

An Ontology for Digital Medicine Outcomes: Development of the Digital medicine Outcomes Value Set (DOVeS)

Benjamin Rosner, Matthew Horridge, Guillen Austria, Tiffany Lee, Andrew
Auerbach

Submitted to: JMIR Medical Informatics
on: October 16, 2024

Disclaimer: © The authors. All rights reserved. This is a privileged document currently under peer-review/community review. Authors have provided JMIR Publications with an exclusive license to publish this preprint on its website for review purposes only. While the final peer-reviewed paper may be licensed under a CC BY license on publication, at this stage authors and publisher expressly prohibit redistribution of this draft paper other than for review purposes.

Table of Contents

Original Manuscript..... 5
Supplementary Files..... 21
 Figures 22
 Figure 1..... 23
 Figure 2..... 24

An Ontology for Digital Medicine Outcomes: Development of the Digital medicine Outcomes Value Set (DOVeS)

Benjamin Rosner¹ MD, PhD; Matthew Horridge² PhD; Guillen Austria³; Tiffany Lee⁴; Andrew Auerbach⁴ MD, MPH

¹Division of Clinical Informatics and Digital Transformation Division of Hospital Medicine University of California San Francisco San Francisco US

²Biomedical Informatics Research Group Stanford University Stanford US

³School of Public Health University of California Berkeley Berkeley US

⁴Division of Hospital Medicine University of California San Francisco San Francisco US

Corresponding Author:

Benjamin Rosner MD, PhD

Division of Clinical Informatics and Digital Transformation

Division of Hospital Medicine

University of California San Francisco

10 Koret Way, Room K-301

UCSF DOM CLIIR Box 0737

San Francisco

US

Abstract

Background: Over the last 10-15 years, U.S. healthcare and the practice of medicine itself have been transformed by a proliferation of digital medicine and digital therapeutic (collectively, digital health tool; DHT) products. While a number of DHT classifications have been proposed to help organize these tools for discovery, retrieval, and comparison by healthcare organizations seeking to potentially implement them, none have specifically addressed that organizations considering their implementation tend to approach the DHT discovery process with one or more specific outcomes in mind. An outcomes-based DHT ontology could therefore be valuable not only for health systems seeking to evaluate tools that influence certain outcomes, but for regulators and vendors seeking to ascertain potential substantial equivalence to predicate devices.

Objective: Our objective was to develop, with the input of industry, healthcare, payer, regulatory, and patient input through the Accelerated Digital Clinical Ecosystem (ADviCE) consortium, an ontology specific to DHT outcomes, the Digital medicine Outcomes Value Set (DOVeS), and to make this ontology publicly available and free to use.

Methods: From a starting point of a 4 generation deep hierarchical taxonomy developed by ADviCE, we developed DOVeS using the Web Ontology Language (OWL) through the open-source ontology editor Protégé, and data from 185 vendors who had submitted structured product information to ADviCE. We adhered to OBO Foundry principles and incorporated the MONDO Disease Ontology and the Ontology of Adverse Events (OAE). After development, DOVeS was field tested between December, 2022 and May, 2023 with 40 additional independent vendors previously unfamiliar with ADviCE or DOVeS; a process that continued until the top 4 generations of classes in the ontology received no further modification suggestions. As a proof of concept, we subsequently developed a prototype DHT Application Finder leveraging DOVeS to enable a user to query for DHT products based on specific outcomes of interest.

Results: In its current state, DOVeS contains 42,320 and 9,481 native axioms and distinct classes respectively. These numbers are enhanced when taking into account the axioms and classes contributed by MONDO and the OAE.

Conclusions: DOVeS is publicly available on BioPortal and Github, and has a Creative Commons license CC-BY-SA that is intended to encourage stakeholders to modify, adapt, build upon, and distribute it. While no ontology is complete, DOVeS will benefit from a strong and engaged user base to help it grow and evolve in a way that best serves DHT stakeholders and the patients they serve.

(JMIR Preprints 16/10/2024:67589)

DOI: <https://doi.org/10.2196/preprints.67589>

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?

Please make my preprint PDF available to anyone at any time (recommended).

✓ **Please make my preprint PDF available only to logged-in users; I understand that my title and abstract will remain visible to all users.**

Only make the preprint title and abstract visible.

No, I do not wish to publish my submitted manuscript as a preprint.

2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?

✓ **Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).**

Yes, but please make my accepted manuscript PDF available only to logged-in users; I understand that the title and abstract will remain visible to all users.

Yes, but only make the title and abstract visible (see Important note, above). I understand that if I later pay to participate in <http://www.jmir.org>, my work will be published in a JMIR journal.

Original Manuscript

An Ontology for Digital Medicine Outcomes: Development of the Digital medicine Outcomes Value Set (DOVeS)

Benjamin I. Rosner, MD, PhD^{1,2}; Matthew Horridge, PhD³; Guillen Austria⁴; Tiffany Lee²;
Andrew Auerbach, MD, MPH²

Author Affiliations:

1. Division of Clinical Informatics and Digital Transformation, University of California San Francisco, San Francisco, CA
2. Division of Hospital Medicine, University of California San Francisco, San Francisco, CA
3. School of Medicine, Stanford University, Stanford, CA
4. School of Public Health, University of California, Berkeley, Berkeley, CA

Abstract

Background: Over the last 10-15 years, U.S. healthcare and the practice of medicine itself have been transformed by a proliferation of digital medicine and digital therapeutic (collectively, digital health tool; DHT) products. While a number of DHT classifications have been proposed to help organize these tools for discovery, retrieval, and comparison by healthcare organizations seeking to potentially implement them, none have specifically addressed that organizations considering their implementation, approach the DHT discovery process with one or more specific outcomes in mind. An outcomes-based DHT ontology could therefore be valuable not only for health systems seeking to evaluate tools that influence certain outcomes, but for regulators and vendors seeking to ascertain potential substantial equivalence to predicate devices.

