



## Position Paper

# Harmonise and integrate heterogeneous areal data with the R package `arealDB`

Steffen Ehrmann<sup>a,d,\*</sup>, Ralf Seppelt<sup>b,c</sup>, Carsten Meyer<sup>a,c,d,\*\*</sup>

<sup>a</sup> German Centre for Integrative Biodiversity Research (iDiv) Halle-Jena-Leipzig, Deutscher Platz 5e, 04103, Leipzig, Germany

<sup>b</sup> UFZ – Helmholtz Centre for Environmental Research, Leipzig, Department Computational Landscape Ecology, Permoserstraße 15, 04318, Leipzig, Germany

<sup>c</sup> Institute of Geosciences and Geography, Martin Luther University Halle-Wittenberg, Halle (Saale), Germany

<sup>d</sup> Institute of Biology, Leipzig University, Leipzig, Germany

## ARTICLE INFO

## Keywords:

Interoperability  
Census data  
Indicator data  
Polygon data  
Data warehouse  
Provenance documentation

## ABSTRACT

Many relevant applications in the environmental and socioeconomic sciences use areal data, such as biodiversity checklists, agricultural statistics, or socioeconomic surveys. For applications that surpass the spatial, temporal or thematic scope of any single data source, data must be integrated from several heterogeneous sources. Inconsistent concepts, definitions, or messy data tables make this a tedious and error-prone process. To date, a dedicated tool to address these challenges is still lacking.

Here, we introduce the R package `arealDB` that integrates heterogeneous areal data and associated geometries into a consistent database, in an easy-to-use workflow. It is useful for harmonising language and semantics of variables, relating data to geometries, and documenting metadata and provenance. We illustrate the functionality by integrating two disparate datasets (Brazil, USA) on the harvested area of soybean. `arealDB` promises quality-improvements to downstream scientific, monitoring, and management applications but also substantial time-savings to database collation efforts.

## 1. Introduction

Areal data are an essential data type in many socioeconomic applications, from visualising characteristics of human populations, to assessing trade statistics or documenting land ownership. They are increasingly used to analyse various environmental variables such as global biodiversity patterns. Areal data are an everyday communication tool in civil society, where they play an important role as illustrative maps in news or education media. Generally, any phenomenon at the level of finite spatial polygons (called geometries henceforth) is recorded as areal data, irrespective of the domain.

Areal data are typically collated and curated by a diverse set of actors, from small and focused projects that build datasets around their particular observations, for example in nature conservation, to national statistical agencies or intergovernmental organisations such as the World Bank Group<sup>1</sup>, the Food and Agriculture Organization<sup>2</sup> and others. The data are presented in formats, arrangements, languages and with

definitions that are primarily adapted to a specific purpose, resulting in many distinct datasets that are by default not interoperable. However, many important downstream applications and analyses surpass the spatial, temporal or thematic scope of any unique data source or organization, and thus rely on combining data from multiple different sources (Otto et al., 2015).

Integrating areal data across many sources comes with a large number of challenges (Table 1) that would, if not addressed, affect database consistency and potentially bias downstream analyses. For example, data tables that refer to the same areas must match one another, both spatially and lexically (Du et al., 2013), for databases to be consistent. Besides, terms that emerge from different languages or ontologies must correspond to one another, so that a variable that has different names in different source data does in fact match (De Giacomo et al., 2018). Moreover, alternative data that describe disputed areas or that have been recorded by different institutions should be acknowledged to avoid bias due to erroneous political assumptions, and final

\* Corresponding author. German Centre for Integrative Biodiversity Research (iDiv) Halle-Jena-Leipzig, Deutscher Platz 5e, 04103, Leipzig, Germany.

\*\* Corresponding author. German Centre for Integrative Biodiversity Research (iDiv) Halle-Jena-Leipzig, Deutscher Platz 5e, 04103, Leipzig, Germany.

E-mail addresses: [steffen.ehrmann@idiv.de](mailto:steffen.ehrmann@idiv.de) (S. Ehrmann), [carsten.meyer@idiv.de](mailto:carsten.meyer@idiv.de) (C. Meyer).

<sup>1</sup> <https://data.worldbank.org/>.

<sup>2</sup> <http://www.fao.org/faostat>.

outputs should be appropriately documented (Henzen et al., 2013).

Several of the challenges in integrating areal data can be addressed individually with specialised tools. For instance, a so-called ETL procedure (extract, transform, load) is typically used in data warehousing, where data from different sources are integrated into a single database (Baumer 2017, Debroy et al., 2018). Moreover, geometries can be matched with GIS software such as QGIS, and ontologies that describe and relate distinct concepts can be created with the software Protégé (Horridge, 2011). Finally, data can be made available via so-called Spatial Data Infrastructures (van den Brink et al., 2017) and data values of the resulting database can be cleaned or validated with R-packages such as dataMaid (Petersen and Ekstrøm, 2019). However, some specific challenges, such as to document input metadata or provenance or to reshape messy data (Wickham, 2014) are typically solved in non-standardised ways (if at all) via custom scripts or macros that are developed for individual use-cases. Oftentimes, important steps of database management are even carried out manually, e.g., by comparing information visually and entering the data by hand into an Excel file.

Notwithstanding the existence of many specialised tools, in practice, their application is not always trivial and may require specific expert knowledge or expensive proprietary software (Debroy et al., 2018). Moreover, none of the existing tools can address the full range of typical problems, so that combinations of independent tools are needed, which comes with further issues of interoperability of the specialised tools. This complexity increases resource and time requirements for a comprehensive workflow, hindering successful implementation and thus scientific progress (Baumer, 2017). To date, a coherent and easy-to-use solution for integrating areal data that considers a large number of frequent issues and that is implemented in a robust and reproducible manner is still lacking.

Here, we introduce the R software package *arealDB* to address the overall challenge of integrating heterogeneous areal data in an easy-to-use framework. The package offers a set of tools that are focused

specifically on harmonisation and integration of areal data, thereby reducing complexity and increasing utility and confidence in data quality for downstream applications. *arealDB* automates complicated and error-prone procedures, such as reshaping data tables, semantic matching from vocabularies or metadata and provenance documentation, with only limited required user-input. Finally, it provides extensive unit testing, ensuring that all tools work as expected.

