

Information Models for the UCL Federated Health Record Server

Kalra D¹, Lloyd, D¹, Austin¹, A, O'Connor¹, A, Patterson, D², Ingram, D¹

¹ University College London, CHIME, Holborn Union Building, Highgate Hill, London, N19 3UA

² The Whittington Hospital, Highgate Hill, London N19 5NF

Correspondence to: d.kalra@chime.ucl.ac.uk

1. Introduction

This paper complements recent publication describing the overall design of the UCL Federated Health Record Server, by describing the information models adopted for the FHR Reference Model, the Archetype Object Dictionary and the representation of persons (patients and users). It should be regarded as a technical Annex to the paper “UCL-TEHRE01-Reprint”, available from the authors.

2. FHR Reference Model

The UCL Federated Health Record Reference Model (FHR-RM) defines a set of classes and attributes that represent the clinical context and medico-legal status of health record entries as a hierarchical set of Record Components. The goal for this model, in contrast to the Archetype Object Dictionary, is to represent the generic and domain-independent characteristics of Record Components.

The UCL FHR-RM is drawn below showing its class inheritance hierarchy (in red), and its aggregation (containment) hierarchy. The diagram conventions are based on the UML notation. The attributes have been omitted from the overall diagram below, and are defined later in this section.

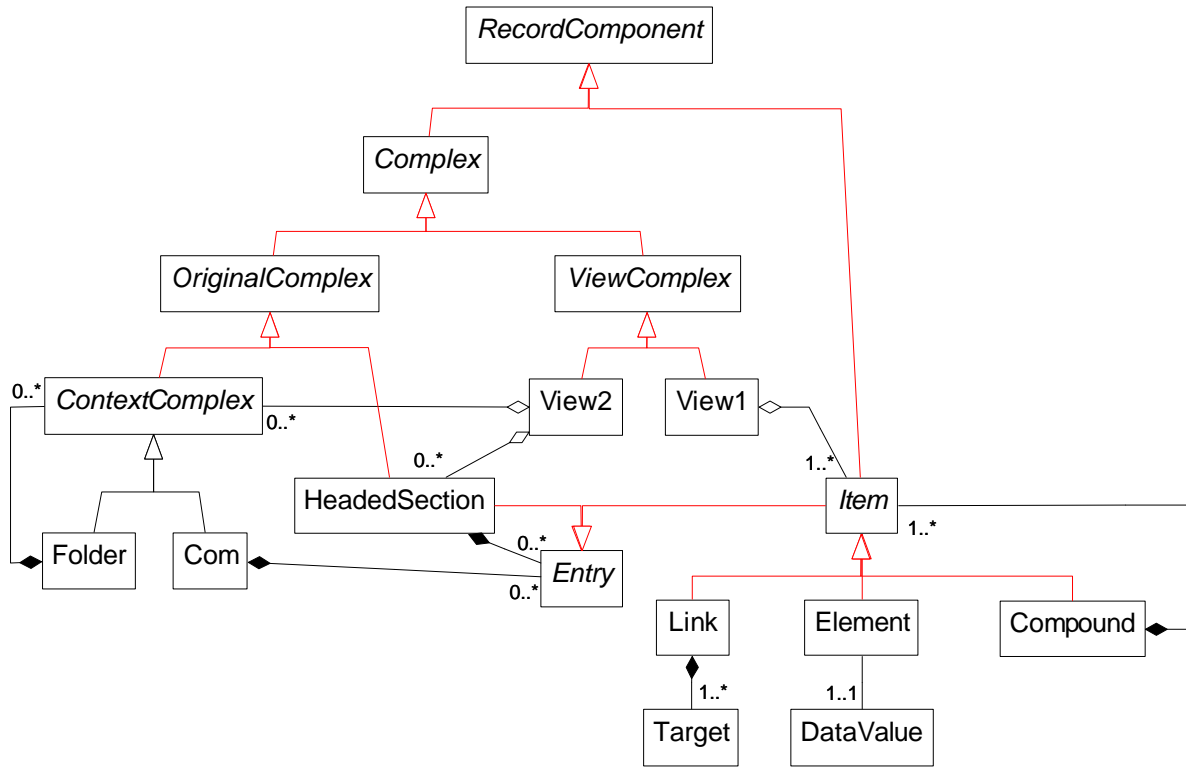


Figure 1: Class Inheritance and Aggregation within the FHR-RM

2.1 Description of the principal FHR-RM Classes

RecordComponent

RecordComponent is the abstract base class for Complex and Item. It defines the common attributes applicable to all of the major classes of the FHR-RM for:

- Record authorship, ownership and duty of care responsibilities
- Subject of care
- Dates and times of health care actions and of their recording
- Version control
- Access rights
- Emphasis and presentation

The complete set of attributes and their data types is presented later in this section. The FHR-RM distinguishes between the aggregation necessary to convey compound clinical concepts and the aggregation within a record that provides a way of grouping observations that relate to the health care activities performed. An example of the former would be *blood pressure*, which is a compound concept composed of *systolic* and *diastolic* values. An example of the latter would be the grouping together of observations under a general heading of *Physical Examination*.

The Complex and Item constructs respectively represent these two broad categories of aggregation.

Complex

In the FHR-RM, Complex is the common abstract super-class for the grouping of observations that relate to the health care activities performed. Two broad categories of Complex are reflected in the FHR-RM through two abstract sub-classes.

1. *OriginalComplex*: this set of classes represents the original organisational structure (grouping) of sets of record entries, as defined by the author(s) of those entries; it provides the medico-legal representation of the underlying information.
2. *ViewComplex*: this set of classes provide the means by which alternative groupings and sub-sets of the original information may be organised and preserved as permanent views in a patient's record, unlike those generic views provided in an ad hoc way by a client system.

