

Gartner Case Study: Entity-Event Knowledge Graph for Powering AI Solutions (Montefiore)

Gartner featured Franz's customer, Montefiore Medical Center, in a research report on Montefiore's Entity-Event Knowledge Graph:

"AI solutions are often hindered by fragmented data and siloed point solutions," according to Gartner's Chief Data and Analytics Officer Research Team. "Montefiore's data and analytics leader used semantic knowledge graphs to power its AI solutions and achieved considerable cost savings as well as improvements in timeliness and the prediction accuracy of AI models." Source: Gartner Case Study: Entity-Event Knowledge Graph for Powering AI Solutions (Montefiore) – Subscription required.

Copy Available from Montefiore/Einstein.

Document Knowledge Graphs with NLP and ML

A core competency for Franz Inc is turning text and documents into Knowledge Graphs (KG) using Natural Language Processing (NLP) and Machine Learning (ML) techniques in combination with AllegroGraph. In this document we discuss how the techniques described in [NLP and ML components of AllegroGraph] can be

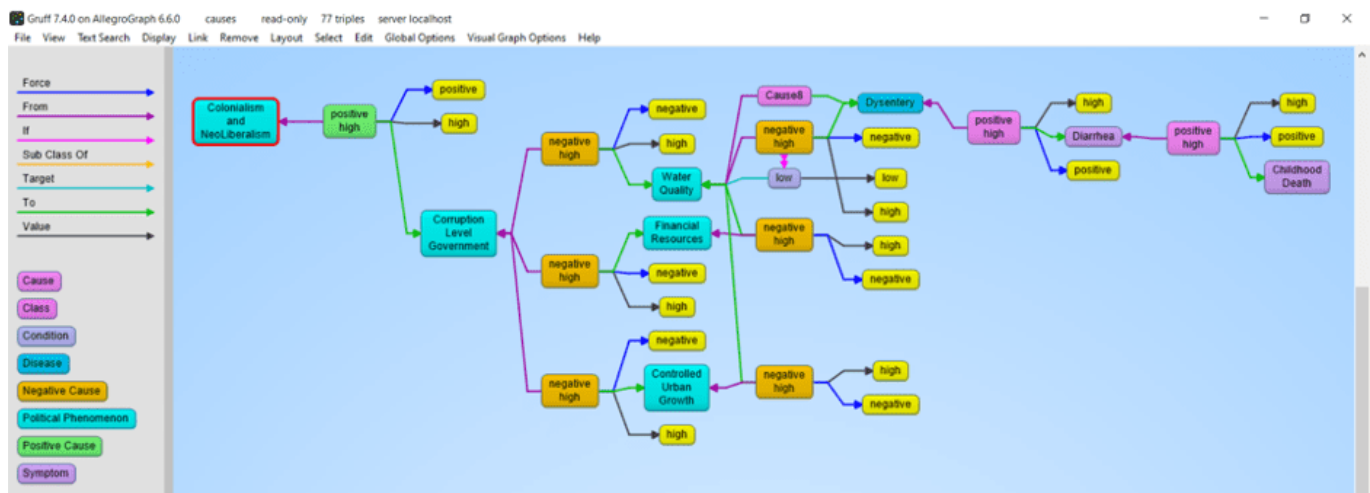
combined with popular software tools to create a robust Document Knowledge Graph pipeline.

We have applied these techniques for several Knowledge Graphs but in this document we will primarily focus on three completely different examples that we summarize below. First is the Chomsky Legacy Project where we have a large set of very dense documents and very different knowledge sources, Second is a knowledge graph for an intelligent call center where we have to deal with high volume dynamic data and real-time decision support and finally, a large government organization where it is very important that people can do a semantic search against documents and policies that steadily change over time and where it is important that you can see the history of documents and policies.

Example [1] Chomsky Knowledge Graph

The Chomsky Legacy Project is a project run by a group of admirers of Noam Chomsky with the primary goal to preserve all his written work, including all his books, papers and interviews but also everything written about him. Ultimately students, researchers, journalists, lobbyists, people from the AI community, and linguists can all use this knowledge graph for their particular goals and questions.

The biggest challenges for this project are finding causal relationships in his work using event and relationship extraction. A simple example we extracted from an author quoting Chomsky is that neoliberalism ultimately causes childhood death.



Example 2: N3 Results and the Intelligent Call Center

This is a completely different use case (See a recent KMWorld Article <https://allegrograph.com/knowledge-graphs-enhance-customer-experience-through-speed-and-accuracy/>). Whereas the previous use case was very static, this one is highly dynamic. We analyze in real-time the text chats and spoken conversations between call center agents and customers. Our knowledge graph software provides real-time decision support to make the call center agents more efficient. N3 Results helps big tech companies to sell their high tech solutions, mostly cloud-based products and services but also helps their clients sell many other technologies and services.

The main challenge we tackle is to really deeply understand what the customer and agent are talking about. None of this can be solved by only simple entity extraction but requires elaborate rule-based and machine learning techniques. Just to give a few examples. We want to know if the agent talked about their most important talking points: that is, did the agent ask if the customer has a budget, or the authority to make a decision or a timeline about when they need the new technology or whether they actually have expressed their need. But also whether the agent reached the right person, and whether the agent talked about the follow-up. In addition, if the customer talks about competing technology we need to recognize that and provide the agent in real-time with a battle card specific to the competing technology. And in order to be able to do the

latter, we also analyzed the complicated marketing materials of the clients of N3.

Example 3: Complex Government Documents

Imagine a regulatory body with tens of thousands of documents. Where nearly every paragraph has reference to other paragraphs in the same document or other documents and the documents change over time. The goal here is to provide the end-users in the government with the right document given their current task at hand. The second goal is to keep track of all the changes in the documents (and the relationship between documents) over time.

The Document to Knowledge Graph Pipeline

Process Name	Input	Output
1. Custom Taxonomy Creation	Corpus Analytics, Taxonomy tool	A SKOS taxonomy containing concepts, concept hierarchy, prefLabels, altLabels.
2. Document Preparation	Documents (pdf, word, ppt, xlsx), Apache Tika, Spacy for XML cleanup	An XML version of each document
3. Extract Document Meta Data	Document + Apache Tika	JSON dictionary of the Document MetaData
4. XML-to-Triples	XML+JSON dictionary, XMLToTriples.py	Graph-based document tree with chapters, sections, and paragraphs as triples. Also includes meta data as triples
5. Entity-Extraction	Paragraphs + taxonomies + AllegroGraph Entity extract or external extractors	Concepts, persons, places, currencies. Connected to paragraphs
6. LOD Enrichment	Paragraphs + IBM Natural Language Understanding.	Concept categories and links to DBpedia and GeoNames, etc.
7. Complex Relationship and Event extraction.	Paragraphs + Taxonomy + Rules in Spacy or AllegroGraph	Complex events and relationships, References to other document sections.
8. NLP and ML	Chapters and paragraphs + all the tools described [here], but also using Spacy, Gensim, BERT, SciKit Learn.	Similarities, sentiment, query answering, smart search, text classification, word embeddings, abstracts
9. Versioning and Document tracking	Old + New document, compare.py	Old document in historic repository, new document in current, changed graph.
10. Statistical Relationships	Concepts + OddRatio.py or OddsRatio.cl	Statistical relationships between concepts.

