

Insights on the development of *PRACTICE*, a research-oriented healthcare platform

Dragan Ahmetovic*, Alessio Angileri*[†], Sara Arcudi*[†], Claudio Bettini*, Gabriele Civitarese*,
Marco Colussi*, Andrea Giachi*[†], Roberta Gualtierotti*[†], Sergio Mascetti*,
Matteo Manzoni*, Flora Peyvandi*[†], Aiman Solyman*, Addolorata Truma*[†]

*Università degli Studi di Milano, Milan, Italy

[†]Fondazione IRCCS Ca' Granda Ospedale Maggiore Policlinico, Milan, Italy

dragan.ahmetovic@unimi.it, alessio.angileri@policlinico.mi.it, sara.arcudi@unimi.it, claudio.bettini@unimi.it,
gabriele.civitarese@unimi.it, marco.colussi@unimi.it, andrea.giachi@unimi.it, roberta.gualtierotti@unimi.it,
sergio.mascetti@unimi.it, matteo.manzoni2@studenti.unimi.it, flora.peyvandi@unimi.it,
aiman.solyman@unimi.it, addolorata.truma@unimi.it

Abstract—This paper describes the development of *PRACTICE*, a distributed healthcare technological platform that supports various research initiatives by the University of Milan and the Angelo Bianchi Bonomi Hemophilia and Thrombosis Center, Fondazione IRCCS Ca' Granda, Ospedale Maggiore Policlinico. *PRACTICE* includes three main components: a mobile app that patients can use to self-acquire ultrasound images at home, a computer-aided diagnosis web application that supports the practitioners through a set of machine learning models, and a set of web tools for image annotation, a prerequisite for training the machine learning models. Although *PRACTICE* was designed in the specific context of supporting the detection of joint recess blood effusions in hemophilic patients, this paper describes the main design and implementation challenges that apply to other applications of a research-oriented health platform.

Index Terms—Hemophilia; point-of-care; deep learning.

I. INTRODUCTION

The application of Artificial Intelligence (AI) methods in the medical domain is a research area investigated by a large number of research groups, due to its potential of revolutionizing diagnosis, treatment, and patient care [1].

Combining these advances with the latest trends in the Internet of Things (IoT) makes it possible to build advanced remote monitoring systems taking advantage of sensing devices such as wearable devices, physiological sensors, and smart home sensors [2]. Such systems have the goal of continuously and unobtrusively monitoring the health status of a patient with the long-term objective of improving the patient's quality of life and reducing health system costs.

Most of the existing studies in this area mainly focus on the data analysis aspects that are indeed crucial to provide clinicians with correct and complete information about the patient's health status. However, real-life deployment of these telemedicine systems requires the development of several tools that are rarely investigated in research papers. For instance, several medical domains require the monitored patient to collaborate in data collection (*e.g.*, self-collecting data) and this requires user-friendly applications. Similarly, clinicians

who receive AI-processed data from their patients require user-friendly applications that help them analyze the results to make informed decisions. Furthermore, in supervised settings, clinicians also need accurate and easy-to-use annotation tools that can be quickly adapted to research needs.

In this paper, we describe *PRACTICE* (Pilot on Remote AutomatiC ulTrasound scan analysIs for hemophiliC patiEnts), a distributed healthcare system designed in collaboration between computer scientists and clinicians to support the application of AI methods in the hemophilia domain. For patients with hemophilia, joint bleeding is a common complication that, if not treated promptly, can lead to recurrent bleeding, which ends with synovial hyperplasia, osteochondral damage, and hemophilic arthropathy [3]. UltraSound imaging is a practical approach to detect bleeding in the joint recess [4]. However, ultrasounds images are usually acquired by medical practitioners in specialized centers during outpatient visits, which can be difficult to schedule for both patients and the specialized centers. In *PRACTICE*, each hemophilic patient is provided with a portable ultrasound system. When necessary (*e.g.*, a routine check or in case of pain), the patient uses the probe to acquire ultrasound images of the joints that are automatically transmitted to the specialized center where a medical practitioner remotely assesses the presence of joint bleeding supported by state-of-the-art AI methods (such as [5], [6]).

