Automated blood volume estimation in surgical drains for clinical decision support

A.A. SAHIN[1], M.A. SAHIN[1], M.E. YUKSEL[2], S.E. ERDEM[1]

[1]Department of Electrical and Electronics Engineering, Hacettepe University, Ankara, Turkiye
[2]Surgical Intensive Care Unit, Ankara Etlik City Hospital, Ankara, Turkiye

Abstract. – OBJECTIVE: Monitoring Jackson Pratt and Hemovac drains plays a crucial role in assessing a patient's recovery and identifying potential postoperative complications. Accurate and regular monitoring of the blood volume in the drain is essential for making decisions about patient care. However, transferring blood to a measuring cup and recording it is a challenging task for both patients and doctors, exposing them to bloodborne pathogens such as the human immunodeficiency virus (HIV), hepatitis B virus (HBV), and hepatitis C virus (HCV). To automate the recording process with a non-contact approach, we propose an innovative approach that utilizes deep learning techniques to detect a drain in a photograph, compute the blood level in the drain, estimate the blood volume, and display the results on both web and mobile interfaces.

MATERIALS AND METHODS: Our system employs semantic segmentation on images taken with mobile phones to effectively isolate the blood-filled portion of the drain from the rest of the image and compute the blood volume. These results are then sent to mobile and web applications for convenient access. To validate the accuracy and effectiveness of our system, we collected the Drain Dataset, which consists of 1,004 images taken under various background and lighting conditions.

RESULTS: With an average error rate of less than 5% in milliliters, our proposed approach achieves highly accurate blood level detection and estimation, as demonstrated by our trials on this dataset. The system also exhibits robustness to variations in lighting conditions and drain shapes, ensuring its applicability in different clinical scenarios.

CONCLUSIONS: The proposed automated blood volume estimation system can significantly reduce the time and effort required for manual measurements, enabling healthcare professionals to focus on other critical tasks. The dataset and annotations are available at: https://www.kaggle.com/datasets/ayenahin/liquid-volume-detection-from-drain-images and the code for the web application is available at https://github.com/itsjustaplant/AwesomeProject.git.

Key Words:
Jackson Pratt drain, Hemovac drain, Liquid volume detection, Clinical decision support, Deep learning, Semantic segmentation.

Introduction

A Jackson Pratt drain and a Hemovac drain are devices designed to collect excess bodily fluids during or after surgery. The measurement of blood in these drains traditionally involves emptying their contents into a measuring cup and meticulously recording the volume and time, a critical aspect of comprehending the surgical process. However, this manual process poses several inconveniences for both doctors and patients. Not only is it susceptible to measurement errors and potential spillage, but it also exposes individuals to bloodborne pathogens such as malaria, syphilis, brucellosis, and most notably, human immunodeficiency virus (HIV), hepatitis B virus (HBV), and hepatitis C virus (HCV). Furthermore, this process consumes valuable time for doctors and can be distressing for recently operated patients, impacting their well-being, as well as that of their caregivers and the healthcare team, due to hygiene concerns and time constraints.

In this study, we introduce a robust, user-friendly, and precise mobile system designed for estimating blood volume in drains through image analysis. Our system allows users to capture a photo with a mobile phone, which is then analyzed with advanced deep-learning techniques to detect the Jackson Pratt and Hemovac drains and also to determine the volume of the blood contained in these drains, as shown in Figure 1. The obtained results are then displayed through cloud storage systems, accessible both on a mobile phone and as a web application. This innovative system not only delivers quick and accurate results but also maintains a comprehensive data-

Corresponding Author: Mehmet Eren Yuksel, MD; e-mail: mehmeteren.yuksel@saglik.gov.tr

