



(12) **DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION**

(13) **A1**

(86) Date de dépôt PCT/PCT Filing Date: 2020/04/13
 (87) Date publication PCT/PCT Publication Date: 2020/10/15
 (85) Entrée phase nationale/National Entry: 2021/10/07
 (86) N° demande PCT/PCT Application No.: US 2020/027980
 (87) N° publication PCT/PCT Publication No.: 2020/210822
 (30) Priorité/Priority: 2019/04/11 (US62/832,755)

(51) Cl.Int./Int.Cl. *G06Q 10/08* (2012.01)
 (71) Demandeurs/Applicants:
 CARNEGIE MELLON UNIVERSITY, US;
 BOSSA NOVA ROBOTICS, US
 (72) Inventeurs/Inventors:
 SKAFF, SARJOUN, US;
 SAVVIDES, MARIOS, US;
 AHMED, UZAIR, US;
 NALLAMOTHU, SREENA, US;
 TAO, RAN, US;
 MOHAN, NIKHIL, US
 (74) Agent: SMART & BIGGAR LLP

(54) Titre : SYSTEME ET PROCEDE POUR ASSOCIER DES PRODUITS ET DES ETIQUETTES DE PRODUITS
 (54) Title: SYSTEM AND METHOD FOR ASSOCIATING PRODUCTS AND PRODUCT LABELS




FIG. 2

(57) **Abrégé/Abstract:**

An automated inventory monitoring system includes an image capture module able to create an image of an aisle of a retail store. Images of products labels are identified in the image and classified as shelf labels or peg labels. For shelf labels, an area of the shelf is defined and associated with the shelf label. Images of products are identified in the image and products on the shelf within an area associated with a shelf label are associated with the shelf label. Products located below a peg label are associated with the peg label. Based on the association between labels and products, out-of-stock products, plugs and spread may be detected and reported to the staff of the retail store.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau

(43) International Publication Date
15 October 2020 (15.10.2020)



(10) International Publication Number
WO 2020/210822 A1

- (51) International Patent Classification:
G06Q 10/08 (2012.01)
- (21) International Application Number:
PCT/US2020/027980
- (22) International Filing Date:
13 April 2020 (13.04.2020)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/832,755 11 April 2019 (11.04.2019) US
- (71) Applicants: **CARNEGIE MELLON UNIVERSITY** [US/US]; 5000 Forbes Avenue, Pittsburgh, Pennsylvania 15213 (US). **BOSSA NOVA ROBOTICS** [US/US]; 610 22nd Street, Suite 250, San Francisco, California 94107 (US).
- (72) Inventors: **SKAFF, Sarjoun**; 2645 Railroad St., Apt 128, Pittsburgh, Pennsylvania 15222 (US). **SAVVIDES, Marios**; 506 Eagle Court, Wexford, Pennsylvania 15090 (US). **AHMED, Uzair**; 5551 Centre Avenue, Apt 602, Pittsburgh, Pennsylvania 15232 (US). **NALLAMOTHU, Sreena**; 5551 Centre Avenue, Apt 411, Pittsburgh, Pennsylvania 15232 (US). **TAO, Ran**; 5701 Centre Avenue, Apt. 801, Pittsburgh, Pennsylvania 15206 (US). **MOHAN, Nikhil**; 845 UN Plaza 20C, New York, New York 10017 (US).
- (74) Agent: **CARLETON, Dennis M.**; Kacvinsky Daisak Bluni PLLC, 2601 Weston Parkway, Suite 103, Cary, North Carolina 27513 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,

(54) Title: SYSTEM AND METHOD FOR ASSOCIATING PRODUCTS AND PRODUCT LABELS




FIG. 2

(57) Abstract: An automated inventory monitoring system includes an image capture module able to create an image of an aisle of a retail store. Images of products labels are identified in the image and classified as shelf labels or peg labels. For shelf labels, an area of the shelf is defined and associated with the shelf label. Images of products are identified in the image and products on the shelf within an area associated with a shelf label are associated with the shelf label. Products located below a peg label are associated with the peg label. Based on the association between labels and products, out-of-stock products, plugs and spread may be detected and reported to the staff of the retail store.



WO 2020/210822 A1

WO 2020/210822 A1 

KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

- (84) Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

- *with international search report (Art. 21(3))*

SYSTEM AND METHOD FOR ASSOCIATING PRODUCTS AND PRODUCT LABELS**Related Applications**

[0001] This application claims the benefit of U.S. Provisional Patent Application No. 62/832,755, filed April 11, 2019, entitled “Shelf Monitoring System and Method”. The contents of this application is incorporated herein in its entirety.

Field of the Invention

[0002] The present invention is related to the field of automated inventory monitoring in a commercial retail setting and, in particular, is directed to systems, processes and methods for automatically tracking products displayed in the retail setting through the use of a mobile robot having a multiple camera sensor suite mounted thereon.

Background of the Invention

[0003] Retail stores, for example, grocery stores, general merchandise stores, dry goods stores or warehouse style stores can have thousands of distinct products that are often concurrently offered for sale. Stores are typically laid out in an aisle configuration wherein each aisle may have shelves of products placed on one or both sides of the aisle. At the ends of the aisle, the shelves will typically have “end caps” which often contain products that the store wishes to promote to its customers. As such, the contents of the end caps may frequently change. In addition, the inventory of the stores may constantly be modified by removing, adding or repositioning the products. As customers purchase the products, products may become out-of-stock and may need to be re-ordered from a wholesaler.

[0004] The shelves in the store are typically provided with shelf labels. The shelf labels serve two purposes. The first is the identification of the product which is to be placed on the shelves in close proximity to the shelf label. The label may comprise bar code or QR code printed on the shelf label identifying the product. The shelf label also typically contains the unit price of the product and may contain other miscellaneous information specific to the particular

store.

- [0005]** The second purpose of the shelf label is to indicate a position on the shelf where the product should be placed. For example, a particular store may place the shelves labels at the far left of the area on the shelf where the associated product is to be positioned (i.e., left justified product placement). It is therefore incumbent on the staff of the store to properly place the products when restocking the shelves. This will also aid the system of the present invention as it attempts to match the actual products on the shelves with the product labels.
- [0006]** Even with frequent restocking schedules, products assumed to be in-stock may be out-of-stock, decreasing both sales and customer satisfaction. Point of sales data can be used to roughly estimate product stock levels, but does not help with identifying misplaced, stolen, or damaged products, all of which can reduce product availability. However, manually monitoring product inventory and tracking product position is expensive and time consuming.
- [0007]** One solution for tracking product inventory relies on planograms which are typically manually created for each individual store, in combination with machine vision technology. Given a planogram, machine vision can be used to assist in shelf space compliance. In such cases, the planogram may need to be manually created and manually updated each time a product is removed, added or repositioned within the store.
- [0008]** To implement machine vision technology relying on a planogram, one or more fixed position cameras can be used throughout a store to monitor aisles, with large gaps in shelf space being checkable against the planogram or shelf labels and flagged as “out-of-stock” if necessary. Alternatively, a number of movable cameras can be used to scan a store aisle. Even with such systems, human intervention is generally required to build an initial planogram that correctly represents the product layout on the fixture, and that includes detailed information relative to a bounding box that can include product identification, placement, and count. Substantial human intervention can also be required to update the planogram, as well as search for misplaced product inventory.

[0009] As such, it would be desirable to be able to automate the tracking of inventory to determine when various products are out-of-stock, have been repositioned, or are otherwise not where they are expected to be. In addition, it would be desirable to be able to implement such a system without the need for the manually created planograms.

Summary of the Invention

[0010] Shelf monitoring and product tracking systems, methods and processes are disclosed herein. In preferred embodiments, a mobile, autonomous robot having a plurality of cameras mounted thereon navigates the aisles of the store to collect images of products on shelves and other fixtures such as pegs in the store. In other embodiments, images of products and fixtures in the store may be collected using any type of camera, including, without limitation, fixed-location cameras, individual images. Images of each aisle may be created and analyzed to determine the identity and status of products on the fixtures, the type and state of the fixtures, and other information about store environment. For example, the system may be capable of determining when products are out-of-stock, miss-positioned with respect to their proper positions on the shelves or wherein a product has been moved to an incorrect position in the store by a customer. In addition, the system is capable of determining when products have been moved by the store to another area of the shelf, removed from stock, or newly added to the store's inventory.

