CompostNet: An Image Classifier for Meal Waste

Sarah Frost
Computational Media
UC Santa Cruz
Santa Cruz, United States
smfrost@ucsc.edu

Bryan Tor
Computer Science
UC Santa Cruz
Santa Cruz, United States
btor@ucsc.edu

Rakshit Agrawal
Computer Science
UC Santa Cruz
Santa Cruz, United States
ragrawa1@ucsc.edu

Angus G. Forbes
Computational Media
UC Santa Cruz
Santa Cruz, United States
angus@ucsc.edu

***Abstract*—Many businesses, cafes, and outdoor spaces provide trash, recycling, and composting bins, requiring consumers to decipher instructional text, icons, or images in order to sort their waste accurately. It can be confusing to know what pieces of waste go in which bin. Moreover, different areas may have different rules for how to separate waste, and people often inadvertently throw their trash in the wrong bin. Machine learning solutions can help us more quickly and accurately choose the proper receptacle for our waste by classifying a photograph of the waste. This paper presents a novel image classification model that categorizes the types of waste produced after eating a meal, which can be used in mobile applications to encourage users to correctly sort waste. Building on recent work in deep learning and waste classification, we introduce CompostNet, a convolutional neural network that classifies images according to how they should be appropriately discarded. We provide details about the design and development of CompostNet, along with an evaluation of its effectiveness in classifying images of waste. Further, we discuss two different approaches to the design of our system, one using a custom model and the other augmenting a pre-trained image classification model (MobileNet) through transfer learning, and and how we achieved greater success with the transfer learning approach. To the best of our knowledge, CompostNet is the first waste classification system that uses a deep learning network to identify compostable, recyclable, and landfill materials. CompostNet is an application of machine learning for social good, and supports United Nations Sustainable Development Goal 12: Responsible Consumption and Production [1].**

***Index Terms*—image classification, neural networks, waste management**

I. Introduction

A. Problem Definition

Junk, waste, rubbish, garbage, trash. Whatever its called, humans produce a lot of it. The United States generates 624,700 metric tons of solid waste each day [2]. Recycling and composting programs exist (see Figure 1), but are vulnerable to recycling contamination. One in four items placed in a recycling container are not actually recyclable [2] which contaminates the surrounding materials, making it unable to be processed. When individuals sort their waste after a meal, they may not know what is recyclable, what is trash, and what is compostable. When people are unsure about how to sort their waste, more of it will be misplaced, resulting in recycling contamination and otherwise recyclable material being sent to landfills.

Fig. 1. Waste receptacles at the National Institutes of Health in Bethesda, MD.

B. Motivation

This work is motivated by a concern about trash. Waste management and organization is a growing concern for many groups. Goal 12 of the United Nations Sustainable Development Goals is “Responsible Consumption and Production” [1], with target 12.5 aiming to “Substantially reduce waste generation through prevention, reduction, recycling and reuse” [1]. Similarly, the European Commission has an environmental policy that sets several priority objectives for waste policy [3]. Extensive research has been done to study waste management across the globe [4]–[7]. In this project, we use machine learning, as it can be used to classify images with a high rate of accuracy. Although efforts to improve waste sorting accuracy must be multifaceted, this system can be used at the first point of differentiating types of waste, and will help people learn how to correctly sort their waste.

C. Terminology

In this paper, we use the term ‘waste’ to refer to all material that is discarded. This term encompasses all of the materials that we are classifying. We use the labels ‘landfill’, ‘recyclable’, and ‘compostable’ to refer to the locations or processes that these materials will go to or undergo after they are correctly disposed. We also use the term ‘trash’ as a synonym to ‘landfill’.

There are different guidelines for recyclable and compostable materials, based on the recycling and composting facilities for a municipality. We use the San Francisco Recology

978-1-7281-1780-5/19/$31.00 ©2019 IEEE

guidelines [8]. Recology is a San Francisco-based integrated resource recovery company that processes compostable waste in an industrial composting facility that can break down fish, meat, dairy, and bio-plastics [9]. In addition, some recycling facilities will not recycle plastic utensils, while Recology will accept plastic utensils [9].

D. Overview

We situate our work in the field of machine learning and image classification, and discuss recent work done to classify waste materials in Section II. We then describe our methodology in Section III, starting with our data collection practices. We begin with the TrashNet dataset [10], which we modify and augment to include compostable products. We then discuss CompostNet, our neural network architectures used in the paper to conduct supervised learning and train two models, one of which is an example of transfer learning. In Section IV we discuss our results. Finally, in Section V, we present a prototype of an iOS application which can assist users in sorting their waste.

