intel®

Intel and Vispera ShelfSight*: Effective On-Shelf Inventory Management for Grocers

Quantifying and preventing the loss of sales opportunity in grocery retail using high-performance image recognition

Authors

Dr. Ceyhun Burak Akgül
General Manager and CTO
Vispera Information Technologies

Dr. Erdem Yörük
Chief Scientist
Vispera Information Technologies

Dr. Farhaan Mohideen
UK Business Unit Lead for Retail
Intel Corporation

Table of Contents

Executive Summary

Out of stock, over stock and shrinkage contributes to nearly USD 1.1 trillion in lost revenue for retailers worldwide[1]. Adding the loss incurred due to sales return, the so-called ghost economy totals USD 1.75 trillion[2]. In a world where the consumer expectation is at an all-time high, product availability and placement are fundamental. There are various strategies to save a sale in the event of stock not being available or misplaced, but these should be considered as a fall-back plan, more appropriate for non-consumables or fashion retailers.

There is also the hard-to-quantify issue of damaged customer loyalty. Shoppers today are increasingly used to finding their favourite products whenever and wherever they want, and they are rather quick to look for the product at the next alternative outlet. Operating in this fiercely competitive market segment, grocers cannot afford to lose customers because products are not available, misplaced, or presented with misleading price tags. Losing out on high value ticket items is a particular problem for low margin grocery retailers.

Much consideration is given to product placement and promotions at the outset of a particular grocery retail implementation. However, if retailers don't ensure that the products shoppers are looking for are available at all times, in the right place, then the investments that they, and their supplier brands make, are futile.

Artificial Intelligence (AI) is opening up various avenues that help retailers analyse essential data at the edge to improve inventory management. AI provides a real-time or near real-time view of all the data required to make smarter decisions much faster. Computer Vision at the edge plays an overwhelmingly important role in enabling an AI architecture, paving the way to intelligent IT platforms like Vispera ShelfSight*. It is predicted these systems will enable businesses to learn and make critical decisions, leading to increased customer satisfaction, higher product sales and greater operational efficiency.

---

[1] Greg Buzek, "Retail's $1.1 Trillion Inventory Distortion Problem," IHL Group,
http://ihlservices.com/wp-content/uploads/2015/06/WeLostAustralia-Outline.pdf

[2] "Retailers and the Ghost Economy $1.75 Trillion Reasons to be Afraid", IHL Group,
https://goo.gl/SVnbGR

Intelligence in the Store

The number of connected devices in our stores is growing and with it the amount of data captured. The Internet of Things (IoT) has enabled the data economy and by bringing AI to these connected devices, retailers are able to gain actionable insights to transform business success. Analysing captured data at the edge provides an enormous opportunity to generate near real-time information to help retailers make better decisions in a timely manner.

Visual data is being captured by surveillance systems, smart cameras, digital signage, and Point of Sale (PoS) systems. It is far more important to act upon this data quickly in the store, rather than sending it back to a remote data centre for processing and subsequent analysis. Captured visual data has a direct impact on the retailer's bottom line. Consider the examples below and how a minute-by-minute or perpetual accuracy on store activity could dramatically increase the revenue and enhance the customer experience:

- Recognising the customer when in the store could provide multiple opportunities for the retailer to provide more customised experiences with relevant products and services. For example, a customer standing in front of a digital display could be offered an immediate discount on a product of interest, or have it bundled with other products to offer a multi-product promotion.

- Cameras are able to pick up persons of interest or images of theft to help security take immediate action. Smart cameras are also able to detect products and prices marked on shelves with insurmountable accuracy in comparison to humans performing the same task.

- Smart systems build upon these capabilities to provide real-time updates if products are missing or misplaced, instead of a customer having to point out unavailability or any discrepancies in price with an ongoing promotion versus standard pricing.

- From a higher-level view, data can be used to test hypotheses about "selling" product placement, assortment and space planning, and to analyse in-store traffic patterns so that layouts and product presentation is done more accurately.

All these scenarios could allow retailers to compete more effectively against those who do not take real-time advantage of the data being generated in their stores[3].

