IEEE Access
Multidisciplinary : Rapid Review : Open Access Journal

Received 23 May 2023, accepted 15 July 2023, date of publication 28 July 2023, date of current version 15 August 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3299866

APPLIED RESEARCH

Technical Feasibility of Implementing and Commercializing a Machine Learning Model for Rare Disease Prediction

GEORGE KOUTITAS[1], (Member, IEEE), KIMBERLY NOLEN[2], SEPIDEH ATTAL[3], ANASTASIOS VENTOURIS[1], YINNON DOLEV[4], AND HANS THIJS VAN DEN BROEK[5]

[1]Pfizer Center for Digital Innovation, 55535 Thessaloniki, Greece
[2]Medical Affairs, Pfizer, New York, NY 10017, USA
[3]Medical Affairs, Pfizer, 75014 Paris, France
[4]Diagnostic and Pathway Informatics, Philips, Haifa 31004, Israel
[5]Diagnostic and Pathway Informatics, Philips Healthcare, 5656 AE Eindhoven, The Netherlands

Corresponding author: George Koutitas (george.koutitas@pfizer.com)

This work was supported by Pfizer.

ABSTRACT Wild-type transthyretin amyloid cardiomyopathy is an under-recognized cause of heart failure. Pfizer previously developed a machine learning model that performed well in identifying wild-type transthyretin amyloid cardiomyopathy vs. nonamyloid heart failure. However, challenges exist when introducing machine learning applications into healthcare, mainly due to restrictions on sharing patient data. This requires the triggering and execution of the model outside the developer's infrastructure using hosts with diverse information technology capabilities. With these barriers in mind, we investigated architectural designs to facilitate the delivery of the model to the customer. We considered manageability and scalability of the model, the host's information technology maturity, maintenance of patient data privacy, and protection of Pfizer's intellectual property. A container-based design, wherein the application is shipped as a container image to a third-party platform, fulfilled these criteria and was piloted on a platform hosted by Philips. Objectives of this pilot included defining and testing the architectural design and technical parameters for sharing the container image, creating a scalable and modular framework to manage multiple applications on different third-party platforms, and exploring a communication pattern based on clinical decision support Hooks/Cards and representational state transfer calls within the Philips platform. Implementing the model may enable earlier identification and treatment of wild-type transthyretin amyloid cardiomyopathy, and learnings from this pilot may lead to improved delivery of other machine learning models to healthcare providers, thereby increasing utilization. This article also presents an overview of architectural designs that may help others adopt new methodologies and ideas for machine learning model commercialization.

INDEX TERMS Artificial intelligence, commercialization, data privacy, machine learning, transthyretin amyloid cardiomyopathy.

I. INTRODUCTION

Transthyretin amyloid cardiomyopathy (ATTR-CM) is a fatal disease caused by the deposition of transthyretin (TTR)-derived amyloid fibrils in the myocardium [1], [2]. ATTR-CM is becoming increasingly recognized as a cause of heart failure, and the wild-type form (ATTRwt-CM) appears to be more common than the hereditary type [3]. However, ATTR-CM is often misdiagnosed or diagnosed late in the disease course due to low disease awareness, fragmented knowledge among different specialists and subspecialists, and the multisystemic nature of the disease [4], [5], [6], [7], [8]. Furthermore, whereas the hereditary form can be confirmed by genotyping, there is no method of systematic identification available for ATTRwt-CM [9]. Treatments that slow or halt disease progression, by suppressing the expression of,

The associate editor coordinating the review of this manuscript and approving it for publication was Bing Li.

or stabilizing, the TTR protein, have recently been made available [10], [11]. Therefore, early identification of ATTR-CM has become an important goal, to facilitate intervention before irreversible heart failure develops.

Pfizer recently developed a random forest machine learning (ML) model that provides a systematic framework for the screening and identification of patients at risk for ATTRwt-CM [12]. The ATTRwt-CM ML model was derived using the medical claims data of 1071 patients with ATTRwt-CM and 1071 nonamyloid heart failure controls; all data were from individuals aged ≥50 years. The model successfully identified ATTRwt-CM vs. nonamyloid heart failure with high sensitivity, specificity, and accuracy in US medical claims data (87%, 87%, and 87%, respectively) from IQVIA, Inc. (Durham, NC, USA) and in US electronic health record (EHR) data (90%, 79%, and 84%, respectively) from Optum, Inc. (Eden Prairie, MN, USA) [12].

Artificial intelligence (AI) and ML have been used to support the delivery of imaging decision support systems [13]. This paper proposes an ML implementation that covers a clinical decision support (CDS) architecture for non-imaging ML cases in clinical practice. For example, deployment of the ATTRwt-CM ML model within EHR systems of healthcare facilities may lead to improved testing and earlier confirmation of ATTRwt-CM in at-risk patients by raising suspicion of diagnosis in patients with heart failure. This could result in earlier treatment and, consequently, improved outcomes [12]. However, successful implementation of ML models relies on the minimization of barriers that may affect their utilization. For example, an essential outcome is that patient privacy is maintained; therefore, to ensure no patient data are transferred to the model's developer (in this example, Pfizer), the model can be run inside the firewalls of a third-party host.

The primary objective of this paper is to present an overview of potential architectural patterns that can bring AI or ML models closer to the consumer (i.e., patients), thereby increasing model utilization. We discuss these architectural patterns using the specific example of our ATTRwt-CM ML model. The secondary objective is to describe a scalable and modular architectural framework that adopts a container-based implementation of the ATTRwt-CM ML model on a third-party platform for use by healthcare providers (HCPs). Both objectives leverage Fast Healthcare Interoperability Resources (FHIR) CDS Cards/Hooks communication protocols to enable interoperable and scalable solutions [14], [15]. Furthermore, we found that adapting the model to FHIR standards required the use of adaptors that are defined in Section III.

By discussing solutions for potential challenges to model commercialization, we aim to enhance understanding of the potential for increased AI or ML model utilization in healthcare. Our proposed solution provides valuable insights and repeatable patterns for use in clinical practice.

A preliminary version of this work has been reported [16].

TABLE 1. Conceptual architectural designs for delivery of ML models.

Scenario	Manageable and scalable	Easily deployed	Data stay with customer	Protection of IP
1. Ship container to host as a standalone container	Yes [a]	Yes [a]	Yes	Yes [b]
2. Ship container to host as a media file (e.g., .vmdk, .ovf, .vdf, .ami)	No	No	Yes	Yes [b]
3. Host container at Pfizer, managed by Kubernetes environment	Yes	Yes	No	Yes
4. Host container at Pfizer, managed by microservice environment	Yes	Yes	No	Yes
5. Create an Epic EHR application	No	Yes	Maybe	Yes [b]
6. Adopt privacy-preserving AI [c]	Maybe	No	Yes	Yes

[a]Assuming continuous integration and delivery. [b]Using encryption. [c]The model runs at the client's browser using privacy-preserving AI principles from OpenMined. AI = artificial intelligence, EHR = electronic health record, IP = intellectual property, ML = machine learning.

II. CONCEPTUAL ARCHITECTURAL DESIGNS FOR DELIVERY OF AI OR ML MODELS

There were four required characteristics of the architectural design for delivery of the ATTRwt-CM model to HCPs. First, the design needed to be manageable and scalable so that it could be applied to future ML or AI applications. Secondly, the application needed to consider the information technology (IT) maturity of the third-party host to allow for varying levels of IT capabilities. Thirdly, patient data needed to remain on the host's platform and not be shared with Pfizer to adhere to local or regional privacy regulations such as General Data Protection Regulation guidelines [17] or the Health Insurance Portability and Accountability Act of 1996 [18]. Finally, Pfizer's intellectual property needed to be protected by limiting code exposure to third parties. The development team weighed six conceptual architectural designs to deliver the model, based on whether or not they met these four requirements (Table 1).