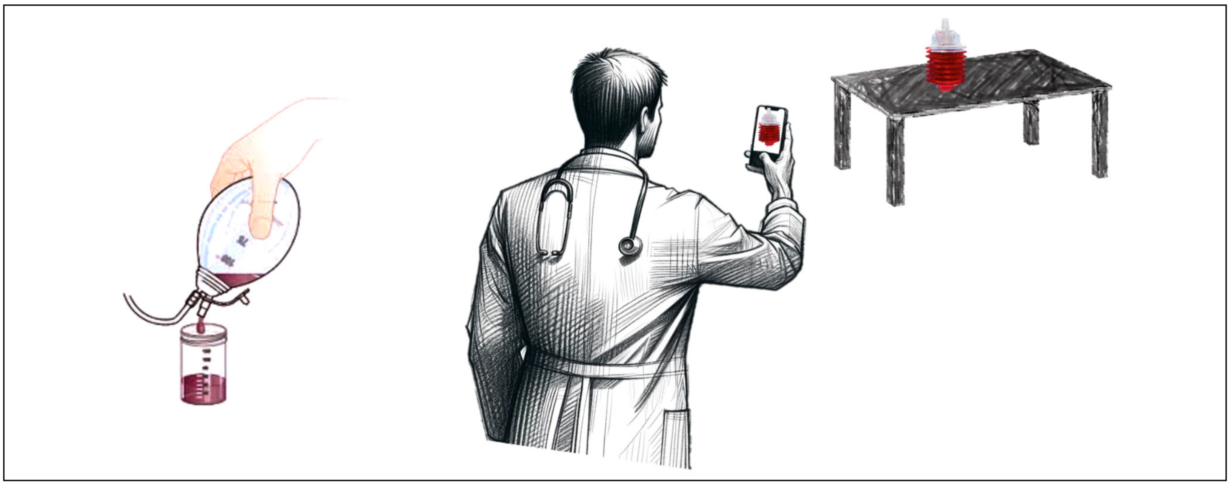

Figure 1. Pouring the contents of the drain into a measuring cup exposes the physicians and the patients to bloodborne pathogens, spillage, and measurement errors (*on the left*). Our proposed system introduces a complete framework to get the measurements from the images automatically and systematically (*on the right*).

base of measurements. Importantly, it performs effectively under various lighting conditions. To the best of our knowledge, this is the first study in the literature that estimates the blood volume in surgical drains using automated image processing techniques. Further, both the dataset and the source code of the system have been made available to contribute to the advancement and accessibility of this approach.

In the literature, quantifying liquid levels is an active research area with applications in various fields such as the chemistry and bottling industries[1,2], medical and hospital settings[3], as well as security and surveillance. Beyond standard cameras, researchers have developed several sensors and robotic systems tailored for different applications. For instance, a millimeter-wave Doppler sensor was designed to measure liquid levels with sub-millimeter accuracy, specifically for recording properties related to centrifuged blood in a blood-collection tube[3]. Subsequently, researchers combined acoustic and visual data to understand the manipulation of a container based on the sounds it generated[4]. Similarly, a motorized system was designed to position the camera close to a glass container, capturing multiple close-contact images[5]. These images were then utilized to quantify liquid levels in glass containers, for biomedicine applications.

On the other hand, the use of standard cameras is more common and less costly. Earlier works used edge, color, and gradient information to compute liquid levels. For transparent vessels, the detection of liquid levels was performed in chemistry applications[6]. For infusion bottles[7], image processing, and motion detection were used to detect the liquid levels. Similarly, a system based on edge detection was proposed to detect the end of drip infusion in a hospital setting, alerting nurses before the liquid runs out in the infusion bag[8]. Furthermore, several researchers have studied the reasoning behind liquid containers, such as content estimation and pouring prediction, especially in robotic settings[9-11].

Several datasets have been recently released to follow up on the demand for liquid detection research. Among these are (1) the Wine data[12]; (2) the general-purpose liquid containers dataset[13]; (3) the Vector-LabPics dataset for chemistry[14]; and (4) the TransProteus Dataset, a computer-generated dataset complemented with real-world data with depth maps[15]. The availability of such datasets paved the way for deep learning, which requires ample data.

Within the deep learning approach, semantic segmentation has been employed to identify liquids and their containers. Semantic segmentation is a computer vision technique that involves classifying and labeling each pixel in an image, providing a pixel-level understanding of the visual content. In two studies[14,16], instance and semantic segmentation were used for chemistry applications, hospitals, and medical labs. Our study also employs semantic segmentation but focuses on detecting and segmenting two types of drains in real-world environments with various back-