Let us first give a quick summary in words of how we turn documents into a Knowledge Graph.

[1] Taxonomy Creation

Taxonomy of all the concepts important to the business using open source or commercial taxonomy builders. An available industry taxonomy is a good starting point for additional customizations.

[2] Document Preparation


We then take a document and turn it into an intermediate XML using Apache Tika. Apache Tika supports more than 1000 document types and although Apache Tika is a fantastic tool, the output is still usually not clean enough to create a graph from, so we use Spacy rules to clean up the XML to make it as uniform as possible.

[3] Extract Document MetaData

Most documents also contain document metadata (author, date, version, title, etc) and Apache Tika will also deliver the metadata for a document as a JSON object.

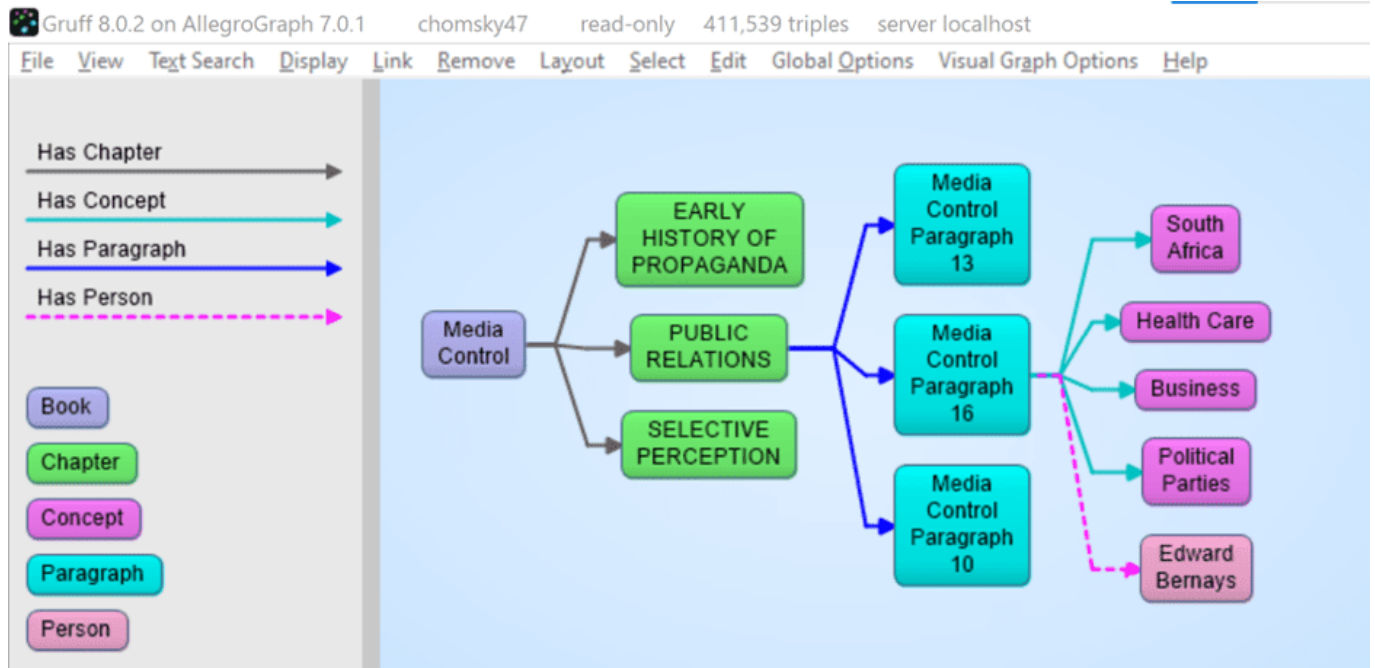
[4] XML to Triples

Our tools ingest the XML and metadata and transform that into a graph-based document tree. The document is the root and from that, it branches out into chapters, optionally sections, all the way down to paragraphs. The ultimate text content is in the paragraphs. In the following example we took the XML version of Noam Chomsky's book Media Control and turned that into a tree. The following shows a tiny part of that tree. We start with the Media Control node, then we show three (of the 11) chapters, for one chapter we show three (of the 6) paragraphs, and then we show the actual text in that paragraph. We sometimes can go even deeper to the level of sentences and tokens but for most projects that is overkill.




[5] Entity Extractor

AllegroGraph's entity extractor takes as input the text of each paragraph in the document tree and one or more of the taxonomies and returns recognized SKOS concepts based on `prefLabels` and `altLabels`. AllegroGraph's entity extractor is state of the art and especially powerful when it comes to complex terms like product names. We find that in our call center a technical product name can sometimes have up to six synonyms or very specific jargon. For example the Cisco product Catalyst 9000 will also be abbreviated as the cat 9k. Instead of developing `altLabels` for every possible permutation that human beings *will* use, we have specialized heuristics to optimize the yield from the entity extractor. The following picture shows 4 (of the 14) concepts discovered in paragraph 16. Plus one person that was extracted by IBM's NLU.



[6] Linked Data Enrichment

In many use cases, AllegroGraph can link extracted entities to concepts in the linked data cloud. The most prominent being DBpedia, wikidata, the census database, GeoNames, but also many Linked Open Data repositories. One tool that is very useful for this is IBM's Natural Language Understanding program but there are others available. In the following image we see that the Nelson Mandela entity (Red) is linked to the dbpedia entity for Nelson Mandela and that then links to the DBpedia itself. We extracted some of his spouses and a child with their pictures.



[7] Complex Relationship and Event Extraction

Entity extraction is a first good step to 'see' what is in your documents but it is just the first step. For example: how do you find in a text whether company C1 merged with company C2. There are many different ways to express the fact that a company fired a CEO. For example: Uber got rid of Kalanick, Uber and Kalanick parted ways, the board of Uber kicked out the CEO, etc. We need to write explicit symbolic rules for this or we need a lot of training data to feed a machine learning algorithm.

[8] NLP and Machine Learning

There are many many AI algorithms that can be applied in Document Knowledge Graphs. We provide best practices for topics like:

- [a] Sentiment Analysis, using good/bad word lists or training data.

- [b] Paragraph or Chapter similarity using statistical techniques like Gensim similarity or symbolic techniques where we just the overlap of recognized entities as a function of the size of a text.

- [c] Query answering using word2vec or more advanced techniques like BERT

- [d] Semantic search using the hierarchy in SKOS taxonomies.


- [e] Summarization techniques for Abstractive or Extractive abstracts using Gensim or Spacy.