The *PRACTICE* system combines several tools:

- GAJA (Guided self-Acquisition of Joint ultrAsound images), an application to guide patients to autonomously acquire joint ultrasound images with a portable probe.
- CADET (Computer-Aided Diagnosis for hEMarThrosis), an application leveraging AI methods to support clinicians in formulating a diagnosis.
- ATOM (Annotation Task Orchestrator Module), a system for the annotation of ultrasound images targeted to clinicians.

In this paper, we report on our experience in designing and implementing PRACTICE and its components. We also report the lessons learned in this ongoing project.

II. ANALYSIS OF REQUIREMENTS

PRACTICE is the result of a multi-year collaboration between two teams of researchers, one from the Computer Science Department of the University of Milan, and the other from the Angelo Bianchi Bonomi Hemophilia and Thrombosis Center, Fondazione IRCCS Ca' Granda, Ospedale Maggiore Policlinico, also affiliated with the Department of Pathophysiology and Transplantation of the University of Milan. The collaboration involved multiple funded projects and various research goals, with the overarching objective of supporting the diagnosis process and the follow-up monitoring of joint recess effusions in patients with hemophilia. In this context, the key functional requirements of the platform are:

- Supporting medical practitioners in the diagnosis process and the follow-up monitoring of joint healthcare of the patients through an interactive computer-aided diagnosis tool that shows ultrasound images collected by the patients and estimates the presence of joint recess effusion.
- Providing guidance to patients with hemophilia for the self-acquisition of ultrasound scans of their joints using a portable ultrasound probe through an application running on a tablet device that detects anatomical markers of the joint and interactively instructs the user on how to move the probe to correctly scan the recess.
- Facilitate the practitioners in annotating the presence of recess effusion and outlining the recess in the images collected by the patients in order to train the computer-aided diagnosis tool to better recognize recess effusion.

In addition to the medical practitioner and the patient, we also identify two supporting figures, along with their roles:

- The system administrator that manages the users of the platform, assigns the annotation tasks to the medical practitioners, and monitors the completion progress of the annotation tasks.
- The data scientist who uses the annotated images to train the machine learning models.

There are also three non-functional requirements that are relevant for the system design:

- The entire decision process, starting with the acquisition of the ultrasound images by the patient and concluding with the determination of the diagnosis by the medical practitioner using the computer-aided diagnosis tool, should not have a longer duration, for the physician, than the usual practice, with the patient going to the hospital for an in-person visit.
- Since the process involves the remote acquisition of ultrasound images, their transmission to the hospital servers, and their usage in the annotation system, the training of the machine learning model, and the computer-aided diagnosis tool, it is crucial to guarantee the patients' privacy at all stages of the process.

- Given the research-oriented nature of the project, the data scientist can be interested in exploring various ML models. This requires high flexibility in the data annotation process.

III. SYSTEM ARCHITECTURE AND TECHNOLOGIES

Figure 1 shows PRACTICE system architecture. The system is composed of the PRACTICE server, the hospital ultrasound device, the GAJA app running on Windows tablet computers and connected to a portable ultrasound probe, and two web applications: CADET and ATOM.

The hospital ultrasound device is a closed system that does not have a publicly available Software Development Kit (SDK). This means that it is not possible to develop ad-hoc applications using the hardware of the ultrasound device. To the best of our knowledge, this is common for most ultrasound devices. Therefore, we integrated this device by leveraging its pre-installed application and configuring it so that, at the end of each visit, it automatically saves the media (images and videos) in a folder on the PRACTICE server. A daemon running on the PRACTICE server watches for changes in that folder and, when it observes a new file, loads the media and its associated metadata (e.g., date of visit) on the database (*main-DB*) through *main-API*, a set of REST APIs implemented through a Node server.

The other three clients (GAJA, CADET, and ATOM) interact directly with *main-API* to store and retrieve data from *main-DB*. All three clients also share a common problem: preserving patients' privacy. To address this issue, the PRACTICE system adopts a pseudonymization approach: all data and media related to a patient are associated with a pseudo-identifier as soon as they are stored in the *main-DB*. All operations related to pseudonymization are implemented by the *pseudonymization-API*, a set of rest APIs that store data in the *identities-DB*, a separate database with higher security (restricted access). In the following processing, the media is associated with the pseudo-identifier, unless the real patient's name is required (e.g., by the practitioner during a visit). In these cases, client applications can access the name through *pseudonymization-API* that implements a role-based access control policy (e.g., the practitioners can access the patients' names, while data scientists cannot).