OriginalComplex

Three concrete classes of *OriginalComplex* are defined in the FHR-RM, to provide for the nested aggregation of original groupings for record entries.

Folder

Folders define the highest-levels of organisation within health records. They will often be used to group large sets of record entries within departments or sites, over periods of time, or to demarcate a prolonged illness and its treatment. Examples of Folders include an episode of care, an inpatient stay, or one stage of a disease process. Folders can contain other Folders, and/or Coms.

Com

A medico-legal set of record entries required by the author to be kept together (to preserve meaning) when information is physically moved or copied to another persistent store. This is to ensure that all persistent EHR stores comprise whole Coms. This explicitly includes caches and cache mechanisms. The Com also defines the medico-legal cohort for the inclusion of new entries within an EHR: any new EHR entry (even if stored on a local feeder) must be a whole Com. Coms cannot contain other Coms or Folders. Examples include:

- the data entered at one date and time by one author (similar to a GEHR Transaction);
- the information gathered through the use of a protocol or template;
- a serialised set of readings taken over time but contributing to one examination;
- the definition of structures corresponding to electronic documents.

HeadedSection

This class is intended for grouping observations under headings *within* a Com. It therefore provides for the fine granularity grouping and labelling of record entries with names that relate the clinical concepts to the health care activities and processes surrounding the patient. Examples of HeadedSection names include presenting history, symptoms, investigations, treatment, drug prescription, needs, or plan. HeadedSections may contain other HeadedSections and/or Items. They cannot contain Coms or Folders.

ViewComplex

Two concrete classes of ViewComplex are defined in the FHR-RM, to provide for two differing mechanisms by which views may be generated.

View1

The View1 provides a means for grouping entries within Coms, at a similar hierarchical level in a record to the HeadedSection. However, the data within a View1 is derived through the use of a predefined query procedure i.e. a View1 comprises a query that generates a set of entries dynamically at the time of a client request. The mechanism by which search criteria can be defined in a generic, durable and portable manner within the View1 class is presently being developed.

View2

The View2 provides a static view of original information, through a set of references to the original entries or to groups of entries (i.e. Items, HeadedSections and/or Coms). It therefore provides a mechanism by which information within one Com may logically appear inside another Com, since the originals of these cannot be nested. This class cannot include object references to other instances of View2, to avoid recursive loops of such references.

Item

This abstract class provides an aggregation construct for clinical concepts that are composed of one or more individual named clinical values (e.g. *pulse*, *blood pressure*, *drug dose*, *heart sounds*). These entries may be aggregated within a hierarchy to represent complex clinical concepts, but such a composition is distinct from the record structure grouping hierarchy provided by the Complex classes. This class also provides a means by which point-to-point linkage or linkage nets within a single FHR can be represented. The Item class hierarchy is described later in this section.

2.2 The Attributes of the RecordComponent Class

The tables below list the attributes of the RecordComponent class. These are inherited throughout the FHR-RM class hierarchy and may acquire instance values at any level of a hierarchy of record entries. Some of these attributes have been defined as mandatory, and must be incorporated within any FHCR in order to comply with this specification. If mandatory information is not present in the underlying feeder system data then a null attribute value must be included within the Record Component object. Other attributes, marked as optional, have been included to meet published requirements or on the basis of implementation and deployment experience. The attribute data types are all of a base type; complex attribute data types have deliberately been avoided to ease implementation and the processability of federated records. The cardinality of all Mandatory attributes is 1, and that of Optional attributes is 0 or 1.

Subject of care

RecordComponent attribute	Mandatory Optional	Description of intended use	Type
SynPatUID	Mandatory	This is the "Subject of Care" attribute and will identify the patient about whom the record component provides information.	STRING
SubjectOf Information	Optional	This will identify the person about whom the information in a record component relates if not the subject of care e.g. if the information is about a family member, such as the patient's father or mother. PERMITTED VALUES: { patient, relative, foetus, mother, donor, personalcontact, otherperson, device } DEFAULT = "patient".	STRING

Note: the values for SubjectOfInformation are taken from ENV13606-2 (Domain Termlist)

Record authorship, ownership and duty of care responsibilities

FHR-RM attribute	Mandatory Optional	Description of intended use	Type
RecordingHealth CareAgent	Mandatory	The healthcare agent responsible for physically including this record component into the patient's source record.	STRING
Responsible HealthCareAgent	Optional	The healthcare agent responsible for effecting the care and for authoring this record component.	STRING
LegallyResponsible HealthCareAgent	Mandatory	The healthcare agent with senior clinical responsibility for the patient at the point of care documented by this record component e.g. Consultant in charge.	STRING
Information Provider	Optional	The person providing healthcare information if not the subject of care (e.g. a family member, friend, another clinician, an electronic device).	STRING

Note 1: information passed to the record server is deemed to be from authenticated sources. Digital signatures are not considered to be part of the FHR information model, but might be stored within an EHR server on an enterprise-specific basis.

Note 2: although countersignature is sometimes required for health record entries, these are usually handled at an application level and do not necessarily form part of the FHR. In cases where more than one actioning healthcare agent needs to be recorded the UCL team have so far proposed that two from the available set of healthcare agent attributes above should be used, such as the RecordingHealthcareAgent and the ResponsibleHealthcareAgent.