The first two architectural designs followed a similar philosophy of shipping the model to the host, either as a container image or a media file. A container is defined as the software package that holds all necessary elements and dependencies to run in any environment. In designs 1 and 2, the host is assumed to have moderate-to-high IT maturity, and the ATTRwt-CM ML application is packaged and shipped as a container image (Figure 1) or media file (Figure 2). Following the logic presented in the figures, the host downloads the container image to a docker or deploys the media file; in both cases, the application runs within the host's domain. These designs ensure that patient data remain with the host and can facilitate both single-patient and bulk implementation.

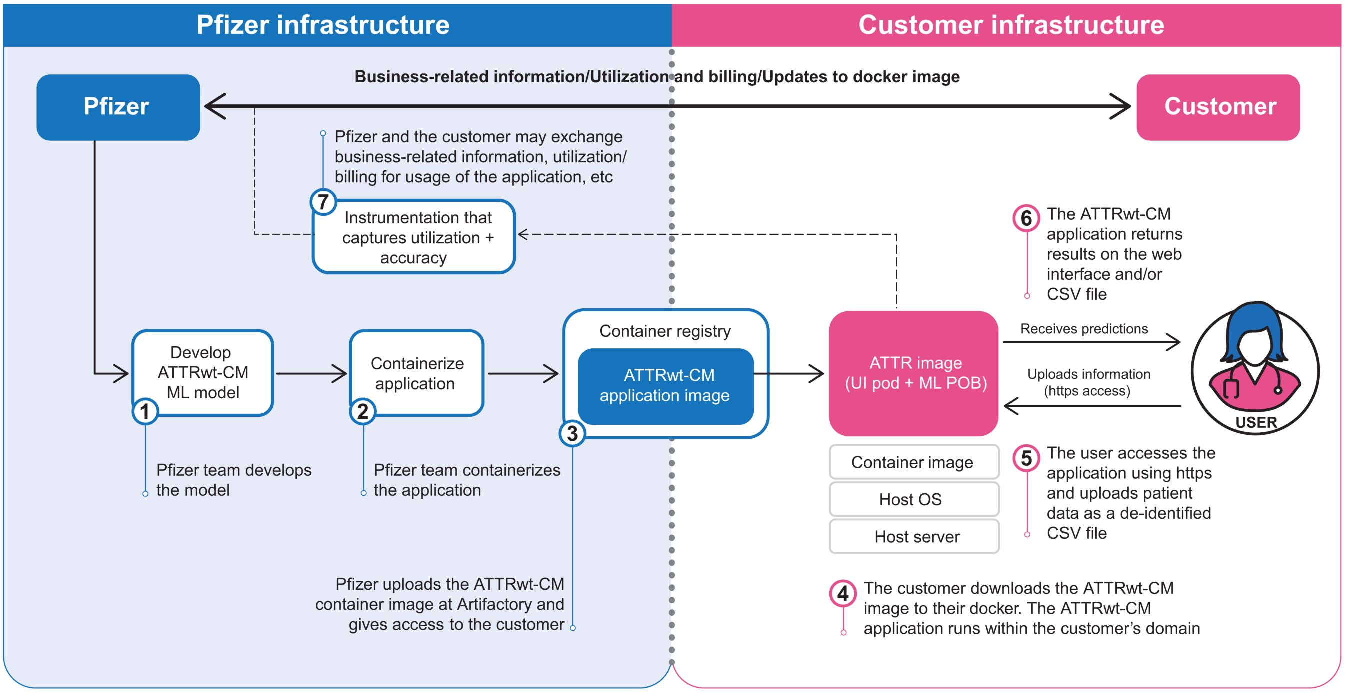

FIGURE 1. **Conceptual architectural design of shipping the container to an external host as a container image. ATTR = transthyretin amyloidosis, ATTRwt-CM = wild-type transthyretin amyloid cardiomyopathy, CSV = comma-separated values, ML = machine learning, OS = operating system, POB = point of business, UI = user interface.**

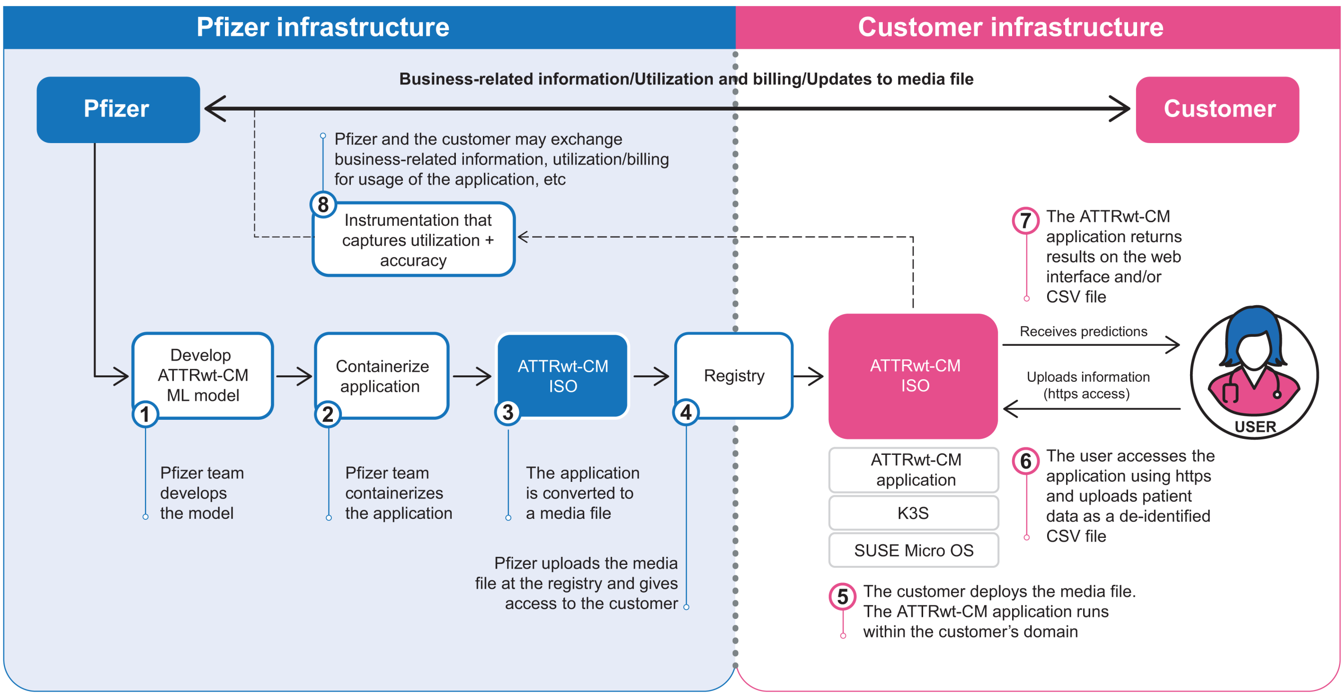

FIGURE 2. **Conceptual architectural design of shipping the container to an external host as a media file. ATTRwt-CM = wild-type transthyretin amyloid cardiomyopathy, CSV = comma-separated values, ISO = International Organization for Standardization, ML = machine learning, OS = operating system.**