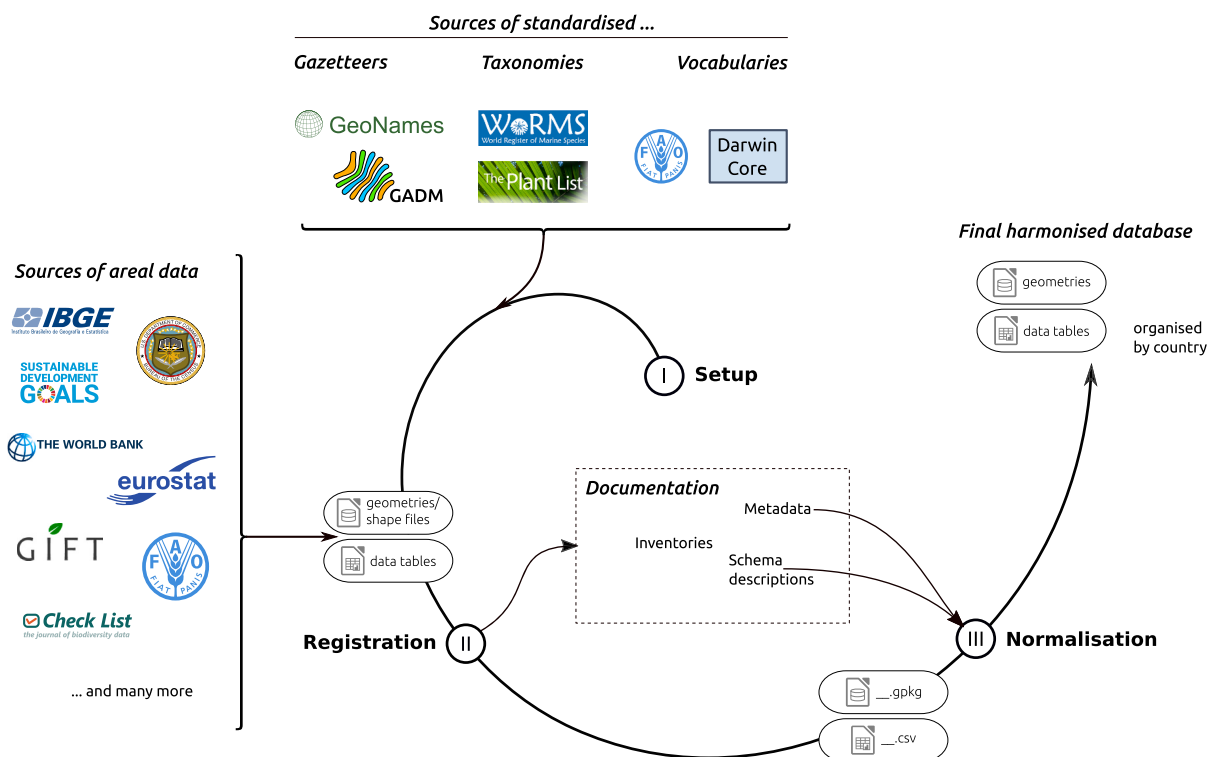
We exemplify the full functionality of *arealDB* by integrating two example datasets on the harvested area of soybean. The first dataset (Brazil) is provided in Portuguese language and accompanied by specific geometries, while the second datasets (USA) is provided in our target language English and does not come with geometries but merely refers to the names of US counties.

## 2. Methods

The R software package *arealDB* contains three groups of tools that reflect three stages of data management (Fig. 1):

- Stage 1 (Initialisation): Set up a database while gathering thematic metadata.
- Stage 2 (Registration): Transform original data from downloaded files into standardised data formats and gather metadata on those files.
- Stage 3 (Normalisation): Harmonise geometries and data tables based on the metadata collected at stage 2, and integrate them in a standardised database.

Technical documentation of any function that comes with this package can be retrieved after installation, for example, via the command `?setPath` for the function of that name.



**Fig. 1. Flow-chart of the general workflow of data integration using *arealDB*.** In stage 1, tables containing standardised terms for all variables of interest are read in to establish an ontological basis. In stage 2, data tables and geometries that have been downloaded from various data sources are registered. In stage 3, those files are harmonised and integrated into the final database. The shown data sources, gazetteers, taxonomies and vocabularies are an exemplary, non exhaustive list of sources that can be handled with *arealDB*.

## 2.1. Project setup

An areal database is started with the function `setPath()`, which creates the standardised directory structure in which the database is stored (Fig. 2a). All further operations within `arealDB` that rely on a path are then relative to this database directory. Any areal database typically contains a set of variables that identify the observed areal units, such as the unit names. However, it would also include other variables that identify the observed phenomenon, such as timesteps, socioeconomic groups of people, agricultural or other commodities or biological species. The function `setVariables()` is used to setup the variables that are used in a project to handle lexical translation and semantic harmonisation of the terms of those variables (Fig. 2b). It creates, by default, the skeleton of two files per variable and database, (1) an index table, which relates the variables' terms to an ID and ancillary information and (2) a translation table, which relates terms in foreign languages and semantics to the target language/ontology. To utilize index tables, an input table that contains term-ID pairs is indispensable and translation tables do not have to, but can be provided with standardised translations of the target variables, which can help improving consistency and data quality (Fig. 1).

Sources of such standardised tabular information may be, for example, gazetteers such as GeoNames (GeoNames, 2019) or the ancillary data in the Global Administrative Areas database (GADM) (Hijmans, 2019), biological taxonomies such as The Plant List (The Plant List, 2013), WoRMS (Horton et al., 2019) or the IUCN Red List (IUCN, 2019), or standardised ontologies, such as those offered by FAOSTAT (FAO, 2019), the Darwin Core (Wieczorek et al., 2012), the Humboldt Core (Guralnick et al., 2018), or the Land Administration Domain Model (Lemmen et al., 2015).

## 2.2. Data registration

An important aspect to ensure the quality of an integrated database is provenance documentation. This allows tracing errors that may show