Dates, times, locations of health care actions and of their recording

FHR-RM attribute	Mandatory Optional	Description of intended use	Type
RecordingDateTime	Mandatory	The date and time this record component was included in the patient's source record (NOT the date and time it was brought into the federation).	DATETIME
HealthcareActivityBegin Time	Optional	The date and time of the health care activity to which this recording relates (this may differ from the RecordingDateTime if a delay occurred before a record could be authored e.g. a home visit at night).	DATETIME
HealthcareActivityEnd Time	Optional		DATETIME
ObservationBeginTime	Optional	The date and time (or intervals) of any health or care acts which occurred in the past but are being recorded at the present e.g. an operation performed several years ago.	DATETIME
ObservationEndTime	Optional		DATETIME

HealthcareActivity Location	Optional	The enterprise, department or other location at which the patient is receiving the care documented in this entry (for audit, management, financial or access rights purposes).	STRING
AcquisitionTimeDate	Optional	The date/time at which this Record Component was added to a Federated Record if its origin was elsewhere e.g. if received as a message from another record system; this attribute is necessary because the RecordingDateTime would represent when the original entry was recorded, not when it was received into the federated health record.	DATETIME
Locale	Optional	To document the time zone and geographical location of the recording clinical system, for example permitting international interpretation of other dates and times recorded.	STRING

Note 1 : the UCL implementation of Healthcare Activity and Observation attributes (using the Java Calendar class) permits the recording of begin or end times to be specified to an arbitrary granularity, permitting an author, for example, to record that observation occurred between 1960 and 1965.

Version control

FHR-RM attribute	Mandatory Optional	Description of intended use	Type
RevisedVersion	Optional	A reference to the version of this Record Component that replaces this version, if it has been revised (referenced via its RC_UID).	STRING
RevisedBy	Optional	A backward reference to the Record Component that this version has replaced, if it has been revised (referenced via its RC_UID).	STRING
Authorisation Status	Mandatory	PERMITTED VALUES: {unattested, attested, obsolete, revision}.	STRING

Access rights

FHR-RM attribute	Mandatory Optional	Description of intended use	TYPE
AccessAmend Rights	Mandatory	PERMITTED VALUES: {admin, audit, clinical, team, profession, hcp} This set of values reflects an ordered set of sensitivity levels. The anticipated default in most EHR systems will be "clinical" i.e. the record component is accessible to all staff involved in the clinical care of the patient. This attribute is used to differentiate sensitivity levels <i>within</i> a single FHR, and are supplementary to any restrictions on overall access to each patient's FHR as a whole.	INTEGER

Note: this attribute permits a sensitivity level to be assigned to Record Components at any level of granularity, as part of a broader approach to access control summarised later in this paper.

Emphasis and presentation

FHR-RM attribute	Mandatory Optional	Description of intended use	Type
Emphasis	Optional	At present this attribute is limited to a Boolean. If set to true the information in this record component was emphasised by the original author.	INTEGER

Note: there is some debate about the importance of representing more detailed aspects of presentation within the FHR. The view taken by the authors is that the specification of presentation characteristics is not necessary nor feasible for all entry instances within the records of individual patients. Where enterprises wish to retain a medico-legal reference to information display characteristics used for a given time period by certain applications, for example through a pointer to an XML Stylesheet, these ought to be retained by each enterprise or by the developers of clinical applications.

Class identifiers

FHR-RM attribute	Mandatory Optional	Description of intended use	Type
Name	Mandatory	This attribute preserves the actual name of the record component used in the original source record; this may be identical to the corresponding Archetype name, but might not be in the case of synonyms.	STRING
RC_UID	Mandatory	An internal reference identifier for each record component, provided by the FHR server.	STRING
SynObjectUID	Mandatory	The unique identifier of the Archetype that provides the template for this set of record components (Note: the Name attribute may not always be identical to the Archetype name).	STRING
ParentRC	Optional	The primary information context, i.e. it is a reference to the record component at the next higher level in a record structure.	STRING
EHCRSource	Optional	The unique identifier of the feeder system contributing this record component to the federated health record; this is important for medico-legal reasons, including the ability to link all parts of the FHR to relevant Data Controllers.	STRING

Other Attributes

FHR-RM attribute	Mandatory Optional	Description of intended use	Type
AuthorsComment	Optional	A free-text comment associated with the record component as a whole (not primarily with its value), intended for use by the author; it might be used by a revisor to explain the rationale for the revision.	STRING
RcuLink	Optional	The RC_UID(s) of other record component(s) in the FHR linked by the author (e.g. to relate an allergic rash to a previous drug prescription). Note: these other components must already be in the record, and therefore the references will be to past or accompanying present entries.	STRING
RcuLinkBack	Optional	This reference represents the reciprocal of the above link, from an historic target record component to the source: it will therefore point forwards in time. Some EHR systems may not permit the retrospective editing of record components to insert this attribute.	STRING

Note: The RcuLink and RcuLinkBack attributes have been implemented using the Java Vector class to permit multiple targets to be specified. The RCU link attributes overlap in function with the Link class described below. This is deliberate to reflect the varying way in which internal links are represented by different feeder systems at present.

Item

The Item abstract class hierarchy provides a means to represent compound and elemental clinical concepts, using the concrete classes Compound and Element respectively. A set of context description attributes is associated with the Item objects, which are largely derived from the CEN EHR Domain Termlist standard ENV 13606-2. Other attributes such as Justification and ProtocolRef permit both a human and a software reference to the rationale behind a clinical entry, including the specification of a protocol or step in a protocol that was used during that part of clinical care. The Item class also inherits the attributes defined in the RecordComponent class, with the option to override the value of any of these at a local level.

