

WASTE MANAGEMENT SYSTEM

TECHNICAL FIELD

5 The present invention relates, in general terms, to a waste management system, waste disposal system and network. More particularly, the invention relates to, but is not limited to, determining whether waste is being discarded into the correct one of a plurality of available bins.

10 BACKGROUND

Waste management, particularly recycling, is becoming increasingly important as the human population expands and attempts to deal with waste. Of more recent focus has been efforts on recycling.

15

Recycling can be labour intensive, particularly in the sorting phase during which waste objects are sorted into items that can be recycled and those that cannot. Some items look similar to other items – e.g., glass bottles and plastic bottles – which can make a difficult task of automated detection of waste streams.

20 Moreover, some packaging can be recycled but is disposed of in a condition where it cannot be recycled – e.g., a recyclable container containing food waste. This can further confound automated attempts at waste management.

It would be desirable to overcome or ameliorate at least one of the above-described problems, or at least to provide a useful alternative.

25

SUMMARY

Disclosed herein is an efficient and automatic add-on for in-situ waste disposal units (bins), in the form of a system for detection of general waste and contamination analyses of general recyclable waste. The system may similarly be used as a new waste disposal system, including bins, rather than as an add-

30

on.

Embodiments of the system are intended to be used in public places with large volumes of people, where general waste is discarded. During training phases,
5 human participation may be used for sorting, to generate ground truth data. A system trained on this data can be used to refine user behaviour and generate records such as waste statistics, barcode details etc. useful for waste managers and third parties.

- 10 As used herein, the term "general waste" will include recyclable waste and non-recyclable waste. Similarly, "recyclable waste" includes cans, paper and plastics (particularly recyclable plastics) waste, and "non-recyclable waste" refers to all other waste, or recyclable waste disposed of in a condition in which it cannot be recycled – e.g. a recyclable container containing food scraps, or a recyclable
15 material that forms a unitary article with non-recyclable material. Together, both recyclable waste in a condition for recycling and not in a condition for recycling will be referred to as "general recyclable waste".

- Disclosed is a waste management system for general waste stream (Cans,
20 Paper, Plastic, non-recyclables) comprising:
an image capture device for capturing an image of a waste object of a user;
at least one processor; and
memory storing instructions that, when executed by the at least one
25 processor, implement an inferencing pipeline for analysing the image to:
identify a waste stream of the waste object;
identify a specific type of the waste object; and
if the waste object belongs to a recyclable waste category, assess
acceptability of the waste object for recycling; and
30 a display for displaying the specific type and acceptability to the user.

Also disclosed is a waste disposal system comprising:

a waste management system as disclosed herein; and

a plurality of bins associated with the waste management system, each bin being for receiving one or more types of waste and including an appropriate
5 bin into which the waste object should be deposited, the display being configured to either:

identify to the user the appropriate bin; and/or

notify the user when the user has deposited the waste object into a said bin other than the appropriate bin.

10

Also disclosed is a waste disposal network comprising:

one or more waste disposal systems as described herein; and

a remote server,

each waste disposal system being configured to transmit transactions to the
15 remote server, each transaction comprising an identifier for one or both waste disposal system and bin into which the waste object was deposited.

Also disclosed is a method for managing waste, comprising:

receiving an image of a waste object of a user;

20 implementing an inferencing pipeline for analysing the image by:

identifying a waste stream of the waste object;

identifying a specific type of the waste object; and

if the waste object belongs to a recyclable waste category,

assessing acceptability of the waste object for recycling; and

25 displaying the specific type and acceptability to the user.

Advantageously, embodiments provide an affordable and green hardware system that can be deployed on existing bins as add-on. The solution is less
invasive and economical compared to existing systems as it is targeted towards
30 digitizing existing waste management infrastructure.

Advantageously, embodiments employ a modular software framework for the

general waste stream. The modular framework involves providing stages of classification, enabling each stage to be trained or retrained on data relevant to a specific locality. The solution can therefore be easily adapted to evolving market standards.

5

Advantageously, embodiments use human-in-loop functionality and avoid robotics. The solution involves less maintenance, promotes behavioural change by nudging the user to confirm to the standards of recycling general waste. The human-in-loop functionality involves using humans to perform most or all steps
10 involving physical movement – e.g. discarding of waste. The present system may then detect and identify the waste through image processing, without needing to physically move.

Advantageously, data collection affords the creation of dashboards for
15 developers and waste managers. The software is periodically updated using transaction feedback. The recorded waste statistics help track and report waste volume and lower operational costs.