[0011] In preferred embodiments of the invention, the system analyzes the panoramic images to detect the presence of and, optionally, to determine the identity of products placed on the fixtures. Additionally, the system can analyze the panoramic images to identify shelf labels indicating which products are expected to be at various positions on the fixtures. The system is then able to match the placement of and, optionally, the identity of the products on the fixtures with the expected positions of the products to determine that the products are shelved properly, are miss-shelved or are out-of-stock. The system is further functional to flag misplaced and out-of-stock products and alert the store's staff such that the misplacement may be corrected or such that the product may be re-stocked.

Brief Description of the Drawings

- [0012] **FIG. 1** is a block diagram of one embodiment of an autonomous robot acting as an inventory monitoring camera system.
- [0013] **FIG. 2** is an illustration of two inventory monitoring camera systems of the type discussed with respect to **FIG. 1** *in situ* in the aisle of a store.
- [0014] **FIG. 3** is an example of panoramic image of the type used by the present invention.
- [0015] **FIG. 4** is an exemplary processing pipeline used to process the panoramic images of the type shown in **FIG. 3**.
- [0016] **FIG. 5** is an example of the output of the product detector of the present invention showing products surrounded by bounding boxes.
- [0017] **FIG. 6** is a block diagram of the two-tier classifier used to classify product images.
- [0018] **FIG. 7** is an example of the output of the label detector of the present invention showing shelf labels surrounded by bounding boxes.
- [0019] **FIG. 8** is an example of the output of the shelf segment classifier of the present invention, showing a binary mask having highlighted areas indicating which pixels in the panoramic image are located on a shelf.
- [0020] **FIG. 9** shows the output of the shelf inference logic, showing bounding boxes around shelves.
- [0021] **FIG. 10** shows the association between products and shelf labels.
- [0022] **FIG. 11** is a schematic representation of the process used to determine product/shelf label association when the products are not completely within the section of a shelf allocated to a specific shelf label.
- [0023] **FIG. 12** is a diagram showing an exemplary matching pipeline for determining a pairwise match between adjacent products on a shelf for purposes of detecting plugs and spreads.
- [0024] **FIG. 13** is a schematic diagram showing a plug.
- [0025] **FIG. 14** shows the output of the matching pipeline and the output of the plug detector.
- [0026] **FIG. 15** is a schematic diagram showing a spread.
- [0027] **FIG. 16** shows the output of the matching pipeline in the output of the spread

detector.

Definitions

- [0028] A “**planogram**” is a list, diagram or map that shows how and where specific products are placed on fixtures (shelves or displays) within a store, including how many facings for each product (distinct rows of the product) and the quantity of each product that sits on the fixture. The planogram is typically manually created.
- [0029] A “**spread**” is defined as group of identical product facings which has spread to encroach the space on the shelf of an adjacent product, where the space allocated to a product is delineated by the placement of the shelf labels.
- [0030] A “**plug**” is defined as a mis-placed product, most likely cause by a customer picking the product and placing it back on the shelf in the wrong spot.
- [0031] A “**fixture**”, as used herein, is defined broadly to include any means of displaying a product, for example, a shelf on which products sit, a peg from which products hang, a pallet sitting on a floor, etc.
- [0032] A “**peg product**” is a product displayed by hanging, usually underneath the price label. The products typically are hanging on a rod, often extending from a pegboard.
- [0033] A “**shelf-ready package**” refers to a box or container, typically a cardboard container in which individual products are shipped, in which the individual products are displayed while in the container by placing the container on the shelf. Often, a portion of the container will be removed to reveal the individual products.
- [0034] A “**ghosted product**” is a product whose image is blurry on the panoramic image.

Detailed Description

- [0035] The present invention is based on the collection of images showing the fixtures of a retail store and the products thereon. Preferably, shelf labels will be visible at some fixed position on the fixtures. Shelf labels define sections of the shelf as being reserved for specific products. Products on the fixtures may be associated with a shelf label and, as such, a determination is able to

be made that a product is in-stock or out-of-stock. The placement of the shelf labels on the fixtures will aid the system of the present invention as it attempts to associate product facings on the fixtures with the shelf labels and to determine when products are out-of-stock. In addition, misplaced products may also be identified based on a comparison of their identity to the shelf label with which they are associated based on their placement on the shelf or peg.

[0036] The images required for analysis of the inventory of the store by the system of the present invention may be collected in any way. For example, the images may be collected manually by photographing sections of the shelves or from stationary or mobile cameras. However, in preferred embodiments of the invention, the images are collected autonomously by a mobile robot which navigates up and down the aisles of the store. In some embodiments, the images are then stitched together to form a panoramic image.

Collection of Images

[0037] The invention is described herein as being based on the analysis of “images” of aisles of products collected by “cameras”. However, as would be realized by one of skill in the art, any representation of an aisle of products could be used. For example, the information required to implement the invention may be obtained from a 3D point cloud or from a planogram. Therefore, the use of the term “image” in the explanation of the invention should be interpreted broadly to include any possible representation. Additionally, the use of the term “camera” should also be interpreted broadly to include any type of sensor used to collect the required information, regardless of whether or not an actual “image” is produced by the sensor.

[0038] An example of such an autonomous robot acting as an inventory monitoring camera system 100 is shown in **FIG. 1** in block form. A camera and optional lighting array 101 are mounted on movable base 102. Movable base 102 may be fitted with drive wheels 104 or may use other forms of locomotion well-known in the robotics field, such as tracks. Movable base 102 is intended to navigate through the aisles of the store to track the status of products on fixtures or other targets 10.

[0039] Movable base 102 can be an autonomous robot having a navigation and object sensing suite 120 that is capable of independently navigating and moving throughout a building, while avoiding obstacles, for example, customers. The autonomous robot preferably has multiple cameras 110 ... 116 attached to movable base 102 by a vertically extending camera support 106. Optional lights 108 are positioned to direct light toward target 10. The object sensing suite may include forward (121), side (122 and 123), top (124) and/or rear (not shown) image and depth sensors to aid in object detection, localization, and navigation. Additional sensors such as laser ranging units 125 and 126 (and respective laser scanning beams 125a and 126a) also form a part of the sensor suite that is useful for accurate distance determination. In certain embodiments, image sensors can be depth sensors that infer depth from stereo images, project an infrared mesh overlay that allows rough determination of object distance in an image, or that infer depth from the time of flight of light reflecting off the target. In other embodiments, simple cameras and various image processing algorithms for identifying object position and location can be used. For selected applications, 3D LIDARs, ultrasonic sensors, radar systems, magnetometers or the like can be used to aid in navigation. In still other embodiments, sensors capable of detecting electromagnetic, light, sound or other location beacons can be useful for precise positioning of the autonomous robot.

[0040] In some embodiments, the depth sensors are associated with image cameras and depth pixels registered to image pixels. This provides depth information for pixels in the image of the shelves. This depth information measures the distances of the image camera to the shelf lip and to the products. In some embodiments, moveable base 102 may also include, either exclusively or in addition to cameras, other types of sensors, for example RADAR, LIDAR, time of flight sensors, etc.

[0041] The camera and depth sensors may produce images rendered in RBD, RGB-D (RGB with depth information), grayscale or black and white. Grayscale may use only one of the R, or G or B channels to make a gray scale-D or a R-D or G-D or B-D. Any other color map transformation may be used, for example, RGB to Z, to make a Z-D map. The camera may render N-channel

images having depth information, For example, RGB + YU + D.

[0042] In alternate embodiments, spatial depth change detection may be used in lieu of absolute depth measurements.

[0043] As seen in **FIG. 1**, various representative camera types useful for constructing an updatable map of product or inventory position may be used. Typically, one or more shelf units, for example, target 10 in **FIG. 1**, would be imaged by a diverse set of camera types, including downwardly (110 and 112) or upwardly (111 and 116) fixed focal length cameras that cover a defined field less than the whole of a target shelf unit, a variable focus camera 115 that adapts its focus to the distance from the imaged target; a wide field camera 113 to provide greater photographic coverage than the fixed focal length cameras; and a narrow field, zoomable telephoto 114 to capture bar codes, product identification numbers, and shelf labels. Alternatively, a high resolution, tilt controllable, height adjustable camera can be used to identify shelf labels. As may be realized, the actual number and type of cameras present in inventory monitoring camera system 100 may vary depending on several factors, including, for instance, the environment in which they are intended to operate.

