MEDINFO 2019: Health and Wellbeing e-Networks for All
L. Ohno-Machado and B. Séroussi (Eds.)

doi:10.3233/SHTI190188

Clinical Text Mining on FHIR

Philipp Daumke[a], Kai U. Heitmann[b,c], Simone Heckmann[b,c], Catalina Martínez-Costa[d], Stefan Schulz[a,d]

[a] Averbis GmbH, Freiburg, Germany, [b] Geyfra GmbH, Hürth, Germany, [c] HL7 Germany, Berlin, Germany
[d] Institute for Medical Informatics, Statistics and Documentation, Medical University of Graz, Austria

Abstract

Semantic standards and human language technologies are key enablers for semantic interoperability across heterogeneous document and data collections in clinical information systems. Data provenance is awarded increasing attention, and it is especially critical where clinical data are automatically extracted from original documents, e.g. by text mining. This paper demonstrates how the output of a commercial clinical text-mining tool can be harmonised with FHIR, the leading clinical information model standard. Character ranges that indicate the origin of an annotation and machine generates confidence values were identified as crucial elements of data provenance in order to enrich text-mining results. We have specified and requested necessary extensions to the FHIR standard and demonstrated how, as a result, important meta-data describing processes generating FHIR instances from clinical narratives can be embedded.

Keywords:

Electronic Health Records, Natural Language Processing, Semantics

Introduction

Semantic Interoperability was defined in 2000 as "…integrating resources that were developed using different vocabularies and different perspectives on the data". Semantic interoperability requires that "systems must be able to exchange data in such a way that the precise meaning of the data is readily accessible and the data itself can be translated by any system into a form that it understands" [1].

Nearly twenty years later, the lack of semantic interoperability continues being an obstacle to a more rational and effective data and information management in healthcare and biomedical research. The authors of [1] had already distinguished between vocabularies and perspectives, highlighting the division between ontology ("what there is") [2] and epistemology ("what we can know") [3]. On the level of current health informatics standards, this has driven the evolution of two genres of semantic resources:

- Terminology systems (vocabularies, thesauri, formal ontologies, classifications), which attach meaning to domain terms and elaborate on necessary and sufficient properties of (classes of) domain entities;
- Information models, which are artefacts that provide standardized structure (section, entry, grouping, etc.) and context (diagnosis, past medical history, medication order) for clinical recording scenarios.

In this paper, we will focus on the latter, particularly on HL7 FHIR [4], a standard for healthcare information exchange and sharing, characterized by its straightforward approach to implement interfaces between Electronic Health Record (EHR) data and data consuming applications.

FHIR is based on interoperable building blocks named *Resources*, small data model components defining sets of properties that describe and provide structure for domain data acquisition. Currently there are approximately 150 resources, uniquely identified with Uniform Resource Identifiers (URIs).

Examples are *MedicationRequest* (prescription), *AdverseEvent*, *Procedure* and *Condition* (problem). They constitute a graph of clinical data by explicit inter-resource references [5]. For instance, a *MedicationRequest* resource explicitly references its prescriber (a FHIR *Practitioner* resource), its patient (*Patient* resource), and the drug prescribed (*Medication* resource). A built-in extensibility mechanism can enrich existing resource definitions.

Semantic interoperability within FHIR is provided by explicit, detail-oriented, prescriptive guidance, with its interoperable meaning anchored in external terminology standards, e.g. SNOMED CT, LOINC, or ICD-10. FHIR provides a granular way to exchange data using a RESTful style approach [6]. Its focus is on providing models for frequently occurring documentation and information exchange tasks. FHIR resources can be serialized in JSON and XML.

Looking at EHR systems in use, there is still a persisting gap between structured and coded information on the one hand and a much larger amount of semi- or unstructured narratives on the other hand. It is unrealistic that this textual documentation will be largely substituted by structured documentation [7]. This is where natural language processing (NLP), in particular text mining, comes in.

This paper addresses the issue how the output of clinical text mining systems can be harmonized in FHIR. We identify gaps in the current FHIR specification related to content essential for text mining. We introduce and discuss specific FHIR extensions that allow text-mining results to be represented in FHIR without loss.