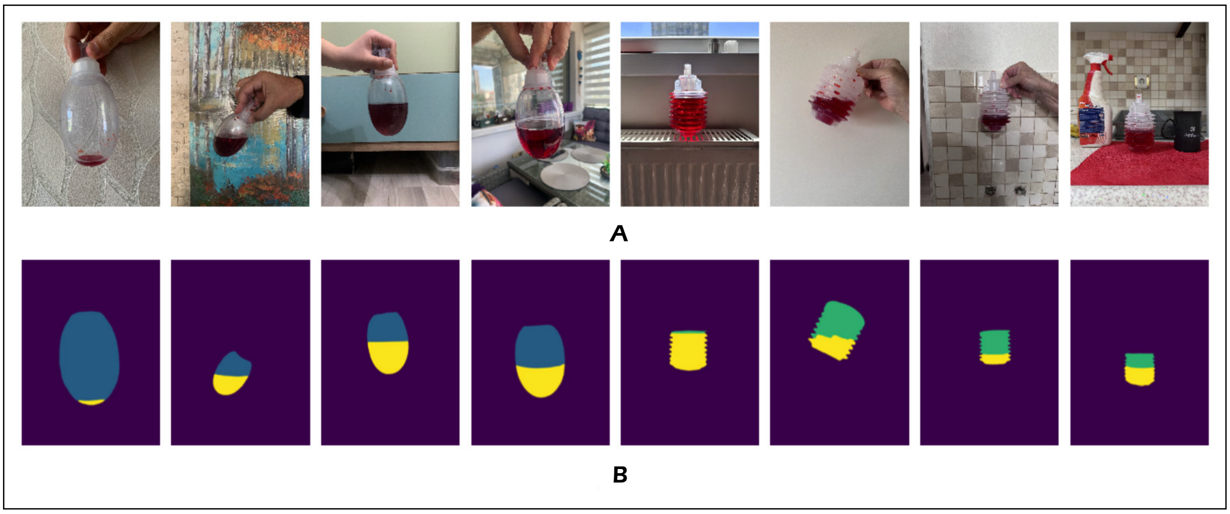

Figure 2. Examples from the Drain dataset with varying amounts of blood in several background and lighting conditions in (**A**), and their segmentation results with our proposed approach in (**B**).

grounds and lighting conditions. Additionally, it quantizes the blood with high precision, along with a web interface that provides online access to the tool for recording the outputs by the patient or their caregiver. In contrast to other studies, our computations were made using segmentation outputs instead of training at volume intervals to achieve the precision needed for medical purposes. Further, our study advances detection to a level where patients and their caregivers can use it through web services.

To compute the volume of the blood in the drain from the image, the exact location of the drain and the blood inside of it must be identified. To use the drain in the picture, all other objects were separated using semantic segmentation methods. Separate annotations were made for the image's drain and blood regions to train the semantic segmentation algorithms. Annotations were then used to make masks. These masks were used to train the pre-trained DeepLabV3[20] fully Convolutional Neural Network (CNN) model[17,18]. 1,004 photos were taken with various backgrounds, distances, and angles to train this model and to achieve high-accuracy results. The details of the dataset, the deep learning method, and the web application are given below.

Materials and Methods

Drain Dataset and Image Annotation

Our Drain Dataset consists of 1,004 pictures collected in two resolutions: 1,536 × 2,048 and 3,024 × 4,032, as shown in Figure 2. The photographs were taken between 10 to 40 cm away from the drain. The dataset was prepared on a variety of backgrounds, both inside and outside, in diverse lighting conditions. Further, to facilitate the generation of a dataset encompassing a broad range of volumes, drains were filled with varying volumes of theatrical blood with short gaps between the fillings. This variability in volume is a critical factor during the subsequent phase of volume calculations. Each image's number and the amount of liquid it contained were recorded. Out of the 1,004 pictures, 532 were from the Jackson Pratt drain and 472 were from the Hemovac drain. Of these, 99 images were reserved for testing the Jackson Pratt drain and 66 for testing the Hemovac drain. The rest of the photos were used in training, as outlined in Table I.

The Visual Geometry Group (VGG) image annotator (VIA)[19] was used to annotate the image manually. The drain and blood were annotated separately in each image using the VIA's poly-

Table I. Properties of the Drain dataset.

Type of drain	Volume (ml)	Training	Test
The Jackson Pratt drain	250	433	99
The Hemovac drain	400	406	66

gon region form. New photos, called masks, were generated to assign a color value of 1 to the marked areas and 0 to the remaining regions. During the blood annotation stage, diverse strategies were implemented to address issues arising from the angle at which the drain was held, resulting in an elliptical structure at the liquid's top while the camera is tilted up or down. This structural distortion presented challenges in the annotation process. To mitigate these challenges, only the front image of the blood, rather than the entire blood's body, is retained in the annotation.

Methodology

This research unfolds on two fronts. Firstly, it delves into an image processing pipeline designed to detect drains from cell phone images. This includes addressing challenges such as cluttered backgrounds, various surroundings with multiple objects, and potential tilting of the drain in different axes. The algorithms identify the drain and discern the blood within, computing its relative volume, ultimately deducing the actual volume in milliliters.

On the second front, the research introduces a cell phone application. This application connects to a cloud server housing the model parameters developed in the first phase. The framework of the system is given in Figure 3A. Within this framework, image processing takes place in the FastAPI part, and the other components are the databases, web servers, and mobile interfaces. The training for image processing follows the steps shown in Figure 3B. These aspects are elaborated further below.

Image Processing Pipeline for Detecting the Drain and Measuring the Blood

In this work, we focus solely on the blood quantization problem in two drain containers using the DeepLabv3 model[20] derived from the DeepLab architecture[17,18], which is a semantic segmentation model that results in annotations at the pixel level. DeepLabv3 is a fully Convolutional Neural Network (CNN) model, but unlike the traditional image classification CNNs, DeepLabv3 employs an encoder-decoder architecture that enables the model to make decisions for every pixel in an image. The encoder is constructed from CNNs, and it extracts the feature maps from the input image. The decoder gradually reconstructs the output to be the same dimension as the input