[0044] To simplify image processing and provide accurate results, the multiple cameras 110 ... 116 are typically positioned a set distance from the shelves during the image collection process. The shelves can be illuminated with LED or other directable lights 108 positioned on or near the cameras. The multiple cameras 110 ... 116 can be linearly mounted in vertical, horizontal, or other suitable orientation on a camera support 106. According to some embodiments, both cameras 110 ... 116 and lights 108 can be movably mounted. For example, hinged, rail, electromagnetic piston, or other suitable actuating mechanisms may be used to programmatically rotate, elevate, depress, oscillate, or laterally or vertically reposition cameras 110 ... 116 or lights 108. In addition, camera support 106 may be movable either horizontally or vertically.

[0045] In some embodiments, to reduce costs, multiple cameras may be fixedly mounted on camera support 106. Such cameras can be arranged to point upward, downward, level, forward or backward with respect to the camera

support and the shelves. This advantageously permits a reduction in glare from products having highly reflective surfaces, because multiple cameras pointed in slightly different directions are more likely to result in at least one image with little or no glare. Angling the camera aids in the avoidance of direct exposure to reflected light. Lights can be mounted along with, or separately from, the sensors, near to or far from the sensors. The lights may be angled forward, backward, upward, downward or level with respect to the light support and the fixtures and can include monochromatic or near monochromatic light sources such as lasers, light emitting diodes (LEDs), or organic light emitting diodes (OLEDs). Broadband light sources may be provided by multiple LEDs of varying wavelength (including infrared or ultraviolet LEDs), halogen lamps or other suitable conventional light sources. Various spectral filters that may include narrowband, wideband, or polarization filters and light shields, lenses, mirrors, reflective surfaces, diffusers, concentrators, or other optics can provide wide light beams for area illumination or tightly focused beams for improved local illumination intensity.

[0046] Electronic control unit 130 contains an autonomous robot sensing and navigation control module 132 that manages robot movements and responses. Electronic control unit 130 may also be provided with communication module 134 which manages data input and output. Robot position localization may utilize external markers and fiducials or may rely solely on localization information provided by robot-mounted sensors. Sensors for position determination may include previously noted imaging, optical, ultrasonic SONAR, RADAR, LIDAR, time of flight, structured light, or other means of measuring distance between the robot and the environment, or incremental distance traveled by the mobile base, using techniques that include but are not limited to triangulation, visual flow, visual odometry wheel odometry and inertial measurements. In preferred embodiments of the invention, the movable base 102 will remain a constant distance from target 10 as movable base 102 traverses the aisles of the store.

[0047] Electronic control unit 130 may also provide image processing using a camera control and data processing module 136. The camera control and data

processing module 136 can include a separate data storage module 138. Data storage model 138 may be, for example, a solid-state hard drive or other form of flash memory. Data storage model 138 is connected to a processing module 140. The communication module 134 is connected to the processing module 140 to transfer product availability and/or identification data or panoramic images to remote locations, including store servers or other supported camera systems, and optionally receive inventory information to aid in product identification and localization. In certain embodiments, data is primarily stored, and images are processed within the autonomous robot. Advantageously, this reduces data transfer requirements, and permits operation even when local or cloud servers are not available. In alternate embodiments, images may be stored and analyzed off-unit on a local server or cloud server.

[0048] The communication module 134 can include connections to either a wired or wireless connect subsystem for interaction with devices such as servers, desktop computers, laptops, tablets, or smart phones. Data and control signals can be received, generated, or transported between varieties of external data sources, including wireless networks, personal area networks, cellular networks, the Internet, or cloud mediated data sources. In addition, sources of local data (e.g. a hard drive, solid state drive, flash memory, or any other suitable memory, including dynamic memory, such as SRAM or DRAM) that can allow for local data storage of user-specified preferences or protocols. In one particular embodiment, multiple communication systems can be provided. For example, a direct Wi-Fi connection (802.11b/g/n/ac/ax) can be used as well as a separate 4G cellular connection.

[0049] Remote servers connectable to inventory monitoring camera system 100 can include, but are not limited to, servers, desktop computers, laptops, tablets, or smart phones. Remote server embodiments may also be implemented in cloud computing environments. Cloud computing may be defined as a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned via virtualization and released with minimal management effort or service provider interaction,

and then scaled accordingly. A cloud model can be composed of various characteristics (e.g., on-demand self-service, broad network access, resource pooling, rapid elasticity, measured service, etc.), service models (e.g., Software as a Service (“SaaS”), Platform as a Service (“PaaS”), Infrastructure as a Service (“IaaS”), and deployment models (e.g., private cloud, community cloud, public cloud, hybrid cloud, etc.).

[0050] In other embodiments the cameras are fixedly mounted to fixtures such as shelves or store infrastructure such as the ceiling. The cameras can optionally be equipped with a motion sensor. The cameras can capture images either continuously, for example at a rate of 10, 15, or 30 frames per second, or intermittently at a set time interval, or when triggered by motion detected by the onboard sensor.

[0051] The camera can further comprise an onboard processor to pre-process the images, for example to detect and blur human faces.

[0052] The camera further comprises a communication module that transmits the images to a local server or to a cloud server.

[0053] **FIG. 2** is an illustration of two inventory monitoring camera systems 100 of the type discussed with respect to **FIG. 1**. Inventory monitoring camera systems 100 are shown inspecting opposite shelves 201 and 202 in an aisle. As shown, each inventory monitoring camera system 100 follows path 205 along the length of an aisle, with multiple cameras capturing images of the shelves 201 and 202.

[0054] In some embodiments, the inventory monitoring camera systems 100 support at least one range finding sensor to measure distance between the multiple cameras and the shelves and products on shelves, with an accuracy of less than 5cm, and with a typical accuracy range between about 5 cm and 1mm. As will be appreciated, LIDAR or other range sensing instruments with similar accuracy can also be used in selected applications. Using absolute location sensors, relative distance measurements to the shelves, triangulation to a known landmark, conventional simultaneous localization and mapping (SLAM) methodologies, or relying on beacons positioned at known locations in a blueprint or a previously built map, the inventory monitoring camera systems 100 can move along a path generally parallel to shelves 201 and 202.

As the movable bases 102 move, vertically positioned cameras are synchronized to simultaneously capture images of the shelves 201 or 202. In certain embodiments, a depth map of the shelves and products is created by measuring distances from the shelf cameras to the shelves and products over the length of the shelving unit using image depth sensors and/or laser ranging instrumentation. The depth map is registered onto the images captured by the shelf cameras, so as the location of each pixel on target can be estimated in 3D.

[0055] As can be seen from **FIG. 1**, each camera is intended to capture a particular vertical portion of the shelf fixture as movable base 102 traverses the aisle. The vertical portion of the shelf fixture maybe captured as one long panoramic image as movable base 102 continuously moves along the aisle, or, alternatively, the vertical portion of the shelf may be captured as a single image or as multiple vertical, overlapping images which may be obtained, for example, if the robot moves a certain distance and then stops to allow imaging of the portion of the shelves currently in front of the cameras. Alternatively, the robot may continuously capture images without stopping.

[0056] For each section of the shelf fixture, multiple images may be captured at varying focal lengths, such as to increase the likelihood of obtaining clear images of products at differing depths from the edge of the shelf. Images from cameras 110 ... 116 may be horizontally and/or vertically stitched together to form a panoramic image needed for analysis of the product status. Using available information, for example, the location of each pixel on target images, consecutive images can be stitched together to create panoramic images that span an entire shelving unit along the entire length of the aisle. The consecutive images can be first stitched vertically among all the cameras, and then horizontally and incrementally stitched with each new consecutive set of vertical images as the inventory monitoring camera systems 100 move along an aisle. If multiple images have been captured for a given section of the aisle, the best image may be selected for inclusion in the stitched-together panoramic image. In this case, the best image may be an image having better focus than other images of the same section of shelf, or, for example, may be an image lacking lighting artifacts or reflections.