[9] Versioning and Document tracking

Several of our customers with Document Knowledge Graphs have noted the one constant in all of these KGs is that documents change over time. As part of our solution, we have created best practices where we deal with these changes. A crucial first step is to put each document in its own graph (i.e. the fourth element of every triple in the document tree is the document id itself). When we get a new version of a document the document ID changes but the new document will point back to the old version. We then compute which paragraphs stayed the same within a certain margin (there are always changes in whitespace) and we materialize what paragraphs disappeared in the new version and what new paragraphs appeared compared to the previous version. Part of the best practice is to put the old version of a document in a historical database that at all times can be federated with the 'current' set of documents.

Note that in the following picture we see the progression of a document. On the right hand side we have a newer version of a document 1100.161 with a chapter -> section -> paragraph -> contents where the content is almost the same as the one in

the following gruff picture that if Noam Chomsky talks about South Africa then the chances are very high he will also talk about Nelson Mandela.



The Knowledge Graph Cookbook

Recipes for Knowledge Graphs that Work:

- Learn why and how to build knowledge graphs that help enterprises use data to innovate, create value and increase revenue. This practical manual is full of recipes and knowledge on the subject.
- Learn more about the variety of applications based on knowledge graphs.
- Learn how to build working knowledge graphs and which technologies to use.
- See how knowledge graphs can benefit different parts of your organization.

- Get ready for the next generation of enterprise data management tools.

Dr. Jans Aasman, CEO, Franz Inc. is interviewed in the Expert Opinion Section.

“KNOWLEDGE GRAPHS AREN’T WORTH THEIR NAME IF THEY DON’T ALSO LEARN AND BECOME SMARTER DAY BY DAY” – Dr. Aasman

INTERVIEWS

The creation of knowledge graphs is interdisciplinary. Good chefs regularly visit other restaurants for inspiration. We have asked experts working in the field of knowledge graphs and semantic data modelling to comment on their experience in this area. They have worked with various stakeholders in different industries, so that you, dear reader, may further develop your understanding of the topic.



JANS AASMAN

FRANZ

Dr. Jans Aasman is CEO at Franz Inc., a leading provider of Knowledge Graph Technologies (AllegroGraph) and AI-based Enterprise solutions. Dr. Aasman is a noted speaker, author, and industry evangelist on all things graph.

“KNOWLEDGE GRAPHS AREN’T WORTH THEIR NAME IF THEY
DON’T ALSO LEARN AND BECOME SMARTER DAY BY DAY”

Click [here](#) to get the book as free PDF or Kindle version.

Franz's 2020 Predictions in the News

Looking to the future of AI, KnowledgeGraph and Semantics we had a number of publications cover our views of where AllegroGraph is headed.

Datanami

20 AI Predictions for 2020

We're still in the midst of a fake news crisis, and with the emergence of deep fakes, it will likely get worse. Luckily, we have the technology available to begin to address it, says Dr. Jans Aasman, the CEO of Franz.

"Knowledge graphs, in combination with deep learning, will be used to identify photos and video that have been altered by superimposing existing images and videos onto source images," Aasman says. "Machine learning knowledge graphs will also unveil the origin of digital information that has been published by a foreign source. Media outlets and social networks will use AI knowledge graphs as a tool to determine whether to publish information or remove it."

DBTA

Ten Predictions for AI and Machine Learning in 2020

AI Knowledge Graphs will Debunk Fake News: "Knowledge Graphs in combination with deep learning will be used to identify photos and video that have been altered by superimposing existing images and videos onto source images. Machine learning knowledge graphs will also unveil the origin of digital

information that has been published by a foreign source. Media outlets and social networks will use AI Knowledge Graphs as a tool to determine whether to publish information or remove it.” – Dr. Jans Aasman, CEO of Franz, Inc.

SD Times

Software predictions for 2020 from around the industry

Jans Aasman, CEO of Franz, Inc.

Digital immortality will emerge: We will see digital immortality emerge in 2020 in the form of AI digital personas for public figures. The combination of Artificial Intelligence and Semantic Knowledge Graphs will be used to transform the works of scientists, technologists, politicians and scholars into an interactive response system that uses the person’s actual voice to answer questions. AI digital personas will dynamically link information from various sources – such as books, research papers and media interviews – and turn the disparate information into a knowledge system that people can interact with digitally. These AI digital personas could also be used while the person is still alive to broaden the accessibility of their expertise.

Dataversity

Semantic Web and Semantic Technology Trends in 2020

“The big-name Silicon Valley companies (LinkedIn, Airbnb, Apple, Uber) are all building knowledge graphs. But more importantly, Fortune 500 companies, especially banks, are also investing in knowledge graph solutions.”

IoT gets into the picture too. Aasman points to “digital twins,” which can be thought of as specialized knowledge graphs, as an exceptionally lucrative element of the

technology with an applicability easily lending itself to numerous businesses. Its real-time streaming data, simulation capabilities, and relationship awareness may well prove to be the 'killer app' that takes the IoT mainstream, he said. As an example, by consuming data transmitted by IoT sensors, digital twins will inform the monitoring, diagnostics, and prognostics of power grid assets to optimize asset performance and utilization in near real-time.

InsideBigData

2020 Trends in Data Modeling: Unparalleled Advancement

Shapes Constraint Language (SHACL): SHACL is a framework that assists with data modeling by describing the various shapes of data in knowledge graph settings, which produces the desirable downstream effect of enabling organizations to automate "the validation of your data," remarked Franz CEO Jans Aasman. SHACL operates at a granular level involving classifications and specific data properties.

Workflow

2020 Trends in CyberSecurity

Software-defined perimeter transmissions also guard information at the data layer by utilizing Datagram Transport Layer Security (DTLS) encryption and Public Key Authentication. Fortifying information assets at the data layer is likely the most dependable method of protecting them, because it's the layer in which the data are actually stored. It's important to distinguish data layer security versus access layer security. The latter involves a process known as security filtering in which, based on particular roles or responsibilities, users can access data. "You can specify

filters where for a particular user or a particular role whether you could see or not see particular [data]," Franz CEO Jans Aasman said. "You could say if someone has the role administrator, we're telling the system 'administrators cannot see [certain data]'. "

Moreover, triple attributes can be based on compliance needs specific to regulations – which is immensely utilitarian in the post-GDPR data landscape. "For the government you could have a feature of whether you're a foreigner or not," Aasman said. "HIPAA doesn't care whether you're a foreigner or not, but you can do a separate mechanism for it."

2020 Trend Setting Products – AllegroGraph

Franz Inc. is proud to announce that it has been named to the 2020 Trend Setting Products in Data Management by Database Trends and Application Magazine.

Database Trends and Applications (DBTA) magazine announced its seventh annual list of trend-setting products in data management and analysis. The list, "DBTA Trend-Setting Products for 2020," recognizes products in the marketplace that are both innovative and effective in helping customers address evolving challenges and opportunities. In all, 100 products are highlighted in the special December edition of *Database Trends and Applications* magazine and on the DBTA website, www.dbta.com.