BRIEF DESCRIPTION OF THE DRAWINGS

20

Embodiments of the present invention will now be described, by way of non-limiting example, with reference to the drawings in which:

Figure 1 illustrates a waste disposal system in accordance with present
25 teachings;

Figure 2 is a high level architecture of the present system;

Figure 3 is a schematic architecture of the backend subsystem employed by a waste disposal network in accordance with present teachings;

Figure 4 shows a welcome screen of a display of the system of Figure 1;

30 Figure 5 shows a waste object detection screen of a display of the system of Figure 1;

Figure 6 shows a user performing a recommended action displayed on a display

of the system of Figure 1;

Figure 7 shows a manual override screen of a display of the system of Figure 1;

Figure 8 shows waste object barcode scanning screen of a display of the system of Figure 1;

5 Figure 9a is a disposal screen displayed on a display of the system of Figure 1, for a single waste object;

Figure 9b. is a disposal screen displayed on a display of the system of Figure 1, for a composite waste object;

10 Figure 10 is a thank you screen is a disposal screen displayed on a display of the system of Figure 1, for a single waste object;

Figure 11 is a dashboard for third party analysis of usage of a waste disposal network; and

15 Figure 13 is a simplified architecture of a neural network architecture for implementing the system process performed by the waste disposal system of Figure 1.

DETAILED DESCRIPTION

20 Disclosed is an efficient and automatic add-on for in-situ detection of general waste and contamination analyses of general recyclable waste. Embodiments provide a computer-vision-based AI-IOT software system that can be deployed flexibly either on edge or cloud. The target for this system is general waste derived in public places with large volumes of people. The system improves user
25 compliance with recycling requirements, without employing automatic sorting using robotics to maximize upstream recyclable waste collection efficiency. Training involves creating a ground truth dataset using human participation to sort correctly. In practice, the process uses humans to attempt to categorise waste (e.g. place waste in the correct bin) and provides notification that the
30 waste was placed in the correct bin or incorrect bin. This nudges user behaviour towards correct waste disposal and generates records such as waste statistics, barcode details etc. useful for waste managers and third party analysis. The

present system combines with the other types of computer vision models that are geared towards waste classification using public datasets. However, the present system can be employed in a hybrid deployment (edge and cloud) scheme, model compression and pruning.

5

Such a waste management system 100 is schematically shown in Figure 1 and comprises:

an image capture device 102 for capturing an image of a waste object of a user;

10

at least one processor (component of system 104); and
memory (component of system 104) storing instructions that, when executed by the at least one processor, implement an inferencing pipeline for analysing the image to:

15

identify a waste stream of the waste object;
identify a specific type of the waste object; and
if the waste object belongs to a recyclable waste category, assess acceptability of the waste object for recycling; and
a display 106 for displaying the specific type and acceptability to the user.

20 **High level Architecture Description**

In some embodiments, the present system uses of an IoT subsystem, data store and machine learning models (models):

IoT subsystem

25

The IoT subsystem 200 (see Figure 2) is a wireless component with an image capture device (e.g. a digital camera 102), a display (e.g. touch screen display 106 for user I/O) and optional sensors (smart lid and smart base 108) for automatic opening of appropriate bin lid, for measuring weight and volume of the waste-objects as well as to detect the fill-level of the trash-bins. This
30 subsystem also encompasses one or more powerful but efficient processors (e.g. a processing engine) that are typically commodity Central Processing Units (CPU) augmented with custom hardware accelerators. The processing engine

also runs the inferencing pipeline, which is triggered whenever a waste-object is captured by the camera 102 based on user's input via the user interface.

Image capture

- 5 The image capture device may comprise a system of cameras to capture images from different angles. This reduces the likelihood of an image of the object (being discarded) being occluded or partially occluded. The camera or cameras may be any suitable cameras.

10 Display

- The display may simply display the specific type of the object, and the acceptability of the object – i.e. whether it was correctly disposed of. Correct disposal includes disposing of a recyclable object in a recycling bin, non-recyclable object in a non-recycling/waste bin, and a recyclable object in a non-recycling/waste bin where the recyclable object is not in a condition for recycling (e.g. a food or drink container that still includes some food or drink).

- In other embodiments, the display is a touch screen 106. A touch screen enables the user to correct an incorrectly classified object – i.e. where the specific type of the object was incorrectly determined. This can involve the user being presented a hierarchy of object categories through which the user traverses to find the specific type – e.g. recyclable or non-recyclable -> (if recyclable) paper or plastic or glass -> (if plastic) drink bottle or food container... and so on. This may also or instead involve providing a field for text input where the user can search for the specific type of the object.

- Where the user interface is a touch screen, two-way interaction between the present system and the user is afforded. The interface may include a set of displays, guiding the user through the process of recycling correctly using a waste management system as mentioned below. The user experience begins with a welcome screen to initiate scan and camera capture of waste object (Figure 4). It is followed by a waste detection screen displaying result of edge

inference for user validation (Figure 5). The user feedback is then used to evaluate the recycling acceptability of the waste object (Figure 6) or to manually override (Figure 7) in the event of inference failure such as those where AI is error-prone in detecting materials using only optical sensors, such as:

- 5 (i) Styrofoam or plastic food containers.
- (ii) Glass or Plastic beverage bottles, jars.