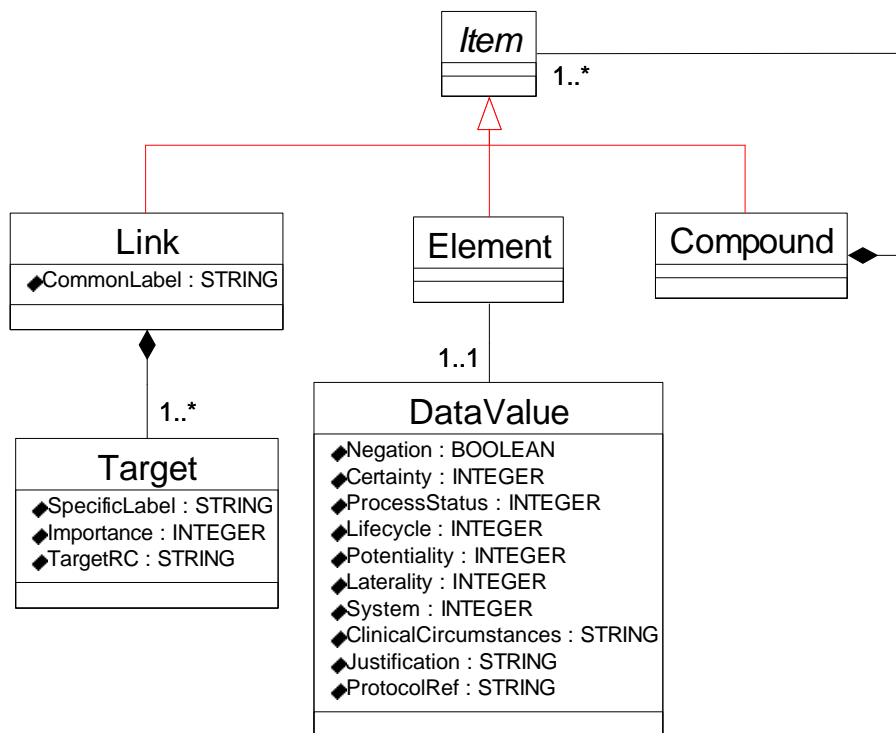


Figure 2: Item Class Hierarchy

An important aspect of the FHR-RM, including the Element, is the binding of a Name attribute (acting as a label) to each content value, providing the individual quantities, dates, images or clinical terms with a primary context in any given record entry.

The Compound class provides an aggregation construct for clinical concepts that are composed of one or more individual named clinical values (e.g. *pulse, blood pressure, drug dose, heart sounds*). These entries may be aggregated within a hierarchy to represent complex clinical concepts, but such a composition is distinct from the record structure grouping hierarchy provided by the RecordItemComplex classes such as the HeadedSection.

An additional child class of RecordItem is Link. This class provides a means by which point-to-point linkage or linkage nets within a single EHR can be represented. From an aggregation perspective, Links behave as Elements: they are leaf nodes in an FHR object hierarchy.

Content Classes

The Element supports a range of data types for the DataValue that may be assigned to any element entry. These generic classes are a distillation of the original foundation work of GEHR, EHCR-SupA, and CEN/TC 251 ENV 13606.

