
CSVW Converters Documentation

Release 1.0.0

Rinke Hoekstra

Nov 16, 2020

Contents

1 Features & Limitations	3
2 Installation	5
2.1 Prerequisites	5
2.2 Installing with pip (preferred)	5
3 Usage	7
3.1 Generating a Skeleton Schema	8
3.2 Converting a CSV file	10
4 The Schema	15
4.1 Differences and Extensions	15
4.2 Short Overview	16
5 FAQ: Frequently Asked Questions	21
5.1 Commonly used Template Formatting	21
6 API Documentation	23
7 Indices and tables	25
7.1 Footnotes	25

This package is a comprehensive tool (CoW¹) for batch conversion of multiple datasets expressed in CSV. It uses a JSON schema expressed using an extended version of the CSVW standard, to convert CSV files to RDF in scalable fashion.

Instead of using the command line tool there is also the webservice [cattle](#), providing the same functionality that CoW provides without having to install it. CSV files can be uploaded to the service and a JSON schema will be created, using that JSON schema cattle is able to create a RDF structured graph. More information about cattle, including how to use it, can be found at: <https://github.com/CLARIAH/cattle>.

[CSV on the Web \(CSVW\)](#) is a W3C standard for metadata descriptions for tabular data. Typically, these data reside in CSV files. CSVW metadata is captured in `.csv-metadata.json` files that live alongside the CSV files that they describe. For instance, a CSV file called `data.csv` and its metadata `data.csv-metadata.json` would be hosted at:

```
http://example.com/data.csv  
http://example.com/data.csv-metadata.json
```

Another feature of CSVW is that it allows the specification of a mapping (or interpretation) of values in the CSV in terms of RDF. The `tableSchema` element in CSVW files defines per column what its properties should be, but may also define custom mappings to e.g. URIs in RDF.

Interestingly, the JSON format used by CSVW metadata is an [extension of the JSON-LD specification](#), a JSON-based serialization for Linked Data. As a consequence of this, the CSVW metadata can be directly attached (as provenance) to the RDF resulting from a CSVW-based conversion.

This is exactly what the CoW converter does.

The rest of this documentation will be fairly technical, for some hands-on examples you can take a look at the [Wiki](#).

¹ COW: **C**SV **O**n the **W**eb.

CHAPTER 1

Features & Limitations

Compared to the CSVW specification, the converter has a number of limitations and extra features. These are:

1. CoW *does not* perform any schema checking, and ignores any and all parts of the CSVW Specification that are not directly needed for the RDF conversion.
2. CoW extends the CSVW specification in several ways:
 - Advanced formatting of URLs and values
 - Dealing with multiple null values and null values for one or more other columns.
 - Simple SKOS support (generating collections and schemes)
 - Optionally skipping/not skipping empty cells
 - A default set of namespace prefixes
3. CoW does some smart guessing:
 - Determining file encoding
 - Determining the delimiter
 - Generating a skeleton schema for any CSV file (see [here](#))
4. CoW produces extensive provenance:
 - Converted data is encapsulated in a Nanopublication
 - The original CSVW schema is encapsulated in the *np:hasProvenance* graph associated with the nanopublication.

CHAPTER 2

Installation

2.1 Prerequisites

- Python 3.8 (installed on most systems)
- pip3
- virtualenv (simply *pip3 install virtualenv*)²

2.2 Installing with pip (preferred)

Open up a terminal (or Command Prompt when you are using Windows) and instantiate a virtual Python environment:

```
virtualenv .
```

Activate the virtual environment:

```
source bin/activate
```

Install CoW in the new environment:

```
pip3 install cow_csvw
```

To upgrade a previously installed version of CoW, do:

```
pip3 install --upgrade cow_csvw
```

(you might need permissions if you're installing outside a virtualenv). To check the version currently installed:

```
cow_tool --version
```

To get help:

² These instructions use `virtualenv` but you can also install all packages globally, or use an alternative such as `conda`.

cow_tool

CHAPTER 3

Usage

The primary command line script for CSVW-based conversion is `cow_tool`. It can be used for two tasks:

1. Generating a *skeleton CSVW JSON-Schema* for a specific CSV file.
2. Using such a schema to *convert a CSV file to RDF* (in NQuads format)

General usage instructions can be obtained by running `cow_tool -h`:

```
usage: cow_tool [-h] [--dataset DATASET] [--delimiter DELIMITER]
                [--quotechar QUOTECHAR] [--processes PROCESSES]
                [--chunksize CHUNKSIZE] [--base BASE]
                {convert,build} file [file ...]
```

The table below gives a brief description of each of these options.