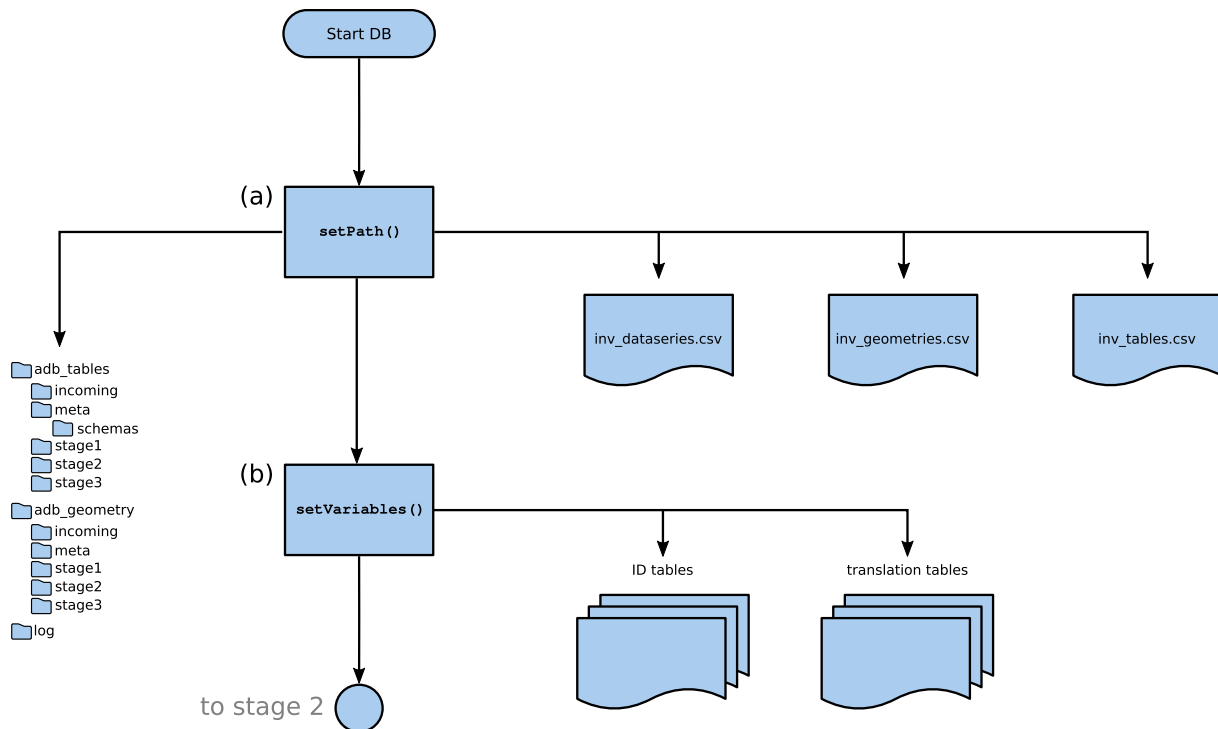
only in the final database to the specific source datasets or to a certain modification process. Documenting provenance requires that the input state of a dataset, as well as procedural metadata that become available as a side-product in the evolution from input to output data, be known.

Thus, the second stage in integrating data with `arealDB` is to create an inventory of the relevant files and to record metadata on the initial state of data, such as original file names, file locations, licenses and the arrangement of data tables (this process is called *registering* in `arealDB`). The arrangement of data tables is managed via the independent R-package `tabshiftr` (Ehrmann and Arnetabshiftr, 2020). Here, so-called schema descriptions are defined, where the table-specific arrangement is described by the position (columns and rows) of data components in the table, which is the basis for automatically reshaping the files in stage 3.

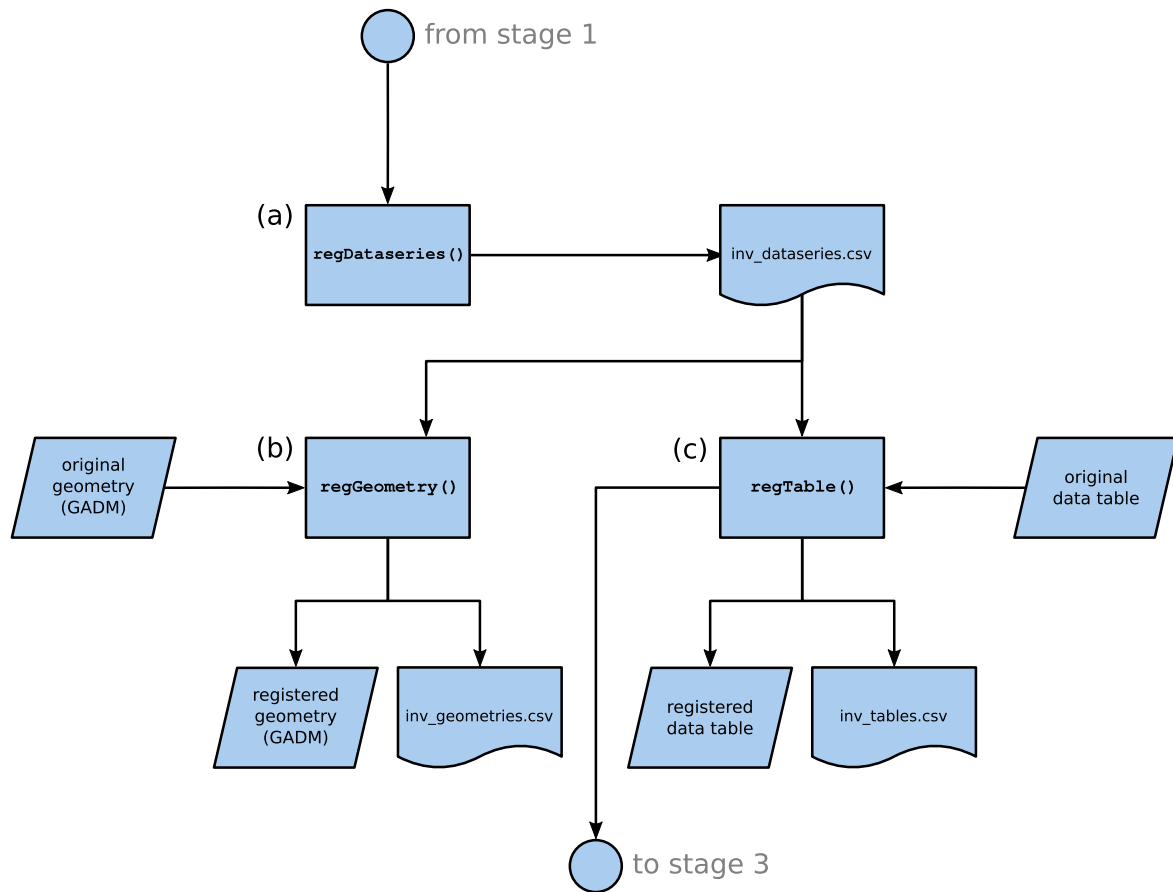
`arealDB` thoroughly documents metadata on three kinds of information, data-series, geometries and data tables. Data-series are collections of data that are provided by the same source and which share more or less the same tabular arrangement and organisational logic. By documenting the data-series, one mostly documents information about a particular data provider and creates a tag that is common for input data that “belong together” and typically share a common structure. That information will eventually be documented in the three inventory tables of the respective names `inv_dataserries.csv`, `inv_geometries.csv` and `inv_tables.csv`. None of these inventory tables ever need to be modified by hand, as all information documented here is managed automatically.

