

**(12) STANDARD PATENT**  
**(19) AUSTRALIAN PATENT OFFICE**

(11) Application No. **AU 2017357645 B2**

(54) Title  
**A robotic fruit picking system**

(51) International Patent Classification(s)  
**A01D 46/30** (2006.01)

(21) Application No: **2017357645**

(22) Date of Filing: **2017.11.08**

(87) WIPO No: **WO18/087546**

(30) Priority Data

(31) Number  
**1618809.6**

(32) Date  
**2016.11.08**

(33) Country  
**GB**

(43) Publication Date: **2018.05.17**

(44) Accepted Journal Date: **2022.11.10**

(71) Applicant(s)  
**Dogtooth Technologies Limited**

(72) Inventor(s)  
**Robertson, Duncan;Cook, Matthew;Herbert, Edward;Tully, Frank**

(74) Agent / Attorney  
**Davies Collison Cave Pty Ltd, Level 15 1 Nicholson Street, MELBOURNE, VIC, 3000, AU**

(56) Related Art  
**US 2006/0213167 A1**  
**US 9468152 B1**  
**JP 2004180554 A**  
**CN 103503639 A**  
**KR 20150105661 A**  
**US 4519193 A**



- (51) International Patent Classification:  
*A01D 46/30* (2006.01)
- (21) International Application Number:  
PCT/GB2017/053367
- (22) International Filing Date:  
08 November 2017 (08.11.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
1618809.6 08 November 2016 (08.11.2016) GB
- (71) Applicant: **DOGTOOTH TECHNOLOGIES LIMITED**  
[GB/GB]; Millside, The Moor, Melbourn, Royston SG8  
6ED (GB).
- (72) Inventors: **ROBERTSON, Duncan**; c/o Dogtooth Technologies Limited, Millside, The Moor, Melbourn, Royston  
SG8 6ED (GB). **COOK, Matthew**; c/o Dogtooth Technologies Limited, Millside, The Moor, Melbourn, Royston SG8  
6ED (GB). **HERBERT, Edward**; c/o Dogtooth Technologies Limited, Millside, The Moor, Melbourn, Royston SG8  
6ED (GB). **TULLY, Frank**; c/o Dogtooth Technologies Limited, Millside, The Moor, Melbourn, Royston SG8 6ED  
(GB).
- (74) Agent: **ORIGIN LIMITED**; Twisden Works, Twisden Road, London NW5 1DN (GB).
- (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, 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,

(54) Title: A ROBOTIC FRUIT PICKING SYSTEM




FIGURE 1

(57) Abstract: A robotic fruit picking system includes an autonomous robot that includes a positioning subsystem that enables autonomous positioning of the robot using a computer vision guidance system. The robot also includes at least one picking arm and at least one picking head, or other type of end effector, mounted on each picking arm to either cut a stem or branch for a specific fruit or bunch of fruits or pluck that fruit or bunch. A computer vision subsystem analyses images of the fruit to be picked or stored and a control subsystem is programmed with or learns picking strategies using machine learning techniques. A quality control (QC) subsystem monitors the quality of fruit and grades that fruit according to size and/or quality. The robot has a storage subsystem for storing fruit in containers for storage or transportation, or in punnets for retail.



SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,  
TR, TT, TZ, UA, UG, US, UZ, VC, VN, 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))*

## A ROBOTIC FRUIT PICKING SYSTEM

### Field of the Invention

5

The field of the invention relates to a robotic fruit picking system comprising an autonomous robot and a method of selectively storing or punnetising the fruit with the optimal flavour or quality by using a robotic fruit picking system –comprising an autonomous robot. For example, the invention relates to systems and methods for robotic fruit picking.

10

### Background of the Invention

A portion of the disclosure of this patent document contains material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

15

Horticultural producers depend critically on manual labour for harvesting their crops. Many types of fresh produce are harvested manually including berry fruits such as strawberries and raspberries, asparagus, table grapes and eating apples. Manual picking is currently necessary because the produce is prone to damage and requires delicate handling, or because the plant itself is valuable, producing fruit continuously over one or more growing seasons. Thus the efficient but destructive mechanical methods used to harvest crops such as wheat are not feasible.

20

25

Reliance on manual labour creates several problems for producers:

30

- Recruiting pickers for short, hard picking seasons is risky and expensive. Domestic supply of picking labour is almost non-existent and so farmers must recruit from overseas. However, immigration controls place a large administrative burden on the producer and increase risk of labour shortage.
- Supply and demand for low-skilled, migrant labour are unpredictable because they depend on weather conditions throughout the growing season and economic

circumstances. This creates significant labour price fluctuations.

- In extremis, this can lead to crops being left un-harvested in the field. E.g. a single 250-acre strawberry farm near Hereford lost more than £200K of produce because of labour shortage in 2007.
- 5 • Human pickers give inconsistent results with direct consequences for profitability (e.g. punnets containing strawberries with inconsistent size or shape or showing signs of mishandling would typically be rejected by customers). Farmers use a variety of training and monitoring procedures to increase consistency but these greatly increase cost.

10

Current technologies for robotic soft fruit harvesting tend to rely on sophisticated hardware and naive robot control systems. In consequence, other soft fruit picking systems have not been commercially successful because they are expensive and require carefully controlled environments.

15

A small number of groups have developed robotic strawberry harvesting technology. However, the robots often come at a high cost and still need human operators to grade and post-process the fruit. Furthermore, the robots are often not compatible with table top growing systems used in Europe and are too expensive to be competitive with

20 human labour. Expensive hardware and dated object recognition technology are used, and lacked the mechanical flexibility to pick other than carefully positioned, vertically oriented strawberries or have been poorly suited to the problem of picking soft fruits that cannot be handled except by their stalks. No product offering has therefore materialized.

25

Hence most of the solutions to date fall down in at least two of the following key areas:

- They require growers to change their working practices significantly, or don't support the table top growing systems used in Europe.
- They rely heavily on human operators. In consequence, they use large machines with
- 30 disproportionately high production cost per unit picking capacity compared to small, autonomous machines manufactured in larger quantities.
- They are too expensive to displace human labour at current prices.

Elsewhere, several academic groups have also applied (mostly quite dated 1980's era)

robotics and computer vision technology to more general harvesting applications. However, the resulting systems have been too limited for commercial exploitation.

5 Farmers need a dependable system for harvesting their crops, on demand, with consistent quality, and at predictable cost. Such a system will allow farmers to buy a high quality, consistent harvesting capacity in advance at a predictable price, thereby reducing their exposure to labour market price fluctuations. The machine will function autonomously: traversing fields, orchard, or polytunnels; identifying and locating produce ready for harvest; picking selected crops; and finally grading, sorting and depositing  
10 picked produce into containers suitable for transfer to cold storage.

It is generally desirable to overcome or ameliorate one or more of the above described difficulties, or to at least provide a useful alternative.

15 Summary of the Invention

According to the present invention, there is provided a robotic fruit picking system comprising an autonomous robot that includes the following subsystems:

20 a positioning subsystem configured to autonomously position the robot using a computer implemented guidance system, such as a computer vision guidance system;

at least one picking arm;

at least one picking head or other type of end effector, mounted on a picking arm to either cut a stem or branch for a specific fruit or bunch of fruits or pluck that fruit or bunch, and then transfer the fruit or bunch;

25 a computer vision subsystem to analyse images of the fruit to be picked or stored;

a storage subsystem for receiving picked fruit and storing that fruit in containers for storage or transportation, or in punnets for retail; and

30 in which the end effector is configured to (a) grip a stem or stalk of a fruit or a bunch of fruits and (b) cut that stem or stalk and/or pluck that fruit or bunch; the end-effector being controlled by the computer vision subsystem to separate the edible and palatable part from at least a part of the stem, branch or stalk without contacting the edible and palatable part;

and in which the system is configured to estimate the statistical probability that a

picking attempt will be successful by taking into account one or more of the following: estimated pose and shape of the target fruit and its stem or stalk, uncertainty associated with the recovered pose and shape estimates, colour of the target fruit's surface, the proximity of detected obstacles, and the range of viewpoints from which the target fruit is visible, such that a fruit is picked when the estimated picking success is greater than a pre-defined threshold.

According to the present invention, there is also provided a method of selectively storing or punnetising the fruit with the optimal flavour or quality by using a robotic fruit picking system –comprising an autonomous robot that includes the following subsystems:

a positioning subsystem configured to autonomously position the robot using a computer implemented guidance system, such as a computer vision guidance system;

at least one picking arm;

at least one picking head or other type of end effector, mounted on a picking arm to either cut a stem or branch for a specific fruit or bunch of fruits or pluck that fruit or bunch, and then transfer the fruit or bunch;

a computer vision subsystem to analyse images of the fruit to be picked or stored;

a storage subsystem for receiving picked fruit and storing that fruit in containers for storage or transportation, or in punnets for retail; and

in which the end effector is configured to (a) grip a stem or stalk of a fruit or a bunch of fruits and (b) cut that stem or stalk and/or pluck that fruit or bunch; the end-effector being controlled by the computer vision subsystem to separate the edible and palatable part from at least part of the stem, branch or stalk without contacting the edible and palatable part;

and in which the system is configured to estimate the statistical probability that a picking attempt will be successful by taking into account one or more of the following: estimated pose and shape of the target fruit and its stem or stalk, uncertainty associated with the recovered pose and shape estimates, colour of the target fruit's surface, the proximity of detected obstacles, and the range of viewpoints from which the target fruit is visible, such that a fruit is picked when the estimated picking success is greater than a pre-defined threshold.

A first preferred embodiment of the invention is a robotic fruit picking system comprising an autonomous robot that includes the following subsystems:

a positioning subsystem operable to enable autonomous positioning of the robot using a computer implemented guidance system, such as a computer vision guidance

5 system;

at least one picking arm;

at least one picking head, or other type of end effector, mounted on each picking arm to either cut a stem or branch for a specific fruit or bunch of fruits or pluck that fruit or bunch, and then transfer the fruit or bunch;

10 a computer vision subsystem to analyse images of the fruit to be picked or stored;

a control subsystem that is programmed with or learns picking strategies;

a quality control (QC) subsystem to monitor the quality of fruit that has been picked or could be picked and grade that fruit according to size and/or quality; and

15 a storage subsystem for receiving picked fruit and storing that fruit in containers for storage or transportation, or in punnets for retail.

We use the term 'picking head' to cover any type of end effector; an end effector is the device, or multiple devices at the end of a robotic arm that interacts with the environment – for example, the head or multiple heads that pick the edible and palatable part of the fruit or that grab and cut the stems to fruits.

20

#### Brief Description of the Drawings

25 Preferred embodiments of the present invention are hereafter described, by way of non-limiting example only, with reference to the accompanying drawings, in which:

**Figure 1** shows a top view (A) and a perspective view (B) of a robot suitable for fruit picking.

30 **Figure 2** shows a perspective view of a robot suitable for fruit picking with arms in picking position.

**Figure 3** shows an example of a robot suitable for fruit picking.

**Figure 4** shows another example of a robot suitable for fruit picking.

**Figure 5** shows another example of a robot suitable for fruit picking.

- 2017357645 05 Oct 2022
- Figure 6** shows an early embodiment of the invention designed for picking tabletop-grown strawberries.
- Figure 7** shows a system for mounting a ‘vector cable’ to the legs of the tables on which crops are grown using metal brackets that simply clip to the legs.
- 5 **Figure 8** shows a number of line drawings with different views of the picking arm and picking head alone or in combination.
- Figure 9** shows a sectional view of a typical embodiment of a QC imaging chamber.
- Figure 10** shows a diagram illustrating a space-saving scheme for storing trays of punnets within a picking robot.
- 10 **Figure 11** shows the main components of the end effector.
- Figure 12** shows the hook extended relative to the gripper/cutter.
- Figure 13** shows the hook retracted relative to the gripper/cutter.
- Figure 14** shows the individual parts of the end effector including the blade above the hook.
- 15 **Figure 15** shows a diagram illustrating the movement of the hook to effect a capture of the plant stem.
- Figure 16** shows a diagram illustrating the plant stem being captured.
- Figure 17** shows a diagram illustrating the produce being gripped and cut from the parent plant.
- 20 **Figure 18** shows a diagram illustrating the release operation.
- Figure 19** shows the sequence of operations that constitute the picking process.
- Figure 20** shows the main mechanical constituents of the loop and jaw assembly.
- Figure 21** shows the loop and jaw assembly, shown with component 1, the loop, extended.
- 25 **Figure 22** shows an exploded diagram of the main constituent parts of the loop and jaw assembly (the loop is omitted for clarity).
- Figure 23** shows the components of the loop actuation mechanism.
- Figure 24** shows the loop/jaw assembly on its approach vector towards the target fruit.
- 30 **Figure 25** shows the loop extended and the assembly moving in parallel to the major axis of the target produce and in the direction of the juncture between stalk and fruit.

- Figure 26** show the loop having travelled past the juncture of the stalk and the fruit, the target produce is now selected and decision step 1 (**Figure 19**) may be applied.
- 5 **Figure 27** shows the loop is retracted to control the position of the target produce and the jaw is actuated so as to grab and cut the stalk of the fruit in a scissor-like motion.
- Figure 28** shows the different elements of the cable management system.
- Figure 29** shows drawings of the cable management system in situ within one of the joints of an arm.
- 10 **Figure 30** shows a sequence of drawings with the cable guide rotating within the cable enclosure.
- Figure 31** shows a cutaway view of cable winding.

#### Detailed Description of Preferred Embodiments of the Invention

15

The invention relates to an innovative fruit picking system that uses robotic picking machines capable of both fully autonomous fruit harvesting and working efficiently in concert with human fruit pickers.

20 Whilst this description focuses on robotic fruit picking systems, the systems and methods described can have a more generalized application in other areas, such as robotic litter picking systems.

25 The picking system is applicable to a variety of different crops that grow on plants (like strawberries, tomatoes), bushes (like raspberries, blueberries, grapes), and trees (like apples, pears, logan berries). In this document, the term *fruit* shall include the edible and palatable part of all fruits, vegetables, and other kinds of produce that are picked from plants (including e.g. nuts, seeds, vegetables) and *plant* shall mean all kinds of fruit producing crop (including plants, bushes, trees). For fruits that grow in clusters or

30 bunches (e.g. grapes, blueberries), *fruit* may refer to the individual fruit or the whole cluster.

Many plants continue to produce fruit throughout the duration of a long picking season and/or throughout several years of the plant's life. Therefore, a picking robot must not

damage either the fruit or the plant on which it grows (including any not-yet-picked fruit, whether ripe or unripe). Damage to the plant/bush/tree might occur either as the robot moves near the plant or during the picking operation.

5 In contrast to current technologies, our development efforts have been directly informed by the needs of real commercial growers. We will avoid high cost hardware by capitalizing on state-of-the-art computer vision techniques to allow us to use lower cost off-the-shelf components. The appeal of this approach is that the marginal cost of manufacturing software is lower than the marginal cost of manufacturing complex  
10 hardware.

An intelligent robot position control system capable of working at high speed without damaging delicate picked fruit has been developed. Whilst typical naïve robot control systems are useful for performing repeated tasks in controlled environments (such as car  
15 factories), they cannot deal with the variability and uncertainty inherent in tasks like fruit picking. We address this problem using a state-of-the-art reinforcement learning approach that will allow our robot control system to learn more efficient picking strategies using experience gained during picking.

20 Key components of the fruit picking robot are the following:

- A tracked rover capable of navigating autonomously along rows of crops using a vision-based guidance system.
- A computer vision system comprising a 3D stereo camera and image processing software for detecting target fruits, and deciding whether to pick them and how to  
25 pick them.
- A fast, 6 degree-of-freedom robot arm for positioning a picking head and camera.
- A picking head, comprising a means of (i) cutting the strawberry stalk and (ii) gripping the cut fruit for transfer.
- A quality control subsystem for grading picked strawberries by size and quality.
- 30 • A packing subsystem for on-board punnetization of picked fruit.

The picking robot performs several functions completely automatically:

- loading and unloading itself onto and off of a transport vehicle;

- navigating amongst fruit producing plants, e.g. along rows of apple trees or strawberry plants;
  - collaborating with other robots and human pickers to divide picking work efficiently;
  - 5 • determining the position, orientation, and shape of fruit;
  - determining whether fruit is suitable for picking;
  - separating the ripe fruit from the tree;
  - grading the fruit by size and other measures of suitability;
  - transferring the picked fruit to a suitable storage container.
- 10 The picking system is innovative in several ways. In what follows, some specific non-obvious inventive steps are highlighted with wording like “A useful innovation is...”.

### 1. System overview

The picking system comprises the following important subsystems:

- Total Positioning Subsystem
- 15 • Picking Arm
- Picking Head
- Computer Vision Subsystem
- Control Subsystem
- Quality Control (QC) Subsystem
- 20 • Storage Subsystem
- Mapping Subsystem
- Management Subsystem

The main purpose of the robot Total Positioning System is physically to move the whole robot along the ground. When the robot is within reach of target fruit, the Picking Arm  
 25 moves an attached camera to allow the Computer Vision Subsystem to locate target fruits and determine their pose and suitability for picking. The Picking Arm also positions the Picking Head for picking and moves picked fruit to the QC Subsystem (and possibly the Storage Subsystem). The Total Positioning System and the Picking Arm operate under the control of the Control Subsystem, which uses input from the  
 30 Computer Vision Subsystem to decide where and when to move the robot. The main purpose of the Picking Head is to cut the fruit from the plant and to grip it securely for

transfer to the QC and Storage subsystems. Finally, the QC Subsystem is responsible for grading picked fruit, determining its suitability for retail or other use, and discarding unusable fruit.

5 **Figure 1** shows a top view (a) and a perspective view (b) of a robot suitable for fruit picking. The robot includes a Positioning Subsystem operable to enable autonomous positioning of the robot using a computer implemented guidance system, such as a Computer Vision Guidance System. Two Picking Arms (100) are shown for this configuration. A Picking Head (101) is mounted on each Picking Arm (100) to either cut  
10 a stem or branch for a specific fruit or bunch of fruits or pluck that fruit or bunch, and then transfer the fruit or bunch. The Picking Head also includes the camera component (102) of the Computer Vision Subsystem, which is responsible for analysing images of the fruit to be picked or stored. A Control Subsystem is programmed with or learns picking strategies. A Quality Control (QC) Subsystem (103) monitors the quality of fruit  
15 that has been picked or could be picked and grade that fruit according to size and/or quality and a Storage Subsystem (104) receives the picked fruit and stores that fruit in containers for storage or transportation, or in punnets for retail. **Figure 2** shows a perspective view of the robot with arms in picking position.

20 **Figures 3-5** show examples of a robot suitable for fruit picking. Three different conceptual illustrations show the robot configured in several different ways, suitable for different picking applications. Important system components shown include the tracked rover, two Picking Arms, and associated Quality Control units (QC). Multiple trays are used to store punnets of picked fruit and are positioned so that a human operator can  
25 easily remove full trays and replace them with empty trays. A discard shoot is positioned adjacent to the Quality Control units for fruit not suitable for sale. When the robot picks rotten or otherwise unsuitable fruit (either by accident or design), it is usually desirable to discard the rotten fruit into a suitable container within the robot or onto the ground. A useful time-saving innovation is to make the container accessible via a discard chute with  
30 its aperture positioned at the bottom of the QC rig so that the arm can drop the fruit immediately without the need to move to an alternative container. A related innovation is to induce positive or negative air pressure in the chute or the body of the imaging chamber (e.g. using a fan) to ensure that fungal spores coming from previously discarded fruit are kept away from healthy fruit in the imaging chamber. The two Picking Arms can

be positioned asymmetrically (as shown in **Figure 5**) to increase reach at the expense of picking speed.

**Figure 6** shows an early embodiment of the invention designed for picking tabletop-grown strawberries and having a single Picking Arm and two storage trays.

These subsystems will be described in more detail in the following sections.

## 2. Total Positioning Subsystem

The Total Positioning Subsystem is responsible for movement of the whole robot across the ground (typically in an intended straight line between a current position and an input target position). The Total Positioning Subsystem is used by the Management Subsystem to move the robot around. The Total Positioning System comprises a means of determining the present position and orientation of the robot, a means of effecting the motion of the robot along the ground, and a control system that translates information about the current position and orientation of the robot into motor control signals.

### 2.1. Pose determination component

The purpose of the pose determination component is to allow the robot to determine its current position and orientation in a map coordinate system for input to the Control Component. Coarse position estimates may be obtained using *differential GPS* but these are insufficiently accurate for following rows of crops without collision. Therefore, a combination of additional sensors is used for more precise determination of heading along the row and lateral distance from the row. The combination may include ultrasound sensors for approximate determination of distance from the row, a magnetic compass for determining heading, accelerometers, and a forwards or backwards facing camera for determining orientation with respect to the crop rows. Information from the sensors is fused with information from the GPS positioning system to obtain a more accurate estimate than could be obtained by either GPS or the sensors individually.