Materials and Methods

The Text Mining System *Health Discovery*

The clinical text mining technology under scrutiny is *Health Discovery* [8] by Averbis GmbH. *Health Discovery* contains over fifty different text-mining annotators for the recognition of diagnostic statements, medical procedures, lab values, drugs, anatomy, morphology, scores and others. It is available for several languages, including English and German. *Health Discovery* bundles annotators in predefined text mining pipelines to facilitate the analysis of text genres such as discharge summaries or pathology reports. It has been successfully used for various use cases, e.g., data driven patient recruitment for clinical trials [9], automated coding and billing [10], filling of tumour registries, rare disease identification [11], antibiotic resistance monitoring [12], radiology report analysis [13] and health data de-identification [14].

UIMA

Health Discovery is based on the *Unstructured Information Management Architecture* (UIMA) [15], a flexible and extensible text-mining framework for the analysis and processing of unstructured text. UIMA allows defining text-mining pipelines consisting of a set of annotators called analysis engines. The analysis engines communicate in a pipeline by adding or modifying meta-information stored in the Common Analysis System (CAS), which, in addition to the metadata, also contains the currently processed document.

In UIMA, all elements to be extracted from texts must be predefined. This is done in a so-called *Type System*, where types and their attributes (features) are defined. The most common type is *Annotation*. It comes with the attributes *begin* and *end*, which specify the text position to which an annotation relates. These attributes are inherited by all subtypes of *Annotation*.

UIMA is only a framework and does not ship a rich selection of components. There are, however, component repositories like DKPro [16], which provide analysis engines of well-known text mining components. Prominent applications in Healthcare built upon UIMA are the DeepQA system Watson [15], the clinical Text Analysis and Knowledge Extraction System (cTAKES) [17], the medication entity and attribute extraction and normalisation tool MedXN [18], and MedTime [19], a software to extract and normalize TIMEX3-based temporal expressions from clinical text.

Making a Text Mining System FHIR Compatible

Currently, no standardized UIMA-based type systems for healthcare have found worldwide acceptance. Each text-mining provider usually defines its own type system. Interoperability is created by exchanging type systems between text mining systems – in UIMA, several type systems can be used simultaneously as long as their namespaces are unique – and mappings between type systems.

To make a text mining system FHIR-compatible, two basic procedures can be chosen:

- Defining a FHIR-compatible type system [20], or
- Exposing a FHIR compatible interface.

We pursue the second approach for the following reason: Healthcare data models continue to be the subject of intensive development. In the last 2-3 decades, different versions of the FHIR standard like Version 2, Version 3, Clinical Document Architecture (CDA) and FHIR have been published. There are also other healthcare standards like OpenEHR [21]. In order to support different standards and their versions in parallel, it is not feasible to create and maintain different versions of the text mining system with a different data model for each standard. Instead, it is sufficient and more effective to provide external interfaces to these standards and versions while keeping the data model of the underlying text mining system stable.

In this paper we describe how a FHIR-compatible interface looks like and how FHIR needs to be extended to meet the requirements of Text Mining. In particular, we will answer the following questions:

1. Which are the right FHIR resources and containers to return a text mining result?
2. How can the document text and text-mining annotations such as conditions and observations be linked together?
3. How can typical text mining features, particularly (i) text span of an annotation and (ii) its confidence score be specified in FHIR?
4. How can the types of our text mining system be mapped to FHIR resources and which challenges occur?

Results

FHIR Resources and Container for Text Mining Results

Typically, a text-mining result consists of the document text and associated text-mining annotations. The document text is usually returned back by the text mining system, mainly for the following reason: Text mining annotations contain a *begin* and an *end* attribute, which specify their exact position in the text at character level. Instead of pure text files, the sender often wants to analyse text-based documents such as PDF or MS Word. The text is then extracted during the text mining analysis. The *begin* and *end* attributes always refer to the text returned by the text mining system, not to the original document. The appropriate resource for the document text is a *Binary* resource. A *Binary* can contain any content, whether text, image, pdf, zip archive, etc.

Text Mining annotations can be represented in a corresponding clinical *FHIR Resource*, such as *Condition*, *MedicationStatement*, *Observation* etc. These resources are used to record detailed information about conditions, diagnoses, consumed drugs or any other observation that has direct or indirect impact on the patient's health.

In order to return a *Binary* together with its clinical *Resources* in a single artefact, it has to be grouped into a *Bundle*. A *Bundle* is a collection of resources useful for a variety of different reasons, e.g. returning search results or sending messages. Accordingly, different types of *Bundle*s exist, e.g., *document*, *searchset* or *collection*. In our case, we need a *Bundle* of type *collection*. A *collection Bundle* is a set of resources collected into a single package for ease of distribution.

It is important to understand that a text mining result is not a *document Bundle* in the FHIR sense. FHIR *document Bundles* represent solely structured documents that are authored and assembled in FHIR.

