\Orchestrating a brighter world

Stool Structure Classification Using Smart Toilets

AUTHORS
Avneet
Raheel Zuberi

NEC CORPORATION
INDIA PRIVATE LIMITED

Abstract

Stool analysis is an important diagnostic tool used in the clinical setting to evaluate gastrointestinal health. The morphology of stool can provide valuable information regarding digestive function, disease, and treatment efficacy. Health monitoring is facilitated by long-term data collection that establishes a health baseline and enables the detection of deviations from it. With the advent of the Internet of Things, monitoring daily excreta from a toilet is emerging as a promising tool to achieve the long-term collection of physiological data. This paper describes a stool structure classification approach that accurately and efficiently tracks stool structure using a smart Toilet. The Smart Toilet captures the image of stools in the toilet plumbing outside the purview of the user. We constructed a stool image dataset with 2350 images, split the dataset into training and testing sets, and several machine learning algorithms were evaluated for their classification performance. The labelling of the dataset was performed by the subject matter experts.

We addressed the following limitations associated with the application of computer vision techniques to a smart toilet system:
(i) uneven separation among the different stool structure categories;
(ii) class imbalance in the dataset;
(iii) limited computational resources in the microcontroller integrated with the smart Toilet.

For the classification of stool form, we achieved a Neighboring accuracy of 95% using a CNN based on MobileNetV2. Our proposed classification system can aid in the diagnosis and monitoring of gastrointestinal diseases, and to improve patient outcomes by facilitating timely and accurate treatment.

© NEC Corporation 2023

\Orchestrating a brighter world NEC

1. Introduction

Biomedical imaging is one of the cornerstones of medical diagnostics and it is being enhanced by sophisticated machine learning techniques. Recent applications of machine learning for health applications have focused on the analysis of physiological data collected over a prolonged time. This analysis provides individualized risk assessment and early warning of disease onset that can be used to trigger interventions. Long-term adherence to precision health monitoring is facilitated by not requiring the user to personally collect the data. Human excreta (urine and stool) are readily available specimens regularly deposited in toilets. With the advent of the Internet of Things (IOT) paradigm, monitoring physiological functions from a toilet during bathroom visits is emerging as an active area of research forprecision health. Research on “smart toilets” for health monitoring has thus far mainly focused on urine analysis; however, important health information is also found in faeces.

Specifically, stool physical characteristics such as form (i.e., consistency) and color contribute to the diagnosis and management of many acute and chronic gastrointestinal (GI) conditions. Stool appearance is one of the early diagnostic indicators for evaluation of irritable bowel syndrome (IBS), (as much as 10-15% of the world population is estimated to suffer from IBS) inflammatory bowel disease (IBD), malabsorption syndromes, and upper and lower GI bleeding (Tanaka et al., 2018). The impact of GI diseases on patients and the health care system is substantial; for example, in the US, GI healthcare cost are higher than the cost associated with heart disease (Peery et al., 2019).

There is no approved clinical method that can reliably and consistently monitor stool frequency, form, and color, either in the home setting or in the hospital. To address this limitation, image capture of the content of a toilet bowl either by the user (Hachuel et al.,2019) or without user intervention (Park et al., 2019) has been proposed.

Our team is developing a smart toilet sensor that can be attached with any toilet and enables discreet imaging of stool in the toilet pot. Stool image analysis is a key enabler of smart toilets for monitoring bowel movement. In this paper, we present a technique used to accurately detect stool structure. For stool structure classification, we use a compact architecture, such as MobileNetV2, as a typical example.

The main contributions of this paper are as follows:
1. We constructed a stool image dataset containing 2,350 stool images spanning all seven Bristol Stool Form Scale (BSFS) types.
2. We present the design of a hierarchical CNN architecture for stool structure classification over seven BSFS values.
3. We present results for stool structure classification using multiple CNN models, such as MobileNetV2 and Resnet50, for training the CNN classifier.

The data centric model development approach also plays a major role in accuracy improvement. Keeping the models same while improving the quality as well amount of data helps us achieve higher accuracy than previous methods.

Generalizable Insights about ML in the Context of Healthcare

We demonstrate machine learning (ML) approach for clinically relevant stool characteristics that is both accurate and computationally efficient. The ML solution enables the classification of stool characteristics and provides objective data to inform improved clinical care. This computational tool is implemented as edge computing near the image data source. We describe an approach that addresses challenges commonly faced by computer vision techniques being applied to medical imaging. First, we use approaches such as hierarchical CNN architecture to overcome the issue of uneven separability between different categories.

Second, by training the CNN using class-balanced loss based on the effective number of samples, we can address the problem of class imbalance in the dataset. Third, by evaluating several recent CNN designs, we select a design that enables image classification that is computationally efficient as defined by metrics of the number of float-point operations (FLOPs) and memory requirement, so that it will be easier to deploy the ML model in a resource-constrained environment, such as the physical smart toilet hardware. Overall, this combination of machine learning and stool specimen imaging enables a new form of physiological monitoring that may provide early warning of disease for timely intervention and improved clinical outcomes.