An innovative means of allowing the robot to estimate its position and orientation with respect to a crop row is to measure its displacement relative to a tensioned cable (perhaps of nylon or other low cost material) that runs along the row (a 'vector cable').

One innovative means of measuring the displacement of the robot relative to the vector cable is to use a computer vision system to measure the projected position of the cable in 2D images obtained by a camera mounted with known position and orientation in the robot coordinate system. As a simplistic illustration, the orientation of a horizontal cable in an image obtained by a vertically oriented camera has a simple linear relationship with the orientation of the robot. In general, the edges of the cable will project to a pair of lines in the image, which can be found easily by standard image processing techniques, e.g. by applying an edge detector and computing a Hough transform to find long straight edges. The image position of these lines is a function of the diameter of the cable, the pose of the camera relative to the cable, and the camera's intrinsic parameters (which may be determined in advance). The pose of the camera relative to the cable may then be determined using standard optimization techniques and an initialization provided by the assumption that the robot is approximately aligned with the row. A remaining one-parameter ambiguity (corresponding to rotation of the camera about the axis of the cable) may be eliminated knowing the approximate height of the camera above the ground.

Another innovative approach to determining position relative to the vector cable is to use a follower arm (or follower arms). This is connected at one end to the robot chassis by means of a hinged joint and at the other to a truck that runs along the cable. The angle at the hinged joint (which can be measured e.g. using the resistance of a potentiometer that rotates with the hinge) can be used to determine the displacement relative to the cable. Two follower arms (e.g. one at the front and one at the back) is sufficient to determine displacement and orientation.

A related innovation is a bracket that allows vector cables to be attached easily to the legs of the tables on which crops such as strawberries are commonly grown. This is illustrated in **Figure 7**. The brackets allow the cable to be positioned a small and consistent distance above the ground, typically 20 cm, so that is easy for humans to step over it whilst ducking underneath elevated tables to move from row to row. One limitation of this approach is that the robot may receive spurious position information if the follower arm falls off the cable. Therefore, a useful additional innovation is to equip the truck with a microswitch positioned so as to break an electrical circuit when the truck loses

contact with the cable. This can be used to allow the control software to detect this failure condition and stop the robot (typically, the control software waits until the duration of detected loss of contact between the truck and the cable is greater than some time threshold to eliminate false detections due to bouncing of the truck on the cable).

5 Since the follower arm might be subjected to significant forces in the event of collision or other kind of failure of control system failure, another useful innovation is to use a magnetic coupling to attach an outer portion of the follower arm to an inner portion. Then in the event of failure the parts of the follower arm can separate without suffering permanent damage. The magnetic coupling can include electrical connections required to  
10 complete an electrical circuit (typically the same circuit that is interrupted by a microswitch in the truck). By this means separation of the follower arm can also trigger the control software to stop the robot. Another benefit of the magnetic coupling arranging is ease of attachment of the follower arm by a human supervisor.

15 An innovative aspect of the pose determination component is a computer vision based system for determining the heading and lateral position of a robot with respect to a row of crops using images obtained by a forwards or backwards facing camera pointing approximately along the row. Such a system can be used to drive the robot in the middle of two crop rows or at a fixed distance from a single crop row. In one embodiment of  
20 this idea, this is achieved by training a regression function implemented using a convolutional neural network (or otherwise) to predict robot heading and lateral position with respect to rows of crops as a function of an input image. Training data may be obtained by driving a robot equipped with multiple forwards and/or backwards facing cameras between representative crop rows under human remote control. Because the  
25 human controller keeps the robot approximately centred between the rows (with heading parallel to the rows), each frame can be associated with approximate ground truth heading and lateral displacement information. Multiple cameras are used to provide a training images corresponding to different approximate lateral displacements from the row. Training images corresponding to different robot headings can be obtained by  
30 resampling images obtained by a forwards looking camera using an appropriate 3-by-3 homography (which can be computed trivially from known camera intrinsic calibration parameters).

A related innovation is to obtain additional training image data at night using an infrared

illuminator and suitable infrared receptive cameras.

In another embodiment of this idea, a computer vision system is designed to detect (in images obtained by a forwards or backwards facing camera) the vertical legs of the tables on which crops are grown. Vertically oriented table legs define vertical lines in the world, which project to lines in the perspective view. Under the assumption that the legs of each table are evenly spaced, vertically oriented, and arranged in a straight line, projected image lines corresponding to a sequence of three or more table legs are sufficient to determine the orientation of a calibrated camera and its lateral displacement with respect to a 3D coordinate system defined by the legs.

**Figure 7** shows a system for mounting a ‘vector cable’ to the legs of the tables on which crops are grown (a) using metal brackets (b) that simply clip to the legs. The cable (e.g. a nylon rope) may be secured to the metal brackets using n-shaped metal spring clips or tape (not shown). The row follower arm (seen in b) comprises a truck that rests on the cable and an arm that attaches the truck to the robot (not shown). The brackets are shaped such that the truck is not impeded as it traverses them.

### ***2.2. Motor control component***

The purpose of the motor control component is to map pose information provided by the pose determination component to motor control signals to move the robot in a given direction. It supports two kinds of motion: (i) moving a given distance along a row of plants and (ii) moving to a given point by travelling in an intended straight line. The motor control system uses a PID controller to map control inputs obtained from the pose determination component to motor control signals.

### ***2.3. Rover***

An important component of the Total Positioning System is the rover, the means by which the robot moves over the ground. Typically, movement over the ground is achieved using powered wheels with tracks. A useful innovation is a mechanism to allow the tracks to be removed so that the robot can also run on rails.

### 3. Picking Arm

The Picking Arm is a robot arm with several (typically 6) degrees of freedom that is mounted to the main body of the robot. Whereas the purpose of the Total Positioning System is to move the whole robot along the ground, the purpose of the Picking Arm is to move the Picking Head (and its computer vision camera) to appropriate positions for locating, localizing, and picking target fruit. Once it is in the picking position, the Picking Head executes a picking routine that comprises a sequence of mechanical actions including separation, gripping, and cutting (see the Picking Head description below). Picking positions are chosen by the Control Subsystem to maximize picking performance according to a desired metric.

Before the Picking Head can be positioned to pick a target fruit the Computer Vision Subsystem must carry out several important operations: (i) detecting the target fruit, (ii) detecting obstructions that might complicate picking of the target fruit (e.g. leaves), (iii) determining the pose and shape of the target fruit. So that the Computer Vision Subsystem can perform these functions, the Picking Arm may be equipped with a monocular or stereo camera, mounted e.g. to the end of the arm. The benefit of having a camera mounted to the arm is the possibility of moving the camera to find viewpoints free from sources of occlusion that would otherwise prevent reliable detection or localization of the target fruit (leaves, other fruits, etc.).

Finally, the Picking Arm must move the Picking Head to an appropriate pose for picking without colliding with the plant, or the support infrastructure used to grow the plant, or itself. This is achieved using a route-planning algorithm described in the Control Subsystem section below.

**Figure 8** shows a number of line drawings with different views of the Picking Arm and Picking Head alone or in combination. The Picking Head shown includes a hook and camera system.

### 4. Picking Head

The purpose of the Picking Head is to sever the target fruit from the plant, to grasp it securely while it is moved to the QC and Storage Subsystems, and to release it. A secondary purpose is to move leaves and other sources of occlusion out of the way so

fruit can be detected and localized, and to separate target fruit from the plant (before it is permanently severed) to facilitate determination of picking suitability.

Picking soft fruits like strawberries is challenging because physical handling of the fruit can cause bruising, reducing saleability. Therefore, such fruits are ideally picked by severing the stem without handling the body of the fruit. An inventive aspect of our system is the use of a stem-severing Picking Head that works in three phases ('grab-grip-cut'):

1. Physical separation of the fruit from the remainder of the plant ('grab'). In advance of permanently severing the fruit from the plant, this step facilitates (i) deciding whether or not the fruit is suitable for picking and (ii) increases the chances that picking in step 2 will proceed successfully without damage to the target fruit of the rest of the plant.
2. Gripping of the picked fruit by its stem ('grip').
3. Severing the stem (above the point at which it is gripped) so as permanently to separate the fruit from the plant ('cut').

The introduction of the physical separation phase (which take place before the fruit is permanently severed from the plant) confers several benefits. Since target fruit may be occluded by leaves or other fruit, pulling it further away from the plant facilitates a better view, allowing the computer vision system to determine more reliably whether the fruit is ready for picking and whether the picking procedure is likely to be successful (for example because other fruits are in the target vicinity). A related innovation is a mechanical gripper that can rotate the gripped fruit during this before-picking inspection phase, e.g. by applying a twisting force to its stalk or otherwise. By this means, a camera or other sensors can obtain information about parts of the fruit that would not otherwise have been visible. One benefit of this innovation is the possibility of deciding to postpone picking a fruit that appears unripe on the once-hidden side.

A possible further innovation is to combine the grip and cut phases (2 and 3) by means of exploiting the gripping action to pull the stem against a cutting blade or blades.

Appendix A describes several innovative mechanical Picking Head designs embodying some of these ideas.

For some soft fruits such as raspberries, it is necessary to remove the fruit from its stem during picking. For such fruits, a useful innovation is to pick the fruit by first severing and gripping its stem and then to remove the body of the fruit from its stem in a subsequent operation. Compared to picking techniques that require holding or gripping the body of the fruit, important benefits if this approach include: (i) minimizing handling of the body of the fruit (which can significantly reduce shelf life, e.g. due to the transference of disease-causing pathogens from fruit to fruit on the surface of the handling device), (ii) the possibility of imaging the body of the picked fruit from all directions for quality control, and (iii) the possibility of removing the stem under controlled conditions.

Various means of removing the picked fruit from the stem are possible. One innovative approach is to pull the fruit from its stem using a jet of compressed air. This allows contact forces to be distributed evenly over a large contact area, minimizing bruising. Another possibility is to pull the fruit by its stem through a collar, shaped to facilitate forcing the body of the fruit off the stem. Depending on the specific type of fruit, collars might be designed to either to distribute the contact force over a large area of the body of the fruit or to concentrate the contact force (possibly via a cutting edge resembling that of a knife or row of needles) in a circular pattern surrounding the stem. A related innovation is to clean the collar after each use or to provide the collar with a disposable surface to reduce the likelihood of transfer of pathogens from fruit to fruit.

Another innovation is to use the inertia of the body of the fruit to separate the body of the fruit from the receptacle. This might be achieved by holding the fruit by its stalk and rotating it about an axis perpendicular to its stalk at a high enough angular velocity. The advantage of this approach is that inertial forces effectively act over the entire mass of the body of the fruit, eliminating the need for contact forces at the surface (more likely to cause bruising because they are applied over a small contact area, increasing localized pressure). One limitation of this approach is that, when the body of the fruit separates from the receptacle, it will fly off at speed on a tangent to the circle of rotation, necessitating some means of arresting its motion sufficiently slowly that it doesn't suffer bruising. Therefore, another innovation is to use a reciprocating back-and-forth motion of the fruit or its stalk in the direction approximately perpendicular to the stalk or an

oscillatory rotary motion with an axis of rotation approximately parallel to the stalk. By performing the motion at an appropriate frequency and with appropriate amplitude it is possible reliably to separate the body of the fruit from the husk without causing it to fly off at high velocity.

5

After picking, the Picking Head grips the fruit as it transferred by the Picking Arm to the Quality Control (QC) Subsystem. In a simple embodiment, the Picking Arm itself might be used to position the picked fruit inside the imaging component of the QC Subsystem before subsequently moving the fruit to the storage. However, time spent transferring  
10 picked fruit to the QC/Storage Subsystem is unproductive because the arm is not being used for picking during the transfer. Therefore, a useful innovation is to include multiple picking units on a single *multiplexed* Picking Head. This means that several fruits in a particular local vicinity can be picked in succession before the arm has to perform the time-consuming movement between the plant and the QC/Storage components and  
15 back. This means that the transfer overhead can be amortized over more picked fruits, increasing productivity. This is particularly advantageous in the common case that fruit are bunched on the plant/tree such the robot arm needs to move only a small distance to pick several targets before transfer.

20 Picking units on the multiplexed Picking Head must be arranged so that inactive picking units do not interfere with the operation of active picking units, or collide with the arm, or other objects. Innovative ways of achieving this include:

- mounting the picking units radially about an axis chosen so that inactive picking units are oriented away from the active picking unit and the fruit being picked.
- 25 • making each picking unit extend independently so it can engage with the fruit while others do not disturb the scene.

Picking units typically have several moving parts e.g. for hooking, cutting, etc., which may need to be driven independently. However, when multiplexing multiple units on a single Picking Head, if each moving part is driven with its own actuator, the arm payload  
30 increases proportionally to the number of picking units, which would adversely affect arm speed, accuracy, and the overall cost of the machine. Several innovative aspects of the implementation of the multiplexed Picking Head keep the overall mass of the multiplexed Picking Head low to allow the arm to move quickly and accurately:

- Multiple picking functions on a picking unit can be driven by a single actuator or motor, selectively engaged by lightweight means, for example electromagnets; an engaging pin; rotary tab; or similar. This is challenging as the different functions may require different actuator characteristics
- 5 • A single motor or actuator can drive one function across all units on the head, selectively engaged by means of an electromagnet; an engaging pin; rotary tab; or similar. This is reasonably straightforward.
- The functions can be driven by lightweight means from elsewhere in the system, for example using a Bowden cable, torsional drive cable/spring, pneumatic or  
10 hydraulic means.

### 5. Computer Vision Subsystem

The purpose of the Computer Vision Subsystem is to locate target fruits, determine their pose in a robot coordinate system, and determine whether they are suitable for picking,  
15 i.e. *before* the fruit is permanently separated from the plant.

To achieve this, the Computer Vision Subsystem uses one or more cameras mounted to the end of the movable Picking Arm (or in general to any other part of the robot such as the chassis). The camera attached to the robot arm is moved under computer control to  
20 facilitate the detection of target fruits, estimation of their pose (i.e. position and orientation), and determination of their likely suitability for picking. Pose estimates and picking suitability indicators associated with each target fruit may be refined progressively as the arm moves. However, this refinement stage takes time, which increases the time required for picking. Therefore, an important innovation is a scheme for moving the arm  
25 efficiently to optimize trade-off between picking speed and picking accuracy (this scheme is described in more detail in the Robot Control Subsystem section, below).

The Computer Vision Subsystem operates under the control of the Robot Control Subsystem, which makes a continuous sequence of decisions about which action to  
30 perform next, e.g. moving the arm-mounted camera to new viewpoints to facilitate discovery of more target fruits, or moving the camera to new points in the local vicinity of a target fruit so as to refine estimates of its position/orientation or indicators of picking suitability.

In outline, the Computer Vision Subsystem works as follows:

1. Under the control of the Robot Control Subsystem, the camera captures images of the scene from multiple viewpoints.
- 5 2. Target fruits are detected in the captured images by pixel-wise semantic segmentation.
3. Approximate estimates of pose and shape are recovered for each detected fruit.
4. More accurate estimates of pose and shape are recovered by combining information in multiple views with statistical prior knowledge. This is achieved by  
10 adapting the parameters of a generative model of strawberry appearance to maximize agreement between the predictions and the images.
5. Picking success probability for each detected fruit is estimated from visual and geometric cues.
6. Under the control of the Control Subsystem, additional images of a particular  
15 target fruit may be captured from new viewpoints so as to increase picking success probability.

The important steps are described in more detail below.

**Image capture.** An important challenge is to control the exposure of the camera system  
20 to obtain images that are consistently correctly exposed. Failure to do this increases the amount of variability in the images, compromising the ability of the machine-learning-based target detection software to detect fruit accurately and reliably. One exposure control strategy is to obtain an image of a grey card exposed to ambient lighting conditions. This image is then analysed to determine the adjustments to exposure time  
25 and/or colour channel gains required to ensure that the grey card appears with a predetermined target colour value. A grey card might be positioned on the robot chassis with reach of the Picking Arms and oriented horizontally to measure ambient illumination arriving from the approximate direction of the sky. However, a potential limitation of this approach is that the illumination of the grey card may not be  
30 representative of the illumination of the plant or target fruit. Therefore, in a system where a (stereo) camera is incorporated within the Picking Head, a useful innovation is to arrange that a part of the Picking Head itself can be used as an exposure control target. A prerequisite is that the exposure control target must appear within the field of view of the

camera. A suitable target could be a grey card imaged from in front or a translucent plastic diffuser imaged from underneath.

5 Real world lighting conditions can compromise image quality, limiting the effectiveness of image processing operations such as target fruit detection. For example, images obtained by a camera oriented directly towards the sun on a cloudless day may exhibit lens flare. Therefore, a useful innovation is to have control system software use the weather forecast to schedule picking operations so that the robot's camera systems are oriented to maximize the quality of the image data being obtained as a function of  
10 expected lighting conditions over a given period. For example, on a day that is forecast to be sunny in a farm where fruit is grown in rows, the robot might pick on one side of the rows in the morning and the other side of the rows in the afternoon. On a cloudy day, robots might more usefully pick on both sides of the row simultaneously to amortize the cost of advancing the robot along the row over more picked fruit at each position. A  
15 related innovation is to adapt viewpoints dynamically as a function of lighting conditions to maximize picking performance. For example, the Picking Head might be angled downwards in conditions of direct sunlight to avoid lens flare even at the expense of a reduction in working volume.

20 **Target detection.** Target fruit is detected automatically in images obtained by a camera mounted to the Picking Arm or elsewhere. A machine learning approach is used to train a detection algorithm to identify fruit in RGB colour images (and/or in depth images obtained by dense stereo or otherwise). To provide training data, images obtained from representative viewpoints are annotated manually with the position and/or extent of  
25 target fruit. Various embodiments of this idea are possible:

1. A decision forest classifier or convolutional neural network (CNN) may be trained to perform semantic segmentation, i.e. to label pixels corresponding to ripe fruit, unripe fruit, and other objects. Pixel-wise labelling may be noisy, and evidence may be aggregated across multiple pixels by using a clustering algorithm.
- 30 2. A CNN can be trained to distinguish image patches that contain a target fruit at their centre from image patches that do not. A sliding window approach may be used to determine the positions of all image patches likely to contain target fruits. Alternatively, the semantic labelling algorithm 1 may be used to identify the likely image locations of target fruits for subsequent more accurate classification by a

(typically more computationally expensive) CNN.

**Target pose determination.** Picking Heads for different types of fruit may work in different ways, e.g. by cutting the stalk or by twisting the fruit until the stalk is severed (see above, and Appendix A). Depending on the Picking Head design, picking a target fruit may necessitate first estimating the position and orientation (or *pose*) of the fruit or its stalk (in what follows, fruit should be interpreted to mean the body of the fruit or its stalk or both). Rigid body pose in general has 6 degrees of freedom (e.g. the X, Y, Z coordinates of a fruit in a suitable world coordinate system and the three angles describing its orientation relative to the world coordinate system's axes). Pose may be modelled as a 4-by-4 homography that maps homogenous 3D points in a suitable fruit coordinate system into the world coordinate system. The fruit coordinate system can be aligned with fruits of specific types as convenient. For example, the origin of the coordinate system may be located at the point of intersection of the body of the fruit and its stalk and the first axis points in the direction of the stalk. Many types of fruit (such as strawberries and apples) and most kinds of stalk have a shape with an axis of approximate rotational symmetry. This means that 5 degrees of freedom typically provide a sufficiently complete representation of pose for picking purposes, i.e. the second and third axes of the fruit coordinate system can be oriented arbitrarily.

The robot determines the pose of target fruit using images obtained from multiple viewpoints, e.g. using a stereo camera or a monocular camera mounted to the moving Picking Arm. For example, the detected position of a target fruit in two or more calibrated views is sufficient to approximately to determine its X,Y,Z position by *triangulation*. The orientation of the fruit or its stalk may then be estimated by assumption (e.g. the assumption that fruits hang vertically) or recovered from image data.

A useful innovation is to use a learned regression function to map images of target fruits directly to their orientation in a camera coordinate system. This can be achieved using a machine learning approach whereby a suitable regression model is trained to predict the two angles describing the orientation of an approximately rotationally symmetric fruit from images (including monocular, stereo, and depth images). This approach is effective for fruits such as strawberries that have surface texture that is aligned with the dominant axis of the fruit. Suitable training images may be obtained using a camera mounted to the

end of a robot arm. First, the arm is moved manually until the camera is approximately aligned with a suitable fruit-based coordinate system and a fixed distance away from the fruit's centroid. The arm is aligned so the fruit has canonical orientation in a camera image, i.e. so that the two or three angles used to describe orientation in the camera coordinate frame are 0. Then the arm moves automatically to obtain additional training images from new viewpoints with different, known relative orientations of the fruit. Sufficiently high quality training data can be obtained by having a human operator judge alignment between the camera and the fruit coordinate system visually by inspection of the scene and the video signal produced by the camera. Typically training images are cropped so that the centroid of the detected fruit appears in the centre of the frame and scaled so that the fruit occupies constant size. Then a convolutional neural network or other regression model is trained to predict fruit orientation in previously unseen images. Various image features are informative as to the orientation of the fruit in the camera image frame (and can be exploited automatically by a suitable machine learning approach), e.g. the density and orientation of any seeds on the surface of the fruit, the location of the calyx (the leafy part around the stem), and image location of the stalk.