A new data-series is registered with the function `regDataserries()`, before geometries and data tables, which are registered with the functions `regGeometry()` and `regTable()` (Fig. 3). The functions carry out the following operations:

- check the arguments for valid values and consistency,
- oversee that the individual items are transformed to the target format with standard names and stored in the correct directory,



**Fig. 2. Flow-chart of the project setup.**(a) The function `setPath()` initiates the project by creating a directory structure in which the files are stored and by creating the inventory tables for data-series, geometries and census tables. (b) The function `setVariables()` creates index and translation tables for all variables that should be handled in this project.



**Fig. 3. Flow-chart of the registration procedure.**(a) The function `regDataseries()` is used to document the various data-series (both, of data tables and geometries) that are provided by the data source. (b) The function `regGeometry()` is then used to register all geometry files that have been downloaded. (c) Finally, the function `regTable()` is used to register all census tables that have been downloaded, and to relate the census tables to data-series and geometries. The registered files are stored in the folders `"/adb_geometries/stage2"` and `"/adb_tables/stage2"`.

- create IDs for all items and insert the provided metadata into the respective inventory tables.

territorial units and the levels of identifying variables to end up with semantically interoperable data tables.

### 2.3. Data normalisation

The third and final step to integrate areal data consists of reshaping and harmonising the output of stage 2 (Fig. 4) and this process is called *normalising* in `arealDB` (Codd, 1990). This final step is crucial, since at stage 2, there is still no guarantee that names of territorial units in geometries are associated with those in data tables or that areal data are georeferenced, that variables are provided in the same language across several sources, or that data tables are provided in a compatible arrangement.

Geometries are typically provided as shape or geopackage files, which have already been optimised for interoperability, and where it is thus sufficient to know which columns in the attribute table contain names of the territorial units (see `regGeometry()`). The function `normGeometry()` builds harmonised geometry collections with regard to coordinate reference system, as well as semantically interoperable attribute tables. The overall procedure is detailed in Fig. A1.

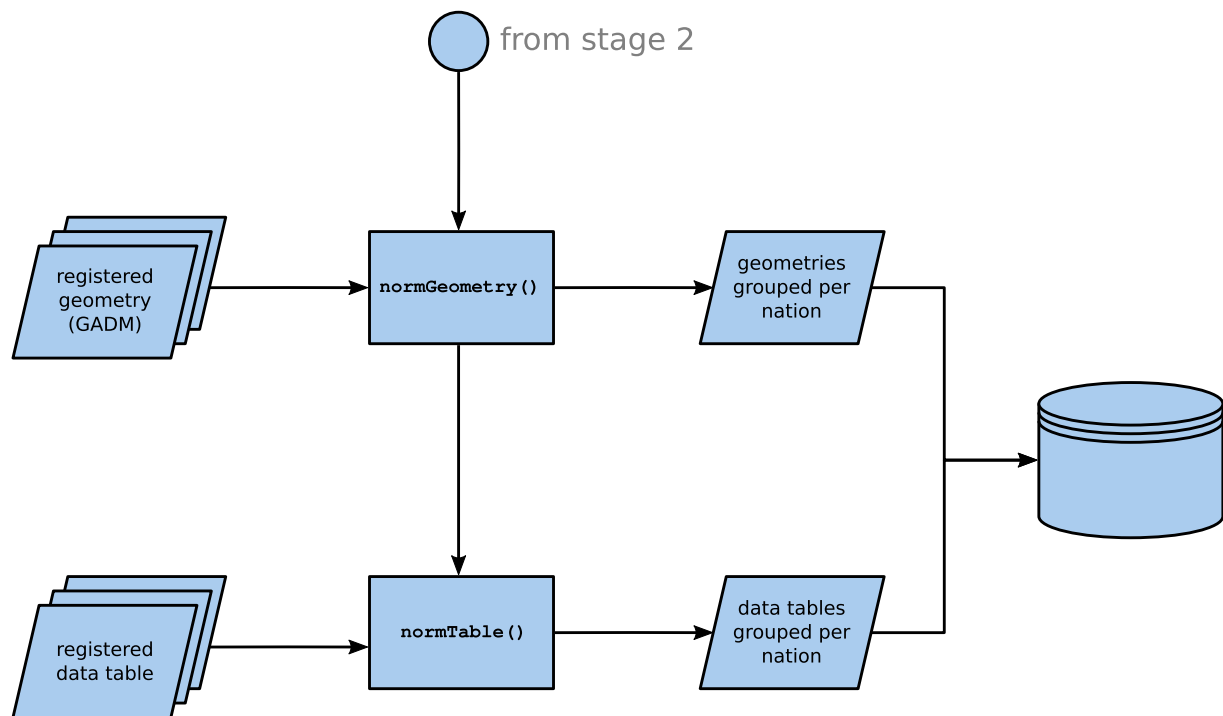
There is no generally accepted way of recording data in tables so that these can be vastly more complex or messy than geometries, and hence require schema descriptions. The schema descriptions that have been recorded at stage 2 document accurate positions of variables held in data tables. The function `normTable()` utilises those schemas, in concert with the function `tabshiftr:reorganise()`, to reshape the data into syntactically interoperable tables. It utilises, furthermore, the functions `matchUnits()` and `matchVars()` to harmonise labels of

### 2.4. The administrative hierarchy

Geometries in `arealDB` are assigned a unique ID for each territorial unit at each administrative level. This requires, first of all, an initial geometry dataset from which this *administrative hierarchy ID* (*ahID*) can be constructed, and which must thus include information on the hierarchical arrangement of the territorial units. At each administrative level a three-digit ID is assigned to the alphabetically sorted unit names. When descending into a lower level, that ID is restarted at 0 within each parent unit. For instance, Tartu County in Estonia has the *ahID* 070013, as Estonia is the 70th country (alphabetically) and Tartu County the 13th county within Estonia.