Finally, there are two other components worth mentioning. The first is *ML-API*, which provides access to the machine learning models through a set of REST APIs available only for local calls and implemented in Python. The second is a set of instances of various annotation tool services. As detailed in Section IV, ATOM orchestrates various third-party annotation tools, each running with its own instance (and possibly its own database) and interacting with *main-API*.

IV. IMAGES ACQUISITION AND ANNOTATION

Deep learning algorithms rely on large datasets to effectively learn to generalize patterns of various pathological conditions or to identify areas of interest. However, a public dataset of ultrasound media is not available for the considered medical

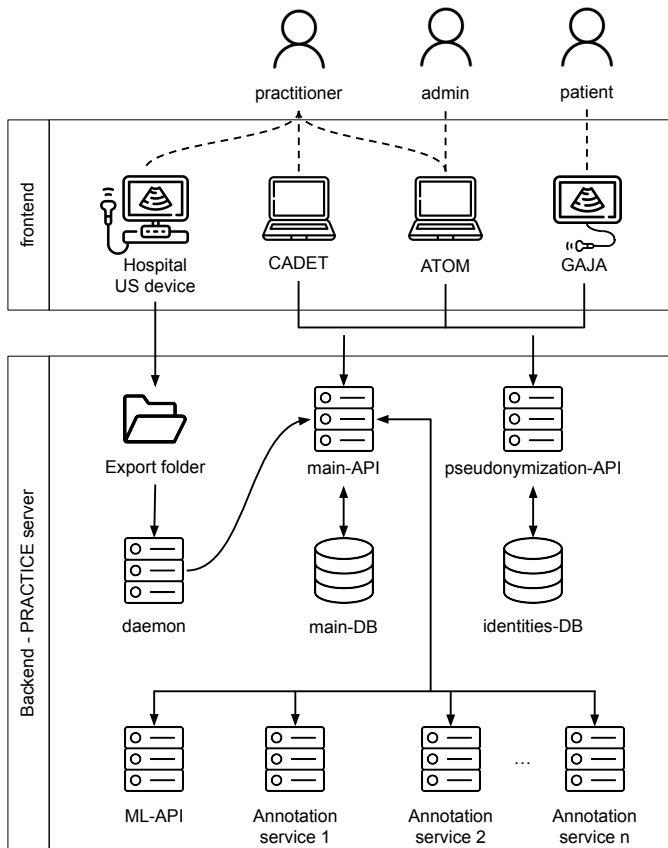


Fig. 1: Overview of PRACTICE architecture

domain. Therefore, we created a new dataset by collecting ultrasound media from hemophilic patients. Data was collected by expert practitioners during hospital visits, using a high-end ultrasound device. Since media is acquired during regular visits, the overall procedure was designed to avoid additional workload for the practitioners and inconvenience to patients.

One problem that emerged during the creation of the dataset is related to the fact that ultrasound imaging is highly dependent on the operator and has a high inter-patient variability hence making the acquired data highly heterogeneous, a factor that can negatively impact the training of the machine learning models. To mitigate this problem, we defined an acquisition protocol based on the following principles [5], [7].

- **Inclusion criteria.** The media of patients with significantly different characteristics (at the level of musculoskeletal ultrasound imaging) are excluded from the dataset. For example, children and patients with prostheses are excluded.
- **Standardization.** By adopting well-established procedures in the medical literature and practice, we defined a standard procedure for image acquisition. This includes, for example, the set of joints to consider and, for each of them, the set of scans¹.

¹A scan defines the probe position and, consequently, which anatomical targets are framed in the ultrasound image.

- **Parameters definition.** When acquiring an ultrasound media, the practitioner can tune several settings (e.g. power and frequency). We selected a fixed value for most of these settings, leaving the practitioner with the ability to select only a few parameters whose value has to be defined specifically for each patient (e.g., the “depth” value).

After dataset acquisition, we defined a set of tools and practices for data annotation. The guiding objectives were to reduce the annotation time and errors. To achieve these objectives, we initially developed an ad-hoc annotation tool. However, we then realized that the research activities frequently required the creation of new annotation tasks. For this reason, we designed the ATOM (Annotation Task Orchestrator Module) system that allows the administrator to quickly create a new annotation task by specifying the following data.