[0057] Image processing to create or improve panoramic image construction can include keypoint detection to find correspondences between overlapping images, image registration using matching features or direct alignment, image calibration to compensate for differing lens and camera combinations, distortion, exposure, or chromatic aberration corrections, and image blending and compositing. Various map projections can be used for arranging stitched images, including rectilinear, cylindrical, equiangular, stereographic, or spherical projection

An example of an image 300 is shown in **FIG. 3**. In preferred embodiments of the invention, one image 300 will be provided for each aisle in the store. The image 300 for each aisle may or may not include end caps. In certain embodiments, images of end caps may be provided as separate images. In preferred embodiments, image 300 is a panoramic image.

Processing Pipeline

[0058] The images 300 collected by inventory monitoring camera system 100 are processed by a processing pipeline which comprises a combination of deep learning detectors and classifiers, as well as logic, to extract the required information from the images. The goal of the pipeline is to detect and flag: (1) out-of-stock items (including products displayed in shelf-ready packages and peg products); (2) plugs; and (3) spreads. It is a further goal of the pipeline to determine shelf label location and content and to identify individual product facings. Additional, optional, goals may include, without limitation: a comparison of product locations to a planogram, classification of the fixture type (e.g., shelves, pegs, etc.), identification of constituent parts of the fixture such as the side counter, and caps, side caps, side stacks, etc., determining the beginning and end of each section of a shelf, and determining the state of the fixture such as broken shelves.

[0059] **FIG. 4** shows one example of a processing pipeline 400 which may be used to process images 300 in accordance with the present invention. The process starts with image 300 of the form shown in **FIG. 3**. Preferably image 300 will comprise an image stitched together from individual, high resolution images captured by an inventory monitoring camera system 100 as shown in **FIG. 1** and as described above. In some embodiments, image 300 may be of lesser

resolution than the original images captured by inventory monitoring camera system 100.

[0060] Product detector 402, shown as a component of pipeline diagram 400 in **FIG. 4**, is a detector which operates on the panoramic image to detect products and place bounding boxes around the products. In current embodiments of the invention, several different types of products that are detected. These include (1) peg products, which are products are not actually on a shelf, but are hanging from pegs; (2) grill products, which are products which may be placed in a bin having as a front surface a metal grating through which products may be observed; (3) shelf products, which are products placed on the shelves ; (4) shelf-ready packages, which are boxes containing the products typically comprising a cardboard box having a top portion that has been removed to transform the cardboard box into a tray. The tray is then placed directly on the shelf; (5) shelf lip products which are products on a shelf equipped with a transparent, translucent, or opaque lip that partially occludes the products or modifies its view; (6) unpackaged products; (7) stacked products; (8) flat products; (9) deformable or bagged products; (10) binned products, which are products that are contained in a bin, that have to be typically reached from the top opening of the bin, and that can typically be viewed from the top; (11) caged products, which are products that are contained in a wireframe cage; (12) pushed products, which are products pushed to the front of the shelf by a pushing fixture; and (13) guarded products, which are products guarded on each side by railing that ensures consistent alignment of the product with its corresponding shelf label.

[0061] **FIG. 5** shows an example of a portion of the image 300 processed by product detector 402, showing bounding boxes 504 surrounding various products 502. It is contemplated that, in other embodiments of the invention, other types of products may be detected, either by using product detector 402 or by using another detector.

[0062] Product detector 402 produces, as an output, the image with a bounding box as shown in **FIG. 5**. In some embodiments, the bounding box is represented in a data structure as a tuple of data of the form $BB = \{x, y, w, h\}$. A tuple may comprise, for example, the x, y coordinates of a corner of the bounding

box as well as the width (w) and the height (h) of the bounding box. Other information may be included in the tuple, for example, depth information. In some embodiments, the tuple may represent bounding boxes within a 3D point cloud. The x,y coordinates of the bounding box may be relative to a coordinate system imposed on the image, having an arbitrary point of origin. In some embodiments of the invention, the incremental units within the coordinate system may be pixels.

[0063] In preferred embodiments, product detector 402 is a machine learning model trained on images of products. Any commonly-known architecture for the machine learning model may be used.

[0064] The number of available images for training product detector 402 is less than the typical number of images required to train a deep neural network. Further, each image, especially if it is panoramic image, may be too large to fit on a single GPU. The solution adopted for product detector 402 is random cropping of images with fixed window size so that each generated training batch is unique. This operation creates big variations from limited data, which allows the detector to generalize well to unseen images. An example of a cropped panoramic image is shown in **FIG. 5**.

[0065] Once products have been identified by product detector 402, the products are classified into one of the various types of products discussed above by product classifier 404. In certain embodiments, only a subset of the product types may be detected. **FIG. 6** shows a two-tiered classifier architecture taking as an input an image of a product 602. Product image 602 may be extracted from the image 300 using the tuples defining the bounding boxes around the products as determined by product detector 402.

[0066] **FIG. 6** shows the architecture of the product classifier 404. A first classifier 604 classifies products into one of the selected categories mentioned above, for example, identifying whether the incoming product image is one of: a peg product, a grill product, a shelf product or a shelf-ready package. If the output of the first model identifies the input image 602 as containing a shelf-ready package, then the input image 602 is forwarded to a second classifier 606, which checks to see if the shelf-ready package identified by classifier 604 is empty or not empty. Classifier 606 outputs a "0" if the shelf-ready package is

empty, and “1” if the shelf-ready package is not empty.

[0067] In certain embodiments, products may be identified as peg products by other means. For example, a product may be determined to be a peg product if the product does not lie above or below a shelf, or, if the product lies in an area associated with a peg label.

[0068] Shelf labels are detected in a similar manner using label detector 406, shown as a component of pipeline 400 in **FIG. 4**. Shelves in stores typically have two kinds of labels. The first type of label is a price label for the products, which are referred to herein as “shelf labels”. An example of a shelf label is shown in **View (A)** of **FIG. 7**. The second type of label is a section label which marks the end of a section of shelving. An example of a section label is shown in **View (B)** of **FIG. 7**. Label detector 406 does not differentiate between the two types of labels but instead detects both types in a single forward pass on the input image 300. In some embodiments, label detector 406 may also detect promotional materials 806 placed along shelf edges. In preferred embodiments, label detector 406 is trained using images of labels. Label detector 406 will also output a tuple describing the x,y location of the label bounding box as well as the width (w) and height (h) of the bounding box. **View (C)** of **FIG. 7** shows bounding box 704 placed around shelf label 702.

[0069] Because section labels and promotional materials are not typically associated with products, they must be removed from the pipeline. This is accomplished by training a classifier 408 to distinguish between shelf labels and section labels and promotional materials. The classifier takes as an input a label image, which can be cropped from the image 300 using the bounding box coordinates generated by label detector 406 and classifies it as a shelf label or not a shelf label. Those labels which are not shelf labels are then ignored for the remainder of pipeline 400.

[0070] A “ghosted product” is a product whose image is out of focus or blurry on the image. As such, the product may not be able to be detected by the product detector 402. The image of the product may be blurry because of one or more of several possible reasons. First, not all products may be at the same depth from the camera. For instance, the camera may be focused to take images at

the edge or near edge of each shelf, while the actual products are sitting further back on the shelf because products near the edge of the shelf have been removed by customers. Additionally, products that are displayed on pegs are often at a depth different from the products on the shelves. That is, the pegs are not as long as the shelves and as such, the products may be farther away from the camera. As such, the image of the product in image 300 may be blurry. In addition, the stitching process which creates image 300 from the horizontal and vertical stitching of individual images of the shelves may leave some products blurry as an artifact of the of the process, especially where the products may appear near the edges of each image being stitched together.

[0071] Box 410 of the processing pipeline 400 shown in **FIG. 4** is a special detector that has been trained to detect ghosted products and to enclose them in a bounding box. As with the product detector in box 402, the ghosted products detector 410 will output a tuple describing the bounding box, having x,y coordinates as well as the width (w) and the height (h) of the bounding box. Ghosted products detector 410 is preferably trained using ghosted data from peg regions of the store aisles. At box 412 of processing pipeline 400, the detected products are combined with the ghosted products to produce a list of the location of all bounding boxes representing products within image 300. The combined listing of products and ghosted products, in some embodiments, may simply be a list of all tuples produced by the product detector 402 and the ghosted product detector 410.