The rest of this paper is organized as follows. Section two describes related prior work and provides further motivation for this research. Section three describes the background of the smart toilet and our system design for stool analysis. Section four describes the proposed methods for stool-form classification. Section five presents the experimental results. Section six describes the limitations of this paper. Finally, Section seven concludes the paper.

2. Related Prior Work

2.1. Bristol Stool Form Classification

The BSFS scale (Lewis and Heaton, 1997) is a standard medical diagnostic tool for categorizing adult stool based on its physical appearance. Normal stool consistency is defined as BSFS type 3, 4 and 5 (Markland et al., 2013). Constipation is defined as Type 1(separate hard lumps, like nuts) or Type 2 (sausage-like, but lumpy). Diarrhoea is defined as a minimum of three loose stools (Type six and Type 7) per day. The BSFS stool chart is shown in Figure 1.

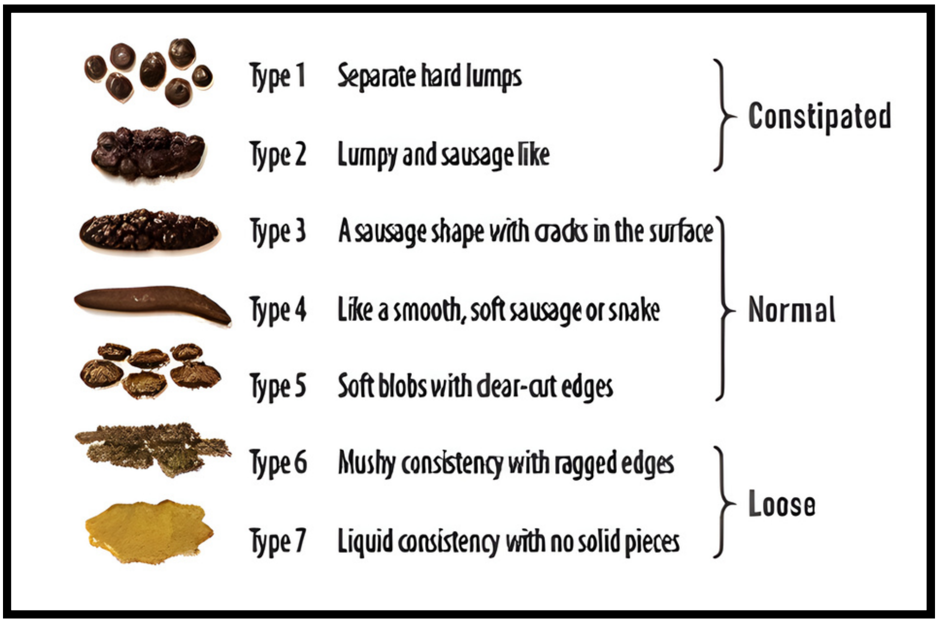

Figure 1: Illustration of the BSFS chart (adapted from http://cdn.intechopen.com/pdfs-wm/46082.pdf)

A 2019 study validated the use of the BSFS by having participants use a printed card tool with graphics to assess the properties of their bowel movements (Ohno et al., 2019).

2.2. Machine Learning Approaches for Stool Image Assessment

Yang et al. (Yang et al., 2019) introduced StoolNet, which combines region of interest (ROI) detection and a shallow CNN for color classification of stool images. Park et al. (Park et al., 2020) used transfer learning to train a classifier on top of a trained deep learning architecture. While these studies provide key insights into automated stool classification, three major challenges remain to be addressed, namely, uneven separability, class imbalance, Variable stool size and model complexity.

Uneven Separability: Visual separability between different BSFS categories is uneven. For example, it is difficult to distinguish type three from type four, while it is easy to tell type one from type three. Traditional CNNs (Krizhevsky et al., 2012; Simonyan and Zis-serman, 2015) use the flat structure to train a N-way classifier and do not consider such uneven separability, which often leads to sub-optimal performance in the task of fine‐grained classification.

A common strategy to address this problem is to predefine a hierarchy or taxonomy of classifiers so that a given testing image can be first evaluated by a coarse classifier and the corresponding fine classifier to make the fine prediction (Murthy et al., 2016; Yan et al., 2015).

Class Imbalance: Medical diagnostic data may have a normal distribution (bell-shaped curve) or a skewed distribution. For instance, in the stool image dataset collected by Park et al. (Park et al., 2020), only a few images report constipated stool, while most images indicate normal stool. A number of solutions have been proposed in the literature to address the problem of class imbalance. The first approach is re-sampling, which aims to alter the training data distribution, usually by random under- and oversampling techniques (Oquab et al., 2014; Chawla et al., 2002). The second approach is cost-sensitive learning, which assigns higher misclassification costs to minority classes compared to the majority classes.