Recommendations displayed by the display 106 to gain recycling acceptability for user to follow are, for example:

- 10 (i) Empty and rinse bottles, cans, jars.
- (ii) Scan waste contents within plastic or brown paper bags individually.
- (iii) Empty and flatten cartons and Tetrapaks.

If the detected waste object belongs to recyclable waste streams (e.g. cans, plastic bottles), the subsequent screen may request for barcode scanning of the barcode of the object (Figure 8). This helps in storing additional product details such as material type, size, and manufacturer. Per step 214 of Figure 2, this screen also allows scanning of a user ID for incentives from third parties (if any) and for behaviour nudging in the future. The system uses edge inference and user inputs to provide disposal feedback (Figure 9a: Simple waste object detection, Figure 9b: Composite waste object detection) and trigger the opening of the optional smart-lid of the specific bin. Finally, a thank you (Figure 10) screen reports fate of the disposed object, disposal transaction statistics, user specific waste statistics and reward points earned (if any).

25

Inference

After image capture, a classification model is used to classify if the images contain recyclable general waste or unacceptable trash that needs to be incinerated. The neural network architecture of these models is chosen such that they are computationally less demanding which is a prerequisite to be ported on edge devices. This avoids the need for central processing (i.e. at a remotely located server) of all images. The classification model uses multiple pre-trained

30

machine learning models that are trained on multiple image sizes to be able to be flexible for deployment. The classification models are trained using a quantization aware strategy to enable on-edge inference. Thus, the feature space may be pruned down to those features that are most influential in
5 accurate object detection, waste stream identification and/or specific type identification.

The classification models are run by the inferencing pipeline 210. The inferencing pipeline (whether run as part of the IoT subsystem or simply run on
10 a processor located locally at the system, or remotely) is reflected in Figure 2. In using the inferencing pipeline 210, a user 201 disposes of an object (at 202) in one of the bins 204. In doing so, the user has provided the system 206 with a waste stream for the object – this waste stream may be correct or incorrect. At 208, the system 206 captures an image of the object, extracts the waste
15 stream (being the selected bin) and allocates a transaction identifier (ID) for the object. The inferencing pipeline 210 is then employed.

The inferencing pipeline 210 is a machine learning (ML) pipeline. The inferencing pipeline 210 determines the waste stream for the object, identifies the specific
20 object and then determines the acceptability of the object. To this end, the inferencing pipeline 210 implements an object detection algorithm that identifies the waste-object in the image.

Waste stream identification

25 Identifying the waste object may comprise, in an image captured by the image capture device, segmenting the waste object pixels from the non-waste object pixels. Segmentation can be performed in a known manner based, for example, on a proximity of the object relative to the image capture device – in general, the waste object will be closer to the image capture device than any other object
30 that is not part of the system itself.

In other embodiments, the object detection model is run first on the entire image with a high resolution to recognise waste streams which are allowed in the bin. The object detection model runs in sliding mode over the high-resolution image following the Slicing Aided Hyper Inference mode.

5

Whether or not segmentation is performed, the object is detected from the image based on a trained object detection model. The object detection model is trained to identify general waste streams (cans, paper, plastic, and non-recyclables) in a high-resolution image. The waste stream of the waste object
10 is thus one of a plurality of waste streams, and the inferencing pipeline comprises a classification model for each waste stream.

Various neural network architectures are available for waste stream identification. For present purposes, the neural network architectures are based
15 on the YOLO family of models and the EfficientDet model family. These architectures have been proven to be very versatile in real-time object detection and have also been trained on large open-source datasets. They have been customized for edge deployment using a pruning strategy for YOLO and a quantization strategy for EfficientDet. In these examples, the object detection
20 algorithm employed by the object detection model is a customized pre-trained Deep Neural Network (henceforth to be known as the detection model) to identify, in the captured image, a waste stream from one of cans, paper, plastic, and non-recyclables. As used herein, the Pruning and Quantization are strategies that balance size, complexity of the model and accuracy. Once the
25 detection model is developed, it is simplified further to determine the optimal network architecture with desired accuracy. The optimal network architecture requires less storage memory, has better processing speed. The strategy is dependent on the edge hardware being used and can vary according to the hardware. The skilled person will appreciate these how to design and/or
30 implement these strategies, in view of present teachings.

In some embodiments, waste stream identification may be performed using a model for each waste stream, that outputs a respective confidence score that the object belongs to the respective waste stream. The waste stream of the object will therefore be the waste stream corresponding to the model with
5 highest confidence. In other embodiments, a single model performs non-binary classification between all waste streams. In addition, or alternatively, where the waste stream identified is non-recyclable, the system may not then proceed to identify the specific type and instead display a notification on the display 106, advising whether the object was disposed of in the correct bin.

10

Specific type identification

After waste stream identification, another set of algorithms identifies the specific type of the waste object – e.g. plastic container, plastic bottle, receipt, note paper and others. The inferencing pipeline 210 of this embodiment therefore
15 identifies the specific type of the object in two steps – waste stream identification using a first machine learning model, then specific type identification using one of a plurality of classification models each being trained to classify objects from a respectively different waste stream.