[0072] As part of the process of identifying pegged products, is necessary to identify where shelves are located on image 300. At box 414 of pipeline 400, the image 300 is processed by a classifier 414 that classifies each pixel of the image 300 determine if the pixel is part of a shelf or not part of a shelf, to produce a binary mask, having pixels located on shelves flagged as a binary "1" in pixels not located on shelves flagged as a binary "0". This results in a binary mask, an example of which shown in **FIG. 8**. A smoothing operation may be applied to the binary mask to smooth the edges of the shelves and merge any breaks in the shelf. The mask in **FIG. 8** is the result of processing the image of **FIG. 3** with the shelf segment classifier 414.

- [0073] The binary mask showing the location of the shelves may be used to determine which of the shelf labels identified by shelf/section label classifier 408 are shelf labels representing product sitting on a shelf or are peg labels representing products hanging from a peg. It is assumed that if a shelf label has a location which overlaps the areas of the binary mask showing the locations of the shelves, then the shelf labels associated with the product sitting on a shelf. Likewise, if the shelf label has a location which does not overlap the areas of the binary mastering location of the shelves in the shelf labels assumed to be associated with the product hanging from a peg. Peg shelf label classifier 416 makes this determination.
- [0074] In alternate embodiments of the invention, shelves may also be localized by *inferring* the location of the shelves from the location of the shelf labels, in box 418 of pipeline 400. It is assumed that if shelf labels are aligned in a horizontal line, as specified by their *x,y* coordinates (discovered by shelf label detector 406), then the shelf labels all lie on a shelf. As such, the presence and dimensions of the shelf can be inferred from the alignment of the shelf labels. In some embodiments of the invention, the output of the shelf segment classifier 414 may also be an input to shelf inference 418. Once it is determined where the shelf is located, a shelf tuple is created defining a bounding box for the shelf. An example of bounding boxes for shelves is shown in **FIG. 9**.
- [0075] Once a location of a shelf is inferred, it is also possible to determine which product bounding boxes, discovered by product detector 402, are positioned on the shelf by comparing the location of the bottom of the products bounding box with the location of the top of the shelves bounding box. The output of shelf inference 418 is a shelf object comprising the location of the shelf, all shelf label tuples associated with the shelf and all product tuples for product bounding boxes located on the shelf.
- [0076] In the image 300, there may be some shelves which are thicker in the vertical direction than other shelves. Such shelves, referred to herein as “stacked shelves”, may have two rows of labels as opposed to one row. In this case, the top row of labels are for products above the shelf and the bottom row of labels is for products below the shelf, which, in some instances, may be

sitting on the floor of the aisle. An example of a stacked shelf is shown by reference number 1002 in **FIG. 10**. The stacked shelves are also inferred in box 420 of pipeline 400.

[0077] The next step in the pipeline is to associate products on the shelves with the respective labels on the shelf edges by product/shelf label association at box 422 of pipeline 400. This is a crucial prerequisite for the detection of out-of-stock products, spreads and plugs. In this step, the products are associated to their respective labels. All products between two neighboring labels, or between a label and shelf end, will be associated to the label on the left (in a left justified configuration).

[0078] On each shelf, a section is marked between the starting coordinates of two neighboring labels along the x-axis. These are referred to as section start points and section end points respectively. In a configuration where the products are left justified with the labels, a “section” would be defined as the area between the left edge of a label and the left edge of the next label to the right. All products falling within this section are associated with the shelf label at the far left of the section. As may be realized by one of skill in the art, in store configurations where the labels are right or center justified, the definitions of the sections and, as a result, the method of determining the product/label associations would be similar, but slightly different. For example, any vertical coordinate of the shelf label may be used to define the section, (area associated with the shelf label) and may include an offset, which may be different for adjacent shelf labels, or which may be 0. An example of a left justified section is shown in **FIG. 11** as area between lines 1102 and 1104.

[0079] An “overlap ratio” is computed for every product within the selected section. If the overlap ratio is above some predefined threshold, then the product gets associated with the label in the selected section. This is illustrated in **FIG. 11** wherein the section under consideration is between the line labeled 1102 and the line labeled 1104. The overlap ratio may be given by the following formula:

$$\text{overlap ratio} = \frac{\text{distance between section start point and product end point}}{\text{product width}}$$

[0080] In preferred embodiments of the invention, the predefined threshold may be 50%. As can be seen in **FIG. 11**, product 1106 has been associated with label 1108 because more than 50% of the product lies within the section defined by lines 1102 and 1104. Once all of the products within the selected section are associated, the next section is selected by moving between the next pair of labels on the shelf. The result of the product/shelf label association is shown in **FIG. 10** showing which products are associated with which labels. Note that products having shelf ready packaging are detected as a single product, as shown by reference number 1004, while products stacked individually on the shelf are detected as individual products as shown by reference number 1006. Also note that products stacked below a shelf which has been determined to be a “stacked shelf” are associated with labels on the shelf above, as shown by reference number 1008. Logically, the output of box 422 is a dictionary in which tuples representing shelf labels are associated with tuples representing the products within the label section. It should be noted that in an out-of-stock situation, the tuple representing the shelf label will not have any product tuples associated therewith.

[0081] Box 424 of pipeline 400 creates the association between labels which have been classified as peg labels and the products associated therewith. This is done simply by associating any products directly below the peg label with the peg label. In addition, products left or right adjacent to the peg label not having a peg label immediately above may be associated as well. The peg label tuples and their associated product tuples are then added to the dictionary created by the product/shelf label association in box 422. This may be accomplished, for example, by determining that the centerline of a bounding box defining the product lies within the horizontal bounds of the bounding box for the peg label. Other criteria may be used to make this determination.

[0082] In box 426 of pipeline 400, those shelf labels which are associated with empty shelf-ready packages are flagged. The empty shelf-ready packages were discovered as part of the two-step classification in **FIG. 6** wherein classifier 604 classifies a product as a shelf-ready package and classifier 606 classifies the shelf-ready package as being empty or not empty. An empty shelf-ready package does not necessarily imply that the product normally present in the shelf-ready package is out-of-stock. Often, an empty shelf-ready package at the edge of a shelf will have full shelf-ready packages

behind the empty shelf-ready package. However, the system may not be able to detect this situation or may interpret the full shelf-ready packages as ghosted products, based on their depth on the shelf. As a result, the empty shelf-ready packages are flagged, as attention is still required by the staff of the store to remove the empty shelf-ready package and reposition the full shelf-ready packages to the edge of the shelf. At box 428 of pipeline 400, this situation is reported to the store.

- [0083]** At box 430 in pipeline 400, it is determined which shelf products are out-of-stock. This happens by consulting the dictionary of shelf label tuples and associated product tuples and determining which shelf label tuples have no associated product tuples. That is, which shelf labels have no products associated therewith. These shelf labels are extracted from the dictionary and placed in a separate out-of-stock list for further processing and eventual reporting to store.
- [0084]** In a manner similar to box 422, at box 436, it is determined if products which have been classified as peg products are out-of-stock. In box 424, the peg labels were associated with peg products. In box 436, those peg label tuples in the dictionary having no associated product tuples (i.e., no products positioned directly under the peg label) are extracted from the dictionary and added to the out-of-stock list.
- [0085]** At box 432 of pipeline 400 a special situation is handled in which a portion of the image 300 is blocked out. This could happen for instance, where the robot is traversing the aisle and comes upon an object (e.g., a person or shopping cart) next to the shelf. In such instances, the robot will navigate around the object but no images of the shelf behind the object are able to be collected. As such, in the image 300, this area of the shelf will appear as an occluded area showing black pixels. This can lead to false reporting of out-of-stock items, as in the situation wherein a label may be visible in image 300, but the section associated with that label is partially within the occluded area of the shelf. In such situations, it may be preferable to ignore the shelf label during the current pass of the robot as opposed to falsely flagging the product as being out-of-stock. As such, in certain embodiments of the invention, the shelf labels found to be in this situation may be removed from the out-of-stock list. At box 438, a similar process detects occluded areas with respect to product labels which have been classified as peg product labels.
- [0086]** At box 434 of pipeline 400 another special situation is handled. In this situation, the stitching process may create an artifact wherein the shelf appears twice within image

300, with one image of the shelf being directly below the other image of the shelf. In such situations, the products on the shelf will be associated with the shelf labels in the top image of the shelf. As such, the shelf labels in the bottom image of the shelf are ignored. At box 440, similar process detects duplicated peg product labels.