Variable Size: Medical diagnostic image possesses unique challenge in terms of size. There may be some cases where the stool is very small (few pixels) as compared to other stool images where major part of image is represented by stool. There are multiple approaches to solve this problem, but we followed the simpler ones where we resized larger images into desired shape. While super resolution is applied to smaller images to enlarge them without losing significant information. Choosing a smaller input (96*96) for CNN models also helps because we do not have to enlarge too many images instead, we have to reduce the size of the image.

Model Complexity: In practice, state-of-the-art CNN models (Simonyan and Zisserman, 2015) incur significant compute overhead, which imposes a barrier to their deployment on devices with limited computational power, e.g., a micro-computer (Raspberry Pi). Many approaches have been proposed to address this challenge, which can be categorized on the basis of techniques that use either model compression or compact architectures. Model compression techniques include parameter pruning and weight quantization (Denton et al., 2014; Cheng et al., 2017). However, these methods require dedicated hardware or software customization for practical implementation. In contrast, compact architecture design methods target more efficient and compact neural network architectures (Iandola et al., 2018; Howard et al., 2017).

3. Smart Toilet System

In this section, we provide an overview of the smart toilet system and formulate stool analysis as a real-time computer vision problem. Smart-toilet approaches have been proposed to obtain health-related information from different configurations, e.g., devices snapped on the toilet bowl (Hall et al., 2020) or integrated in the toilet seat (Park et al.,2020; Conn et al., 2019). Notably, Park et al. (Park et al., 2020) introduced a defecation monitoring module that uses sensors and computer vision to acquire basic properties of human excreta from sensors integrated in a commercially available electronic bidet. The acquired images are fed offline to machine learning algorithms for analysis. However, cameras and illumination devices in the toilet seat create an uncomfortable environment for the user, as highlighted by the results of a user survey regarding the technology (Park et al.,2020).

An alternative approach that avoids the adoption barrier due to user discomfort, is a technology that integrates sensors in the toilet plumbing where they are not visible to users. A toilet manufacturer reported such a configuration for urinalysis.

We have developed a novel approach to image feces in the toilet seat. Our design offers a unique opportunity for a real - time inline sensing approach specific to excretion without engendering user discomfort. Also have face recognition option when user uses the toilet. The main advantage of this design is easily integration with any toilet seat.

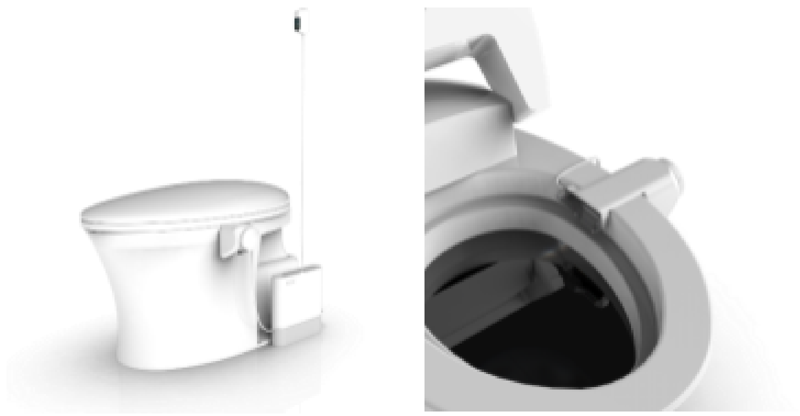

(a) (b)

Figure 2: (a) The setup for stool image analysis.
(b) Camera placement inside the toilet pot.

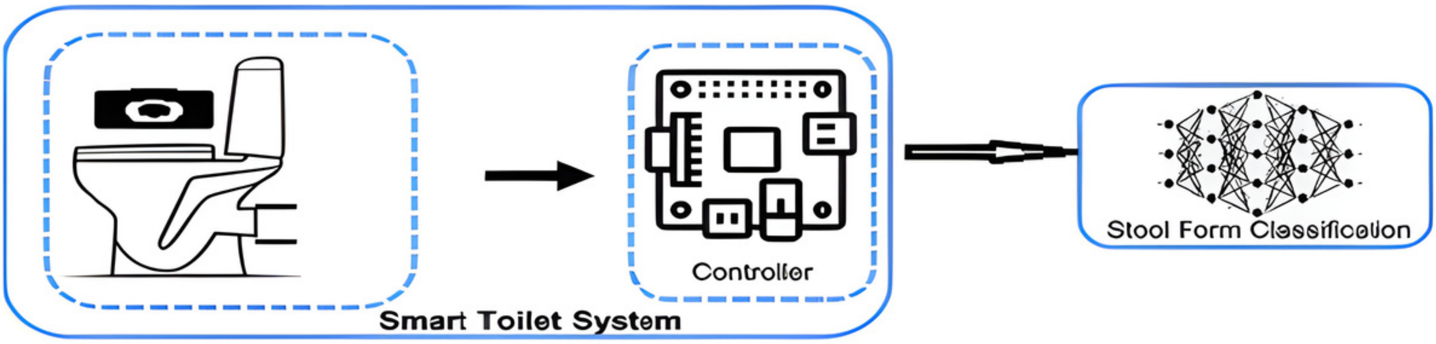

Figure 3: Framework for stool image analysis in a smart toilet system. (a) Images are captured when the stool is in the toilet pot (b) Images are processed by the controller and fed to machine learning algorithms.