Because knowledge of the orientation of the stalk may be very important for picking some types of fruits (or otherwise informative as to the orientation of the body of the fruit), another useful innovation is a stalk detection algorithm that identifies and delineates stalks in images. A stalk detector can be implemented by training a pixel-wise semantic labelling engine (e.g. a decision forest or CNN) using manually annotated training images to identify pixels that lie on the central axis of a stalk. Then a line growing algorithm can be used to delineate visible portions of stalk. If stereo or depth images are used, then stalk orientation can be determined in a 3D coordinate frame by matching corresponding lines corresponding to a single stalk in two or more frames. Solution dense stereo matching problem is considerably facilitated by conducting semantic segmentation of the scene first (stalks, target fruits). Assumptions about the range of depths likely to be occupied by the stalk can be used to constrain the stereo matching problem.

Given an approximate pose estimate for a target fruit, it may be that obtaining an additional view will improve the pose estimate, for example by revealing an informative part of the fruit such as the point where the stalk attaches. Therefore, a useful innovation

is an algorithm for predicting the extent to which additional views out of a set of available viewpoints will most significantly improve the quality of an initial pose estimate. Pose estimates obtained using multiple views and statistical prior knowledge about the likely shape and pose of target fruits can be fused using an innovative model fitting approach (see below).

**Size and shape determination and pose estimate refinement.** Whether a target fruit is suitable for picking may depend on its shape and size, e.g. because a customer wants fruit with diameter in a specified range. Furthermore, certain parameters of the picking system may need to be tuned considering the shape and size of the fruit, e.g. the trajectory of the Picking Head relative to the fruit during the initial ‘grab’ phase of the picking motion (see above). Therefore, it may be beneficial to estimate the shape and size of candidate fruits before picking as well as to refine (possibly coarse) pose estimates determined as above. This can be achieved using images of the fruit (including stereo images) obtained from one or more viewpoints.

An innovative approach to recovering the 3D shape of a candidate fruit from one or more images is to adapt the parameters of a generative model of the fruit’s image appearance to maximize the agreement between the images and the model’s predictions, e.g. by using Gauss-Newton optimization. This approach can also be used to refine a coarse initial estimate of the fruit’s position and orientation (provided as described above). A suitable model could take the form of a (possibly textured) triangulated 3D mesh projected into some perspective views. The shape of the 3D mesh could be determined by a mathematical function of some parameters describing the shape of the fruit. A suitable function could be constructed by obtaining 3D models of a large number of fruits, and then using Principal Component Analysis (or other dimensionality reduction strategy) to discover a low-dimensional parameterization of the fruit’s geometry. Another simpler but effective approach is to hand craft such a model, for example by assuming that the 3D shape of fruit can be explained as a volume of revolution to which parametric anisotropic scaling has been applied in the plane perpendicular to the axis. A suitable initialization for optimization can be obtained by using the 2D image shape (or the mean 2D image shape of the fruit) to define a volume of revolution. The pose parameters can be initialized using the method described above.

A key benefit of the model fitting approach is the possibility of combining information from several viewpoints simultaneously. Agreement between the real and predicted image might be measured, e.g. using the distance between the real and predicted silhouette or, for a model that includes lighting or texture, as the sum of squared differences between pixel intensity values. A useful innovation is to use the geometric model to predict not only the surface appearance of the fruit but the shadows cast by the fruit onto itself under different, controlled lighting conditions. Controlled lighting can be provided by one or more illuminators attached to the end of the robot arm. Another useful innovation is to model agreement using a composite cost function that comprises terms reflecting agreement both between silhouette and stalk.

Another benefit of the model fitting approach is the possibility of combining image evidence with statistical prior knowledge to obtain a *maximum likelihood* estimate of shape and pose parameters. Statistical prior knowledge can be incorporated by penalizing unlikely parameter configurations that are unlikely according to a probabilistic model. One valuable innovation is the use for this purpose of a statistical prior that model the way that massive fruits hang from their stalks under the influence of gravity. In a simple embodiment, the prior might reflect our knowledge that fruits (particularly large fruits) tend to hang vertically downwards from their stalks. Such a prior might take the simple form of a probability distribution over fruit orientation. A more complex embodiment might take the form of the joint distribution over the shape and size of the fruit, the pose of the fruit, and the shape of the stalk near the point of attachment to the fruit. Suitable probability distributions are usually formed by making geometric measurements of fruit growing under representative conditions.

Some Picking Head designs make it possible physically to separate a candidate fruit further from the plant and other fruits in the bunch before picking (see above). For example, the ‘hook’ design of Picking Head (see Appendix A) allows a candidate fruit to be supported by its stalk so that it hangs at a predictable distance from a camera mounted to the robot arm. One benefit of this innovation is the possibility of capturing an image (or stereo image) of the fruit from a controlled viewpoint, thereby facilitating more accurate determination of the size and shape, e.g. via shape from silhouette.

**Determination of picking suitability.** An attempt to pick a target fruit might or might

not be successful. Successful picking usually means that (i) the picked fruit is suitable for sale (e.g. ripe and undamaged) and delivered to the storage container in that condition, (ii) no other part of the plant or growing infrastructure is damaged during picking, and (iii) the Picking Arm does not undergo any collisions that could interfere with its  
5 continuing operation. However, in the case of rotten fruits that are picked and discarded to prolong the life of the plant, it is not a requirement that the picked fruit is in saleable condition.

A valuable innovation is to determine the picking suitability of a target fruit by estimating  
10 the statistical probability that an attempt to pick it will be successful. This probability can be estimated before attempting to pick a the target fruit via a particular approach trajectory and therefore can be used by the Control Subsystem to decide which fruit to pick next and how to pick it. For example, the fruits that are easiest to pick (i.e. those most likely to be picked successfully) might be picked first to facilitate subsequent  
15 picking of fruits that are harder to pick, e.g. because they are partly hidden behind other fruits. The picking success probability estimate can also be used to decide *not* to pick a particular target fruit, e.g. because the expected cost of picking in terms of damage to the plant or picked fruit will not outweigh the benefit provided by having one more picked fruit. The Control Subsystem is responsible for optimizing picking schedule to achieve  
20 the optimal trade-off between picking speed and failure rate (see below).

An important innovation is a scheme for estimating the probability of picking success using images of the scene obtained from viewpoints near the target fruit. For example, we might image the surface of the fruit by moving a camera (possibly a stereo or depth  
25 camera) mounted to the Picking Arm's end effector in its local vicinity. Various image measurements might be used as indicators of picking success probability, including e.g. (i) the estimated pose and shape of the fruit and its stalk, (ii) the uncertainty associated with the recovered pose and shape estimates, (iii) the colour of the target fruit's surface, (iv) the proximity of detected obstacles, and (v) the range of viewpoints from which the  
30 candidate fruit is visible.

A suitable statistical model for estimating picking success probability might take the form of a multivariate histogram or Gaussian defined on the space of all picking success indicators. An important innovation is to learn and refine the parameters of such a model

using picking success data obtained by working robots. Because the Quality Control Subsystem provides accurate judgments about the saleability of picked fruits, its output can be used as an indicator of ground truth picking success or failure. An online learning approach can be used to update the model dynamically as more data are generated to  
5 facilitate rapid adaptation of picking behaviour to the requirements of a new farm or phase of the growing season. Multiple robots can share and update the same model.

Since the Picking Head might approach a target fruit via a range of possible trajectories (depending on obstacle geometry and the degrees of freedom of the Picking Arm), the  
10 probability of picking success is modelled as a function of hypothesized approach trajectory. By this means, the Control Subsystem can decide how to pick the fruit to achieve the best trade-off between picking time and probability of picking success. The probability of collision between the Picking Arm and the scene can be modelled during the path planning operation using an explicit 3D model of the scene (as described in the  
15 Control Subsystem section below). However, an alternative and innovative approach is to use an *implicit* 3D model of the scene formed by the range of viewpoints from which the target fruit can be observed without occlusion. The underlying insight is that if the target fruit is wholly visible from a particular viewpoint, then the volume defined by the inverse projection of the 2D image perimeter of the fruit must be empty between the  
20 camera and the fruit. By identifying one or more viewpoints from which the target fruit appears un-occluded, obstacle free region of space is found. Provided no part of the Picking Head or Arm strays outside of this region of space during picking, there should be no collision. Occlusion of the target fruit by an obstacle between the fruit and the camera when viewed from a particular viewpoint can be detected by several means  
25 including e.g. stereo matching.

Another important innovation is a Picking Head that can pull the target fruit away from the plant, to facilitate more reliable determination of the fruit's suitability for picking  
*before* the fruit is permanently severed from the plant. Novel Picking Head designs are  
30 described in Appendix A.

## 6. Quality Control Subsystem

The primary function of the Quality Control Subsystem is to assign a measure of quality to individual picked fruits (or possibly individual bunches of picked fruits for fruits that

are picked in bunches). Depending on the type of fruit being picked and the intended customer, quality is a function of several properties of the fruit, such as ripeness, colour, hardness, symmetry, size, stem length. Picked fruit may be assigned a grade classification that reflects its quality, e.g. grade 1 (symmetric) or grade 2 (shows significant surface creasing) or grade 3 (very deformed or otherwise unsuitable for sale). Fruit of too low quality for retail sale may be discarded or stored separately for use in other applications, e.g. jam manufacture. An important implementation challenge is to ensure that the QC step can be carried out quickly to maximize the productivity of the picking robot.

A secondary function of the QC Subsystem is to determine a more accurate estimate of the fruit's size and shape. This estimate may be used for several purposes, e.g.

- for quality grading, since any asymmetry in the 3D shape for the fruit may be considered reason to assign a lower quality grade;
- as a means of estimating the fruit's mass and thereby of ensuring that the required mass of fruit is placed in each punnet according to the requirements of the intended customer for average or minimum mass per punnet;
- to facilitate more precise placement of the fruit in the storage container, and therefore to minimize the risk of bruising due to collisions.

The QC Subsystem generates a quality measure for each picked piece of fruit by means of a computer vision component comprising some cameras, some lights, and some software for image capture and analysis. Typically, the cameras are arranged so as to obtain images of the entire surface of a picked fruit that has been suitably positioned, e.g. by the Picking Arm. For example, for fruits like strawberries, which can be held so as to hang vertically downwards from their stalks, one camera might be positioned below the fruit looking upwards and several more cameras might be positioned radially about a vertical axis looking inwards. However, one limitation of this scheme is that a large amount of volume would be required to accommodate cameras (allowing for the camera-object distance, the maximum size of the fruit, and the tolerance inherent in the positioning of the fruit). One solution might be to rotate the fruit in front of a single camera to obtain multiple views – however any undamped motion of the fruit subsequent to rotation might complicate imaging. Therefore, another useful innovation is to use mirrors positioned and oriented so as to provide multiple virtual views of the fruit to a single camera mounted underneath the fruit. This scheme considerably reduces the both the cost and the size of the QC Subsystem. Cameras and/or mirrors are typically

arranged so that the fruit appears against a plain background in all views to facilitate segmentation of the fruit in the images.

Another useful innovation is to obtain multiple images under different lighting  
5 conditions. For example, this might be achieved by arranging a series of LED lights in a  
circle around the fruit and activating them one at a time, capturing one exposure per  
light. This innovation considerably increases the informativeness of the images because  
directional lights induce shadows both on the surface of the fruit and on a suitability  
10 positioned background screen. Such shadows can be used to obtain more information  
about the 3D shape of the fruit and e.g. the positions of any surface folds that could  
reduce saleability.

Using these images, image analysis software measures the fruit's 3D shape and detects  
various kinds of defect (e.g. rot, bird damage, spray residue, bruising, mildew, etc.). A  
15 useful first step is a semantic labelling step that is used to segment fruit from background  
and generate per pixel labels corresponding to the parts of the fruit (e.g. calyx, body,  
achene, etc.). In the same manner as the Computer Vision Subsystem (which makes  
crude 3D geometry measurements before picking) 3D geometry can be recovered by the  
QC Subsystem by adapting the parameters of a generative model to maximize the  
20 agreement between the model and the image data. Again, a statistical prior can be used to  
obtain a *maximum likelihood* estimate of the values of the shape parameters. A useful  
innovation is to use an estimate of the mass density of the fruit to determine an estimate  
of weight from an estimate of volume. By this means, we obviate the need to add the  
extra complexity of a mechanical weighing device.

25 Most aspects of quality judgement are somewhat subjective. Whilst human experts can  
grade picked fruit reasonably consistently, it may be hard for them to articulate exactly  
what factors give rise to a particular quality label. Therefore, a useful innovation is to use  
quality labelling data provided by human experts to train a machine learning system to  
30 assign quality labels automatically to newly picked fruit. This may be achieved by training  
an image classifier with training data comprising (i) images of the picked fruit obtained by  
the QC hardware and (ii) associated quality labels provided by the human expert. A  
variety of models could be used to map the image data to a quality label such as a simple  
linear classifier using hand-crafted features appropriate to the type of fruit in question.

E.g. in the case of strawberries, appropriate features might be intended to capture information about geometric symmetry, seed density (which can indicate dryness of the fruit), ripeness, and surface folding. With enough training data, it would also be possible to use a convolutional neural network to learn a mapping directly from images to quality labels.

**Figure 9** shows a sectional view of a typical embodiment of a QC imaging chamber. Cameras (and possibly other sensors, such as cameras sensitive to specific (and possibly non-visible) portions of the EM spectrum including IR, (ii) cameras and illuminators that use polarised light, and (iii) sensors specific to particular chemical compounds that might be emitted by the fruit.) positioned around the walls of the cylindrical imaging chamber provide a view of every part of the surface of the picked fruit. The picked fruit is gripped by a suitable end effector (the one shown here is the hook design described in Appendix A) and lowered into the chamber by the Picking Arm (of which only the head is shown). A useful innovation is to create a ‘chimney’ to reduce the amount of ambient light entering the QC rig – this is a small cylinder on top of the imaging chamber’s aperture that blocks unwanted light from the sides. One difficulty associated with imaging fruit inside the QC imaging chamber is that fruit debris can collect inside the chamber. Therefore a valuable innovation is to equip the chamber with a base that can be pulled out and wiped clean after a period of use.

## 7. Storage Subsystem

The purpose of the Storage Subsystem is to store picked fruit for transport by the robot until it can be unloaded for subsequent distribution. Because some types of fruit can be damaged by repeated handling, it is often desirable to package the fruit in a manner suitable for retail immediately upon picking. For example, in the case of fruits like strawberries or raspberries, picked fruits are typically transferred directly into the punnets in which they will be shipped to retailers. Typically, punnets are stored in trays, with 10 punnets per tray arranged in a 2-by-5 grid. When all the punnets in each tray are filled, the tray is ready for removal from the robot and replacement with an empty tray.

Since some fruit can be bruised easily by vibration caused by the motion of the robot over the ground, a useful innovation is to mount the trays via a suspension system (active

or passive) designed to minimize their acceleration under motion of the robot over rough terrain.

5 Unloading full trays of picked fruit may necessitate the robot travelling to the end of the row – so it is advantageous for the robot to accommodate more trays to amortize the time cost of travelling to and from the end of the row over more picked fruit. However, it is also advantageous for the robot to be small so that it can manoeuvred and stored easily. Therefore a useful innovation is to equip the robot with tray-supporting shelves that extend outwards at each end but detach or rotate (up or down) out of the way to  
10 reduce the robot’s length when it needs to be manoeuvred or stored in a confined spaced.

Another useful innovation is to store trays in two vertically oriented stacks inside the body of the robot as illustrated in **Figure 3**. One stack contains trays of yet-to-be-filled  
15 punnets, the other stack contains trays of full punnets. The Picking Arm transfers picked fruit directly into the topmost trays. Once the punnets in a tray are filled, the stack of full trays descends by one tray depth to accommodate a new a tray, a tray slides horizontally from the top of the stack of yet-to-be-filled trays to the top of the stack of full trays, and the stack of yet-to-be-filled trays ascends by one tray depth to bring a new yet-to-filled  
20 tray within reach of the robot arm. This design allows the robot to store multiple trays compactly within a limited footprint – which is important in a robot that must traverse narrow rows of crops or be transported on a transport vehicle.

25 Refrigerating picked fruits soon after picking may dramatically increase shelf life. One advantage of the compact arrangement of trays described above is that the full trays can be stored in a refrigerated enclosure. In practice however, the power requirements of a refrigeration unit on board the robot may be greater than can be met readily by convenient portable energy sources such as rechargeable batteries. Therefore, another  
30 useful innovation is to use one of various means of remote power delivery to the fruit picking robot. One possibility is to use electrified overhead wires or rails like a passenger train. Another is to use an electricity supply cable connected at one end to the robot and at the other to a fixed electrical supply point. The electricity supply cable might be stored in a coil that is wound and unwound automatically as the robot progresses along crop rows and such a coil might be stored on a drum that is located inside the robot or at the

end of the crop row. As an alternative to delivering electrical power directly to the robot, a coolant liquid may be circulated between the robot and a static refrigeration unit via flexible pipes. In this case, the robot can use an internal heat exchanger to withdraw heat from the storage container.

5

**Figure 10** shows a diagram illustrating a space-saving scheme for storing trays of punnets within a picking robot. Trays of full punnets are stored in one stack (shown right), trays of empty punnets are stored in another (shown left). The Picking Arm can place picked fruit only in the topmost trays. Once the right-hand side topmost tray is full (A), the stack of trays of full punnets descends downwards, a tray slides sideways (B), and the stack of empty trays ascends (C).

10

Tray removal/replacement may be achieved by a human operator or by automatic means. A useful innovation is a means of drawing the human operator's attention to the need for tray replacement via the combination of a strobe light on the robot itself and a corresponding visual signal in the Management User Interface. A strobe light that flashes in a particular colour or with a particular pattern of flashes may be advantageous in allowing the operator to relate the visual signal in the UI to the specific robot that requires tray replacement or other intervention. Another useful innovation is the idea of using a small, fast moving robot to work in concert with the larger, slower moving picking robot. The small robot can remove trays (or full punnets) automatically from the picking robot and deliver them quickly to a refrigeration unit where they can be refrigerated for subsequent distribution.

15

20

25

For some types of fruit, it is common for the customer (supermarket etc.) to define requirements on the size and quality of fruit in each punnet (or in each tray if punnets). Typical requirements include:

- i. each full punnet has total weight within some allowable tolerance of a nominal value;
- 30 ii. less than some proportion of the fruits in each punnet differ in size by more than some threshold percentage from the mean; and
- iii. less than some proportion of the fruit in each punnet exhibits unusual shape or blemish.

Depending on the contract between the grower and the customer, punnets not meeting these requirements (or e.g. trays containing one or more punnets not meeting these requirements) may be rejected by the customer, reducing the grower's profit. The commercial requirements can be modelled by a cost function that is a monotonically decreasing function of the grower's expected profit from supplying a punnet or tray to a customer. For example, a simple punnet cost function might depend on a linear combination of factors as follows:

$$\text{Cost} = w_0 e + w_1 u + w_2 c + w_3 d$$

Where  $e$  means the excess weight of strawberries in the punnet compared to the target weight,  $u$  is an indicator variable that is 1 if the punnet is underweight or 0 otherwise, and  $c$  is a measure of the number of strawberries that are outside of the desired size range. Finally,  $d$  is a measure of how long it will take to place a strawberry in a particular punnet, which is a consequence of how far the arm will have to travel to reach the punnet. The weights  $w_i$  reflect the relative importance of these factors to profitability, e.g.  $w_1$  reflects the cost of a tray containing an underweight punnet being rejected, weighted by the risk that an additional underweight punnet will cause the tray to be rejected; similarly,  $w_3$  reflects the impact on overall machine productivity of spending more time placing strawberries in more distant punnets.

An interesting observation is that distributing the exact same picked fruits differently between punnets could give rise to a different total cost according to the cost function described above. For example, because meeting the punnet weight or other packaging requirements more precisely means that less margin for error is required, so that a greater number of punnets can be filled with the same amount of fruit, or because placing similarly sized fruits in each punnet reduces the likelihood that a tray will be rejected by the customer. Therefore, a useful innovation is a strategy for automating the allocation of picked fruit into multiple punnets (or a discard container) based on size and quality measures to minimize the statistical expectation of total cost according to the metric described earlier, i.e. to maximize expected profitability for the grower. Compared to human pickers, a software system can maintain a more accurate and more complete record of the contents of many punnets simultaneously. Thus, the robot can place picked

fruit in any one of many partially filled punnets (or discard picked fruit that is of insufficient size or quality). However, the task is challenging because:

- there may be room for only a limited number of partially filled punnets;
- as the punnets are filled up, the amount of space available for additional fruit is reduced;
- moving picked fruits from punnet to punnet is undesirable because it is time consuming and may damage the fruit; and
- the size and quality of yet-to-be-picked fruits is generally not known a priori, and so it is necessary to optimize over *possible* sequences of picked fruits and associated quality and size grading.

