# FHIR FLI: An Open Source Platform for Storing, Sharing and Analysing Lifestyle Data

Kiran Gopinathan[1], Nikolaos Alexandros Kaloumenos[1], Kinnari Ajmera[1], Alexandru Matei[1], Ian Williams[2] and Andrew Davis[2]
[1]*University College London, United Kingdom*
[2]*Nuffield Health, United Kingdom*

Keywords: Personal Health Record, Lifestyle Data, FHIR.

Abstract: Consumers and healthcare organisations alike are increasingly interested in using digital health solutions to reduce the risk of chronic conditions or to help manage these conditions outside hospitals. Equally, there is a strong public health benefit in helping individuals adopt and improve healthy lifestyle behaviours. The first step in this direction is the ability to record and analyse lifestyle data. Currently, lifestyle logging platforms use proprietary data formats. Data is segregated among different platforms, impacting consumers, service providers, research institutes and public health bodies. Our aim is to facilitate the transfer of information between individuals and organisations that hold or require their lifestyle data. We demonstrate that an open source platform based on a clinically recognised interoperability standard - Fast Healthcare Interoperability Resources (FHIR) - can meet both consumers and industry needs. We use as an example the case of people managing arthritis. Our contributions are: (i) an extension of the FHIR standard for lifestyle data, (ii) a reference architecture for a Personal Lifestyle Record, (iii) integration with voice-enabled digital assistants for lifestyle data capture and (iv) an open source implementation of this architecture that retrieves, saves and analyses lifestyle data from wearable devices.



## 1 INTRODUCTION

The UK population is living longer, but at the same time people spend more years in ill-health, managing chronic conditions (Public Health England, 2015).

To cope with the increasing burden of disease, more and more resources are dedicated to prevention and management of chronic conditions within the community. For example, a recent study in the UK has shown a remission in Type 2 diabetes for participants following a low-calorie, diet-based, weight management programme. Throughout the process, the trial participants had interaction with their GP and could access a dietician (Lean et al., 2017)

As fitness, well-being and lifestyle data becomes more relevant in the provision of care, consumers and healthcare professionals need the ability to aggregate disparate information about one's lifestyle into a consolidated view spanning across the years, across providers, and covering many data sources. Improved access to lifestyle data can lead to better quality of care, better health outcomes and improved customer satisfaction.

We consider lifestyle data any measurement related to:

- lifestyle risk factors (e.g. physical activity)
- mental health (e.g quality of sleep, stress levels)
- monitoring of chronic conditions (e.g blood glucose)
- maintaining or improving ability (e.g physiotherapy, rehabilitation, mental resilience, elite sports).

The current landscape in the industry is very fragmented - each organisation that generates lifestyle data uses proprietary formats. The challenge we aim to address in this paper is the lack of a true platform approach, in the well-being and lifestyle domain.

In a platform approach, the analysis of data (for day-to-day monitoring, service personalisation or behaviour change) is separated from data collection. For this to happen, there is a need for an encompassing and open standard to describe lifestyle data in a consistent manner, agnostic of the source systems generating the data. This will open the way for services specifically designed to deliver value by interpreting



Gopinathan, K., Kaloumenos, N., Ajmera, K., Matei, A., Williams, I. and Davis, A.
FHIR FLI: An Open Source Platform for Storing, Sharing and Analysing Lifestyle Data.
DOI: 10.5220/0006791302270233
In *Proceedings of the 4th International Conference on Information and Communication Technologies for Ageing Well and e-Health (ICT4AWE 2018)*, pages 227-233
ISBN: 978-989-758-299-8

the data in new, useful ways. This is the path followed in the medical community, where much effort has been put into defining interoperability standards to describe clinical data.

In this paper, we propose an architecture for a Personal Lifestyle Record that offers capabilities to retrieve, store and analyse lifestyle data in a provider-agnostic environment, illustrating the potential of a platform approach. To inform our proposal, we use as an illustrative case study the management of arthritis.

The paper is organised as follows. Section 2 presents the current technology landscape for managing and using lifestyle data. Section 3 discusses an illustrative customer journey, in arthritis management. Section 4 proposes a method for standardisation of lifestyle data and offers two specific examples. Section 5 presents a reference architecture that decouples data generation from data analysis. Section 6 discusses emerging user interaction channels, such as voice-enabled digital assistants for capturing lifestyle data. Section 7 describes the implementation work done to date and Section 8 summarises our contributions and discusses future work.