The hardware setup used for stool image analysis is shown in Figure 2.

4. Smart Toilet System

In this section, we describe the dataset used for analysis, as well as the machine learning techniques used for classification.

4.1. Stool Image Data Set Preparation

No publicly available stool image dataset exists; thus, we developed our own image dataset. The stool image dataset used in (Yang et al., 2019) for their StoolNet model only contained 110 images (each rotated and used four times) and, these images are not publicly available.

Our work leverages a dataset of 2,350 stool images spanning all seven BSFS types obtained from test sites which has been setup on various locations in Japan. A total of 2,350 unique images were obtained through real test locations. Different masks were added to ensure the privacy of patients. Because the data was small and not evenly distributed therefore we have to use augmentation technique to increase the number of images and improve the class balance. Fig 4 shows original as well as augmented images per class.

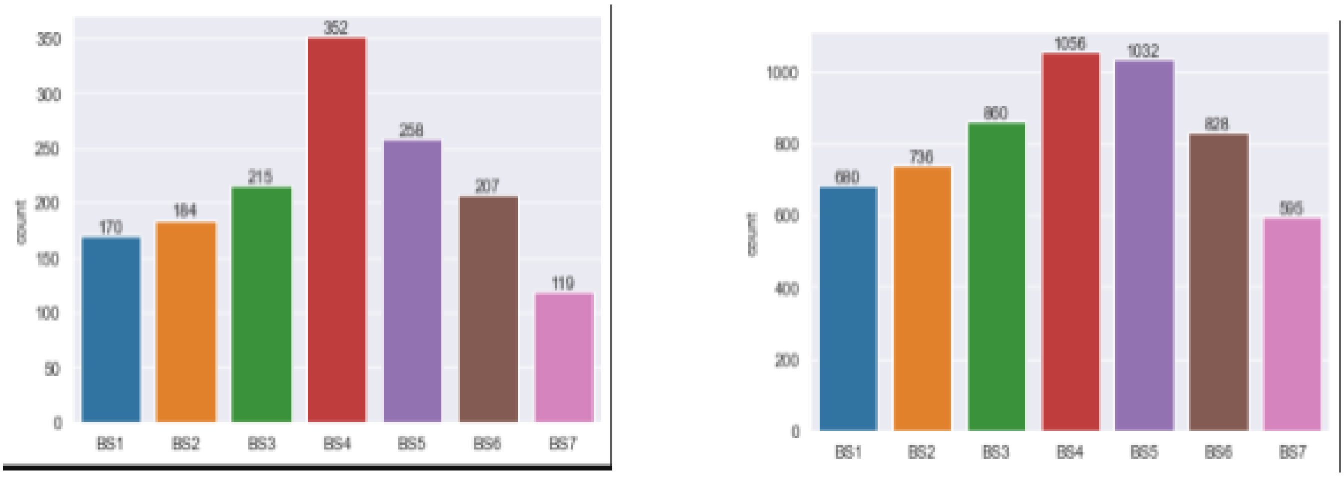

(A) (B)

Figure 4: (A) Bar graph of original train dataset.
(B) Bar graph of augmented train dataset.