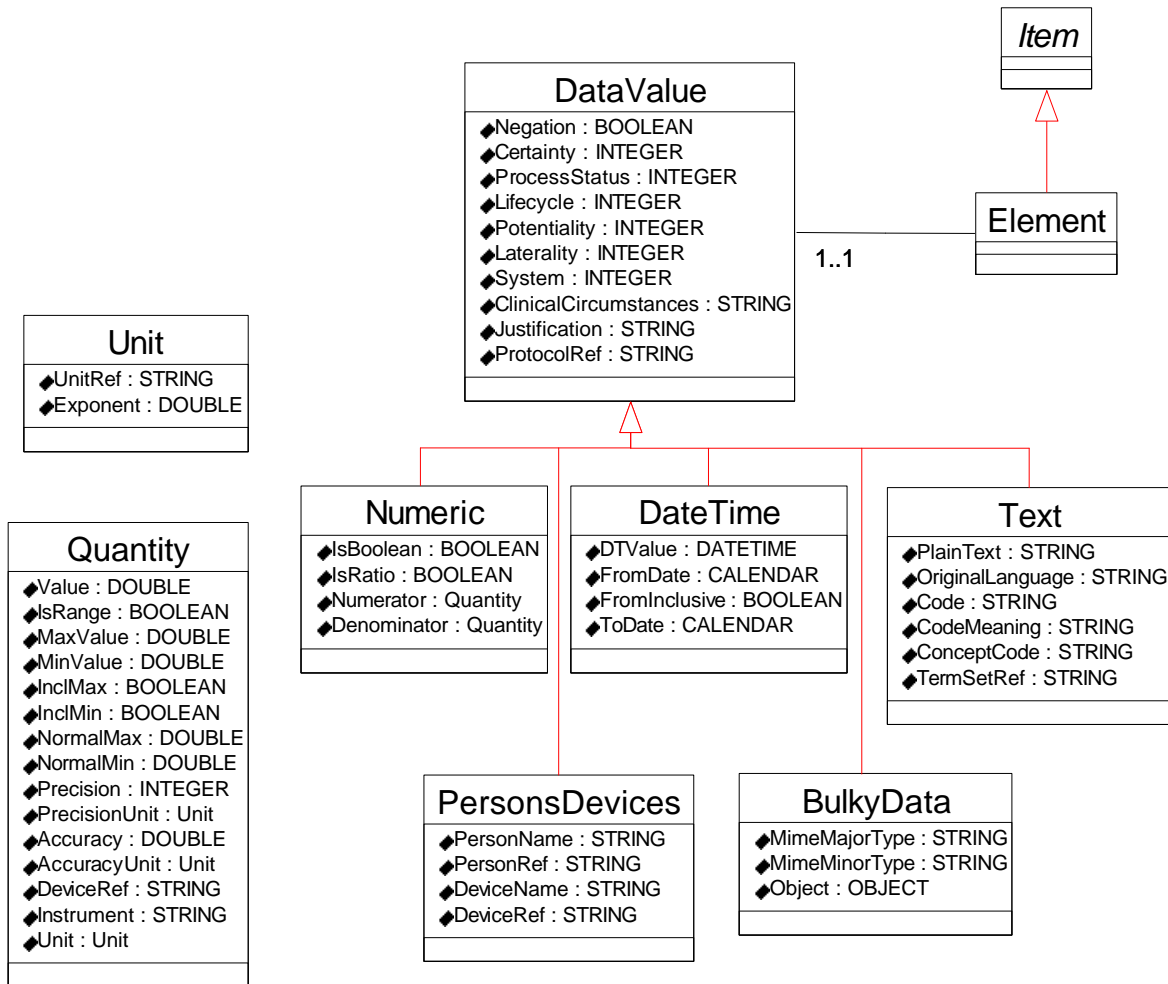


Figure 3: Object model of Element content

Separate dictionaries for units and for referencing terminology systems are under development. The model for persons and devices above will reference the richer information objects in the Persons Directory Service (see below), which will later also include a register of devices. The name strings are also included in the PersonsDevices class for medico-legal safety, to ensure that these attributes of a record component’s content can be interpreted even if that Directory Service is somehow unavailable.

It should be noted that ENV 13606-4 defines a set of specific content models for commonly used objects such as drug prescriptions. The UCL FHR-RM deliberately does not define specific record objects of this nature: they are instead capable of being defined in and implemented through the Archetype Object Dictionary. This approach attempts to separate the most stable aspects of a health record model (through the

FHR-RM) from those where local variation or evolution over time are most likely to occur (via the Archetype Object Dictionary).

3. FHR Archetype Object Dictionary

The classes and attributes of the Reference Model, described in the previous section, are deliberately defined at a high level of abstraction to provide an information model that can be applied to any potential health record entry. However, the individual feeder systems providing data to the FHR server are likely to be highly specific to the local requirements of individual sites, to specialities and to groups of professionals.

The Archetype Object Dictionary provides the formalism by which the specific clinical data sets and aggregates normally found in health records and in contemporary feeder systems can be defined. Archetype entries utilise the FHR-RM classes as basic building blocks, using the Name attribute of each class instance to generate specific clinical hierarchies that can be directly mapped to feeder system data schemata and can be the target of a client request.

The Archetypes can be mapped onto the data representations used in each individual feeder system through a set of access methods. These might be defined jointly by the developers of each feeder system and the developers of the FHR server at each installation, or might be derived from published interface specifications. The references to the access methods are logically integrated within the Archetype Object Dictionary during the “sign-up process” by which each feeder system is connected to the federation. In a “live” federation, a request by a client application or middleware service for a set of Record Components will result in the invocation of the relevant method(s) by the FHR service in order to retrieve the necessary health care record data from a feeder system.

Other features of the Archetype Object Dictionary are the mapping of Archetypes to clinical concept tags, and the inclusion of validation criteria that might be used to verify the instantiation of a Record Component’s candidate data value. These are shown diagrammatically below.

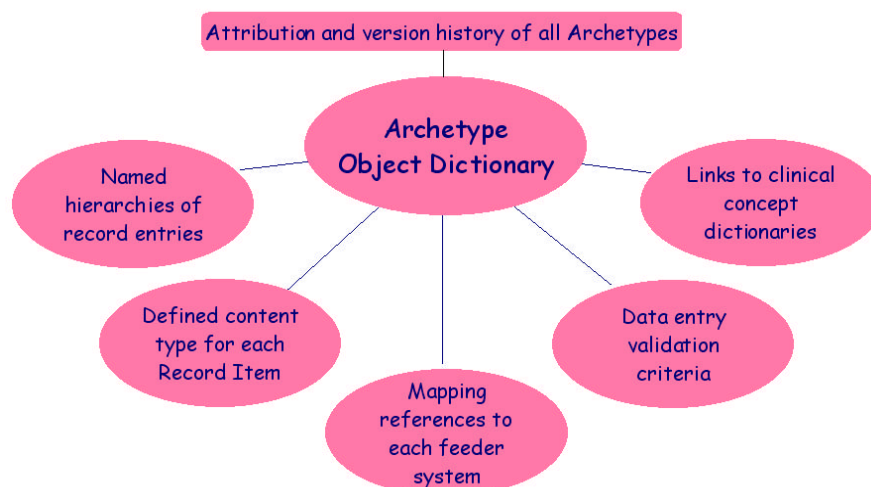


Figure 4: Functional sub-components of the Archetype Object Dictionary

The Archetype Object Dictionary Client component is described in section 3.2 of this paper

3.1 Object Model of the Archetype Object Dictionary

The formal object model of the Archetype Object Dictionary is closely related to the FHR Reference Model. It extends the RecordComponent class of the FHR-RM through the addition of one compound attribute that is used to represent the information about the creation, versioning and use of each library definition, and supports the mapping of that definition to a set of medical knowledge concept tags.