## 2 BACKGROUND

Smart-phones have become a leading consumer choice for managing lifestyle data. They are a natural hub to bring together data coming from built-in sensors on the phone, data coming from wearable devices that synchronise with the phone and data coming from mobile apps installed on the phone. There are a number of limitations when using a smart-phone based lifestyle record:

- it is not an inclusive approach: Apple Healthkit data can only be accessed from an iPhone.
- it does not cover the needs of service providers, who need to access and aggregate data from a population of clients

Solutions created by wearable companies have the advantage of being accessible across platforms, mobile and web. They focus on analysing and providing insight on data collected through the wearable. They have two main limitations:

- the data model of each platforms is controlled by the vendor and it only contains items relevant within the context of use for each wearable - third parties cannot extend the data model.
- the data is still locked, as Fitbit users can only compare themselves against other Fitbit users for example.

Cloud-based platforms like Open MHealth or Microsoft Health Vault have cross-platform accessibility and allow developers to extend the data model. Information can be shared with service providers. However, these platforms also have some limitations.

- at the moment, they still use proprietary data formats
- third-party developers cannot develop custom visualisations or predictive models

There are two conclusions we can draw from reviewing the current landscape. First, there is a clear direction towards enabling data movement. Many wearable companies allow customers to download their own data for example. Conversely, other companies position themselves as data repositories, offering the necessary integrations for consumers to bring data into their platforms.

Second, there is no open, standard information model for lifestyle data, with clear governance supporting it, which is not tied to a specific commercial organisation.

We argue that such an information model could well be developed following an open and inclusive approach, as opposed to being developed internally by any one company. There is a large body of work which could inform these efforts, stemming from the clinical community pursuit to standardise the way medical data is represented.

FHIR (Bender and Sartipi, 2013) is such a standard, supported by HL7 and adopted by a growing number of healthcare organisations, including the NHS (Kavanagh, 2017).

The FHIR specification defines a set of core resources (*Person*, *Observation*, *Diagnostic* etc). and an infrastructure for handling resources. FHIR Observations are already used to describe:

- Vital signs: e.g. temperature
- Laboratory Data: e.g. blood glucose
- Devices Measurements: e.g. pulse oximetry
- Clinical assessments: e.g. Glasgow Coma Score
- Personal characteristics: e.g. eye colour

The FHIR standard can be extended to suit particular contexts of use - including management of lifestyle data. Extending FHIR to suit particular scenarios of use is achieved by creating FHIR *Profiles*. A profile extends or restricts core FHIR resources, based on the domain-specific needs.

A FHIR extension for lifestyle data is being developed as part of the Finnish national Personal Health Record infrastructure (Kanta, 2017). However, their approach is to create a limited number of FHIR observations, describing fitness information aggregated over a period of time - for example 'Distance reached in 24h'. We favour an approach where any data point created at the source system is mapped to a FHIR observation. In our approach, three distinct walking / running episodes in a 24 hour window would be represented as three FHIR observations, not a single one.

## 3 ILLUSTRATIVE CASE STUDY

Lifestyle data can be used to improve the health and well-being of people of all ages.

To discuss our proposal for managing lifestyle data, we consider the case of arthritis, a group of chronic conditions that affect the joints in the body. As of 2017, 400,000 adults in the UK have rheumatoid arthritis, and prevalence increases with age. Further, 8.75 million people in the UK have sought treatment for osteoarthritis, where prevalence also increases with age (Arthritis Research UK, 2017).

For the management of rheumatoid arthritis, the National Institute for Health and Care Excellence (NICE) in the UK recommends access to physiotherapy services, to improve fitness and encourage regular exercise (Deighton et al., 2009). Similarly, for the management of osteoarthritis NICE recommends physical exercise for local muscle strengthening and general aerobic fitness (Conaghan et al., 2008).

In term of assessing pain intensity - a key measurement for arthritis management - there are different scales available, with good degree of correlation among them (Downie et al., 1978). Once could use a descriptive scale (nil, mild, moderate, severe, very severe), a numeric scale (0 - no pain, 10 - worst possible pain) (Farrar et al., 2001) or a visual analogue scale (100 mm in length anchored by the two extremes) (Hawker et al., 2011).