- The set of media from the dataset.
- The annotation tool, a third-party application. For example, for some annotation tasks, we used *Label Studio* [8] that, for privacy reasons, we configured to run on our server. These tools automatically transmit the annotations to PRACTICE, which stores them.
- The type of annotation (e.g., the set of classes).
- The set of annotators (i.e., practitioners).

The system was designed to interact with any compatible annotation tool, including those for creating image class annotations, bounding boxes, and segmentations. In addition to creating annotation tasks, ATOM also provides two main functions. One function is designed for annotators, who can access the list of tasks assigned to them and run the annotation tool. The other function is designed for the administrator to monitor the completion of the annotation tasks and to check the inconsistencies among the annotators. Specifically, for each annotation task, the administrator can define one or more equality functions. Then ATOM uses these functions to create a confusion matrix that shows, for each pair of annotators, the percentage of images (among those annotated by both annotators) that have the same annotation (according to the equality functions). Figure 2 shows an example screen of the admin panel, On the left we can see the list of active tasks, with the progress. On the right a detailed view of a single selected task, where the progress is divided among the different practitioners and the tables on the bottom report statistics of the currently annotated data for various equality functions.

V. GAJA: SELF-ACQUISITION OF ULTRASOUND IMAGES

GAJA is an application that assists patients in self-acquiring ultrasound images of their joints for remote diagnosis by medical practitioners. Current approaches for patient-acquired ultrasound images are based on two paradigms: patient training to acquire ultrasound images on their own [9] or real-time remote guidance of patients by an expert during ultrasound image acquisition [10]. However, both approaches have limitations that impede their applicability in the considered scenario.

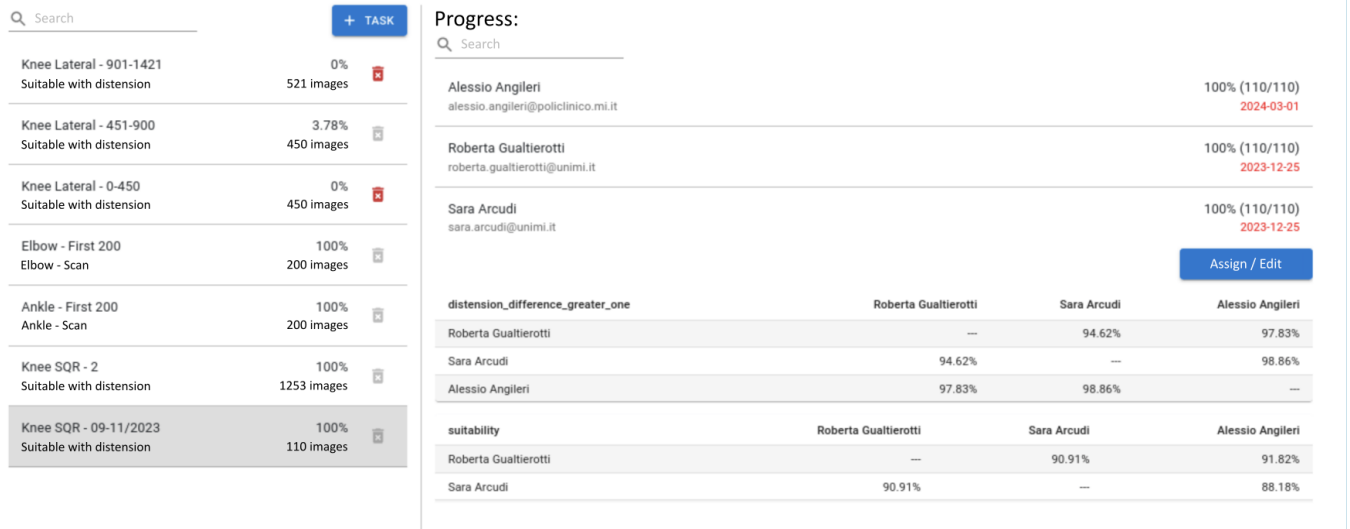


Fig. 2: ATOM admin example screen

The first approach requires extensive training for patients to acquire good images independently, and patients tend to forget the correct practices over time [11]. In the second approach, real-time guidance is a time-consuming process, and it requires the concurrent availability of a medical practitioner, which is a major constraint and expense for the hospital. The motivation behind GAJA is to address these limitations, providing real-time automated guidance for ultrasound image acquisition to the patient, without requiring support from medical practitioners.