20 For example, a single machine learning model may be used to identify the waste stream. A classification model is then selected to identify the specific type of the object, the classification model being selected based on the waste stream from a plurality of classification models each trained to identify the specific type of objects from a respectively different waste stream. In other embodiments, a
25 single waste stream may correspond to a plurality of classification models for classifying the specific type of the object, and each said classification model is run to produce a confidence score for classification of the specific type of the object. The specific type having highest confidence across all said classification models is then assumed to be the specific type of the object.

30

A family of image classification models is developed for specific type identification. The classification models are based on convolutional neural

network architectures that are tailored to mobile deployment on the edge and embedded applications but can also be deployed on remote cloud instances. The classification models classify the pre-recognised waste streams as individual waste objects (distinct categories of general recyclables and non-recyclables (if
5 desired)). The image classification models are fine-tuned and pruned to strike a balance between accuracy and computational complexity for each waste stream and individual or group of waste objects.

The classification models for identifying the specific type of the object may also
10 determine whether the object is contaminated – e.g. a clear plastic container may look dirty, indicating there may be food inside. Thus, the inferencing pipeline also determines the acceptability of the object for disposal in the bin into which it was discarded.

15 Acceptability

After identifying the specific type of the object, the system checks the acceptability of the object to be placed in the bin into which it was discarded. This may involve augmenting user feedback with results of acceptability detection from the system. The user feedback may be considered corrective –
20 i.e. if the user's feedback conflicts with the results of acceptability detection, the user's feedback may be assumed to be accurate, thereby producing data for re-training the classification model for identifying the specific type.

Acceptability can be determined using an acceptability model. The acceptability
25 model may detect one or more properties of the object to determine its acceptability. In some instances, where the classification model detects cleanliness (which may be limited by transparent objects) the acceptability model uses the output of the classification model to display to the user whether or not the object was appropriately disposed of – e.g. a clean recyclable
30 container in correct cycling bin is acceptable, an unclean (contaminated) recyclable container in any bin other than a general waste bin is unacceptable.

A weighing system in the smart base 108 (see Figure 1) may also be used to determine acceptability. After identifying the specific type using a classification model, the inferencing pipeline may extract (e.g. from memory in the system or from a remote database) a weight corresponding to the specific type when
5 empty – this may not be applied where the specific type is note paper or some other object that is not a container. The weighing system then weighs the object and the acceptability model determines if the weight of the object corresponds to the weight extracted by the inferencing model – e.g. the weight of the object is the same as, or is within a predetermined threshold of, the weight extracted
10 by the inferencing pipeline. If the weight of the object does not correspond to the weight extracted by the inferencing pipeline, then the object is not acceptable.

Thus the inferencing pipeline comprises an object detection model for identifying
15 the waste stream of the object (e.g. plastic, paper, can), and is configured to identify a specific type of the object by selecting the classification model for the waste stream of the object and using the selected classification model to determine the specific type (e.g. plastic container, plastic bottle, receipt, note paper) of the object. The contamination detection is yet another customized
20 pre-trained Deep Neural Network (henceforth to be known as acceptability model which may, for example, assess the acceptability of the waste object based on one or both of cleanliness of the waste object and emptiness of the waste object).

25 The inferencing pipeline is reflected in Figure 13, in which input images 1300 are received or captured, features are extracted by the waste stream identification model, and a waste stream is detected for the object (1302). A fully connected layer or other network or networks are employed to classify the object into a specific type (1304). Rule-based learning for acceptability
30 assessment is then used to determine if the object is acceptable for disposal in a particular bin, even if the object itself is deposited into a bin of the appropriate waste stream (1306).

The ML inferencing using Deep Neural Networks happens on the edge, so there is no heavy data transfer (like images and inference annotations) between the backend system and the IoT system. This significantly reduces the communication requirements, which could be a significant fraction of the operation cost.

Both the detection and acceptability models are compressed and pruned. Compressed models have the advantage of requiring much less power footprint than traditional "fat" models, thereby enabling their deployment in battery operated portable edge devices. The ML pipeline also includes the necessary logic to pre-process the input and post-process the outputs.

The machine learning models described above are, with the exception of analysing weight of the object, specific computer vision models for object detection and classification. A baseline models are pre-trained on large open-source datasets, which are then customized further for the data store described in the previous section and optimized for deployment on edge systems. Optimisation can be continually refined as more user feedback is obtained. Refinement may be performed on deep neural networks with large feature spaces, at a central server. The deep neural networks can then be pruned and compressed for edge deployment.