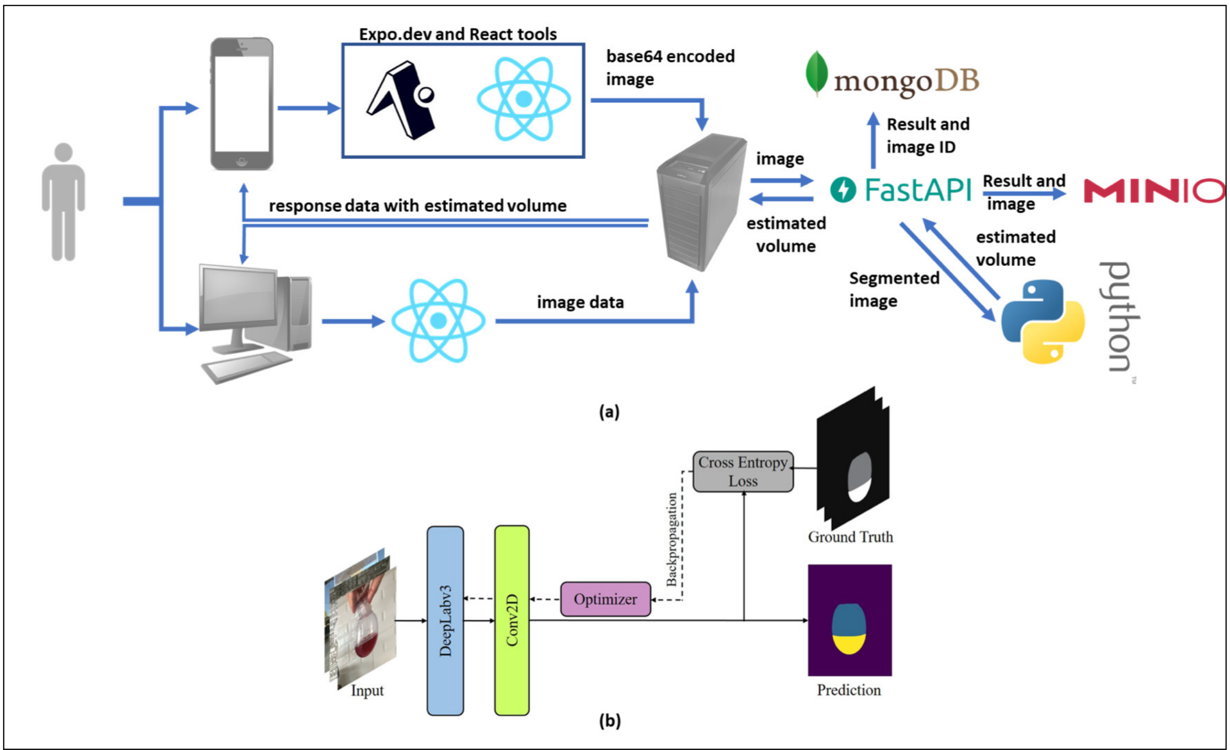

Figure 3. Framework of the system. The components of the web application are given in (**a**), and the architecture of training is shown in (**b**).

and uses upsampling to recover the details from the low-dimensional feature maps[17]. Within this process, the uniqueness of DeepLabv3 lies in its use of atrous convolutions that allow the model to capture multi-scale contextual features efficiently. The model further incorporates Atrous Spatial Pyramid Pooling, a mechanism for classifying regions of arbitrary scale. Since the drains in our database also appear on multiple scales, our solution uses the DeepLabv3 network to detect the drain and blood. Additionally, ResNet-50 served as the backbone. A well-known ResNet family architecture ResNet-50 includes 50 layers total depth, 48 convolution layers, 1 max-pooling layer, and an average pooling layer[21].

For training the DeepLabv3 model, both the blood and drain masks, as well as the original photos, were used. Four classes were obtained: Class-0 for the background, Class-1 for the Jackson Pratt drain, Class-2 for the Hemovac drain, and Class-3 for the blood. Figure 3B shows the architecture we utilize for training. The model was given images in the JPG file format, all set to 900x900 resolutions during the training. The model is trained with fixed-size images, but there are no size specifications for test images, as the model can handle images of any size in the test stage.

The learning rate is set to 1e-5, representing the step size of the gradient descent. Batch Size, or the quantity of photos used in each training iteration, is selected as 3. Annotation masks for randomly selected photos in each cycle are updated based on the class type: areas with blood are labeled as 1, Jackson Pratt drain locations as 2, and Hemovac drain locations as 3 in the annotation maps. Subsequently, 4D matrices are created by stacking multiple images. The last layer of the pre-trained model network is a convolution layer with 256 layers of input and 21 layers of output. We add one additional convolutional layer after the model because we only have 4 classes in our dataset, and we want to replace it with a new convolutional layer that has 4 outputs. The Adam algorithm is employed as an optimizer to manage gradient rates during the backpropagation step. Autograd is performed to initiate backpropagation from a variable, converting the data into gradient variables usable by the network. Predictions from the network are compared with ground truth data, and the loss is computed between predictions and annotations. Based on this loss function, gradients are calculated through backpropagation, and weights are updated. The model obtained after 10,000 rotations in the training loop is utilized in the outputs. The Tesla A100 GPU (https://www.nvidia.com/en-us/data-center/a100/) is employed in the Colab environment to train the model. The saved model weights are then utilized in our backend Application Programming Interface (API) to estimate the volume of blood in the Jackson Pratt or Hemovac drain.