Drawbacks of these designs include that management of the application is complex (e.g., it involves artifact image updates and communication with the host for every update) and the host is required to have the IT infrastructure for hosting the container. As illustrated in the figures, the overall process of these designs is similar, with the main difference being that the container-based design (Figure 1) assumes that a container can run in the docker environment of the host,

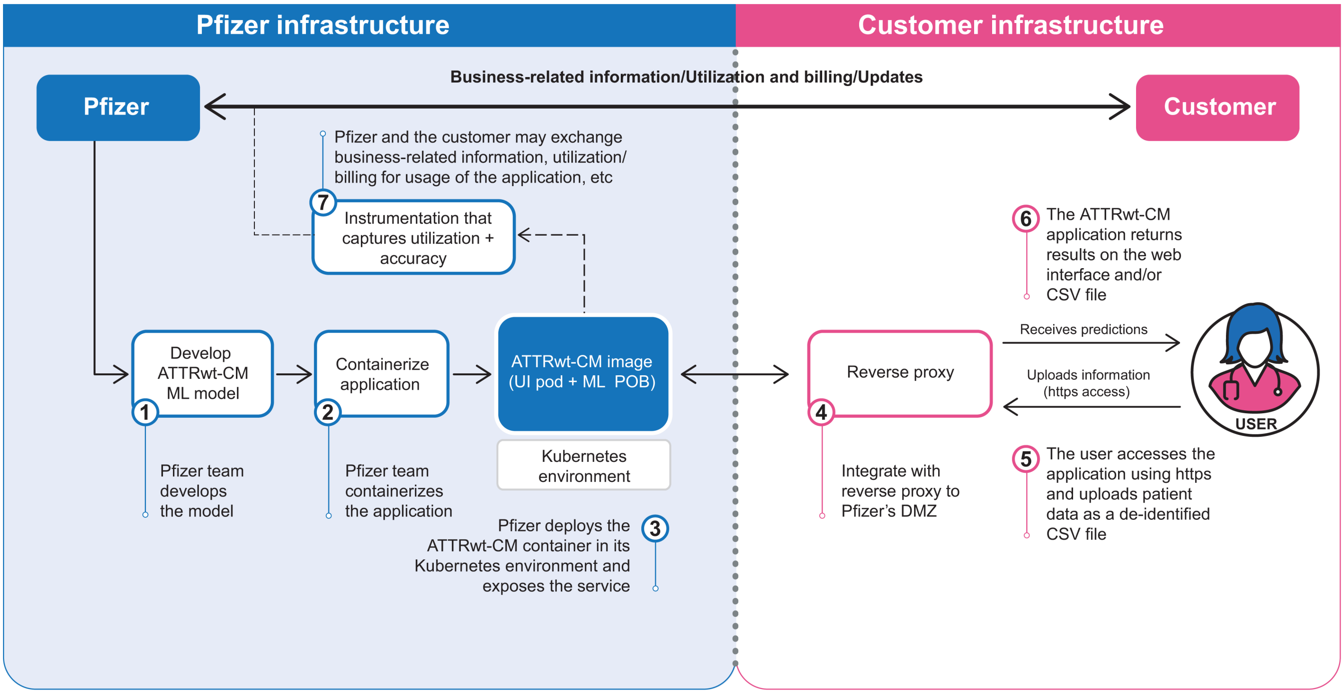

FIGURE 3. **Conceptual architectural design of deploying the container in Pfizer's Kubernetes environment. ATTRwt-CM = wild-type transthyretin amyloid cardiomyopathy, CSV = comma-separated values, DMZ = demilitarized zone, ML = machine learning, POB = point of business, UI = user interface.**

whereas the media file design (Figure 2) may add complexity as the host must run the executable file in different computing machines.

In designs 3 and 4, the customer is assumed to have low IT maturity and to be unable to host or manage the application. In these designs, Pfizer hosts the application either in their Kubernetes environment, exposing the service to a reverse proxy server (Figure 3) or in their microservice environment (Figure 4). These designs are easily manageable and scalable and ensure the protection of Pfizer's intellectual property; however, a disadvantage is that patient data are shared with Pfizer, thereby requiring systems to de-identify and re-identify data (e.g., through allocation of patient identification numbers).

In design 5, an ATTRwt-CM application is developed for release on the Epic App Orchard (Verona, WI, USA) marketplace (Table 1). In this case, Pfizer creates a representational state transfer RESTful application programming interface (API) for the model that aligns with FHIR, whereas the model itself runs on Pfizer's platform (Figure 5). Data are sent through FHIR to the model and the ATTRwt-CM API pushes the results to the Epic EHR user interface. Although this design provides a simple user experience, one downside is the requirement for information exchange between the customer and Pfizer. In addition, expanding the delivery of the application to HCPs who do not use the Epic App Orchard marketplace may add complexity. The advantage is that the design follows FHIR data exchange standards, making it expandable to EHR systems and promoting interoperability.

Design 6 sees the ATTRwt-CM ML model being trained and run in the HCP's environment using privacy-preserving and federated learning AI principles, such as those from OpenMined [19], [20], [21] (Table 1; Figure 6). The main principle of this design is to revert the process of ML model training (i.e., the model is sent for partial training to datasets that are not in the administrative domain of the model producer). As illustrated in Figure 6, all information stays with the customer, and models can be trained on a dataset outside Pfizer's administrative domain. This design maximizes the potential data outreach of the model; however, OpenMined and federated learning are new concepts and technologies and require further exploration before applying in this context.

Other designs were considered, but they failed to meet the required outcomes and were not pursued further. These included an open-access option, where Pfizer offers the model's source code at GitHub as open access and requests that users either reference Pfizer or improve or comment on the model. Another design that was not pursued was the utilization of dedicated hardware at customers' premises. For this method, the container is deployed and hosted on an edge hardware device, such as Raspberry Pi, and resides within the customers' physical environment and network; however, this process creates scalability and manageability problems.

The team opted to pilot the application using design 1 (Figure 1)—shipping it to a third-party host as a standalone container. This approach would be manageable, scalable, and easily deployed. It would also protect the privacy of patient data by keeping the data with the customer and would protect

