dHealth 2024
D. Hayn et al. (Eds.)

doi:10.3233/SHTI240028

Automated Clinical Trial Cohort Definition and Evaluation with CQL and CDS-Hooks

Anja SCHWAB[a,1], Anna M. LIN[a] and Stephan WINKLER[a]

[a] *Department of Medical and Bioinformatics, University of Applied Sciences Upper Austria, Hagenberg, Austria*

Abstract. Background: Patient recruitment for clinical trials faces major challenges with current methods being costly and often requiring time-consuming acquisition of medical histories and manual matching of potential subjects. Objectives: Designing and implementing an Electronic Health Record (EHR) and domain-independent automation architecture using Clinical Decision Support (CDS) standards that allows researchers to effortlessly enter standardized trial criteria to retrieve eligibility statistics and integration into a clinician workflow to automatically trigger evaluation without added clinician workload. Methods: Cohort criteria are translated into the Clinical Quality Language (CQL) and integrated into Measures and CDS-Hooks for patient- and population-level evaluation. Results: Successful application of simplified real-world trial criteria to Fast Healthcare Interoperability Resources (FHIR®) test data shows the feasibility of obtaining individual patient eligibility and trial details as well as population eligibility statistics and a list of qualifying patients. Conclusion: Employing CDS standards for automating cohort definition and evaluation shows promise in streamlining patient selection, aligning with increasing legislative demands for standardized healthcare data.

Keywords. Clinical Trial, Clinical Quality Language, Health Level 7, Clinical Quality Measures, Clinical Decision Support Systems

1. Introduction

1.1. Problem statement

Recruitment of patients for early-stage clinical trials is a complex and challenging process. Traditional methods often include manual evaluation of patients' eligibility by clinicians. The idea of automation is not new, but many solutions are specific to individual clinical domains and highly dependent on specific Electronic Health Record (EHR) systems. [1–4]

[1] Corresponding Author: Anja Schwab, University of Applied Sciences Upper Austria, Hagenberg, Austria, anja.schwab939@gmail.com

1.2. Background

Recent studies have highlighted the growing adoption of Clinical Decision Support (CDS) standards in tool development, underlining their potential to transform clinical trial recruitment processes [5]. Related Fast Healthcare Interoperability Resources (FHIR®)-based publications include estimations on the expected number of eligible patients [6], an automated recruitment system for a cardiology department [7], a preliminary trial eligibility assessment [8], and a recruitment registry [9]. The proposed solutions are either domain-specific or focus on a single aspect of the recruitment process.

1.3. Objective

The objective of this work is to design and implement an (a) EHR-independent, non-domain specific architecture that (b) enables researchers to enter standardized study criteria without technical authoring knowledge, (c) generates overall eligibility statistics and (d) obtains a list of qualifying patients to potentially contact them. For clinicians who participate in the enrolment process (e) all studies that a specific patient qualifies for should automatically be evaluated without any required knowledge of existing trials, and (f) provide automated evaluation and resulting display of relevant information integrated directly into their workflow avoiding any additional administrative effort.

2. Methods

The choice of standards is essential for ensuring interoperability while building EHR-independent, non-domain-specific systems. As a quickly evolving, EHR-independent standard for healthcare data exchange, the FHIR® standard [10] was chosen as the center of this architecture to ensure seamless, standardized data exchange for accurate clinical decision-making to attain objective (a).

2.1. CQL for Logic Definitions of Inclusion and Exclusion Criteria

The Clinical Quality Language (CQL) [11] allows researchers to define highly complex and specific inclusion and exclusion criteria that integrate directly with FHIR®-based medical records. In contrast to the limitations of assessing eligibility by regular FHIR® search [8], the CQL enables flexible and extensive authoring capabilities with the option of reusability of the same logic definitions for patient- and population-level evaluations. Narrative inclusion and exclusion criteria of the study cohorts can be authored into arbitrarily complex CQL expressions that are evaluated to assess an overall eligibility status of either *true* – a patient qualifies or *false* – a patient is not eligible for the trial. Existing CQL authoring tools allow for the generation of such CQL expressions without explicit knowledge of the underlying data retrieval syntax, e.g. via User Interface

(UI) prompts [12] to achieve the objective (b). FHIR® R4 is being used because the most current FHIR® release R5 is not yet integrated into a CQL engine.

2.2. CDS-Hooks for Patient-Level Triggers in the Clinician's Workflow

CDS rules can be defined by CQL expressions that are automatically applied to FHIR® data. CDS-Hooks [13] are a CDS standard that can trigger logic within an EHR system upon certain actions such as opening a patient chart and therefore, enabling a direct integration of clinical knowledge within a real-time healthcare system [11].

With a patient-view CDS-Hook implementation, the evaluation is triggered upon opening a patient file – allowing for automated checks that run in the background of the clinician's workflow. A dialog window could inform the clinician about available trials in real-time for one specific patient without additional interaction, and without the requirement that clinicians have explicit knowledge about existing trials accomplishing objectives (e) and (f) - ultimately increasing awareness of existing trials. CDS-Hooks act as the bridge between the abstract decision logic and its practical application within real-time clinical environments.

2.3. FHIR® Measures for Automated Population-Level Evaluations

The defined CQL expressions can be applied to an entire population by using a FHIR® Measure, invoked by the FHIR® Operation *$evaluate-measure*[2]. The resulting FHIR® MeasureReport returns a total number of eligible patients as well as a list of qualifying research subjects [14] targeting the objectives defined in (c) and (d).

2.4. Application of the Standards to the Clinical Trial Use Case

The UI is realized through a JavaFX prototype, which serves to demonstrate potential application within a research environment. Clinician workflows are simulated using manually created *hookInstances* that mimic automatic triggers in a real-world clinical context. The core component, the "CQL Factory," uses Java to generate necessary FHIR® resources and CQL through a template-oriented approach. This process is enhanced by additional functionalities for converting user-input criteria into executable CQL expressions. The architecture facilitates adaption to any language or tool that can interact with a FHIR® server via RESTful API. Figure 1 provides a detailed visualization of the architecture.

[2] http://hl7.org/fhir/R4/measure-operation-evaluate-measure.html

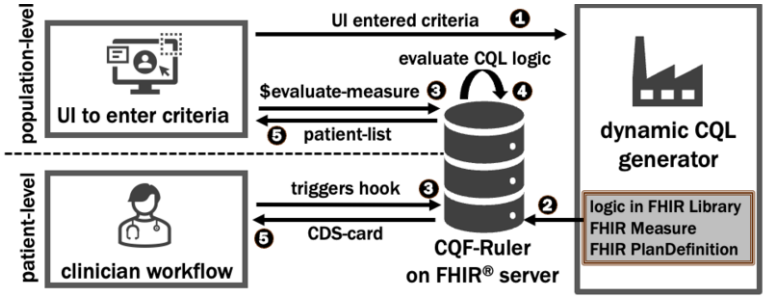

Figure 1. Architecture: New trials can be added by the researcher via the UI (1). The tool automatically translates the criteria into the corresponding CQL expressions targeting objective (b). The FHIR® Measure resource for population-level evaluation and the PlanDefinition[3] for the CDS-Hook are generated and include references to the CQL logic library. The resources are posted to a FHIR® server, along with any dependency libraries (2). The open-source project CQF-Ruler[4] generates the CDS-Hook from the corresponding FHIR® PlanDefinition with an embedded action trigger (3) to evaluate eligibility based on CQL expressions (4). A CDS-card returns any relevant information that should be displayed to the clinician (5). Within the clinical trial UI, the *$evaluate-measure* operation can be invoked (3). The CQF-Ruler again evaluates the CQL expressions (4) and generates statistics as well as optionally a list of all eligible Patient ids (5) depending on the report type.

The architecture uses key FHIR® resources for different aspects: Patient and related resources for CQL evaluations applicable to all use cases, the Library resource for CQL execution across all scenarios, PlanDefinition for CDS-Hooks in patient-level, and Measure for the *$evaluate-measure* operation in population-level evaluations.

3. Results

The proposed standards were applied to a real-world clinical trial scenario demonstrating the architecture's technical feasibility in patient eligibility evaluation and its potential for integration into clinical workflows by achieving all objectives defined in Section 1.

3.1. Application to a Real-World Clinical Trial[5]

The criteria of trial NCT04753502 as illustrated in Table 1 exhibit an ideal level of detail to demonstrate the framework's technical ability to accurately process and interpret the specific requirements of clinical studies. They are sufficiently detailed to test the system's capabilities but not so intricate as to overshadow the framework's functionality with the complexities of CQL expression generation. These five types of criteria have been selected: Age, Gender, Pregnancy Status, Tobacco Use, Condition inclusion and exclusion codes. The cohort criteria were simplified as the focus is on the application of CQL and not on authoring the logic itself. Therefore, the logic definition of the location of care, for example, is

[3] http://hl7.org/fhir/R4/cdshooksserviceplandefinition.html
[4] https://github.com/cqframework/cqf-ruler
[5] Relevant FHIR® Resources, CQL expressions as well as a postman collection with a complete walkthrough are publicly available on GitHub: https://github.com/1anja1/ClinicalTrialCQL

disregarded. The same workflow can be applied to trial criteria in any medical domain and any EHR that is FHIR®-based consequently achieving objectives (a) and (b). The Agency for Healthcare Research and Quality created a CDS-authoring tool that shows that the generation of more complex definitions is possible [12]. Table 2 gives an overview of three example FHIR® test patients.

Table 1. Simplified criteria of trial NCT04753502 from https://clinicaltrials.gov

Min Age	Max Age	Gender	Pregnancy Status	Tobacco Use	Inclusion Criteria	Exclusion Criteria
18	-	female	true	Not relevant	Appendicitis	-

Table 2. Overview of example test patients

Patient ID	Age	Gender	Pregnancy	Tobacco Use	Conditions
Patient1	5	male	false	false	Asthma (J45.909), Acute appendicitis (K35)
Patient2	25	female	true	false	Acute appendicitis (K35)
Patient3	44	female	true	true	Malignant neoplasm of breast (C50)

{"cards": [{ "summary": "Clinical Trial Qualification",
"detail": "The Patient qualifies for trial: Laparoscopic Treatment for Appendicitis During Pregnancy",
"source": {"label": "Clinical Trial Details",
"uri": "https://clinicaltrials.gov/study/NCT04753502"}, ...} (1)

GET server-base/Measure/ID/$evaluate-measure?reportType=subject-list (2)

{"resourceType": "MeasureReport",
"contained":[{…"entry": [{"item":{"reference":
"Patient2"}}]}], …
"group": [{"population": [{ … "count": 1, … }] … }]} (3)

Listing 1 shows the card that gets returned when the hook is invoked on the Patient resource with the id *Patient2* – who is eligible for the trial. In a clinician workflow, the CDS-Hook would get triggered upon opening a patient chart and the card would be displayed as a dialog window. This action gets simulated by sending a *hookInstance* of a *patient-view* hook to the FHIR® server which also serves as a CDS-Hooks endpoint consequently achieving the defined objectives (e) and (f). The HTTP request to the FHIR® server for population-level evaluation is shown in Listing 2. The reportType *subject-list* supports returning

a list of Patient ids. If a reportType is not specified, a summary report results in counting an anonymous total of eligible subjects. Listing 3 demonstrates the successful implementation of objectives (c) and (d) by supplying the number of qualifying patients ("count") as well as a list of their ids upon calling the *$evaluate-measure* operation on all three test patients shown in Table 2. Using this, researchers can gain insights into eligibility statistics and, if the research platform's consent guidelines permit, they may also proactively reach out to eligible participants.

4. Discussion

It is important to acknowledge that criteria for real-world clinical trials can be exceedingly complex. CQL possesses the capability to handle the level of complexity required for comprehensive clinical knowledge. Existing CQL authoring tools [12] already showed the creation of CQL definitions based on user input. The novelty of the approach presented in this paper, however, extends beyond merely generating CQL. It is primarily aimed to evaluate the feasibility of utilizing CDS standards for automating patient selection in clinical trials. It encapsulates the entire workflow from the translation of criteria to its application in clinical workflows. This comprehensive process, involving a FHIR® Measure for population assessment and a CDS-service for integration into a clinician workflow underlines the power and versatility of CDS standards. This work lays the foundation for a CDS-based architecture for clinical trial patient selection with the potential to expand this framework in the future to accommodate the intricate complexities and specific cohort requirements prevalent in various medical domains.

4.1. Limitations & Outlook

As demonstrated in Chapter 3, it is technically feasible to create CDS logic with CQL which can be applied both at the patient and population level. However, in the absence of a generic mechanism for parameterizing measure logic within the FHIR® R4 *$evaluate-measure* operation, the utilization of diverse criteria remains unstructured. This limitation necessitates direct modifications to the CQL expressions. As a workaround, a CQL template script gets directly modified by replacing the specified default values of the parameters with the values defined by the researchers for the cohort.

The ability to pass the FHIR® Parameters Resource to *$evaluate-measure* in a future FHIR® version is anticipated to streamline this process significantly. It will enable the input of criteria into the CQL logic via the Parameters Resource, thereby possibly eliminating the need for a new library for each clinical trial. It is important to note that this impacts only the architecture implementers, not the researchers entering trial criteria.

There is potential for combining the CDS-Hook with a *Substitutable Medical Applications, Reusable Technologies (SMART)* application. This integration could facilitate seamless utilization within EHRs and health portals, enhancing the utility and applicability of the presented approach. Future research should explore the integration of the framework into a real-world FHIR®-based hospital information system, providing valuable insights into its applicability and effectiveness in a clinical setting.

5. Acknowledgements

Special thanks to Bryn Rhodes and his clinical reasoning team for their exceptional mentorship and expertise in knowledge authoring and their feedback, which have been crucial to the development of this paper.

References

[1] B. Kankova and G. Duftschmid, Reusing Data of an EU-Wide EHR System for Clinical Trials: A Capabilities Analysis., *dHealth* (2020), 17–22
[2] M. Dugas et al., Routine data from hospital information systems can support patient recruitment for clinical studies, *Clinical Trials* **7**(2) (2010), 183–189
[3] G. Melzer et al., Leveraging real-world data for the selection of relevant eligibility criteria for the implementation of electronic recruitment support in clinical trials, *Applied Clinical Informatics* **12**(01) (2021), 17–26
[4] G. De Moor et al., Using electronic health records for clinical research: the case of the EHR4CR project, *Journal of biomedical informatics* **53** (2015), 162–173
[5] P. Taber et al., New standards for clinical decision support: a survey of the state of implementation, *Yearbook of medical informatics* **30**(01) (2021), 159–171
[6] A. Banach et al., APERITIF–automatic patient recruiting for clinical trials based on HL7 FHIR, in *Public Health and Informatics*, IOS Press, (2021), 58–62.
[7] C. Scherer et al., Implementation of a clinical trial recruitment support system based on fast healthcare interoperability resources (FHIR) in a cardiology department, *European Heart Journal-Digital Health* **3**(4) (2022), ztac076–2795
[8] C. Gulden et al., Investigating the Capabilities of FHIR Search for Clinical Trial Phenotyping, *Studies in health technology and informatics* **253** (2018), 3–7
[9] A. C. Cheng et al., ResearchMatch on FHIR: Development and evaluation of a recruitment registry and electronic health record system interface for volunteer profile completion, *Journal of Clinical and Translational Science* **7**(1) (2023), e222
[10] Health Level Seven International, FHIR v4.0.1, (2019),. https://hl7.org/fhir/R4
[11] Health Level Seven International, Clinical Quality Language (CQL) v1.5.2, (2023),. https://cql.hl7.org/
[12] Agency for Healthcare Research and Quality, CDS Connect Authoring,. https://cds.ahrq.gov/cdsconnect/authoring (accessed January 20, 2024)
[13] HL7 & Boston Children's Hospital, CDS Hooks 2.0 Release, (2022),. https://cds-hooks.hl7.org/2.0/ (accessed January 20, 2024)
[14] A. M. Lin et al., From Authoring to Evaluating an Electronic Health Quality Measure–Applying Logic to FHIR® with CQL for Calculating Immunization Coverage, in *dHealth 2023*, IOS Press, (2023), 12–17.