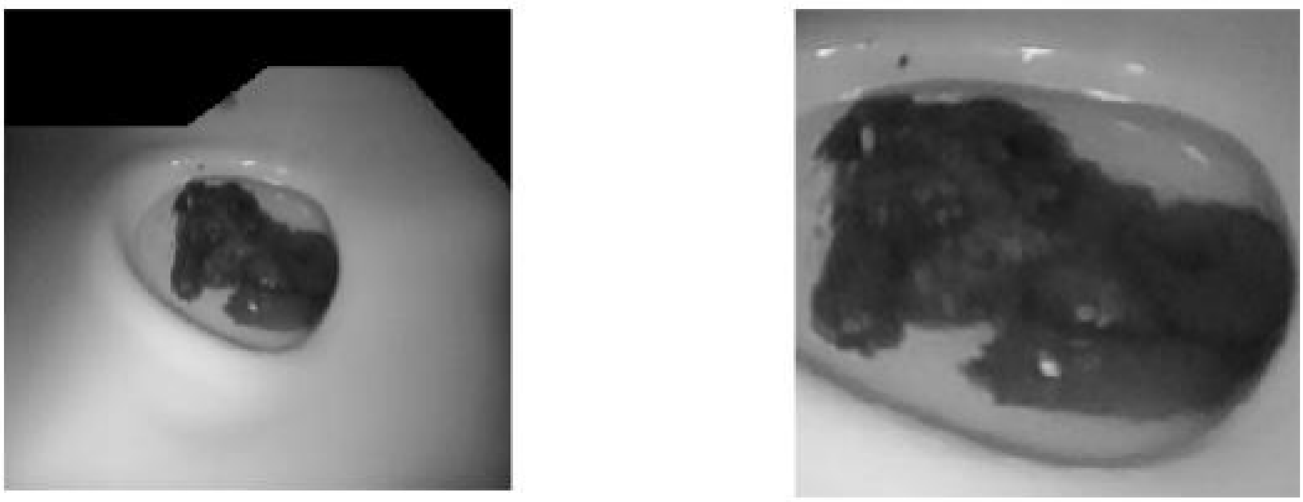

Figure 5: (left) Real image with a mask obtained from toilet camera, (Right) cropped poo part from image.

We used the online platform Labelbox (Labelbox, 2019) to label the images and crop specific part of images that contain the stool and assigned to each image a BSFS score from one through 7. Despite being a clinical standard, the BSFS score does not capture the full variety of stool forms and does not account for the presence of stools of more than one BSFS category in the same image. From a clinical point of view, the important information is whether the bowel movement is normal (types 3,4,5) or abnormal, that is, constipated (types 1,2) and diarrhoea (types 6,7).

We also found that when working with stool images without toilet background (Figure 6), it generates better results in training and validation, while in real world scenarios, stool images with toilet background generate better.


Figure 6: Stool pixels without toilet background (left), Stool pixels with toilet background (Right).

Another problem we faced during data preparation was the small size of cropped stool. To overcome this, we used super resolution techniques. If the image size is lower than decided threshold, then super resolution techniques were applied which results in larger cropped and better resolution image.

4.2. Stool Structure Classification

The first part of our problem is to determine whether there is stool present in toilet pot or not and this has been achieved using segmentation techniques.

When there is stool in the pot, we crop that part of image using segmentation and do further stool structure classification on the cropped part.

4.2.1 Architecture

Figure 7: Model architecture for classification.

4.2.2 Base CNN Design

Various CNN designs have been proposed over the past few years for various applications (Simonyan and Zisserman, 2015; Howard et al., 2017; Gatys et al., 2015). We consider a single-board computer, i.e., ASUS Tinker Board, to load our CNN models. The computational resources available on Tinker Board are limited compared to a server, therefore deep CNNs such as VGG16 are not feasible in this application scenario. In this paper, we explore CNN designs, namely MobileNetV2 (Sandler et al., 2018), VGG16, and Resnet50.

MobileNetV2. To reduce computation cost, MobileNetV1 (Howard et al., 2017) replaces the standard convolutional filters by two layers: depth wise convolution and 1 × 1 pointwise convolution, where depth wise convolution only extracts spatial features for each independent channel and pointwise convolution extracts channel-wise information. Furthermore, MobileNetV2 uses an inverted bottleneck structure to increase representational power. The base CNN is used for feature extraction. We use these CNN designs as the base CNN by removing their last classification layers. Specifically, we removed the last two layers (one dropout layer and one fully-connected layer) from MobileNetV2 and added our custom layers on top of it to generate better results.

5. Experiments & Results

Experiments were conducted to evaluate the effectiveness of the proposed approach for classifying stool forms and detecting images of stool. We pre-processed the stool images by cropping them and selecting the parts of the images where stool pixels were majorly presented.

5.1. Results of Stool Structure Classification

Visual differences among various BSFS categories are uneven. That is, it is difficult to distinguish BS3 from BS4, while it is easy to tell BS1 from BS3. To leverage the hierarchical structure of stool-form categories, we have come make neighbouring class accuracy. In the neighbouring class accuracy matrix, accuracy of adjacent classes is also taken into account.

For example: if we are calculating the accuracy of class BS3, then images that have been predicted as BS2 or BS4 are also considered accurate due to corner/edge cases.

5.1.1. Hierarchical Architecture

The training process for hierarchical CNN includes three steps. We first initialize the base CNN with pre-training on ImageNet (Deng et al., 2009). After initialization, we train the classifier and the base CNN together over seven consolidated categories. We used 96*96 pixels as input size to the CNN.

We used Tensorflow and Keras (Paszke et al., 2019) to implement and train the CNN using stochastic gradient descent with momentum. Experiments were executed on Windows platform. The training was performed with mini batches of size 128. We also tried to batch size of 64, 128 and 256.

