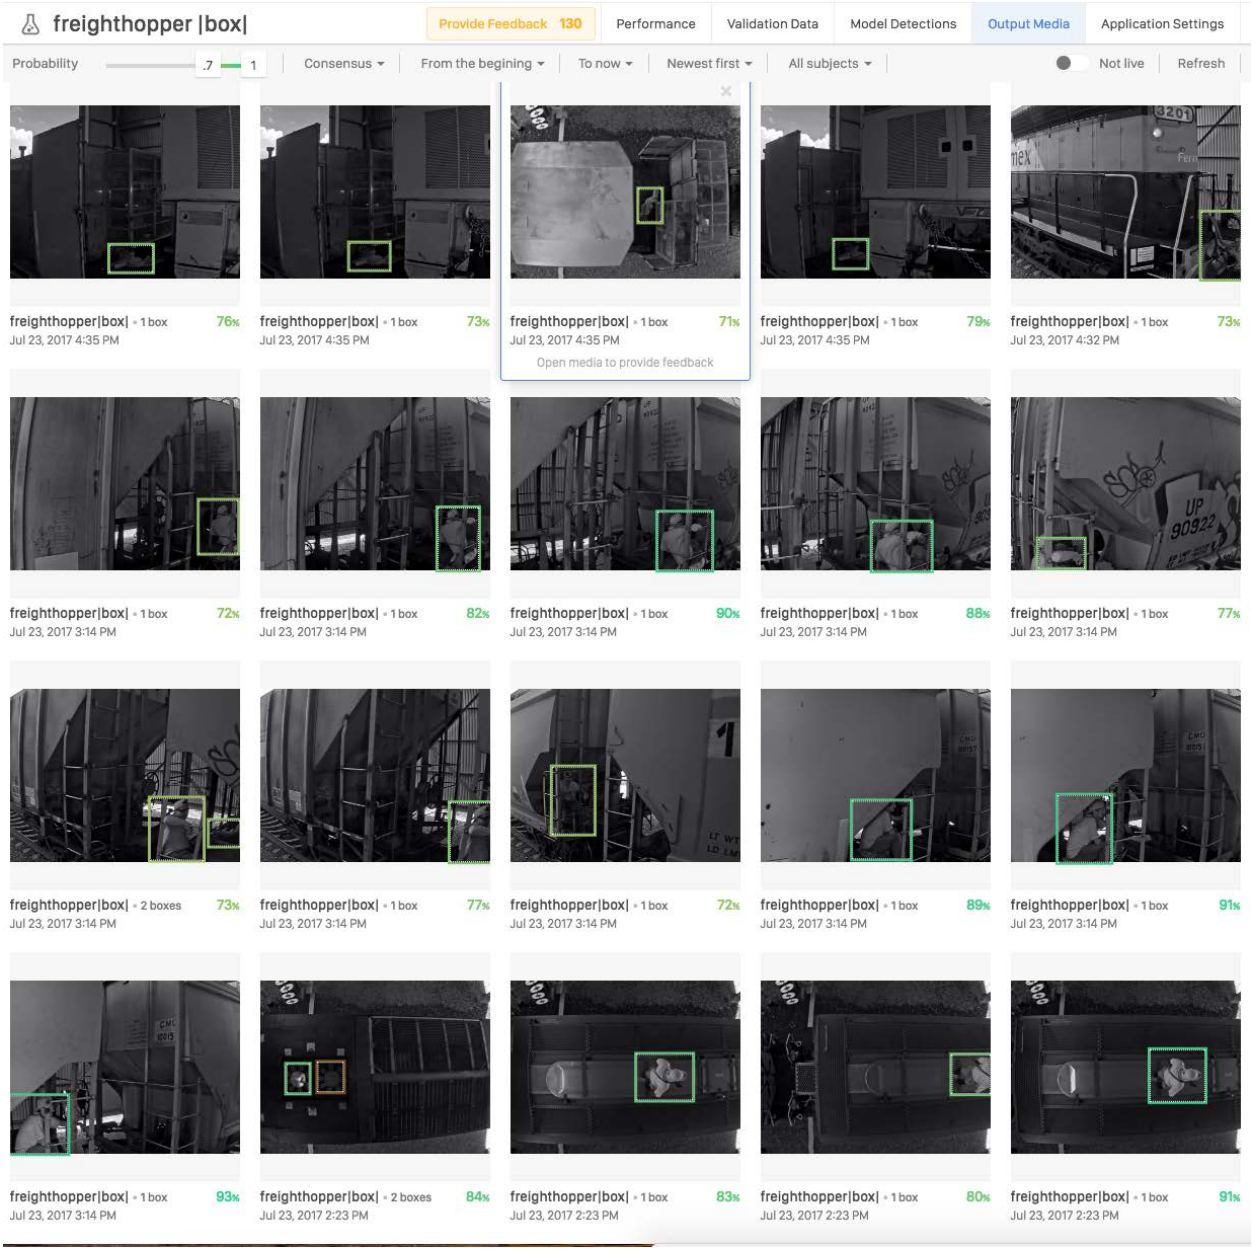


*Figure 5 - Application Detections (Web). This view provides a real-time interface to review application predictions and associated confidence levels.*

## Industrial / Manufacturing Defect Detection

The Cogniac Platform can be used to process imagery from assembly lines or during a manufacturing process to detect the presence of defects or product flaws. Some example

applications are:

**Defect Detection** - Identifying defects in products

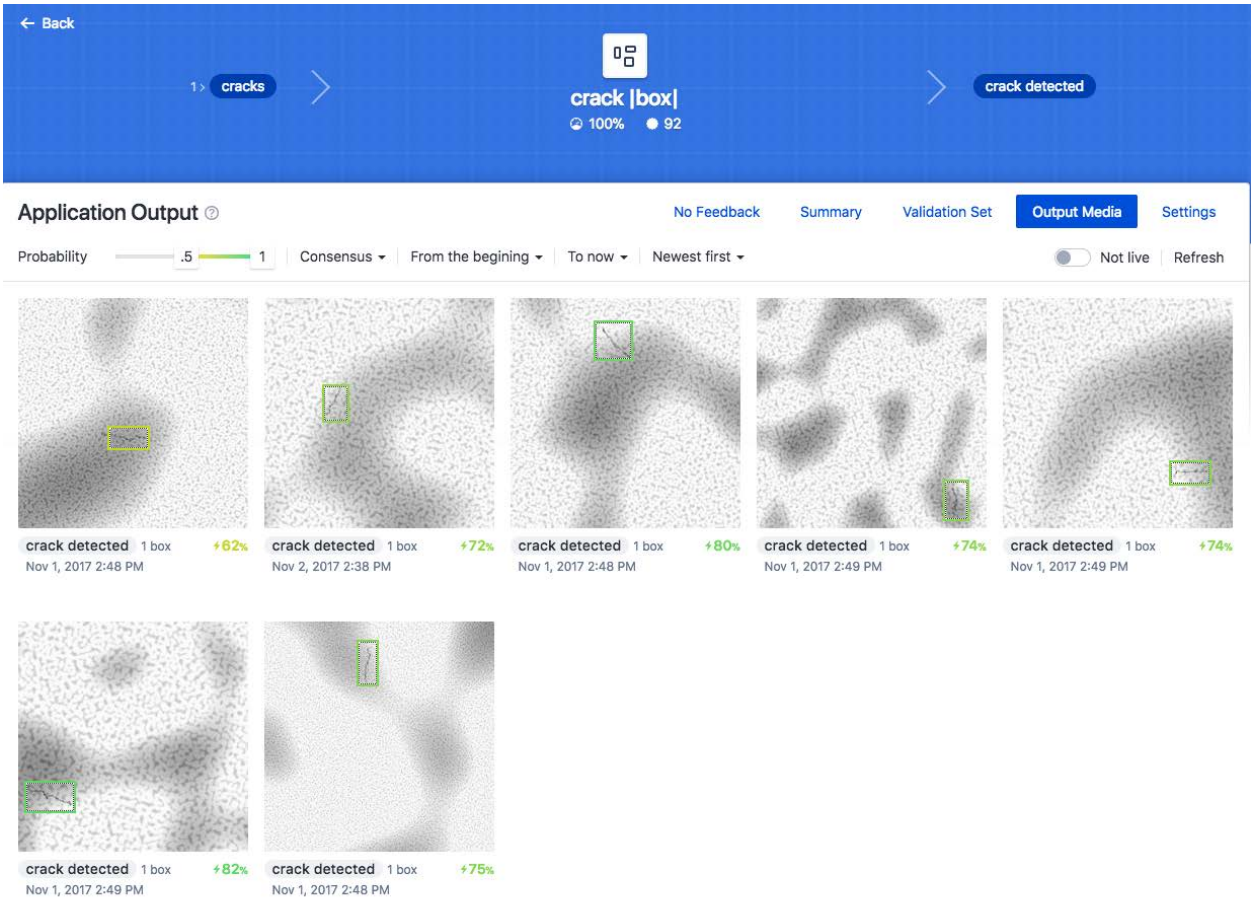


*Figure 6 - Application Detections (Web). This view provides a real-time interface to review application predictions and associated confidence levels. Bounding boxes isolate specific areas where the model thinks "defects" are present.*

## Integrating Cogniac into your solution

The full functionality of the Cogniac system is exposed through a restful API which enables seamless integration into existing systems and workflows. Documentation for the Cogniac public API is available here (https://cogniac.readme.io/docs/).

## Cogniac Advantages

### 1. Flexible & Adaptable System

The Cogniac system generates customized deep convolutional network models for individual tasks based on example imagery and sparse user annotations or feedback. Once deployed, it can be easily retargeted to perform new tasks using the existing imagery or new tasks on different imagery. If training imagery is available, the time to create new applications using the system is typically from 2 to 10 days. A few hundred images of each category is all that is required to begin training a new application model.

### 2. Continuous Improvement, with human in the loop to minimize errors

While Cogniac can operate in both fully automatic mode as well as semi-automatic mode, for most real-world tasks the semi-automatic mode is more appropriate. In semi-automatic mode the system continuously monitors the confidence level of every new image or video prediction. Predictions which have a low confidence level are prioritized based on the confidence for review by a human expert user. This accomplishes several important goals:

1. Fully automatic processing of the 'normal' images or otherwise 'easy' images in which the deep convolutional neural network model is highly confident in its prediction

2. Preserving the human expert's attention for reviewing only the most confusing images for the neural network

3. Generating new training data for the neural network based on the results of the review feedback from the human expert user. This is precisely the data that is most valuable to extending the model performance, ensuring continuous improvement
4. Ultimately provides better than human-level accuracy by optimally combining strengths of man and machine to avoid well-understood limits of human attention span on repetitive observation tasks

### 3. Superior Technical Approach

The Cogniac system is based purely on deep convolutional neural networks, which provide over 100x higher model capacity compared to previous computer vision approaches. This means that the applications models can absorb significantly higher levels of real-world image variance and provide application performance levels in the 90-99% range prior to any human

corrections.

Unlike other workflows, the Cogniac system does not require manual intervention by data scientists to create each model. The Cogniac hyperparameter optimization system is constantly searching for variants of world-class deep convolutional models, model configuration and image preprocessing to find and train optimal models.

## Contact us

Eric Stresen-Reuter
ericsr@cogniac.co
407-920-1152

Please visit our website for additional information and application examples for our solution - www.cogniac.co