In a simple embodiment of the above idea, each successive picked fruit might be placed to maximize incremental cost decrease according to the cost metric described earlier. However, this greedy local optimization approach will not produce a globally optimal distribution of fruit. A more sophisticated embodiment works by optimizing over the expected future cost of the stream of yet-to-be-picked strawberries. Whilst it may not be possible to predict the size or quality of yet-to-be-picked strawberries, it is possible to model the statistical distribution over these properties. This means that global optimization of fruit placement can be achieved by Monte Carlo simulation or similar. For example, each fruit can be placed to minimize total cost considering (i) the known existing placement of strawberries in punnets and (ii) expectation over many samples of future streams of yet-to-be-picked strawberries. A probability distribution (Gaussian, histogram, etc.) describing the size of picked fruits and possibly other measures of quality can be updated dynamically as fruit is picked.

Note that the final term in the above cost function ( $w_{3,d}$ ) can be used to ensure that the robot tends to place larger strawberries in more distant punnets. Since punnets containing larger strawberries require fewer strawberries, this innovation minimizes the number of time-consuming arm moves to distant punnets.

Sometimes, the Storage Subsystem cannot place picked fruit into any available punnet without increasing the expected cost (i.e. reducing expected profitability), for example because a strawberry is too large to be placed in any available space, or because its quality or size cannot be determined with high statistical confidence. Therefore, another useful

innovation is to have the robot place such fruits into a separate storage container for subsequent scrutiny and possible re-packing by a human operator.

For fruits that are picked in bunches comprising several individual fruits on the same  
5 branch structure, e.g. table grapes or on-the-vine tomatoes, it may be important that  
none of the individual fruits is damaged or otherwise blemished, for example because a  
single rotten fruit can shorten the life or spoil the appearance of the entire bunch.  
Therefore, a valuable innovation is a two-phase picking procedure in which first the  
entire bunch is picked and second unsuitable individual fruits are removed from it. In  
10 one embodiment, this might work as follows:

1. A first robot arm picks the bunch by severing the stalk.
2. Visual inspection of the bunch is performed to determine the positions of any  
blemished individual fruits. During visual inspection, the first robot arm continues to  
hold the bunch for visual inspection.
- 15 3. A second robot arm trims blemished fruits from the bunch. A picking head similar to  
that used to pick individual fruits like strawberries singly can also be used to trim  
individual fruits from a bunch.

In another embodiment of this idea, a first robot arm might transfer the bunch to a static  
support for subsequent inspection and removal of unwanted individual fruits. By this  
20 means, it is possible to use only a single robot arm.

As well as deciding into which punnet (or other container) picked fruit should be placed,  
it may also be necessary or beneficial for a robot to decide where in the target punnet to  
place the fruit. A key challenge is to place picked fruit so as to minimize bruising (or  
25 other kinds of damage) due to collisions with the walls of floor of the punnet or with  
other fruit already in the punnet. In the context of fruit picking systems that work by  
gripping the stalk of the fruit, another challenge is to determine the height at which the  
fruit should be released into the punnet – too high and the fruit may be bruised on  
impact, too low and the fruit may be squashed between the gripper and the base of the  
30 punnet. Additionally, picked fruit doesn't necessarily hang vertically because the stalk is  
both non-straight and somewhat stiff. Therefore, a useful innovation is to measure the  
vertical displacement between the base of the fruit and the point at which it is gripped or  
the pose of the picked fruit in an end effector coordinate system, so that the fruit can be

released at the optimal height. This can be achieved by using a monocular or stereo camera to determine the position of the bottom of the picked fruit relative to the (presumed known) position of the gripper. A related innovation is to use an image of the punnet (obtained by the cameras in the Picking Head or otherwise) to determine the position of other picked fruit already in the punnet. Then the position or release height can be varied accordingly. For fruits such as strawberries that may be usefully held by their stalks, another useful innovation is to orientate the gripper such that the stalk is held horizontally before the fruit is released into the storage container. This allows the compliance of the section of stalk between the gripper and the body of the fruit to be used to cushion the landing of the fruit when it is placed into the container.

A related innovation is to position and orientate the fruit automatically to maximize visual appeal. This might be achieved, for example, by placing fruit with consistent orientation.

Because the robot knows which fruit was placed in each punnet, it can keep a record of the quality of the punnet. Therefore, a useful innovation is to label each punnet with bar code that can be read by the robot and therefore used to related specific back to the record of which fruit the punnet contains.

## 8. Mapping Subsystem

A human supervisor uses the Management User Interface to indicate on a map where robots should pick (see below). A prerequisite is a geo-referenced 2D map of the environment, which defines both (i) regions in which robots are free to choose any path (subject to the need to traverse the terrain and to avoid collision with other robots) and (ii) paths along which the robot must approximately follow, e.g. between rows of growing plants. The robot can pick from plants that are distributed irregularly or regularly, e.g. in rows.

A suitable map may be constructed by a human supervisor using the Mapping Subsystem. To facilitate map creation, the Mapping Subsystem allows a human operator easily to define piecewise-linear paths and polygonal regions. This is achieved by any of several means:

- Using geo-referenced aerial imagery and image annotation software. The UI allows the user to annotate the aerial imagery with the positions of the vertices of polygonal regions and sequences of positions defining paths, e.g. via a series of mouse clicks. When annotating the start and end points of rows of crops, an integer-based row indexing scheme is used to facilitate logical correspondence between the start and end locations.
- Using a physical survey device, the position of which can be determined accurately, e.g. via differential GPS. The user defines region boundaries by positioning the surveying tool manually, e.g. at a series of points along a path, or at the vertices of a polygonal region. A simple UI device such as a button allows the user to initiate and terminate definition of a region. The survey device may be used to define the physical locations of (i) waypoints along shared paths, (ii) the vertices of polygonal regions in which the robot can choose any path, (iii) at the start and end of a row of crops. The survey device may be a device designed for handheld use or a robot vehicle capable of moving under radio remote control.

In the context of farming, an important concern is that heavy robots may damage soft ground if too many robots take the same route over it (or if the same robot travels the same route too many times). Therefore, a useful innovation is to choose paths within free regions to distribute routes over the surface of the ground as far as possible. A tuneable parameter allows for trade-off between travel time and distance and the degree of spread.

## 9. Management Subsystem

The Management Subsystem (including its constituent Management User Interface) has several important functions:

- It allows a human supervisor to define which rows of crops should be picked using a 2D map (created previously using the Mapping Subsystem).
- It allows a human supervisor to set the operating parameter values to be used during picking and QC, e.g. the target ranges of fruit size and ripeness, the quality metric to be used to decide whether to discard or keep fruit, how to distribute fruits between punnets, etc.
- It facilitates the movement of robots around the farm.

- It divides work to be done amongst one or more robots and human operators.
- It controls the movement of the robot along each row of strawberries.
- It allows robots to signal status or fault conditions to the human supervisor.
- It allows the human supervisor immediately to put any or all robots into a powered down state.
- It allows the human supervisor to monitor the position and progress of all robots, by displaying the position of the robots on a map.

If there are multiple robots, then they collaborate to ensure they can move around in the same vicinity without colliding.

In case fully autonomous navigation of the robots around a site is undesirable for safety or other reasons, it may be desirable for robots to be capable of being driven temporarily under human remote control. Suitable controls might be made available by a radio remote control handset or by a software user interface, e.g. displayed by a tablet computer.

To obviate the need for the human operator to drive several robots separately (e.g. from a storage container to the picking site), a valuable innovation is a means by which a chain of several robots can automatically follow a single 'lead' robot driven under human control. The idea is that each robot in the chain follows its predecessor at a given approximate distance and takes approximately the same route over the ground.

A simple embodiment of this idea is to use removable mechanical couplings to couple the second robot to the first and each successive robot to its predecessor. Optionally a device to measure the direction and magnitude of force being transmitted by a robot's coupling to its predecessor might be used to derive a control signal for its motors. For example, a following robot might always apply power to its wheels or tracks in such a way as to minimize or otherwise regulate the force in the mechanical coupling. By this means all the robots in the chain can share responsibility for providing motive force.

More sophisticated embodiments of this idea obviate the need for mechanical couplings by using a combination of sensors to allow robots to determine estimates of both their absolute pose and their pose relative to their neighbours in the chain. Possibly a

communication network (e.g. a WiFi network) might be used to allow all robots to share time-stamped pose estimates obtained by individual robots. An important benefit is the possibility of combining possibly-noisy relative and absolute pose estimates obtained by many individual robots to obtain a jointly optimal estimate of pose for all robots. In one such embodiment, robots might be equipped with computer vision cameras designed to detect both absolute pose in the world coordinate system and their pose relative to their neighbours. Key elements of this design are described below:

- The robot is designed to have visually distinctive features with known position or pose in a standard robot coordinate frame. For example, visually distinctive markers might be attached to each robot in certain pre-determined locations. The markers are typically designed for reliable automatic detection in the camera images.
- A camera (or cameras) is (are) attached to each robot with known pose in a standard robot coordinate frame. By detecting the 2D locations in its own camera image frame of visually distinctive features belonging to a second robot, one robot can estimate its pose relative to that of the second robot (e.g. via the Discrete Linear Transformation). Using visually distinctive markers that are unique to each robot (e.g. a bar code or a QR code or a distinctive pattern of flashes made by a flashing light) provides means by which a robot can uniquely identify the robot that is following it or being followed by it.
- One or more robots in the chain also maintain an estimate of their absolute pose in a suitable world coordinate system. This estimate may be obtained using a combination of information sources, e.g. differential GPS or a computer-vision based Simultaneous Localization and Mapping (SLAM) system. Absolute position estimates from several (possibly noisy or inaccurate) sources may be fused to give less noisy and more accurate estimates.
- Inter-robot communications infrastructure such as a wireless network allows robots to communicate with each another. By this means robots can interrogate other robots about their current pose relative to the robot in front. Pose information is provided along with a time stamp, e.g. so that the moving robots can compensate for latency when fusing pose estimates.
- In a chain of robots, the absolute and relative position estimates obtained by all robots are fused to obtain a higher quality estimate of the pose of all robots.
- A PID control system is used by each robot to achieve a desired pose relative to

the trajectory of the lead robot. Typically, a target position for the control system is obtained by finding the point of closest approach on the lead robot's trajectory. The orientation of the target robot when it was previously at that point defines the target orientation for the following robot. Target speed may be set e.g. to preserve a constant spacing between all robots.

When picking, teams of robots may become dispersed over a large area. Because robots are visually similar, this may make it very difficult for human supervisors to identify individual robots. To allow the human supervisor to relate robot positions displayed on a 2D map in the software UI to robot positions in the world, a useful innovation is to equip each robot with a high visibility strobe light that gives an indication in response to a mouse click (or other suitable UI gesture) on the displayed position in the UI. Individual robots can be made more uniquely identifiable by using strobes of different colours and different temporal sequences of illumination. A related innovation is to direct the (possibly coloured) light produced by the strobe upwards onto the roof of the polytunnel in which the crops are being grown. This facilitates identification of robots hidden from view by tall crops or the tables on which some crops (like strawberries) are grown.

Because a single human supervisor may be responsible for multiple robots working simultaneously, it is useful if the UI exposes controls (e.g. stop and start) for each robot in the team. However, one difficulty is to know which remote control setting is necessary to control which robot. In a situation where an emergency stop is needed, therefore, the system is typically designed so that pressing the emergency stop button (for example on the supervisor's tablet UI) will stop all robots for which the supervisor is responsible. This allows the supervisor to determine which robot was which after ensuring safety.

While picking, it is possible that robots will encounter so-called fault conditions that can be only be resolved by human intervention. For example, a human might be required to remove and replace a full tray of picked fruit or to untangle a robot from an obstacle that has caused it to become mechanically stuck. This necessitates having a human supervisor to move from robot to robot, e.g. by walking. To allow a human supervisor to do this efficiently, a useful innovation is to use information about the position of the robots and the urgency of their fault conditions (or impending fault conditions) to plan the human supervisor's route amongst them. Route planning algorithms can be used e.g. to

minimize the time or that human supervisors spend moving between robots (and therefore to minimize the number of human supervisors required and their cost). Standard navigation algorithms need to be adapted to account for the fact that the human operator moves at finite speed amongst the moving robots.

5

## 10. Robot Control Subsystem

### 10.1. Overview

Whilst the robot is picking, the Robot Control Subsystem makes a continuous sequence of decisions about which action to perform next. The set of actions available may include

10 (i) moving the whole robot forwards or backwards (e.g. along a row of plants), (ii) moving the Picking Arm and attached camera to previously unexplored viewpoints so as to facilitate detection of more candidate fruit, (iii) moving the Picking Arm and attached camera in the local vicinity of some candidate fruit so as to refine an estimate of its pose or suitability for picking, and (iv) attempting to pick candidate fruits (at a particular

15 hypothesized position/orientation). Each of these actions has some expected cost and benefit. E.g. spending more time searching for fruit in a particular vicinity increases the chances that more fruit will be picked (potentially increasing yield) but only at the expense or more time spent (potentially decreasing productivity). The purpose of the Robot Control System is to schedule actions, ideally in such a way as to maximize

20 expected profitability according to a desired metric.

In a simple embodiment, the Robot Control Subsystem might move the whole robot and the picking head and camera in an alternating sequence of three phases. In the first phase, the arm moves systematically, e.g. in a grid pattern, recording the image positions

25 of detected fruits as it does so. In the second phase, the picking head and camera move to a position near each prospective target fruit in turn, gathering more image data or other information to determine (i) whether or not to pick the fruit and (ii) from what approach direction to pick it. During this second phase, the system determines how much time to spend gathering more information about the target fruit on the basis of a

30 continuously refined estimate of the probability of picking a suitable fruit successfully (i.e. picking success probability). When the estimated picking success probability is greater than some threshold, then picking should take place. Otherwise the control system might continue to move the arm until either the picking success probability is greater than the threshold or some time limit has expired. In the third phase, once all

detected fruits have been picked or rejected for picking then the whole robot might move along the row of plants by a fixed distance.

### **10.2. Total Position Control**

During picking, the Control Subsystem uses the Total Positioning Subsystem to move  
5 the whole robot amongst the plants and within reach of fruit that is suitable for picking,  
e.g. along a row of strawberry plants. Typically, the robot is moved in a sequence of  
steps, pausing after each step to allow any fruit within reach of the robot to be picked. It  
is advantageous to use as few steps as possible, e.g. because time is required for the robot  
to accelerate and decelerate during each step. Furthermore, because the time required to  
10 pick the within-reach fruit depends on the relative position of the robot to the fruit, it is  
advantageous to position the robot so to minimize expected picking time. Thus, a  
valuable innovation is to choose the step sizes and directions dynamically so as to try to  
maximize expected picking efficiency according to a suitable model.

15 In a simple embodiment of this idea, a computer vision camera might be used to detect  
target fruit that is nearby but outside of the robot's present reach. Then the robot can be  
either (i) repositioned so as to minimize expected picking time for the detected fruit or  
(ii) moved a greater distance if no suitable fruit was detected in its original vicinity.  
Additionally, a statistical model of likely fruit positions might be used to tune step size.  
20 The parameters of such a statistical model might be refined dynamically during picking.

### **10.3. Robot Arm Path Planning**

The Picking Arm and attached camera moves under the control of the Control  
Subsystem to locate, localize, and pick target fruits. To enable the arm to move without  
colliding with itself or other obstacles, a path planning algorithm is used to find collision-  
25 free paths between an initial configuration (i.e. vector of joint angles) and a new  
configuration that achieves a desired target end effector pose. Physics simulation based  
on a 3D model of the geometry of the arm and the scene can be used to test whether a  
candidate path will be collision free. However, because finding a collision-free path at  
runtime may be prohibitively time consuming, such paths may be identified in advance  
30 by physical simulation of the motion of the robot between one or more pairs of points in  
the configuration space – thereby defining a graph (or 'route map') in which the nodes  
correspond to configurations (and associated end effector poses) and edges correspond  
to valid routes between configurations. A useful innovation is to choose the cost (or

‘length’) assigned to each edge of the graph to reflect a weighted sum of factors that reflect the overall commercial effectiveness of the picking robot. These might include (i) the approximate time required, (ii) the energy cost (important in a battery powered robot), and (iii) the impact of component wear on expected time to failure of robot components (which influences service intervals, downtime, etc.). Then path planning can be conducted as follows:

1. Search over the nodes of the graph to find a configuration  $C_0$  that can be reached without collision by a linear move from the initial configuration  $C_i$ .
2. Search over the nodes of the graph to find a configuration  $C_1$  from which the target pose can be reached without collision by a simple linear move. Configurations corresponding to the target pose can be determined by inverse kinematics, possibly using the configuration associated with each candidate configuration as starting point for non-linear optimization.
3. Find the shortest path in the graph between nodes  $C_0$  and  $C_1$ , e.g. using Dijkstra’s algorithm or otherwise.

One limitation of this approach is that the graph can only be precomputed for known scene geometry – and in principle the scene geometry could change every time the robot moves, e.g. along a row of crops. This motivates an interesting innovation, which is to build a mapping between regions of space (‘voxels’) and the edges of the route map graph corresponding to configuration space paths that would cause the robot to intersect that region during some or all of its motion. Such a mapping can be built easily during physical simulation of robot motion for each edge of the graph. By approximating real and possibly frequently changing scene geometry using voxels, those edges corresponding to paths which would cause the arm to collide with the scene can be quickly eliminated from the graph at runtime. A suitable voxel-based model of approximate scene geometry might be obtained using prior knowledge of the geometry of the growing infrastructure and the pose of the robot relative to it. Alternatively, a model may be formed dynamically by any of a variety of means, such as depth cameras, ultrasound, or stereo vision.

In the context of fruit picking, some kinds of collision may not be catastrophic, for example, collision between the slow-moving arm and peripheral foliage. Therefore, another useful innovation is to use a path planning algorithm that models obstacles not

as solid object (modelled e.g. as bounding boxes), but as probabilistic models of scene space occupancy by different types of obstacle with different material properties, e.g. foliage, watering pipe, grow bag. Then the path planning algorithm can assign different costs to different kinds of collision, e.g. infinite cost to a collision with an immovable  
 5 object and a lower (and perhaps velocity-dependent) cost to collisions with foliage. By choosing the path with the lowest expected cost, the path planning algorithm can maximize motion efficiency, e.g. by adjusting a trade-off between economy of motion and probability collision with foliage.

#### ***10.4. Learning Control Policies by Reinforcement Learning***

10 Given an estimate of the approximate location of target fruit, the system can gain more information about the target (e.g. its shape and size, its suitability for picking, its pose) by obtaining more views from new viewpoints. Information from multiple views can be combined to create a more accurate judgement with respect to the suitability of the fruit for picking and the suitability of a particular approach vector. As a simplistic example,  
 15 the average colour of a target fruit in multiple views might be used to estimate ripeness. As another example, the best viewpoint might be selected by taking the viewpoint corresponding to the maximum confidence estimation of stalk orientation. However, obtaining more views of the target may be time-consuming. Therefore, it is important (i) to choose the viewpoints that will provide the most useful additional information for  
 20 the minimum cost, and (ii) to decide when to stop exploring more viewpoints and attempt to pick the target or abandon it. For illustration:

If a target fruit (or its stem) is partially occluded (by foliage, other fruits, etc.) then it may be valuable to move in the direction required to reduce the amount of occlusion.  
 25 Generally, it is desirable to find a viewpoint from which the whole fruit is visible without occlusion because such a viewpoint defines, via the back-projected silhouette, a volume of space in which the picking head (and picked fruit) can be moved towards the target fruit without colliding with any other obstacles.

30 If a target fruit is observed from a viewpoint that makes it hard to determine its pose (or that of its stalk) for picking purposes, then it may be valuable to move to a viewpoint from which it would be easier to determine its pose.

If the target fruit is positioned near other target fruits so that which stalk belongs to

which fruit is ambiguous then it may be valuable to obtain another view from a viewpoint in which the stalks can be more easily associated with target fruits

If approaching the fruit to pick it from the current viewpoint would require a time-consuming move of the Picking Arm (e.g. because some moves require significant reconfiguration of the robot arm's joint angles) then it would be desirable to localize the strawberry from a viewpoint corresponding to a faster move.

Sometimes multiple views will be required to determine with sufficient confidence that the fruit should be picked. Sometimes, it will be unambiguous but rather than obtaining more views it would be better to move on, e.g. leaving the fruit can be picked by human pickers instead. However, designing an effective control routine manually may be prohibitively difficult.

A strategy for doing this is to use reinforcement learning to learn a control policy that decides on what to do next. The control policy maps some state and the current input view to a new viewpoint. The state might include observations obtained using previous views and the configuration of the arm, which could affect the cost of subsequent moves.