Web Application Architecture

We propose an integrated system featuring Nginx, MinIO, MongoDB, FastAPI, React, and React Native that seamlessly combines web serving, data storage, databases, and interface development. This architecture is designed to provide a platform that efficiently manages large-scale data, facilitates image analysis, and delivers real-time results. Upon saving the model as a tensor, we deploy it to a cloud server to establish a robust and user-friendly web application system. Leveraging MinIO for file storage in the file system, FASTAPI for efficient RESTAPI creation, and React for a simple frontend, our backend evaluates uploaded images stored in the MinIO file system. The system then returns a response indicating the estimated blood level. Detailed explanations of each technology used in this work are provided below:

- Nginx (https://www.nginx.com) is a web server and reverse proxying software. As a web server, it plays a crucial role in responding to client requests and ensuring reliable page delivery. Additionally, Nginx operates as a reverse proxy, sitting between clients and web servers, effectively distributing incoming traffic and enhancing the overall performance and security of web applications.
- MinIO (https://min.io) is a high-performance object store built for large-scale data lakes. It can handle extensive volumes of data and serves a pivotal role in our system by efficiently storing both analyzed and uploaded images. By recording each data point on the system, all images can be used to detect faulty points and improvements.
- MongoDB (https://www.mongodb.com) is a flexible and scalable document database. It is used to store analyzed data in object IDs for further analysis of the system.
- FastAPI (https://fastapi.tiangolo.com) is a modern, high-performance microframework. Due to Python's speed and simplicity, we chose Python as our backend. When the FastAPI thread

is spawned, it runs our pre-trained model to detect blood and drain in the uploaded image. When the process is completed, each image is stored on MinIO, and our algorithms run to estimate the blood volume in the image.
- React (https://react.dev) is a library for creating native user interfaces. We chose to create our initial web prototype with React due to its simplicity. Our interface is simple yet ideal for uploading an image of the drain.
- React Native (https://reactnative.dev) is a library for creating native user interfaces for mobile applications. We chose React Native to reuse the components from the web application developed with React.

All of these components are connected, as shown in Figure 3A. The codes to connect these architectures are given at the GitHub link provided with the paper.

Results

We used the Intersection over Union (IOU) metric in Equation 1 to evaluate our model, measuring the similarity between the prediction and the ground truth. To calculate the IOU, both the ground truth area and the prediction area are necessary. In our setup, the masks act as our ground truth area, and the segmented output of a test image serves as our prediction area. The ratio of the overlap of these areas to their unions is presented in Table II. IOU values are computed individually for blood and for each type of drain.

$$\mathrm{IOU} = \frac{\text{Area of Overlap}}{\text{Area of Union}} \tag{1}$$

Jackson Pratt, Hemovac drains and blood can be easily identified by separating from the background, as shown by the outputs. Our four classes are also segmented in different colors. Volume calculation activities are carried out by proportioning for 250 ml Jackson Pratt and 400 ml Hemovac utilizing the segment colors.

Table II. IOU results of drains.

Type of drain	Drain	Blood
The Jackson Pratt drain	0.99242	0.99015
The Hemovac drain	0.98042	0.96788

As depicted in Figure 4, the pre-trained model underwent training for Jackson Pratt and Hemovac drains, utilizing blood and drain ground truth masks derived from the original photographs. Consequently, our model produces segmented outputs. The drains and blood can be readily identified by their clear separation from the background, as illustrated by the outputs. Our model assigns distinct colors to the four classes during segmentation. Volume calculation activities are performed by correlating the segment colors, specifically allocating 250 ml for Jackson Pratt and 400 ml for Hemovac drains.

Since we meticulously incorporate data from diverse angles into our dataset, tilting the drain does not impede detection. However, during volume estimation, the tilting angle influences the margin of deviation. Results of 170, 175, 177, 179, and 182 ml were obtained from images taken by altering the slope under consistent environmental conditions for a drain containing 167 ml of blood. In this scenario, angles range from 0 to 45. It is evident that the result deviates from the real value as the angle increases; therefore, it is recommended to avoid excessive tilting of the drain when using the application.

We generated test datasets by defining small volume intervals. Each volume is represented by a prediction. The link between the volume values generated by our model and the volume values of the images in the test dataset is depicted in Figure 5.

By comparing the ground truth value x_i and the predicted values x, we determined the residual standard deviation S_{res} as in Equation 2:

$$S_{res} = \sqrt{\frac{\sum(x_i - \bar{x})^2}{n-2}} \tag{2}$$

The discrepancy between the actual results' trend and the predicted data is known as the residual value. The challenges experienced during volume estimate are confirmed by the fact that the high-accuracy results displayed in our IOU table outperform the sigma value.