To correctly acquire ultrasound images of a body part, an operator should follow established protocols that define how to set the probe parameters, how to position the probe, and how to interpret the ultrasound video output [12]. GAJA aims to automate some of these actions and to provide real-time guidance for actions that cannot be automated, which normally require specialized training and medical knowledge. Furthermore, it provides reminders and tutorials for those steps that are easy to explain. This makes GAJA a functional and usable system for remote self-acquisition of ultrasound images. Currently, GAJA runs on a Windows tablet and uses an SDK to access the stream of frames from a portable ultrasound probe.

To meet GAJA objectives, we defined a collaborative process between patients and practitioners based on two main steps [13].

a) Reference image acquisition: During the patient's initial visit at the hospital, a medical practitioner collects a *reference image* for each target joint. Using an object detection model, GAJA detects, for each reference image, a predefined set of anatomical markers (e.g. patella and femur for the knee). GAJA stores this information together with the probe parameters used by the practitioner, such as scan depth and gain. Finally, the patient receives a short training (approximately 10 minutes) to learn how to use GAJA.

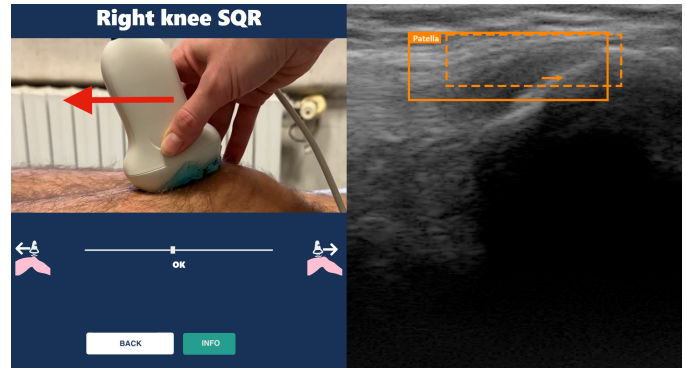


Fig. 3: Ultrasound images acquisition guidance interface

b) Self-acquisition: The patient can initiate remote visits at home, as periodical checkups or on-demand in case of pain or trauma. During a remote visit, the patient selects the target joint and completes a short questionnaire on their health status. Brief tutorial videos (a few seconds long) are used to remind the patient of the key steps of the process, like putting the gel on the probe and positioning the probe correctly. The patient then proceeds to acquire ultrasound images, guided by the interface shown in Fig. 3. The interface screen is divided into two areas: the right section displays the feed acquired from the probe and shows the anatomical markers detected in real time, the target positions to which they need to be aligned, and the arrows indicating the direction of the movement needed to align the markers; the left section of the screen shows example images illustrating how to position and move the probe. When the anatomical markers are aligned, a short sound cue is played, and their bounding box is coloured green. The user is advised to hold still to avoid blur in the acquired images, and, after a few seconds, ultrasound images are collected and

sent to the PRACTICE server.

VI. CADET: COMPUTER-AIDED DIAGNOSIS

CADET is a web-based interface that supports clinicians in formulating the diagnosis; it manages both in-presence and remote visits. To design CADET we first analyzed the habitual visit procedure adopted by practitioners without the support of a computer-aided diagnosis system. The physicians used to collect media with the ultrasound probe and then enter the diagnosis of blood effusion in a word processor file, following a template that defines a set of information for each joint [12]. The diagnosis was finally uploaded to the national health system server and, after printing, stored in the patient's physical medical record.

This procedure had several limitations. First, the media and exam data were not linked, making access to the patient's medical history (complete with diagnosis and the media) impractical. This affects practitioners, who need to review the stored data during follow-up visits, and also makes it impossible to use the data for the training of ML models. Second, some operations required the practitioner's intervention although, in principle, they could be automated. This included, for example, the creation of the diagnosis on the word processor. Finally, no CAD system was implemented and remote visits were not possible.

We initially designed a first CADET prototype in which the practitioner could use the web app to automatically acquire media from the ultrasound probe. This solution was designed with the idea that the practitioner could quickly switch from CADET to the probe. However, due to technical limitations of the ultrasound probe (no SDK is available), this was not possible. Therefore, we designed a solution in which the practitioner first acquires media using the ultrasound probe and then interacts with CADET to formulate the diagnosis. The practitioner first completes an initial general medical history through a guided questionnaire and then selects the joint, one at a time. The diagnosis of each joint is divided into four steps: media selection (Figure 4a), joint-specific history, a questionnaire related to the standard HEAD-US procedure [12], and a guided questionnaire for the diagnosis of blood effusion (Figure 4b).