To train a control policy via reinforcement learning, it is necessary to define a utility function that rewards success (in this case picking saleable fruit) and penalizes cost (e.g. time spent, energy consumed, etc.). This motivates the idea that a control policy could be trained whilst robots operate in the field, using high quality picking success information using their on-board QC rig to judge picking success. An interesting innovation is that multiple picking robots can be used to explore the space of available control policies in parallel, sharing results amongst themselves (e.g. via a communication network or central server) so that all robots can benefit from using the best-known control policy. However, one limitation of that approach is that it may take a great deal of time to obtain enough training data for an effective control policy to be learned. This gives rise to an important innovation, which is to use for training purposes images of the scene obtained from a set of viewpoints arranged on a grid in camera pose space. Such a dataset might be captured by driving the picking robot under programmatic control to visit each grid point in turn, acquiring a (stereo) image of the scene at each one. Using a training set acquired in this way, reinforcement learning of a control policy can be achieved by simulating the

movements of the robot amongst the available viewpoints, accounting for the costs associated with each movement. In simulation, the robot can move between to any of the viewpoints on the grid (under the control of current control policy), perform image processing on each, and decide to pick a target fruit along a hypothesized physical path.

5 It is not possible to be certain that picking a target fruit along a particular path have succeeded in the physical world. However, for some target fruits, merely identifying the correct stalk in a stereo viewpoint gives a high probability of picking success. Therefore, we use correct identification of the stalk of a ripe fruit as a proxy for picking success in reinforcement learning. Ground truth stalk positions may be provided by hand for the

10 training set.

Central to reinforcement learning is some means evaluating the effectiveness of a particular control policy on the dataset. In outline, this is achieved as follows:

1. Set cost = 0
- 15 2. For each target fruit in dataset
3. Start at a randomly selected nearby viewpoint on grid
4. Update state (includes current pose and suitability for picking estimate) using current view
5. Use current control policy to map state to action (in {Move, Pick, Abandon})
- 20 6. Switch(action)

Move:

- i. Move Picking Arm to new viewpoint (determined by current policy)
- ii. Increase cost by cost of move (function of time taken, power consumed, etc.)
- 25 iii. Goto 3

Pick:

- i. Increase cost by cost of moving Picking Arm along picking trajectory (determined by current pose estimate)
- ii. If picking was successful, decrease cost by value of successful pick
- 30 iii. Goto 1 and select next target

Abandon:

i.Goto 1 and select next target

Using this cost evaluation scheme, we can compare the effectiveness of multiple control policies and select the best, e.g. by exhaustive search over available policies.

### 10.5. *Holistic Robot Control*

5 The reinforcement learning strategy described above relates to control policies for localizing a fruit and determining its suitability for picking given an approximate initial estimate of its position. Note however that the same approach can also be used to train a holistic control policy for the whole robot. This means expanding space of available actions to include (i) moving the Picking Arm to more distant viewpoints (to find coarse  
10 initial position estimates for target fruits) and (ii) moving by a given amount along the row of crops. An additional innovation is to extend the reinforcement learning scheme to include actions carried out by human operators, such as manual picking of hard to reach fruit.

## 11. Miscellaneous Innovations

15 1. Because picking robots maintain a continuous estimate of their position in a map coordinate system, they can gather geo-referenced data about the environment. A useful innovation is therefore to have robots log undesirable conditions that might require subsequent human intervention along with a map coordinate and possibly a photograph of the scene. Such conditions might include:

- 20
- damage to the plant or growing infrastructure (e.g. caused by a failed picking attempts or otherwise);
  - the decision to leave ripe fruit unpicked because picking would incur too great a risk of failure or because the fruit is out of reach of the Picking Arm.

25 2. A useful and related idea is to have picking robots store the map coordinate system locations of all detected fruit (whether ripe or unripe) in computer memory. This makes possible several innovations:

a. One such innovation is to perform yield mapping to enable the farmer to identify problems such as disease or under- or over-watering early. Of interest might be  
30 e.g. the density of fruit production or the proportion of unripe fruits that subsequently mature into ripe fruits.

b. Another innovation is yield prediction. To pick ripe fruits in good time, picking robots must typically traverse every crop row every few days. By acquiring images

of ripe and unripe fruits, they can measure the size of individual target fruits as they grow and ripen. Since most unripe fruits will ripen in time, such data facilitates learning predictive models of future crop yield, e.g. in the next day, week, etc. A suitable model might map current and historic ripeness and size estimates for individual target fruits and weather forecast information (e.g. hours of sunlight, temperature) to a crop yield forecast.

5 c. Another related innovation reduces the amount of time that robots must spend searching for target fruit in the target detection phase of picking. During target detection, the system determines the approximate position of ripe and unripe target fruits. The system typically finds target fruits by moving a camera on the end of the Picking Arm(s) to obtain images from a wide range of viewpoints. However, spending more time searching for target fruits may increase yield but only at the expense of a reduction in picking rate. Therefore, a useful innovation is to store the map coordinate system position of unripe fruits that have been detected but not picked in computer memory so that a robot can locate not-yet-picked target fruits more quickly on a subsequent traversal of the crop row. In a simple embodiment, the position of previously detected but not yet picked fruits is stored so the robot can return directly to the same position on the subsequent traversal without spending time searching. In a more complex embodiment, previous detections are used to form a probability density estimate that reflects the probability of finding ripe fruit at a particular location in map coordinate system space (e.g. by kernel density estimation of otherwise). This density estimate might be used to help a search algorithm prioritize regions of space where ripe fruits are likely to be found and deprioritize regions where they are not. E.g. a simplistic search algorithm might obtain views from a denser sample of viewpoints in regions of space where ripe fruits are likely to be found. In a related innovation, yield prediction (as in (b) above) might be used to account of the impact on time on the ripeness of previously unripe fruits.

30 3. The picking robot can be equipped to perform several functions in addition to picking, including the ability to spray weeds or pests with suitable herbicides and pesticides, or to reposition trusses (i.e. stalk structures) to facilitate vigorous fruit growth or subsequent picking.

4. Using adjustable arms that can be repositioned to maximize picking efficiency for a particular variety of strawberry or phase of the growing season. Picking Arms can also be positioned asymmetrically to increase the reach of the two arms working in concert (see **Figure 5**).

5

5. In use, it is likely that a robot will lean to one side or other due to ground slope or local non-smoothness of the terrain. This effect may be exacerbated by the use of a suspension system (intended to reduce shock as the robot travels over bumps) because of the compliance of the suspension. An unfortunate consequence of leaning over is that the position of the robot's Picking Arms will not be positioned as designed with respect to the plants (they may be e.g. closer or further, higher or lower). This compromises performance because (i) pre-defined camera poses chosen in the chassis coordinate frame may no longer be optimal if the chassis is rotated and (ii) chassis-relative models of environment geometry (which are used to prevent collisions between Picking Arms and the environment) may also be wrong. One obvious strategy for compensating for the impact of lean on the environment geometry model is to dilate the environment geometry in 3D to provide some tolerance to error, however this reduces performance by compromising the available working volume for the arm (in narrow crop rows, space may already be at a premium). Therefore a valuable innovation is a system both to measure the degree to which the robot is leaning over and then to compensate for the degree of the lean by adapting models of the scene's geometry and camera viewpoints accordingly. The degree of lean might be measured directly using an accelerometer (which measures the vertical direction) or indirectly by measuring the position of a part of the robot (e.g. using a vector cable follower arm or GPS) in a coordinate frame based on the crop row. The lean may then be allowed for by applying an appropriate 3D transformation to pre-defined camera poses and environment geometry. Another means of correcting for lean induced by sloping but smooth terrain is to adjust dynamically the lateral position of the robot's tracks in the row so that the Picking Arms are closer to their design position despite the lean.

10

15

20

25

30

6. If a picked fruit is gripped by its stalk, then it may swing like a pendulum following repositioning by a robot arm. This may incur a productivity cost, since for some operations (e.g. placing the fruit into a QC imaging chamber or punnet) the pose of

the fruit must be carefully controlled to avoid collisions - which may necessitate pausing arm motion following repositioning until the amplitude of such oscillations decreases to an acceptable level. Therefore a valuable innovation is to use a means of damping to reduce the amplitude the oscillations as soon as possible. In one  
5 embodiment, damping is achieved passively, using a soft gripper. In another embodiment, damping is achieved actively by modulating the velocity of the robot arm's end-effector in such a way as to bring the oscillation to a stop as quickly as possible. An estimate of the mass and pendulum length of the strawberry can be used to design a deceleration profile (dynamically or otherwise) required to minimize the  
10 amplitude or duration of oscillation.

7. Many of the diseases and other defects that affect soft fruits (and reduce their quality grade) affect their visual appearance in the images obtained by our QC imaging chamber. The first image processing step is to segment in each view the  
15 body of the fruit from the calyx (and the background of the imaging chamber). Typically this is achieved by using a decision forest to label pixels. It may also be useful to segment achenes (seeds) at the same time as described below. Then we characterise the appearance of the body of the fruit using a variety of quality measures (which may be innovative in isolation or combination). Some features not  
20 previously mentioned include:

- a. **The spatial distribution of achenes (e.g.in strawberries) or drupelets (e.g. in raspberries).** For healthy fruits, achenes and drupelets are generally arranged quite regularly, i.e. the distances between neighbouring achenes and drupelets are generally similar locally.  
25 However, some diseases and other problems have the effect of disrupting this regular arrangement in the developing fruit. Therefore an innovative idea is to use a measure of the regularity of the spatial arrangement of achenes or drupelets as an indicator of the quality of the fruit. One approach is to first detect (in images obtained by our QC imaging  
30 chamber) the image positions of achenes or drupelets using computer vision and then assign a cost at each such point using an energy function that assigns lowest energy to regularly arranged points. Achenes can be detected by semantic labelling, e.g. by using a decision forest classifier to assign a table to each pixel. Drupelets may be detected e.g. by using a

point light source to induce a specular reflection on the shiny surface of the fruit and then detecting local maxima in image brightness. Then the sum of costs over points (or local clusters of points) can be used to give an indication of fruit health.

- 5           b.    **The colour of the achenes.** In e.g. strawberries the achenes become redder when the fruit becomes overripe and black when the fruit becomes rotten. Labelled training images of fruits exhibiting particular defects can be used to determine thresholds of acceptability for colour.
- 10           c.    **The colour of the flesh of the fruit** (excluding the achenes). This is a good indicator of under-ripeness, over-ripeness, and localised bruising. Again labelled training data can be used to determine thresholds of acceptability.
- 15           d.    **The 3D shape of the fruit.** An explicit 3D model of the fruit may be obtained by model fitting to the silhouette and image intensity information as described in the existing provisional application. However an 'implicit' 3D model of the shape may be obtained more simply by characterising the 2D shape in each of multiple views. An effective strategy for doing this proceeds one 2D view at a time by first determining the approximate long axis of the fruit and then by characterising shape in a coordinate system aligned with the long axis.
- 20           The long axis position may be estimated e.g. by the line joining the centroid of the body of the fruit and the image position of the gripper that is used to hold the fruit. Shape may be characterised e.g. by measuring the distance from the centroid of the fruit to the edge of the fruit in each of the clock face directions between 2 o'clock and 10 o'clock given a 9-element vector (noting that we ignore the top of the clock face to avoid the calyx). We can distinguish good (grade 1) shapes from bad (grade 2) shapes using a body of training examples with associated expert-derived ground truth labels. A suitable strategy is k nearest neighbours in the 9D shape space. Robustness to errors in long axis localisation can be achieved by generating multiple 9D shape vectors for randomly perturbed versions of the detected long axis in the training images.
- 25
- 30

8.    Typically, a picked fruit is considered grade 2 (sub-par) or grade 3 (reject) if any of

our several independent quality measures give a grade 2 or grade 3 score (although other means of combining defect scores are possible). Note that an important advantage of using image features specifically designed to detect specific kinds of defect instead of a more black-box machine learning approach is that doing so allows us to give the grower a clear and intuitive explanation (or, better, visual indication) for why a particular picked fruit received a particular quality grade. This allows the grower to adjust meaningful thresholds of acceptability for different kinds of defect. In practice, it is commercially very important for growers to make different decisions about the acceptability of different kinds of defect at different times of the season, considering the requirements of different customers, and as a function of productivity and demand.

9. Another application for agricultural robots is the targeted application of chemicals such as herbicides and pesticides. By using computer vision to locate specific parts of the plant or instances of specific kinds of pathogen (insects, dry rot, wet rot, etc.), robots can apply such chemicals only where they are needed, which may be advantageous (in terms of cost, pollution, etc.) compared to treatment systems that require spraying the whole crop. To facilitate doing this kind of work using a robot that is otherwise used for picking it would obviously be advantageous for the robot to support interchangeable end effectors, including one that can be used for picking and one that can be used for spraying. The latter would usually require routing pipes along the length of the robot arm, which may be difficult. However, a useful innovation is to have the spraying end effector contain a small reservoir of liquid chemical - thereby obviating the need for routing pipes. This might be achieved using a cartridge system where the robot arm would visit a station on the chassis to gather a cartridge of chemicals (which might be similar to the cartridges used in inkjet printers). Alternatively the arm might visit a cartridge in the chassis to suck the required liquid chemicals from a cartridge into its reservoir or to expel unused chemical from its reservoir back into a cartridge. It might be that several different types of chemical are combined in dynamically programmable combination to achieve more optimal local treatment, or that multiple cartridges contain multiple different chemical combinations.

## Appendix A: Produce-Picking End Effectors

This appendix describes several innovative Picking Head designs.

### *Background*

5 Picking fruit necessitates having an end effector (such as may be affixed to a robot arm) that is sufficiently small, strong, selective (so as to pick only fruit that is suitable for sale), exclusive (so as not to pick other fruit). The picking operation must not damage plants or growing infrastructure, e.g. grow bags. Additional constraints (low power, low cost, lightweight, durable) further constrain the design.

10

This invention solves these problems with a lightweight end effector capable of reliable picking. The compact nature of the designs considerably facilitates separating the desired items of fruit from unwanted items.

15 In general picking of produce comprises selection of the fruit to be picked, exclusion of items not intended to be picked (including unripe fruits and growing infrastructure), gripping of the fruit or its stem and separating the fruit from its parent plant. Here we disclose a number of embodiments of the invention to mechanically select, grasp, and cut produce from the host plant.

### 20 *Hook Embodiment*

The first embodiment of the invention selects and excludes with a hook that has a dynamically programmable trajectory. Under the control of the Picking Arm (or otherwise) the hook is mechanically swept through a (in general dynamically) chosen volume of space, and any stems within this swept volume are gathered into the  
25 hook. With the stem thus captured, the hook may be used to pull the target fruit away from the plant (and potential sources of occlusion like leaves or other fruits) so that measurements of picking suitability may be made (including visual, olfactory and tactile measurements). Such measurements inform a decision to pick or release the target fruit.

30 The hook might be made long and narrow (e.g. in the shape of the letter J) so as to minimize the volume that needs to be swept out as it is moved (e.g. along its own long axis) towards the target fruit, thereby minimizing the size of the gap needed (e.g. between leaves, stems, or other fruit) for the hook to reach the target fruit without collision.

Note that it may be advantageous to position the long axes of the hook nearly coincident with the optical axis of the Picking Head camera (or nearly in the middle of the optical axes of the two eyes of a stereo camera). Doing so simplifies the problem of finding a non-colliding path to target fruit because any target that appears from the camera's viewpoint to be unoccluded (by other fruit or foliage) can usually be safely approached by moving along the line corresponding to the ray between the camera's optical centre and the target. This reasoning obviates the need to obtain a 3D model of the environment to plan safe routes towards target fruit.

At this stage the hook motion may be reversed to release the produce without damage, or the gripper and cutter mechanism may be actuated to hold the produce and separate it from the stem prior to transport and release.

More detailed information about the mechanical aspect of the invention is now presented. Referring to **Figure 11**, the main features of the invention as assembled are shown. The main components of the end effector are specifically: (1) a hook, (2) a gripper that fits within the hook, (3) a lower support and (4) a blade. The hook is square in cross section and in conjunction with the blade forms a scissor cutting action. Screws and other support structures not shown.

**Figure 12** shows the hook extended relative to the grip/cut mechanism (prior to selecting or picking – blade omitted). **Figure 13** shows the hook retracted relative to the gripper/cutter, the gripper fitting against the hook to grip the plant stem and the cutting mechanism in its post actuation configuration (blade omitted). **Figure 14** shows an exploded view of the main parts in this embodiment, including the individual parts of the end effector including (4) the blade above the hook.

Referring to **Figures 12-14**, the main features are: (item 1) the hook (comprising a long thin section and a hook, which is semi-circular in this embodiment), which (through movement in space, **Figure 15** item 5) allows produce selection and exclusion of unwanted items. The inside of the hook is shaped to form one half of the gripper mechanism, and forms one half of the scissor style cutting surface. The tip is pointed in

this embodiment, maximising selectivity and exclusivity and assisting location within the gripper/blade mechanism when retracted.

5 There is a lower support (**Figures 11-13** item 3) which constrains the retracted hook, strengthening the gripping action and constraining the blade to be adjacent to the hook cutting surface, increasing cutting reliability. In this embodiment the gripper (2) is shaped to fit into the inside surface of the hook, allowing the mechanism to be maximally compact. This embodiment uses a flexure spring.

10 **Figure 15** is an example of the invention showing the movement of the hook (5) to effect a capture of the plant stem. The movement of the hook (5) captures the stem within the hook. Instead of having a cutting/gripping mechanism of fixed size, in the invention the hook is small (allowing good selection with maximal exclusion of unwanted stems) while the hook movement is of variable size (allowing selectivity in the presence  
15 of stem positional uncertainty).

Varying the capture volume (per item of produce) allows an optimum trade-off between selectivity (of the desired produce) and exclusivity (of unwanted items).

20 The grip and cut actions are performed in the same movement by actuation of the device (refer to **Figures 12 and 13**). As the hook retracts relative to the gripper, the stem is gripped first between the hook and the gripper (**Figures 11-13**, items 1 and 2). The gripper includes a spring, allowing for a range of stem size and allowing the hook to continue to retract. After the grip has been achieved, the continued movement of the  
25 hook against the blade separates the produce from the parent plant. Release of the fruit is achieved by extending the hook once more.

Construction method: Referring to **Figure 11**, in the present embodiment, the hook (1) is made of metal, e.g. steel. The gripper (2) is made of plastic (perhaps acetal if integrated  
30 flexure is required). The blade (4) is made of knife steel and the lower support (3) is plastic.

**Figure 16** shows the plant stem thus captured. **Figure 17** shows the produce gripped and cut from the parent plant (optional operation). **Figure 18** shows the release operation (optional operation).

5 **Figure 19** shows the sequence of operations that constitute the picking process. These are as follows:

- *Detect Produce* is performed by the Control and Computer Vision subsystems described in the main body of this document.
- *Approach* of the plant stem is made, locating the hook close to the plant stem of the produce desired to be picked.  
10
- *Select Item* is performed using a looping movement that is determined according to the location of the desired produce and unwanted produce and infrastructure. An example of this movement is shown in **Figure 15** item 5.
- *Decision Step 1* is performed with the selected item captured in the hook (as shown in  
15 **Figure 16**). With the fruit thus isolated (but still attached to the parent plant), a check (visual, olfactory, tactile) is made to detect the nature (and correctness) of the fruit. The result of this step is a decision to pick (produce for harvest or disposal) or don't-pick (unripe fruit, captured infrastructure, nothing captured, incorrect capture). Initial grading and quality control of the item is performed and stored and used if Decision Step 2 is reached.  
20
- *Release* is performed in the case of undesired items being captured by the hook, with a motion that is the reverse of capture (**Figure 18**).
- *Pick* is performed in the case of successful capture of produce ready to harvest. The hook is retracted relative to the gripper/cutter (transitioning from configuration as in  
25 **Figure 12** to configuration as in **Figure 13**). This effects a grip operation followed by a cut operation, separating the fruit from the parent plant (**Figures 16 and 17** illustrate before and after state).
- *Decision Step 2* is a secondary sensing operation performed with the fruit picked and transferred to another part of the machine. The result of this decision is to store or  
30 dispose of the fruit. This decision is performed with different sensors to that in Decision Step 1. This forms a more detailed assessment of the fruit thus picked. The result of this step (in conjunction with information from Decision Step 1) is a decision to store or dispose of the fruit. In the case of store, size and quality grading is used to determine the location of storage.

- *Store* the fruit is moved to the storage area (according to size/quality grading from Decision Steps 1 and 2) and released with the reversal of the picking operation (extending the hook relative to the gripper, releasing the fruit).
- *Dispose* is performed in the case of picked fruit that is mouldy or otherwise unsuitable for sale. The release operation is as per *Store* above, with the exception that the transfer is to a disposal area.