Discussion

Our results demonstrate that the prototype system can be used effectively for blood-level detection of various drains. This shows that deep

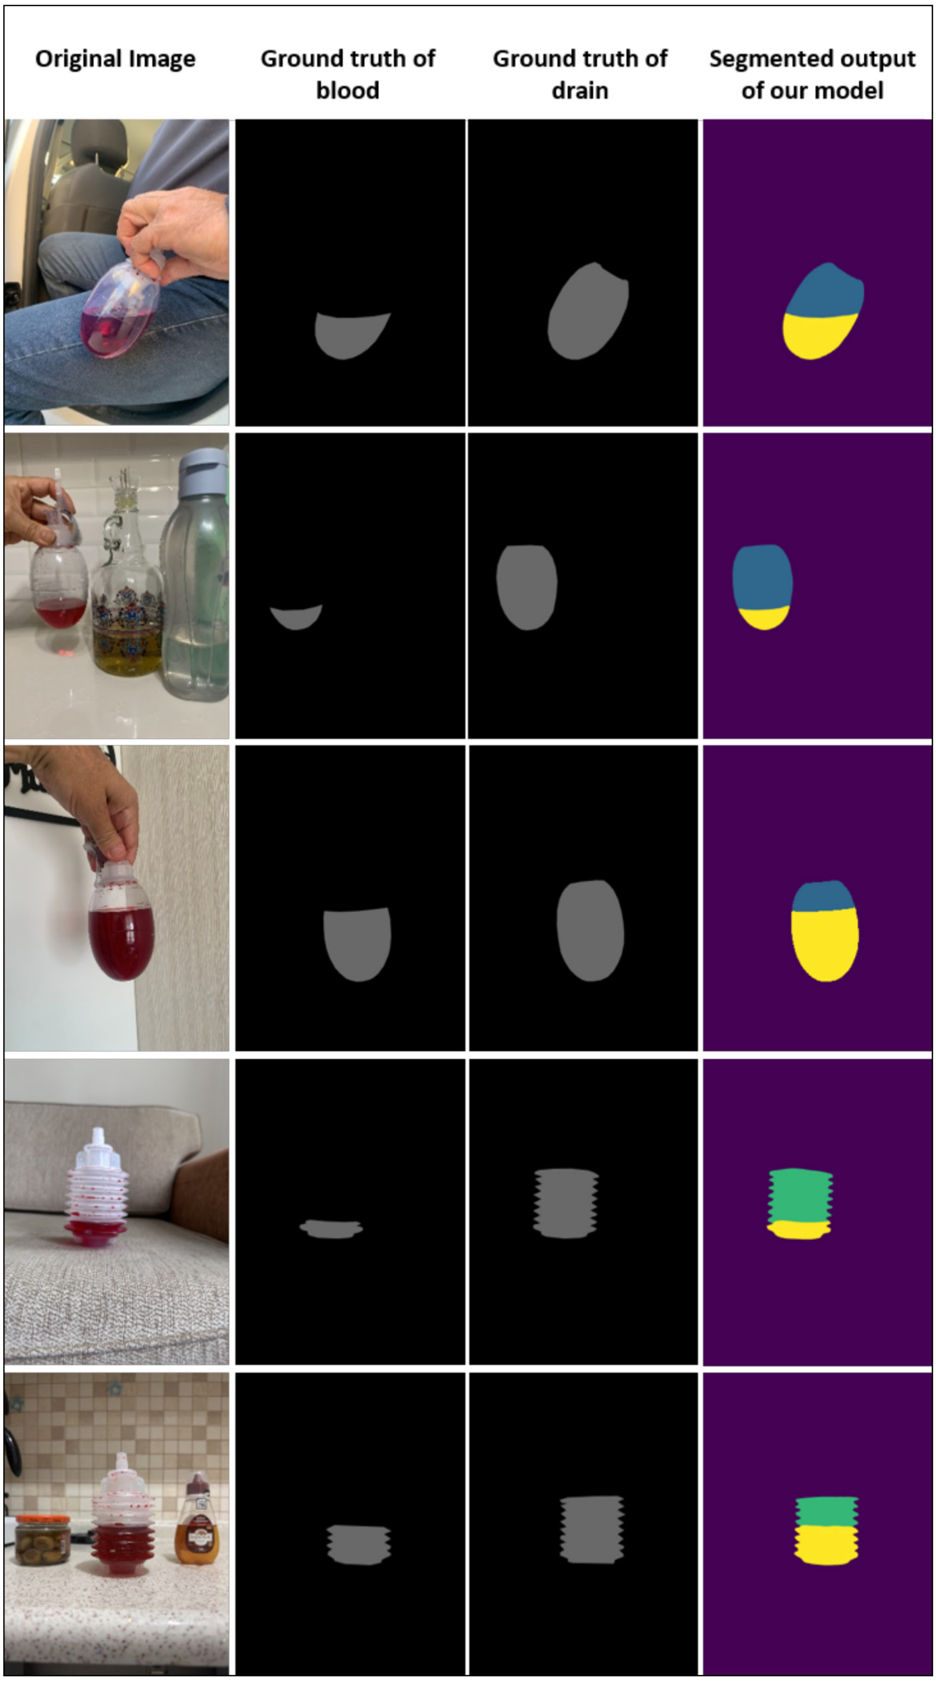

Figure 4. Summary view of segmentation.