[0087] The depth information from the depth sensors can be used in two ways. The first way is to complement the out-of-stock pipeline, by confirming the absence of a product. This is done by measuring the distance between the shelf lip and the product above it, and if that distance is equal to the distance to the back of the shelf, the product is determined to be absent. This information can be combined to the out-of-stock logic in the pipeline to avoid reporting out of stocks in cases when the product detector would have not detected an existing product.

[0088] The second way that the depth information can be used is to way to create N-channel images, for example, RGB-D, by adding depth information D. For example, RGB + YU + D images may be created. CNNs, other types of neural networks, machine learning classification or AI algorithms may then be trained on the N-channel images to capture the 3D and other features in addition to the conventional 2D features. Using the N-channel images, out-of-stock products, plugs and spreads detection substantially follows the same described pipeline except that all images are N-channel instead of just RGB.

[0089] Matching pipeline 450 of pipeline 400 is used in the detection of plugs and spreads.

FIG. 13 shows a plug situation. Products 1304 are associated with label 1302.

However, product 1306 is in the shelf section that should only be populated by products 1304. Product 1306 is therefore a plug.

[0090] **FIG. 15** shows a spread situation. Products 1504 are associated with label 1502.

However, products 1506 are identical products 1504 and, as such, should occupy the section of the shelf between label 1502 and 1508. Instead, product 1506 are infringing on the section of the shelf associated with label 1508.

[0091] To detect plugs and spreads, it is necessary to determine if one product on the shelf matches another product. For example, in **FIG. 13**, if products 1304 do not match product 1306, and product 1306 is in the section of the shelf allocated to label 1302, then a plug has been detected. Likewise, in **FIG. 15**, if products 1504 match products 1506, then a spread has been detected.

[0092] In one embodiment of the invention, a pairwise matching process is undertaken to

determine if products next to each other on shelves match each other. In this embodiment, the purpose of matching pipeline 450 is to determine if one product matches an adjacent product. A more detailed version of the matching pipeline 450 is shown in **FIG. 12**. Two preliminary checks are first performed to identify products that match or do not match by doing a size check and a color check.

[0093] In box 1202 a size check is performed. In this check, each product on the shelf is checked with the product to its immediate right (in a left justified configuration) to determine if the products are the same size. If any difference in the size of each product in the pair of products falls within a certain predefined range, then it is determined that the products are the same size and as such may be possibly be the same product. For example, if the predefined range is five pixels, then if the size of the products falls within five pixels of each other it is determined that they are of the same size. If the size check falls outside of the predefined range, then it is determined that the products are not the same size and, as such, are definitely different products. In this case, no further processing is performed by matching pipeline 450 for this pair of products.

[0094] In some embodiments, the range is a function of the camera's distance to the product. This distance can be measured by depth sensors coupled to the cameras, wherein each depth pixel is registered to a pixel in the color camera. If the products are measured to be at a different distance from the camera, then the range is adjusted accordingly.

[0095] If it is determined that the products are the same size, then an analysis of the color distribution of the product is performed in box 1204. The analysis of color distribution could be performed in one of several ways. For instance, in one embodiment, the average color of all the pixels may be ascertained to determine a match. In other embodiments, patch-wise matching may be performed. In yet other embodiments, a histogram of the color distribution may be obtained and compared with the histogram of the other product in the pair. If it is determined that the color distribution of the products does not match, then it is determined that the products are different, and the processing of that pair of products in the matching pipeline 450 ends.

[0096] If it is determined that both the size and color distribution of adjacent products indicate a match, as determined by boxes 1202 and 1204 respectively, the pair of products is next sent to deep feature matching 1205. In one embodiment of the

invention, deep feature matching takes images of two products and feeds them to a deep learning CNN, which determines if the images match or do not match. In another embodiment of the invention, features may be extracted from the images and feature-wise matching may be done by a deep learning CNN. In another embodiment of the invention an additional deep learning CNN may be used to perform optical character recognition (OCR) of any writing on the front of the products to determine if the products match.

[0097] Auto encoder 1206 can involve use of deep models where deep features are learned from the images and matched. In auto encoder 1206, embeddings for each of the images are learned and followed with training a pair-wise deep classifier 1208 on the autoencoder features. The pair-wise classifier 1208 provides a decision of “1” if the pair of images match and “0” if they don’t.

[0098] In alternate embodiments of the invention, deep learning neural network classifiers may be used to directly identify the product from an image of the product. The deep learning neural network classifiers may operate on images of the products extracted from image 300 or may operate on higher resolution images originally captured by inventory monitoring camera system 100 and used to form image 300.

[0099] The output of the matching pipeline is then sent to spread logic 1210, which is used to detect spreads, and plug logic 1212, which is used to detect plugs, as described below.

[0100] **FIG. 14** shows a situation in which a plug has been detected. Two adjacent products are determined not to match, as shown by the red line between the products. Specifically, is determined at product 1402 and 1404 do not match each other and that product 1406 and 1408 also do not match each other. Because products 1402 and 1404 both reside in the shelf section associated with label 1410, a plug situation has been detected as indicated by the red box around products 1402 and 1404. Although products 1406 and 1408 also do not match, because product 1406 is within the shelf section associated with label 1412 and product 1408 is within the shelf section associated with label 1414, no plug situation is indicated.

[0101] **FIG. 16** shows a situation in which a spread has been detected. It is been determined that product 1602 and 1604 match each other and, as such, should be in the same section of the shelf associated with shelf label 1606. However, product 1604 is mostly within the shelf section associated with shelf label 1608 and, as such, is associated with shelf label 1608 in the dictionary. Because product 1604 should have

been associated with shelf label 1606, it is determined that a spread has been detected, as indicated by the yellow box in **FIG. 16**.

[0102] Also, it should be noted that if no other products are associated with shelf label 1608, shelf label 1608 may be flagged as an out of stock product.

[0103] In another embodiment of the invention, plugs and spreads can be detected by positively identifying each product on the shelf and determining if it is in the correct section of the shelf. In this embodiment, a deep learning classifier is trained to take as input images of products and output the identity of the product.

[0104] Once it has been determined that a product is out-of-stock, that is, there is a shelf label or peg label having no associated products, it is necessary to identify those products to the store. It should be noted that the out-of-stock list only knows that a particular shelf or peg label has no products associated with it. There is no knowledge at this point of the identity of the product referred to by the information on the shelf or peg label. To determine the identity of the product, the system uses a mapping between the shelf and peg label positions on the image and the shelf and peg labels in the original high-resolution images captured by the inventory monitoring camera system 100. The high-resolution images are those images which were stitched together to create image 300 of the shelf. In the high-resolution images, the system is able to read the content of the shelf and peg labels, for example, bar codes or text and is thus able to identify the product referred to by the contents of the shelf or peg label. The identity of the out-of-stock items can thus be identified to the store.

[0105] In box 442 of **FIG. 4**, visualization on a web page may be provided to indicate the location and identity of out-of-stock items, plugs and spreads.

[0106] Many modifications and other embodiments of the invention will come to the mind of one skilled in the art having the benefit of the teachings presented in the foregoing descriptions and the associated drawings. Therefore, it is understood that the invention is not to be limited to the specific embodiments disclosed, and that modifications and embodiments are intended to be included within the scope of the appended claims. It is also understood that other embodiments of this invention may be practiced in the absence of an element/step not specifically disclosed herein.

Claims:

1. A computer-implemented method comprising:
 - obtaining a representation of an aisle containing stocked products;
 - identifying a product in the representation;
 - identifying a product label in the representation; and
 - associating the product with the product label.

2. The method of claim 1, wherein associating the product with the product label comprises:
 - identifying an area of the representation of the aisle associated with the product label; and
 - determining that the product lies within the area associated with the product label.