### ***Loop Embodiment***

In a second embodiment of the invention the select and exclude phases of the picking sequence described earlier is performed by means of actuating a loop e.g. of wire. The diameter, position and orientation of the loop is programmatically controlled and actuated such that it may:

- i. have an arbitrarily small volume on its approach to the target fruit;
- ii. increase in diameter, to be larger than the estimated diameter of the target fruit, as it is moved in parallel with, and centred on, the major axis of the target fruit and in the direction of the juncture between the target fruit and its stalk and;
- iii. have an arbitrarily small diameter once it has moved past the juncture of the stalk and target fruit.

In this way the loop will select the stalk of the target fruit from other objects in the environment (e.g. other produce, leaves, growing infrastructure etc.) the stalk can then be manipulated such that the target fruit may be moved away from other objects in the environment.

The fruit, if rejected after Decision Step 1, may be released by:

- i. increasing the diameter of the loop to at least the estimated diameter of the target produce; and
- ii. moving the loop in parallel, and centred on, the major axis of the target produce and away from the juncture of the stalk and target produce.

Once the produce is released a new picking operation may be started.

### *Jaw Embodiment*

In a third embodiment of the invention, the select, exclude, and grip phases of the picking sequence described earlier are performed by a set of jaws. The position, orientation and attitude (by which is meant whether opened, partially closed or closed) of the jaws are programmatically controlled and actuated such that

- i. they approach the stalk of the target fruit in the closed attitude so as to minimize their swept volume;
- ii. when arbitrarily near the stalk they are actuated to an open attitude;
- 10 iii. when adjacent to the stalk they are in a partially closed attitude so as to perform the select, exclude and grip phases simultaneously; the jaws may be further actuated to a closed attitude, such that a mounted blade is moved perpendicular to the stalk, cutting it using a scissor-like motion against the opposing side of the jaw.

15

The fruit, if rejected after Decision Step 1, may be released in the third embodiment of the invention by, following step (iii) in the above paragraph:

- i. actuating the jaw to an open attitude; and
  - ii. moving the jaw along the previous approach vector and in the opposite direction.
- 20 The produce is then released and the picking operation may be restarted.

### *Loop-and-Jaw Embodiment*

As shown in **Figures 24-27**, in a fourth embodiment of the invention, the second and third embodiments are combined such that a loop performs the select and exclude steps and the jaws perform the grip and the cut steps.

25

**Figure 20** shows the main mechanical constituents of the loop and jaw assembly (the actuation mechanism of the jaw is omitted for clarity).

30 **Figure 21** shows the loop and jaw assembly, shown with component 11, the loop, extended.

**Figure 22** shows an exploded diagram of the main constituent parts of the loop and jaw assembly (the loop is omitted for clarity).

**Figure 23** shows the components of the loop actuation mechanism. The loop is extended and retracted by means of the rotation of the drum (16). Both the drum and loop (11) sit within housing (12) and (17), to constrain the motion of the loop during actuation.

**Figure 24** shows the loop/jaw assembly on its approach vector towards the target fruit.

**Figure 25** shows the loop extended and the assembly moving in parallel to the major axis of the target produce and in the direction of the juncture between stalk and fruit.

**Figure 26** shows the loop having travelled past the juncture of the stalk and the fruit, the target produce is now selected and decision step 1 (**Figure 19**) may be applied.

**Figure 27** shows the loop is retracted to control the position of the target produce and the jaw is actuated so as to grab and cut the stalk of the fruit in a scissor-like motion.

More detailed information about the mechanical aspect of the loop-and-jaw embodiment of the invention are now presented. Figure 20 shows the main mechanical components of the loop and jaw assembly, being: (11) a loop which is sufficiently flexible in its plane of operation – being the plane perpendicular to the major axis of the target produce; (12) a housing for the loop to constrain the degrees of freedom of the loop to its plane of operation; (13) a blade mounted statically w.r.t one side of the jaw; (14) one side of the jaw, which may be sprung so as to allow stalks of various diameters to be gripped; (15) the other side of the jaw, which forms the opposing side of grabbing mechanism and the scissor-like cutting mechanism; (16) a drum, to which the loop is mounted, that by its rotation extends or retracts the loop; and (17) a housing for the drum which constrains the loop to move within the grooves of the drum.

These components (except 11 and 12) are displayed with greater clarity in the exploded diagram, **Figure 22**.

The loop itself (11) may be embodied in several different ways. In common they have the properties that they are flexible in the plane of operation, and are resistant to plastic deformation. This may be in the form of a single strand or multi-strand metal or plastic wire, a series of static links joined such that adjoining links pivot about a common axis,  
5 or a series of static links that overlap adjoining links and pivot about a flexible, continuous member that is attached to each link.

The housing for the loop (12) is shown as a hollow tube. It is mounted statically w.r.t the superstructure of the loop assembly. It has the properties that its inner surface has a low  
10 coefficient of friction and that it has a cross sectional shape that matches that of the loop. The housing may be made of plastic, rubber, or metal and may be reinforced with strands of metal wire and may be lined with another material. It may also be in the form of a groove cut within the superstructure of the loop assembly and may be of arbitrary cross-sectional shape.

15 Components 16 and 17 constitute one embodiment of the actuation of the loop. It is comprised of a drum (16) which rotates within a housing (17). Figure 18 shows detail of the drum and its assembly within the housing. The drum has two grooves in its outer circumference and the housing has matching grooves. When assembled these grooves  
20 constrain the movement of the loop to be static with respect to the drum and slide along the inner circumference of the housing. In this embodiment the housing is shown mounted perpendicular to the plane of operation of the loop. However, it may also be mounted in the plane of operation of the loop so as to accommodate embodiments of the loop that are inflexible in other planes (e.g. a series of pivoted links). The actuation  
25 may have other embodiments for example a linear actuation that pulls the loop to retract it and which is opposed by a sprung arm to extend it.

The jaws of the jaw assembly (14 and 15), may be constructed of acetal, for integrated flexure, or of less flexible material, either plastic or metal, and include a sprung and  
30 pivoted subcomponent or have a rubberised inner surface to provide compliance. This compliance is a necessary property of the jaw assembly to accommodate the gripping of stalks of varying diameter with a gripping force of similar magnitude. The jaws are mounted onto a superstructure about which they may pivot in 1 degree of freedom in the same plane. They may be actuated in a variety of ways, one embodiment of which is a

Bowden cable, another is by means of a servo motor driven toothed-gear. A blade (13) is mounted rigidly to the compliant jaw (14) but this compliance is not static w.r.t the blade. When the jaws are moved to the closed attitude the blade slides over the surface of the opposing jaw (15) creating a scissor-like cutting mechanism between them.

5

### *Miscellaneous innovations*

Variations of the hook movement (**Figure 15** item 5) may be used. E.g. rotation of the hook along its major axis by approximately 90 degrees allows the hook itself to be approximately parallel to the produce stem, allowing for greater selectivity amongst stems that are close to each other. This rotation is reversed before the gripping and cutting actuation.

10

To increase the speed of picking, an end effector holding picked fruit (or a part of the end effector that is responsible for holding the fruit) may be detached and transferred to another part of the machine, and another copy (or variation) of the end effector may be used to pick more produce. In this way, picking can work in parallel with the storage or disposal operations, thus increasing speed. Additionally, one of a variety of end effectors may be selected according to which is best suited to the task of picking a particular item of fruit to be picked.

20

The hook and gripper may capture the produce into a smaller gripped pallet that may be removed from the end effector. This comprises the gripper part (**Figure 11** Item 2) along with the hook, a subset of the hook, or an additional part (that may fit inside the hook as presented) that may be released from the machine.

25

The hook may be reconfigured to be a dual bifurcating hook, with gripper/cutter on each side. This allows the capturing hook movement to be clockwise or counter clockwise (a variation on **Figure 15** item 5).

30

The blade may be in two places: above and below the gripper (separately actuated). By twisting the hook through approximately 180 degrees before operation and using the appropriate blade, the hook movement may again be clockwise or counter clockwise while still allowing the blade to be above the gripper, effecting correct holding of the produce by the stem.

After picking, fruit may be lowered into an imaging chamber for grading purposes (see the Quality Control section, above). It is usually desirable that the Picking Head can be tilted downwards whilst holding the fruit in the imaging chamber (as illustrated in **Figure 9**) because otherwise the design of the imaging chamber would be compromised by the need to avoid mechanical interference with the Picking Head. Therefore, for end effectors that hold picked fruit by its stalk, a useful innovation is to have the surfaces that contact the stalk oriented at an angle to the vertical direction when the Picking Head is oriented such that there is no downwards tilt. This allows the picked fruit to hang vertically when the picking head is pitched downwards compared to the horizontal.

The hook may be reshaped to be square, triangular, or other shapes. Instead of a hook, a simple 'L' shape suffices and allows for easier release of the fruit or other items during the (optional) release stage (**Figure 19**, "Release").

The cutter may be replaced with a less sharp blade, for a more scissor-like rather than cutting action.

The gripper may be made of rubber (allowing use without a spring) or other materials, and positioned above or below the hook, although this is less optimal for compactness.

The cross section of the hook may be varied. Generally, a flat inside surface is preferred to ensure a reliable cutting action.

A related innovation is a hooking apparatus designed to allow the produce to be rotated for more complete inspection at the point of deciding whether to pick (**Figure 19**). For example by the mechanism of twisting its stem to allow the reverse of the produce to be imaged prior to committing to picking.

Should the Picking Arm collide with immovable objects (e.g. infrastructure) during picking, it may be able to stop automatically, for example by detecting that its intended position is different to its actual position. However, the collision may still cause damage to the arm or its end effector, or necessitate a time-consuming intervention by a human supervisor to move the arm safely away from tangled obstacles. The probability of

collision can be considerably reduced by designing the end effector and its motion path to minimize the volume of 3D space swept out whilst moving towards the target. However, it is still possible that the front of the end effector may collide with an immovable obstacle. In this event, the following innovations significantly help to reduce  
5 the likelihood of damage to the robot or to infrastructure and the requirement for human intervention:

- The end effector may be designed to deform under compressive force. The Loop design (above) self-evidently embodies this idea if the wire loop is sufficiently deformable. If the Hook design is used, the hook can be designed to buckle  
10 elastically or plastically if a sufficient compressive force is applied in the longitudinal direction.
- By approaching target fruit by moving the end effector approximately in the direction of the normal to the plane in which its cross-sectional area is minimized, the probability that any obstacles will collide first with the deformable hook instead of  
15 other non-deformable parts of the robot is increased.
- As an alternative to making the end effector deformable, a rigid end effector can be mounted to a spring so that it will move backwards into the picking head if enough force is applied. Furthermore, a microswitch may be used to detect backwards  
20 movement of the end effector so that the arm can stop moving (or reverse its direction of motion) as soon as a collision occurs. The length and stiffness of the spring should be calibrated so that the robot arm can stop harmlessly (and subsequently reverse) before excessive force is applied to the obstacle.
- For the Hook embodiment, another useful innovation is to retract the hook into its support when not actively picking to reduce the chances of snagging.

25

## Appendix B: A Rotary Cable Management System for a Robot Arm

In what follows, we describe an innovative solution to the problem of running various types of cable through the articulating joints of robot arms. Here, *cable* should be interpreted to mean any flexible object intended to guide matter or energy along its path, either as a means of providing power or transmitting information or moving material. This definition obviously includes (but is not limited to) electrical cables and wires, optical fibres, and pipes.

An important challenge here is to allow a wide enough range of angular motion at each joint. This is especially important in robots that do not merely repeat pre-programmed motion paths but instead determine dynamically where to move as a function of observations of the environment. In the former case, it is usually possible to design the motion paths so that the joints are never driven past their limits. But in the latter, the desired motion path might not be predictable in advance. Sometimes it may not be possible to move the robot arm directly from its current configuration to a desired target pose because the range of motion at one or more joints is insufficient. When this happens, it may be necessary to make a more complex ‘reconfiguration’ move so that the joint in question is oriented further away from its end stop. However, reconfiguration moves may be expensive in power or time, because they may require the robot to make large moves at all joints. Increasing the range of motion at the joints, decreases the probability that a reconfiguration move will be required for the robot to reach a new target pose.

### *Design Requirements*

Key design requirements for the rotary cable management system:

- Allows the joint to achieve large changes of rotation angle.
- Has high reliability (equates to low stress in the cable components and low reversal of stress – which accelerates fatigue and plastic yield).
- Capable of carrying complex cables – self-guided robots sense their environment and act based on this, meaning high data rates are required to transfer enough information for effective action, fast enough for efficient use of time. Typically, high

data rates require either twisted pair cables or fibre optic, both of which are particularly sensitive to twisting and coiling actions.

- Compact - self-guided robotic systems that interact with their environment need a small footprint to nimbly negotiate around the environment.
- 5 • Protected from the environment

Alternative solutions to these requirements are sub-optimal and include:

- WiFi (or other wireless communications channel). However, limited network bandwidth may make this difficult in environments where multiple robots are working nearby and it is not a solution for transmitting power or conducting matter.
- 10 • Externally guided cables – this requires an ‘umbilical’ that can get caught and damaged by the environment and severely limits the ability of the robot end effector to move with respect to the umbilical.
- Optical data transmission and inductive power transfer - high cost.
- Slip rings (very high cost and size for reliable versions).

15

### *Description*

**Figure 28** shows the different elements of the cable management system. The system consists of a cable enclosure (a) and a central cable guide (b) that twists relative to the enclosure (c), allowing a coil of cable to expand and contract much like a clock spring as the cable guide twists. The system may also be configured to support one or more coils of cable (d).

20

**Figure 29** shows drawings of the cable management system in situ within one of the joints of an arm. The one or more cables run through articulating joints of the robot arm. The cable guide is designed such that the cable is ducted away through the centre to the next stage of the arm. The cable is well supported by the cable guide, so only the well-defined coil of cable moves. This is illustrated in **Figure 30** with a sequence of drawings showing the cable guide rotating within the cable enclosure. This arrangement minimises cyclical stresses on the cable as it never reverse-bends, but merely bends slightly more or less, to accommodate the twist.

25

30

At one end, the twisting motion is limited by the coil pulling tight. At the other end, the cable unwinds itself to the point where the inner-most section of cable starts to rub on

the inside of the next coil and bend backwards (a form of capstan lock), rather than the whole coil continuing to unwind. The design must be arranged such that there is sufficient margin for error both in manufacture, assembly and operation, such that these limits are never reached, to prolong the life of the cable.

5

**Figure 31** shows a cutaway view of cable winding. When wound at one end of stroke the cable is pulled tight and there are more windings; and at the other the cable is pushed out to the edge of the enclosure and there are fewer windings. At both extremes (and between them) the change of curvature of the cable remains low, so the strain rate and fatigue seen by e.g. the copper and plastic components of the cable is low - giving a long life for large overall displacements.

10

Specific innovations:

- Use of the arrangement of a central cable guide and enclosure to define a coil of cable that can accommodate relative twisting of one to the other with low fatigue of the cable.
- Use of this arrangement with wires, electrical cable, optical fibres, fibre optic cables, pipes, ribbon cables, individual cores.
- Use of this arrangement especially with "flat" twisted pair cable, which is ideally sized to be flexible in one direction but self-supporting in the other direction, so forming a stable coil
- Multiple stacks of coils can be arranged on top of each other enabling many cables to be managed in a similar footprint, with the cable guide ducting all cables away through the central axis.
- The enclosure can be made to open and close easily to facilitate build of the cables. Among the many ways of doing this is to include a hinge in the enclosure, and make the enclosure out of many parts and build them around the coil.
- Shelves can be added between the cables to provide a smooth running-surface.
- Can lubricate the cables to minimise wear and friction; with added lubricant, or using inherently lubricious materials in the cable, or an added membrane coiled with the cable to lubricate the surfaces.
- Can also add an extra element to the central cable guide or inner-most section of cable to reduce the angle of contact with the next coil along and prevent capstan lock, thus increasing the stroke and reducing stress.

15

20

25

30

- 5 • Screened cables are often used in data lines to improve resistance to emission of noise; however screened cables are often less flexible and available in fewer configurations than unscreened. It is possible to benefit from the flexibility and wide variety of unscreened cables, while maintaining high integrity, by modifying this assembly to perform the screening function.
- 10 • Screening the assembly involves making the key components (enclosure, cable guide, shelves) from metal, conductively-filled material (e.g. carbon-loaded plastic) or conductively-coated material (e.g. metallised plastic), and earthing them; or by placing the whole assembly within a conductive shell. Effective screening is likely to require continuity of an earth between rotating halves of the assembly which can be implemented in various ways including an earthing conductor in the coiled cable; by adding conductive compliant surfaces to either the inner cable guide or enclosure, (e.g. brushes or conductive compliant seal elements); and capacitive coupling, by reducing the gaps between components to a minimum.
- 15 • At low temperatures, cables often suffer degradation faster due to materials nearing their glass transition temperatures. Low temperature operation of this cable management system can be achieved by using the cables themselves as heaters, running current through them to keep them warm.
- 20 • Can pump hot and cold fluid through pipes to control temperature.
- Can make some or all components transparent for ease of fabrication, inspection, and maintenance.

## Appendix C: Features Summary

This section summarises the most important high-level features; an implementation of the invention may include one or more of these high-level features, or any combination of any of these. Note that each feature is therefore potentially a stand-alone invention and may be combined with any one or more other feature or features; the actual invention defined in this particular specification is however defined by the appended claims.

10 The high level features are organized into the following categories:

- Robot hardware features or core robot features
- Operational optimization features
- End effector features
- Computer Vision features
- 15 • AI/machine learning features
- Picking process features
- Methods or applications

20 There is inevitably a degree of overlap between these features. This approach to organising the features is therefore not meant to be a rigid demarcation, but merely a general high level guide.

### **Robot hardware features or core robot features**

25 In this section, we summarise features which are robot hardware features or core robot features. The primary feature is a robotic fruit picking system comprising an autonomous robot that includes the following subsystems:

a positioning subsystem operable to enable autonomous positioning of the robot using a computer implemented guidance system, such as a computer vision guidance system;

30 at least one picking arm;

at least one picking head, or other type of end effector, or other type of end

effector, mounted on each picking arm to either cut a stem or branch for a specific fruit or bunch of fruits or pluck that fruit or bunch, and then transfer the fruit or bunch;

a computer vision subsystem to analyse images of the fruit to be picked or stored;

5 a control subsystem that is programmed with or learns picking strategies;

a quality control (QC) Subsystem to monitor the quality of fruit that has been picked or could be picked and grade that fruit according to size and/or quality; and

a storage subsystem for receiving picked fruit and storing that fruit in containers for storage or transportation, or in punnets for retail.

10

Whilst the primary application for this system is in picking strawberries, raspberries and tomatoes, this approach may be re-purposed outside of the fruit picking context. For example, it may be used for litter picking or collecting other kinds of items. One can therefore generalise the system as follows:

15

A robotic picking system comprising an autonomous robot that includes the following subsystems:

a positioning subsystem operable to enable autonomous positioning of the robot using a computer implemented guidance system, such as a computer vision guidance

20 system;

at least one picking arm;

at least one end effector, or other type of end effector, mounted on each picking arm to pick or collect an item, and then transfer that item;

a computer vision subsystem to analyse images of the item to be picked or

25 stored;

a control subsystem that is programmed with or learns picking or collecting strategies;

a quality control (QC) subsystem to monitor the item that has been picked/collected or could be picked/collected; and

30

a storage subsystem for receiving picked/collected items and storing that item in containers for storage or transportation.

There are multiple, optional features that can be used in such a system, or that could constitute a stand-alone feature, that can be used independently of the system defined

above. We list these as follows. Whilst we specifically reference a fruit picking system, all of the following features can be used outside of that context, for example for litter collecting or indeed collecting other items; generalizing beyond fruit is explicitly envisaged in all that follows, throughout this Appendix C.

5

A robotic fruit picking system that comprises a tracked or wheeled rover or vehicle capable of navigating autonomously using a computer vision-based guidance system.

10

A robotic fruit picking system in which the computer vision subsystem comprises at least one 3D stereo camera.

15

A robotic fruit picking system in which the computer vision subsystem that analyses fruit images comprises image processing software for detecting a fruit, and the control subsystem comprises software for deciding whether to pick the fruit and the optimal strategy for picking the fruit, based on automatically updateable strategies, such as reinforcement learning based strategies, including deep reinforcement learning.

20

A robotic fruit picking system in which the control subsystem automatically learns fruit picking strategies using reinforcement learning.

25

A robotic fruit picking system in which the picking arm has 6 degrees-of-freedom.

A robotic fruit picking system in which the picking arm positions the end effector and a camera, each mounted on the picking arm.

30

A robotic fruit picking system in which the end effector comprises a means of (i) cutting the fruit stalk or stem and (ii) gripping the cut stalk or stem to transfer the fruit to the QC and storage subsystems.

A robotic fruit picking system in which the robot automatically loads and unloads itself onto, and off of, a storage container or a transport vehicle.

A robotic fruit picking system in which the robot automatically navigates amongst fruit producing plants, such as along rows of apple trees or strawberry plants, including table grown strawberry plants, or raspberry plants.