II. Related Work

A. Image Classification

Image classification is one of the major applications of artificial intelligence. Recent image classification models often rely on deep neural networks, specifically Convolutional Neural Networks (CNNs). CNNs are neural network variants that learn by performing convolutions and have shown stellar performance for image classification tasks [11]. Image classification relies on supervised learning, which requires labeled data to train networks. After the network is trained, it can classify images into discrete classes [11]. Using image classification, these networks can answer questions related to the visual properties of an image [12].

B. Waste Classification

Waste classification can be addressed in many ways - from educating individuals about sorting household trash [5] to using a hyperspectral imaging system to analyze attributes of waste products at compost or recycling facilities [7]. Few researchers have studied the use of CNNs to develop image recognition models for classifying waste. However, there are two examples of waste classifying CNNs. Gary Thung and Mindy Yang built the CNN “TrashNet” to classify waste into five classes of recyclable content and trash [10]. SpotGarbage [13] is a mobile application designed by researchers at the Indian Institute of Technology. This app allows users to identify garbage in the street around Indias urban centers. SpotGarbage uses a CNN called GarbNet, which has been trained on an annotated dataset called Garbage In Images. Although both are examples of CNNs used for waste classification, they do not classify food waste, or compostable material, separately.

III. Technical Structure

A. Dataset

We began with the data collected by Thung and Yang for ”TrashNet”. Their dataset consists of 2527 images in six classes: glass, plastic, cardboard, metal, paper, and trash [10]. We wanted to train our models on images of compostable content, in addition to recyclable and landfill content. We kept the subcategories of recyclables, but were more concerned with the three categories of ‘trash’, ‘recyclable’, and ‘compostable’. We augmented their dataset by adding 175 photos of food waste and 49 photos of landfill waste, for a total of 2751 images. Three example images are seen in Figure 2. We also moved images previously classified as ”trash” to the ”compostable” class, bringing the total image count to 177. We followed the methods for data collection that Thung and Yang outline in their project. We took photos of the waste against a white poster board, used natural or overhead lights, and focused on one piece of waste in each photo. Although Thung and Yang resized their images to 512×384 pixels, we resized our images to 400×300 pixels for our version B model. Our version A model required images to be resized to 224×224 pixels.

Fig. 2. Three images from our compostable class.

B. Models

1. CompostNet - Version A

Our first version of CompostNet utilized a pretrained MobileNet model that was partially retrained on our dataset. MobileNet is a lightweight mobile-first convolutional neural network trained on the ImageNet database [14]. We chose MobileNet because it balances efficiency and accuracy. We trained CompostNet - Version A on Tensorflow v. 2. This pretrained model serves as the first layer in our model. After the MobileNet model layer we have 4 layers which have been trained on our dataset. The last dense layer has a softmax function to present us with the seven outputs corresponding to the waste classes. The system architecture, including the MobileNet layer and our retrained layers, is shown in Figure 3.

2. CompostNet - Version B

This version of CompostNet was built around three convolutional layers. Our first layer breaks down the input image into 32 output matrices. It then goes through a max pooling layer with a 2x2 filter, halving the dimensionality of the output matrices, which reduces computation costs and helped us avoid overfitting. We apply a dropout rate of 30% after this step to further avoid overfitting. In our testing, we found that this

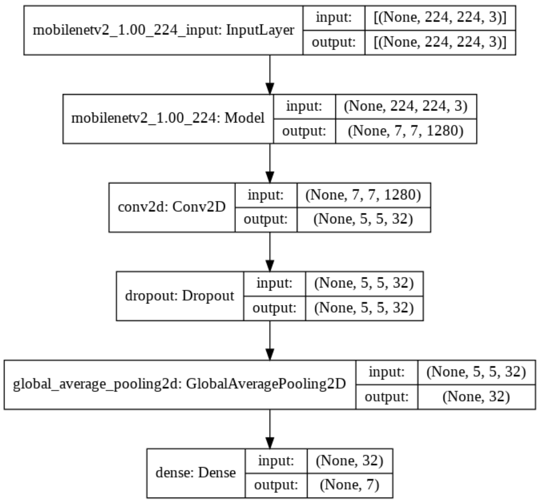

Fig. 3. Architecture for CompostNet - Version A. [14]

dropout rate was most effective in improving the accuracy of this model. These layers are replicated twice, but with different output matrices of 64 and 128 respectively. Finally, we flattened the output and passed it through a densely-connected neural network layer. The network architecture can be viewed on our Github.