Cost of Out of Stock and How to Prevent it

There are two facets to a product being unavailable (or misplaced on the shelf or presented with the wrong price tag):

- Firstly, *an out of stock event associated with an especially frequently selling item directly translates to loss of sales opportunity, which can be quantified quite easily.* A reasonable estimate for the loss of sales opportunity associated with a frequently sold item can be calculated as *pNK*, where:

  - *p = unit price of the product*

  - *N = average number of units sold during a certain timeframe (determined by seasonality, promotions, or other factors)*

  - *K = number of timeframes during which out of stock events are observed*

The value of revenue loss expression above scales up for higher priced products (increasing p), for frequently selling fast mover products (increasing N) and by the number of cycles during which the retailer failed to replenish the product (increasing K).

The situation is further aggravated when out of stock events associated with other fast movers are also taken into account.

The global annual loss figure for retailers accumulated over all such out of stock events, amounts to USD 634.1 billion, more than half of the total USD 1.1 trillion annual loss due to inventory distortions[1].

A 2007 report by Gruen and Corsten[4] provides a detailed account of various out of stock metrics quantifying different aspects of the situation along with an examination of root causes and remedies for out of stock reduction, if not total prevention.

- A further important point emphasised in the Gruen and Corsten report constitutes the second, maybe even more important, facet of out of stock. Simply put, *the cost of out of stock extends well beyond the lost sales of the out of stock item alone.* A host of strategic and operational hard-to-quantify costs apply to both retailers and suppliers, including damaged customer loyalty as well as store and brand equity, and reduced impact of promotions. Furthermore, out of stock distorts demand and leads to inaccurate forecasts.

[3] One Door Merchandising Research Reveals Five Areas Where Retail Winners Excel, online on January 15, 2018, https://www.businesswire.com/news/home/20180115005273/en/Door-Merchandising-Research-Reveals-Areas-Retail-Winners

[4] Thomas W. Gruen, Daniel Corsten, "A Comprehensive Guide To Retail Out-of-Stock Reduction In the Fast-Moving Consumer Goods Industry", 2007, https://www.nacds.org/wp-content/uploads/2017/02/Retail-Out-of-Stock-Reduction.pdf

For retailers, an additional cost is the time employees spend trying to satisfy shoppers who ask about a specific out of stock item. In Gruen and Corsten's report, the latter cost is reported to be USD 800 per week for a typical US grocery store. For shoppers, valuable time is lost waiting for resolution instead of continuing shopping - the net effect for retailer being loss of sales opportunity regarding other possibly available items.

More than a decade has passed since the first authoritative research study on the effects of out of stock and its potential remedies, yet the issue hasn't diminished for grocery retailers. The key to the solution of the out of stock problem in retail, as in virtually all walks of business where similar operational problems occur, is *sustainable and scalable* execution driven by *timely and accurate* measurement at scale. While effective and efficient supply chain management solutions are in place in all modern grocers to mitigate the *distribution centre* and *store out of stock* to a large extent, approaches to prevent *shelf out of stock* remain deceptively basic, inefficient, and ineffective.

Mobilise store personnel on a periodical basis to go check shelves and replenish so that out of stock events can be detected by the human eye. The same personnel should also arrange the products on the shelf according to a planogram – a visual diagram providing sheer amount of detail on the placement of every product on the shelf, check the integrity and position of price tags, and ensure promotion actions. All these operations are to be performed multiple times every single day at each shelf at each store. Given the circumstances, no wonder there is so much distortion on retail shelves hence so much loss of opportunity.

In a recent pilot with a retailer growing its operations across Turkey and transitioning to planogram implementation at its outlets, Vispera found that, on an average day, more than two thirds of all missing products were actually available in the store (in the depot) but were not brought to shelf – an example clearly illustrating how much shelf out of stock weighs over store out of stock.

The root cause is clear. The store personnel cannot catch up with the speed and level of detail necessitated by everyday operations. The problem is not one of insufficient workforce. On the contrary, deploying more people would only add to operational inefficiency at the questionable benefit of a marginal effectiveness improvement. Humans are not to blame, as the ultimate task at hand, of detecting several missing objects (products) in a somewhat ordered arrangement of hundreds with speed and accuracy, is not one where humans excel, never did and never will.