"The world of data management and analytics continues to

evolve rapidly with new technologies and strategies,” remarked Thomas Hogan, Group Publisher of *Database Trends and Applications*. “Cutting through the hype and identifying products that deliver results in the real world is more important than ever. This list highlights products that are truly transformative in bringing greater agility, efficiency and innovation to market.”

“We are honored to receive this acknowledgement for our efforts in delivering Enterprise Knowledge Graph Solutions,” said Dr. Jans Aasman, CEO, Franz Inc. “In the past year, we have seen demand for Enterprise Knowledge Graphs take off across industries along with recognition from top technology analyst firms that Knowledge Graphs provide the critical foundation for artificial intelligence applications and predictive analytics. Our AllegroGraph Knowledge Graph Platform Solution offers a unique comprehensive approach for helping companies accelerate the creation of Enterprise Knowledge Graphs that deliver new value to their organization.”

The Importance of FAIR Data in Earth Science

Franz’s CEO, Jans Aasman’s recent Marine Technology News:

Data’s valuation as an enterprise asset is most acutely realized over time. When properly managed, the same dataset



supports a plurality of use cases, becomes almost instantly available upon request, and is exchangeable between departments or organizations to systematically increase its yield with each deployment.

These boons of leveraging data as an enterprise asset are the foundation of GO FAIR's Findable Accessible Interoperable Reusable (FAIR) principles profoundly impacting the data management rigors of geological science. Numerous organizations in this space have embraced these tenets to swiftly share information among a diversity of disciplines to safely guide the stewardship of the earth.

According to Dr. Annie Burgess, Lab Director of Earth Science Information Partners (ESIP), the "most pressing global challenges cannot be solved by a single organization. Scientists require data collected across multiple disciplines, which are often managed by many different agencies and institutions." As numerous members of the earth science community are realizing, the most effectual means of managing those disparate data according to FAIR principles is by utilizing the semantic standards underpinning knowledge graphs.

Read the full article at [Marine Technology News](#)

Ontology Summit 2020 –

Knowledge Graphs

The Ontology Summit is an annual series of events that involves the ontology community and communities related to each year's theme chosen for the summit. The Ontology Summit was started by Ontolog and NIST, and the program has been co-organized by Ontolog, NIST, NCOR, NCB0, IA0A, NCO_NITRD along with the co-sponsorship of other organizations that are supportive of the Summit goals and objectives.

Knowledge graphs, closely related to ontologies and semantic networks, have emerged in the last few years to be an important semantic technology and research area. As structured representations of semantic knowledge that are stored in a graph, KGs are lightweight versions of semantic networks that scale to massive datasets such as the entire World Wide Web. Industry has devoted a great deal of effort to the development of knowledge graphs, and they are now critical to the functions of intelligent virtual assistants such as Siri and Alexa. Some of the research communities where KGs are relevant are Ontologies, Big Data, Linked Data, Open Knowledge Network, Artificial Intelligence, Deep Learning, and many others.

Dr. Jans Aasman presented – ***“Why Knowledge Graphs Hit the Hype Cycle and What they have in common”***

Presentation Page

Presentation Slides



SHACL – Shapes Constraint Language in AllegroGraph

SHACL is a SHApE Constraint Language. It specifies a vocabulary (using triples) to describe the shape that data should have. The *shape* specifies things like the following simple requirements:

- How many triples with a specified subject and predicate should be in the repository (e.g. at least 1, at most 1, exactly 1).
- What the nature of the object of a triple with a specified subject and predicate should be (e.g. a string, an integer, etc.)

See the specification for more examples.

SHACL allows you to validate that your data is conforming to desired requirements.

For a given validation, the shapes are in the *Shapes Graph* (where *graph* means a collection of triples) and the data to be validated is in the *Data Graph* (again, a collection of triples). The SHACL vocabulary describes how a given shape is linked to *targets* in the data and also provides a way for a Data Graph to specify the Shapes Graph that should be used for validation. The result of a SHACL validation describes whether the Data Graph conforms to the Shapes Graph and, if it does not, describes each of the failures.

Namespaces Used in this Document

Along with standard predefined namespaces (such

as `rdf:` for `<http://www.w3.org/1999/02/22-rdf-syntax-ns#>` and `rdfs:` for `<http://www.w3.org/2000/01/rdf-schema#>`), the following are used in code and examples below:

```
prefix fr: <https://franz.com#>
prefix sh: <http://www.w3.org/ns/shacl#>
prefix franz: <https://franz.com/ns/allegrograph/6.6.0/>
```

A Simple Example

Suppose we have a *Employee* class and for each *Employee* instance, there must be exactly one triple of the form

```
emp001 hasID "000-12-3456"
```

where the object is the employee's ID Number, which has the format is [3 digits]-[2 digits]-[4 digits].

This TriG file encapsulates the constraints above:

```
@prefix sh: <http://www.w3.org/ns/shacl#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .

<https://franz.com#Shapes> {
  <https://franz.com#EmployeeShape>
    a sh:NodeShape ;
    sh:targetClass <https://franz.com#Employee> ;
    sh:property [
      sh:path <https://franz.com#hasID> ;
      sh:minCount 1 ;
      sh:maxCount 1 ;
      sh:datatype xsd:string ;
      sh:pattern "^[0-9][0-9][0-9]-[0-9][0-9]-[0-9][0-9][0-9][0-9]$" ;
    ] .
}
```

It says that for instances of `fr:Employee` (`sh:targetClass <https://franz.com#Employee>`), there must be exactly 1 triple with predicate (path) `fr:hasID` and the object of that triple

must be a string with pattern [3 digits]-[2 digits]-[4 digits]
(sh:pattern "^[0-9][0-9][0-9]-[0-9][0-9]-[0-9][0-9][0-9][0-9]\$").

This TriG file defines the Employee class and some employee instances:

```
@prefix fr: <https://franz.com#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
```

```
{
  fr:Employee
    a rdfs:Class .
  fr:emp001
    a fr:Employee ;
    fr:hasID "000-12-3456" ;
    fr:hasID "000-77-3456" .
  fr:emp002
    a fr:Employee ;
    fr:hasID "00-56-3456" .
  fr:emp003
    a fr:Employee .
}
```

Recalling the requirements above, we immediately see these problems with these triples:

1. *emp001* has two *hasID* triples.
2. The value of *emp002*'s ID has the wrong format (two leading digits rather than 3).
3. *emp003* does not have a *hasID* triple.