Based on the clinical context, there is a demonstrable need for data related to physical activity and pain levels to be shared between users, GPs and physiotherapists. The hypothetical customer journey we use as a source of requirements is then: a patient diagnosed with arthritis has been referred by his GP to a physiotherapist. The physiotherapist recommends a regime of moderate physical activity, which includes outside running sessions and pilates classes at a local gym. The user would like to share a full record of his physical activity with the physiotherapist, as well as be able to track how his symptoms (i.e. pain levels) change over time, as a result of his efforts. In turn, the physiotherapist would like to share summary data with the patients GP.

To support this journey, we need to enable three types of information flows.

The first scenario we aim to support is a person centric view: allowing one individual to aggregate data about himself which is held by different platforms and providers. For example, somebody managing osteoarthritis could attend physiotherapy sessions, he could run in the park during the weekend and attend a T'ai chi class in the gym once a week. He should have access, combine and own all of this data - information generated by a specialist, by a wearable, by a fitness service provider, or self reported.

The second scenario we aim to support is an organisation centric view: allowing service providers to receive data shared with them by their clients. For example, a physiotherapist should be able to access data about physical activity levels and self reported pain levels from each individual under treatment. Equally, a research institution may want to accumulate large datasets of lifestyle data for clinical research, from patients willing to share their data.

Third, we are looking to support organisations that hold lifestyle data about their clients and want to offer access to this data to the individuals themselves. For example, a gym chain looking to share individual attendance data with the gym members, so that they can further share this data with a physiotherapist.

## 4 FHIR PROFILES FOR LIFESTYLE DATA

Our first contribution is a demonstration of how the FHIR standard may be used to describe lifestyle data in a provider agnostic fashion. There are many potential data sources for lifestyle data:

- self-recorded observations
- observations captured by providers of a certain service: health assessments, personal coach
- smartphones, wearables and body sensors
- smart home devices: smart meters, environment sensors, smart speakers with voice-enabled digital assistants

In order to capture this variety of data we needed to identify a set of profiles capable of representing it. Following FHIR best practices (Furore, 2017), we chose to extend the Observation profile, prioritizing removing fields irrelevant for our context of use, and only introducing new fields when no existing fields could possibly represent the data. The constructed observation profile is shown in Figure 1. This could be used for example to create a recording of self reported pain intensity, in reference to our arthritis case study.

| Name | Flags | Card. | Type |
|---|---|---|---|
| Observation | | 0..* | Observation |
| identifier | Σ | 0..1 | Identifier |
| status | Σ ?! | 1..1 | code Binding |
| code | Σ | 1..1 | CodeableConcept |
| subject | Σ | 1..1 | Reference(Patient \| Device) |
| effective[x] | Σ | 0..1 | |
| issued | Σ | 1..1 | instant |
| performer | Σ | 1..1 | Reference(Organization \| Patient) |
| reliability | | 0..1 | Extension |
| sensitivity | | 0..1 | Extension |
| component | Σ | 0..* | BackboneElement |
| code | Σ | 1..1 | CodeableConcept |
| value[x] | Σ | 0..1 | |
| interpretation | | 0..1 | CodeableConcept Binding |
| referenceRange | | 0..* | see (referenceRange) |

Figure 1: FHIR Observation Profile.

The **code** field is used to specify the type of observation (in our case, a pain intensity recording). The **subject** field is a reference to the person the observation concerns. The **performer** field is a reference to the source system (e.g. device, organisation or person). To record the measurement result, we use a set of related fields, grouped under **Component**. These include the actual value, but also the admissible range and the method of measurement. Finally, we record both the time period described by this observation - the **effective** time (in our case, when did the pain occur) and also the date the FHIR observation has been created - the **issued** date.

The fields discussed so far are applicable to any type of observation. There are some additional considerations that are specific to the lifestyle domain.

First, we need to account for the varying degree of reliability for the data sources: looking at a simple measure such as weight measured in kilograms - this could be self-reported, retrieved from a smart scale, or measured during a health assessment. To capture this 'trustworthiness' of the data, we introduce a field for **reliability**, that takes numerical values. Each organisation using these FHIR profile may specify rules that assign reliability scores to source systems. Every FHIR observation created based on data coming from a given source would then carry the same reliability value as that assigned to the source system.