3. The method of claim 2, the representation comprising an image of the aisle.

4. The method of claim 3, further comprising:
 - establishing a coordinate system superimposed on the representation of the aisle;
 - determining the positions of the product label and the product with respect to the coordinate system;
 - determining the area associated with the product label with respect to the coordinate system; and
 - associating the product with the product label when the position of the product falls within the area associated with the product label.

5. The method of claim 4, the product labels and products being identified by tuples, the tuples comprising at least the coordinates within the coordinate system of a bounding box placed around the product label or product within the image, and the width and height of the bounding box.

6. The method of claim 5 wherein identifying a product label in the representation of the aisle comprises:

using a trained label detection model to detect images of labels within the representation of the aisle, the label detection model outputting the tuple describing the label;

submitting the label to a label classifier to determine if the label is a product label or another type of label; and

adding the tuple to a list of product labels if the image of the label is classified as a product label.

7. The method of claim 6, wherein identifying a product in the representation of the aisle comprises:

using a trained product detection model to detect products within the representation of the aisle, the product detection model outputting the tuple describing the product;

determining the type of the product; and

adding the tuple and a type of the product to a list of products.

8. The method of claim 7 wherein the type of the product is determined by submitting the product to a product classifier.

9. The method of claim 7, further comprising determining positions of shelves within the representation of the aisle.

10. The method of claim 9, wherein determining positions of shelves comprises:

using a trained shelf detection model to detect shelves within the representation of the aisle, the shelf detection model outputting a tuple describing a bounding box containing the shelf.

11. The method of claim 9, wherein determining positions of shelves comprises:

inferring the positions of the shelves by determining that a plurality of product labels are horizontally aligned.

12. The method of claim 9, further comprising:

determining that one or more product labels are shelf labels by determining

that the bounding boxes for the product labels overlap a bounding box for a shelf;
and

determining that one or more product labels are peg labels by determining that the bounding box for the product labels do not overlap a bounding box for a shelf.

13. The method of claim 12 wherein defining the area associated with the peg label comprises:

defining an area directly below the bounding box for the peg label and areas adjacent to the area directly below the bounding box for the peg label having no associated peg labels as the area associated with the peg label.

14. The method of claim 13 wherein determining that the product lies within the area associated with the peg label comprises:

defining that the bounding box for the product overlaps at least partially with the area associated with the peg label.

15. The method of claim 14, further comprising:

defining the area associated with a shelf label as an area on a shelf defined by the shelf label and an adjacent shelf label.

16. The method of claim 15 wherein the area associated with a shelf label is defined as an area above or below a shelf between an offset from a vertical coordinate for the shelf label and an offset from a corresponding vertical coordinate for an adjacent shelf label.

17. The method of claim 16, further comprising determining that a product lies within the area associated with the shelf label by

determining that the width of the bounding box for the product overlaps the area associated with the shelf label by a predetermined percentage.

18. The method of claim 17 wherein determining the overlap between the width of the bounding box for the product and the area associated with a shelf label

comprises dividing the distance between the leftmost coordinate of the area associated with the shelf label and the rightmost coordinate of the product bounding box by the overall width of the product bounding box.

19. The method of claim 15, further comprising:
 - determining that two products match or do not match.

20. The method of claim 19, further comprising:
 - determining that two adjacent products on the same shelf do not match;
 - determining that the two adjacent products are within the area associated with the same shelf label; and
 - flagging the detection of a plug.

21. The method of claim 19, further comprising:
 - determining that two adjacent products on the same shelf match;
 - determining that the two adjacent products are not within the area associated with the same shelf label; and
 - flagging the detection of a spread.

22. The method of claim 19 wherein determining that two products match or do not match comprises:
 - determining if sizes of the bounding boxes associated with the adjacent products match within a predetermined threshold and if not, determining that the adjacent products do not match;
 - determining that color distributions of the products within the bounding boxes match within a predetermined threshold and, if not, determining that the adjacent products do not match; and
 - submitting adjacent products to a trained model to perform deep feature matching between the products, the trained model providing an indication that the products match or do not match.

23. The method of claim 19 wherein determining that two products match or do not match comprises:

submitting each product to a product identification classifier;
receiving as output from the product identification classifier an identification
of each product; and
determining that the identifications of the products match or do not match.

24. The method of claim 3 wherein obtaining an image of an aisle comprises:
capturing images of sections of the aisle; and
stitching images for two or more sections together to create a stitched image.