After completing the process for each joint, the practitioner can access the final diagnosis that follows the same format as the word processor template. This report can then be uploaded to the national health system server and possibly printed for physical storage.

The remote diagnosis procedure is similar, with the main difference that some information is already available (media and history).

CADET adopts two main solutions to support the practitioner. First, it implements a knowledge-based system to guide the practitioner in diagnosis formulation. This solution was first designed in terms of a decision tree in which each node is a Boolean condition and each leaf is the joint-specific medical report. CADET implements this decision tree through a questionnaire (see Figure 4b) in which some

answers are automatically provided based on the data inserted in the previous steps (*e.g.*, whether the patient has pain) and the remaining are provided by the practitioner. The second solution adopted to support the diagnosis is to automatically detect recess distention [5], which is a necessary condition for blood effusion. Taking into account the media available for a given joint, the system suggests a distention value on a scale of four possible alternatives (see Figure 4b). The practitioner can then decide to accept the suggested value or to change it.

Several solutions were also adopted to speed up the process. First, CADET automatically pre-selects the media obtained from a visit based on a ML solution that identifies, for each media, the scan, the joint name, and its laterality. For each joint, some data are precomputed on the basis of previous visits and the patient's medical history. For example, for each joint, the practitioner has to specify whether there is a prosthesis. If the practitioner specifies that there is one during a visit, the system automatically loads the same value during the following visits. Finally, CADET automatically generates the diagnosis file that can be uploaded to the national system.

VII. CONCLUSIONS

This paper presents PRACTICE, a healthcare system specifically designed to support hemophilic patients and the medical practitioners assisting them. The system was also designed with a third main actor in mind: the data scientist who uses the collected data to train new ML models. This required defining medical procedures and technical solutions for the acquisition, annotation, and storage of ultrasound media.

Two PRACTICE components are currently being used: CADET supports the practitioner during visits and ATOM makes it possible to assign annotation tasks to practitioners. The third component, GAJA, is currently in an advanced prototyping phase, and experiments with patients are expected to begin in the next few weeks.

ACKNOWLEDGMENTS

This research is partially supported by the MUSA (Multi-layered Urban Sustainability Action) and by the FAIR (Future Artificial Intelligence Research) projects both funded by the NextGeneration EU program. The study was partially supported by the Italian Ministry of Health – Bando Ricerca Corrente. The Fondazione IRCCS Ca' Granda Ospedale Maggiore Policlinico is a member of the European Reference Network (ERN) EuroBloodNet.

REFERENCES

- [1] P. Rajpurkar, E. Chen, O. Banerjee, and E. J. Topol, "Ai in health and medicine," *Nature medicine*, vol. 28, no. 1, pp. 31–38, 2022.
- [2] N. Y. Philip, J. J. Rodrigues, H. Wang, S. J. Fong, and J. Chen, "Internet of things for in-home health monitoring systems: Current advances, challenges and future directions," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 2, pp. 300–310, 2021.