Store personnel need a power tool telling them what is missing and where, with unparalleled speed and accuracy, so that they can go pick the missing product from back of the store and then replenish the shelf immediately. This is the simplest yet most effective means to prevent shelf out of stock. That much needed power tool is brought to retailers by a particular form of AI: high-performance image recognition.

Anatomy of Vision-Based Shelf Measurement

A typical retail aisle is an immensely data-rich environment. The information it contains is sheer and dispersed into shelves, displays and promotions that carry hundreds to thousands of products stacked together, looking very similar but belonging to almost as many distinct labels and associated to almost as many price tags. Despite being entirely visual, extraction of this information to its every detail for an accurate shelf measurement is obviously a daunting and often error-prone task for humans. Fortunately, Computer Vision can be of great use, as it has reached levels on par with human perfection with the advent of deep learning technologies, in particular the so-called deep convolutional neural networks (CNNs) a class of feed-forward artificial Neural Networks (NNs).

While the theory and early practice of CNNs already date back to 1989[5] (and even to 1940s and onwards in the case of NNs), their reign in Computer Vision has been indisputably sealed only in the 2010s, once the levels of available training data and processing power had caught up with their potential. Today, almost all automated visual tasks, ranging from self-driving vehicles to face recognition owe their success to the use of CNNs and related deep learning technologies that continue to emerge.

Here is how CNNs work in a nutshell. The input to a typical convolutional network is a digitised imagery of the scene, given in raw pixel values, and the output is any semantic abstraction of the input (e.g. labels of present objects), which the network is specifically trained to predict. The black box in between hides a deep sequence of "convolutional" layers and other certain nonlinearities, each cascaded to process the output of the layer before.

Each convolutional layer carries a large variety of image filters - or neurons, which fire whenever they detect a certain pattern on the visual input they receive. When trained on real visual data, these patterns end up being very primitive and common (e.g. oriented edges, corners) in the shallow layers, where receptive fields i.e. spatial scopes of neurons are small, while they get more and more complex and specialised (e.g. shapes, faces) as one goes deeper, where semantically meaningful configurations can emerge within larger receptive fields.

[5] Y. LeCun, L. D. Jackel, B. Boser, J. S. Denker, H. P. Graf, I. Guyon, D. Henderson, R. E. Howard and W. Hubbard: Handwritten Digit Recognition: Applications of Neural Net Chips and Automatic Learning, IEEE Communication, 41-46, invited paper, November 1989
https://link.springer.com/chapter/10.1007/978-3-642-76153-9_35

Arguably, such deep hierarchy of "parts-to-whole" relations resemble the organisation in the human visual cortex, and unlike earlier conventional methods does not require any hand-engineering. Rather it emerges as such while all filter patterns and parameters are learned end-to-end from visual data. Moreover, once learned they can be transferred across different visual tasks. Such aspects, along with the exceptional performances put deep CNNs in the centre stage of today's visual AI.

Nonetheless, applying CNNs to the problem of retail shelf measurement with (near) real-time actionable results, presents whole new challenges. To begin with, we have the sheer complexity of any given retail scene. Tasking a single-shot deep network for the whole measurement process may in theory be perfectly possible, but the detail in the output and the requirement for a combinatorically rich training data would rather mandate a coarse-to-fine division of inference, which a typical human eye would also do.

First, determine the area of interest in the scene, namely the sections and rows of the shelf, then segment and locate the individual products present on them, and finally recognise their labels and read their prices from the nearest detected price tag.

As a result, each subtask calls for the design of a dedicated CNN, which now is to be conveniently trained on a much more constrained data. In particular, for retail shelf measurement we would deploy networks that perform:

- The detection tasks, i.e. segmentation of shelves into regions of interest, localisation of the products and their price tags

- The recognition tasks, i.e. fine-grained classification of the detected product crops and reading their prices.

Depending on the positioning of the camera, another task would be to resolve for occlusions due to the dynamics of a retail scene with shoppers, carts and floor staff in the camera field of view. For each of these tasks, the state-of-the-art provides specific CNN models including region proposal networks (e.g. RCNN), and recurrent neural networks (e.g. LSTM).