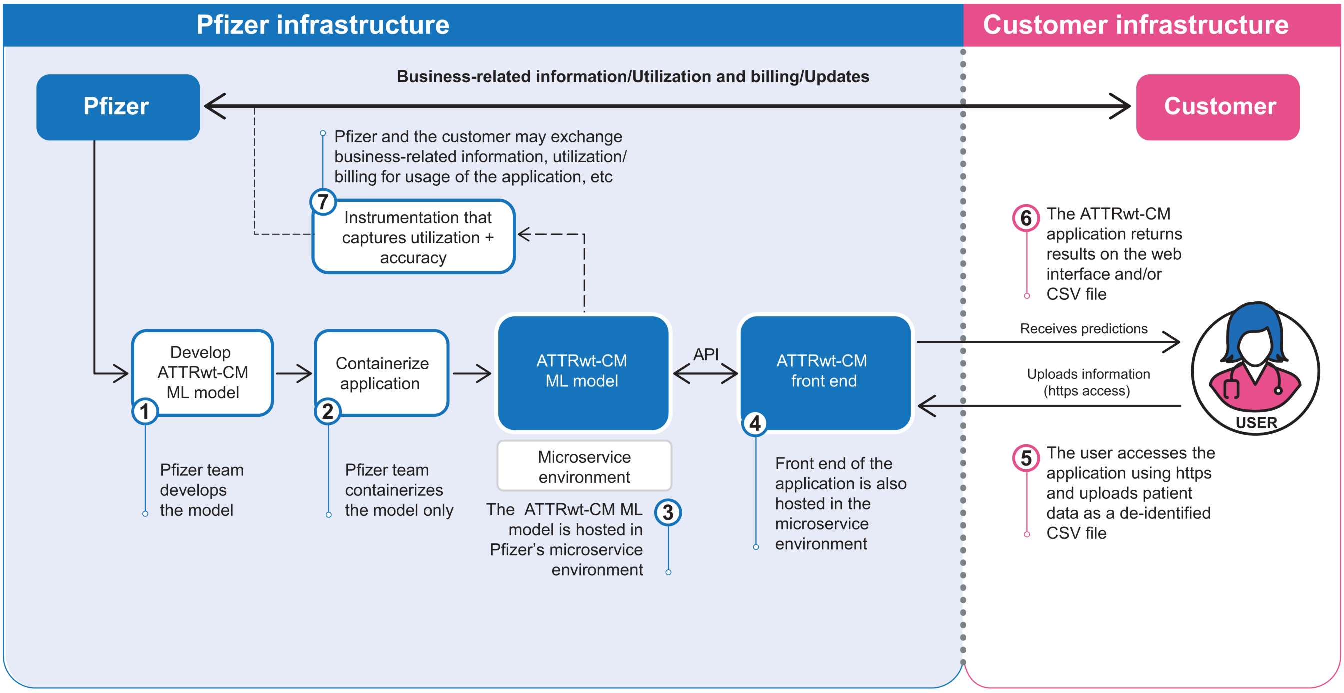

FIGURE 4. **Conceptual architectural design of running the application in Pfizer's microservice environment. API = application programming interface, ATTRwt-CM = wild-type transthyretin amyloid cardiomyopathy, CSV = comma-separated values, ML = machine learning.**

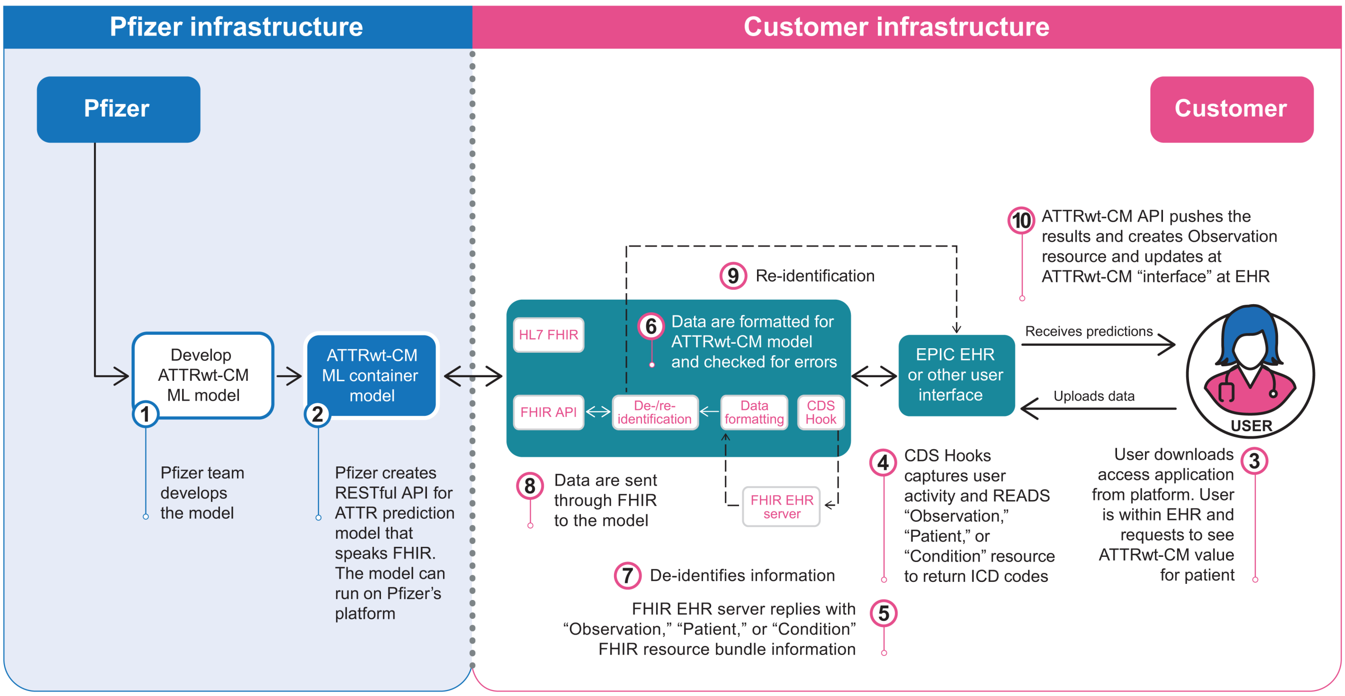

FIGURE 5. **Conceptual architectural design of releasing an application on the Epic App Orchard Marketplace. API = application programming interface, ATTR = transthyretin amyloidosis, ATTRwt-CM = wild-type transthyretin amyloid cardiomyopathy, CDS = clinical decision support, EHR = electronic health record, FHIR = Fast Healthcare Interoperability Resources, ICD = International Classification of Diseases, ML = machine learning.**

the intellectual property of the model using encryption. The container would be a single image or be based on a microservice architecture, depending on the number and complexity of the delivered services.

III. PILOTING THE ATTRwt-CM ML MODEL ON THE PHILIPS PLATFORM

This section presents the proposed implementation that requires a connection of the ML operations of the model

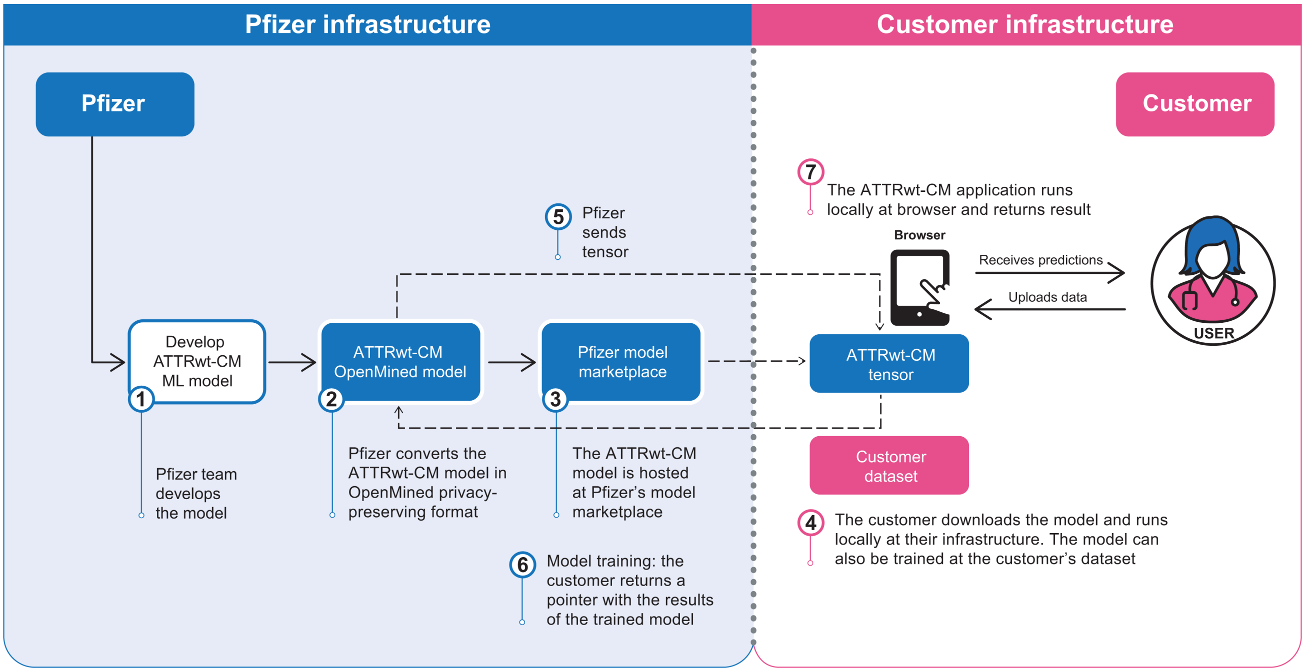

FIGURE 6. Conceptual architectural design of running the model locally using OpenMined privacy-preserving principles. ATTRwt-CM = wild-type transthyretin amyloid cardiomyopathy, ML = machine learning.

producer with an external platform (the host) that employs an FHIR-based communication protocol.