We load the two TriG files into our repository, and end up with the following triple set. Note that all the employee triples use the default graph and the SHACL-related triples use the graph `<https://franz.com#Shapes>` specified in the TriG file.

s	p	o	g
Employee	rdf:type	rdfs:Class	
emp001	rdf:type	Employee	
emp001	hasID	"000-12-3456"	
emp001	hasID	"000-77-3456"	
emp002	rdf:type	Employee	
emp002	hasID	"00-56-3456"	
emp003	rdf:type	Employee	
EmployeeShape	property	_:b7A1D241Ax1	Shapes
_:b7A1D241Ax1	pattern	"^[0-9][0-9][0-9]-[0-9][0-9]-[0-9][0-9][0-9]\$"	Shapes
_:b7A1D241Ax1	datatype	xs:string	Shapes
_:b7A1D241Ax1	maxCount	"1"	Shapes
_:b7A1D241Ax1	minCount	"1"	Shapes
_:b7A1D241Ax1	path	hasID	Shapes
EmployeeShape	targetClass	Employee	Shapes
EmployeeShape	rdf:type	NodeShape	Shapes

Now we use **agtool shacl-validate** to validate our data:

```
bin/agtool shacl-validate --data-graph default --shapes-graph
https://franz.com#Shapes shacl-repo-1
```

```
Validation report:          Does not conform
Created:                   2019-06-27T10:24:10
Number of shapes graphs:   1
Number of data graphs:     1
Number of NodeShapes:      1
Number of focus nodes checked: 3
```

3 validation results:

Result:

```
Focus node:      <https://franz.com#emp001>
Path:            <https://franz.com#hasID>
Source Shape:    _:b7A1D241Ax1
                  Constraint                      Component:
<http://www.w3.org/ns/shacl#MaxCountConstraintComponent>
Severity:        <http://www.w3.org/ns/shacl#Violation>
```

Result:

```
Focus node:      <https://franz.com#emp002>
Path:            <https://franz.com#hasID>
Value:           "00-56-3456"
Source Shape:    _:b7A1D241Ax1
                  Constraint                      Component:
<http://www.w3.org/ns/shacl#PatternConstraintComponent>
Severity:        <http://www.w3.org/ns/shacl#Violation>
```

Result:


```

Focus node:      <https://franz.com#emp003>
Path:           <https://franz.com#hasID>
Source Shape:   _:b7A1D241Ax1
                Constraint
Component:      <http://www.w3.org/ns/shacl#MinCountConstraintComponent>
Severity:       <http://www.w3.org/ns/shacl#Violation>

```

The validation fails with the problems listed above. The **Focus node** is the subject of a triple that did not conform. **Path** is the predicate or a property path (predicates in this example). **Value** is the offending value. **Source Shape** is the shape that established the constraint (you must look at the shape triples to see exactly what **Source Shape** is requiring).

We revise our employee data with the following SPARQL expresssion, deleting one of the emp001 triples, deleting the emp002 triple and adding a new one with the correct format, and adding an emp003 triple.

```
prefix fr: <https://franz.com#>
```

```
DELETE DATA {fr:emp002 fr:hasID "00-56-3456" } ;
```

```
INSERT DATA {fr:emp002 fr:hasID "000-14-1772" } ;
```

```
DELETE DATA {fr:emp001 fr:hasID "000-77-3456" } ;
```

```
INSERT DATA {fr:emp003 fr:hasID "000-54-9662" } ;
```

Now our employee triples are

s	p	o	g
emp002	hasID	"000-14-1772"	
emp003	hasID	"000-54-9662"	
Employee	rdf:type	rdfs:Class	
emp001	rdf:type	Employee	
emp001	hasID	"000-12-3456"	
emp002	rdf:type	Employee	
emp003	rdf:type	Employee	

We run the validation again and are told our data conforms:

```
% bin/agtool shacl-validate --data-graph default --shapes-graph https://franz.com#Shapes shacl-repo-1
```

Validation report:	Conforms
Created:	2019-06-27T10:32:19
Number of shapes graphs:	1
Number of data graphs:	1
Number of NodeShapes:	1
Number of focus nodes checked:	3

When we refer to this example in the remainder of this document, it is to the un-updated (incorrect) triples.

SHACL API

The example above illustrates the SHACL steps:

1. Have a data set with triples that should conform to a shape
2. Have SHACL triples that express the desired shape
3. Run SHACL validation to determine if the data conforms

Note that SHACL validation does not modify the data being validated. Once you have the conformance report, you must modify the data to fix the conformance problems and then rerun the validation test.

The main entry point to the API is **agtool shacl-validate**. It takes various options and has several output choices. Online help for **agtool shacl-validate** is displayed by running `agtool shacl-validate --help`.

In order to validate triples, the system must know:

1. What triples to examine
2. What rules (SHACL triples) to use
3. What to do with the results

Specifying what triples to examine

Two arguments to **agtool shacl-validate** specify the triples to evaluate: `--data-graph` and `--focus-node`. Each can be specified

multiple times.

- The `--data-graph` argument specifies the graph value for triples to be examined. Its value must be an IRI or default. Only triples in the specified graphs will be examined. default specifies the default graph. It is also the default value of the `--data-graph` argument. If no value is specified for `--data-graph`, only triples in the default graph will be examined. If a value for `--data-graph` is specified, triples in the default graph will only be examined if `--data-graph default` is also specified.
- The `--focus-node` argument specifies IRIs which are subjects of triples. If this argument is specified, only triples with these subjects will be examined. To be examined, triples must also have graph values specified by `--data-graph` arguments. `--focus-node` does not have a default value. If unspecified, all triples in the specified data graphs will be examined. This argument can be specified multiple times.

The `--data-graph` argument was used in the simple example above. Here is how the `--focus-node` argument can be used to restrict validation to triples with subjects `<https://franz.com#emp002>` and `<https://franz.com#emp003>` and to ignore triples with subject `<https://franz.com#emp001>` (applying **agtool shacl-validate** to the original non-conformant data):

```
% bin/agtool shacl-validate --data-graph default \
  --shapes-graph https://franz.com#Shapes \
  --focus-node https://franz.com#emp003 \
  --focus-node https://franz.com#emp002 shacl-repo-1
Validation report:          Does not conform
Created:                   2019-06-27T11:37:49
Number of shapes graphs:   1
Number of data graphs:     1
Number of NodeShapes:      1
Number of focus nodes checked: 2
```

2 validation results:

Result:

```
Focus node:      <https://franz.com#emp003>
Path:           <https://franz.com#hasID>
Source Shape:    _:b7A1D241Ax2
                  Constraint
Component:       <http://www.w3.org/ns/shacl#MinCountConstraintComponent>
Severity:        <http://www.w3.org/ns/shacl#Violation>
```

Result:

```
Focus node:      <https://franz.com#emp002>
Path:           <https://franz.com#hasID>
Value:          "00-56-3456"
Source Shape:    _:b7A1D241Ax2
                  Constraint
Component:       <http://www.w3.org/ns/shacl#PatternConstraintComponent>
Severity:        <http://www.w3.org/ns/shacl#Violation>
```

Specifying What Shape Triples to Use

Two arguments to **agtool shacl-validate**, analogous to the two arguments for data described above, specify Shape triples to use. Further, following the SHACL spec, data triples with predicate `<http://www.w3.org/ns/shacl#shapeGraph>` also specify graphs containing Shape triples to be used.