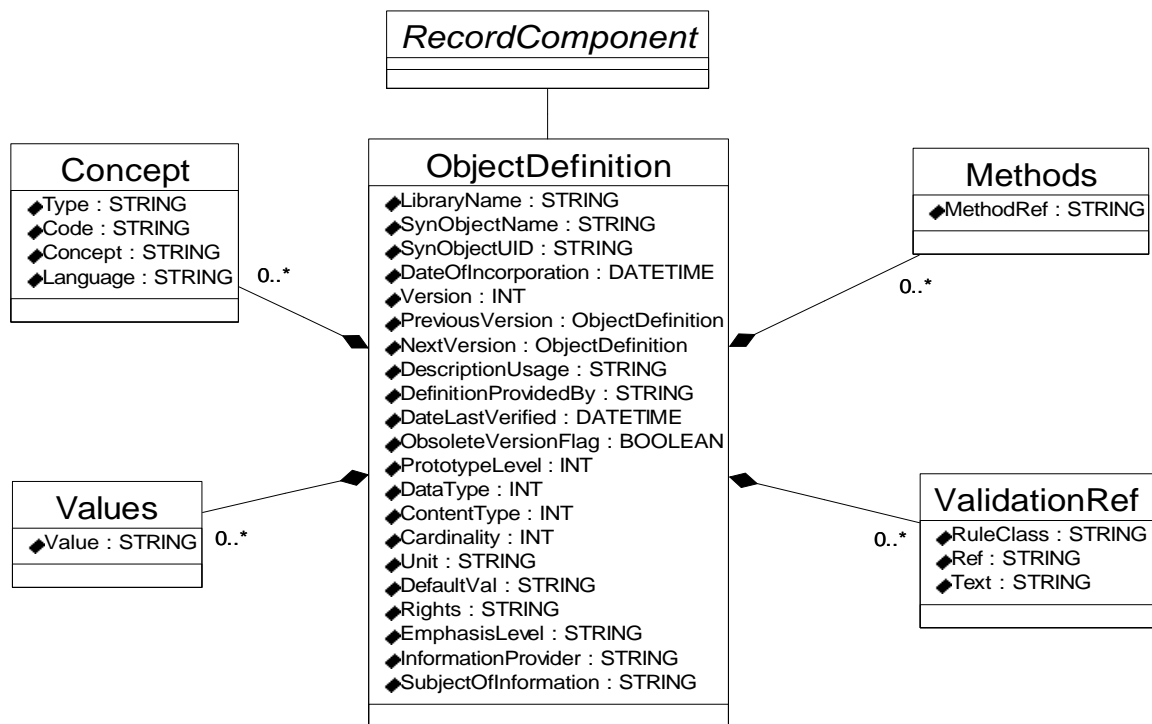


Figure 5: Information Model of the Archetype Object Dictionary

ObjectDefinition Class

The ObjectDefinition class contains the attributes relevant to managing the library entries associated with each Archetype. This includes the formal definition, author identification and version of any local or national standardised data sets within the Dictionary. In addition, some descriptive text (a definition or explanation) may be provided to clarify the intended clinical use of the object. It will also be necessary to store information about changes that occur to Archetypes over time; this might mean recording if this particular object is the current definition, and the identification of its predecessors and/or successors. The individual attributes of ObjectDefinition are described below.

ObjectDefinition attribute	Mandatory Optional	Description of intended use	Type
LibraryName	Mandatory	Archetypes are authored within libraries to permit traceability and the managed distribution of these within multi-agency domains.	STRING
SynObjectName	Mandatory	This is the standard preferred name by which the Archetype is known.	STRING
SynObjectUID	Mandatory	This UID is used to uniquely identify this Archetype within Record Components.	STRING
DateOfIncorporation	Mandatory	When the Archetype was authored in this Library.	DATE
Version	Mandatory	The version number.	INT
PreviousVersion	Optional	A reference to the previous version if this is a revision.	ObjectDefiniton
NextVersion	Optional	A reference to the successor version if this Archetype has been revised.	ObjectDefiniton
DescriptionUsage	Optional	A textual description of how this Archetype was intended to be used for record entries, intended as guidance for those mapping feeder systems or clinical applications.	STRING
DefinitionProvidedBy	Mandatory	The reference source guiding this Archetype definition, such as a clinical guideline.	STRING
DateLastVerified	Mandatory	When the reference source was last checked to confirm this Archetype is still valid.	DATETIME
ObsoleteVersionFlag	Optional	To permit Archetypes to be marked as obsolete even if a revision has not been authored.	STRING
PrototypeLevel	Mandatory	This attribute permits selective sharing of parts of an Archetype library to others. PERMITTED VALUES: {0-2} (2=PRIVATE, 1=PRIVATE_SHARABLE, 0=PUBLIC).	INT
DataType	Mandatory	The FHR-RM class to which this Archetype applies. Permitted Values: {0-7} (0=Folder, 1=Com, 2=HeadedSection, 3=Compound, 4=Element, 5=Link, 6=View1, 7=View2).	INT
ContentType	Mandatory	Specifying the Data Value type for Archetypes whose DataType is Element. Permitted Values: {0-5} (0=No_Content, 1=Text, 2=Numeric, 3=Date_Time, 4=Persons_Devices, 5=Bulky).	INT

Cardinality	Mandatory	Indicating the number of instances of this Archetype that may be created within any one instance of its parent e.g. 1 to many.	STRING
Unit	Optional	Specifying the unit of recording for Archetypes whose DataType is Element.	STRING
DefaultVal	Optional	Providing a default value on instantiation for Archetypes whose DataType is Element.	STRING
Rights	Optional	Permitted values for these Record Component attributes may be specified in the Archetype definition, for example in the case of a Family History Archetype to indicate that the SubjectOfInformation may not be the patient.	STRING
EmphasisLevel	Optional		STRING
InformationProvider	Optional		STRING
SubjectOfInformation	Optional		STRING

Values Class

This class permits the author of the Archetype to specify a fixed possible value list for Archetypes whose DataType is Element.

MethodRef Class

This class stores a set of method references that may be used to identify feeder system data relating to this Archetype.