The first implementation of Pfizer's ATTRwt-CM ML model was with the Philips HealthSuite Diagnostics (HSD); the architectural design of the pilot is illustrated in Figure 7. The Philips platform was chosen for the pilot because their existing solutions are used by hospitals and include dashboards called Care Orchestrators, which aggregate patient data for the application of CDS [22] and provide a natural landing spot for consumption and testing of the model. Furthermore, due to data restrictions, any cloud infrastructure in the administrative domain of Pfizer was not an option, and the Philips HSD platform provided the required infrastructure to host the service.

The primary objectives of the pilot study were to *(i)* define and test the architectural design and technical parameters for sharing the container image; *(ii)* capture metrics, including engagement and performance data, after running the container image in a test environment; *(iii)* create a scalable framework to manage multiple applications on different third-party platforms; *(iv)* explore security and intellectual property problems when running the application in a third-party environment; *(v)* understand the effort and validity of the integration of container images on third-party platforms; and *(vi)* employ a standards-based approach utilizing CDS Hooks specifications, to enable scaling up. Pilot engagement data and post pilot results will be reported in a future publication.

For this specific test implementation, Pfizer containerized the application as a single image and delivered the container to a Philips repository. Philips downloaded the container image to host on their HSD (the container is considered the CDS service). This process ensures that Pfizer cannot access the patient data; however, they will eventually be able to retrieve data on engagement (e.g., usage and number of calls) and performance (e.g., execution time and accuracy), as generated by the instrumentation engine, as well as information from quality checks (e.g., data format errors).

The Philips HSD contains three main components relevant for implementing the model; namely, the Care Orchestrator web application, which houses the user interface for the end user and acts as the CDS Client; the Care Orchestrator backend, which, upon request, executes the model by means of a CDS Hooks mechanism; and a clinical data repository (CDR), a database that stores all relevant patient-specific data.

The process for execution of the application is provided in Figure 8. Our design utilizes algorithm wrappers to adapt the execution of the model outside Pfizer's administrative domain and in an FHIR-based system and provides a reproducible framework for implementation that can be further studied and explored. When the user opens a patient encounter in the Care Orchestrator Web Application, this automatically triggers a CDS Hook to perform data prefetching, and a REST call invokes the algorithm in the container image. Patient data prefetched from the CDR include a unique identifier; age in years; sex; International Classification of Diseases, Tenth Revision (ICD-10) diagnosis codes; the dates of the ICD-10 diagnosis codes; and the dates of encounter visits. The application checks that data are in the proper format and, if required, maps diagnostic codes to the correct ICD-10 version on which the model was trained. An HTTP 400 error reply is returned for data that fail this initial quality check, ICD-10 code mapping, or data validation. When patient data

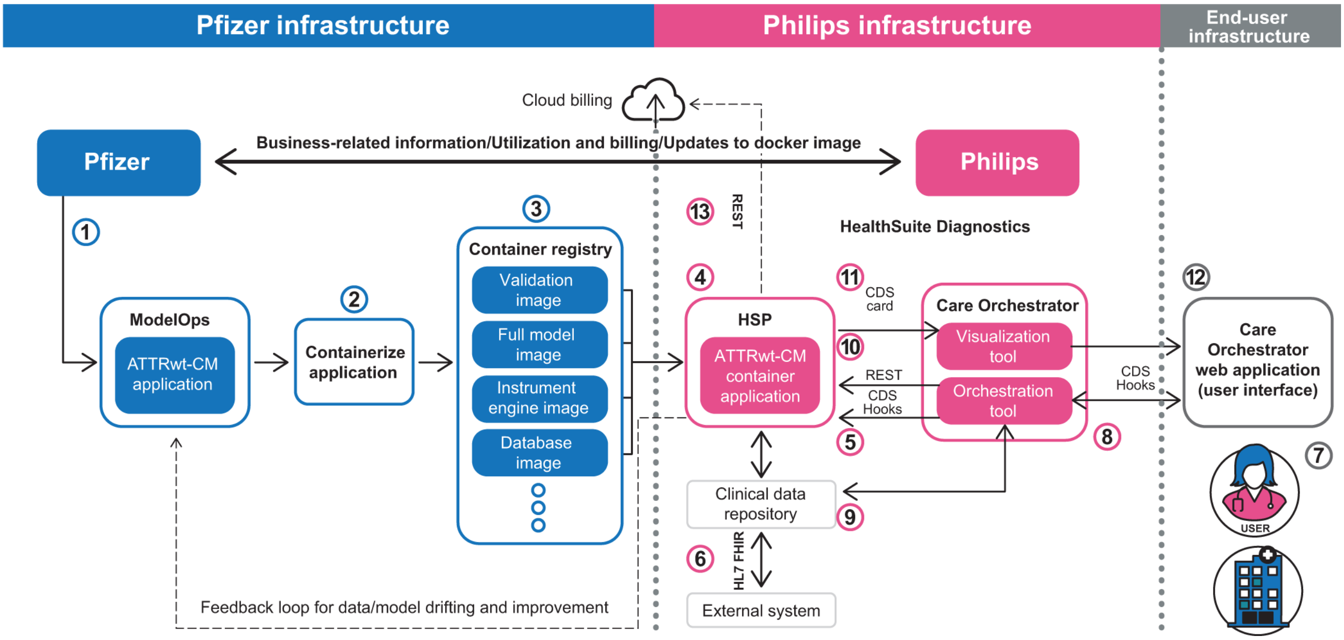

FIGURE 7. Architectural design of the Pfizer/Philips pilot study. (1) Pfizer develops the ATTRwt-CM application; (2) Pfizer containerizes the application; (3) the application images are stored within a container registry that helps with the management of multiple versions of the application with several different hosts (for the collaboration with Philips, a single image of the application was shipped to the Philips repository); (4) the container runs within Philips' HSP and is considered the CDS Service following the CDS Hooks, FHIR HL7 specification; (5) data prefetching information is shared between the Care Orchestrator and application in the form of a CDS Hook, which invokes the data prefetching template form to collect data from the Care Orchestrator (the CDS Client) and the FHIR server, which is represented as the CDR; (6) the CDR calls the FHIR EHR server and retrieves the required information to be available at the Care Orchestrator; (7) the user launches the web application when opening a patient encounter; (8) the CDS Hook triggers an event at the orchestration tool; (9) the orchestration tool calls the CDR to prefetch data for the ATTRwt-CM application; (10) the orchestration tool triggers the ATTRwt-CM algorithm using REST; (11) the ATTRwt-CM algorithm runs and returns the results as a CDS Card for visualization; (12) the user can access the visualization at the web application; (13) utilization results may be captured by the instrumentation engine and posted to the billing system; it should be noted that for the Philips implementation, the instrumentation engine was not considered and the HSP platform periodically shares usage analytics data with Pfizer. ATTRwt-CM = wild-type transthyretin amyloid cardiomyopathy, CDR = clinical data repository, CDS = clinical decision support, FHIR = Fast Healthcare Interoperability Resources, HSP = HealthSuite Platform, REST = representational state transfer.