Linking Document Text and Text Mining Annotations: FHIR Provenance

In FHIR, resources in a *Bundle* have an independent existence – they can, e.g., also be accessed using the RESTful API. If the *Binary* was not linked to the clinical resources, it would not be possible to identify the origin of a text-mining determined resource in subsequent processing steps. Therefore, we need another FHIR resource to link both resources, the *Provenance*. *Provenance* contains, among others, three attributes: *entity* gives details about the source of the data (like a PDF, text…), *agent* gives details about the algorithm that extracted the data (software name, vendor, version...) and *target* points at the created resource (e.g. *Condition*, *Observation*, *MedicationStatement*...). For each clinical resource, we also specify a *Provenance* resource. Figure 1 illustrates the resulting *collection Bundle* including clinical and associated *Provenance* resources.

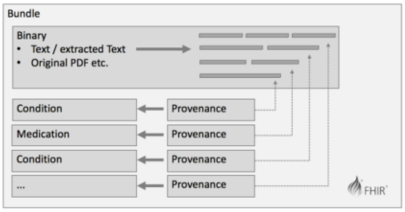

Figure 1– Simplified View of the Provenance Information Conveyed in a FHIR Bundle of Type Collection. The Result Is Conveyed in a FHIR Bundle of Type Collection. A Binary Resource Holds the Source Text (Character Stream) And/or the Extracted Text with the Original PDF Where the Text Was Extracted. Subsequently All Clinical Resources That Come out of the NLP Analysis Are listed. Each of these Resources Is Accompanied by a Provenance Resource.

Specification of Text Spans and Confidence in FHIR

As mentioned above, text mining annotations usually contain begin and end attributes that specify the exact position of an annotation in the text. Since there is no FHIR representation for these attributes yet, we specified two FHIR extensions, currently under approval by the international HL7/FHIR committee.

Extension *Character Range*

The rationale for this extension is to specify, within a source text, the place in which an annotation was found. Thus, the origin of a text mining annotation can be specified not only at document level, but also at word and character level. Traceability is an important feature of text mining systems to help users gain confidence in text mining applications and clinical research based upon text mining results.

The extension is named *character-range*. It contains an attribute named *valueRange* of type *Range*. *valueRange* itself contains the attributes *low* and *high* of type *SimpleQuantity* with the purpose to specify begin and end of annotations in text. *SimpleQuantities* always need a unit. Since characters are countable, their unit in FHIR is "1":

```
"extension": [{
  "url":     "http://example.com/StructureDefinition/character-
range/0.9",
  "valueRange": [{
    "low": {"value": 201,
  "unit": 1},
    "high": {"value": 219,
   "unit": 1}}]}]
```

Extension Confidence

Text mining results may be biased, incomplete or erroneous, their enrichment with a confidence score is therefore crucial. Especially text-mining systems based on machine learning, often provide a confidence value, which estimates the probability for the correctness of the annotation. Thus, the requested extension (ii) is *confidence*, specialising the type *Quantity*, stating the accuracy of annotations. The value range of a confidence is often, but not necessarily, between 0 and 1:

```
 "extension": [{
  "url": "http://example.com/StructureDefinition/confidence/0.9",
   "valueQuantity": {"value": 0.1}}]
```

Mapping Text Mining Types to FHIR Resources

Now that the container, the resources and the necessary extensions have been defined, the mapping of text mining annotations to FHIR resources can be performed in a straightforward way. Both *Health Discovery* and *FHIR* know the notation of a *Concept* or *CodeableConcept*. Concepts exist to encode a medical statement with a term from a terminology. They typically consist of a code, a preferred term, a code system and the version of the code system. In *Health Discovery*, the synonym that is responsible for a match in a text (=matchedTerm) is also an attribute of *Concept*. Table 1 shows the mapping between *HealthDiscovery.Concepts* and *FHIR.CodeableConcepts*. Table 2 gives an overview of the terminology bindings used by *Health Discovery* for different *Concept* types in US English and German. In this work, we focus on the mapping of diagnoses, medications and laboratory values. Table 3 gives an overview of the mappings between *Health Discovery* and *FHIR*.