The arguments to **agtool shacl-validate** are the following. Each may be specified multiple times.

- The `--shapes-graph` argument specifies the graph value for shape triples to be used for SHACL validation. Its value must be an IRI or default. default specifies the default graph. The `--shapes-graph` argument has no default value. If unspecified, graphs specified by data triples with the `<http://www.w3.org/ns/shacl#shapeGraph>` predicate will be used (they are used whether or not `--shapes-graph` has a value). If `--shapes-graph` has no value and there are no data triples with

the `<http://www.w3.org/ns/shacl#shapeGraph>` predicate, the data graphs are used for shape graphs. (Shape triples have a known format and so can be identified among the data triples.)

- The `--shape` argument specifies IRIs which are subjects of shape nodes. If this argument is specified, only shape triples with these subjects and subsidiary triples to these will be used for validation. To be included, the triples must also have graph values specified by the `--shapes-graph` arguments or specified by a data triple with the `<http://www.w3.org/ns/shacl#shapeGraph>` predicate. `--shape` does not have a default value. If unspecified, all shapes in the shapes graphs will be used.

Other APIs

There is a lisp API using the function `validate-data-graph`, defined next:

```
validate-data-graphdb &key data-graph-iri/s shapes-graph-iri/s shape/s focus-node/s verbose conformance-only?
function
```

Perform SHACL validation and return a validation-report structure.

The validation uses `data-graph-iri/s` to construct the `dataGraph`. This can be a single IRI, a list of IRIs or `NIL`, in which case the default graph will be used. The `shapesGraph` can be specified using the `shapes-graph-iri/s` parameter which can also be a single IRI or a list of IRIs. If `shape-graph-iri/s` is not specified, the SHACL processor will first look to create the `shapesGraph` by finding triples with the predicate `sh:shapeGraph` in the `dataGraph`. If there are no such triples, then the `shapesGraph` will be assumed to be the same as the `dataGraph`.

Validation can be restricted to particular shapes and focus nodes using the `shape/s` and `focus-node/s` parameters. Each of these can be an IRI or list of IRIs.

If `conformance-only?` is `true`, then validation will stop as soon as any validation failures are detected.

You can use `validation-report-conforms-p` to see whether or not the `dataGraph` conforms to the `shapesGraph` (possibly restricted to just particular `shape/s` and `focus-node/s`).

The function `validation-report-conforms-p` returns `t` or `nil` as the validation struct returned by `validate-data-graph` does or does not conform.

`validation-report-conforms-preport`
function

Returns `t` or `nil` to indicate whether or not `REPORT` (a `validation-report` struct) indicates that validation conformed. There is also a REST API. See HTTP reference.

Validation Output

The simple example above and the SHACL examples below show output from **agtool validate-shacl**. There are various output formats, specified by the `--output` option. Those examples use the plain format, which means printing results descriptively. Other choices include `json`, `trig`, `trix`, `turtle`, `nquads`, `rdf-n3`, `rdf/xml`, and `ntriples`. Here are the simple example (uncorrected) results using `ntriples` output:

```
% bin/agtool shacl-validate --output ntriples --data-graph
default --shapes-graph https://franz.com#Shapes shacl-repo-1
```

```
_ : b271983AAx1
<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
<http://www.w3.org/ns/shacl#ValidationReport> .
```

```

_:b271983AAx1      <http://www.w3.org/ns/shacl#conforms>
"false"^^<http://www.w3.org/2001/XMLSchema#boolean> .
_:b271983AAx1      <http://purl.org/dc/terms/created>
"2019-07-01T18:26:03"^^<http://www.w3.org/2001/XMLSchema#dateT
ime> .
_:b271983AAx1      <http://www.w3.org/ns/shacl#result>
_:b271983AAx2 .
_:b271983AAx2
<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
<http://www.w3.org/ns/shacl#ValidationResult> .
_:b271983AAx2      <http://www.w3.org/ns/shacl#focusNode>
<https://franz.com#emp001> .
_:b271983AAx2      <http://www.w3.org/ns/shacl#resultPath>
<https://franz.com#hasID> .
_:b271983AAx2      <http://www.w3.org/ns/shacl#resultSeverity>
<http://www.w3.org/ns/shacl#Violation> .
_:b271983AAx2
<http://www.w3.org/ns/shacl#sourceConstraintComponent>
<http://www.w3.org/ns/shacl#MaxCountConstraintComponent> .
_:b271983AAx2      <http://www.w3.org/ns/shacl#sourceShape>
_:b271983AAx3 .
_:b271983AAx1      <http://www.w3.org/ns/shacl#result>
_:b271983AAx4 .
_:b271983AAx4
<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
<http://www.w3.org/ns/shacl#ValidationResult> .
_:b271983AAx4      <http://www.w3.org/ns/shacl#focusNode>
<https://franz.com#emp002> .
_:b271983AAx4      <http://www.w3.org/ns/shacl#resultPath>
<https://franz.com#hasID> .
_:b271983AAx4      <http://www.w3.org/ns/shacl#resultSeverity>
<http://www.w3.org/ns/shacl#Violation> .
_:b271983AAx4
<http://www.w3.org/ns/shacl#sourceConstraintComponent>
<http://www.w3.org/ns/shacl#PatternConstraintComponent> .
_:b271983AAx4      <http://www.w3.org/ns/shacl#sourceShape>
_:b271983AAx3 .
_:b271983AAx4 <http://www.w3.org/ns/shacl#value> "00-56-3456"
.
_:b271983AAx1      <http://www.w3.org/ns/shacl#result>
_:b271983AAx5 .

```

```
_:b271983AAx5
<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
<http://www.w3.org/ns/shacl#ValidationResult> .
_:b271983AAx5      <http://www.w3.org/ns/shacl#focusNode>
<https://franz.com#emp003> .
_:b271983AAx5      <http://www.w3.org/ns/shacl#resultPath>
<https://franz.com#hasID> .
_:b271983AAx5      <http://www.w3.org/ns/shacl#resultSeverity>
<http://www.w3.org/ns/shacl#Violation> .
_:b271983AAx5
<http://www.w3.org/ns/shacl#sourceConstraintComponent>
<http://www.w3.org/ns/shacl#MinCountConstraintComponent> .
_:b271983AAx5      <http://www.w3.org/ns/shacl#sourceShape>
_:b271983AAx3 .
```

You can have the triples added to the repository by specifying the `--add-to-repo` option `true`.