### 2.5. Translating terms

When handling data from sources that span large spatial extents, these are likely only present in different languages. However, terms may be provided not only in different languages (*sensu stricto*), but also with distinct semantic meanings. For example, the concept of *patch of land that is dominated by grassy vegetation and on which cattle graze*, could be called “pasture” (British English), but also “rangeland” (American English), and is called “pastagem” in Portuguese. The language translation “pastagem  $\Leftrightarrow$  pasture” from Portuguese to English is functionally similar to the semantic translation “rangeland  $\Leftrightarrow$  pasture”. Both examples are cases of many-to-one translations because terms in different languages



**Fig. 4. Flow-chart of the normalisation procedure.** The function `normGeometry()` groups geometries per nation and creates the *administrative hierarchy ID* (*ahID*). The function `normTable()` reshapes the data tables into tidy format, calls `matchUnits()` to assign *ahID* to the areal data and groups the tables per nation. The normalised files are stored in the folders `"/adb_geometries/stage3"` and `"/adb_tables/stage3"`

(*sensu lato*) refer to the same term in the target language. Additionally, in a particular dataset, the term “pasture” might refer to anything from *artificially maintained grassland managed for livestock grazing* to *natural grassland* (i.e., with or without livestock). Ideally, these different meanings will become evident from available metadata. In *arealDB*, an individual term is allowed to refer to different concepts, depending on where it originates, which constitutes a one-to-many semantic translation.

Many-to-one translations are handled quite straightforwardly, in

**Table 1**

List of issues, which have to be considered when building a database of areal data from distinct sources.

challenge	class	required activity
reproject geometries	georeferencing	harmonise spatial projection of distinct input geometries to be able to match them spatially
match geometries and unit names	georeferencing	connect names of territorial units to the correct geometries
territorial changes	alternative data	match areal data associated to territorial units that change through time
data sources that disagree	alternative data	match areal data associated with the same territorial units provided by different data sources
disputed areas	alternative data	identify data that belong to territorial units that are claimed by different authorities
territorial unit names	translation	translate territorial unit names into a common language
distinct concepts	translation	map variable names and values of categorical variables to standardised concepts
disorganised messy data	documentation	arrange all data in the same format
metadata	documentation	document dataset characteristics when input data are retrieved from the source
data provenance	documentation	document the procedure by which the final data product was derived

that the target value is repeated in the column *target* for each translation and terms that refer to it are recorded in the column *origin* (Table 2). One-to-many translations are provided, in the column *source*, either with *geoID* or *tabID*, depending on whether the terms originate from geometries or tables, and in *ID* with the respective ID.

The function `translateTerms()` manages all translations by comparing new terms individually and explicitly with the translation tables that have been created in stage 1. The user is provided with an interface that suggests a range of terms pre-selected from the translation table via approximate string matching (fuzzy matching). The missing translations then have to be provided by the user, so that the new terms can finally be compared against the look-up section of the translation table to check for consistent translations.

### 3. Results

We have tested the described functionalities on several dozen agricultural and forestry census datasets from most countries of the American continent. These data represent many of the challenges outlined in Table 1, including different table arrangements of the areal data, dissimilar variables provided in different languages and associations to geometries, if they were provided at all, in distinct spatial projections. In Appendix A, we exemplarily showcase the integration of data on the soybean harvested area for Brazil and the US. Downstream applications, such as maps of the spatial patterns (Fig. 5 and Appendix B) or

**Table 2**

A translation table that includes (a) many-to-one translations (lines 1 and 2) and a case of one-to-many translations (line 3), as well as (b) language (line 1) and semantic (lines 2 and 3) translations of the term ‘pasture’.

Origin	target	source	ID	notes
Pastagem	pasture			
Rangeland	pasture			
Pasture	grassland	tabID	3	
...				

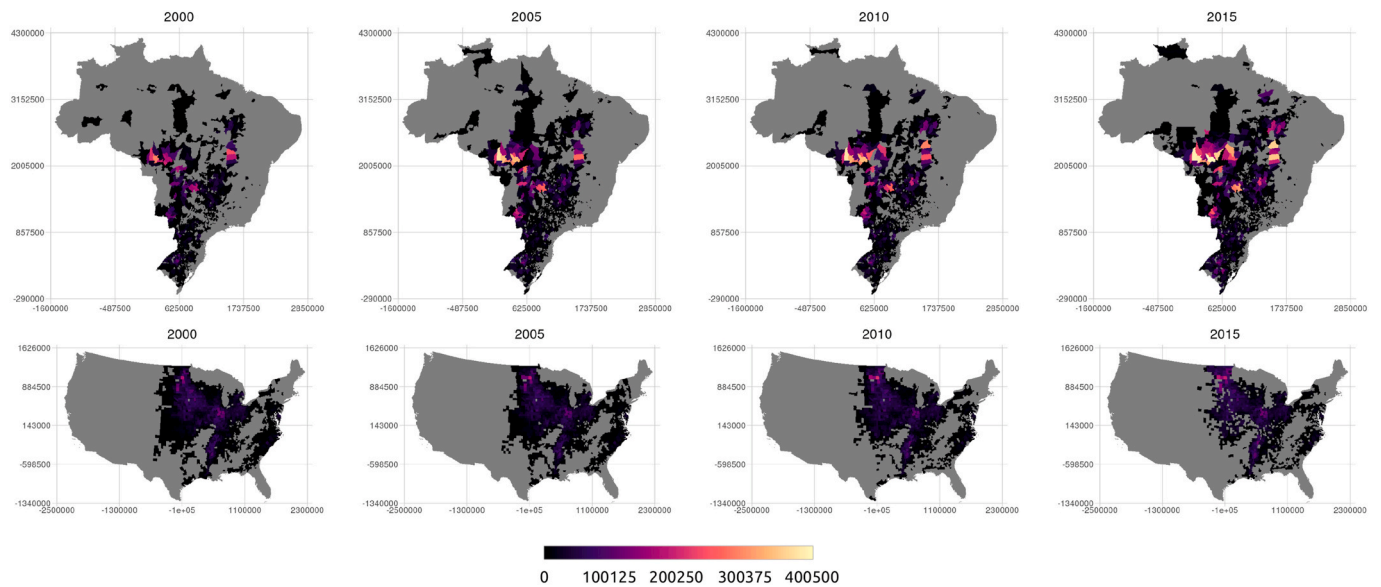


Fig. 5. Choropleth maps of the final, integrated database. The maps have been created with the simple code snippet that is shown in [Appendix B](#). The standardised geometries and data tables were read in with a loop through all countries of interest and the maps were created with another loop through the years of interest to plot the data via the R-package *geomtr* ([geomtr, 2020](#)).

explorative analysis, can profit from the relative ease of accessing all data in a database that has been built with *arealDB* at once.

The schema descriptions for each input dataset are central to this workflow. They record the *type* of each variable, the location in a

spreadsheet (*row*, *col*) at which the variable is and other metadata that are relevant to reshape the table.