Second, the same event (measurement) may be captured in more than one system. The same physical workout (e.g. a pilates class) could be captured by a consumer wearable device and by the class booking system of the gym. However, the event would be described differently in the two platforms. Our design approach is to avoid any reconciliation or aggregation at this stage. We recommend creating two distinct FHIR observations. The relation between these observations (duplicate data, contradictory data, complementary data) shall be established if and when necessary to answer a specific query.

Third, depending on the source system, data may be more or less sensitive, so different access rules would apply to data. To enable differentiated access rules for different types of observation, we add to the profile a field to record how **sensitive** the data is.

In reference to our arthritis case study, the other type of lifestyle information that should be saved relates to physical workout sessions. To document an exercise session, we propose using a **Diagnostic Report** profile, shown in Figure 2. This is another type of FHIR resource that groups observations related to the same episode. First, our profile on Diagnostic Report specifies the same fields we have discussed for individual observations: a code, subject, performer, time. Second, the **result** field contains reference to additional observations. The list we provide is not exhaustive. Depending on what information is available from the source system, implementers may add additional references to observations, so that all available information is captured.

Many lifestyle devices also make the distinction between aggregate and trace data - running watches are able to record a user's heart rate every few seconds during a session (trace), and also calculate an average heart rate at the end (aggregate). To reflect this, where applicable, fields for both types of data have been added to the profile - heartRate and averageHeartRate.

| Name | Flags | Card. | Type |
|---|---|---|---|
| DiagnosticReport | | 0..* | DiagnosticReport |
| identifier | Σ | 0..* | Identifier |
| status | Σ ?! | 1..1 | code Binding |
| category | Σ | 0..1 | CodeableConcept |
| code | Σ | 1..1 | CodeableConcept Binding |
| subject | Σ | 1..1 | Reference(Patient) |
| effective[x] | Σ | 0..1 | |
| effectivedateTime | | | dateTime |
| effectivePeriod | | | Period |
| issued | Σ | 1..1 | instant |
| performer | Σ | 1..1 | BackboneElement |
| role | Σ | 0..1 | CodeableConcept |
| actor | Σ | 1..1 | Reference(Organization \| Observation) |
| *result* | | *0..** | *Reference(Observation)* |
| exerciseDuration | | 0..* | Reference(Observation) |
| exerciseDistance | | 0..* | Reference(Observation) |
| caloriesBurned | | 0..* | Reference(Observation) |
| location | | 0..* | Reference(Observation) |
| altitude | | 0..* | Reference(Observation) |
| heartRate | | 0..* | Reference(Observation) |
| speed | | 0..* | Reference(Observation) |
| cadance | | 0..* | Reference(Observation) |
| averageHeartRate | | 0..* | Reference(Observation) |
| averageSpeed | | 0..* | Reference(Observation) |

Figure 2: FHIR FLI Profile - Workout Session.

# 5 REFERENCE ARCHITECTURE

A range of EHR platforms already use the FHIR standard as a data exchange mechanism. However, FHIR messages are mapped to an internal clinical model (for example, OpenEHR) and then saved in a structured database. Data can be interrogated either through structured query languages on the internal clinical model.

For example, SMART ON FHIR (Mandel et al., 2016) is an EHR platform built to support importing and exporting healthcare data through FHIR compliant messages - however, all data is stored internally as a set of key value pairs in a relational database.

In comparison, we aim to investigate whether an infrastructure based solely on FHIR messages can be used for reporting and analytics. By designing and implementing such a platform, we aim to evaluate the role of the FHIR standard when it comes to storing data, as opposed to transmitting data (as it has been traditionally seen).

In this section, we discuss the components of a reference architecture built around the FHIR profiles for lifestyle data.

First, the platform should have a *Connector* component to allow consumers to retrieve lifestyle data from different sources. The connector is required because lifestyle data is usually spread across a multitude of systems and platforms, each offering an incomplete view of the individual. We aim to offer consumers an easy way to retrieve their own data from these various platforms. The connector comes with built-in integrations to the main consumer-oriented lifestyle tracking platforms. The connector can also be extended by each organisation implementing the reference architecture so that it also retrieves data from company-specific internal databases or systems. The connector is also responsible to retrieve and send data in the FHIR format. This ensures interoperability between different organisations that have committed to exposing data in FHIR format.