- 5 A robotic fruit picking system in which the system automatically collaborates with other robotic systems and human pickers to divide picking work efficiently.

A robotic fruit picking system in which the system automatically determines the position, orientation, and shape of a target fruit.

10

A robotic fruit picking system in which the system automatically determines whether a fruit is suitable for picking based on factors which are automatically updateable in the quality control subsystem.

- 15 A robotic fruit picking system in which the end effector separates the edible and palatable part of a ripe fruit from its stem or stalk without contacting the edible part.

A robotic fruit picking system in which the system automatically grades a fruit by size and other measures of suitability that are programmed in to, or learnt by, the QC subsystem.

20

A robotic fruit picking system in which the system automatically transfers a picked fruit to a suitable storage container held in the storage subsystem without handling the edible and palatable part of the fruit or other sensitive parts of the fruit that could be bruised by handling.

25

A robotic fruit picking system in which the control subsystem minimises the risk of the end effector or other part of the robot damaging a fruit or plant on which the fruit grows using machine learning based picking strategies.

30

A robotic fruit picking system in which the picking arm moves an attached camera to allow the computer vision subsystem to locate target fruits and determine their pose and suitability for picking.

A robotic fruit picking system in which the picking arm is a light weight robotic arm with at least some joints that exhibit a range of motion of +/- 275 degrees that positions the end effector for picking and moves picked fruit to the QC subsystem.

- 5 A robotic fruit picking system in which the control subsystem operates the total positioning system and the picking arm.

10 A robotic fruit picking system in which the control subsystem uses input from the computer vision subsystem that analyses fruit images to decide where and when to move the robot.

15 A robotic fruit picking system in which the QC subsystem is responsible for grading picked fruit, determining its suitability for retail or other use, and discarding unusable fruit.

20 A robotic fruit picking system in which the robot picks rotten or otherwise unsuitable fruit (either by accident or design), and then discards that fruit into a suitable container within the robot or onto the ground, and that container is accessible via a discard chute with its aperture positioned at the bottom of the QC rig so that the arm can drop the fruit immediately without the need to move to an alternative container.

25 A robotic fruit picking system in which positive or negative air pressure is induced in a discard chute or an imaging chamber (e.g. using a fan) to ensure that fungal spores coming from previously discarded fruit are kept away from healthy fruit in the imaging chamber.

30 A robotic fruit picking system in which the system comprises one or more 6-axis light weight robotic picking arms with some or all joints that exhibit a range of motion of +/- 275 degrees.

A robotic fruit picking system in which the system comprises two or more picking arms and the picking arms are positioned asymmetrically on the robot.

A robotic fruit picking system in which the robot has tracks that are removable and the

robot can run on rails if the tracks are removed.

A robotic fruit picking system in which the robot is equipped with fruit holding trays that are suspension mounted.

5

A robotic fruit picking system in which the robot is equipped with fruit holding trays that are mounted on movable arms that move from a first extended position to a second, more compact position.

10 A robotic fruit picking system in which the robot is equipped with fruit holding trays arranged as two or more vertically oriented stacks.

A robotic fruit picking system in which the robot is powered from a remote power source.

15

A robotic fruit picking system in which the robot has one or more lights (e.g. strobe lights) that activate when a fruit tray or holder needs to be replaced.

20 A robotic fruit picking system in which a fast moving robot removes trays or holders automatically from a slower moving robot that does the fruit picking.

A robotic fruit picking system in which the robot has one or more lights (e.g. strobe lights) that activate in response to a user input and shine an identifying signal above the robot.

25

A robotic fruit picking system in which the system includes an imaging or analysis chamber in which fruit is placed by the picking arm and is then imaged or analysed for grading or quality control purposes.

30 A robotic fruit picking system in which the system includes an imaging or analysis chamber in which fruit is imaged or analysed for grading or quality control purposes and in which the imaging or analysis chamber includes an aperture and a chimney or cylinder or lid or baffle on top of the imaging or analysis chamber's aperture that is designed to

block unwanted light from entering the chamber whilst still permitting fruit to be lowered or passed into it.

5 A robotic fruit picking system in which the system includes an imaging or analysis chamber in which fruit is imaged or analysed for grading or quality control purposes and in which the imaging or analysis chamber includes one or more cameras and/or other sensors, such as cameras sensitive to specific (and possibly non-visible) portions of the EM spectrum including IR, (ii) cameras and illuminators that use polarised light, and (iii) sensors specific to particular chemical compounds that might be emitted by the fruit.)

10 A robotic fruit picking system in which the system includes a cable management system for cables that run through the articulating joints of a robot, the cable management system including a cable enclosure and a central cable guide that twists relative to the enclosure, allowing a coil or spiral of cable to expand and contract as the joints rotate.

15 The robot may include a picking arm made up of several individual rigid bodies, each attached to another rigid body at an articulating joint, and there is a cable enclosure associated with one or more of each of the articulating joints. The cable guide may be configured such that the cable is ducted away through the centre of the cable enclosure to the next body. The cable management system may be configured to minimise the change of local curvature of the cable as the articulating joints move through their full range of motion. The cables may be unscreened and the enclosure then provides screening. The cables may also serve to provide sufficient heat to reduce cable degradation

20 25 A robotic fruit picking system in which the picking arm is adjustable and can be repositioned to maximize picking efficiency for a particular crop variety or growing system, such as for the height of a specific table top growing system.

30 A robotic fruit picking system in which the system is configured to perform several functions in addition to picking, including the ability to spray weeds or pests with suitable herbicides and pesticides, or to reposition or prune trusses to facilitate vigorous fruit growth or subsequent picking.

A robotic fruit picking system in which the robot estimates its position and orientation with respect to a crop row by measuring its position and/or orientation displacement relative to a tensioned cable.

- 5 A robotic fruit picking system in which the robot estimates its position and orientation with respect to a crop row by measuring its displacement relative to a tensioned cable that runs along the row (a 'vector cable').

- 10 A robotic fruit picking system in which the robot includes one or more follower arms that are mounted to follow the robot.

A robotic fruit picking system in which the follower arm is connected at one end to the robot chassis by means of a hinged joint and at the other to a truck that runs along the cable.

- 15 A robotic fruit picking system in which the angle at the hinged joint is measured to determine the displacement relative to the cable.

- 20 A robotic fruit picking system in which the angle is measured from the resistance of a potentiometer.

A robotic fruit picking system in which two follower arms are used to determine displacement and orientation relative to the vector cable.

- 25 A robotic fruit picking system in which a computer vision guidance system measures the displacement of the robot relative to the vector cable.

- 30 A robotic fruit picking system in which the computer vision system measures the projected position of the cable in 2D images obtained by a camera mounted with known position and orientation in the robot coordinate system.

A robotic fruit picking system in which a bracket allows the vector cable to be attached to the legs of tables on which crops grow.

A robotic fruit picking system in which a truck is equipped with a microswitch positioned so as to break an electrical circuit with the truck loses contact with the cable.

5 A robotic fruit picking system in which magnetic coupling is used to attach an outer portion of the follower arm to an inner portion of the follower arm, such that in the event of a failure or other event the portions of the follower arm can separate without damage.

10 A robotic fruit picking system in which the separation of an outer portion and an inner portion of the follower arm triggers a control software to stop the robot.

### **Operational optimization features**

15 In this section, we summarise features which contribute to the operational effectiveness of the system.

A robotic fruit picking system in which the picking arm is controlled to optimize trade-off between picking speed and picking accuracy.

20 A robotic fruit picking system in which the control sub-system determines the suitability of a specific target fruit or bunch for picking via a particular approach trajectory by determining the statistical probability that an attempt to pick that target will be successful.

25 A robotic fruit picking system in which estimating the probability of picking success is determined from images of the scene obtained from viewpoints near a particular target fruit.

30 A robotic fruit picking system in which determining the statistical probability is based on a multivariate statistical mode, such as Monte Carlo simulation.

A robotic fruit picking system in which the statistical model is trained and updated from picking success data obtained by working robots.

- 5 A robotic fruit picking system in which the control sub-system determines the probability of collision between a picking arm and an object using an implicit 3D model of the scene formed by the range of viewpoints from which a target fruit can be observed without occlusion.
- 10 A robotic fruit picking system in which the control sub-system determines one or more viewpoints from which the target fruit appears un-occluded, and hence identifies an obstacle free region of space.
- 15 A robotic fruit picking system in which the computer vision subsystem uses a statistical prior to obtain a maximum likelihood estimate of the values of the shape parameters of a target fruit and the system then calculates an estimate of the volume of that target, and from that estimates the weight of the target.
- 20 A robotic fruit picking system in which the computer vision subsystem determines the fruit's size and shape.
- 25 A robotic fruit picking system in which the computer vision subsystem determines the fruit's size and shape as a means of estimating the fruit's mass and thereby of ensuring that the require mass of fruit is placed in each punnet according to the requirements of the intended customer for average or minimum mass per punnet.
- 30 A robotic fruit picking system in which picked fruit is automatically allocated into specific punnets or containers based on size and quality measures of the picked fruit to minimize the statistical expectation of total cost according to a metric of maximizing the expected profitability for a grower.
- 35 A robotic fruit picking system in which a probability distribution describing the size of picked fruits and other measures of quality is updated dynamically as fruit is picked.
- 40 A robotic fruit picking system in which the picking arm places larger strawberries in punnets that are more distant from the base of the picking arm, in order to minimize the number of time-consuming arm moves to distant punnets.

A robotic fruit picking system in which the picking arm places selected fruits in a separate storage container for subsequent scrutiny and re-packing by a human operator, if the quality control subsystem identifies those selected fruits as requiring scrutiny by a human operator.

5

A robotic fruit picking system in which the control subsystem implements a two-phase picking procedure for bunches of fruit, in which first the entire bunch is picked and second, unsuitable individual fruits are removed from it.

10 A robotic fruit picking system in which the robot measures the pose of the picked fruit, so that the fruit can be positioned at the optimal pose for imaging or analysis or for release at the optimal height to fall into a punnet or container.

A robotic fruit picking system in which the robot determines the position of other picked fruit already in a punnet or container and varies the release position or height into the punnet or container accordingly for new fruit to be added to the punnet or container.

15

A robotic fruit picking system in which the robot automatically positions or orients picked fruit in a punnet or other container to maximize visual appeal.

20

A robotic fruit picking system in which the robot automatically generates a record of the quality or other properties of a fruit in a specific punnet and adds a machine readable image to that punnet that is linked back to that record.

25 A robotic fruit picking system in which the robot chooses paths within free regions of the ground so as to distribute routes over the surface of the ground in a way the optimizes the trade off between journey time and damage to the ground.

A robotic fruit picking system in which a chain of several robots automatically follow a single 'lead' robot that is driven under human control.

30

A robotic fruit picking system in which information about the position of several robots and the urgency of any fault condition, or impending fault condition, affecting one or more robots is used to plan a human supervisor's route amongst them.

- 5 A robotic fruit picking system in which the position of the robot with respect to target fruit is controlled to optimize picking performance, such as minimizing expected picking time.
- 10 A robotic fruit picking system in which collision-free paths or obstacle free trajectories are identified in advance of run-time by physical simulation of the motion of the robot between one or more pairs of points in the configuration space, thereby defining a graph (or 'route map') in which the nodes correspond to configurations (and associated end effector poses) and edges correspond to valid routes between configurations.
- 15 A robotic fruit picking system in which a robot arm path planning is built by mapping between regions of space ('voxels') and edges of a route map graph corresponding to configuration space paths that would cause the robot to intersect that region during some or all of its motion,
- 20 A robotic fruit picking system in which the system logs undesirable conditions in the environment that might require subsequent human intervention along with a map coordinate.
- A robotic fruit picking system in which the system stores locations of all detected fruit (whether ripe or unripe) in computer memory in order to generate a yield map.
- 25 A robotic fruit picking system which the yield map enables a farmer to identify problems such as disease or under- or over-watering.
- A robotic fruit picking system in which the system stores a map coordinate system position of unripe fruits that have been detected but not picked in computer memory.
- 30 A robotic fruit picking system in which the yield map takes into account the impact on time on the ripeness of previously unripe fruits.

A robotic fruit picking system in which the system measures the degree to which the robot is leaning over and compensates for the degree of lean by adapting models of the scene's geometry and camera viewpoints accordingly.

A robotic fruit picking system in which the system includes an accelerometer.

5

A robotic fruit picking system in which the degree of lean is directly measured using the accelerometer or indirectly by measuring the position of a part of the robot in a coordinate frame based on the crop row.

10 A robotic fruit picking system in which an appropriate 3D to 3D transformation required to correct for the lean is applied to pre-defined camera poses and environment geometry.

A robotic fruit picking system in which the lateral position of the robot's tracks in the row is dynamically adjusted so that the picking arms are closer to their design position despite the lean.

15

A robotic fruit picking system in which oscillations of a fruit caused by picking the fruit are reduced by damping achieved by a soft gripper.

20 A robotic fruit picking system in which oscillations of a fruit caused by picking the fruit are reduced by damping achieved by modulating the acceleration or velocity or movement of the robot arm's end-effector.

A robotic fruit picking system in which the system estimates the mass and pendulum length of the fruit.

25

A robotic fruit picking system in which the system designs a deceleration or acceleration profile (dynamically or otherwise) required to minimize the amplitude or duration of oscillations.

30

**End effector features**

In this section, we summarise features which relate to the end effector; we refer to the end effector as the 'end effector'.

5 A robotic fruit picking system in which the end effector uses at least the following phases:

(i) a selection phase during which the target fruit is physically partitioned or separated from the plant or tree and/or other fruits growing on the plant/tree and/or growing infrastructure and

10 (ii) a severing phase during which the target fruit is permanently severed from the plant/tree.

A robotic fruit picking system in which the end effector includes a hook.

15 A robotic fruit picking system in which the fruit is moved away from its original growing position during the selection phase.

A robotic fruit picking system in which a decision phase is introduced after the selection phase and before the severing phase.

20 A robotic fruit picking system in which the decision phase includes rotation of the fruit by its stem or otherwise.

A robotic fruit picking system in which the decision phase is used to determine whether or not to sever the fruit, or the manner in which the fruit should be severed.

25 A robotic fruit picking system in which the selection phase is made reversible.

A robotic fruit picking system in which the reversibility is accomplished by a change of shape of the hook.

30 A robotic fruit picking system in which the reversibility is accomplished by movement or rotation of the hook.

A robotic fruit picking system in which the system is capable of simultaneously gripping and cutting the stem of a target fruit.

A robotic fruit picking system in which the system includes multiple picking units located on a single multiplexed end effector.

5 A robotic fruit picking system in which multiple picking functions on a picking unit are driven off a single actuator or motor, selectively engaged by lightweight means, such as: electromagnets, an engaging pin, rotary tab, or similar.

10 A robotic fruit picking system in which a single motor or actuator drives one function across all units on the head, selectively engaged by means such as: an electromagnet, an engaging pin, rotary tab, or similar.

A robotic fruit picking system in which the functions are driven by lightweight means from elsewhere in the system, such as using: a bowden cable, torsional drive cable/spring, pneumatic or hydraulic means.

15 A robotic fruit picking system in which the end effector pulls the target fruit away from the plant in order to determine the fruit's suitability for picking before the fruit is permanently severed from the plant.

20 A robotic fruit picking system in which the end effector comprises a hook with a dynamically programmable trajectory.

A robotic fruit picking system in which the end effector uses at least the following phases:

25 (i) a selection phase during which the target fruit is physically partitioned or separated from the tree and/or other fruits growing on the tree and/or growing infrastructure;

(ii) a severing phase during which the target fruit is permanently severed from the tree; and

30 in which the selection and severing phases are performed by the actuation of a loop, wherein the loop diameter, position and orientation are programmatically controlled.

A robotic fruit picking system in which the end effector uses at least the following phases:

(i) a selection phase during which the target fruit is physically partitioned or separated from the tree and/or other fruits growing on the tree and/or growing infrastructure;

(ii) a severing phase during which the target fruit is permanently severed from the tree; and

and in which the selection and severing phases are performed by a set of jaws, wherein the jaws diameter and position are programmatically controlled.

A robotic fruit picking system in which the jaw attitude such as opened, partially closed or closed is programmatically controlled.

### **Computer Vision features**

In this section, we summarise features which relate to the computer vision system used for autonomous navigation and the computer vision subsystem used for fruit imaging.

A robotic fruit picking system in which the computer vision based system is used in order to determine heading and lateral positions of the robot with respect to a row of crops using images obtained by a forwards or backwards facing camera pointing approximately along the row.

A robotic fruit picking system in which the computer vision based subsystem detects a target fruit and in which the robot includes an end effector wherein part of the end effector is used as an exposure control target.

A robotic fruit picking system in which the control system software uses lighting conditions inferred or derived from the weather forecast as an input to the control subsystem or computer vision subsystem to control picking strategies or operations.

A robotic fruit picking system in which the computer vision based subsystem detects a target fruit and an end effector is able to physically separate a candidate fruit further from the plant and other fruits in the bunch before picking.

- 5 A robotic fruit picking system in which mirrors are positioned and oriented to provide multiple virtual views of the fruit.

- 10 A robotic fruit picking system in which the computer vision based system obtains multiple images of a target fruit under different lighting conditions and infers information about the shape of the target fruit.

A robotic fruit picking system in which the computer vision based subsystem uses an image segmentation technique to provide an indication of a fruit health.

- 15 A robotic fruit picking system in which the computer vision based subsystem detects the positions or points of the fruit achenes or drupelets and assigns a cost to those positions or the arrangement of those positions using an energy function that assigns a lowest energy to regularly arranged positions.

- 20 A robotic fruit picking system in which a semantic labelling approach is used to detect achenes, such as a decision forest classifier.

A robotic fruit picking system in which the sum of costs over points provides an indication of fruit health.

- 25 A robotic fruit picking system in which the indication of the fruit health is provided from analysing one or more of the following: colour of the achenes, colour of the flesh of the fruit or 3D shape of the fruit.

- 30 A robotic fruit picking system in which a neural network or other machine learning system, trained from a database of existing images with associated expert-derived ground truth labels, is used.

A robotic fruit picking system in which the computer vision based subsystem is used to classify a fruit, and in which the system allows a grower to adjust thresholds for

classifying the fruit.

A robotic fruit picking system in which the computer vision based subsystem locates specific parts of a plant or specific plants that require a targeted localized application of chemicals such as herbicides or pesticides.

A robotic fruit picking system in which the computer vision based subsystem detects instances of specific kinds of pathogen such as: insects, dry rot, wet rot.

A robotic fruit picking system in which the computer vision subsystem detects drupelets or achenes using specularities induced on the surface of a fruit by a single point light source.

A robotic fruit picking system in which an end effector is used for picking and another end effector is used for spraying.

A robotic fruit picking system in which an end effector is a spraying end effector that contains a small reservoir of liquid chemical.

A robotic fruit picking system in which the picking arm visits a station on the chassis to gather a cartridge of chemicals.

A robotic fruit picking system in which the picking arm visits a cartridge in the chassis to suck the required liquid chemicals from a cartridge into its reservoir or to expel unused chemical from its reservoir back into a cartridge.

A robotic fruit picking system in which several different types of chemical are combined in dynamically programmable combination to achieve more optimal local treatment.

A robotic fruit picking system in which multiple cartridges contain multiple different chemical combinations.

**AI/machine learning features**

In this section, we summarise features which relate to AI or machine learning features.

A robotic fruit picking system in which a machine learning approach is used to train a detection algorithm to automatically detect a target fruit.

5

A robotic fruit picking system in which the system identifies fruit in RGB color images obtained by a camera.

A robotic fruit picking system in which the system identifies fruit in depth images obtained by dense stereo or otherwise.

10

A robotic fruit picking system in which the training data is a dataset in which the position and orientation of target fruit is annotated by hand in images of plants that are representative of those likely to be obtained by the camera.

15

A robotic fruit picking system in which a detection algorithm is trained to perform semantic segmentation on the images captured by the camera.

A robotic fruit picking system in which the semantic segmentation labels each image pixels such as ripe fruit, unripe fruit or other object.

20

A robotic fruit picking system in which a clustering algorithm aggregates the results of the semantic segmentation.

A robotic fruit picking system in which the machine learning approach is a decision forest classifier.

25

A robotic fruit picking system in which the machine learning approach is a convolutional neural network.

30

A robotic fruit picking system in which a convolutional neural network is trained to distinguish image patches containing a target fruit at their center from image patches that do not.

A robotic fruit picking system in which a sliding window approach is used to determine the positions of all images likely to contain a target fruit.

- 5 A robotic fruit picking system in which a semantic segmentation is used to identify the likely image locations of a target fruit for subsequent more accurate classification or pose determination by a CNN or other form of inference engine.

- 10 A robotic fruit picking system in which a machine learning approach with a regression model is used to predict the angles describing the orientation of approximately rotationally symmetric fruit from images, including monocular, stereo and depth images.