At stage 2, the Brazilian dataset is organised in a systematic way, however the arrangement is rather complicated (Listing 1). First of all,

```
library(readr)
# input and schema for Brazil
read_csv(file = paste0(dbPath, "/adb_tables/stage2/processed/bra_3_soy_2000_2018_
  ibge.csv"), col_names = FALSE)
#> # A tibble: 5,569 x 20
#>   X1      X2      X3      X4      X5      X6      X7      X8      X9     X10
#>   <chr> <chr> <chr> <chr> <chr> <chr> <chr> <chr> <chr> <chr>
#> 1 Tabel  NA    NA    NA    NA    NA    NA    NA    NA    NA
#> 2 Varia  NA    NA    NA    NA    NA    NA    NA    NA    NA
#> 3 Munic Ano x  NA    NA    NA    NA    NA    NA    NA    NA    NA
#> 4 NA     2000  2001  2002  2003  2004  2005  2006  2007  2008
#> 5 NA     Soja  Soja  Soja  Soja  Soja  Soja  Soja  Soja  Soja
#> 6 Alta   -      -      -      -      100   -      -      -      -
#> 7 Ariqu  -      -      450   -      -      -      50   -      -
#> 8 Cabi x  200    486    600    1500   1500   5370   7500   6000   7000
#> 9 Cacao  -      -      -      -      -      -      -      -      -
#> 10 Cerej  2700   3353   3400   4516   7184   8000   18000  16200  18000
#> # with 5,559 more rows, and 10 more variables: X11 <chr>,
#> # X12 <chr>, X13 <chr>, X14 <chr>, X15 <chr>, X16 <chr>,
#> # X17 <chr>, X18 <chr>, X19 <chr>, X20 <chr>

read_rds(path = paste0(dbPath, "/adb_tables/meta/schemas/schema_1.rds"))
#> 1 cluster
#> origin: 4/1 (top/left)
#>
#> variable      type      row      col      rel      dist
#> -----
#> al2           id           1         1         F         F
#> al3           id           1         1         F         F
#> year          id           4        2:20        F         F
#> commodities   id           5        2:20        F         F
#> harvested     measured      2:20        F         F
```



the spreadsheet contains a metadata header in the first three rows so that the origin of the table is at the fourth row in the first column. The values of the identifying variables `years` and `commodities` and the target variable `harvested` are stored in the same columns. Moreover, states (`a12`) and municipalities (`a13`) are combined in the first column, so that this column has to be split. The names of the states are abbreviated, requiring translation to correct names (e.g., “RO  $\Leftrightarrow$  Rondônia”).

**Listing 1.** Schema description of the Brazilian dataset.

The US dataset, on the other hand, is already “tidy” at stage 2, i.e., all variables are recorded in individual columns that simply have to be selected, and no lexical or ontological translations are required (Listing 2).

**Listing 2.** Schema description of the US dataset.

After normalising (at stage 3), both tables share the same arrangement, where every information is encoded by IDs that point either to metadata (`tabID`, `geoID`), territorial units (`ahID`), or to commodities (`faoID`).

**Listing 3.** Summary of the core table of the harmonised and integrated database.

## 4. Discussion

The present R package `arealDB` provides so far missing software for harmonising and integrating areal data across multiple heterogeneous sources into a single, consistent database (Fig. 6). By guiding users through the three stages ‘project setup’, ‘data registration’ and ‘data normalisation’, embedded into a collaborative and transparent software environment (Lowndes et al., 2017), `arealDB` substantially improves the speed, scientific accuracy, and reproducibility of integrating areal data.

Lacking such a tool, many data integrators have thus far gone through a highly time-consuming and error-prone process of opening individual input datasets in Excel, translating variables “by hand”, and

reshaping tables by copying relevant data from one spreadsheet into another, typically without any provenance documentation. More sophisticated workflows based on advanced tools, such as lookup and pivot tables or programmable statistical/data management tools often require custom-scripted solutions for the different, heterogeneous input data-sets. `arealDB`, by contrast, merely requires that users document metadata, provide translation tables and specify a schema description for each dataset.

### 4.1. Handling geometries

Using `arealDB`, each individual areal data table may be linked to a different geometry dataset. This helps to avoid political and other assumptions where more than one source of areal data exists for territorial units, for example in cases of disputed areas or administrative changes. Data that refer to such disparate geometry sources can coexist in a database without biasing downstream analyses that are sensitive to the areas of measurement units (e.g., when estimating ecological scaling relationships from species checklists) (Kreft and Jetz, 2007; Keil and Chase, 2019). Where such considerations are not an issue, datasets may be linked to standardised geometries, such as GADM (Hijmans, 2019).

Currently, `arealDB` matches areal data via an internal assessment of the geographical overlap of geometries (using the R package `sf` (Pebesma, 2018); see the flowchart in Appendix C). This procedure could be further refined by additionally incorporating functions that can automatically detect and handle temporal changes of territorial units (Bernard et al., 2018). However, as all raw input geometries are retained in the final database, more sophisticated procedures dedicated to matching recorded geometries may be applied alongside `arealDB`’s own functions.