Display recommendation action or feedback

The recommended action of the waste-object is communicated to the user via user interface – e.g., the touch-screen display 106, or a display and associated keyboard – and opening of the optional smart lid triggered by ML inferencing on the edge. The user interface may also be configured to receive feedback from the user in relation to the waste object. For example, the touch screen will also collect additional information about the waste-object from the user, which would be helpful to refine the accuracy of both the detection and acceptability models. One or more labels may be applied to the image based on the feedback – this

helps build a dataset to be sent to a central server for updating global models and/or local models for waste stream identification and specific type identification, that are then pruned and compressed for deployment on the edge (i.e. on the systems in accordance with present teachings). The exact action of the user can be recorded by a combination of camera, optional weight sensor and optional fill sensors. Fill sensors detect fill level of the waste bin and alert the management (e.g. waste services provider or local council) to schedule collection when the fill level reaches a certain threshold. These sensors can be optical, infrared or ultrasound based. The IoT subsystem 200 subsystem can also accurately track the waste disposal behaviour of each user, if a user ID is either obtained through the user interface (e.g. touch screen 106) or is detected – e.g. through facial recognition. The users can receive both instantaneous feedback and their daily/weekly/monthly waste disposal behaviour which can nudge their waste-disposal behaviour.

The IoT subsystem communicates with a backend subsystem, wirelessly over cellular (like 5G/4G NB IoT)/wireless LAN(WiFi6) networks. The information communicated to the backend subsystem includes the captured images of the waste-object, the trash-bin identifier (which could be the identifier of the edge subsystem), the user identifier (optional) and the transaction identifier. Since most of the processing happens on edge, the data transfer volume per day is quite small. The on-edge processing also does not incur significant power consumption, due to use of compressed models in the ML pipeline.

To further reduce data transfer only a subset of images may be sent for use in training – e.g., images that are sufficiently clear, or sufficiently well-labelled. Thus, the system may collect an image for each waste object of a plurality of waste objects and, for each image, identify the image as an acceptable image for transmission to a remote server, based suitability of the image for training the models described herein.

Data store

A data store can be provided, whether in memory locally on the system 100, or in a remote server. The data store contains the images data and corresponding annotations that are required to train the machine learning models – i.e. the models for performing waste stream and specific type identification, and
5 potentially also acceptability assessment unless acceptability is determined based on a decision tree or decision structure – e.g. clean as determined by classification model Yes/No, or weight corresponding to weight extracted by inference model Yes/No. Suitable images, both in terms of quality and quantity, that are sufficient to build accurate machine learning models are challenging to
10 acquire. This is especially true in the case of models built for waste segregation since compared to the variety of real-world waste situations, the datasets are quite limited. Presently, proper datasets for object detection and classification models can be generated that can be used to train models to sort waste, deposited or discarded at public collection bins, in terms of non-recyclable and
15 recyclable general waste (cans, paper, plastic). Acceptable, recyclable general waste is then suitable for a corresponding upcycling pipeline. The dataset building method is applicable for two-dimensional views from the camera perspective. To train baseline machine-learning models, we have two methods for building the dataset. The methods can be described as:

20 (i) Synthetic: The method is based on creating datasets that utilizes self-acquired and multiple sources of existing image datasets. The existing datasets are collected and labelled under different open license terms. The datasets consist of waste images of cans, paper, plastic, and non-recyclable objects. The possible materials that are not useful for upcycling include cans (non-empty),
25 plastics (non-empty bottles and containers, straws, cutlery), paper (soiled paper, non-empty Tetrapaks). The following stages are used to generate the files (a) Background and waste object extraction using segmentation on waste images, (b) Data augmentation is then performed on background and waste object images using different techniques such as flip, rotate and shear,
30 brightness, contrast, saturation, colour and lighting changes, blurring, noise additions, (c) merging and overlaying the augmented backgrounds and waste object images. To emulate realistic conditions as closely as possible, different

blending methodologies are adopted including partial occlusion effects where waste objects are not completely visible. The blended images are then resampled or cropped and smoothened to remove heterogeneity of sources. The synthetic dataset is then re-annotated for both object detection (waste stream)

5 and classification (specific type).

(ii) In situ: the image capture devices are deployed in varying real-world environments, and the corresponding systems will begin to collect user images that are meant for machine learning inference and decision making. However, in an active learning framework using human-in-the-loop approach, user
10 feedback should be accounted for as a measure of accuracy of the model and generate classification labels. Since it is not possible for the user to label objects in detail, the second stage is about using semi-supervised learning to label objects for experts review and update. The images on the edges will be anonymized to a certain extent by removing background pixels that are not
15 relevant. The image comprising the remaining pixels together form the content of the live in situ database that will be used to continuously improve the machine learning models.

Waste Disposal system

20 The above embodiments are described with reference to a waste management system. A waste disposal system can include a waste management system as described above and, in addition, a plurality of bins associated with the waste management system. This is reflected by reference 212 in Figure 2. Each bin receives one or more types of waste. Clearly, for each object there will be an
25 appropriate one of those bins into which the waste object should be deposited. The display associated with the waste management system of the waste disposal system is configured to:

identify to the user the appropriate bin; and/or

30 notify the user when the user has deposited the waste object into a said bin other than the appropriate bin.

