Interdisciplinary Platform for Bruise Image Research

Janusz Wojtusiak, PhD[1], Mohammad Qodrati, MD[1], Michał Markiewicz, PhD[2], Kiyarash Aminfar, MS[1], David Lattanzi, PhD[1], Katherine Scafide, RN, PhD[1]
[1]George Mason University, Fairfax, VA; [2]Jagiellonian University, Kraków, Poland

Introduction

WHO estimates 1 in 3 women worldwide (736 million) aged 15 years or older have experienced physical trauma at the hands of an intimate partner or stranger in their lifetime. In the US, CDC estimates over 10 million women and men experience intimate partner violence (IPV) annually. Bruises and other types of soft tissue trauma are the most common injuries reported among survivors of violence across the lifespan (children to elderly). Information learned from the assessment and documentation of these injuries has a significant impact on a violence survivor's health and legal outcomes. Undetected or misinterpreted injuries may go untreated or undiagnosed, with potentially dangerous effects. Existing research and clinical practice clearly indicate inequity due to difficulties in detecting and categorizing not only bruises but potentially any lesions on dark skin. That detection can be greatly improved using alternate light sources (ALS), and more specifically violet or blue light in the range 415nm-450nm[1].

Current research activities on bruise analysis and ALS are hampered by lack of data, difficulties in collecting bruise images and linked clinical data, and inability to collaborate across institutions. Access to very large datasets is needed especially when deep learning methods are applied to construct models for classifying images, but also when traditional statistical analyses are performed. To address this issue, we developed a platform that allows researchers to browse, search and compare bruise images, and that integrates with deep learning capabilities.

Methods

Functionality: The created interactive platform is designed to serve several purposes: (1) provide researchers and practitioners access to reference sets of bruise images; (2) provide access to structured data associated with images; (3) enable data collection as bulk imports from partner institutions as well as individual bruises through direct upload or a mobile app; (4) link to image annotation tools; (5) connect to deep learning-based decision support. The platform's functionality has been determined via structured interviews with five clinicians and reinforced by further user study. The clinicians also indicated the need to examine injury history, mechanism, cause, and age. Further, there is a need for structured data, including age, gender, body mass index, co-morbidities and medications that affect bruising.

Architecture: Shown in Figure 1, the platform architecture consists of web-based frontend (implemented in Next.js), main service/façade responsible for directing requests from frontend (NestJS), specialized microservices responsible for most of the platform functionality (NestJS and Python), data integrator service (NestJS), database systems (currently PostgreSQL), authentication service (NestJS), data loaders (Python) including EHR connectivity, deep learning system (Python), and external data labeling system (LabelBox). All components communicate through APIs and use JSON to encode data.

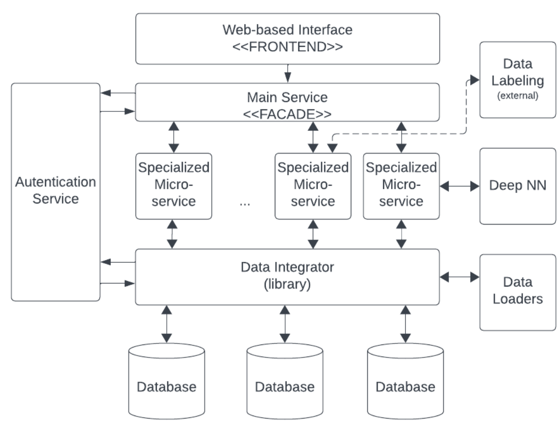

Figure 1: The bruise research platform architecture.

HIPAA compliance, security and privacy of patient data are the guiding principle in the design of the platform. It provides a 3-step process of authenticating access to data. The authentication service is called by the façade each time a request is made, then again when a call to database is made to prevent possible data leaks due to programming errors while implementing microservices. Further, the system implements permissions on the level of individual data elements stored in database. Only data elements that a user has access can leave the database system to reach the data integrator service.

Data from different organizations/studies are physically separated, thus the need for multiple databases, all of which have the same structure. One database keeps "public" data that is open to researchers. Access to organization specific databases can be granted. Data integrator service ensures seamless access to all databases without user knowing where specific data come from as long as access is granted.