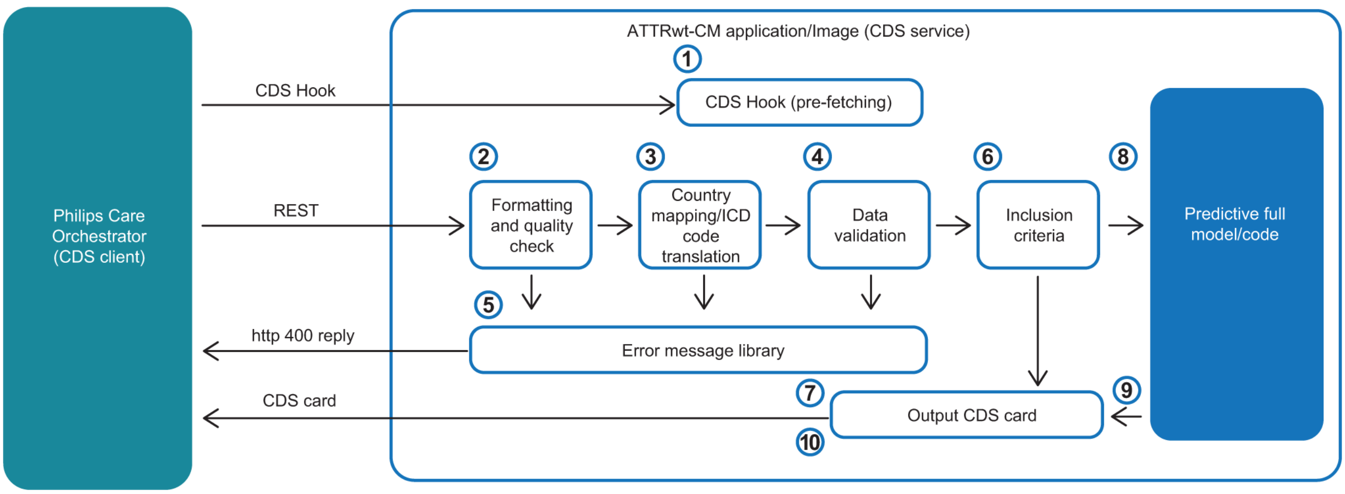

FIGURE 8. Execution of the ATTRwt-CM application. (1) Upon opening a patient encounter, a CDS Hook performs data prefetching; (2) a REST call invokes the algorithm in the container image and the data undergo formatting and a quality check; (3) diagnostic codes are mapped to the correct format, if needed; (4) data are validated; (5) an error message is returned if data are not valid; (6) eligibility is assessed; (7) if a patient is not eligible, a CDS Card is returned stating "not eligible"; (8) data are entered into the model code; (9) the model code produces a result; (10) the result is returned as a CDS Card for visualization at the web application. ATTRwt-CM = wild-type transthyretin amyloid cardiomyopathy, CDS = clinical decision support, ICD = International Classification of Diseases, REST = representational state transfer.

pass these checks, the application determines whether the patient fulfills the model's inclusion/exclusion criteria (previously published [12]) prior to executing the model code. The application returns the results as one of six CDS Cards: (*i*) no available data due to the patient not having ICD-10 codes for evaluation, in which case the model code is not executed; *(ii)* the patient does not meet eligibility criteria, another situation in which the model code is not executed; *(iii)* high suspicion of ATTRwt-CM (risk score $\geq$ 47.5%), with the actual risk score shown; *(iv)* low suspicion of ATTRwt-CM (risk score $<$ 47.5%), with the actual risk score shown); (*v*) no prediction because the data entered the model but did not return a valid response due to an unhandled error; or *(vi)* no prediction because not enough patient data were available to assign a risk score. The CDS Cards are rendered for visualization at Philips on the Care Orchestrator. The user accesses the visualizations at the web application, and the instrumentation engine posts utilization metrics to the billing system.

System end-to-end delays are to be minimized to provide an acceptable user experience, defined as $<$ 1 s to 2 s from entering a patient's name in the dashboard to retrieving a result. Any Philips platform-related delays involving data retrieval, internal data queries, and triggering functions on the platform are not to exceed 3 s, and these delays are to be controlled with precise database searches. The execution time of the ATTRwt-CM algorithm for one user, assuming all required inputs are delivered in the REST triggering call, is to be no more than 500 ms and is controlled by use of a code-efficient algorithm. Communication delays, which include the CDS Hooks and REST calls between the Care Orchestrator, CDR, and ATTRwt-CM application, are not to exceed 500 ms; and the presentation of results is to take no longer than 1 s.

IV. DISCUSSION

Pfizer's ATTRwt-CM ML model has previously performed well in identifying ATTRwt-CM vs. nonamyloid heart failure [12]. Utilization of the model in the clinical setting can facilitate earlier identification of patients with heart failure at risk for ATTRwt-CM and, subsequently, earlier treatment, which is associated with improved outcomes [23]. This pilot project with Philips presents an architectural solution for delivery of the model to HCPs by using a container-based design and communication patterns based on CDS Hooks, CDS Cards, and REST calls. A major concern with the use of any ML application in healthcare is the privacy of patient data. In this pilot, the application was hosted on a third-party platform, and patient data remained with the customer/host. Overall, this project represents a breakthrough achievement in the field of ML solutions in healthcare by delivering advanced technology to HCPs in a simple, user-friendly application without compromising patient confidentiality.

ML and AI in healthcare is a rapidly progressing field and represents a unique opportunity for improving the early identification of diseases and supporting clinical decision making. In addition to the ATTRwt-CM ML model described here, other models that predict risk and/or outcomes for conditions such as kidney disease, type 2 diabetes, myocardial infarction, chronic obstructive pulmonary disease, and coronary artery disease have been developed [24], [25], [26], [27], [28]. However, a research gap remains between the development of predictive models and the technical implementation of these models in the clinical setting [29]. While the technology to facilitate the implementation of these ML or AI models exists (e.g., FHIR and CDS Hooks; see example uses in [30], [31], and [32]), research into how these technologies can be integrated in a manageable and scalable way to deliver these models to the consumer is limited. Our paper bridges the gap between model development and implementation by presenting potential architectural solutions for model delivery and providing a concrete example of model implementation at a third-party host. Our solution not only demonstrates how this advanced technology can be delivered to different clinical settings for use by HCPs, but it also provides the specifications needed for continued diligence of the algorithm and easy deployment.

A limitation of this pilot is that it tested only the components that were needed to deploy the application and the technical feasibility of the design; model drift and performance were not tested. Another limitation was that synthetic data were used to test technical feasibility. Although these data were intended to simulate real-world data, they are not actual data. Future studies will assess the clinical feasibility of the design prior to deployment and will examine external validation. Further implementation research and post-implementation evaluations will also be conducted, as well as assessment of the model as Software as a Medical Device through the appropriate healthcare authorities.

V. CONCLUSION

This article presents a container-based design for the implementation of the ATTRwt-CM ML model on a third-party platform, which may enable earlier identification and treatment of this fatal disease and lead to improved outcomes. Learnings from this pilot may also inform the delivery of other related ML models to HCPs, thereby increasing utilization and potentially improving identification and management of other diseases.