Another common challenge in matching territorial units lies in names that are shared by multiple territorial units located in different nations and at different administrative levels. For example, the term

```
read_csv(file = paste0(dbPath, "/adb_tables/stage2/processed/usa_3_soy_2000_2018_
  usda.csv"), col_names = FALSE)
#> # A tibble: 30,143 x 21
#>   X1      X2      X3      X4      X5      X6      X7      X8      X9      X10
#>   <chr> <chr> <chr> <chr> <chr> <chr> <chr> <chr> <chr> <chr>
#> 1 Progr Year Period Week Geo State Stat Ag D Ag D Coun
#> 2 SURVEY 2018 YEAR NA COUN ALAB 1 BLAC 40 DALL
#> 3 SURVEY 2018 YEAR NA COUN ALAB 1 BLAC 40 ELMO
#> 4 SURVEY 2018 YEAR NA COUN ALAB 1 BLAC 40 OTHE
#> 5 SURVEY 2018 YEAR NA COUN ALAB 1 BLAC 40 PERRY
#> 6 SURVEY 2018 YEAR NA COUN ALAB 1 BLAC 40 SUMT
#> 7 SURVEY 2018 YEAR NA COUN ALAB 1 COAS 50 BALD
#> 8 SURVEY 2018 YEAR NA COUN ALAB 1 COAS 50 OTHE
#> 9 SURVEY 2018 YEAR NA COUN ALAB 1 MOUN 20 BLOU
#> 10 SURVEY 2018 YEAR NA COUN ALAB 1 MOUN 20 CHER
#> # with 30,133 more rows, and 11 more variables: X11 <chr>,
#> # X12 <chr>, X13 <chr>, X14 <chr>, X15 <chr>, X16 <chr>,
#> # X17 <chr>, X18 <chr>, X19 <chr>, X20 <chr>, X21 <chr>

read_rds(path = paste0(dbPath, "/adb_tables/meta/schemas/schema_2.rds"))
#> 1 cluster (whole spreadsheet)
#>
#>   variable      type      row      col      rel      dist
#>   -----
#>   a12          id          6          F          F
#>   a13          id         10          F          F
#>   year         id          2          F          F
#>   commodities  id         16          F          F
#>   harvested    measured    20          F          F
```

```
read_csv(file = paste0(dbPath, "/adb_tables/stage3/brazil.csv"), col_types = "
iiiiidi") %>%
summary()
#>      id      tabID      geoID      ahID
#> Min.    : 1      Min.    :1      Min.    :5      Min.    :32001001
#> 1st Qu.:24648    1st Qu.:1      1st Qu.:5      1st Qu.:32011029
#> Median :49296    Median :1      Median :5      Median :32015197
#> Mean   :49296    Mean   :1      Mean   :5      Mean   :32015793
#> 3rd Qu.:73944    3rd Qu.:1      3rd Qu.:5      3rd Qu.:32021306
#> Max.   :98591    Max.   :1      Max.   :5      Max.   :32027137
#>
#>      year      harvested      faoID
#> Min.    :2000      Min.    : 0      Min.    :236
#> 1st Qu.:2004      1st Qu.: 460      1st Qu.:236
#> Median :2009      Median : 2950      Median :236
#> Mean   :2009      Mean   : 12388      Mean   :236
#> 3rd Qu.:2014      3rd Qu.: 11500      3rd Qu.:236
#> Max.   :2018      Max.   :411224      Max.   :236
#>
#>      NA's      :65529

read_csv(file = paste0(dbPath, "/adb_tables/stage3/united states of america.csv"),
col_types = "iiiiidi") %>%
summary()
#>      id      tabID      geoID      ahID
#> Min.    : 1      Min.    :2      Min.    :3      Min.    :238001001
#> 1st Qu.: 6871    1st Qu.:2      1st Qu.:3      1st Qu.:238016074
#> Median :13740    Median :2      Median :3      Median :238025054
#> Mean   :13740    Mean   :2      Mean   :3      Mean   :238026647
#> 3rd Qu.:20610    3rd Qu.:2      3rd Qu.:3      3rd Qu.:238036065
#> Max.   :27480    Max.   :2      Max.   :3      Max.   :238050074
#>
#>      year      harvested      faoID
#> Min.    :2000      Min.    : 40.47      Min.    :236
#> 1st Qu.:2004      1st Qu.: 2954.21      1st Qu.:236
#> Median :2008      Median : 12788.07      Median :236
#> Mean   :2009      Mean   : 20302.90      Mean   :236
#> 3rd Qu.:2013      3rd Qu.: 33108.34      3rd Qu.:236
#> Max.   :2018      Max.   :218125.56      Max.   :236
```

‘Santa Cruz’ may refer to city districts, municipalities, departments, or other units in over twenty different nations. *arealDB* addresses this issue by matching territorial unit names to geometries hierarchically.

#### 4.2. Documenting metadata

Even *arealDB* is still prone to some human error, mostly related to correctly reading in the required files and information. *arealDB*’s system of unique IDs ensures that all data values can be linked to metadata on their input tables, associated geometries, and original data sources. Hence, inconsistencies that may still exist in an output database are fully traceable, due to the metadata collected while collating the database (Henzen et al., 2013). Moreover, the relational setup of the resulting database allows that database management software can retrieve metadata on all harmonised datasets (e.g., the recorded number of entities or observations, or value ranges) and make those and other metadata available also to other tools and environments.

#### 4.3. Interoperability

*arealDB* was primarily designed for the purpose of integrating heterogeneous areal data within a given knowledge domain. However, by storing all variables in the same data structure, the tools provided here enable areal databases that are syntactically interoperable across

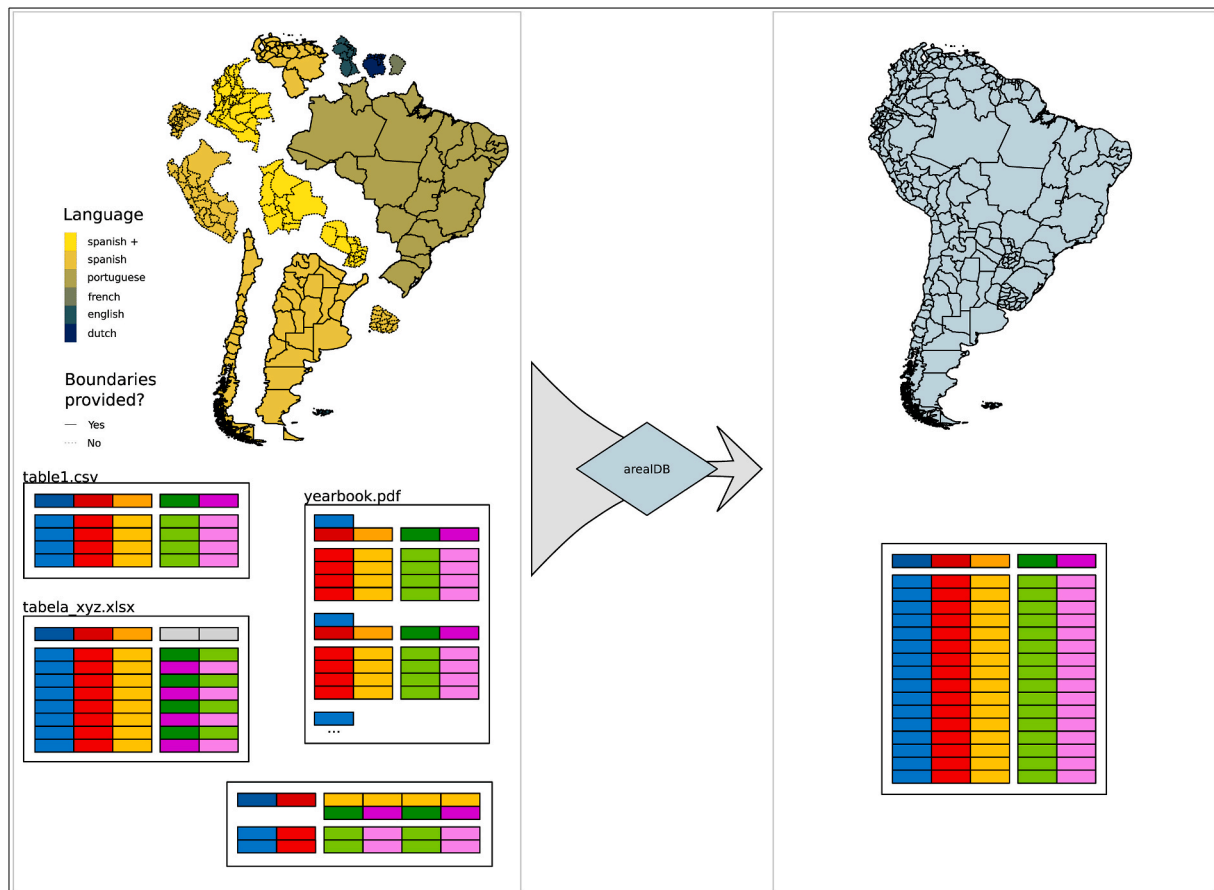
domains. The areal data in each and every output table of *arealDB* have a value of *ahID* (Listing 3) that is by default always derived from the same spatial basis, the GADM dataset. Merely areal data for which *ahID* was derived from specific non-GADM geometries may deviate from a common list of *ahID* values. This means that areal databases from several distinct domains (for instance, human health and biodiversity) that have all been built with the same geometries dataset can simply be joined via *ahID* to combine information on the distinct topics within a consistent database and facilitate interdisciplinary applications (Otto et al., 2015).