A robotic fruit picking system in which a machine learning approach is used to train a detection algorithm to identify and delineate stalks in images captured by a camera.

15

A robotic fruit picking system in which a machine learning approach is used to train a prediction algorithm to predict how much improvement to an initial pose estimate for a target fruit is likely to be revealed by a given additional viewpoint.

- 20 A robotic fruit picking system in which the system predicts which additional information, including which viewpoints out of a set of available viewpoints, is likely to be the most valuable, including the most beneficial to overall productivity.

A robotic fruit picking system in which the additional information is the location or point where the stalk attaches to the target fruit.

25

A robotic fruit picking system in which the additional information is the knowledge that the fruit is visible without occlusion from a particular viewpoint.

- 30 A robotic fruit picking system in which the additional information is the knowledge that the space between the camera and the fruit is free of obstacles from a particular viewpoint.

A robotic fruit picking system in which the system recovers a 3D shape of a target fruit from one or more images of the target fruit obtained from one or more viewpoints, and in which a generative model of the target fruit's image appearance is used.

- 5 A robotic fruit picking system in which a geometric and/or photometric model fitting approach is used to predict the surface appearance of a target fruit as well as the shadows cast by the target fruit onto itself under different, controlled lightning conditions.

- 10 A robotic fruit picking system in which the cost function, namely the measure of agreement between images, is made robust to occlusion or the fruit is physically separated from sources of occlusion.

- 15 A robotic fruit picking system in which a machine learning approach is used to train a labeling algorithm to automatically assign a label to an image captured by the system, wherein pre-labeled images provided by human experts are used to train the system.

- 20 A robotic fruit picking system in which labelling data provided by human experts is used to train a machine learning system to assign quality labels automatically to newly picked fruit, by training an image classifier with training data comprising (i) images of the picked fruit obtained by the QC subsystem and (ii) associated quality labels provided by the human expert.

A robotic fruit picking system in which the control policy subsystem is trained via reinforcement learning while the robot is operating.

- 25 A robotic fruit picking system in which the control subsystem is trained via reinforcement learning, and wherein training is done by simulating the movements of the robot using images of the real-world environment captured amongst available viewpoints.

- 30 A robotic fruit picking system in which the control system is trained to predict picking success via reinforcement learning, and in which training is done in a simulated picking environment.

A robotic fruit picking system in which the system is trained to predict picking success

via reinforcement learning, in which a predictor of picking success is to ensure that the predicted path of the end effector sweeps through a 3D volume that encompasses the target stalk but not other stalks.

- 5 A robotic fruit picking system in which the control subsystem is trained via reinforcement learning that includes actions carried out by human operators.

A robotic fruit picking system in which a machine learning approach is used to train a model to predict yield forecast.

10

A robotic fruit picking system in which the system records a map coordinate system location along with an image of all detected fruit and the recorded data is used to train a model to estimate a crop yield forecast.

- 15 A robotic fruit picking system in which the system uses picking success data obtained by working robots to learn and refine the parameters of a dynamically updateable statistical model for estimating picking success probability.

### **Picking process features**

20

In this section, we summarise features which relate to the picking process used by the system.

- 25 A robotic fruit picking system in which the system is operable to cut the stem of a fruit, in which the fruit is picked by first severing and gripping its stem, and the body of the fruit is removed from its stem in a subsequent operation.

30 A robotic fruit picking system in which the system is operable to sever the fruit from its stem using a jet of compressed air, without requiring the handling of the body of the fruit.

A robotic fruit picking system in which the robot includes a collar, shaped to facilitate forcing the body of the fruit off the stem.

A robotic fruit picking system in which the system is operable to sever the stem from its fruit by using the inertia of the body of the target fruit to separate the body of the fruit from its stem.

- 5 A robotic fruit picking system in which the system is operable to cut or sever the stalk of a fruit using a reciprocating back-and-forth motion of the fruit in the direction approximately perpendicular to its stalk or an oscillatory rotary motion with an axis of rotation approximately parallel to the stalk.

- 10 A robotic fruit picking system in which a path planning algorithm is used to model obstacles as probabilistic models of scene space occupancy by different types of obstacle with different material properties.

- 15 A robotic fruit picking system in which the end effector is operable to cut the stem of a fruit, in which the system includes a deformable end effector that is designed to deform under compressive force.

A robotic fruit picking system in which the system is operable to cut the stem of a fruit without handling the body of the fruit.

20

A robotic fruit picking system in which the robot is operable at night with a computer vision system that operates at night, and picks fruit when they are cooler and hence firmer to minimize bruising.

- 25 A robotic fruit picking system in which the end effector is operable to cut a fruit stem cleanly and without tearing to increase fruit productivity.

- 30 A robotic fruit picking system in which the quality control subsystem predicts the flavour or quality of a fruit and places the fruit in a specific storage container according to the flavour or quality prediction.

A robotic fruit picking system in which the prediction of the flavour or quality of a fruit depends on the analysis of a growth trajectory data measured over time for the fruit.

**Methods or applications**

In this section, we summarise features which relate to methods or applications of the system.

5 A method of optimizing fruit yield prediction by imaging each fruit using the robotic fruit picking system defined above, to determine ripeness or suitability for picking.

A method of optimizing fruit yield mapping across a fruit farm or multiple fruit farms, by imaging each fruit using the robotic fruit picking system defined above, to determine  
10 ripeness or suitability for picking.

A method of maximizing fruit shelf life by using the robotic fruit picking system defined above.

15 A method of selectively storing or punnetising the fruit with the optimal flavor or quality by using the robotic fruit picking system defined above.

A final aspect is the fruit when picked using the robotic fruit picking system defined above. The fruit can be strawberries, including strawberries grown on a table. The fruit  
20 can be raspberries. The fruit can be apples, or pears, or peaches, or grapes, plums, cherries, or olives or tomatoes.

Throughout this specification and the claims which follow, unless the context requires otherwise, the word “comprise”, and variations such as “comprises” and “comprising”,  
25 will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

The reference in this specification to any prior publication (or information derived from it), or to any matter which is known, is not, and should not be taken as an acknowledgment or admission or any form of suggestion that that prior publication (or  
30 information derived from it) or known matter forms part of the common general knowledge in the field of endeavour to which this specification relates.

**Note**

It is to be understood that the above-referenced arrangements are only illustrative of the application for the principles of the present invention. Numerous modifications and  
5 alternative arrangements can be devised without departing from the spirit and scope of the present invention. While the present invention has been shown in the drawings and fully described above with particularity and detail in connection with what is presently deemed to be the most practical and preferred example(s) of the invention, it will be apparent to those of ordinary skill in the art that numerous modifications can be made  
10 without departing from the principles and concepts of the invention as set forth herein.

Claims Defining the Invention

1. A robotic fruit picking system comprising an autonomous robot that includes the following subsystems:

- 5 a positioning subsystem configured to autonomously position the robot using a computer implemented guidance system, such as a computer vision guidance system;
- at least one picking arm;
- at least one picking head or other type of end effector, mounted on a picking arm to either cut a stem or branch for a specific fruit or bunch of fruits or pluck that fruit or
- 10 bunch, and then transfer the fruit or bunch;
- a computer vision subsystem to analyse images of the fruit to be picked or stored;

a storage subsystem for receiving picked fruit and storing that fruit in containers for storage or transportation, or in punnets for retail; and

- 15 in which the end effector is configured to (a) grip a stem or stalk of a fruit or a bunch of fruits and (b) cut that stem or stalk and/or pluck that fruit or bunch; the end-effector being controlled by the computer vision subsystem to separate the edible and palatable part from at least a part of the stem, branch or stalk without contacting the edible and palatable part;

- 20 and in which the system is configured to estimate the statistical probability that a picking attempt will be successful by taking into account one or more of the following: estimated pose and shape of the target fruit and it's stem or stalk, uncertainty associated with the recovered pose and shape estimates, colour of the target fruit's surface, the proximity of detected obstacles, and the range of viewpoints from which the target fruit
- 25 is visible, such that a fruit is picked when the estimated picking success is greater than a pre-defined threshold.

2. The robotic fruit picking system of Claim 1 in which the system further includes the following subsystem:

- 30 a quality control (QC) subsystem to monitor the quality of fruit that has been picked or could be picked and grade that fruit according to size and/or quality.

3. The robotic fruit picking system of any preceding Claim that further includes the following subsystem:

a control subsystem that is programmed with or learns picking strategies.

4. The robotic fruit picking system of any preceding Claim in which the end effector separates the edible and palatable part of a ripe fruit from its stem or stalk.

5

5. The robotic fruit picking system of Claim 1 in which the end effector maintains the edible and palatable part of a ripe fruit with at least part of the stem or stalk.

10

6. The robotic fruit picking system of any preceding Claim in which the system automatically transfers a picked fruit to a suitable storage container held in the storage subsystem such as to minimize handling the edible and palatable part of the fruit or other sensitive parts of the fruit that could be bruised by handling.

15

7. The robotic fruit picking system of any preceding Claim in which the picking arm moves an attached camera to allow the computer vision subsystem to locate target fruits and determine their pose and suitability for picking, in which suitability for picking is determined by estimating the statistical probability that a picking attempt will be successful by taking into account one or more of the following: the estimated pose and shape of the target fruit and its stem or stalk, uncertainty associated with the recovered pose and shape estimates, colour of the target fruit's surface, the proximity of detected obstacles, and the range of viewpoints from which the target fruit is visible.

20

8. The robotic fruit picking system of any preceding Claim in which the end effector uses at least the following phases:

25

(a) a selection phase during which the target fruit is physically partitioned or separated from the plant or tree and/or other fruits growing on the plant/tree and/or growing infrastructure and

(b) a severing phase during which the target fruit is permanently severed from the plant/tree.

30

9. The robotic fruit picking system of any preceding Claim in which the end effector includes a hook.

10. The robotic fruit picking system of Claim 8 – 9, in which the fruit is moved away from its original growing position during the selection phase.

5 11. The robotic fruit picking system of Claim 8 - 10 in which a decision phase is introduced after the selection phase and before the severing phase.

12. The robotic fruit picking system of Claim 11 in which the decision phase includes rotation of the fruit by its stem or otherwise.

10 13. The robotic fruit picking system of Claim 11 - 12 in which the decision phase is used to determine whether or not to sever the fruit, or the manner in which the fruit should be severed.

15 14. The robotic fruit picking system of Claim 8 - 13 in which the selection phase is made reversible so as to release the target fruit.

20 15. The robotic fruit picking system of Claim 14 in which the end effector includes a hook and the reversibility of the selection phase is accomplished by a change of shape of the hook.

16. The robotic fruit picking system of Claim 14 - 15 in which the end effector includes a hook and reversibility of the selection phase is accomplished by movement or rotation of the hook.

25 17. The robotic fruit picking system of any preceding Claim in which the system includes multiple end effectors located on a single multiplexed picking head.

30 18. The robotic fruit picking system of Claim 17 in which multiple picking functions on a end effector are driven off a single actuator or motor, selectively engaged by lightweight means, such as: electromagnets, an engaging pin, rotary tab, or similar.

19. The robotic fruit picking system of Claim 17 - 18 in which a single motor or actuator drives one function across all end effectors on the head, selectively engaged by means such as: an electromagnet, an engaging pin, rotary tab, or similar.

5 20. The robotic fruit picking system of Claim 18 - 19 in which the functions are driven by lightweight means from elsewhere in the system, such as using: a bowden cable, torsional drive cable/spring, pneumatic or hydraulic means.

21. The robotic fruit picking system of any preceding Claim in which the end  
10 effector pulls the target fruit away from the plant in order to determine the fruit's suitability for picking before the fruit is permanently severed from the plant.

22. The robotic fruit picking system of any preceding Claim in which the end effector comprises a hook with a dynamically programmable trajectory.

15

23. The robotic fruit picking system of any preceding Claim in which the end effector uses at least the following phases:

(a) a selection phase during which the target fruit is physically partitioned or separated from the tree and/or other fruits growing on the tree and/or growing  
20 infrastructure;

(b) a severing phase during which the target fruit is permanently severed from the tree; and

in which the selection phase is performed by the actuation of a loop, wherein the loop diameter, position and orientation are programmatically controlled.

25

24. The robotic fruit picking system of any preceding Claim in which the end effector uses at least the following phases:

(a) a selection phase during which the target fruit is physically partitioned or separated from the tree and/or other fruits growing on the tree and/or growing  
30 infrastructure;

(b) a severing phase during which the target fruit is permanently severed from the tree; and

and in which the selection and severing phases are performed by a set of jaws, wherein the jaws diameter and position are programmatically controlled.

25. The robotic fruit picking system of Claim 24 in which the jaw attitude such as opened, partially closed or closed is programmatically controlled.

5 26. The robotic fruit picking system of any preceding Claim in which the grip and cut phases are combined.

27. The robotic fruit picking system of Claim 26 in which in which the grip and cut phases are combined by means of exploiting the gripping action to pull the stem and/or stalk against a cutting blade or blades.

10

28. A method of selectively storing or punnetising the fruit with the optimal flavour or quality by using a robotic fruit picking system –comprising an autonomous robot that includes the following subsystems:

15 a positioning subsystem configured to autonomously position the robot using a computer implemented guidance system, such as a computer vision guidance system;

at least one picking arm;

at least one picking head or other type of end effector, mounted on a picking arm to either cut a stem or branch for a specific fruit or bunch of fruits or pluck that fruit or bunch, and then transfer the fruit or bunch;

20 a computer vision subsystem to analyse images of the fruit to be picked or stored;

a storage subsystem for receiving picked fruit and storing that fruit in containers for storage or transportation, or in punnets for retail; and

25 in which the end effector is configured to (a) grip a stem or stalk of a fruit or a bunch of fruits and (b) cut that stem or stalk and/or pluck that fruit or bunch; the end-effector being controlled by the computer vision subsystem to separate the edible and palatable part from at least part of the stem, branch or stalk without contacting the edible and palatable part;

30 and in which the system is configured to estimate the statistical probability that a picking attempt will be successful by taking into account one or more of the following: estimated pose and shape of the target fruit and it's stem or stalk, uncertainty associated with the recovered pose and shape estimates, colour of the target fruit's surface, the proximity of detected obstacles, and the range of viewpoints from which the target fruit is visible, such that a fruit is picked when the estimated picking success is greater than a

pre-defined threshold.




FIGURE 1




FIGURE 2




FIGURE 3




FIGURE 4




FIGURE 5




FIGURE 6




FIGURE 7




FIGURE 8

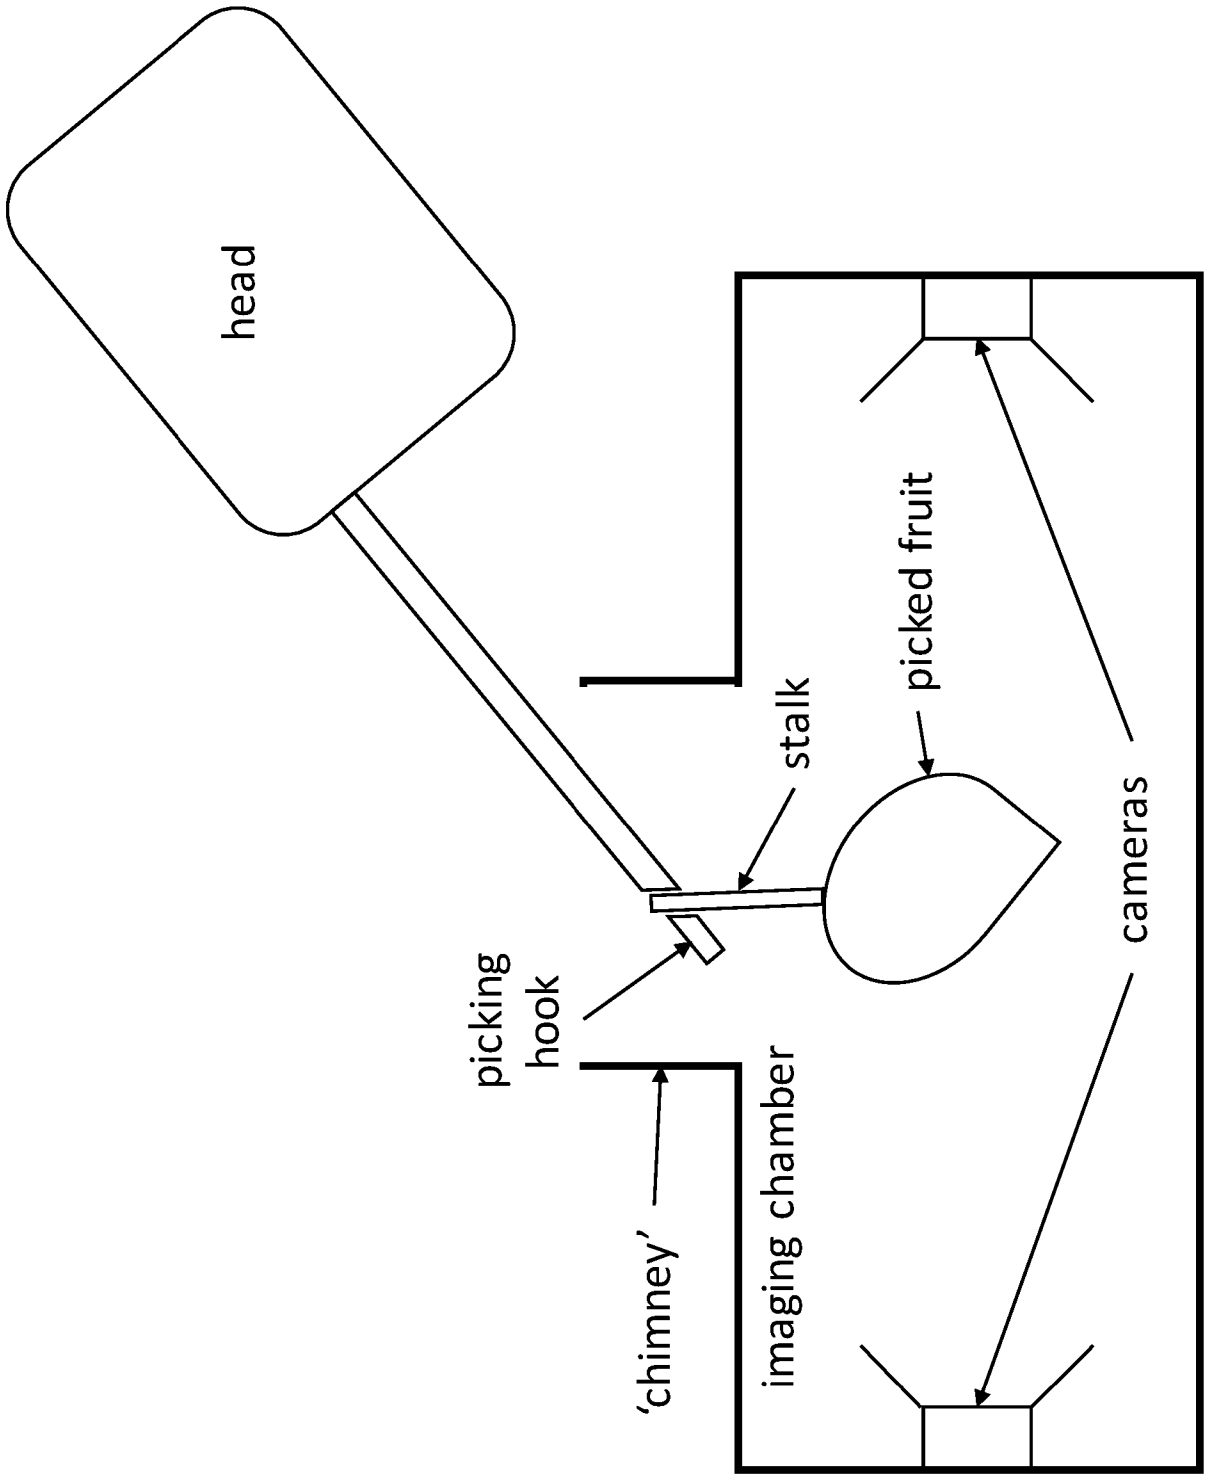


FIGURE 9

10/31




FIGURE 10




FIGURE 11




FIGURE 12




FIGURE 13




FIGURE 14




FIGURE 15




FIGURE 16




FIGURE 17




FIGURE 18

19/31




FIGURE 19




FIGURE 20




FIGURE 21




FIGURE 22




FIGURE 23




FIGURE 24




FIGURE 25




FIGURE 26




FIGURE 27






FIGURE 28

29/31



(b)

CABLE MANAGEMENT SYSTEMS  
ENABLE HIGHLY COMPACT ARRANGEMENT FOR  
FLEXIBLE MANIPULATION OF END EFFECTOR



(a)

FIGURE 29

CABLE GUIDE ROTATES WITHIN ENCLOSURE  
ABLE TO ROTATE THROUGH ANGLE GREATER THAN 360°.   
540° SHOWN HERE




FIGURE 30

31/31





FIGURE 31