In the plain output information is provided about how many data graphs are examined, how many shape graphs were specified and node shapes are found, and how many focus nodes are checked. If zero focus nodes are checked, that is likely not what you want and something has gone wrong. Here we mis-spell the name of the shape graph (`https://franz.com#shapes` instead of `https://franz.com#Shapes`) and get 0 focus nodes checked:

```
% bin/agtool shacl-validate --data-graph default --shapes-
graph https://franz.com#shapes shacl-repo-1
Validation report:           Conforms
Created:                     2019-06-28T10:34:22
Number of shapes graphs:     1
Number of data graphs:       1
Number of NodeShapes:        0
Number of focus nodes checked: 0
```

SPARQL integration

There are two sets of magic properties defined: one checks for basic conformance and the other produces validation reports as triples:

- `?valid franz:shaclConforms (?dataGraph [?shapesGraph])`
- `?valid franz:shaclFocusNodeConforms1 (?dataGraph ?nodeOrNodeCollection)`
- `?valid franz:shaclFocusNodeConforms2 (?dataGraph ?shapesGraph ?nodeOrNodeCollection)`
- `?valid franz:shaclShapeConforms1 (?dataGraph ?shapeOrShapeCollection [?nodeOrNodeCollection])`
- `?valid franz:shaclShapeConforms2 (?dataGraph ?shapesGraph ?shapeOrShapeCollection [?nodeOrNodeCollection])`
- `(?s ?p ?o) franz:shaclValidationReport (?dataGraph [?shapesGraph])`
- `(?s ?p ?o) franz:shaclFocusNodeValidationReport1 (?dataGraph ?nodeOrNodeCollection) .`
- `(?s ?p ?o) franz:shaclFocusNodeValidationReport2 (?dataGraph ?shapesGraph ?nodeOrNodeCollection) .`
- `(?s ?p ?o) franz:shaclShapeValidationReport1 (?dataGraph ?shapeOrShapeCollection [?nodeOrNodeCollection]) .`
- `(?s ?p ?o) franz:shaclShapeValidationReport2 (?dataGraph ?shapesGraph ?shapeOrShapeCollection [?nodeOrNodeCollection]) .`

In all of the above `?dataGraph` and `?shapesGraph` can be IRIs, the literal 'default', or a variable that is bound to a SPARQL collection (list or set) that was previously created with a function

like `https://franz.com/ns/allegrograph/6.5.0/fn#makeSPARQLList` or `https://franz.com/ns/allegrograph/6.5.0/fn#lookupRdfList`. If a collection is used, then the SHACL processor will create a temporary RDF merge of all of the graphs in it to produce the data graph or the shapes graph.

Similarly, `?shapeOrShapeCollection` and `?nodeOrNodeCollection` can be bound to an IRI or a SPARQL collection. If a collection is used, then it must be bound to a list of IRIs. The SHACL

processor will restrict validation to the shape(s) and focus node(s) (i.e. nodes that should be validated) specified.

The `shapesGraph` argument is optional in both of the `shaclConforms` and `shaclValidationReport` magic properties. If the `shapesGraph` is not specified, then the `shapesGraph` will be created by following triples in the `dataGraph` that use the `sh:shapesGraph` predicate. If there are no such triples, then the `shapesGraph` will be the same as the `dataGraph`.

For example, the following SPARQL expression

```
construct { ?s ?p ?o } where {  
  # form a collection of focusNodes  
  bind(<https://franz.com/ns/allegrograph/6.6.0/fn#makeSPARQLList>(  
    <http://Journal1/1942/Article25>,  
    <http://Journal1/1943>) as ?nodes)  
    (?s ?p ?o)  
<https://franz.com/ns/allegrograph/6.6.0/shaclShapeValidationReport1>  
  ('default' <ex://franz.com/documentShape1> ?nodes) .  
}
```

would use the default graph as the Data Graph and the Shapes Graph and then validate two focus nodes against the shape `<ex://franz.com/documentShape1>`.

SHACL Example

We build on our simple example above. Start with a fresh repository so triples from the simple example do not interfere with this example.

We start with a TriG file with various shapes defined on some classes.

```
@prefix sh: <http://www.w3.org/ns/shacl#> .  
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .  
@prefix fr: <https://franz.com#> .
```

```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
```

```
<https://franz.com#ShapesGraph> {
fr:EmployeeShape
  a sh:NodeShape ;
  sh:targetClass fr:Employee ;
  sh:property [
    ## Every employee must have exactly one ID
    sh:path fr:hasID ;
    sh:minCount 1 ;
    sh:maxCount 1 ;
    sh:datatype xsd:string ;
    sh:pattern "[0-9][0-9][0-9]-[0-9][0-9]-[0-9][0-9][0-9][0-9]$" ;
  ] ;
  sh:property [
    ## Every employee is a manager or a worker
    sh:path fr:employeeType ;
    sh:minCount 1 ;
    sh:maxCount 1 ;
    sh:datatype xsd:string ;
    sh:in ("Manager" "Worker") ;
  ] ;
  sh:property [
    ## If birthyear supplied, must be 2001 or before
    sh:path fr:birthYear ;
    sh:maxInclusive 2001 ;
    sh:datatype xsd:integer ;
  ] ;
  sh:property [
    ## Must have a title, may have more than one
    sh:path fr:hasTitle ;
    sh:datatype xsd:string ;
    sh:minCount 1 ;
  ] ;

  sh:or (
    ## The President does not have a supervisor
    [
      sh:path fr:hasTitle ;
```

```

        sh:hasValue "President" ;
    ]
    [
        ## Must have a supervisor
        sh:path fr:hasSupervisor ;
        sh:minCount 1 ;
        sh:maxCount 1 ;
        sh:class fr:Employee ;
    ]
) ;

sh:or (
    # Every employee must either have a wage or a salary
    [
        sh:path fr:hasSalary ;
        sh:datatype xsd:integer ;
        sh:minInclusive 3000 ;
        sh:minCount 1 ;
        sh:maxCount 1 ;
    ]
    [
        sh:path fr:hasWage ;
        sh:datatype xsd:decimal ;
        sh:minExclusive 15.00 ;
        sh:minCount 1 ;
        sh:maxCount 1 ;
    ]
)
.
}

```

This file says the following about instances of the class `fr:Employee`:

1. Every employee must have exactly one ID (object of `fr:hasID`), a string of the form `NNN-NN-NNNN` where the `Ns` are digits (this is the simple example requirement).
2. Every employee must have exactly one `fr:employeeType` triple with value either `"Manager"` or `"Worker"`.

3. Employees may have a `fr:birthYear` triple, and if so, the value must be 2001 or earlier.
4. Employees must have a `fr:hasTitle` and may have more than one.
5. All employees except the one with title "President" must have a supervisor (specified with `fr:hasSupervisor`).
6. Every employee must either have a wage (a decimal specifying hourly pay, greater than 15.00) or a salary (an integer specifying monthly pay, greater than or equal to 3000).