IV. Results and Evaluation

Thung and Yang achieved an accuracy of 75% using their CNN 'TrashNet'. We took that as our baseline and aimed to exceed that level of accuracy.

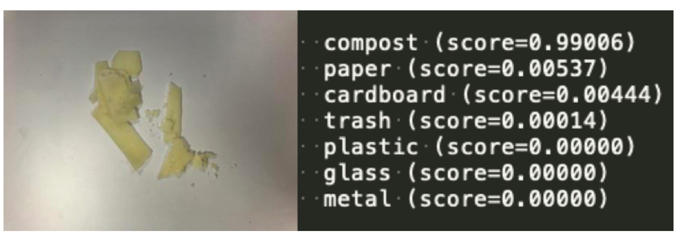

Fig. 4. A test image and the confidence scores returned by Version A model.

A. CompostNet - Version A

We split the dataset into two groups, 80% of images were used to train the network, and 20% of the images were excluded from training to test the network after it had finished training. We trained our model for 20 epochs, with a batch set of 64. We then tested the accuracy of the model on test image classification. Our network returned confidence scores for each of our seven classes. These values add up to one, and demonstrate how the model has classified the object. Once the network has been trained, it can classify images in approximately two seconds. An example test photo and the network output scores is shown in Figure 4. It is clear that the dataset must be expanded because of the amount of items that will need to be classified if a hardware system is deployed with this model.

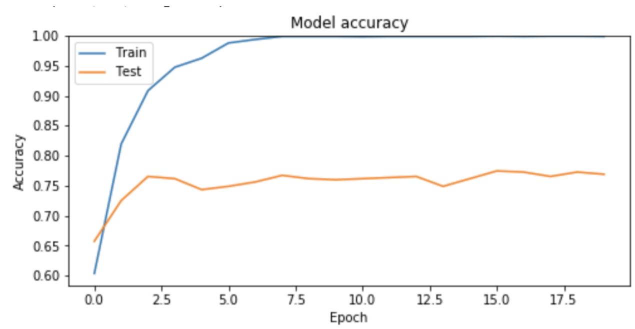

Fig. 5. Test and train accuracy for CompostNet - Version A.

After achieving a high training and test accuracy, the final test accuracy was 77.3%, as shown in Figure 5.

B. CompostNet - Version B

For this model, our train-test split was also 80/20. We ran the model for 20 epochs, with a batch size of 8. Training the model any further significantly decreased the accuracy, likely because it was overfitting to the data. The model's final test accuracy was 22.695%, see Figure 6.

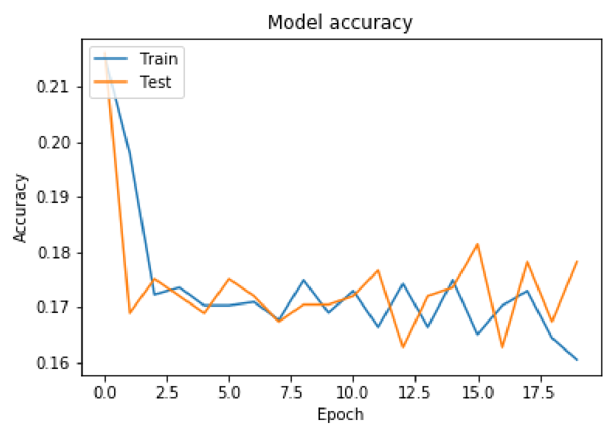

Fig. 6. Test and train accuracy for CompostNet - Version B.

Version A had an accuracy of 77.3% while Version B had an accuracy of 22.695%. This difference is not surprising. Version A has 159 layers and has been trained on the ImageNet dataset, which has over 1.3 million images [15]. Version B is significantly less robust. Expanding our dataset and adjusting the hyperparameters of either model may improve the accuracy.

V. Deployment

We have validated that our system classifies images with an accuracy above 75%, which was our baseline. To deploy this model, we have built a prototype of an iPhone app that users can use to identify if their waste can be recycled.

We saved our Version A model and converted the saved model to a TensorFlow lite compatible format. We used our model in the TensorFlow Lite image classification iOS project developed by Google [16]. This code project provides the user interface for the ML model. The user can open the app and hover the camera over a object. Our model is then used to classify the image, and the top three classes are returned for the user to see, as shown in Figure 7. This app has limited functionality, for example, it does not allow the user to choose a photo from his or her photo library. Although this iOS application is fairly simple, this deployment is exciting and represents a proof of concept.

Fig. 7. The iOS app after the system has classified an image.