Table 1 – Mappings between HealthDiscovery.Concept and FHIR.CodeableConcept

Health Discovery	FHIR Resource
Concept	**CodeableConcept**
.conceptid	.code.coding.code
.dictCanon	.code.coding.display
.matchedTerm	.code.text
.source.split('_')[0]	.code.coding.system
.source.split('_')[1]	.code.coding.version

Table 2 – Terminology Binding in Health Discovery for US-English and German

Type	US English	German
Diagnosis	ICD-10 Clinical Modification [22]	ICD-10 German Modification [23]
Drugs	RxNorm [24]	ABDAMed [25]
Ingredients	WHO ATC [26]	German ATC [27]
Lab Values	LOINC [28]	LOINC [28]

Table 3 – HealthDiscovery Types Mapped to FHIR. Instances of (Codeable) Concept Are Marked with an Asterisk (*). The Mapping Between Medication.doseFrequency and MedicationStatement.dosage.timing (**) Is Not Further Specified, as It Is Rather Complex and Beyond this Paper

Averbis Health Discovery	FHIR Resource
Diagnosis	**Condition**
.concept*	.code*
.clinicalStatus	.clinicalStatus
.verificationStatus	.verificationStatus
.belongsTo	.subject.display
.side	.bodySite.text
Medication	**Medication**
.drug.ingredientConcept*	.ingredient*
.doseForm*	.form*
	MedicationStatement
.status	.status
.drug.strength.value	.dosage.doseQuantity.value
.drug.strength.unit	.dosage.doseQuantity.unit
.doseFrequency**	.dosage.timing**
.administrations*	.dosage.method*
LaboratoryValue	**Observation**
.parameter*	.code*
.fact.value	.valueQuantity.value
.fact.unit	.valueQuantity.unit
.interpretation	.interpretation.text

.upperLimit.normalizedValue	.referenceRange.high.value
.upperLimit.normalizedUnit	.referenceRange.high.unit
.lowerLimit.normalizedValue	.referenceRange.low.value
.lowerLimit.normalizedUnit	.referenceRange.low.unit

Assembling the JSON Response

Figure 2 shows an excerpt of a text mining response for the text “The patient had a cold”. “Cold” is recognized as entity of the type *Condition* and annotated with the ICD-10 code “*J00 Acute nasopharyngitis* [common cold]”.

```
{"resourceType": "Bundle",
 "id": "628320",
 "meta": {
  "versionId": "1",
  "lastUpdated": "2018-11-24T15:00:45.182+00:00"},
 "type": "collection",
 "entry": [{
   "fullUrl": "urn:uuid:f844ec9a-ef45-11e3-8bb6-
00aa004d0001",
   "resource": {
    "resourceType": "Binary",
    "contentType": "text/plain",
    "content": "UGF0aWVudCBoYWQgYSBjb2xk"}},
  {"fullUrl": "urn:uuid:f844ec9a-ef45-11e8-8bb6-
00aa004d0001",
   "resource": {
    "resourceType": "Device",
    "manufacturer": "Averbis GmbH",
    "model": "Health Discovery",
    "version": "5.6.0"}},
  {"fullUrl": "urn:uuid:f844ec9d-ef45-11e8-96f5-
00aa004d0001",
   "resource": {
    "resourceType": "Condition",
    "clinicalStatus": "active",
    "verificationStatus": "unknown",
    "code": {
     "coding": [{
      "system": "http://hl7.org/fhir/sid/icd-10-cm",
      "version": "2018",
      "code": "J00",
      "display": "Acute nasopharyngitis [common cold]"}]},
    "subject": {
     "display": "anonymous patient"}}},
  {"fullUrl": "urn:uuid:f844ec9d-ef45-11e8-96f3-
00aa004d0001",
   "resource": {
    "resourceType": "Provenance",
    "target": [{
     "reference": "urn:uuid:f844ec9d-ef45-11e8-96f5-
00aa004d0001"}],
    "recorded": "2018-11-23T18:34:03.184859+01:00",
    "agent": [{
     "whoReference": {
      "reference": "urn:uuid:f844ec9a-ef45-11e8-8bb6-
00aa004d0001"}}],
    "entity": [{
     "extension": [{
       "url":
"http://example.com/StructureDefinition/confidence/0.9",
       "valueQuantity": {
        "value": 0.1}},
      {"url": "http://example.com/StructureDefinition/character-
range/0.5",
       "valueRange": {
        "low": {
         "value": 15,
         "unit": "1"},
        "high": {
         "value": 19,
         "unit": "1"}}}],
     "role": "source",
     "whatReference": {
      "reference": "urn:uuid:f844ec9a-ef45-11e3-8bb6-
00aa004d0001"}}]}}]}
```

Figure 2– JSON Response.