ACKNOWLEDGMENT

Medical writing support was provided by Emily Balevich, Ph.D., of Engage Scientific Solutions, and was funded by Pfizer.

COMPETING INTERESTS

George Koutitas, Kimberly Nolen, Sepideh Attal, and Anastasios Ventouris are employees of and own stock/stock options in Pfizer. Yinnon Dolev and Hans Thijs van den Broek are employees of Philips.

REFERENCES

[1] F. L. Ruberg, M. Grogan, M. Hanna, J. W. Kelly, and M. S. Maurer, "Transthyretin amyloid cardiomyopathy: JACC state-of-the-art review," *J. Am. Coll. Cardiol.*, vol. 73, no. 22, pp. 2872–2891, 2019.

[2] M. M. Kittleson, F. L. Ruberg, A. V. Ambardekar, T. H. Brannagan, R. K. Cheng, J. O. Clarke, L. M. Dember, J. G. Frantz, R. E. Hershberger, M. S. Maurer, J. Nativi-Nicolau, V. Sanchorawala, and F. H. Sheikh, "2023 ACC expert consensus decision pathway on comprehensive multidisciplinary care for the patient with cardiac amyloidosis: A report of the American college of cardiology solution set oversight committee," *J. Amer. College Cardiol.*, vol. 81, no. 11, pp. 1076–1126, 2023.

[3] M. S. Maurer et al., "Genotype and phenotype of transthyretin cardiac amyloidosis: THAOS (transthyretin amyloid outcome survey)," *J. Amer. College Cardiol.*, vol. 68, no. 2, pp. 161–172, 2016.
[4] R. M. Witteles, S. Bokhari, T. Damy, P. M. Elliott, R. H. Falk, N. M. Fine, M. Gospodinova, L. Obici, C. Rapezzi, and P. Garcia-Pavia, "Screening for transthyretin amyloid cardiomyopathy in everyday practice," *JACC, Heart Failure*, vol. 7, no. 8, pp. 709–716, Aug. 2019.
[5] L. H. Connors, F. Sam, M. Skinner, F. Salinaro, F. Sun, F. L. Ruberg, J. L. Berk, and D. C. Seldin, "Heart failure resulting from age-related cardiac amyloid disease associated with wild-type transthyretin: A prospective, observational cohort study," *Circulation*, vol. 133, no. 3, pp. 282–290, Jan. 2016.
[6] J. H. Pinney, C. J. Whelan, A. Petrie, J. Dungu, S. M. Banypersad, P. Sattianayagam, A. Wechalekar, S. D. J. Gibbs, C. P. Venner, N. Wassef, C. A. McCarthy, J. A. Gilbertson, D. Rowczenio, P. N. Hawkins, J. D. Gillmore, and H. J. Lachmann, "Senile systemic amyloidosis: Clinical features at presentation and outcome," *J. Amer. Heart Assoc.*, vol. 2, no. 2, Mar. 2013, Art. no. e000098.
[7] M. Grogan, C. G. Scott, R. A. Kyle, S. R. Zeldenrust, M. A. Gertz, G. Lin, K. W. Klarich, W. L. Miller, J. J. Maleszewski, and A. Dispenzieri, "Natural history of wild-type transthyretin cardiac amyloidosis and risk stratification using a novel staging system," *J. Amer. College Cardiol.*, vol. 68, no. 10, pp. 1014–1020, Sep. 2016.
[8] M. S. Maurer, S. Bokhari, T. Damy, S. Dorbala, B. M. Drachman, M. Fontana, M. Grogan, A. V. Kristen, I. Lousada, J. Nativi-Nicolau, C. C. Quarta, C. Rapezzi, F. L. Ruberg, R. Witteles, and G. Merlini, "Expert consensus recommendations for the suspicion and diagnosis of transthyretin cardiac amyloidosis," *Circulation, Heart Failure*, vol. 12, no. 9, Sep. 2019, Art. no. e006075.
[9] Y. Ando, T. Coelho, J. L. Berk, M. W. Cruz, B.-G. Ericzon, S.-I. Ikeda, W. D. Lewis, L. Obici, V. Planté-Bordeneuve, C. Rapezzi, G. Said, and F. Salvi, "Guideline of transthyretin-related hereditary amyloidosis for clinicians," *Orphanet J. Rare Diseases*, vol. 8, no. 1, p. 31, 2013.
[10] M. S. Maurer, M. B. Sultan, and C. Rapezzi, "Tafamidis for transthyretin amyloid cardiomyopathy," *N. Engl. J. Med.*, vol. 380, no. 2, pp. 196–197, 2019.
[11] M. L. Müller, J. Butler, and B. Heidecker, "Emerging therapies in transthyretin amyloidosis—A new wave of hope after years of stagnancy?" *Eur. J. Heart Failure*, vol. 22, no. 1, pp. 39–53, Jan. 2020.
[12] A. Huda, A. Castaño, A. Niyogi, J. Schumacher, M. Stewart, M. Bruno, M. Hu, F. S. Ahmad, R. C. Deo, and S. J. Shah, "A machine learning model for identifying patients at risk for wild-type transthyretin amyloid cardiomyopathy," *Nature Commun.*, vol. 12, no. 1, p. 2725, May 2021.
[13] S. Mandal, A. B. Greenblatt, and J. An, "Imaging intelligence: AI is transforming medical imaging across the imaging spectrum," *IEEE Pulse*, vol. 9, no. 5, pp. 16–24, Sep. 2018.
[14] D. Bender and K. Sartipi, "HL7 FHIR: An agile and RESTful approach to healthcare information exchange," in *Proc. 26th IEEE Int. Symp. Comput.-Based Med. Syst.*, Jun. 2013, pp. 326–331.
[15] M. L. Braunstein, "Healthcare in the age of interoperability: The promise of fast healthcare interoperability resources," *IEEE Pulse*, vol. 9, no. 6, pp. 24–27, Nov. 2018.
[16] G. Koutitas, K. Nolen, and S. Attal, "Implementation of a machine learning model to assess transthyretin amyloid cardiomyopathy risk in an external platform," in *Proc. 18th Int. Symp. Amyloidosis*, Heidelberg, Germany, Sep. 2022, p. 327.
[17] EUR-Lex. (2016). *Regulation (EU) 2016/679 on the Protection of Natural Persons With Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation)*. [Online]. Available: http://data.europa.eu/eli/reg/2016/679/oj
[18] *Health Insurance Portability and Accountability Act of 1996, Pub. L. No. 104–191*, Aug. 1996. Accessed Aug. 7, 2023. [Online]. Available: https://www.congress.gov/bill/104th-congress/house-bill/3103
[19] (2022). *OpenMined*. [Online]. Available: https://www.openmined.org/
[20] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 50–60, May 2020.
[21] N. Rieke, J. Hancox, W. Li, F. Milletarã, H. R. Roth, S. Albarqouni, S. Bakas, M. N. Galtier, B. A. Landman, K. Maier-Hein, S. Ourselin, M. Sheller, R. M. Summers, A. Trask, D. Xu, M. Baust, and M. J. Cardoso, "The future of digital health with federated learning," *NPJ Digit. Med.*, vol. 3, no. 1, p. 119, Sep. 2020.
[22] (2019). *HL7® International CDC Hooks*. [Online]. Available: https://cds-hooks.hl7.org/
[23] P. Elliott, B. M. Drachman, S. S. Gottlieb, J. E. Hoffman, and S. L. Hummel, "Long-term survival with tafamidis in patients with transthyretin amyloid cardiomyopathy," *Circulat., Heart Failure*, vol. 15, no. 1, 2022, Art. no. e008193.
[24] K. C. Howlader, M. S. Satu, M. A. Awal, M. R. Islam, S. M. S. Islam, J. M. W. Quinn, and M. A. Moni, "Machine learning models for classification and identification of significant attributes to detect type 2 diabetes," *Health Inf. Sci. Syst.*, vol. 10, no. 1, p. 2, Feb. 2022.
[25] Y. Zou, L. Zhao, J. Zhang, Y. Wang, Y. Wu, H. Ren, T. Wang, R. Zhang, J. Wang, Y. Zhao, C. Qin, H. Xu, L. Li, Z. Chai, M. E. Cooper, N. Tong, and F. Liu, "Development and internal validation of machine learning algorithms for end-stage renal disease risk prediction model of people with type 2 diabetes mellitus and diabetic kidney disease," *Renal Failure*, vol. 44, no. 1, pp. 562–570, Dec. 2022.
[26] G. Y. H. Lip, A. Genaidy, G. Tran, and P. Marroquin, "Incident and recurrent myocardial infarction (MI) in relation to comorbidities: Prediction of outcomes using machine-learning algorithms," *Eur. J. Clin. Invest.*, vol. 52, no. 8, 2022, Art. no. e13777.
[27] X. Wang, H. Ren, J. Ren, W. Song, Y. Qiao, Z. Ren, Y. Zhao, L. Linghu, Y. Cui, Z. Zhao, L. Chen, and L. Qiu, "Machine learning-enabled risk prediction of chronic obstructive pulmonary disease with unbalanced data," *Comput. Methods Programs Biomed.*, vol. 230, Mar. 2023, Art. no. 107340.
[28] I. S. Forrest, B. O. Petrazzini, Á. Duffy, J. K. Park, C. Marquez-Luna, D. M. Jordan, G. Rocheleau, J. H. Cho, R. S. Rosenson, J. Narula, G. N. Nadkarni, and R. Do, "Machine learning-based marker for coronary artery disease: Derivation and validation in two longitudinal cohorts," *Lancet*, vol. 401, no. 10372, pp. 215–225, Jan. 2023.
[29] L. Petersson, I. Larsson, J. M. Nygren, P. Nilsen, M. Neher, J. E. Reed, D. Tyskbo, and P. Svedberg, "Challenges to implementing artificial intelligence in healthcare: A qualitative interview study with healthcare leaders in Sweden," *BMC Health Services Res.*, vol. 22, no. 1, p. 850, Dec. 2022.
[30] L. Rubin, N. P. Lopez, A. Gaiera, F. Campos, D. Luna, and F. B. G. de Quiros, "Development, implementation and preliminary results of an electronic reminder for HIV screening using a service oriented architecture," *Stud. Health Technol. Inform.*, vol. 264, pp. 763–767, Aug. 2019.
[31] K. Kawamoto, P. Kukhareva, J. H. Shakib, H. Kramer, S. Rodriguez, P. B. Warner, D. Shields, C. Weir, G. D. Fiol, T. Taft, and C. H. Stipelman, "Association of an electronic health record add-on app for neonatal bilirubin management with physician efficiency and care quality," *JAMA Netw. Open*, vol. 2, no. 11, Nov. 2019, Art. no. e1915343.
[32] S. Sinha, M. Jensen, S. Mullin, and P. L. Elkin, "Safe opioid prescripting: A SMART on FHIR approach to clinical decision support," *Online J. Public Health Informat.*, vol. 9, no. 2, p. e193, Sep. 2017.