Concept Class

This class enables a client application to reference an Archetype through the use of a locally-defined label, an abbreviated name or a language translation of it. It will also enable an application to identify the set of available objects that correspond to a clinical subject heading. This class is a place-holder for the methodology by which Archetype definitions can be appropriately linked to, for example, GALEN ontology or terminology services.

Concept attribute	Mandatory Optional	Description of intended use	Type
Type	Optional	The classification system or ontology from which the code has been derived.	STRING
Code	Optional	A code referencing the clinical concept within that classification system or ontology.	STRING
Concept	Optional	A rubric for that code, included for safety and to permit searches to utilise this class of information if that classification system or ontology is not available as a live look-up service.	STRING
Language	Optional	The natural language used for the rubric.	STRING

ValidationRef Class

This class, which is still undergoing evaluation, is a place-holder for the expression of rules regarding the validation of instance values for element objects, or the interdependence of values on other components of an Item or Complex. These rules would be used primarily during data entry rather than retrieval. For example, an entry value may be drawn from a pick-list or reference database (such as *drug name*), it may be subject to upper and lower limits (such as *height*), or its value may be restricted by other values in the record (such as the patient's age or gender).

This class contains a set of rules that must be evaluated against any candidate value for an Element conforming to this Archetype. A string text message can be returned to the clinical application if a condition is met. This provides a useful means of providing messages back to end users:

- if the value they have offered is not permitted;
- if they need to re-affirm the value (e.g. it is a rather unusual value, but not impossible);
- if the value is accepted but some further action advice needs to be communicated back to the user.

The three situations map to three sub-types of rule, reflected in three values for the RuleClass attribute: REJECT, CONFIRM, ACCEPT. If more than one rule has been defined for an Archetype, the provisional intention for the service implementing this class is to evaluate rules in the order:

1. REJECT
2. CONFIRM
3. ACCEPT

This class is a place-holder for the methodology by which Archetype definitions can be appropriately linked to electronic guidelines and to other decision support services.

ValidationRef attribute	Mandatory Optional	Description of intended use	Type
RuleClass	Optional	Action to be performed if the rule condition is met. PERMITTED VALUES: {0-2} (0=ACCEPT, 1= CONFIRM, 2= REJECT}	INT
Ref	Optional	The rule string to be evaluated against a candidate value for an Element of that Archetype.	STRING
Text	Optional	A string to be returned by the Federated Health Record server to the calling application if this rule is met.	STRING

3.2 Archetype Object Dictionary Client

The UCL Archetype Object Dictionary Client (ODC) component:

- provides an authoring tool for Archetypes in terms of their constituent compound clinical concepts;
- includes the formal definition, author identification and version of any local or national standardised data sets within the Dictionary;
- incorporates pointers to access methods which can extract data held on feeder systems to which the FHR services are connected;
- ensures adequate version control and maintenance procedures to accommodate revisions of Archetypes over time.

The Archetype Object Dictionary Client component has been written entirely using Java Foundation classes and Swing, allowing true cross-platform deployment. It utilises an object database PSE Pro, from Object Design Inc., which is also a Java application and is similarly capable of installation on any platform that supports a Java Virtual Machine. The licence for PSE Pro permits the distribution of run-time versions alongside the Archetype Object Dictionary application, removing the need to purchase any additional third-party software. The ODC permits the structure of the record object definitions to be captured in a way that the user originally intended for maximum performance and flexibility.

The core features of the ODC are listed below.

ODC Class Hierarchy
ODC Archetype Properties
Creating New Archetype Entries
Cardinality on Instantiation
Validation Criteria
Data Retrieval Methods
Copying and Pasting Archetypes in the Hierarchy
Publicising Archetypes
Deleting an Archetype
Marking an Archetype Obsolete
Revising an Archetype Definition
Reviewing the Version History
Tracking Archetypes having Multiple Parents
Exporting the Database
Saving the Database
Help about screen

Future work will enable synonyms for clinical object names to be identified and linked to preferred terms, and offer a multi-lingual set of clinical object names. Data entry validation criteria may also be incorporated, and their linkage to run-time protocol components is being explored.

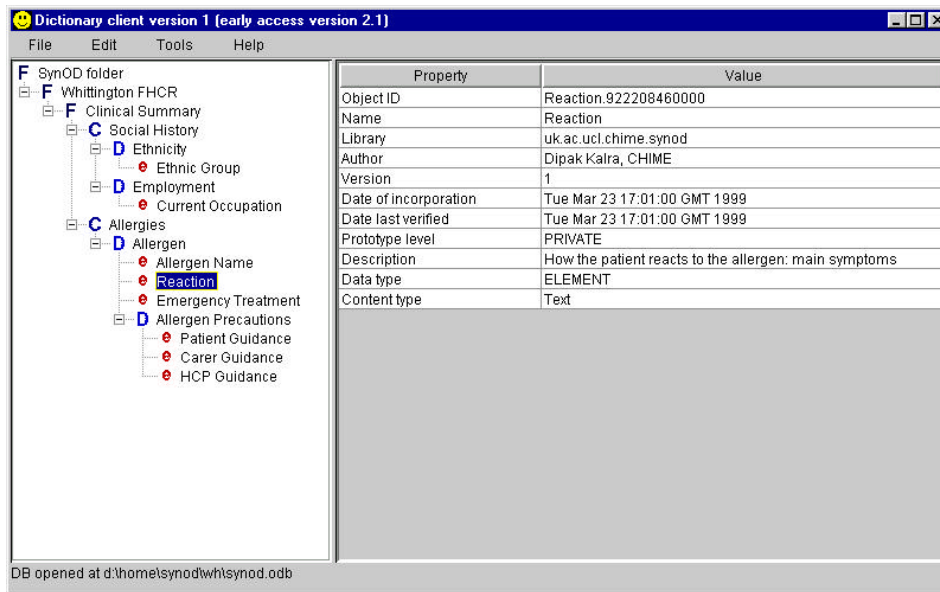


Figure 6: Example screen from the Archetype Object Dictionary Client

4. FHR Persons Directory Service

The Persons Directory Service is a component providing information on the identification of patients, healthcare professionals and other staff to the other FHR services. It provides a repository of person names and other demographic information, together with their access rights status, that can be used to identify persons within an FHR or to authenticate access rights to a given set of record components.