Second, the architecture requires a *Converter* component, which maps data from proprietary formats onto FHIR profiles for lifestyle data. The converter is also responsible to record the source of the information for each FHIR observation created. It also enacts the business rules, assigning values to the reliability and sensitivity fields of the FHIR profiles.

The converter component relies on a library of selected FHIR profiles suitable to capture common data items in the lifestyle domain. Organisations implementing the infrastructure for internal use could develop additional FHIR profiles for data they hold internally. We propose a development process whereby whenever a new FHIR profile for lifestyle data is created by organisations implementing this architecture, the converter component is updated accordingly.

Third, the reference architecture requires a document database, to store each FHIR observation in its native format - a JSON message.

Fourth, the reference architecture includes a data visualisation component, which comes with built-in visualisations for common lifestyle data. We propose a process whereby whenever a new FHIR profile is created, the data visualisation component is extended so that it can graphically represent the new type of data. Organisations implementing the architecture may develop custom data visualisations to suit their specific needs. One clear benefit of building visualisations on top of a standard data model such as FHIR is that visualisation over the data will continue to function as expected, even if the systems generating the data change.

Fifth, the reference architecture allows for an advanced analytics component. This enables machine learning experts to build predictive models against a standard representation of data. These models can be evaluated and used independently of the underlying data source. The type of models particularly suitable to lifestyle data is based on time series analysis of behavioural traces, in order to predict trends in behaviours of interest or occurrence of specific events.

# 6 USER EXPERIENCE

In this section, we discuss new ways for customers and organisations to interact with lifestyle records. Whereas health records are accessed by clinicians in controlled environments (a GP clinic, a hospital etc), interaction with lifestyle records are more diverse, in terms of users, context and purpose.

One such interaction model is that of voice-enabled services, which, aside from increasing convenience, could remove a host of accessibility barriers from users.

A personal lifestyle record could allow consumers to create self-reported observations (for example, related to pain levels) using voice-enabled digital assistants. Through this, elderly people facing accessibility problems when a using a smartphone can more easily manage a chronic condition using voice to track their symptoms. Conversely, when developing these interfaces, it is important to ensure that they respond logically as otherwise could lead to frustration.

Similarly, lifestyle records can also by queried by voice, asking for example for the next physiotherapy session booked at the gym.

Due to the nature of the interfaces provided by voice enabled services, the types of data that can be recorded must be predefined. For example, the Amazon Alexa interface does not allow users to submit voice recordings as inputs, hence it would not be possible to record qualitative data. However most lifestyle data can be recorded quantitatively. For those that can't, quantitative mappings can be defined (a numerical scale in the case of pain) mitigating this issue.

# 7 FHIR FLI - OPEN SOURCE IMPLEMENTATION

Based on the reference architecture, we have started development of an open-source implementation - FHIR FLI. This acts as a starting point for organisations wishing to deploy a software solution that conforms to the proposed reference architecture.

FHIR FLI contains a number of FHIR profiles specifically developed for the lifestyle domain. For example, a profile to standardise sleep information, or for describing a workout session. The full list is available at https://simplifier.net/FhirFli/ resources

FHIR FLI offers out of the box connectors that allow consumers to authenticate and pull lifestyle data from Fitbit, Google Fit and Apple Healthkit. Data is then converted into the new FHIR profiles and saved into a document database.

Finally, FHIR FLI contains a data visualisation component and an Amazon Alexa skill under current development.

The FHIR FLI data visualization component provides an interface for both corporate and individual users to view and analyse FHIR data. Using the dashboard, individual users can view their own personal data, whereas corporate users can view aggregate anonymised data for all users registered with their companies. The dashboard draws its data from a document database, but also reroutes requests through an Apache Drill connector to allow for larger scale queries (in the case of corporate users).

The FHIR FLI Amazon Alexa skill provides users a fluid interface for making specific predefined queries on FHIR data stored within a document database - queries such as "What was my average heart rate yesterday?" will retrieve and return any matching data through the Amazon Alexa device. The skill also allows users to record data into a document database through an Amazon Alexa, and once a user has submitted the data, it is converted into the new FHIR profiles and stored into the document database.