In some embodiments, the object is detected while being deposited into a bin. The display will therefore show whether the choice of bin was correct. In other embodiments, the waste disposal system will capture the image of the object prior to it being disposed of, and identify to the user the appropriate bin either
5 by displaying the appropriate bin on the display or by opening a lid over the appropriate bin. Either way, the user is notified if they deposit the object into the wrong bin.

Waste disposal network

10 While there may be a single waste disposal system used, in practice there will be multiple such systems. The waste disposal systems, together with a remote server, form a waste disposal network. This enables models to be updated far more rapidly due to data being captured at multiple edge systems (waste disposal system) rather than only on a single such system.

15 As reflected by 216 in Figure 2, each waste disposal system 212 is configured to transmit transactions 218 to the remote server (not shown). Each transaction includes an identifier for the waste disposal system 212 and/or waste stream of the object and bin into which the object was deposited. The transaction may
20 also include one or more of a user action, transaction ID and user ID. In some embodiments, only transactions with acceptable images are sent to the remote server, or the image is only sent when it is acceptable (i.e. acceptable for use in re-training machine learning models stored at the remote server). The term "user action" refers to user's response to the messages on UI screen or display.
25 It can include user's response/action to acceptability related instructions and response/action towards recommended colour code of the bin to deposit the waste, and other actions.

The server uses data supplied by the waste disposal systems 212 to update the
30 machine learning models that are then sent to the systems 212, at 220 in Figure 2 – this includes other images and other mechanisms for creating a ground truth dataset to be used to improve the detection model(s) offline. For any

transaction, the information sent to the system 212 may also include information on reward points for the transaction, the points being displayed to the user. The information may also include an output of an analytics system of the waste disposal network, that assesses waste disposal behaviour of the user
5 (where a used ID is provided) for each occasion the user deposits waste into a waste disposal system 212. This can provide the user feedback on the consistency or frequency with which they accurately identify the bin into which their waste objects should be deposited, and/or actions the user can take to improve their use of the bins.

10

Backend subsystem

The backend system (300 of Figure 3) hosts the data store for the images of waste and their corresponding annotations/labels. This data store is built from an initial set of waste images, thereafter it is continuously augmented with
15 operationally generated images it receives from multiple edge deployments. The initial annotation/labelling of images can be done manually. Thus the data store include initial images (e.g. images from public databases), captured images from waste disposal systems, synthetic images, user annotations and inferred annotations (e.g. annotations applied by machine learning models of waste
20 disposal systems).

The data store is connected to the ML training loop 310 through a data filter 314, that filters data depending on various factors such as the suitability of the data for training a particular model – e.g. only data relating to a particular waste
25 stream should be used to train classification models for that waste stream.

The backend system 300 may also store user information (304) such as a user ID and behavioural information (e.g. correct bins usage, disposal frequency, types of waste disposed of), and a transaction information store (306) for
30 storing information about all transactions across a waste disposal network.

- The backend system also hosts the detection, classification and acceptability models for general waste, stored in a model store (308). These models are full models (i.e. not pruned or compressed) in various of formats (TensorflowV2, Pytorch, etc.) which are amenable to continuous training. The backend also
- 5 incorporates the ML training loop (310) which uses the images in the datastore to automatically update the model parameters in presence of new data (e.g. new images from the waste disposal systems on the network). Besides detection and acceptability models, the backend also incorporates some generative models that are used to generate synthetic images of waste-objects for better
- 10 training performance. Periodically, the backend compresses a snapshot of the detection and acceptability model and pushes it to IoT subsystem via TCP/IP based protocol (like gRPC/MQTT/REST etc). Pushing may involve pruning and compression of the models.
- 15 Besides, the infrastructure for model training and data store. The backend also maintains an inferencing pipeline to identify barcode/brand/logo, to obtain trackable information from the captured waste object(s). This information is valuable to link the waste-object to the product manufacturer, enabling policies like EPR (Extended Producers' Responsibility) and other incentive structures.
- 20
- The backend receives the per trash-bin record of waste-object transactions and trash-bin status from the edge subsystem which can enable accurate accounting of waste-object generated on a daily/monthly/yearly basis.
- 25 The backend maintains a database of the signed-up users (user store 304). It receives the transaction records for each of the user in-order to better educate and nudge their recycling behaviour by rewarding them with a suitable incentive.
- 30 The backend 300 may also comprise an I/O interface or user interface 213 for cleaning data stored in the data store – e.g. accept/reject images, generate labels and other actions.

The backend system is deployed on on-premises or public cloud infrastructure, which can offer sufficient storage and compute capability for ML training and data storage.

5

Dashboard/Alert component

A dashboard interface (Figure 11) can be provided to allow sustainability stakeholders (property managers etc) to monitor their key performance indicators like the gross tonnage of trash collected, recycling rate, bin status and general waste disposal behaviour of the populace. The backend system maintains a log of each user-waste-object-bin transaction which could be used to present reasonably accurate analytics to the end user.