Discussion

Provenance aspects of clinical data have been awarded major emphasis in recent years. The growing amount of EHRs provides unprecedented opportunity for its re-use in many tasks. However, there are various caveats to the use of such data, including inaccuracy, incompleteness and unknown provenance [29]. This is even more the case when EHR content is furthermore processed, such as by text mining systems. A dataset in FHIR created by structured data entry done by physicians requires a different interpretation compared to a seemingly identical dataset produced by text mining analysis of an unstructured narrative. For instance, the former one might meet the quality level required for triggering clinical decision support algorithms, whereas the latter one might not, although it may be perfectly suited for supporting cohort selection or outcome analyses. *Confidence* values that qualify each text mining generated data element can here be used as a valuable filter.

Quality assurance of text mining systems also requires formative assessment cycles in which human experts iteratively check extracted data elements against the source in order to assess their accuracy. If FHIR-based resources were not able to preserve this link between source and extract, this gap would have to be bridged outside the FHIR standard. The new *character-range* specification allows embedding text mining quality assurance completely within FHIR.

Our work is original in the sense that it extends FHIR to incorporate provenance and quality information. With the two extensions for *confidence* and *character-range*, FHIR now contains all specifications to represent text mining generated EHR extracts. Apart from this, two recent publications are worth mentioning, that combined FHIR with text mining. In [30], a FHIR data model automatically structures prostate pathology reports in several natural languages. In [18], the FHIR data model was also used to automatically structure free text and, in addition, to integrate structured data. In contrast to these systems, we argue that – due to different HL7 dialects and other standards – it is not necessary to base the underlying text mining system internally on FHIR, but to ensure that its external interface provides FHIR compatibility, as we have demonstrated. Thus, different text mining systems can operate in parallel or in combination on different HL7 dialects, feeding their results in a standardised information model.

Conclusions

FHIR is increasingly accepted as a universal clinical information model. It supports provenance assertions of data as a critical foundation for authenticity, trust and reproducibility. Provenance is especially relevant if datasets are created by machines, e.g. by text mining and machine learning. We have demonstrated how the output of a commercial clinical text mining tool can be harmonised with FHIR, preserving important provenance information like character span and confidence value. Necessary extensions to the FHIR standard are currently being processed by the international HL7/FHIR committee.

Acknowledgements

The consortia SMITH and MIRACUM are funded by the German Federal Ministry of Education and Research (BMBF) within the Medical Informatics Funding Scheme (FKZ 1ZZ1606A and 01ZZ1609A).