We consider two compact CNN designs and two traditional CNN designs, namely, VGG16 (Simonyan and Zisserman, 2015), ResNet50 as the base CNN in the hierarchical architecture. We evaluate the performance of classification over the BSFS scale with seven values. The results are shown in Table 2 and summarized as follows:

© NEC Corporation 2023

Models	Accuracy (Fine-Grained)	Accuracy (Neighbour)	FLOPs	Memory Required
VGG16	82.78%	93.68%	15.39 G	61.5 MB
ResNet50	84.66%	94.72%	4.16 G	104.8 MB
MobileNetV2	85.31%	95.10%	0.35 G	15.6 MB
Custom CNN	79.61%	90.15%	0.32 G	15.2 MB

Table 2: Accuracy matrix with various base CNN designs on the stool image dataset.

- Classification for ResNet50 and for MobileNetV2. The hierarchical architecture brings a slight increase in required memory and FLOPs for inferencing, because it has three more classifiers than the flat architecture.
- Hierarchical architectures with MobileNetV2 as the base CNN achieve the best performance Moreover, MobileNetV2 only requires 0.35 GFLOPs for conferencing and the memory requirement is only 15.6 MB.

5.1.2. Prediction Analysis

Visual separability between different BSFS categories is uneven. For example, it is difficult to distinguish type three from type four, while it is easy to tell a type one from type three. To leverage the hierarchical structure of stool-form categories, we have come make neighbouring class accuracy. In neighbouring class accuracy, the matrix accuracy of adjacent classes also taken into account.

For example: we are calculating the accuracy of class three then images which have predicted class two or class four also considered accurate due to edge/Multi-class cases. This can be better explained using the confusion matrix.

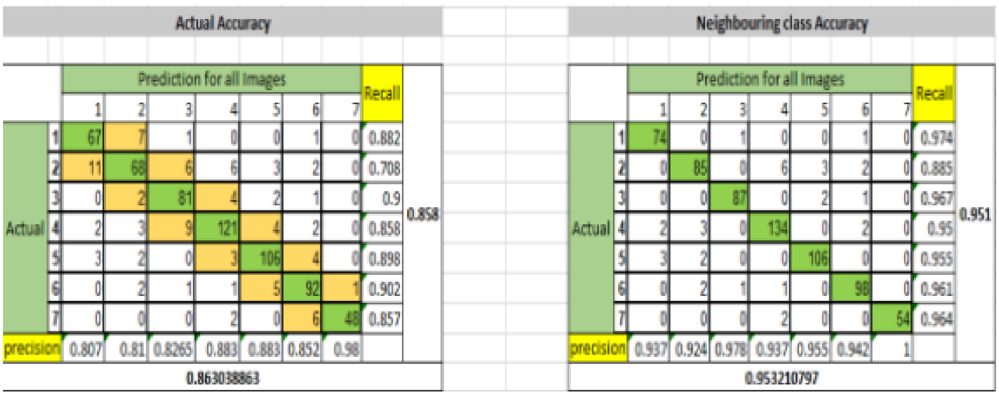

Actual Accuracy

	Prediction for all Images: 1	2	3	4	5	6	7	Recall
Actual 1	67	7	1	0	0	1	0	0.882
2	11	68	6	6	3	2	0	0.708
3	0	2	81	4	2	1	0	0.9
4	2	3	9	121	4	2	0	0.858
5	3	2	0	3	106	4	0	0.898
6	0	2	1	1	5	92	1	0.902
7	0	0	0	2	0	6	48	0.857
precision	0.807	0.81	0.8265	0.883	0.883	0.852	0.98	

0.858

0.863038863

Neighbouring class Accuracy

	Prediction for all Images: 1	2	3	4	5	6	7	Recall
Actual 1	74	0	1	0	0	1	0	0.974
2	0	85	0	6	3	2	0	0.885
3	0	0	87	0	2	1	0	0.967
4	2	3	0	134	0	2	0	0.95
5	3	2	0	0	106	0	0	0.955
6	0	2	1	1	0	98	0	0.961
7	0	0	0	2	0	0	54	0.964
precision	0.937	0.924	0.978	0.937	0.955	0.942	1	

0.951

0.953210797

Figure 8: Confusion matrix with actual accuracy and neighbouring accuracy (MobileNet V2).

© NEC Corporation 2023

As illustrated in section 5.1.1, MobileNetV2 outperforms VGG16 and ResNet50 m in terms of the metrics computational costs (FLOPs and required memory).