The source code is available at https://github.com/fhirfli

# 8 CONCLUSIONS

FHIR FLI is an open source Personal Lifestyle Record that helps both consumers and organisations to combine and analyse lifestyle data, independently of the source systems.

Using FHIR ensures interoperability, allowing consumers to share data with service providers or research institutes and empowering organisations to share more data with customers.

The next step we are pursuing is assessing the feasibility and performance of a FHIR based data repository compared with data repositories based on clinical models such as OpenEHR.

Future work should also concentrate on the implementation of access control policies and development of the machine learning module. This would demonstrate how large data sets of lifestyle data can be assembled from individuals willing to anonymously share their data with research institutes, to build new predictive models linking lifestyle behaviours to clinical risks and outcomes.

# ACKNOWLEDGEMENTS

We are indebted to Dr. Dean Mohamedally, our project supervisor, for guiding our research.



# REFERENCES

Arthritis Research UK (2017). State of musculoskeletal health 2017. https://www.arthritisresearchuk.org/arthritis-information/data-and-statistics/state-of-musculoskeletal-health.aspx.

Bender, D. and Sartipi, K. (2013). Hl7 fhir: An agile and restful approach to healthcare information exchange. In *Computer-Based Medical Systems (CBMS), 2013 IEEE 26th International Symposium on*, pages 326–331. IEEE.

Conaghan, P. G., Dickson, J., and Grant, R. L. (2008). Care and management of osteoarthritis in adults: summary of nice guidance. *Bmj*, 336(7642):502–503.

Deighton, C., OMahony, R., Tosh, J., Turner, C., Rudolf, M., Group, G. D., et al. (2009). Management of rheumatoid arthritis: summary of nice guidance. *BMJ*, 338(mar16_1):b702–b702.

Downie, W., Leatham, P., Rhind, V., Wright, V., Branco, J., and Anderson, J. (1978). Studies with pain rating scales. *Annals of the rheumatic diseases*, 37(4):378–381.

Farrar, J. T., Young, J. P., LaMoreaux, L., Werth, J. L., and Poole, R. M. (2001). Clinical importance of changes in chronic pain intensity measured on an 11-point numerical pain rating scale. *Pain*, 94(2):149–158.

Furore (2017). Fhir best practices. https://simplifier.net/guide/ProfilingAcademy/Best-practices.

Hawker, G. A., Mian, S., Kendzerska, T., and French, M. (2011). Measures of adult pain: Visual analog scale for pain (vas pain), numeric rating scale for pain (nrs pain), mcgill pain questionnaire (mpq), short-form mcgill pain questionnaire (sf-mpq), chronic pain grade scale (cpgs), short form-36 bodily pain scale (sf-36 bps), and measure of intermittent and constant osteoarthritis pain (icoap). *Arthritis care & research*, 63(S11).

Kanta (2017). Phr for developers. https://www.hl7.org/documentcenter/public_temp_B8B93A79-1C23-BA17-0CD130736D1F56A7/wg/mobile/Kanta%20PHR%20for%20Developers.pdf.

Kavanagh, R. (2017). How is nhs digital using fhir. http://www.hl7.org.uk/doc_store/FHIR/NHSDigitalFHIR.pdf.

Lean, M. E., Leslie, W. S., Barnes, A. C., Brosnahan, N., Thom, G., McCombie, L., Peters, C., Zhyzhneuskaya, S., Al-Mrabeh, A., Hollingsworth, K. G., Rodrigues, A. M., Rehackova, L., Adamson, A. J., Sniehotta, F. F., Mathers, J. C., Ross, H. M., McIlvenna, Y., Stefanetti, R., Trenell, M., Welsh, P., Kean, S., Ford, I., McConnachie, A., Sattar, N., and Taylor, R. (2017). Primary care-led weight management for remission of type 2 diabetes (direct): an open-label, cluster-randomised trial. *The Lancet*.

Mandel, J. C., Kreda, D. A., Mandl, K. D., Kohane, I. S., and Ramoni, R. B. (2016). Smart on fhir: a standards-based, interoperable apps platform for electronic health records. *Journal of the American Medical Informatics Association*, 23(5):899–908.

Public Health England (2015). England has the potential to have the lowest disease burden in the world. https://www.gov.uk/government/news/england-has-the-potential-to-have-the-lowest-disease-burden-in-the-world.