Convolutional Neural Networks for Sustainable Waste Classification

Sarah Ciresi, Anand Natu

Background / Problem Statement

Motivation: Municipal recycling is complicated by inconsistent waste disposal practices - educating consumers on proper disposal methods can significantly enhance the efficiency and safety of recycling processes

Automated waste classification: Given an input image of some piece of waste, determine the appropriate disposal treatment as a multiclass classification task

Mobile use case: Target end use is a in a mobile app that should allow consumers to perform real-time waste classification using their phone camera. Therefore, we seek an optimal combination of classification performance and model efficiency (compute- and memory-intensity)

Problem Formulation: Input image is classified into one of five recyclable categories, compost, or landfill-bound trash

Methods

Model Selection: Problem feasibility established with best-in-class performance models (VGG-16, ResNet101), and then optimized for a mobile hardware setting (MobileNetV2). Results compared to RBF SVM baseline

Transfer Learning: ImageNet pretrained models are used, with output layer swapped to match our classification task. Models are finetuned (re-trained on our data) using 3 different learning methods:

1. **No finetuning** other than final linear layer
2. **Partial finetuning** of subset of model layers
3. **Full finetuning** of all model layers

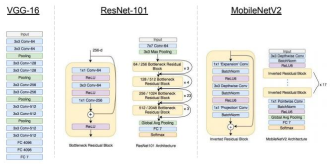

Specification of Convolutional Models
Tuned HP values and compute / memory performance

Model	Learning Method	α	γ	η
VGG-16	(1)	0.00012	0.1	11
VGG-16	(2)	0.0001	0.1	7
VGG-16	(3)	0.0001	0.1	7
ResNet101	(1)	0.00015	0.1	10
ResNet101	(2)	0.0001	0.1	7
ResNet101	(3)	0.0001	0.1	7
MobileNetV2	(1)	0.01	0.1	7
MobileNetV2	(2)	0.0005	0.1	7
MobileNetV2	(3)	0.0008	0.1	7

Model	Input Size	# Params	Memory (MB)	GFLOPs
VGG-16	224x224	133M	528	16
ResNet101	224x224	44M	170	8
MobileNetV2	224x224	3.4M	13	0.32

Discussion

Classification Performance

- Fully finetuned ResNet101 model achieves the highest test set accuracy
- However, comparable performance from the lightest and most compute- and memory-efficient model, MobileNetV2 (90.1% test accuracy)

Model Diagnostics

- Accuracies correlate positively with extent of fine-tuning, which comports with general expectation
- Analysis of misclassified images confirms that the model errors resemble challenging cases for human-level reasoning on this classification task, or indicate data volume as a root-cause

Error Analysis and Deep-Dives

Object ambiguity: Plastic cup, but is also compostable (denoted in the green text)

Incomplete image: Only part of the bottle is visible, so its paper label is misinterpreted

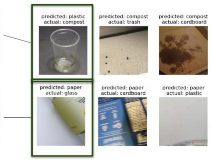

Data Collection and Preprocessing

- We use the **TrashNet** dataset collected by Yang and Thung as the base of our dataset
  - **2,527** color images of waste labelled as one of six classes: paper, plastic, metal, cardboard, glass, trash
- We augment the dataset with an additional **1049** images
  - **926** images across the six TrashNet categories
  - **300** images of waste in the novel **compost** class
- Data is hand-collected by taking photographs of waste using cell phone cameras to ensure that images resemble input from the intended mobile use case
- All images are resized to dimension 224x224x3 to match standard input size of pretrained model architectures
- Combined dataset size of **3,753** images is split **70/15/15** train/validation/test
- Employ data augmentation techniques (random crops and random flips) on images in the training set

Experiments and Results

Confusion Matrices for Test-Set Classification
classifications, by combination of predicted vs. ground-truth class label

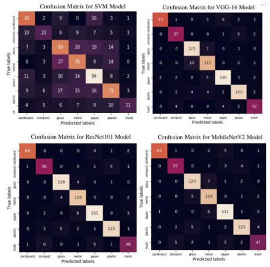

Summary of Model Performance
Accuracy (% correct), by model, by learning method

Model	Learning Method	Train Acc.	Val Acc.	Test Acc.
SVM	--	99.8%	56.5%	54.6%
VGG-16	(1)	85.7%	82.4%	--
VGG-16	(2)	94.6%	90.2%	--
VGG-16	(3)	94.5%	90.3%	89.7%
ResNet101	(1)	73.9%	78.6%	--
ResNet101	(2)	97.6%	93.7%	--
ResNet101	(3)	97.7%	94.4%	91.6%
MobileNetV2	(1)	77.1%	81.7%	--
MobileNetV2	(2)	95.5%	91.1%	--
MobileNetV2	(3)	95.2%	92.1%	90.1%

- Convolutional models **significantly outperform the SVM baseline,** and do so exhaustively across all class categories
- Of all learning methods surveyed, **full finetuning consistently achieves the best results** for all model architectures used

Conclusions

- High accuracies and baseline clearance show that **robust waste classification can be achieved by leveraging CNNs and deep learning-based modeling**
- Consistently strong performance for lightweight, compute-optimized models suggests that **classification performance can be feasibly adapted to a mobile hardware setting**
- High visual variance in certain waste types (e.g. compost) suggests that **augmenting data is an important determinant of robustness**

Future Work

Expand and enrich dataset, especially for trash and compost classes

Measure on-line classification time for model usage on mobile hardware

Integrate model into mobile app for common mobile operating system(s)

Implement instructive guidance to help users properly dispose of classified waste