Deep neural networks alone do not complete the shelf measurement, they merely provide the raw output in the form of present product labels, prices and their positions on the shelf. Curating these into a coherent shelf measurement with more involved semantics, and presenting them as leanly as possible, with instant alerts whenever a shelf replenishment is required, are further tasks that can be quite demanding depending on the target key performance indicators (KPIs). For instance, measuring the planogram compliance alone requires dedicated smart solutions that register the vision output against a generic target layout that is abstracted before.

Aside from the aforementioned design issues, probably the most prominent engineering question involves the hardware choices and optimisation, especially when it comes to cost efficiency, speed, system and data maintenance.

Deep learning methods are infamous for their thirst for computing power and the common choice has often been Graphics Processing Units (GPUs) at some distant compute server. However, real-time shelf measurement sits at a crucial junction, where edge or near-edge computing is preferable. Possible alternatives include an integrated GPU, or Field Programmable Gate Array (FPGA) on the camera gateway device or compute sticks like the Intel® Movidius™ Neural Compute Stick (NCS). Each option presents its own pros and cons.

In general, it is predicted that the GPU option presents the highest performance gains over CPU, especially with new model optimisation solutions provided by their manufacturers. At the same time, using custom models with such GPU-based solutions and common platforms such as TensorFlow is estimated to be more flexible. However, a major concern with the GPU is the high system cost.

On the other hand, options such as FPGA-based Intel® Deep Learning Inference Accelerator (Intel® DLIA) and the Intel Movidius NCS are thought to be much cheaper, and easier to scale as an on-premise solution. To achieve gains over a CPU, these new generation compute devices can be optimised using the Intel® Distribution of OpenVINO™ toolkit.
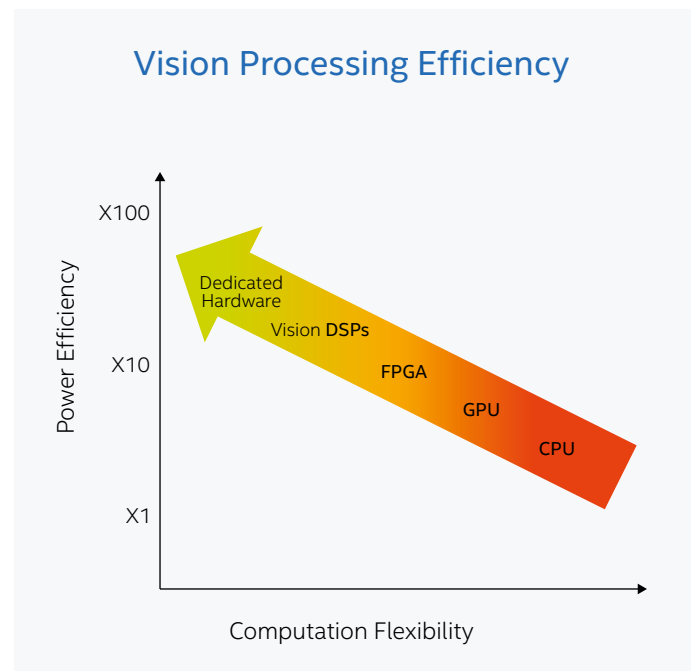

Vision Processing Efficiency

Figure 1. Power efficiency versus computation flexibility of major compute options in vision processing.

The main driver for using an FPGA is the efficiency it brings to the edge device especially on memory bound workloads and the ability to maintain accuracy with higher performance. FPGAs are important tools for many deep learning applications, from search, to speech recognition and to visual content analysis. FPGAs reside also on midway amongst the competing options in terms of power efficiency versus computation flexibility trade-off – see figure 1.

Vispera ShelfSight*: A Visually Smart Solution for Grocery Shelves

Vispera ShelfSight is a real-time shelf monitoring and management system that can analyse retail shelves for out of stock detection, planogram compliance and empty space detection using state of the art deep learning algorithms developed by Vispera. The system is capable of detecting and recognising individual products to produce highly accurate shelf measurements in real-time utilising IoT cameras mounted on the aisles.

With Vispera ShelfSight, retailers can now have access to the digitised view of the store in real-time. In today's competitive environment, ShelfSight is a powerful solution for retailers, giving an edge on operational excellence by enabling immediate to-the-point store actions.