Database: All data are stored in key-value database implemented within a relational database system. This hybrid approach allows for keeping relationships in relational form, and data in large field-less database. Key structure of the database is defined by a dedicated ontology constructed specifically for storing injury data, and that links to external terminologies (SNOMED, ICD-10, LOINC, RxNorm). Finally, the database is designed for longitudinal and multimodal bruise data in which the same bruise can be photographed and measured at different time points and using different ALS. Access to database is allowed only through stored procedures which allow for standardization of data, but also ensure security and fine-grained access control.

Deep machine learning: images and structured data from the platform are available for deep learning model construction and application. Neural networks are trained and applied on a separate computing platform with an API that connects them. Data from the platform can be processed in batches during the network training, but also individually when trained networks are being applied.

Results

The core functionality of the platform provides users with the ability to perform tasks described above. Figure 2 shows a screenshot from the main screen (a "tile view") of the platform used to browse through and search images of bruised with certain characteristics. It allows users to filter data by any characteristics of subject, bruise, measurement, and image. Users can also browse through subjects, and access specific bruise timelines.

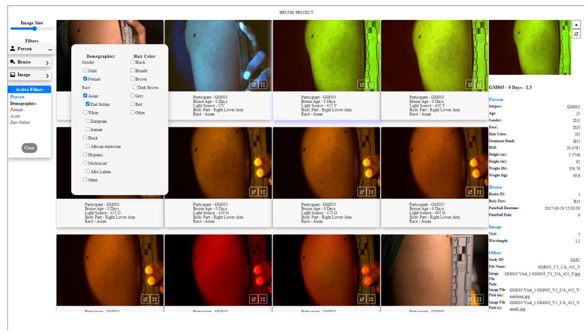

Figure 2: Tile view of bruises in the platform.

The platform is currently populated with about 30,000 images of bruises along with structured data[2]. This includes repeated images of bruises over 21 days post injury and control images from before bruising. Bruises are photographed using eleven combinations of light sources (natural light, ultraviolet, 415nm, 450nm, 515nm) and filters (yellow or orange). Baseline structured data currently included in database include subject demographics, subject measurements (height, weight, BMI, arm fat, arm circumference, etc.), and detailed measurements of skin color (bruise area and opposite side). At each visit, bruises are measured, and their color recorded (coded as L*a*b*). In total, about 90 structured data elements are included in addition to images.

Initial results showed suitable performance in classification and object detection tasks through transfer learning on the existing database. The initial framework for object detection employed YOLO V5 network. During the inference phase, the pipeline exhibited a notable capacity for rapid prediction, indicating its potential suitability for mobile devices. The initial pilot model, which was trained on 2,000 samples, achieved a mean average precision (AP) of 0.75 on a testing set and 0.95 on training set. In scenarios involving challenging test images, where even expert human analysis faced difficulty in identifying bruises, the model demonstrated a superior capability in certain intricate cases.

Discussion and Conclusions

Previous studies clearly indicate the need for using ALS in detecting bruises and collecting evidence. The presented work described a large project intended to build a platform for collecting, integrating, storing and analyzing image and structured data about bruises, with particular focus on integration with deep learning algorithms. The current database of about 30,000 images is being expanded by laboratory-collected control data as well as data from two hospital systems. In order to address equity and accuracy of bruise detection and assessment, even more data must be integrated into the platform and analyzed. It is particularly important to accommodate large variation of skin color, but also medical histories, medications, bruise types and physical characteristics of people. For detection and classification purposes, future phases of the study will pursue several avenues: the expansion of the image dataset utilizing more advanced augmentation techniques and the integration of consensus-based evaluation methods.

Acknowledgements. The authors express their gratitude to anonymous study participants who allowed to be induced with bruises to be photographed over a long period of time. Many aspects of the project have been completed by members of our research and development team at Mason. This project is supported by Award No. 15PNIJ-21-GG-04145-SLFO awarded by the National Institute of Justice, Office of Justice Programs, U.S. Department of Justice. The opinions, findings, and conclusions or recommendations expressed in this presentation are those of the author and do not necessarily reflect those of the Department of Justice. Work towards deployment of the platform is supported by a generous gift from an anonymous donor.

References

1. Scafide KN, Downing NR, Kutahyalioglu NS, Sheridan DJ, Langlois NE, Hayat MJ. Predicting alternate light absorption in areas of trauma based on degree of skin pigmentation: Not all wavelengths are equal. Forensic Sci Int. 2022 Oct;339:111410.
2. Equitable and Accessible Software for Injury Detection website: https://bruise.gmu.edu