Objective: Our objective was to develop, with the input of industry, healthcare, payer, regulatory, and patient input through the Accelerated Digital Clinical Ecosystem (ADviCE) consortium, an ontology specific to DHT outcomes, the Digital medicine Outcomes Value Set (DOVeS), and to make this ontology publicly available and free to use.

Methods: From a starting point of a 4 generation deep hierarchical taxonomy developed by ADviCE, we developed DOVeS using the Web Ontology Language (OWL) through the open-source ontology editor Protégé, and data from 185 vendors who had submitted structured product information to ADviCE. We adhered to OBO Foundry principles and incorporated the MONDO Disease Ontology and the Ontology of Adverse Events (OAE). After development, DOVeS was field tested between December, 2022 and May, 2023 with 40 additional independent vendors previously unfamiliar with ADviCE or DOVeS; a process that continued until the top 4 generations of classes in the ontology received no further modification suggestions. As a proof of concept, we subsequently developed a prototype DHT Application Finder leveraging DOVeS to enable a user to query for DHT products based on specific outcomes of interest.

Results: In its current state, DOVeS contains 42,320 and 9,481 native axioms and distinct classes respectively. These numbers are enhanced when taking into account the axioms and classes contributed by MONDO and the OAE.

Conclusions: DOVeS is publicly available on BioPortal and Github, and has a Creative Commons license CC-BY-SA that is intended to encourage stakeholders to modify, adapt, build upon, and distribute it. While no ontology is complete, DOVeS will benefit from a strong and engaged user base to help it grow and evolve in a way that best serves DHT stakeholders and the patients they serve.

Keywords: Digital health; digital medicine; digital therapeutics; ontology; medical informatics

Introduction

Over the last 10-15 years, the U.S. healthcare industry and the practice of medicine itself have been transformed by a proliferation of digital health software applications. It has been estimated, for example, that more than 350,000 health and wellness apps are available in mainstream app stores,[1] and that more than 300 million people have used them in one form or another.[2] Along with this breadth of applications comes a variety of terminology describing them. The colloquially used umbrella term “digital health,” for example, encapsulates many concepts including wellness apps, consumer grade mobile health apps, personal health tracking devices, remote patient monitoring applications, telemedicine platforms, and software as a medical device (SaMD)[3].

Driven in part by regulatory agencies and industry advocacy groups in the U.S. and abroad, efforts have been undertaken to more clearly delineate the boundaries between these types of applications. For example, “digital health” has been defined as technologies, platforms, and systems that engage consumers for lifestyle, wellness, and health-related purposes; capture, store, or transmit health data; and/or support life science and clinical operations.[4] As such, digital health products typically do not require clinical evidence, and do not meet the regulatory definition of a medical device. Digital *medicine* products, on the other hand, include evidence-based software and/or hardware products that measure and/or intervene in the service of human health. These products require clinical evidence and may (or may not) be classified as medical devices. At the most regulated end of this spectrum, digital *therapeutics* are health software intended to treat or alleviate a disease, disorder, condition, or injury by generating and delivering a medical intervention that has a demonstrable positive therapeutic impact on patient health, and produce real-world outcomes.[5] Digital therapeutics typically do fall under regulatory oversight, with prescription digital therapeutics as those that are prescribed by a licensed healthcare professional.

While these definitions organize digital health products into regulatory and potentially clinical versus non-clinical categories, they do not necessarily enable the side-by-side outcomes-based comparison of similar products for the purpose of understanding relative efficacy or safety, nor do they facilitate establishing — for regulators and for vendors filing 510(k) applications — whether certain products may be “substantially equivalent” to predicate devices.[6,7] Furthermore, most digital health tool end users (e.g. clinicians and patients) or decision makers (e.g. health system administrators making subscription or purchasing decisions, digital pharmacy and therapeutic committees making formulary decisions, or payers making coverage decisions), have one or more *outcomes* in mind that they are hoping the digital application can be used to influence. In these respects, a system to organize digital medicine and digital therapeutics products (which we will collectively refer to as digital health tools – DHTs) according to the outcomes they intend to influence, and a common set of outcome metrics, would be valuable.

The purpose of this work is to describe the development of an *outcomes-based* ontology, the Digital medicine Outcomes Value Set (DOVeS). This ontology came into being as part of the Accelerated Digital Clinical Ecosystem (ADviCE) consortium, a group of health systems, vendors, payers, and policymakers whose mission is to enable and scale the safe and effective adoption of DHTs in clinical practice.[8]

Methods

Phase 0: Defining scope and purpose

DOVeS emanated from needs of the University of California San Francisco-Stanford-FDA Center of Excellence in Regulatory Science and Innovation (CERSI)-funded ADviCE consortium to identify real-world performance metrics with which the safety and efficacy of DHTs could be characterized, whether for healthcare system leaders seeking to comparatively assess and implement DHTs, or for regulators seeking to classify them into similar outcomes-based groupings.[9] The purpose of DOVeS was to establish common sets of DHT real-world performance measure (RWPM) outcomes to facilitate more straightforward comparison between DHT products. The scope of DOVeS was to capture a body of outcomes with the potential to be influenced by DHTs, preferably but not exclusively by products capable of integrating into the electronic health record (EHR). DOVeS was intended to satisfy questions posed to vendors in a Digital Health Common Application (DHCA, see below), and to address competency questions for use cases such as: What DHT applications influence outcomes pertaining to disease X, What DHT applications satisfy healthcare administrative outcome Y, What DHT applications potentially mitigate adverse event Z?

Phase 1: Pre-DOVeS development – Background context and knowledge capture

DOVeS emanated from ADviCE, about which background context will be helpful. Briefly, in order to inform a RWPM framework, leaders from 5 health systems participating in ADviCE formulated a strategy to learn from the DHT marketplace and collaboratively developed a Digital Health Common Application. The purpose of the DHCA was to collect from vendors a standardized set of intake and discovery information about their products and the populations and disease areas these products serve. This common application could then enable health system decision makers with a variety of common metrics in one place as their health systems — or any other — embark upon initial DHT vendor discovery for potential implementation decisions. Fields captured in the DHCA included: Whether or not the DHT integrates with the EHR (and if so, which ones); mechanisms of integration (e.g. application programming interfaces (APIs), flat files, data warehouses); current or planned regulatory status of the DHT product; stage of company development (e.g. conceptual, proof of concept, beta, commercial deployment); security, privacy, and data sharing policies; disease indications; intended patient population(s); device integration; safety and efficacy evidence; accessibility, equity, feature types; and the outcomes that the product seeks to improve.

Between April, 2019 and June, 2020, 185 vendors voluntarily completed and submitted DHCA data to ADviCE. Based on these data, on an expansion of the FDA's Developing Software Precertification Program: A Working Model,[10] and consensus-building work carried out by the ADviCE consortium, a RWPM taxonomical hierarchy was initially developed, with 3 top-level classes: "Healthcare outcomes", "Non-clinical product performance outcomes", and "User experience outcomes." Beneath these top-level classes were 3 levels of descendant outcome classes. The taxonomy resided in a spreadsheet, but without the functionalities of an ontology.

Phase 2: Development of DOVeS ontology

Content addition

Beginning with the 3 top-level classes and the collection of descendant outcomes described above, we used the Web Ontology Language (OWL) through the open-source ontology editor Protégé¹¹ to create an ontology, and to populate subclasses, definitions, and axioms related to

DHTs. Based on data from the 185 vendors in the ADviCE database, input from members of the ADviCE consortium, one investigator's direct knowledge of the digital health space (BR), and two investigators' involvement in the University of California San Francisco's Digital Diagnostics and Therapeutics Committee — a digital Pharmacy and Therapeutics committee [11] (BR and AA) — we added more structure including subclasses, definitions, relationships, and refining values.

Technical considerations

The DOVeS ontology was developed specifically employing the OWL2EL fragment. This fragment is commonly used in many biomedical ontologies due to its balance between expressivity and computational efficiency.[12] Throughout the development process, we aimed to follow the OBO Foundry principles in order to ensure that the processes for creating and maintaining DOVeS would be robust, scalable, and aligned with widely accepted ontology development best practices.[13] Since DOVeS is specific to digital medicine and therapeutic applications, and not a general-purpose biomedical domain ontology, there are no plans to submit it to the OBO Foundry Library.

Of the OBO Foundry's 15 principles, the following 6 directly influenced our development environment and setup (Table 1), while other principles such as P8: Documentation, P10: Commitment to Collaboration, and several others were used for management and maintenance, [13] and are accommodated through our use of Github and BioPortal.[13]

Table 1 OBO Foundry Principles directly influencing DOVeS development environment and setup.

OBO Foundry Principle	Description
P1: Open	DOVeS ontology sources and releases are publicly available on Github under the open and permissive CC-BY-SA Creative Commons license.
P2: Common Format	We represented DOVeS using OWL. OWL supports serialization into several accepted concrete syntaxes, including the OBO Syntax, allowing interoperability with various ontology tools. This choice aligns with the OBO principle of adhering to a common format across ontologies.
P3: URI/Identifier Space	Uniform naming conventions were applied to the DOVeS ontology International Resource Identifier (IRI) and all associated terms. Each term is assigned a universally unique alphanumeric identifier, prefixed by 'https://w3id.org/doves/'. This approach follows the OBO Foundry recommendation to use a meaningless alphanumeric identifier for ontology terms, with human-readable labels encoded separately, as noted above
P4: Versioning	GitHub served as the primary platform for ontology development, providing out-of-the-box versioning and release functionality. Our development approach closely followed the established ontology engineering practices popularized by the Gene Ontology (GO) Consortium.[14,15] Released versions of DOVeS are made publicly available on BioPortal to support accessibility and dissemination.[16]
P6: Textual Definitions	Terms within DOVeS are accompanied by textual definitions,

	which follow the Aristotelian genus-differentia descriptive format. This format defines a term by specifying its broader category (genus) and its distinguishing characteristics (differentia). These definitions are provided using the IAO:0000115 (definition) annotation property.
P12: Naming Conventions	We applied consistent naming conventions across all DOVeS terms. Specifically, terms are named in lowercase unless they are proper nouns, and each term has exactly one primary label specified by an <code>`rdfs:label`</code> annotation. Additionally, synonyms and alternative labels are provided using <code>`skos:altLabel`</code> annotations.

Secondary ontology integrations

Because many DHTs are meant to influence disease or disease trajectories, we incorporated the Mondo Disease Ontology (MONDO), an OWL ontology unifying multiple disease ontologies and terminologies. MONDO is a publicly available ontology with an active maintenance userbase that harmonizes multiple disease resources.[17] It consists of 54,000 classes of which the human disease branch has 22,500 classes. We extracted the human disease branch (MONDO:0700096) using the *minimum information to reference an external ontology term* (MIREOT) method[18] provided by the ROBOT command line tool[19]. We chose MONDO over the International Classification of Diseases-10 (ICD-10), for example, because MONDO is a richly axiomatized OWL ontology whereas ICD-10 is a structured taxonomy that is optimized for healthcare reporting and billing. Using MONDO, therefore, allowed us to leverage the built-in semantics of OWL and an OWL reasoner to infer relationships between disease terms and for runtime querying. Furthermore, terms in MONDO are, to a greater degree than ICD-10, richly annotated with synonyms and definitions, including synonyms from ICD-10. This makes MONDO more useful for DOVeS at runtime where applications may need to search for DHTs by synonyms. Finally, MONDO worked well with our build tools. We used the Ontology Development Kit (ODK), which works best with OBO-style ontologies, like MONDO.

Because many DHTs are also used to mitigate adverse events, we also incorporated the Ontology of Adverse Events (OAE).[20]²³ This allows DHTs to be described in terms of outcomes that correspond to mitigation of certain types of adverse events such the number or rate of diabetic hypoglycemic events (OAE:0001057).

Phase 3: Testing and refinement

After initial ontology development based on input from the ADviCE consortium, internal expertise, and DHCA data of vendors in the ADviCE database, the DOVeS ontology was then externally field tested with an independent sample of DHT vendors who had not been through the DHCA intake process and were not otherwise familiar with it. Lists of potential vendors for outreach were provided by the FDA, Providence Health, the Digital Therapeutics Alliance, and the Organization for the Review and Care of Health Apps (ORCHA). Inclusion criteria consisted of vendors serving the U.S. market with digital medicine or digital therapeutic products that in some way served patients and/or clinicians for clinical care and were not purely administrative in nature.

To carry out field testing, we reached out via email to 206 DHT vendors between December 2022 and May 2023. Non-responses were followed up with an additional 2 outreach emails of varying invitation language. Forty vendors (19.4%) expressed interest and were invited to one-on-one semi-structured interviews via web meetings to allow them to visually inspect the ontology and to suggest additions and modifications. Thirty-seven vendors (92.5%) — 17 of which were producers of FDA-listed Artificial Intelligence/Machine Learning Medical Devices — accepted this invitation and participated in the web meetings. To plan for these meetings and focus on the areas within the ontology of greatest relevance to each of the vendors, vendors were sent REDCap surveys in advance, asking them to describe the areas in which their products operated and the types of outcomes their products aimed to influence, including: Clinical outcomes, educational outcomes (e.g. patient comprehension), engagement and adherence outcomes (e.g. user engagement), health economic outcomes, healthcare operations outcomes (e.g. operational efficiency), healthcare utilization outcomes (e.g. hospital readmission), patient reported outcomes (e.g. patient reported outcome response rates), process of care outcomes (e.g. adherence to HEDIS measures), cybersecurity (e.g. security and privacy certifications), interoperability, EHR integration, user experience (e.g. Net Promoter Score), product uptime/downtime outcomes, and any peer reviewed manuscripts describing outcomes. Of the 40 vendors sent REDCap surveys, 26 (65%) completed them. Those that did not complete them, but that nevertheless participated in one-on-one web meetings (n=11), had opportunities to describe their outcomes and areas of focus verbally. Information gleaned from the web meetings were captured in a database for potential DOVeS modifications.

We concluded one-on-one web meetings with vendors once we had achieved a steady state at which none of the top-level classes and their 3 descendent generations (i.e. the original hierarchy from ADviCE as well as additions to these levels we had made in the development process) resulted in modification suggestions from no less than 5 sequential vendors. Although some vendors late in the process had no ontology modification suggestions at all (suggesting some degree of ontology maturity), many had highly use-case specific outcomes suggestions particular to their products at class levels greater than 3 descendants deep from the top class. Although no ontology can be considered “complete,” we took this trend in modification suggestions as a sign of reasonable high-level class content and structural maturity.

Phase 4: Prototype development for real-world application demonstration

To demonstrate the potential utility of DOVeS for practical application, we built a prototype *DHT Application Finder* that allows a user to search for DHTs based on filter settings, outcomes-based search terms, conditions, specialties, procedures, types of technology, regulatory status, and EHR integration capabilities. The backend of the DHT Application Finder made use of OWL reasoning for querying, with the fields and filters on the front end of the DHT Application Finder being ontology-driven. Every item in the DHT Application Finder query form either corresponded to an ontology term from DOVeS or allowed a user to specify a list of ontology terms from DOVeS.

To identify specific DHTs, the DHT Application Finder referenced the ADviCE database of 185 vendors' products. The records in the database were annotated with ontology terms from DOVeS that described the conditions treated by the application, outcomes, etc. At runtime we constructed an OWL class expression that ontologically described each DHT application. When

a user submits a search request, the criteria in the search request are translated to a complex OWL class expression that describes a general class of DHT applications. Then OWL reasoning is used to classify and retrieve the DHT applications that are entailed to be subclasses of the general class.

Results

Phases 0-3 Scope, background, development, refinement and testing - Principal results

In its current state, DOVeS contains 42,320 and 9,481 native axioms and distinct classes respectively. (Table 2) These numbers are enhanced when taking into account the axioms and classes contributed by MONDO and the OAE.

Table 2. DOVeS summary metrics.

Metric	Count Native to DOVeS	Count including MONDO and OAE
Axioms	42,320	286,097
Number of distinct, named classes	9,481	36,412
Object property	12	32

Phase 4 Prototype development - Principal results

The DHT Application Finder front end and sample search results are shown in Figure 1. The DHT Application Finder's functionality is best illustrated with an example. Because of the availability of multiple DHTs in the ADviCE database pertaining to diabetes, we describe a use case in which a health system decision maker seeks to find candidate DHTs to help patients with diabetes in their health system achieve glycemic control (i.e. glycemic control is the outcome of interest). The example is most informative by illustrating an end-user seeking to identify DHTs first without, and then with outcomes-based search terms. In this example, the end user initially searches for DHTs based only on "diabetes" in the Conditions field. (Figure 2A) The user selects "diabetes mellitus" from the type-ahead list that appears, and presses "Search." The search results show a list of DHTs associated with diabetes mellitus including those for "type 2 diabetes mellitus" as well as "diabetic retinopathy." However, neither of these provides a list of DHTs exclusively related to the outcome of glycemic control.

Therefore, the user additionally (or could have done so as the first step) enters the term "Glyc" into the "Outcomes to be influenced" entry field (Figure 2B), and the type-ahead offers two choices: "Glycemic Control" and "Reduction in Glycemic Events". Choosing "Glycemic Control Outcome" and performing the search again, the user obtains a list now of only those DHTs specific to achieving glycemic control (e.g. this time not including DHTs associated specifically with diabetic retinopathy). In the case of the ADviCE database, this outcomes-based search reduced the field of products that the user might otherwise have needed to explore further by 33%.

Figure 1 consists of two panels, A and B, showing the ADviCE Application Finder interface.

Panel A: Search for outcome "glycemic control"

The interface is titled "ADviCE Application Finder (powered by the DOVeS E-Catalog)". It has three main sections: "Patient age group", "Patient role", and "End user". Each section has a list of checkboxes for selection. Below these are sections for "Outcomes to be influenced", "Conditions, specialties and procedures", "Technology", and "Regulatory". The "Regulatory" section includes a "DHT Integration" dropdown menu.

Panel B: A set of search results

The interface shows search results for the outcome "glycemic control". It lists several applications, each with a brief description and a "Details" link. The results are organized into a table with columns for "Application", "Description", and "Details".

Figure 1. A. Prototype DHT Application Finder with underlying DOVeS ontology. A. Search for outcome "glycemic control" B. A set of search results. (Company/product names redacted)

Figure 2 shows the entry fields for the DHT Application Finder.

Panel A: Conditions, specialties and procedures

The field is labeled "Conditions" and contains the text "diabetes". Below the field, a list of suggestions is shown, including "diabetes mellitus (T1M, diabetes mellitus, diabetes, diabetes mellitus (diabetes))", "type 2 diabetes mellitus (type 2 diabetes mellitus, non-insulin-dependent diabetes mellitus, diabetes mellitus, non-insulin dependent, association with, non-insulin dependent diabetes, type 2 diabetes, type II diabetes, non-insulin dependent diabetes mellitus, noninsulin dependent diabetes, NIDDM, T2DM - type 2 diabetes, severe, diabetic, non-insulin-dependent diabetes mellitus, diabetes mellitus, type 2, diabetes mellitus, type 2, diabetes mellitus, type 2)", and "gestational diabetes (maternal gestational diabetes mellitus, diabetes in pregnancy, gestational diabetes)".

Panel B: Outcomes to be influenced

The field is labeled "Desired outcomes" and contains the text "Glyc". Below the field, a list of suggestions is shown, including "Glycemic Control Outcome" and "Reduction in Glycemic Events Outcome".

Figure 2. Entry fields in the DHT Application Finder for A. Conditions, specialties, and procedures, and B. Outcomes to be influenced. The fields show type-ahead completion options for conditions known to the system, along with their synonyms that are drawn directly from the MONDO disease ontology.

Discussion

Principal results

In developing DOVeS, we describe an ontology framed by outcomes specific to DHTs, rather than technical or implementation features. Because our goal was to address products that would potentially be used in clinical practice that might fall under regulatory oversight, we focused on digital medicine and digital therapeutic products, and did not intend to capture outcomes associated with the much broader category of digital health.

The 3 top-level classes of DOVeS, derived from the background work done by the ADviCE

consortium, consist of “Healthcare outcomes,” “Non-clinical product performance outcomes,” and “User experience outcomes.” The top-level class “Healthcare outcomes” is central to DHTs, as most such products intend to influence something in this class. It is composed of descendent classes including, for example, clinical care outcomes (e.g. health benefits and mitigation of adverse events), healthcare utilization outcomes, healthcare operational outcomes, process of care outcomes, medico-legal outcomes, education outcomes, healthcare economic outcomes, socioeconomic outcomes, healthcare utilization outcomes, family and caregiver outcomes, and engagement and adherence outcomes.

Despite the large set of classes pertaining directly to healthcare outcomes, “Non-clinical product performance outcomes” and “User experience outcomes” were nevertheless felt to be important not only to health systems and regulators, but to DHT vendors themselves. These classes consisted respectively of: 1. Technical outcomes associated with the product (cybersecurity and privacy certification outcomes), interoperability, technical certification, and system availability outcomes to name a few, and 2. A variety of outcomes associated with user satisfaction and usability.

Efforts to organize DHTs to date have been described in part based on the specialties they serve, disease areas, user types, and certain functionalities,[21] but this can limit the way in which DHTs that are platforms serving across multiple specialties, or that span diseases of multiple organ systems (e.g. diabetes) are organized. The World Health Organization developed the Classifications of Digital Health Interventions as a non-ontological framework according to different user types (persons, healthcare providers, health management and support personnel), data services, and services and application types, but not on outcomes, and not in a computable format²³.

Although some limited outcomes-based taxonomies[22] and classifications exist,[23] to our knowledge, none are specific to digital medicine or digital therapeutics applications. Organizing DHTs into common categories based on the outcomes they influence then gives the marketplace an opportunity to compare different DHTs side by side for their impacts on common outcomes, and with the appropriate outcomes results, may make regulatory, purchasing, formulary, and coverage decisions more efficient and objective.

While a few taxonomies do exist that categorize DHTs on a limited set of functional categories and characteristics,[24,25] these are narrow in scope, and lack the interrelationships of an ontology. Establishing DHT organizational frameworks as ontologies provides a solution to heterogeneity problems like this, thereby defining the concepts and relationships that make interoperability possible.[26] It is common, for example, for ontologies to have “off the shelf” tooling; logic-based, precisely defined semantics; modeling languages to describe complex relationships; machine processing capability allowing for the computation of relationships; rigorous relationship structures, and in many cases the ability to be incorporated into other ontologies.

Limitations

By design, DOVeS was informed by DHTs excluding those in the “digital health” class (that class of products focused on wellness that does not fall under regulatory oversight). Furthermore, while the classes for the ontology were developed with academic, regulatory, community

health, payer, and patient stakeholders, and included field testing with an independent set of DHT vendors, DOVeS — like most ontologies — is not complete, but will rely and depend on an active base of engaged supporters to continue to build and maintain it. In addition, outcomes and the metrics to measure those outcomes are not necessarily universal in the DHT industry. In the process of outcome and metrics discovery, we attempted to select those through both internal and external consensus. However, these outcomes and metrics may need modification in the future as the ontology grows and is further informed by the broader community of stakeholders and the evolving industry.

Conclusions

An ontology for organizing DHTs into outcomes-based groupings is an important step as the digital medicine and digital therapeutics industry matures. While the DOVeS ontology provides an open-source framework as a starting point to do this, real-world value from the ontology will be achieved through the development of robust tools similar to the prototype DHT Application Finder that will leverage the ontology for specific use cases. These may include, for example, use by regulators to streamline the process for determining substantial equivalence to predicate devices, and use by health system decision makers to identify products that influence the same or related outcomes so that the universe of potential DHT products being considered for implementation can be rapidly narrowed to only those that are most relevant. While it is not known whether the 33% reduction in products that were more specifically identified with an outcomes-based search relative to a disease-based search described in the example from the ADviCE database above will necessarily translate if the entire marketplace of DHT products were to be indexed, any reasonable reduction would nevertheless offer savings in time and effort by decision makers attempting to identify DHTs for deeper discovery for potential clinical use.

DOVeS is publicly available on BioPortal[27] and Github[28], and has a Creative Commons license CC-BY-SA that is intended to encourage stakeholders to modify, adapt, build upon, and distribute it. Only with an engaged userbase will the ontology continue to grow and evolve, and — when applied to the marketplace — help organize DHTs in a way that is beneficial to stakeholders and ultimately to the patients they serve.

Acknowledgements

We thank Leeda Rashid and Vinay Pai of the FDA for thoughtful project feedback. We acknowledge the Digital Therapeutics Alliance, ORCHA, the FDA, and Providence Health for assistance in identifying potential vendors to review the ontology.

Funding: This project was supported by the Food and Drug Administration (FDA) of the U.S. Department of Health and Human Services (HHS) as part of a Center of Excellence in Regulatory Science and Innovation grant to University of California, San Francisco, U01FD005978. The contents are those of the author(s) and do not necessarily represent the official views of, nor an endorsement, by FDA/HHS, or the U.S. Government.

Conflicts of Interest

BR reports equity in Kuretic, Inc.
MH none declared

GA none declared
TL none declared
AA reports equity in Kuretic, Inc.

Abbreviations

ADviCE – Accelerated digital clinical ecosystem
API – Application programming interface
CERSI - Center of Excellence in Regulatory Science and Innovation
DHCA – Digital health common application
DHT – Digital health tool
EHR – Electronic health record
FDA – Food and Drug Administration
ICD – International classification of diseases
IRI - International Resource Identifier
DOVeS – Digital medicine Outcomes Value Set
MIREOT - Minimum information to reference an external ontology term
OAE – Ontology of adverse events
ODK - Ontology Development Kit
OWL – Web ontology language
RWPM – Real world performance metrics

References

1. Olsen E. Digital health apps balloon to more than 350,000 available on the market, according to IQVIA report | MobiHealthNews [Internet]. mobihealthnews. 2021 [cited 2024 Apr 8]. Available from: <https://www.mobihealthnews.com/news/digital-health-apps-balloon-more-350000-available-market-according-iqvia-report>
2. Wylie L. Health App Revenue and Usage Statistics (2024) - Business of Apps [Internet]. Business of Apps. [cited 2024 Apr 8]. Available from: [https://www.businessofapps.com/data/health-app-market/#:~:text=2023%20\(%24mm\),Health%20App%20Usage,that%20used%20them%20in%202021.](https://www.businessofapps.com/data/health-app-market/#:~:text=2023%20(%24mm),Health%20App%20Usage,that%20used%20them%20in%202021.)
3. IMDRF Software as a Medical Device (SaMD) Working Group. "Software as a Medical Device": Possible Framework for Risk Categorization and Corresponding Considerations [Internet]. International Medical Device Regulators Forum; 2014. Report No.: IMDRF/SaMD WG/N12FINAL:2014. Available from: <http://www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-140918-samd-framework-risk-categorization-141013.docx>
4. Goldsack J, Coder M, Fitzgerald C, Navar-Mattingly N, Coravos A, Atreja A. Digital Health, Digital Medicine, Digital Therapeutics (DTx): What's the difference? [Internet]. 2019 [cited 2024 Apr 7]. Available from: <https://www.healthxl.com/blog/digital-health-digital-medicine-digital-therapeutics-dtx-whats-the-difference>
5. ISO/TR 11147:2023 - Health informatics — Personalized digital health — Digital therapeutics health software systems [Internet]. [cited 2024 Apr 9]. Available from: <https://www.iso.org/standard/83767.html>
6. US Food and Drug Administration. The 510(k) Program: Evaluating Substantial Equivalence in Premarket Notifications [510(k)] Guidance for Industry and Food and Drug Administration Staff [Internet]. U.S. Department of Health and Human Services; 2014 Jul. Available from: <https://www.fda.gov/media/82395/download?attachment>
7. Wreh E. Premarket notification [510(k)]. Medical Device Regulation Elsevier; 2023. p. 57–89. [doi: 10.1016/B978-0-323-95354-2.00006-2]
8. ADVICE: Curated Health Apps for Health Systems [Internet]. [cited 2024 Apr 9]. Available from: <https://www.advicehealth.org>
9. U.S. Food and Drug Administration. Developing Frameworks and Tools for Integration of Digital Health Tools into Clinical Practice [Internet]. US Food and Drug Administration. 2024 [cited 2024 Oct 1]. Available from: <https://www.fda.gov/science-research/advancing-regulatory-science/developing-frameworks-and-tools-integration-digital-health-tools-clinical-practice>
10. U.S. Food and Drug Administration. Developing Software Precertification Program: A Working Model (v0.2 - June 2018) [Internet]. U.S. Food and Drug Administration; 2018

Jun. Available from: <https://www.fda.gov/media/113802/download?attachment>

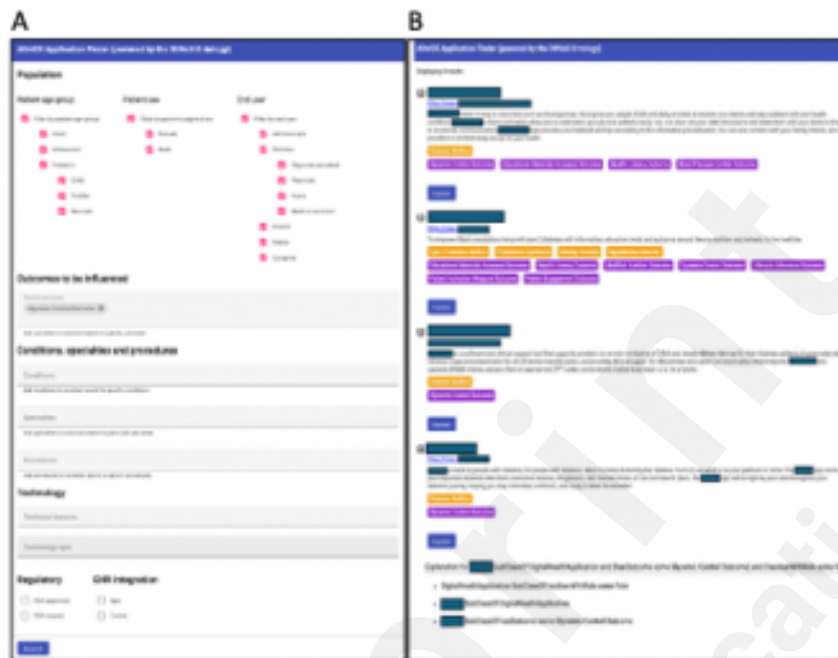
11. Auerbach AD, Neinstein A, Khanna R. Balancing innovation and safety when integrating digital tools into health care. *Ann Intern Med* 2018 May 15;168(10):733–734. PMID:29582072
12. Institut für Theoretische Informatik TU Dresden, Baader F, Brandt S, Institut für Theoretische Informatik TU Dresden, Lutz C, Institut für Theoretische Informatik TU Dresden. Pushing the EL envelope. Dresden: Technische Universität Dresden; 2005. [doi: 10.25368/2022.144]
13. OBO Foundry. OBO Foundry Principles: Overview [Internet]. Principles: Overview. [cited 2024 Sep 13]. Available from: <https://obofoundry.org/principles/fp-000-summary.html>
14. Gene Ontology Consortium, Aleksander SA, Balhoff J, Carbon S, Cherry JM, Drabkin HJ, Ebert D, Feuermann M, Gaudet P, Harris NL, Hill DP, Lee R, Mi H, Moxon S, Mungall CJ, Muruganugan A, Mushayahama T, Sternberg PW, Thomas PD, Van Auken K, Ramsey J, Siegele DA, Chisholm RL, Fey P, Aspromonte MC, Nugnes MV, Quaglia F, Tosatto S, Giglio M, Nadendla S, Antonazzo G, Attrill H, Dos Santos G, Marygold S, Strelets V, Tabone CJ, Thurmond J, Zhou P, Ahmed SH, Asanitthong P, Luna Buitrago D, Erdol MN, Gage MC, Ali Kadhum M, Li KYC, Long M, Michalak A, Pesala A, Pritazahra A, Saverimuttu SCC, Su R, Thurlow KE, Lovering RC, Logie C, Oliferenko S, Blake J, Christie K, Corbani L, Dolan ME, Drabkin HJ, Hill DP, Ni L, Sitnikov D, Smith C, Cuzick A, Seager J, Cooper L, Elser J, Jaiswal P, Gupta P, Jaiswal P, Naithani S, Lera-Ramirez M, Rutherford K, Wood V, De Pons JL, Dwinell MR, Hayman GT, Kaldunski ML, Kwitek AE, Laudederkind SJF, Tutaj MA, VEDI M, Wang S-J, D'Eustachio P, Aimo L, Axelsen K, Bridge A, Hyka-Nouspikel N, Morgat A, Aleksander SA, Cherry JM, Engel SR, Karra K, Miyasato SR, Nash RS, Skrzypek MS, Weng S, Westerfield M. The Gene Ontology knowledgebase in 2023. *Genetics* 2023 May 4;224(1):iyad031. PMID:36866529
15. Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, Davis AP, Dolinski K, Dwight SS, Eppig JT, Harris MA, Hill DP, Issel-Tarver L, Kasarskis A, Lewis S, Matese JC, Richardson JE, Ringwald M, Rubin GM, Sherlock G. Gene Ontology: Tool for the unification of biology. *Nat Genet* 2000 May;25(1):25–29. PMID:10802651
16. Whetzel PL, Noy NF, Shah NH, Alexander PR, Nyulas C, Tudorache T, Musen MA. BioPortal: enhanced functionality via new Web services from the National Center for Biomedical Ontology to access and use ontologies in software applications. *Nucleic Acids Res* 2011 Jul;39(Web Server issue):W541-5. PMID:21672956
17. Vasilevsky NA, Matentzoglou NA, Toro S, Flack JE, Hegde H, Unni DR, Alyea G, Amberger JS, Babb L, Balhoff JP, Bingaman TI, Burns GA, Callahan TJ, Carmody LC, Chan LE, Chang GS, Dumontier M, Failla LE, Flowers MJ, Garrett HA, Gration D, Groza T, Hanauer M, Harris NL, Helbig I, Hilton JA, Himmelstein DS, Hoyt CT, Kane MS, Köhler S, Lagorce D, Larralde M, Lock A, Santiago IL, Maglott DR, Malheiro AJ, Meldal BH, McMurry JA, Munoz-Torres M, Nelson TH, Ochoa D, Oprea TI, Osumi-Sutherland D, Parkinson H, Pendlington ZM, Rath A, Rehm HL, Remennik L, Riggs ER, Roncaglia P, Ross JE, Shadbolt MF, Shefchek KA, Similuk MN, Sioutos N, Sparks R, Stefancsik R, Stephan R, Stupp D, Sundaramurthi JC, Tammen I, Thaxton CL, Valasek E, Wagner AH, Welter D, Whetzel PL,

- Whiteman LL, Wood V, Xu CH, Zankl A, Zhang XA, Chute CG, Robinson PN, Mungall CJ, Hamosh A, Haendel MA. Mondo: Unifying diseases for the world, by the world. medRxiv 2022 Apr 16; [doi: 10.1101/2022.04.13.22273750]
18. Courtot M, Gibson F, Lister AL, Malone J, Schober D, Brinkman RR, Ruttenberg A. MIREOT: The minimum information to reference an external ontology term. Appl Ontol 2011;6(1):23–33. [doi: 10.3233/AO-2011-0087]
 19. Jackson RC, Balhoff JP, Douglass E, Harris NL, Mungall CJ, Overton JA. ROBOT: A tool for automating ontology workflows. BMC Bioinformatics 2019 Jul 29;20(1):407. PMID:31357927
 20. He Y, Sarntivijai S, Lin Y, Xiang Z, Guo A, Zhang S, Jagannathan D, Toldo L, Tao C, Smith B. OAE: the ontology of adverse events. J Biomed Semantics 2014 Jul 5;5:29. PMID:25093068
 21. Digital.Health [Internet]. [cited 2024 Apr 9]. Available from: <https://digital.health/new-products>
 22. Dodd S, Clarke M, Becker L, Mavergames C, Fish R, Williamson PR. A taxonomy has been developed for outcomes in medical research to help improve knowledge discovery. J Clin Epidemiol 2018 Apr;96:84–92. PMID:29288712
 23. Library of Digital Endpoints – Digital Medicine Society (DiMe) [Internet]. [cited 2024 Apr 9]. Available from: <https://dimesociety.org/get-involved/library-of-digital-endpoints/>
 24. Glöggl M, Ammenwerth E. Development and validation of a useful taxonomy of patient portals based on characteristics of patient engagement. Methods Inf Med 2021 Jun;60(S 01):e44–e55. PMID:34243191
 25. Ritchie J, Welch B. Categorization of Third-Party Apps in Electronic Health Record App Marketplaces: Systematic Search and Analysis. JMIR Med Inform 2020 May 29;8(5):e16980. PMID:32469324
 26. Abbes H, Gargouri F. MongoDB-Based Modular Ontology Building for Big Data Integration. J Data Semant 2017 Oct 27;7(1):1–27. [doi: 10.1007/s13740-017-0081-z]
 27. Rosner B. Digital medicine Outcomes Value Set (DOVeS) Ontology | NCBO BioPortal [Internet]. NCBO BioPortal. [cited 2024 Sep 27]. Available from: <https://bioportal.bioontology.org/ontologies/DOVES>
 28. Horridge M. dovesproject · GitHub [Internet]. Github. [cited 2024 Sep 28]. Available from: <https://github.com/dovesproject>

Supplementary Files

Figures

Prototype DHT Application Finder with underlying DOVeS ontology. A. Search for outcome “glycemic control” B. A set of search results. (Company/product names redacted).



Entry fields in the DHT Application Finder for A. Conditions, specialties, and procedures, and B. Outcomes to be influenced. The fields show type-ahead completion options for conditions known to the system, along with their synonyms that are drawn directly from the MONDO disease ontology.

A
Conditions, specialties and procedures

Search:

diabetes mellitus (DM, diabetes mellitus, diabetes, diabetes mellitus (diabetes))

type 2 diabetes mellitus (type 2 diabetes mellitus, non-insulin-dependent diabetes mellitus, diabetes mellitus, non-insulin-dependent, association with, non-insulin-dependent diabetes, type 2 diabetes, type 2 diabetes mellitus, non-insulin-dependent diabetes mellitus, non-insulin-dependent diabetes, NIDDM, T2DM - type 2 diabetes, severe, rigid, non-insulin-dependent diabetes mellitus, diabetes mellitus, type 2, diabetes mellitus, type 2)

gestational diabetes (intrauterine gestational diabetes mellitus, diabetes in pregnancy, gestational diabetes)

B
Outcomes to be influenced

Desired outcomes:

Glycemic Control Outcome

Reduction in Glycemic Events Outcome