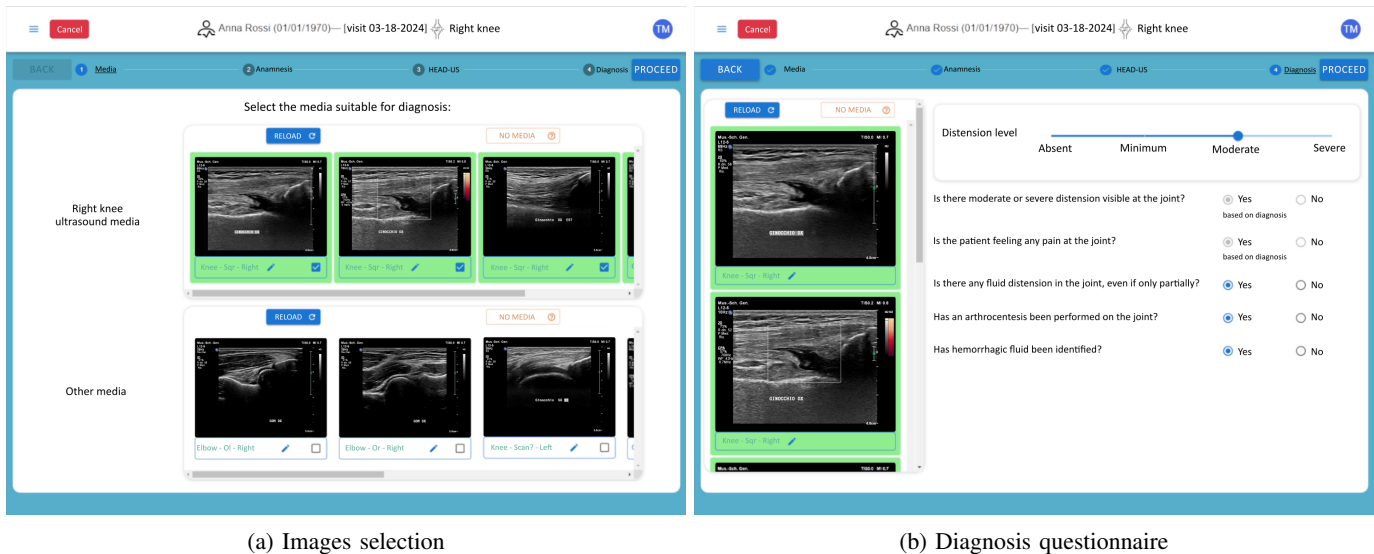


Fig. 4: CADET interface

- [3] M. W. Hilgartner, “Current treatment of hemophilic arthropathy,” eng, *Current opinion in pediatrics*, vol. 14, no. 1, pp. 46–49, 2002, ISSN: 1040-8703.
- [4] D. Plut, B. F. Kotnik, I. P. Zupan, *et al.*, “Diagnostic accuracy of haemophilia early arthropathy detection with ultrasound (head-us): A comparative magnetic resonance imaging (mri) study,” eng, *Radiology and oncology*, vol. 53, no. 2, pp. 178–186, 2019, ISSN: 1581-3207.
- [5] M. Colussi, G. Civitarese, D. Ahmetovic, *et al.*, “Ultrasound detection of subquadricepsal recess distension,” *Intelligent Systems with Applications*, p. 200 183, 2023.
- [6] R. Gualtierotti, S. Arcudi, A. Ciavarella, *et al.*, “A computer-aided diagnosis tool for the detection of hemarthrosis by remote joint ultrasound in patients with hemophilia,” *Blood*, vol. 140, no. Supplement 1, pp. 464–465, 2022.
- [7] R. Gualtierotti, L. P. Solimeno, F. Peyvandi, *et al.*, *Ultrasound evaluation of hemophilic arthropathy: A proposal of definitions in a changing landscape*, 2024.
- [8] M. Tkachenko, M. Malyuk, A. Holmanyuk, and N. Liubimov, *Label Studio: Data labeling software*, Open source software available from <https://github.com/heartexlabs/label-studio>, 2020-2024. [Online]. Available: <https://github.com/heartexlabs/label-studio>.
- [9] N. M. Duggan, N. Jowkar, I. W. Ma, *et al.*, “Novice-performed point-of-care ultrasound for home-based imaging,” *Scientific Reports*, vol. 12, no. 1, p. 20 461, 2022.
- [10] Y. Baribeau, A. Sharkey, O. Chaudhary, *et al.*, “Hand-held point-of-care ultrasound probes: The new generation of pocus,” *Journal of cardiothoracic and vascular anesthesia*, vol. 34, no. 11, pp. 3139–3145, 2020.
- [11] P. Aguero, R. F. Barnes, A. Flores, and A. von Drygalski, “Teleguidance for patient self-imaging of hemophilic joints using mobile ultrasound devices: A pilot study,” *Journal of Ultrasound in Medicine*, vol. 42, no. 3, pp. 701–712, 2023.
- [12] C. Martinoli, O. D. C. Alberighi, G. Di Minno, *et al.*, “Development and definition of a simplified scanning procedure and scoring method for haemophilia early arthropathy detection with ultrasound (head-us),” *Thrombosis and haemostasis*, vol. 109, no. 06, pp. 1170–1179, 2013.
- [13] M. Colussi, S. Mascetti, D. Ahmetovic, *et al.*, “Gaja-guided self-acquisition of joint ultrasound images,” in *International Workshop on Advances in Simplifying Medical Ultrasound*, Springer, 2023, pp. 132–141.