How it works: visual data capture and processing

Vispera ShelfSight is a hybrid on premise and cloud system. In ShelfSight, visual shelf data captured by image acquisition units supporting multiple pre-calibrated shelf-facing cameras are transferred to the on-premise backend via an edge application hosted on a lightweight PC serving as a gateway. The backend administers the gateway applications hence cameras, manages and schedules measurements, serves the image data to the deep NN–based visual inference applications, pushes the output measurement results and alerts to local front-end applications, and detailed shelf and product data to the cloud.

While local front-end applications are dedicated for real-time shelf monitoring in the form of out of stock and non-compliance alerts guiding the store personnel what to replenish and where via friendly user interfaces, detailed shelf data are transferred to the cloud for further aggregation, visualisation and analytics – see figure 2.

The whole system is based on a microservice architecture where applications communicate and exchange data via Rest APIs. As it can serve third-party applications with its data and consume data from them quite easily, ShelfSight qualifies as a versatile in-store IoT platform potentially integrating a rich battery of applications relevant to the retail outlet environment. These include but are not limited to people counting and tracking, analysis of shopping habits and patterns, understanding of shopper-product interaction, automated ordering, and fusion with other in-store sensors, readers and devices, such as RFID, beacons, and digital signage.
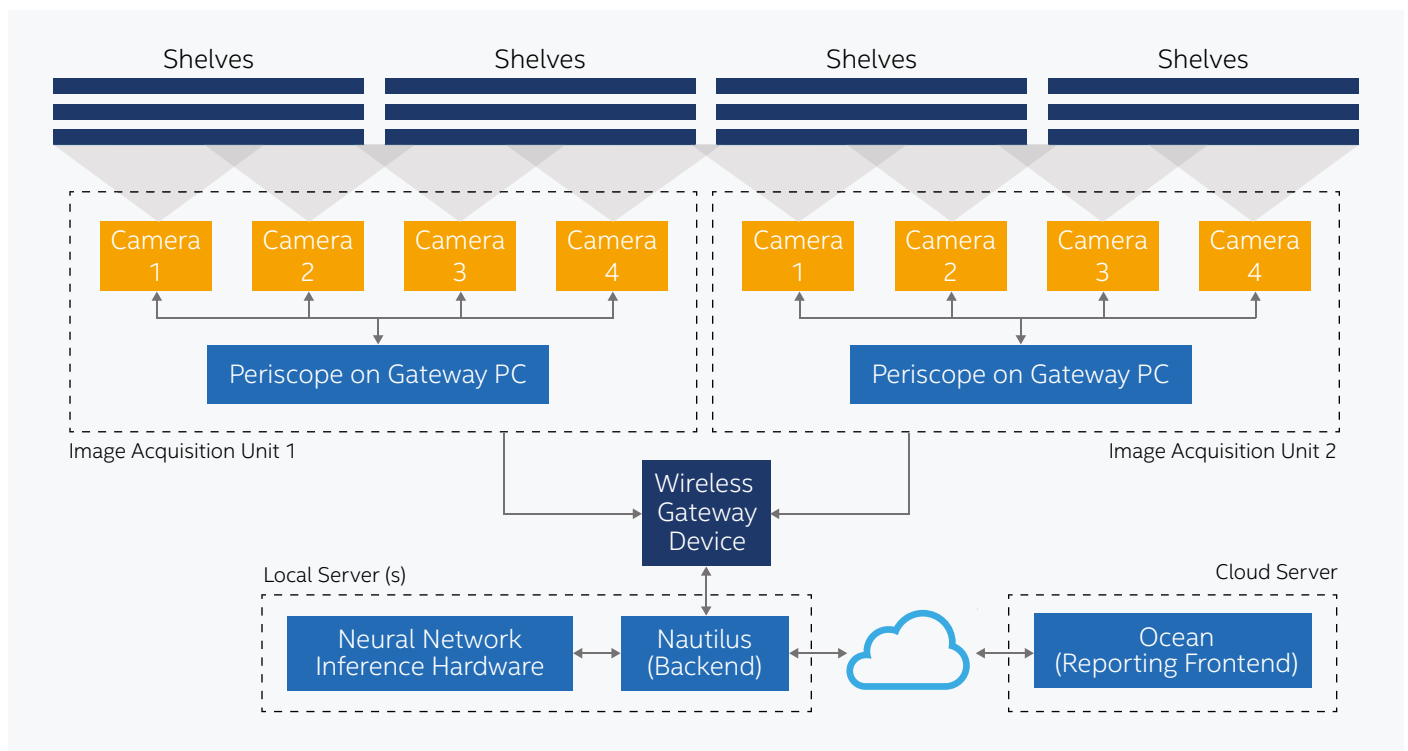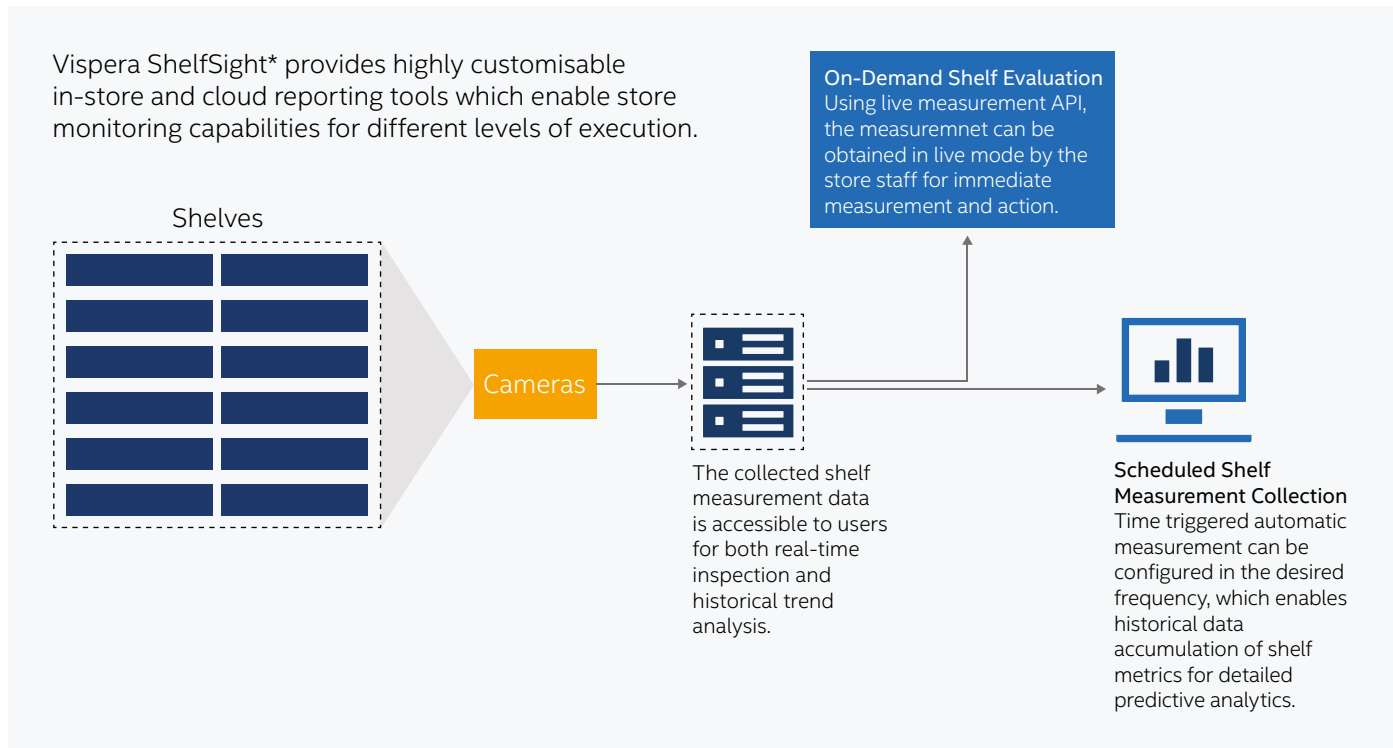

Figure 2. Layout of Vispera ShelfSight* components and their interconnection.

Vispera ShelfSight* provides highly customisable in-store and cloud reporting tools which enable store monitoring capabilities for different levels of execution.