GEORGE KOUTITAS (Member, IEEE) holds a B.Sc. degree in physics from Aristotle University of Thessaloniki in Greece, an M.Sc. degree (with distinction) in mobile and satellite communications from the University of Surrey, UK, and a Ph.D. in electrical engineering under EPSRC scholarship from the Centre for Communications Systems Research, UK. His post-doctoral research was in Smart Grids and Optimization algorithms under an EU funded program. He also holds various professional certifications such as Machine Learning from MIT.

He is currently a Director in Architecture, Digital Medicines & ML at Pfizer and holds more than 15 years of academic and entrepreneurial experience in the US, UK, and EU. He has founded two startup companies in Austin, Texas, in the areas of digital innovation in the energy sector and virtual/augmented reality in the medical response sector. He was also an Assistant Professor and Director of the XReality Research Lab at Texas State University. His main research interests include system modeling and optimization.

Dr. Koutitas is the inventor of one US nonprovisional patent and has published over 60 scientific publications in peer reviewed journals and international conferences and 2 books that received over 1400 citations from the international scientific community.

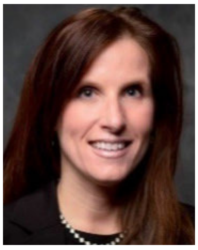

KIMBERLY NOLEN holds a B.S. degree in statistics from the University of Georgia and a Pharm.D. degree from Nova Southeastern University, and completed a Certification in Health Information Technology, Clinician Leader, from The George Washington University in 2012. Post-doctorate, she completed a Pharmacy Practice Residency at Grady Memorial Hospital in Atlanta, Georgia, and an Infectious Diseases Residency at the University of Mississippi Medical Center in Jackson, Mississippi.

She currently works in Medical Affairs at Pfizer.

Dr. Nolen has been recognized as a trusted Health Informatician thought leader by past appointments to a Federal Advisory Committee, the Health IT Standards Committee, in the US.

SEPIDEH ATTAL holds an M.D. degree with a specialization in emergency medicine from Paris University, France, and is a board-certified physician. She also graduated from Paris Descartes University with a degree in statistics and clinical trial methodology, infectious diseases. She most recently completed an executive M.B.A. degree in finance and management from Paris Dauphine University.

She is currently a Rare Disease Category Lead V.P. at Pfizer and has extensive experience in Medical Affairs at Pfizer and other large pharmaceutical companies.

ANASTASIOS VENTOURIS received a B.S. degree in mathematics from Aristotle University of Thessaloniki in 2012, followed by M.S. degrees in web science and data science from the same institution in 2013 and 2020, respectively.

He started his career as a Data Scientist in the game design industry from 2014 to 2021. He then joined Pfizer Inc in the Analytics Innovation Lab and later moved to a role as the Analytics Engineering Lead in January 2023.

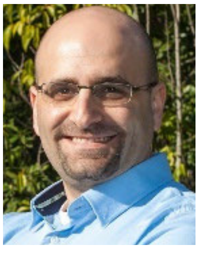

YINNON DOLEV received a B.Sc. degree in computer science and biology in 2006 and an M.Sc. degree in neurobiology in 2010, both from The Hebrew University of Jerusalem, Israel. His doctoral thesis titled 'Visual Perception in Jumping Spiders (Araneae, Salticidae)' resulted in a Ph.D. degree awarded in 2016 from the University of Canterbury in Christchurch, New Zealand.

He is currently a Solutions Architect of the EDI Advanced Development Team in Philips Medical Systems.

HANS THIJS VAN DEN BROEK holds a B.Sc. and M.Sc. degrees in technical medicine focused on 'Medical Imaging and Interventions' from the University of Twente. His doctoral thesis entitled 'Advancing medical devices for novel image-guided percutaneous cardiac interventions' resulted in a Ph.D. degree from the University Medical Center Utrecht.

He is currently a Clinical Scientist in the cardiovascular division at Philips Medical Systems.