Here is some employee data:

```
@prefix fr: <https://franz.com#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
```

```
{
  fr:Employee
    a rdfs:Class .

  fr:emp001
    a fr:Employee ;
    fr:hasID "000-12-3456" ;
    fr:hasTitle "President" ;
    fr:employeeType "Manager" ;
    fr:birthYear "1953"^^xsd:integer ;
    fr:hasSalary "10000"^^xsd:integer .

  fr:emp002
    a fr:Employee ;
    fr:hasID "000-56-3456" ;
    fr:hasTitle "Foreman" ;
    fr:employeeType "Worker" ;
    fr:birthYear "1966"^^xsd:integer ;
    fr:hasSupervisor fr:emp003 ;
    fr:hasWage "20.20"^^xsd:decimal .

  fr:emp003
    a fr:Employee ;
```

```
fr:hasID "000-77-3232" ;
fr:hasTitle "Production Manager" ;
fr:employeeType "Manager" ;
fr:birthYear "1968"^^xsd:integer ;
fr:hasSupervisor fr:emp001 ;
fr:hasSalary "4000"^^xsd:integer .
```

```
fr:emp004
  a fr:Employee ;
  fr:hasID "000-88-3456" ;
  fr:hasTitle "Fitter" ;
  fr:employeeType "Worker" ;
  fr:birthYear "1979"^^xsd:integer ;
  fr:hasSupervisor fr:emp002 ;
  fr:hasWage "17.20"^^xsd:decimal .
```

```
fr:emp005
  a fr:Employee ;
  fr:hasID "000-99-3492" ;
  fr:hasTitle "Fitter" ;
  fr:employeeType "Worker" ;
  fr:birthYear "2000"^^xsd:integer ;
  fr:hasWage "17.20"^^xsd:decimal .
```

```
fr:emp006
  a fr:Employee ;
  fr:hasID "000-78-5592" ;
  fr:hasTitle "Filer" ;
  fr:employeeType "Intern" ;
  fr:birthYear "2003"^^xsd:integer ;
  fr:hasSupervisor fr:emp002 ;
  fr:hasWage "14.20"^^xsd:decimal .
```

```
fr:emp007
  a fr:Employee ;
  fr:hasID "000-77-3232" ;
  fr:hasTitle "Sales Manager" ;
  fr:hasTitle "Vice President" ;
  fr:employeeType "Manager" ;
  fr:birthYear "1962"^^xsd:integer ;
  fr:hasSupervisor fr:emp001 ;
```

```
    fr:hasSalary "7000"^^xsd:integer .  
}
```

Comparing these data with the requirements, we see these problems:

1. emp005 does not have a supervisor.
2. emp006 is pretty messed up, with (1) employeeType "Intern", not an allowed value, (2) a birthYear (2003) later than the required maximum of 2001, and (3) a wage (14.40) less than the minimum (15.00).

Otherwise the data seems OK.

We load these two TriG files into an empty repository (which we have named **shacl-repo-2**). We specify the default graph for the data and the `https://franz.com#ShapesGraph` for the shapes. (Though not required, it is a good idea to specify a graph for shape data as it makes it easy to delete and reload shapes while developing.) We have 101 triples, 49 data and 52 shape. Then we run **agtool shacl-validate**:

```
%    bin/agtool    shacl-validate    --shapes-graph  
https://franz.com#ShapesGraph --data-graph default shacl-  
repo-2
```

There are four violations, as expected, one for emp005 and three for emp006.

Validation report:	Does not conform
Created:	2019-07-03T11:35:27
Number of shapes graphs:	1
Number of data graphs:	1
Number of NodeShapes:	1
Number of focus nodes checked:	7

4 validation results:

Result:

Focus node:	<https://franz.com#emp005>
Value:	<https://franz.com#emp005>
Source Shape:	<https://franz.com#EmployeeShape>

Constraint Component:
<https://www.w3.org/ns/shacl#OrConstraintComponent>
Severity: <https://www.w3.org/ns/shacl#Violation>

Result:
Focus node: <https://franz.com#emp006>
Path: <https://franz.com#employeeType>
Value: "Intern"
Source Shape: _:b19D062B9x221

Constraint Component:
<http://www.w3.org/ns/shacl#InConstraintComponent>
Severity: <http://www.w3.org/ns/shacl#Violation>

Result:
Focus node: <https://franz.com#emp006>
Path: <https://franz.com#birthYear>
Value: "2003"^^<http://www.w3.org/2001/XMLSchema#integer>
Source Shape: _:b19D062B9x225

Constraint Component:
<http://www.w3.org/ns/shacl#MaxInclusiveConstraintComponent>
Severity: <http://www.w3.org/ns/shacl#Violation>

Result:
Focus node: <https://franz.com#emp006>
Value: <https://franz.com#emp006>
Source Shape: <https://franz.com#EmployeeShape>
Constraint Component:
<http://www.w3.org/ns/shacl#OrConstraintComponent>
Severity: <http://www.w3.org/ns/shacl#Violation>

Fixing the data is left as an exercise for the reader.

Turn Customer Service Calls

into Enterprise Knowledge Graphs

Franz's CEO, Jans Aasman's recent Destination CRM article:

The need for text analytics and speech recognition has broadened over the years, becoming more prevalent and essential in the sales, marketing, and customer service departments of various types of businesses and industries. The goal is simple for these contact center use cases: provide real-time assistance to human agents interacting with potential customers to close sales, initiate them, and increase customer satisfaction.

Until fairly recently, the rich array of unstructured data encompassing client texts, chats, and phone calls was obscured from contact centers and organizations due to the sheer arduousness of speech recognition and text analytics. When readily integrated into knowledge graphs, however, these same sources become some of the most credible for improving agent interactions and achieving business objectives.

Powered by the shrewd usage of organizational taxonomies, machine learning, natural language processing (NLP), and semantic search, knowledge graphs make speech recognition and text analytics immediately accessible, enabling real-time customer interactions that can maximize business objectives—and revenues.

Taxonomies

Taxonomies are the foundation of the knowledge graph approach to rapidly conveying results of speech recognition and text analytics for timely customer interactions. Agents need three types of information to optimize customer interactions: their personas (such as an executive or a purchase department representative, for example), their reasons for contacting

them, and their industries. Taxonomies are instrumental to performing these functions because they provide a hierarchy of relevant terms to organizations.

Read the full article at [Destination CRM](#)

Creating Explainable AI With Rules

Franz's CEO, Jans Aasman's recent Forbes article:

There's a fascinating dichotomy in artificial intelligence between statistics and rules, machine learning and expert systems. Newcomers to artificial intelligence (AI) regard machine learning as innately superior to brittle rules-based systems, while the history of this field reveals both rules and probabilistic learning are integral components of AI.

This fact is perhaps nowhere truer than in establishing explainable AI, which is central to the long-term business value of AI front-office use cases.

Granted, simple machine learning can automate backend processes. However, the full extent of deep learning or complex neural networks – which are much more accurate than basic machine learning – for mission-critical decision-making and action requires explainability.

Using rules (and rules-based systems) to explicate machine learning results creates explainable AI. Many of the far-reaching applications of AI at the enterprise level – deploying it to combat financial crimes, to predict an individual's immediate and long-term future in health care,

for example – require explainable AI that's fair, transparent and regulatory compliant.

Rules can explain machine learning results for these purposes and others.

Read the full article at Forbes