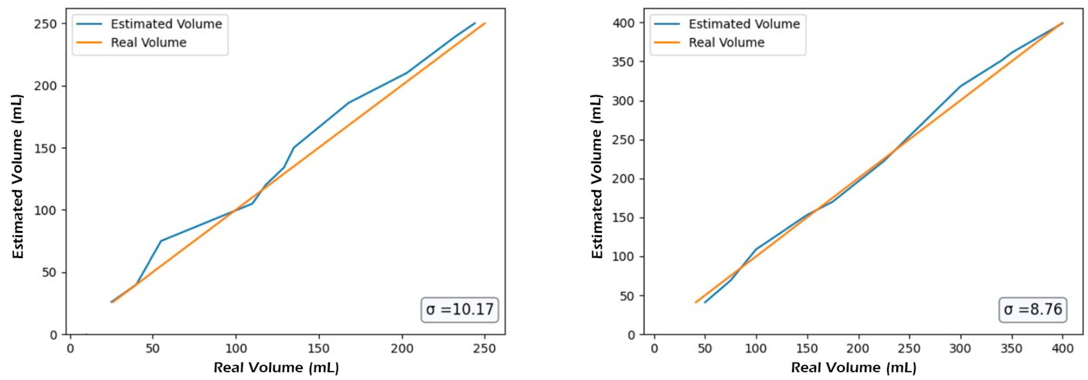

(a) Jackson Pratt Drain (b) Hemovac Drain

Figure 5. Results of blood volume estimation. **a**, The real volume of the Jackson Pratt drain and its estimated value. **b**, The real volume of the Hemovac drain and its estimated value. The residual standard deviation is shown with σ.

learning can significantly improve the monitoring and management of drains like Jackson Pratt and Hemovac, providing an automated and efficient alternative to manual measurements.

On the other hand, it is crucial to acknowledge the constraints of the prototype. The system's performance may be influenced by factors such as lighting conditions, drain positioning, tilt angle, network speed, and server-side issues, which would adversely affect the accuracy of the blood-level estimates. Our system can be further studied to overcome these constraints, both in developing deep learning to enhance its robustness to diverse scenarios and in developing a more robust real-time system to eliminate the listed limitations.

Conclusions

This paper presents a full framework utilizing a deep-learning approach for automated blood volume estimation in Jackson Pratt and Hemovac drains and also introduces the Drain Dataset to develop this framework. The system successfully estimates the fluid levels in these drains with a slight margin of error attributed to the drain's tilt angle. The user interface presents the fluid-level information clearly, facilitating easy interpretation and decision-making. Overall, our prototype highlights the potential of deep learning techniques in improving the monitoring and management of drains, offering a more efficient and automated solution for healthcare professionals. The proposed system demonstrates high accuracy, robustness, and potential to enhance patient care in postoperative settings. Future work involves refining and optimizing the system further to ensure its reliability and suitability for clinical applications, as well as integrating real-time alerts and notifications for timely interventions and improved clinical decision-making.

ORCID ID

Aleyna Aysen Sahin: 0009-0004-8729-8686
Mehmet Alperen Sahin: 0009-0003-1424-6917
Mehmet Eren Yuksel: 0000-0002-7110-0717
Seniha Esen Erdem: 0000-0002-8868-1132

Authors' Contributions

Aleyna Aysen Sahin and Mehmet Alperen Sahin collected the data, conducted the experiments, and drafted the work. Mehmet Eren Yuksel proposed the idea, designed the work, analyzed the results, and contributed to the manuscript. Seniha Esen Erdem conceptualized the idea, supervised the experiments, and reviewed the work. All the authors approved the submitted version of the manuscript.

Availability of Data and Materials

All the data and the codes of the web application are available at the following links, respectively: https://www.kaggle.com/datasets/ayenahin/liquid-volume-detection-from-drain-images and https://github.com/itsjustaplant/AwesomeProject.git.

Conflict of Interest

The authors have no conflict of interest.

Funding

We did not receive any funding for the study.

Ethics Approval

Not applicable. Theatrical blood was used, and no patient data was involved.

Informed Consent

Not applicable.