The dashboard 1100 can present aggregate statistics across a waste disposal network (1102), the locations of individual waste disposal systems in a waste disposal network (1104), the proportions of each waste stream deposited at each waste disposal system in a network (1106), the live bin status of each bin in each waste disposal system (1108) to enable bins to be emptied when needed, and analytics on proportions of recyclable and general waste produced by a population (1110) .

System Hardware Components.

The waste management system 206 of Figure 2 is an AI-IoT add-on that can be attached to existing general waste bins to form a waste disposal system. As explained previously, embodiments of the technology provide a computer vision-based AI software pipeline that can detect, classify, and analyze general waste and give on-site feedback to a user in real-time. The pipeline can operate in low-latency, power, and communication constrained environments. The key components of some embodiments of the waste management system are :

1. The Edge Module (112 in Figure 1) hosts the on-board computer, on which an AI software pipeline is deployed. The AI pipeline uses compressed and

pruned models which are specifically designed to run on resource constrained systems.

2. The Touch Screen accepts additional inputs from the user (if any) and displays the user feedback.
- 5 3. The Camera captures the video/image feed of the “waste” object.
4. The Aluminium Profiles (110 in Figure 1) provide the support framework to hold components in their positions.
5. The Smart Lid (optional) covers the opening of the bin and prevents the disposal of unacceptable waste items.
- 10 6. The Smart Base (optional) contain additional load cell sensors to monitor the weight of deposited objects.
7. The Fill Sensor (optional) detects the fill level of the bin and alerts the concerned staff when a certain threshold is reached.

15 **User interaction with the waste disposal system 212**

The following summarizes the waste disposal system workflow of some embodiments:

1. The waste disposal system is initiated when the user clicks on the welcome touch screen (Figure 4).
- 20 2. The system uses a low-power high resolution camera, for capturing an image of an object being disposed of, which is triggered when user initiates transactions after pressing the start screen.
3. A pre-processing logic (waste stream identification model) is applied to extract the general waste stream from the captured image.
- 25 4. The system deploys an ML inferencing pipeline with compressed models on energy efficient edge hardware. It consists of general waste stream detection model, waste stream specific object (specific type) classification and waste stream specific acceptability models.
5. The generic objection detection model recognizes the waste stream
- 30 (Figure 2). The general waste stream specific classification model detects the waste object (cans, notes, receipts, Tetrapak, plastic bottle, used tissues).

6. The waste detection screen displays detected waste object with user options to choose and approve – i.e. provide feedback. It can be challenging for AI to detect waste object material with absolute certainty using optical sensors, therefore in the event of inference failure, the user is also presented with
- 5 manual override (Figure 7) option. Once the detected object is approved for disposal, user is presented with a screen for evaluating waste acceptability (Figure 6) specific to the waste object such as empty and rinse bottles and cans, scan waste contents within plastic or brown paper bags individually, empty and flatten cartons and Tetrapaks.
- 10 7. Optional: If the detected waste object belongs to a recyclable stream, the subsequent screen requests the user for barcode scanning (Figure 8). The system includes an optional barcode detection model for extracting additional information on waste object from image data such as material type, size, and manufacturer. For the eligible classes (recyclable streams), the system (or a
- 15 user identifier collector of the system, which may be an input field on the touch screen display, an identification scanner or other system) allows the scan or input of user particulars (e.g. a user identifier) for incentives/rewards (by third parties, if any) and to keep track of user waste disposal behaviour for customized periodical nudging in the future. The user identifier is associated
- 20 with the transaction sent to the remote server.
8. Using the user inputs and waste object specific binary classification model (acceptability model) (Figure 2), the system displays the feedback of the appropriate bin (Figure 9a and Figure 9b) and triggers the opening of the optional smart lid for disposal of the detected waste object if present.
- 25 9. The thank you screen of the display 106 of the waste management system (Figure 10) reports transaction attributes that include fate of the disposed object, disposal transaction statistics, user specific waste statistics and reward points earned (if any).
10. The optional smart base (108) equipped with load cells measures gross
- 30 tonnage of trash collected, and optional fill sensors monitor bin status to alert the concerned staff via a live dashboard application (Figure 1).

1. The backend system maintains a log of each user-waste-object-bin transaction used to present recycling analytics to the property managers. The system also can trigger automatic emails with waste audit reports to concerned authorities at regular intervals.

- 5 12. All the inferences and evaluations take place at the edge module that hosts the onboard computer and only valuable images, and metadata are sent to the cloud at fixed interval to trigger software update.

The present systems enable waste object segregation at the source – i.e. point
10 of disposal – in a manner that improves user compliance for future disposal actions (i.e. transaction). Segregating at the source using sustainable and cost-effective solutions, results in:

1. More uniform waste categories upstream (cans, paper, plastic, and non-recyclables) that can be used in streamlining/optimization downstream
15 recycling/upcycling activities with reduced labour, operational and maintenance costs.
2. Government and Sustainability-linked financial grants, tax incentives for organizations that adopt green measures.