The Persons Directory provides a means of registering staff and patients within a consistent repository as part of the FHR. This model has been proposed, and implemented as the Persons Directory Service, in order to provide a means of searching for patients, confirming the correct patient has been chosen, and providing a basic demographic data-set as part of each patient's federated health record. In many situations where an FHR server is deployed there is likely also to be a regional or national directory of patients and also of healthcare agents, which would replace the service described here. The overall engineering approach to the FHR middleware would permit the replacement of the Persons Directory Service with a local alternative quite easily.

The information model builds on the early work of GEHR and Synapses, which has been refined by the EHCR-SupA project. The models proposed here by UCL are a simplified but consistent representation of the Healthcare Agent subsystem defined in CEN/TC 251 ENV 13606 (EHCR Communication). This model is deliberately not intended to mimic a full patient demographic server such as a hospital PAS.

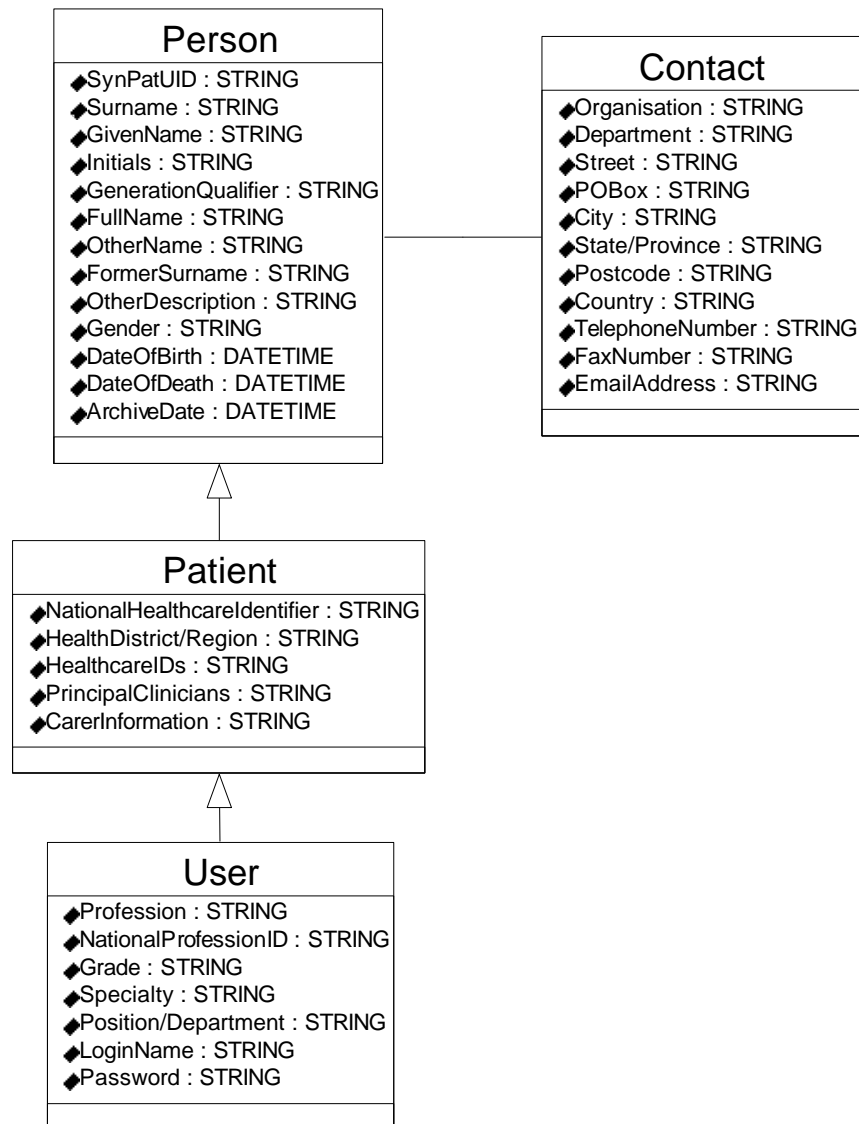


Figure 7: Information Model of the FHR Persons Directory

The data repository uses and extends Novell NDS objects and its metadirectory, and is accessed via Java Naming and Directory Interface (JNDI) APIs. This entails some configuring of the NDS tree and its class models to optimise it as an object repository for patient and staff identification. For deployment purposes, Novell eDirectory has been used as the product to provide and manage the NDS services.

A Software and Devices Directory is also being developed using NDS, and is intended to provide a registry of all electronic sources of FHR information (such as monitoring devices and decision support software) that might be referenced within a patient's record.

5 Access Control

A combination of internal services is used to deliver an overall access control framework governing access to FHR information, reflecting enterprise policies by:

- determining user profiles from available authentication and certification services;
- limiting patient searching within organisational contexts;
- limiting access to sub-categories of the record based on roles e.g. a department or speciality.

Specific structured parts of each patient's FHR can record individual patient consent to:

- map a user's role-based privilege to the sensitivity of individual record components;
- permit access to sub-categories of the record based on roles e.g. for research or teaching;
- exclude named persons from adopting certain roles for accessing individual patient records.

These services are in the process of being implemented and tested, and will be published later.