Shelves

Cameras

The collected shelf measurement data is accessible to users for both real-time inspection and historical trend analysis.

On-Demand Shelf Evaluation
Using live measurement API, the measuremnet can be obtained in live mode by the store staff for immediate measurement and action.

Scheduled Shelf Measurement Collection
Time triggered automatic measurement can be configured in the desired frequency, which enables historical data accumulation of shelf metrics for detailed predictive analytics.

Figure 3. Vispera ShelfSight* provides shelf measurements in scheduled mode or on-demand.

How it works: alarms, KPIs, and analytics

ShelfSight provides shelf measurements in scheduled mode or on-demand – see figure 3. As pointed out earlier, the key to out of stock prevention is to solve the situation as soon as the event occurs.

In scheduled mode, ShelfSight issues measurements regularly at desired pre-set frequency, detects out of stock events (along with other secondary non-compliance issues) and alerts the store personnel via a push notification on the personnel's mobile device and on screens placed in the store. Products known to be unavailable in the whole store due to supply reasons can be flagged as "don't care" at any time to make sure the personnel get always the relevant notification.

In on-demand mode, the measurement request can be directly issued by the store personnel whenever they want. Out of stock and other non-compliance issues such as excess, missing or misplaced facings, shelf gaps, and inadequate shelf space allocated to brands are reported to the productivity front-end in the form of a task list visually guiding the personnel to make the perfect shelf specified by a reference planogram.

While planograms can be specified by enterprise planogram software used by the retailer, actual images captured by ShelfSight after replenishment can be assigned as reference. This last feature enabled by ShelfSight is very handy especially when it is impractical and expensive to maintain and distribute store-level planograms prepared in the retailer's back office.

Retail execution KPIs and views provided by Vispera ShelfSight*

Out-of-stock SKUs – List of SKUs that are present in the planogram but missing from the shelf

Realogram view – The realogram highlights excess and misplaced items on a digitised version of the shelf scene, thus enabling quick correction of product placement errors

Planogram view – The planogram highlights the missing items with respect to the ideal product placement of the shelf scene, thus enabling quick replenishment of missing products

Planogram compliance – Performance metric related to the number of correctly placed, misplaced or missing items indicating the deviation of the shelf from ideal product placement

Product shelf share – Percentage of brands displayed in the shelf, calculated from occupied shelf space

Product facing share – Percentage of brands displayed in the shelf, calculated from facing counts

Facing score – Performance metric indicating the shelf utilisation level compared to the ideal product placement

Note that along with live measurement features presented to the store personnel, ShelfSight sends all shelf data to the cloud in the background for the calculation of additional retail execution KPIs, data aggregation, group and trend visualisation as well as data analytics enabled by Vispera's web-based reporting applications accessible anytime from anywhere with any Internet-connected device having a web browser. Concurrently, ShelfSight reporting API can be used to integrate the data to the retailer's enterprise systems.

Conclusion

Out of stock poses tremendous challenges to retailers and it is impacting their bottom line. Progressive retailers are adopting new digital strategies to understand their customers, streamline their business operations, increase profitability and stay competitive.

Computer Vision has emerged as a great digital strategy, providing retailers with a vast amount of store data. Computer Vision coupled with the advances in deep learning algorithms provides an excellent opportunity to automate many activities in the store that was previously done manually. Out of stock and other non-compliance issues can be automatically alerted by building intelligence at the edge.

Vispera's ShelfSight solution provides an integrated end-to-end smart shelving solution backed by Intel's comprehensive portfolio of hardware and software tools to enable retailers to address store operations effectively.

Find the solution that is right for your organization. Contact your Intel representative or visit **https://www.intel.com/ content/www/us/en/internet-of-things/vision-and-ai-to-analize-data.html.**

Learn More

You may find the following resources useful:

- Vispera Information Technologies: http://vispera.co/
- Intel® Vision Products: https://www.intel.com/ content/www/us/en/internet-of-things/vision-and-ai-to-analize-data.html
- Intel® Technology for Retail: https://www.intel.com/ content/www/us/en/retail/retailsolutions.html
- Intel Digital Signage solutions: https://www.intel. com/content/www/us/en/retail/digital-signage/ overview.html

Solution Provided By: