

Using Computable Phenotypes in Point-of-Care Clinical Trial Recruitment

Martin CHAPMAN^{a,1}, Jesús DOMÍNGUEZ^a, Elliot FAIRWEATHER^a,
Brendan C. DELANEY^b and Vasa CURCIN^a

^aKing's College London, London, United Kingdom

^bImperial College, London, United Kingdom

Abstract. A key challenge in point-of-care clinical trial recruitment is to autonomously identify eligible patients on presentation. Similarly, the aim of *computable phenotyping* is to identify those individuals within a population that exhibit a certain condition. This synergy creates an opportunity to leverage phenotypes in identifying eligible patients for clinical trials. To investigate the feasibility of this approach, we use the Transform clinical trial platform and replace its archetype-based eligibility criteria mechanism with a computable phenotype execution microservice. Utilising a phenotype for *acute otitis media with discharge* (AOMd) created with the *Phenoflow* platform, we compare the performance of Transform with and without the use of phenotype-based eligibility criteria when recruiting AOMd patients. The parameters of the trial simulated are based on those of the REST clinical trial, conducted in UK primary care.

Keywords. clinical trials, phenomics, learning healthcare system, microservices

1. Introduction

Randomised clinical trials (RCT) remain the standard for the approval of new treatments in the context of healthcare. However, the process of translating research findings into clinical care is lengthy, as RCTs are often separated from the clinical care context. This is reflected in, for example, the use of dedicated Electronic Data Capture (EDC) trial systems. The concept of *eSource* has emerged to address this issue by embedding clinical trial functionality directly into routine electronic health record (EHR) systems; the 'no-data-wasted' philosophy of the Learning Health Systems (LHS) paradigm.

The *Transform* platform—developed as part of the EU-funded FP7 programme—provides *eSource* connectivity for RCTs by integrating with a range of EHR systems to perform automated eligibility checking at the point-of-care. Transform was successfully piloted in Europe, where it was integrated with diverse EHR systems and recruited and managed the follow-up of a total of 600 patients [1]. More recently, Transform was used to facilitate the UK-based Runny Ear Study (REST), a pragmatic RCT exploring alternative treatments to antibiotics for *acute otitis media with discharge* (AOMd) in children, which currently costs the UK NHS over £3M each year. The study aimed to recruit 399 participants across 175 general medical practices [2].

Transform specifies eligibility criteria (EC) using a subset of openEHR-based archetypes, bound to EHR-agnostic elements from the clinical data integration model

¹ Corresponding Author: Martin Chapman; E-mail: martin.chapman@kcl.ac.uk

(CDIM) to enable cross-source compatibility [3]. Upon encountering a vendor-specific EHR, each archetype is translated to a suitably formatted set of queries that can be executed against the record to determine eligibility (e.g. XQueries for an XML schema). To represent EC in REST, two key archetypes were created, one to specify participant age, which must be between 1 and 15, and another to specify a diagnosis of AOMd using the Read coding scheme, including *acute right (F527.)*, *acute with effusion (F510.)* and *acute suppurative (F520.)*. To represent participant age, the former uses a fixed date of birth range, in response to the fact that direct age information is often omitted from a patient's EHR. Similarly, the latter uses a finite list of codes representing AOMd, including those given above. As an example, the DOB archetype is shown in Figure 1, binding the *date of birth value* attribute (part of the DOB event)—with a matching constraint corresponding to the date range specified—to its CDIM representation, *CDIM_000007*.

```
<Archetype><![CDATA[ archetype (adl_version=1.4) concept [at0000]
definition EVENT[at0000] matches { attributes cardinality matches {1..1; ordered} matches {
  ATTRIBUTE[at0001] matches { value matches {|2003-12-01..2018-12-01|} } } }
ontology terminologies_available = <"CDIM", ...> term_definitions = < [{"en"} = < items = <
  [{"at0000"} = < text = <"date of birth"> type = <"Event"> >
  [{"at0001"} = < text = <"date of birth value"> type = <"Date"> > > >
term_binding = < [{"CDIM"} = < items = <
  [{"at0001"} = <[CDIM::CDIM_000007]> > > > ]]></Archetype>
```

Figure 1. Transform's date of birth (DOB) archetype.

While beneficial in respect of data source flexibility, the requirement to define EC archetypes from scratch represents an overhead in trial preparation. A more expedient option is to utilise standardised, peer-reviewed, criteria definitions, which can be found as patient *phenotypes*, a term that in this context describes a set of individuals who all exhibit the same condition. *Phenotyping* is then the process of defining a set of rules that enable the identification of such a cohort, and *computable phenotypes* are the realisations of these rules in a machine-executable form, such as SQL or FHIR queries. Significant efforts are underway to develop large, validated collections of phenotypes, e.g. [4]. In addition, the benefits of using phenotypes to determine eligibility have been shown in offline trial recruitment [5]. Therefore, in this work, we investigate the suitability of using phenotypes as the basis for determining patient eligibility within Transform to run a trial such as REST.

2. Methods

To evaluate our approach, a novel AOMd phenotype was developed to represent the EC from the REST study. For this, we utilised the *Phenoflow* authoring platform, which allows the logic of a phenotype to be captured as a multi-layer, high-level representation, which can then be connected to different implementation units, and is executable across different platforms [6]. The ability for phenotypes to be executed across different platforms is essential in the Transform system, where components are executed within a variety of different clinical environments (e.g. general practitioner (GP) surgeries).

To construct a phenotype's logic within Phenoflow, one specifies a series of steps, each of which contains a piece of logic required to determine a patient's membership of

a particular condition cohort. To construct our AOMd phenotype, we took the logic of the two existing EC archetypes (e.g. Figure 1), and represented each as a distinct step within our definition, along with an appropriate implementation unit realising the phenotype logic itself. As an example, the step for determining eligibility via the clinical codes present in a patient's record is shown in Figure 2. Notably, in the construction of this step, we were able to leverage the flexibility provided by Phenoflow to obtain a dynamic list of codes relating to AOMd, from sources such as the Unified Medical Language System (UMLS).

To plug the AOMd phenotype into Transform, we developed a computable phenotype execution microservice (PhEM), shown in Figure 3 alongside the platform's existing services. When a patient record is examined for eligibility, the platform either utilises the standard archetype approach, or utilises this new service (passing to it the record as a specialised single patient population), depending on how it is configured. Note that the insight into the structure of individual EHR schema provided by the translated archetypes is also used by PhEM in order to retain data source flexibility.

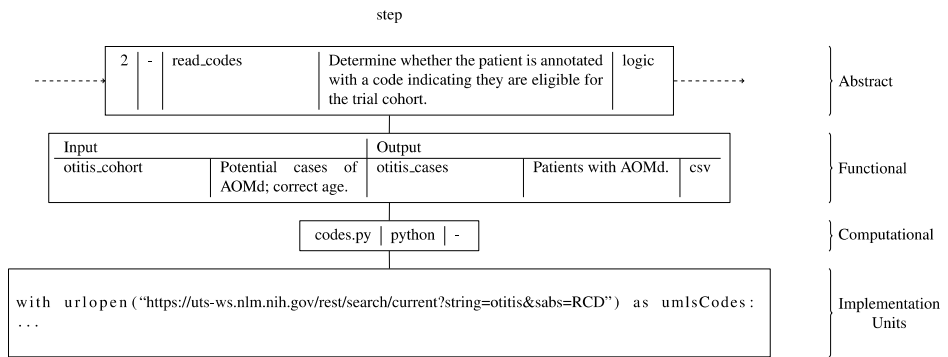


Figure 2. Step of Phenoflow-based AOMd phenotype, under multi-layer representation, and implementation.




Figure 3. Overview of Transform architecture, with additional phenotype execution component.

2.1. Trial recruitment

To evaluate the use of PhEM in clinical trial recruitment, we simulated the REST trial using synthetic data generated by the *Synthea* patient population generator [7]. 10258 patients were generated, to reflect the size of the real population presenting across GP practices (Section 1). Similarly, to ensure the data was suitably representative, we used

real demographics and condition progression distributions, provided by Synthea as a part of a module specific to AOMd. Those generated patients with AOMd exhibited one of five forms of the condition: the three forms specified earlier, as well as ‘caused by COVID-19’ (*Y20ff*) and *acute bilateral (F528.)*, the latter representing patients *not* targeted by the trial. Transform was also augmented with a novel simulation mode, which interfaces with a FHIR-compliant simulated EHR system to access population records, and emulates the interactions of a GP, such as the completion of electronic Case Report Forms (eCRFs) to confirm patient eligibility. The execution of Transform within a portion of the simulations was tracked using debugging tools (e.g. stepwise execution), to ensure that the recruitment or omission of patients was not due to software error.

3. Results

Table 1. Population demographics, and patients recruited by the archetype and PhEM mechanisms.

		All	1-10	10-17	F527.	F520.	F510..	Y20ff	F528.
Total		10258	1085	883	166	201	153	181	179
Eligible		631	528	103	148	185	137	161	0
Recruited	Archetype	455	375	80	148	176	131	0	0
	PhEM	631	528	103	148	185	137	161	0




Figure 4. Transform’s recruitment by age.




Figure 5. Transform’s recruitment by code.

Our results show that PhEM enables Transform to recruit at least the same number of patients as the archetype mechanism (Table 1). Moreover, we can see that PhEM actually recruits more accurately than the archetypes, across a number of controlled sub-populations. These include those patients with the clinical code *F527.* (Figure 4) and those born in 2010 (Figure 5). Figure 4 shows that Transform's archetype mechanism misses patients at the lower end of the recruitment age range (false negatives), e.g. age 1, and erroneously recruits patients that exceed the age range (false positives), e.g. age 16. This is likely because Transform's archetype definitions are only able to express a fixed birth date range, which does not offer the same dynamic understanding of age that the phenotype implementation is able to extract. Figure 5 shows that the archetype mechanism misses patients with key clinical codes in the population that are not statically defined at the start of the trial, such as COVID-19 codes (e.g. *Y20ff*), which emerged during REST. In contrast, the phenotype implementation, with its dynamic access to updated code sets, is able to sufficiently identify patients exhibiting these codes.

As execution time is a key factor in point-of-care clinical trial recruitment, to complement our results we also recorded how long it took both the archetypes and PhEM to determine eligibility. Median (with quartiles) execution times were 12.37 (11.74; 12.49) and 22.22 (20.15; 24.30) seconds, respectively, which represents a statistically significant increase ($p < 0.05$). We attribute this to an increase in computational overhead when PhEM is used.

4. Discussion and Conclusion

In this work, we implement a phenotype execution microservice (PhEM) for an existing eSource platform, Transform, and compare its performance to the use of archetypes when determining trial eligibility in a population of patients targeted by the REST study. Our results show that PhEM offers improved accuracy when recruiting patients, suggesting that existing phenotypes can be leveraged in the definition of EC, to avoid their novel definition in trial-specific formats. For example, different phenotypes can now be plugged into Transform to be used as a part of future trials. This reduces trial preparation burden, but does assume the existence of standardised phenotype definitions for relevant conditions, which may not always be the case.

While the improved accuracy introduced by PhEM is attributable to static archetype definitions, it is worth noting that archetypes can be updated during a trial. However, this requires expert intervention and introduces the potential for error. Moreover, while it may not be common for terminologies to be updated during recruitment (outside of the advent of COVID-19), it is important to account for this possibility in trials such as REST, where the efficacy of the recruitment procedure is key. Despite this increase in accuracy, our results also show that the cost of introducing PhEM is an increase in execution time. As such, employing phenotypes is a trade-off between the two.

References

- [1] Delaney BC, Curcin V, Andreasson A, Arvanitis TN, Bastiaens H, Corrigan D, et al. Translational Medicine and Patient Safety in Europe: TRANSFoRm—Architecture for the Learning Health System in Europe. *BioMed research international*. 2015;2015:36–43.
- [2] National Institute For Health Research (NIHR). Immediate oral, immediate topical or delayed oral antibiotics for acute otitis media with discharge (the Runny Ear Study: REST); 2018, Available at <https://www.journalslibrary.nihr.ac.uk/programmes/hta/168501>.
- [3] Ethier JF, Curcin V, Barton A, McGilchrist MM, Bastiaens H, Andreasson A, et al. Clinical data integration model. Core interoperability ontology for research using primary care data. *Methods of information in medicine*. 2015;54(1):16–23.
- [4] Denaxas S, Gonzalez-Izquierdo A, Direk K, Fitzpatrick NK, Fatemifar G, Banerjee A, et al. UK phenomics platform for developing and validating electronic health record phenotypes: CALIBER. *JAMIA*. 2019;26(12):1545–1559.
- [5] Ahmad FS, Rickett IM, Hammill BG, Eskenazi L, Robertson HR, Curtis LH, et al. Computable Phenotype Implementation for a National, Multicenter Pragmatic Clinical Trial: Lessons Learned from ADAPTABLE. *Circulation: Cardiovascular Quality and Outcomes*. 2020;13(6):355–364.
- [6] Chapman M, Rasmussen L, Pacheco J, Curcin V. Phenoflow: A Microservice Architecture for Portable Workflow-based Phenotype Definitions. *AMIA Summits on Translational Science* (forthcoming). 2021.
- [7] Walonoski J, Kramer M, Nichols J, Quina A, Moesel C, Hall D, Duffett C, Dube K, Gallagher T, McLachlan S. Synthea: An approach, method, and software mechanism for generating synthetic patients and the synthetic electronic health care record. *J Am Med Inform Assoc*. 2018 Mar 1;25(3):230–238.