Table 1: Commandline options for `cow_tool`

Option	Explanation
<code>dataset</code>	Specifies the name of the dataset, if it is different from the filename with the <code>.csv</code> extension stripped.
<code>delimiter</code>	Forces the use of a specific delimiter when parsing the CSV file (only used with <code>build</code> option)
<code>quotechar</code>	Forces the use of a specific quote character (default is <code>"</code> , only used with <code>build</code> option)
<code>encoding</code>	Forces the use of a specific file encoding when parsing the CSV file (only used with <code>build</code> option)
<code>processes</code>	Specifies the number of parallel processes to use when converting a CSV file (default is 4)
<code>chunksize</code>	Specifies the number of lines that will be passed to each process (default is 5000)
<code>base</code>	The base for URIs generated with the schema (only used with <code>build</code> option, the default is <code>http://data.socialhistory.org</code>)
<code>{convert, build}</code>	The <code>convert</code> option triggers a conversion to RDF for the files specified in <code>file [file ...]</code> . The <code>build</code> option generates a skeleton JSON schema for the files specified.
<code>file [file ...]</code>	A list of files to be converted (or “built”); any unix-style wildcards are allowed.

3.1 Generating a Skeleton Schema

Since JSON is a rather verbose language, and we currently do not have a convenient UI for constructing CSVW schema files, CoW allows you to generate a skeleton schema for any CSV file.

Suppose you want to build a skeleton schema for a file `imf_gdppc.csv` (from⁴) that looks like:

```
1 Rank;Country;GDP_Per_Capita
2 1;Qatar;131,063
3 2;Luxembourg;104,906
4 3;Macau;96,832
5 4;Singapore;90,249
6 5;Brunei Darussalam;83,513
7 6;Kuwait;72,675
8 7;Ireland;72,524
9 8;Norway;70,645
```

Make sure you have your virtual environment enabled (if applicable), and run:

```
cow_tool build imf_gdppc.csv --base=http://example.com/resource
```

The `--base` option specifies the base for all URIs generated through the schema. This is `https://iisg.amsterdam/` by default (see <http://datalegend.net>)

This will generate a file called `imf_gdppc.csv-metadata.json` with the following contents:

```
1 {
2   "dialect": {
3     "quoteChar": "\",
4     "delimiter": ";",
5     "encoding": "ascii"
6   },
7   "dcat:keyword": [],
8   "dc:license": {
9     "@id": "http://opendefinition.org/licenses/cc-by/"
10   },
11   "dc:publisher": {
12     "schema:name": "CLARIAH Structured Data Hub - Datalegend",
13     "schema:url": {
14       "@id": "http://datalegend.net"
15     }
16   },
17   "url": "imf_gdppc.csv",
18   "@context": [
19     "http://csvw.clariah-sdh.eculture.labs.vu.nl/csvw.json",
20     {
21       "@base": "http://example.com/resource/",
22       "@language": "en"
23     },
24     {
25       "@owl": "http://www.w3.org/2002/07/owl#",
26       "napp-eng81": "https://iisg.amsterdam/napp/dataset/englandwales1881/",
27       "dbo": "http://dbpedia.org/ontology/",
28       "clioctr": "https://iisg.amsterdam/clio/country/",
29       "hisclass": "https://iisg.amsterdam/hisclass/",
30       "hisco-product": "https://iisg.amsterdam/hisco/product/",
```

(continues on next page)