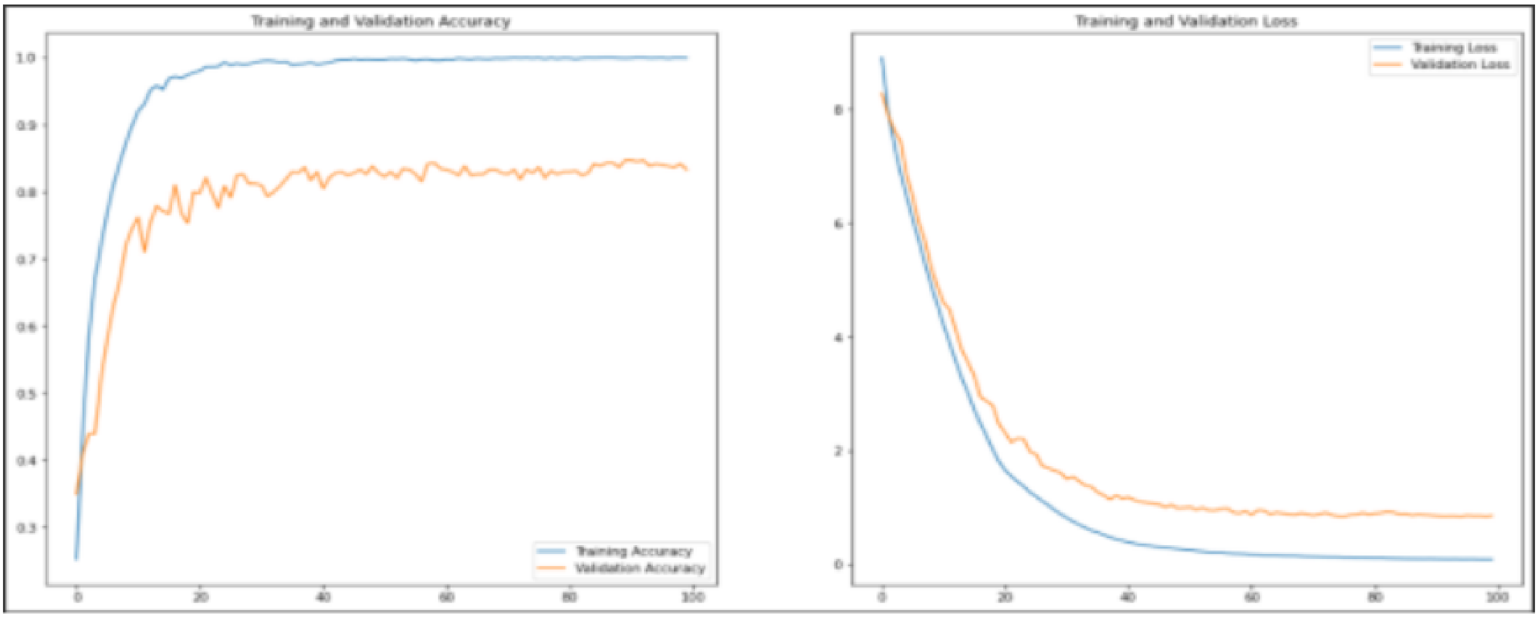

Figure 9: Training graphs (accuracy and loss) of MobileNetV2.

Stool Class	1	2	3	4	5	6	7
Recall	0.974	0.885	0.967	0.95	0.955	0.961	0.964
Precision	0.937	0.924	0.978	0.937	0.955	0.942	1

Table 3: Recall and precision of MobileNetV2 considering neighbouring classes.

Stool Class	1	2	3	4	5	6	7
Recall	0.882	0.708	0.9	0.858	0.898	0.902	0.857
Precision	0.98	0.87	0.81	0.826	0.883	0.852	0.98

Table 4: Recall and precision of MobileNetV2 with fine-grained classes.

6. Limitations & Discussion

The development of an ML-AI program to automatically classifies stool images for form (Bristol scale) requires a large number of annotated photos of stool in a toilet. We annotated the dataset with more than 2300 images. A limitation of this approach is that the photos had no clinical data associated with them, and while they spanned the full spectrum of the Bristol scale, the representation of associated gastrointestinal conditions or symptoms was unknown Additionally, while the use of the Bristol scale helps standardize stool evaluation, there remains some variability in assessment even among gastrointestinal specialists.

We envision that with the future deployment of the Smart Toilet hardware prototype for use by human subjects, we will be able to collect time series data from individual subjects. We expect that stool image data collection from the controlled environment will result in more consistent lighting and even background that will enhance the model accuracy. A smart toilet with machine learning image analysis capability to determine stool frequency and form will provide important diagnostic data that can help identify specific food intolerance (e.g., foods that exacerbate IBS or chronic diarrhoea) and effects of medication (e.g., medications taken for diarrhoea or constipation), and can trigger timely evaluation. We envision that the Smart Toilet time series data collected from individuals will be integrated with machine learning predictive models and provide a valuable diagnostic and surveillance tool for GI, infectious disease, and other specialties.

7. Conclusion

We have developed an automated technique for stool classification using a combination of a Smart Toilet and machine learning. We have developed a comprehensive stool image dataset for assessing the classification approach. We have used CNN architectures that can be used for stool image analysis in a resource-limited computational environment. Specifically, we showed that the CNN based on MobileNetV2 can achieve a neighbouring accuracy of 95%, with the memory requirement of only 15.6 MB and 0.32 GFLOPs for inferencing. Our results open up an interesting new re - search direction on privacy-preserving and real-time stool classification for health assessment.