References

[1] J. Heflin and J. Hendler, *Semantic interoperability on the web*. Maryland Univ. College Park Dept. of Computer Science, 2000.
[2] D. Kalra and D. Lloyd, EN13606 *Electronic Health Record Communication Part 1: Reference Model*. CEN TC/251, Brussels, 2007.
[3] T. Beale, Archetypes: Constraint-based domain models for future-proof information systems. OOPSLA 2002 workshop on behavioural semantics (Vol. 105).
[4] HL7 FHIR Resource list: https://www.hl7.org/fhir/resourcelist.html [accessed March 13, 2019]
[5] J.C. Mandel, D.A. Kreda, K.D. Mandl, I.S. Kohane, and R.B Ramoni, SMART on FHIR: a standards-based, interoperable apps platform for electronic health records, *JAMIA* 23 (2016), 899-908.
[6] D. Bender and K. Sartipi, HL7 FHIR: An agile and RESTful approach to healthcare information exchange. *IEEE 26th International Symposium on Computer-Based Medical Systems* (2013), 326-331.
[7] S.M. Meystre, G.K. Savova, K.C. Kipper-Schuler, and J.F. Hurdle, Extracting information from textual documents in the electronic health record: a review of recent research. *IMIA Yearbook of Medical Informatics* (2008), 128-144.
[8] https://averbis.com/health-discovery/ [accessed March 13, 2019]
[9] http://www.trinetx.com/nlp/ [accessed March 13, 2019]
[10] S. Schulz, T. Seddig, S. Hanser, A. Zaiss, and P. Daumke. Checking coding completeness by mining discharge summaries. *Stud Health Technol Inform.* (2011), 594-598.
[11] C. Catarino, A. Grandjean, S. Doss, M. Mücke, S. Tunc, K. Schmidt, J. Schmidt, P. Young, T. Bäumer, C. Kornblum, M. Endres, P. Daumke, T. Klopstock, and B. Schoser (2017), mineRARE: Semantic text-mining of electronic medical records as diagnostic decision support tool to search for rare neurologic diseases such as Pompe disease, Fabry disease and Niemann-Pick type C disease. *European Journal of Neurology* 24 (2017), 75.
[12] D. Schober, R. Choquet, K. Depraetere, F. Enders, P. Daumke, M.C. Jaulent, D. Teodoro, E. Pasche, C. Lovis, and M. Boeker, DebugIT: Ontology-mediated layered Data Integration for real-time Antibiotics Resistance Surveillance. *Semantic Web Applications and Tools for Life Sciences*, Berlin 2014.
[13] A. Gerstmair, P. Daumke, K. Simon, and E. Kotter. Intelligent image retrieval based on radiology reports. *Eur Radiol*. 22 (2012), 2750-2758
[14] K. Tomanek, P. Daumke, F. Enders, J. Huber, K. Theres, and M.L. Müller, An Interactive De-Identification-System. *SMBM – 5th International Symposium on Semantic Mining in Biomedicine* (2012), 82-86.
[15] D.A. Ferrucci, E.W. Brown, J. Chu-Carroll, J. Fan, D. Gondek, A. Kalyanpur, A. Lally, J.W. Murdock, E. Nyberg, J.M. Prager, N. Schlaefer, Nico and C.A. Welty. Building Watson: An overview of the DeepQA project." *AI magazine* 31.3 (2010), 59-79.
[16] I. Gurevych, M. Mühlhäuser, C. Müller, J. Steimle, M. Weimer, and T. Zesch, Darmstadt Knowledge Processing Repository Based on UIMA. *Proceedings of the First Workshop on Unstructured Information Management Architecture at Biannual Conference of the Society for Computational Linguistics and Language Technology.* Heidelberg: Springer (2007), 89.
[17] G.K. Savova, J.J. Masanz, P.V. Ogren, J. Zheng, S. Sohn, K.C. Kipper-Schuler, and C.G. Chute, Mayo clinical Text Analysis and Knowledge Extraction System (cTAKES): Architecture, Component Evaluation and Applications. Journal of the American Medical Informatics Association: *JAMIA* 17 (2010), 507-513.
[18] S. Sohn, C. Clark, S.R. Halgrim, S.P. Murphy, C.G. Chute CG, and H. Liu. MedXN: an open source medication extraction and normalization tool for clinical text. *JAMIA* 21 (2014), 858-865.
[19] Y.K. Lin, H. Chen, and R.A. Brown, MedTime: a temporal information extraction system for clinical narratives. *J Biomed Inform* 46 (2013), Suppl:S20-8.
[20] N. Hong, A. Wen, F. Shen, S. Sohn, S. Liu, H. Liu, and G. Jiang, Integrating structured and unstructured EHR data using an FHIR-based type system: a case study with medication data. *AMIA Summits on Translational Science Proceedings* (2017), 74.
[21] https://openehr.org/ [accessed March 13, 2019]
[22] https://www.cdc.gov/nchs/icd/icd10cm.htm [accessed March 13, 2019]
[23] https://www.dimdi.de/dynamic/de/klassifikationen/icd/icd-10-gm/ [accessed March 13, 2019]
[24] https://www.nlm.nih.gov/research/umls/rxnorm/ [accessed March 13, 2019]
[25] http://abdata.de/datenangebot/abdamed/ [accessed March 13, 2019]
[26] https://www.whocc.no/atc_ddd_index/ [accessed March 13, 2019]
[27] https://www.dimdi.de/dynamic/de/arzneimittel/atc-klassifikation/ [accessed March 13, 2019]
[28] https://loinc.org [accessed March 13, 2019]
[29] W.R. Hersh, M.G. Weiner, P.J. Embi, J.R. Logan, P.R. Payne, E.V. Bernstam, H.P. Lehmann, G. Hripcsak, T.H. Hartzog, J.J. Cimino, and J.H. Saltz, Caveats for the use of operational electronic health record data in comparative effectiveness research. *Med Care* 51 (2013), S30-37.
[30] P. Prinsen, R. van Ommering, M. Janssen, M. Loos, A. Tahmasebi, S. Thompson, S. Coons, P. Henrik Vincent, C. Lindh, O. Akre, A. Rastinehad, L.A. Mynderse, and P.J. Chang, A fully automatic solution for structuring prostate pathology reports. *Journal of Clinical Oncology* 36 (2018), e18586-e18586.

Address for Correspondence

Philipp Daumke, Averbis GmbH, Averbis GmbH, Salzstraße 15, 79098 Freiburg, Germany, info@averbis.com