⁴ https://en.wikipedia.org/wiki/List_of_countries_by_GDP_%28PPP%29_per_capita

(continued from previous page)

```

31 "ldp": "http://www.w3.org/ns/ldp#",
32 "clio": "https://iisg.amsterdam/clio/",
33 "occhisco": "https://iisg.amsterdam/napp/OCCHISCO/",
34 "dbr": "http://dbpedia.org/resource/",
35 "skos": "http://www.w3.org/2004/02/skos/core#",
36 "xml": "http://www.w3.org/XML/1998/namespace/",
37 "sdmx-concept": "http://purl.org/linked-data/sdmx/2009/concept#",
38 "napp": "https://iisg.amsterdam/napp/",
39 "prov": "http://www.w3.org/ns/prov#",
40 "sdmx-code": "http://purl.org/linked-data/sdmx/2009/code#",
41 "napp-can91": "https://iisg.amsterdam/napp/dataset/canada1891/",
42 "hiscam": "https://iisg.amsterdam/hiscam/",
43 "dbpedia": "http://dbpedia.org/resource/",
44 "np": "http://www.nanopub.org/nschema#",
45 "hisclass5": "https://iisg.amsterdam/hisclass5/",
46 "canfam-auke": "https://iisg.amsterdam/canfam/auke/",
47 "dcterms": "http://purl.org/dc/terms/",
48 "schema": "http://schema.org/",
49 "foaf": "http://xmlns.com/foaf/0.1/",
50 "sdv": "http://example.com/resource/vocab/",
51 "hisco": "https://iisg.amsterdam/hisco/",
52 "bibo": "http://purl.org/ontology/bibo/",
53 "sdmx-dimension": "http://purl.org/linked-data/sdmx/2009/dimension#",
54 "hsn": "https://iisg.amsterdam/hsn2013a/",
55 "dc": "http://purl.org/dc/terms/",
56 "hisco-relation": "https://iisg.amsterdam/hisco/relation/",
57 "hisco-status": "https://iisg.amsterdam/hisco/status/",
58 "dbp": "http://dbpedia.org/property/",
59 "clioprop": "https://iisg.amsterdam/clio/property/",
60 "csvw": "http://www.w3.org/ns/csvw#",
61 "clioind": "https://iisg.amsterdam/clio/indicator/",
62 "dc11": "http://purl.org/dc/elements/1.1/",
63 "qb": "http://purl.org/linked-data/cube#",
64 "canfam-dimension": "http://data.socialhistory.org/vocab/canfam/dimension/",
65 "rdfs": "http://www.w3.org/2000/01/rdf-schema#",
66 "canfam": "https://iisg.amsterdam/canfam/dataset/canada1901/",
67 "napp-sct81": "https://iisg.amsterdam/napp/dataset/scotland1881/",
68 "sdmx-measure": "http://purl.org/linked-data/sdmx/2009/measure#",
69 "rdf": "http://www.w3.org/1999/02/22-rdf-syntax-ns#",
70 "sdr": "http://example.com/resource/",
71 "xsd": "http://www.w3.org/2001/XMLSchema#",
72 "time": "http://www.w3.org/2006/time#",
73 "napp-dimension": "http://data.socialhistory.org/vocab/napp/dimension/"
74 }
75 ],
76 "dc:title": "imf_gdppc.csv",
77 "@id": "http://example.com/resource/imf_gdppc.csv",
78 "dc:modified": {
79 "@value": "2018-11-14",
80 "@type": "xsd:date"
81 },
82 "tableSchema": {
83 "aboutUrl": "{_row}",
84 "primaryKey": "Rank",
85 "columns": [
86 {
87 "datatype": "string",

```

(continues on next page)

(continued from previous page)

```

88 "titles": [
89   "Rank"
90 ],
91 "@id": "http://example.com/resource/imf_gdppc.csv/column/Rank",
92 "name": "Rank",
93 "dc:description": "Rank"
94 },
95 {
96   "datatype": "string",
97   "titles": [
98     "Country"
99   ],
100  "@id": "http://example.com/resource/imf_gdppc.csv/column/Country",
101  "name": "Country",
102  "dc:description": "Country"
103 },
104 {
105   "datatype": "string",
106   "titles": [
107     "GDP_Per_Capita"
108   ],
109   "@id": "http://example.com/resource/imf_gdppc.csv/column/GDP_Per_Capita",
110   "name": "GDP_Per_Capita",
111   "dc:description": "GDP_Per_Capita"
112 }
113 ]
114 }
115 }
```

The exact meaning of this structure is explained in [the section below](#).

3.2 Converting a CSV file

If we now want to convert our example file `imf_gdppc.csv`, you first make sure you have your virtual environment enabled (if applicable), and run:

```
cow_tool convert imf_gdppc.csv
```

This will produce a file `imf_gdppc.csv.nq` that holds an NQuads serialization of the RDF.

This is also the preferred method for converting multiple files at the same time. For instance, if you want to convert *all* CSV files in a specific directory, simply use unix-style wildcards:

```
cow_tool convert /path/to/some/directory/*.csv
```

Going back to our running example, the resulting RDF will be serialized as N-Quads. This is a computer friendly but not so much human friendly serialization so for the benefit of (human) readability below the RDF will be represented in the TriG serialization:

```

1 @prefix ns1: <http://www.w3.org/ns/prov#> .
2 @prefix ns2: <http://www.w3.org/ns/csvw#> .
3 @prefix ns3: <http://schema.org/> .
4 @prefix ns4: <http://purl.org/dc/terms/> .
5 @prefix ns5: <urn:uuid:5> .
```

(continues on next page)

(continued from previous page)

```

6 @prefix ns6: <http://www.nanopub.org/nschema#> .
7 @prefix ns7: <https://iisg.amsterdam/vocab/> .
8 @prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
9 @prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
10 @prefix xml: <http://www.w3.org/XML/1998/namespace> .
11 @prefix xsd: <http://www.w3.org/2001/XMLSchema#> .

12 <https://iisg.amsterdam/imf_gdppc/pubinfo/48422b27/2018-11-14T10:59> {
13   <https://iisg.amsterdam/imf_gdppc/nanopublication/48422b27/2018-11-14T10:59>_
14   ↳ ns1:generatedAtTime "2018-11-14T10:59:00"^^xsd:dateTime ;
15     ns1:wasGeneratedBy <https://github.com/CLARIAH/wp4-converters> .
16 }

17 <https://iisg.amsterdam/imf_gdppc/provenance/48422b27/2018-11-14T10:59> {
18   <https://iisg.amsterdam/imf_gdppc/assertion/48422b27/2018-11-14T10:59>_
19   ↳ ns1:generatedAtTime "2018-11-14T10:59:00"^^xsd:dateTime ;
20     ns1:wasDerivedFrom <http://example.com/resource/imf_gdppc.csv>,
21       <https://iisg.amsterdam/48422b27cba4a0e68c9c66d0f7ca614ec688dfcb> .
22
23   <http://example.com/resource/__row__> ns1:wasDerivedFrom "http://example.com/
24   ↳ resource/{_row}"^^xsd:string .

25   <http://example.com/resource/imf_gdppc.csv> ns4:license <http://opendefinition.
26   ↳ org/licenses/cc-by/> ;
27     ns4:modified "2018-11-14"^^xsd:date ;
28     ns4:publisher [ ns3:name "CLARIAH Structured Data Hub - Datalegend"@en ;
29       ns3:url <http://datalegend.net/> ] ;
30     ns4:title "imf_gdppc.csv"@en ;
31     ns2:dialect [ ns2:delimiter ";" ;
32       ns2:encoding "ascii" ;
33       ns2:quoteChar "\"" ] ;
34     ns2:tableSchema [ ns2:aboutUrl <http://example.com/resource/__row__> ;
35       ns2:column ( <http://example.com/resource/imf_gdppc.csv/column/Rank>
36         ↳ <http://example.com/resource/imf_gdppc.csv/column/Country> <http://example.com/
37         ↳ resource/imf_gdppc.csv/column/GDP_Per_Capita> ) ;
38       ns2:primaryKey "Rank" ] ;
39     ns2:url "imf_gdppc.csv"^^xsd:anyURI .
40
41   <http://example.com/resource/imf_gdppc.csv/column/Country> ns4:description
42   ↳ "Country"@en ;
43     ns2:datatype xsd:string ;
44     ns2:name "Country" ;
45     ns2:title "Country"@en .

46   <http://example.com/resource/imf_gdppc.csv/column/GDP_Per_Capita> ns4:description
47   ↳ "GDP_Per_Capita"@en ;
48     ns2:datatype xsd:string ;
49     ns2:name "GDP_Per_Capita" ;
50     ns2:title "GDP_Per_Capita"@en .

51   <http://example.com/resource/imf_gdppc.csv/column/Rank> ns4:description "Rank"@en_
52   ↳ ;
53     ns2:datatype xsd:string ;
54     ns2:name "Rank" ;
55     ns2:title "Rank"@en .
}

```

(continues on next page)

(continued from previous page)

```

54 ns5:db490c7-50c3-4ad6-b0df-d48fe3dfa984 {
55     <https://iisg.amsterdam/48422b27cba4a0e68c9c66d0f7ca614ec688dfcb> ns7:path "/tmp/
56     ↵V2RY7QULW9/web_interface/91a7c0a271826cf3e7e5b470dfd5e345/imf_gdppc.csv"^^
57     ↵xsd:string ;
58     ns7:sha1_hash "48422b27cba4a0e68c9c66d0f7ca614ec688dfcb"^^xsd:string .
59
60     <https://iisg.amsterdam/imf_gdppc/nanopublication/48422b27/2018-11-14T10:59> a_
61     ↵ns6:Nanopublication ;
62     ns6:hasAssertion <https://iisg.amsterdam/imf_gdppc/assertion/48422b27/2018-11-
63     ↵14T10:59> ;
64     ns6:hasProvenance <https://iisg.amsterdam/imf_gdppc/provenance/48422b27/2018-
65     ↵11-14T10:59> ;
66     ns6:hasPublicationInfo <https://iisg.amsterdam/imf_gdppc/pubinfo/48422b27/
67     ↵2018-11-14T10:59> .
68
69     <https://iisg.amsterdam/imf_gdppc/assertion/48422b27/2018-11-14T10:59> a_
70     ↵ns6:Assertion .
71
72     <https://iisg.amsterdam/imf_gdppc/provenance/48422b27/2018-11-14T10:59> a_
73     ↵ns6:Provenance .
74
75     <https://iisg.amsterdam/imf_gdppc/pubinfo/48422b27/2018-11-14T10:59> a_
76     ↵ns6:PublicationInfo .
77 }
78
79 <https://iisg.amsterdam/imf_gdppc/assertion/48422b27/2018-11-14T10:59> {
80     <http://example.com/resource/0> ns7:Country "Qatar"^^xsd:string ;
81     ns7:GDP_Per_Capita "131,063"^^xsd:string ;
82     ns7:Rank "1"^^xsd:string .
83
84     <http://example.com/resource/1> ns7:Country "Luxembourg"^^xsd:string ;
85     ns7:GDP_Per_Capita "104,906"^^xsd:string ;
86     ns7:Rank "2"^^xsd:string .
87
88     <http://example.com/resource/2> ns7:Country "Macau"^^xsd:string ;
89     ns7:GDP_Per_Capita "96,832"^^xsd:string ;
90     ns7:Rank "3"^^xsd:string .
91
92     <http://example.com/resource/3> ns7:Country "Singapore"^^xsd:string ;
93     ns7:GDP_Per_Capita "90,249"^^xsd:string ;
94     ns7:Rank "4"^^xsd:string .
95
96     <http://example.com/resource/4> ns7:Country "Brunei Darussalam"^^xsd:string ;
97     ns7:GDP_Per_Capita "83,513"^^xsd:string ;
98     ns7:Rank "5"^^xsd:string .
99
100    <http://example.com/resource/5> ns7:Country "Kuwait"^^xsd:string ;
101    ns7:GDP_Per_Capita "72,675"^^xsd:string ;
102    ns7:Rank "6"^^xsd:string .
103
104    <http://example.com/resource/6> ns7:Country "Ireland"^^xsd:string ;
105    ns7:GDP_Per_Capita "72,524"^^xsd:string ;
106    ns7:Rank "7"^^xsd:string .
107
108    <http://example.com/resource/7> ns7:Country "Norway"^^xsd:string ;
109    ns7:GDP_Per_Capita "70,645"^^xsd:string ;
110    ns7:Rank "8"^^xsd:string .

```

(continues on next page)

(continued from previous page)

102

}

What does this mean?

- Everything in https://iisg.amsterdam/imf_gdppc/provenance/48422b27/2018-11-14T10:59 is the RDF representation of the CSVW JSON schema.
- Everything in https://iisg.amsterdam/imf_gdppc/assertion/48422b27/2018-11-14T10:59 is the RDF representation of the CSV file.

Since the global `aboutUrl` is set to `{_row}`, every row is represented in RDF as a resource with the base URI concatenated with the row number. The column names are used as predicates to relate the row resource to a string literal representation of the value of a cell in that row.

- The graph `ns5:db490c7-50c3-4ad6-b0df-d48fe3dfa984` is the default graph that contains the Nanopublication.

CHAPTER 4

The Schema

The CoW converter uses the CSWV standard syntax for defining mappings from CSV to RDF graphs. These mappings are all defined in the `tableSchema` dictionary. For a full reference of the things you can do, we refer to the [CSV on the Web \(CSWV\)](#) specification and in particular to the document on [Generating RDF from Tabular Data on the Web](#).

Important: CoW does not purport to implement the full CSWV specification, nor has it been tested against the [official test suite](#). In fact, CoW extends and deviates from the CSWV specification in several important ways.

We document the most important differences in the section below, and give a [short overview](#) of how schemas can be defined.

4.1 Differences and Extensions

1. While CSWV allows only for simple references to values in a column using the curly-brackets syntax (e.g. `{name}` to refer to the value of the name column at the current row), CoW interprets the strings containing these references in two ways:

1. as Python Format Strings, and
2. as Jinja2 Templates

This allows for very elaborate operations on row contents (e.g. containing conditionals, loops, and string operations).³.

2. CSWV allows only to specify a single `null` value for a column; when the cell in that column is equal to the `null` value, it is ignored for RDF conversion. CoW extends the CSWV treatment of `null` values in two ways:

1. multiple potential `null` values for a column, expressed as a JSON list, and
2. conditional on values in *another* column, as a JSON-LD list (using the `@list` keyword)
3. CoW allows the use of `csvw:collectionUrl` and `csvw:schemeUrl` on column specifications. This will automatically cast the value for `valueUrl` to a `skos:Concept`, and adds it to the collection or scheme respectively indicated by these urls using a `skos:member` or `skos:inScheme` predicate.

³ In the future we may enable the Jinja2 plugin mechanism. This will allow running custom Python functions as filters over values.

4. By default CoW skips cells that are empty (as per the CSVW specification), setting the `csvw:parseOnEmpty` attribute to `true` overrides this setting. This is useful when an empty cell has a specific meaning.
5. Column specifications with a `xsd:anyURI` datatype are converted to proper URIs rather than Literals with the `xsd:anyURI` datatype. This allows for conditionally generating URIs across multiple namespaces using Jinja2 templates, see [issue #13](#).
6. Column specifications in CoW should have a JSON-LD style `@id` attribute. This ensures that all predicates generated through the conversion are linked back to the RDF representation of the CSVW JSON schema that informed the conversion.
7. CoW converts column names to valid Python dictionary keys. In general this means that spaces in column names will be replaced with underscores.
8. For convenience, CoW uses a default set of namespaces, specified in the `src/converter/namespaces.yaml` file, that will be used to interpret namespace prefix use in the JSON schema. Any namespace prefixes defined in the JSON schema will override the default ones.

4.2 Short Overview

A very simple `tableSchema` may have the following structure:

```
1 "tableSchema": {  
2     "aboutUrl": "{_row}",  
3     "primaryKey": "Rank",  
4     "columns": [  
5         {  
6             "@id": "http://example.com/resource/imf_gdppc.csv/column/Rank",  
7             "dc:description": "Rank",  
8             "datatype": "string",  
9             "name": "Rank"  
10        }  
11    ]  
12 }
```

For the conversion to RDF, only the `aboutUrl` and `columns` attributes are of importance.

4.2.1 `aboutUrl`

The `aboutUrl` attribute defines a template for all URIs that occur in the *subject* position of triples generated by the converter. It may appear in the `tableSchema` or in one of the `columns`. If defined in the `tableSchema`, it acts as a *global* template that may be overridden by individual columns.

We explain URL template expansion [here](#).

4.2.2 `columns`

The `columns` array defines a schema for each column, and any additional virtual columns. The distinction between the two is important, as non-virtual columns must actually be present in the CSV (schema compliance) while virtual columns only instruct the conversion to RDF.

In the schema above, we state that the column identifiable with the `name` `Rank` specifies a literal value, with the `datatype` of `string` (a shorthand for `xsd:string`). The `titles` array gives a number of alternative

4.2.3 Column Attributes

Every column is a dictionary that may have the following attributes.

Table 1: Attributes usable in column specifications

Attribute	Explanation
name	Specifies the column to which this column specification applies. If no <code>propertyUrl</code> is defined on the column, the value for <code>name</code> will be used to generate the URL for the <i>predicate</i> position of the triple generated.
virtual	If set to <code>true</code> , the column specification is not taken into account when validating a CSV file against this schema.
aboutUrl	Overrides the <i>global</i> <code>aboutUrl</code> template defined for the schema. This template will be used to generate the <i>subject</i> URL of the triple.
valueUrl	If present, this template will be used to generate the <i>object</i> URL of the triple. Otherwise, the value for <code>name</code> is used to retrieve the value for that cell, to generate a URL.
datatype	Specifies that this column should result in a triple where the <i>object</i> is a Literal with the datatype specified here (for common XML Schema datatypes, it is possible to drop the <code>xsd:</code> prefix). The value of the literal is then the value of the cell in this row indicated by the value of <code>name</code> . Special case: when the datatype is <code>xsd:anyURI</code> CoW creates a URI rather than a literal value.
csvw:value	Specifies that this column should result in a triple where the <i>object</i> is a Literal with the default <code>xsd:string</code> datatype (unless otherwise specified in the <code>datatype</code> attribute). The literal value for this cell is determined by applying the <code>ref::template expansion <template-expansion></code> rule to this row. Can only be used in <code>virtual</code> columns.
csvw:process	When set to <code>true</code> , specifies that this column should be processed even when the cell corresponding to this column in this row is empty.
null	Specifies that this template does not apply if the cell in this column in this row corresponds to the value specified here. Can take a single value (as per the CSVW spec) or an array of values.
lang	Specifies the language tag for the literal in the <i>object</i> position, but only if the <code>datatype</code> is set to be <code>string</code> .
collectionUrl	Specifies that the <code>valueUrl</code> (or equivalent) should be of type <code>skos:Concept</code> and that it is a <code>skos:member</code> of the URL generated by applying the <code>collectionUrl</code> template.
schemeUrl	Specifies that the <code>valueUrl</code> (or equivalent) should be of type <code>skos:Concept</code> and that it is <code>skos:inScheme</code> the URL generated by applying the <code>schemeUrl</code> template.

4.2.4 Template Expansion with Jinja2 templates and Python format strings

When a CSV file is processed, CoW does this row by row in the file, producing a dictionary where key/value pairs correspond to column headers and the value of the cell. So for:

```
Rank;Country;GDP_Per_Capita
1;Qatar;131063
```

the first row becomes⁵

```
row = { 'Rank': 1, 'Country': 'Qatar', 'GDP_Per_Capita': 131063}
```

For each row, CoW then applies each column definition in the `columns` array in the JSON-LD file (i.e. which does not have to mean each column in the CSV file).

The URL templates in the attributes `aboutUrl`, `propertyUrl`, `valueUrl`, and the regular template in the `csvw:value` are used to generate URLs and Literal values from the values of the cells in a specific row.

⁵ Assuming that you have the proper locale settings that instructs Python to interpret the comma as a thousands separator.

The values for the URL templates that the parser receives are *interpreted as URLs*. This means that they are expanded relative to the @base URI of the CSVW JSON schema file, unless they are explicitly preceded by a defined namespace prefix.

The names of Jinja2 or Python formatting field names should correspond to the keys of the dictionary (i.e. to the column names). CoW supports a special CSVW field name `_row` that inserts the row number. This means that our row now becomes:

```
row = {'Rank': 1, 'Country': 'Qatar', 'GDP_Per_Capita': 131063, '_row': 1}
```

With this preparation of the row data the template expansion can begin. CoW always first applies: * the Jinja2 template (see documentation), * and then the Python format strings (see documentation).

For instance (assuming a @base of `http://example.com/`), we define an `aboutUrl` with the special `_row` key as a Python string formatting field name, and `Country` as a Jinja2 field name:

```
"aboutUrl": "{_row}/{Country}"
```

the JSON-LD parser interprets the value for `aboutUrl` as the following URI:

```
"http://example.com/{_row}/{Country}"
```

we then apply the Jinja2 formatting (`Template("http://example.com/{_row}{{Country}}").render(**row)`):

```
"http://example.com/_row/Qatar"
```

followed by the Python formatting (`"http://example.com/{_row}/{Country}"".format(**row)`):

```
"http://example.com/1/Qatar"
```

For `csvw:value` attributes this works similarly, with the exception that the JSON-LD parser will not interpret these fields as URIs:

```
"csvw:value": "{_row}/{Country}"
```

is parsed as:

```
"{_row}/{Country}"
```

This means that one can use Jinja2 conditional formatting on `csvw:value` attributes in combination with an `xsd:anyURI` value for `datatype` to generate custom URIs that do not fit within a defined namespace.

Jinja2 is a very expressive templating language. To give a small example, we could define a virtual column that allows us to specify whether a country is `http://example.com/rich` or `http://example.com/poor` depending on whether the GDP is over 100k.

Our virtual column may look as follows:

```
1 {
2     "virtual": "true",
3     "aboutUrl": "{Country}",
4     "propertyUrl": "rdf:type",
5     "valueUrl": "{% if GDP_Per_Capita > 100000 %}rich{% else %}poor{% endif %}"
6 }
```

This will produce, for Qatar and Singapore, the respective triples:

```
<http://example.com/Qatar>    rdf:type <http://example.com/rich> .  
<http://example.com/Singapore> rdf:type <http://example.com/poor> .
```

If you happen to be a bit experienced with the Python3 or ipython shell, then you could also quickly test Jinja templates like so:

```
1 from jinja2 import Template  
2 my_jinja_template = "{% if GDP_Per_Capita > 100000 %}rich{% else %}poor{% endif %}"  
3 row = {'Rank': 1, 'Country': 'Qatar', 'GDP_Per_Capita': 131063}  
4 Template(my_jinja_template).render(row)  
5 # returns 'rich'
```


CHAPTER 5

FAQ: Frequently Asked Questions

Please refer to our [wiki](#) for questions on specific topics.

5.1 Commonly used Template Formatting

- Leading zeroes: `{} '%05d' | format(variable|int) {}`, where 5 is the number of digits to fill up to.
 - If-else statements: `{} % if conditional_variable=="something" {} value_if {} else {} value_else {} endif {}.`
 - Convert to string and concatenate: `{}{{variable ~ 'string'}}{}`, e.g. if variable has value “Hello” then the result would be “Hello string”. Note the double braces.
 - Arithmetic: use double braces and cast as numeric first, e.g. `{}{{variable|float() * 1000}}{}`.
 - Lowercase, uppercase, etc.: `{}{{variable|lower()}}``. Note the double brace.
 - String slices: `{}{{variable[n:m]}}{}` as described [here](#).
-

CHAPTER 6

API Documentation

- code

CHAPTER 7

Indices and tables

- genindex
- modindex
- search

7.1 Footnotes