References

1) Music J, Stancic I, Dzaja B, Pekic V. Image-Based Sensor for Liquid Level Monitoring during Bottling with Application to Craft and Home-Brewing. Sensors (Basel) 2023; 23: 7126.
2) Cobo M, Heredia I, Aguilar F, Iglesias LL, García D, Bartolomé B, Moreno-Arribas MV, Yuste S, Pérez-Matute P, Motilva MJ. Artificial Intelligence to Estimate Wine Volume from Single-View Images. Heliyon 2022; 8: e10557.
3) Nakagawa T, Hyodo A, Kogo K, Kurata H, Osada K, Oho S. Contactless Liquid-Level Measurement with Frequency-Modulated Millimeter Wave Through Opaque Container. IEEE Sensors Journal 2012; 13: 926-933.
4) Xompero A, Donaher S, Iashin V, Palermo F, Solak G, Coppola C, Ishikawa R, Nagao Y, Hachiuma R, Liu Q, Feng F, Lan C, Chan R, Christmann G, Song JT, Neeharika G, Reddy CKT, Jain D, Rehman B, Cavallaro A. Multi-Modal Estimation of the Properties of Containers and Their Content: Survey and Evaluation. IEEE Transactions on Multimedia 2021; 1: 1-12.
5) Shi J, Cheng Y, Li B, Xu J, Hou X, Meng Y. Research on the Horizontal View Scheme of Automatic Capacity Verification of Working Glass Container. IEEE ICCCR 2021; 325-328.
6) Eppel S, Kachman T. Computer Vision-Based Recognition of Liquid Surfaces and Phase Boundaries in Transparent Vessels, with Emphasis on Chemistry Applications. arXiv preprint 2014; 1404.7174. Available at: https://arxiv.org/pdf/1404.7174.
7) Zhu H. New Algorithm of Liquid Level of Infusion Bottle Based on Image Processing. International Conference on Information Engineering and Computer Science 2009; 1-4.
8) Pranjoto H, Agustine L, Lestariningsih D, Werdani YDW, Andyardja W, Filipe JML, Atoum MFM, Yaro A. A Novel Approach of Controlling Stoppage of Drip Infusion Using Image Processing on Raspberry PI Platform. PPASA 2021; 58: 65-72.
9) Schenck C, Fox D. Detection and Tracking of Liquids with Fully Convolutional Networks. arXiv preprint 2016; 1606.06266. Available at: https://arxiv.org/pdf/1606.06266.
10) Mottaghi R, Schenck C, Fox D, Farhadi A. See the Glass Half Full: Reasoning about Liquid Containers, Their Volume and Content. Proceedings of the IEEE International Conference on Computer Vision 2017; 1871-1880.
11) Schenck C, Fox D. Perceiving and Reasoning about Liquids Using Fully Convolutional Networks. Int J R R 2018; 37: 452-471.
12) Cano MC. Liquid Volume Estimation from RGB Images (software) 2021. Available at: https://github.com/MiriamCobo/BrainGut-WineUp.
13) Wu Y, Ye H, Yang Y, Wang Z, Li S. Liquid Content Detection in Transparent Containers: A Benchmark. Sensors 2023; 23: 6656.
14) Eppel S, Xu H, Bismuth M, Aspuru-Guzik A. Computer Vision for Recognition of Materials and Vessels in Chemistry Lab Settings and the Vector-Labpics Data Set. ACS Central Science 2020; 6: 1743-1752. Available at: https://pubs.acs.org/doi/10.1021/acscentsci.0c00460.
15) Eppel S, Xu H, Wang YR, Aspuru-Guzik A. Predicting 3D Shapes, Masks, and Properties of Materials Inside Transparent Containers, Using the TransProteus CGI Dataset. Digital Discovery 2022; 1: 45-60. Available at: https://pubs.rsc.org/en/content/articlelanding/2022/dd/d1dd00014d.
16) Eppel S, Xu H, Aspuru-Guzik A. Computer Vision for Liquid Samples in Hospitals and Medical Labs Using Hierarchical Image Segmentation and Relations Prediction. arXiv preprint 2021; 2105.01456. Available at: https://arxiv.org/abs/2105.01456.
17) Chen LC, Papandreou G, Schroff F, Adam H. Rethinking Atrous Convolution for Semantic Image Segmentation. arXiv preprint 2017; 1706.05587. Available at: https://arxiv.org/abs/1706.05587.
18) Chen LC, Papandreou G, Kokkinos I, Murphy K, Yuille AL. Deeplab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. IEEE Transactions on Pattern Analysis and Machine Intelligence 2017; 40: 834-848.
19) Dutta A, Zisserman A. The Via Annotation Software for Images, Audio, and Video. Proceedings of the 27th ACM International Conference on Multimedia 2019; 2276-2279.
20) Pytorch: DeepLabv3. Available at: https://pytorch.org/hub/pytorch_vision_deeplabv3_resnet101/ (Accessed: 2 January 2023).
21) He K, Zhang X, Ren S, Sun J. Deep Residual Learning for Image Recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016; 770-778.