VI. Conclusion AND Future Work

A. Future Work

We would like to expand the types of compostable materials (bio-plastic, paper plates, bamboo utensils, etc.) in our dataset to improve the accuracy of the CompostNet model. The Office of Sustainability at the University of California, Santa Cruz has expressed interest in this project as it aligns with the University of California sustainability goal of 90% of waste diverted from landfill by 2020 [17]. In an email sent to the UCSC community, Associate Chancellor Ashish Sahni wrote to students: “Campus recycling is temporarily being landfilled due to high rates of contamination, with the highest rates in our residential and dining halls…If you are unsure if something is recyclable, it is better to throw it into the landfill bin!” [18]. Our CompostNet system can help educate students about what can be recycled and composted at the end of a meal, to reduce the amount of material that is incorrectly sorted, preventing further cases of recycling contamination.

B. Conclusion

We have analyzed two models for image classification of three categories of waste - landfill, recyclable, and compostable. We demonstrate that with a small dataset, a model based on transfer learning returns good results for recognizing images of waste. The growing amount of waste sent to landfills is a rising concern addressed by the United Nations Sustainable Development Goal 12.5, as well as other countries and organizations. It is incumbent upon all of us to divert waste from landfills, and reduce the amount of waste we create.

Our dataset and code are available for download at https://github.com/sarahmfrost/compostnet.

VII. Acknowledgements

We would like to thank Manu Mathew Thomas at UC Santa Cruz for his assistance and support. We are grateful to Dan Frost, Professor Emeritus at UC Irvine, for his constructive criticism throughout the development of this project.

References

[1] “Goal 12: Responsible consumption and production.” [Online]. Available: https://www.globalgoals.org/12-responsible-consumption-and-production

[2] “The world’s trash crisis, and why many americans are oblivious,” Apr 2016. [Online]. Available: https://www.latimes.com/world/global-development/la-fg-global-trash-20160422-20160421-snap-htmlstory.html

[3] “Waste.” [Online]. Available: http://ec.europa.eu/environment/waste/index.htm

[4] S. Lebersorger and F. Schneider, “Discussion on the methodology for determining food waste in household waste composition studies,” *Waste management*, vol. 31, no. 9-10, pp. 1924–1933, 2011.

[5] Y. Zhuang, S.-W. Wu, Y.-L. Wang, W.-X. Wu, and Y.-X. Chen, “Source separation of household waste: a case study in china,” *Waste management*, vol. 28, no. 10, pp. 2022–2030, 2008.

[6] A. Demirbas, “Waste management, waste resource facilities and waste conversion processes,” *Energy Conversion and Management*, vol. 52, no. 2, pp. 1280–1287, 2011.

[7] G. Bonifazi, S. Serranti, A. Bonoli, and A. DallAra, “Innovative recognition-sorting procedures applied to solid waste: the hyperspectral approach,” *WIT Transactions on Ecology and the Environment*, vol. 120, pp. 885–894, 2009.

[8] Recology, “What goes where.” [Online]. Available: https://www.recology.com/recology-san-francisco/what-goes-where/

[9] S. Sweeney, “Industrial composting: What it is and how it works,” May 2018. [Online]. Available: https://www.urthpact.com/industrial-composting-what-it-is-and-how-it-works/

[10] G. Thung and M. Yang, “Trashnet.” [Online]. Available: https://github.com/garythung/trashnet

[11] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *nature*, vol. 521, no. 7553, p. 436, 2015.

[12] J. Shihadeh, A. Ansari, and T. Ozunfunmi, “Deep learning based image classification for remote medical diagnosis,” in *2018 IEEE Global Humanitarian Technology Conference*. IEEE, 2018, pp. 1–8.

[13] G. Mittal, K. B. Yagnik, M. Garg, and N. C. Krishnan, “Spotgarbage: smartphone app to detect garbage using deep learning,” in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 2016, pp. 940–945.

[14] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, “Mobilenets: Efficient convolutional neural networks for mobile vision applications,” *arXiv preprint arXiv:1704.04861*, 2017.

[15] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, “Imagenet large scale visual recognition challenge,” *International journal of computer vision*, vol. 115, no. 3, pp. 211–252, 2015.

[16] Tensorflow, “tensorflow/examples,” Jul 2019. [Online]. Available: https://github.com/tensorflow/examples/tree/master/lite/examples/image_classification/ios#tensorflow-lite-image-classification-ios-example-application

[17] “Sustainability.” [Online]. Available: https://www.ucop.edu/sustainability/policy-areas/waste-reduction-and-recycling/index.html

[18] A. C. A. Sahni, “We need your help with campus recycling,” May 2019.