8. Acknowlegement

This project was developed under the guidance o f NEC Japan (NEC Platforms, Ltd.).
A product based on this paper has been launched in Japan. For further product details please refer to this link:
https://www.necplatforms.co.jp/press/202205/20220523_01.html

References

- S. Lewis and K. Heaton. Stool form scale as a useful guide to intestinal transit time. *Scandinavian Journal of Gastroenterology*, 32(9):920–924, 1997.
- A. Markland, O. Palsson, P. Goode, K. Burgio, J. Busby-Whitehead, and W. Whitehead. The association of low dietary intake of fiber and liquids with constipation: evidence from the national health and nutrition examination survey. *The American Journal of Gastroenterology*, 108(5):796, 2013.
- N. Chawla, K. Bowyer, L. Hall, and W. Kegelmeyer. Smote: the synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16(1):321–357, June 2002.
- M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.
- Y. Cheng, D. Wang, P. Zhou, and T. Zhang. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv*:1710.09282, 2017.
- J. Deng, W. Dong, R. Socher, L. Li, K. Li, and F. Li. ImageNet: A Large-Scale Hierarchical Image Database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- E. Denton, W. Zaremba, J. Bruna, Y. LeCun, and R. Fergus. Exploiting the linear structure within convolutional networks for efficient evaluation. In *Advances in Neural Information Processing Systems*, 2014.
- L. Gatys, A. Ecker, and M. Bethge. Texture synthesis using convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2015.
- D. Hachuel, A. Jha, D. Estrin, A. Martinez, K. Staller, and C. Velez. Augmenting gastrointestinal health: a deep learning approach to human stool recognition and characterization in macroscopic images. *arXiv preprint arXiv*:1903.10578, 2019.
- D. Hall, D. Crismon, J. Larsen, K. Campbell, J. Reynolds, and J. Blake. Us20200225121a1: Toilet equipped to provide faecal analysis, 2020.
- N. Conn, K. Schwarz, and D. Borkholder. In-home cardiovascular monitoring system for heart failure: comparative study. *JMIR mHealth and uHealth*, 7(1), 2019.
- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In IEEE *Conference on Computer Vision and Pattern Recognition*, 2016.

- A. 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. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- F. Iandola, M. Moskewicz, K. Ashraf, S. Han, W. Dally, and K. Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and < 1 mb model size. In *International Conference on Learning Representations*, 2018.
- V. Murthy, V. Singh, T. Chen, R. Manmatha and D. Comaniciu. Deep decision network for multiclass image classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- H. Ohno, H. Murakami, K. Tanisawa, K. Konishi, and M. Miyachi. The validity of an observational assessment tool for multifaceted evaluation of faecal condition. *Scientific Reports*, 9(1):1–9, 2019.
- Labelbox. https://labelbox.com, 2019.
- M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.
- S. Park, D. Won, B. Lee, D. Escobedo, A. Esteva, A. Aalipour, T. Ge, J. Kim, S. Suh, E. Choi, A. Lozano, C. Yao, S. Bodapati, F. Achterberg, J. Kim, H. Park, Y. Choi, W. Kim, J. Yu, A. Bhatt, J. Lee, R. Spitler, S. Wang, and S. Gambhir. A mountable toilet system for personalized health monitoring via analysis of excreta. *Nature Biomedical Engineering*, 4:624–635, 2020.
- A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019.
- A. Peery, S. Crockett, C. Murphy, J. Lund, E. Dellon, J. Williams, E. Jensen, N. Shaheen, A. Barritt, S. Lieber, B. Kochar, E. Barnes, Y. Fan, V. Pate, J. Galanko, T. Baron, and R. Sandler. Burden and cost of gastrointestinal, liver, and pancreatic diseases in the United States: Update 2018. *Gastroenterology*, 156(1):254–272.e11, 2019.
- S. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012.
- M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- S. Tanaka, K. Yabunaka, M. Matsumoto, N. Tamai, H. Noguchi, M. Yoshida, G. Nakagami,
- J. Sugama and H. Sanada. faecal distribution changes using colorectal ultrasonography in older people with physical and cognitive impairment living in long-term care facilities: a longitudinal observational study. *Healthcare*, 6(2):55, 2018.
- Z. Yang, L. Leng, and B. Kim. Stoolnet for color classification of stool medical images. *Electronics*, 8:1464, 2019.

\Orchestrating a brighter world

NEC

Acknowledgement

Avneet
Data Scientist
Public Solutions, NEC India

Raheel Zuberi
Member Technical Staff
Public Solutions, NEC India

Contact Us:

marketing@india.nec.com

in.nec.com

Copyright © NEC Corporation India Private Limited 2023
All rights reserved.