25. The method of claim 24, the stitched image being a panoramic image showing
the entirety of the aisle.

1/16




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5

6/16

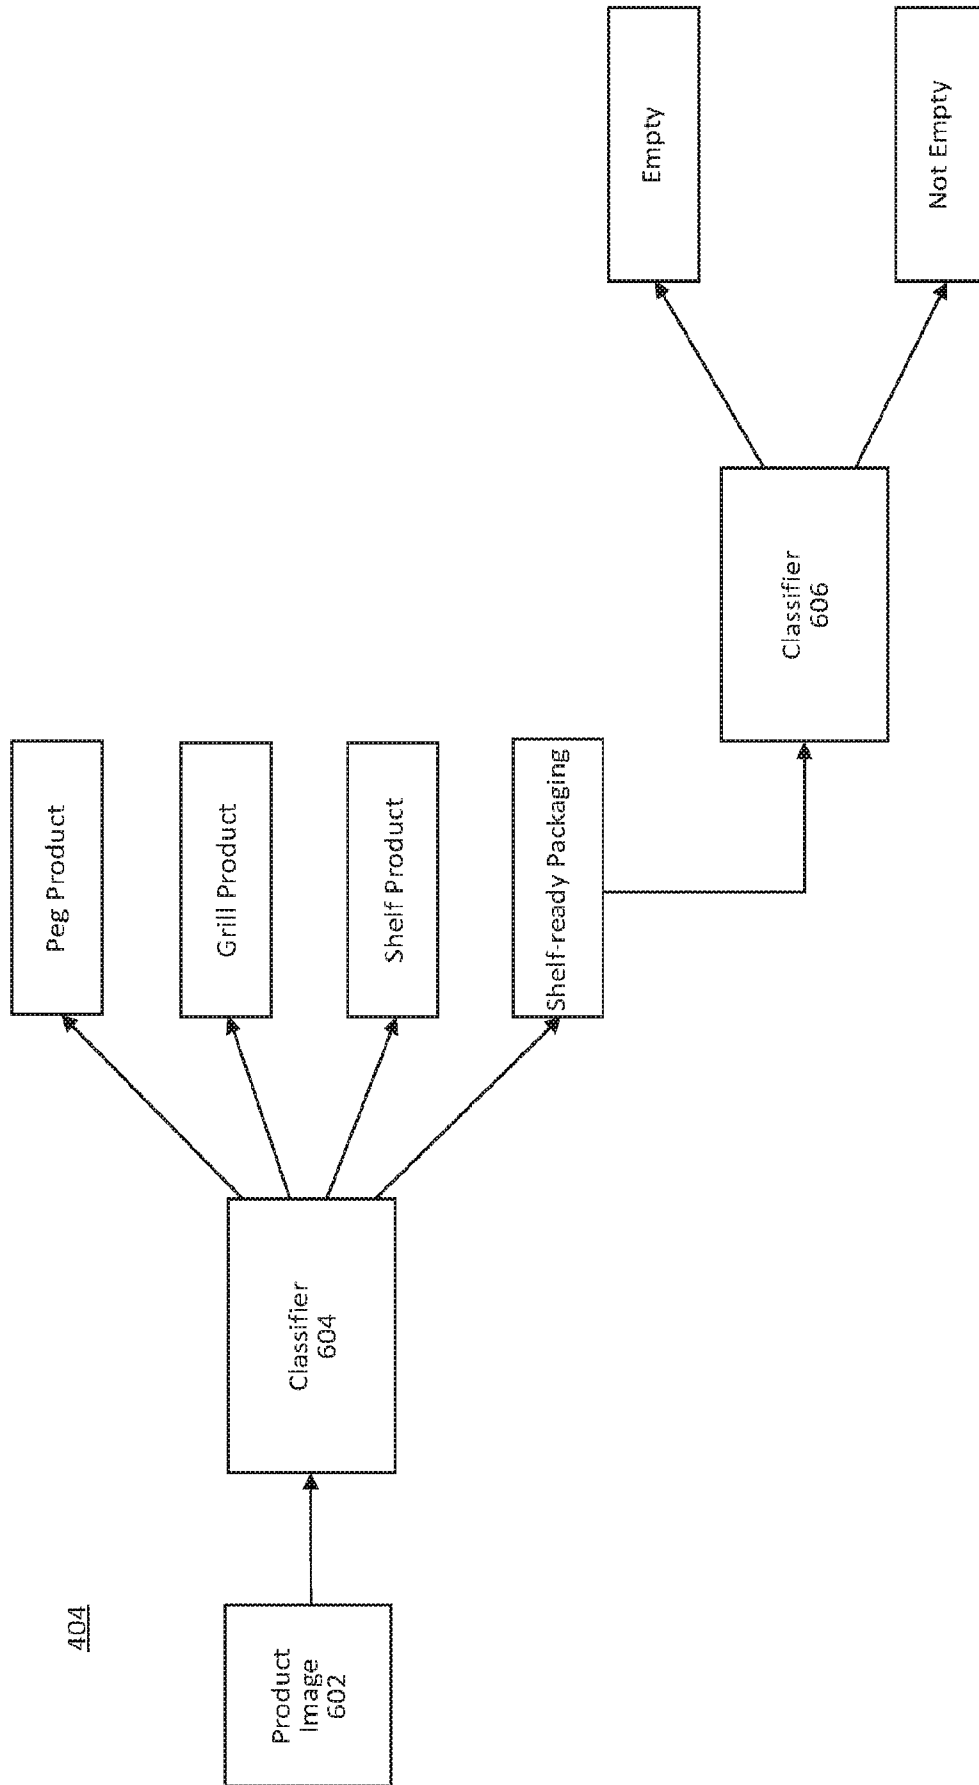




FIG. 6

7/16



(A)



(B)



706

702

704

(C)

FIG. 7




FIG. 8

9/16



FIG. 9




FIG. 10

11/16




FIG. 11

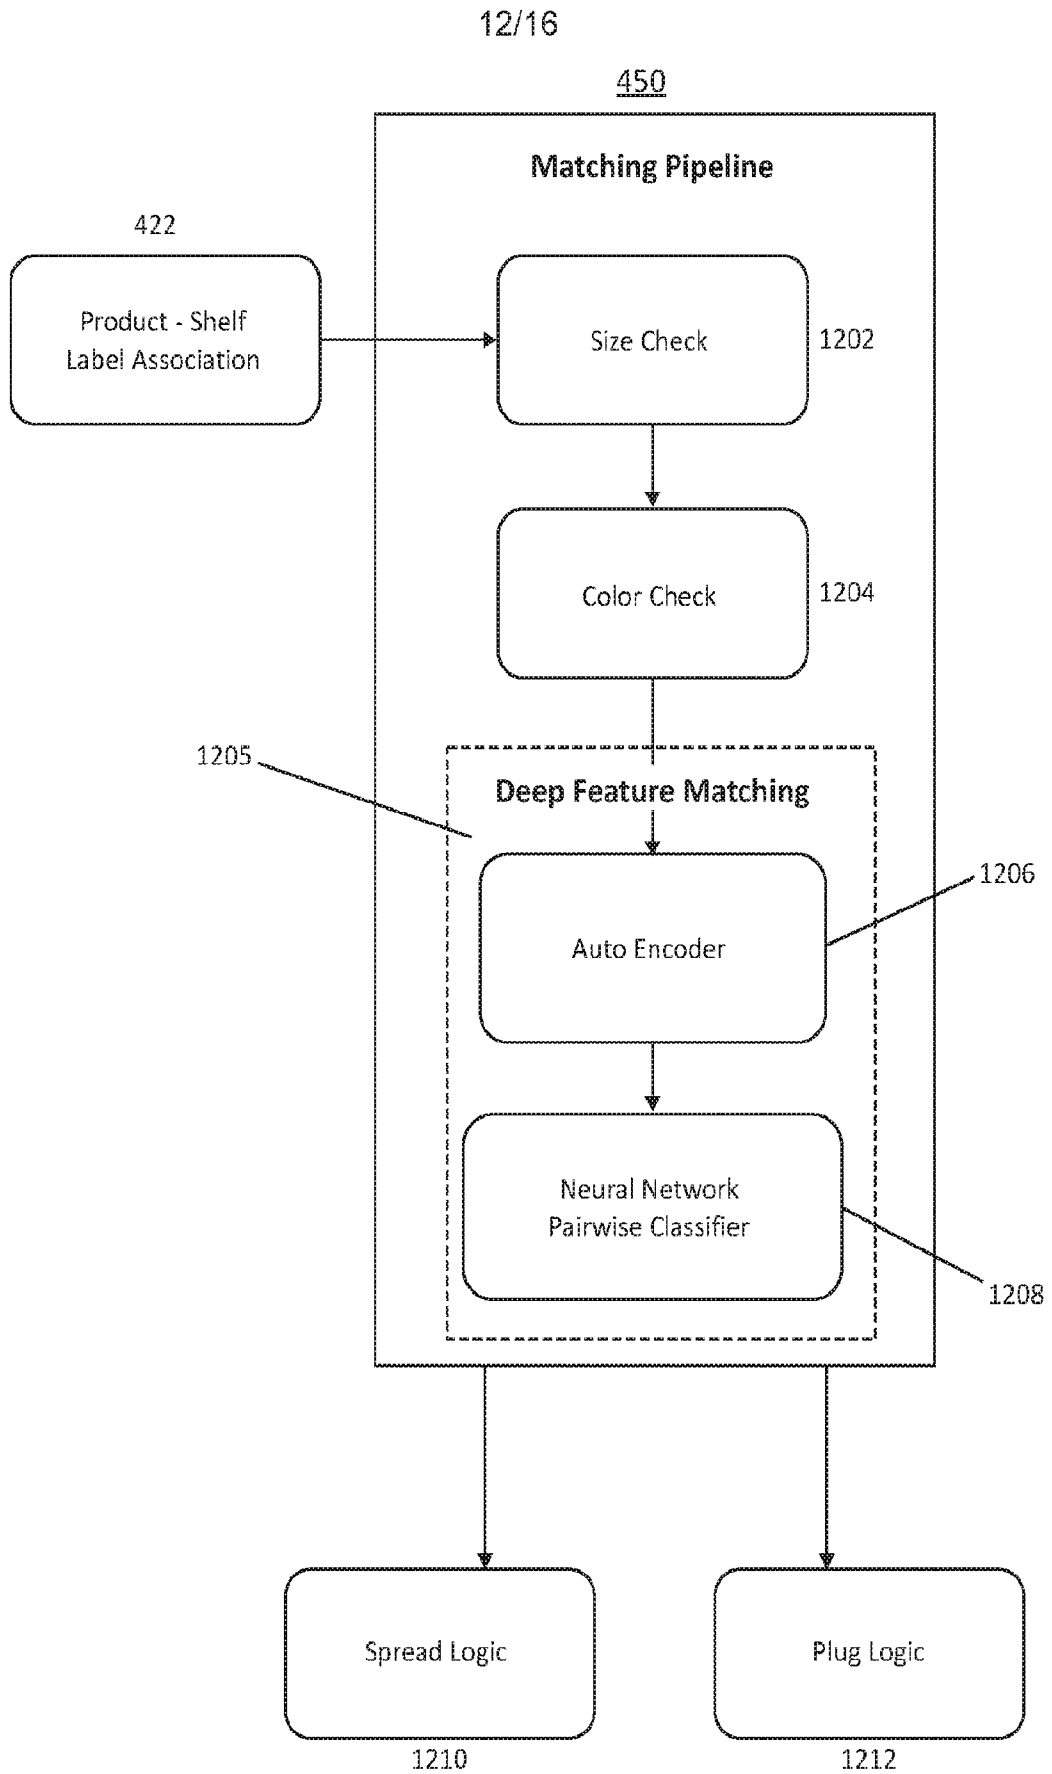


FIG. 12




FIG. 13

14/16




FIG. 14




FIG. 15




FIG. 16




FIG. 2