20 It will be appreciated that many further modifications and permutations of various aspects of the described embodiments are possible. Accordingly, the described aspects are intended to embrace all such alterations, modifications, and variations that fall within the spirit and scope of the appended claims.

25 Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise," and variations such as "comprises" and "comprising," 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.

30

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 information derived from it) or known matter forms part of the common general knowledge in the field of endeavor to which this specification relates.

CLAIMS

1. A waste management system for general waste stream comprising:
an image capture device for capturing an image of a waste object
5 of a user;
at least one processor; and
memory storing instructions that, when executed by the at least one
processor, implement an inferencing pipeline for analysing the image to:
identify a waste stream of the waste object;
10 identify a specific type of the waste object; and
if the waste object belongs to a recyclable waste category, assess
acceptability of the waste object for recycling; and
a display for displaying the specific type and acceptability to the
user.
15
2. The waste management system of claim 1, wherein the waste stream is
one of a plurality of waste streams, and the inferencing pipeline
comprises a classification model for each waste stream.
- 20 3. The waste management system of claim 2, wherein the inferencing
pipeline comprises an object detection model for identifying the waste
stream of the object, the inferencing pipeline being configured to identify
a specific type of the waste object by selecting the classification model
for the waste stream of the waste object.
25
4. The waste management system of any one of claims 1 to 3, wherein the
inferencing pipeline comprises an acceptability model for assessing
acceptability of the waste object based on one or both of cleanliness of
the waste object and emptiness of the waste object.
30
5. The waste management system of any one of claims 1 to 4, comprising
a user interface comprising the display, the user interface being

configured to receive feedback from the user in relation to the waste object.

- 5 6. The waste management system of claim 5, wherein the at least one processor is configured to apply one or more labels to the image based on the feedback.
7. The waste management system of any one of claims 2 to 4, wherein the processor is configured to:
- 10 collect an image for each waste object of a plurality of waste objects; and
- for each image, identifying the image as an acceptable image for transmission to a remote server, and based on the suitability of the image for training models for identifying a waste stream of the waste object,
- 15 identifying a specific type of the waste object, and assessing acceptability of the waste object for recycling.
8. A waste disposal system comprising:
- a waste management system according to any one of claims 1 to
- 20 7; and
- a plurality of bins associated with the waste management system, each bin being for receiving one or more types of waste and including an appropriate bin into which the waste object should be deposited, the display being configured to either:
- 25 identify to the user the appropriate bin; and/or
- notify the user when the user has deposited the waste object into a said bin other than the appropriate bin.
9. The waste disposal system of claim 8, further comprising a user identifier collector for obtaining an identifier of the user, the at least one processor associating the user identifier with data specifying whether the user deposited the waste object into the appropriate bin.
- 30

10.A waste disposal network comprising:

one or more waste disposal systems according to claim 8 or 9; and
a remote server,

5 each waste disposal system being configured to transmit transactions to the remote server, each transaction comprising an identifier for one or both waste disposal system and bin into which the waste object was deposited.

10 11.The waste disposal network of claim 10 when dependent on 8 when dependent on 7, wherein each said transaction corresponds to a said acceptable image.

12.The waste disposal network of claim 10 or 11, comprising a user analytics
15 system for assessing waste disposal behaviour of the user, for each occasion the user deposits waste into a bin associated with a said waste management system.

13.A method for managing waste, comprising:

20 receiving an image of a waste object of a user;
implementing an inferencing pipeline for analysing the image by:
identifying a waste stream of the waste object;
identifying a specific type of the waste object; and
if the waste object belongs to a recyclable waste category,
25 assessing acceptability of the waste object for recycling; and
displaying the specific type and acceptability to the user.

14.The method of claim 13, wherein the waste stream is one of a plurality
of waste streams and the inferencing pipeline comprises a classification
30 model for each waste stream, wherein implementing the inferencing pipeline comprises identifying the specific type of the waste object by:
selecting the classification model for the waste stream of the waste

object; and

identifying the specific type of the waste object using the classification model.

- 5 15.The method of claim 13 or 14, wherein assessing acceptability of the waste object comprises applying an acceptability model to assess acceptability of the waste object based on one or both of cleanliness of the waste object and emptiness of the waste object.
- 10 16.The method of any one of claims 13 to 15, further comprising receiving feedback from the user in relation to the waste object.
- 17.The method of claim 16, further comprising applying one or more labels to the image based on the feedback.
- 15 18.The method of any one of claims 13 to 17, further comprising:
collecting an image for each waste object of a plurality of waste objects; and
for each image, identifying the image as an acceptable image for
20 transmission to a remote server, based on the suitability of the image for training models for identifying a waste stream of the waste object, identifying a specific type of the waste object, and assessing acceptability of the waste object for recycling.