While *arealDB* enables such data integration technically, the endeavour of actually integrating databases across knowledge domains in a meaningful way hinges crucially on advances in ontological standardisation to come up with concepts that are valid simultaneously across knowledge domains (van den Brink et al., 2017).

## 5. Conclusions

*arealDB* provides a range of tools that standardise the process of integrating heterogeneous areal data sources into a single, harmonised database, removing hurdles that come with high resource and time requirements. It enables users that may lack the necessary background knowledge to address the majority of the anticipated issues in setting up a coherent workflow of data integration. This helps in homogenising





**Fig. 6. Schematic overview of the purpose of *arealDB*.** Both, data tables that are disorganised and messy and that are based on different languages and concepts, as well as distinctly organised geometric datasets are harmonised and integrated into one consistent database.

data collection methodologies, which is urgently needed to tackle data management strategies that are able to deal with vast amounts of heterogeneous data (Otto et al., 2015).

The tools presented here allow to process any areal data, such as socioeconomic census and survey data, ecological checklist data, data on infectious diseases, (sub-)national indicator data, cadastral parcel data, and many more. *arealDB* lowers the effort to surpass barriers of the spatial, temporal or thematic scope of single data sources. This can enable many exciting and easy to set-up applications at the pan-regional and global scale, supporting data-integration needs across application domains and progressing towards multiple Sustainable Development Goals.

### Software availability

*arealDB* is an R package that is available from CRAN via the function `install.packages("arealDB")`. This paper is based on version v0.3.6, which is installed via `devtools::install_github("EhrmannS/arealDB@v0.3.6")`. All programming was performed by the authors unless stated otherwise.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgements

CM acknowledges funding by the Volkswagen Foundation via a.

Freigeist Fellowship SE acknowledges funding by iDiv via the Flexpool mechanism (FZT-118, DFG).

### Appendix. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.envsoft.2020.104799>.

### References

- Baumer, Benjamin S., 2017. A grammar for reproducible and painless extract-transform-load operations on medium data. *J. Comput. Graph Stat.* 28 (2), 256–264. <https://doi.org/10.1080/10618600.2018.1512867>.
- Bernard, Camille, Plumejeaud-Perreau, Christine, Villanova-Oliver, Marlène, Gensel, Jérôme, Dao, Hy, 2018. An ontology-based algorithm for managing the evolution of multi-level territorial partitions. In: *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, pp. 456–459. <https://doi.org/10.1145/3274895.3274944>.
- Codd, E.F., 1990. *The Relational Model for Database Management: Version 2*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, ISBN 0-201-14192-2.
- De Giacomo, Giuseppe, Lembo, Domenico, Lenzerini, Maurizio, Poggi, Antonella, Rosati, Riccardo, 2018. Using ontologies for semantic data integration. In: *A Comprehensive Guide through the Italian Database Research over the Last 25 Years*. Springer, ISBN 978-3-319-61892-0, pp. 187–202. [https://doi.org/10.1007/978-3-319-61893-7\\_11](https://doi.org/10.1007/978-3-319-61893-7_11), 05.
- Debroy, Vidroha, Brimble, Lance, Matt, Yost, 2018. Newtl: engineering an extract, transform, load (etl) software system for business on a very large scale. In: *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*. ACM, pp. 1568–1575.
- Du, Heshan, Alechina, Natasha, Jackson, Mike, Hart, Glen, 2013. Matching formal and informal geospatial ontologies. In: *Geographic Information Science at the Heart of Europe*. Springer, pp. 155–171. [https://doi.org/10.1007/978-3-319-00615-4\\_9](https://doi.org/10.1007/978-3-319-00615-4_9). [https://link.springer.com/chapter/10.1007/978-3-319-00615-4\\_9](https://link.springer.com/chapter/10.1007/978-3-319-00615-4_9).
- Ehrmann, Steffen, Rümmler, Arne, 2020. *tabshifr: Reshape Disorganised Messy Data*. <https://cran.r-project.org/web/packages/tabshifr/index.html>. R package version 0.1.2.

- Ehrmann, Steffen, 2020. *geometr: Generate and Modify Interoperable Geometric Shapes*. <https://cran.r-project.org/web/packages/geometr/index.html>. R package version 0.2.5.
- Faostat FAO, 2019. <http://www.fao.org/faostat>.
- Geonames GeoNames, 2019. <http://www.geonames.org>.
- Guralnick, Robert, Walls, Ramona, Walter, Jetz, 2018. Humboldt core-toward a standardized capture of biological inventories for biodiversity monitoring, modeling and assessment. *Ecography* 41 (5), 713–725. <https://onlinelibrary.wiley.com/doi/10.1111/ecog.02942>.
- Henzen, Christin, Mäs, Stephan, Bernard, Lars, 2013. Provenance information in geodata infrastructures, 05. In: *Geographic Information Science at the Heart of Europe*, vol. 2013. Springer International Publishing, ISBN 978-3-319-00615-4, pp. 133–151. [https://doi.org/10.1007/978-3-319-00615-4\\_8](https://doi.org/10.1007/978-3-319-00615-4_8). doi: 10.1007/978-3-319-00615-4\_8.
- Hijmans, Robert J., 2019. Global administrative areas. [www.gadm.org](http://www.gadm.org).
- Horridge, M., 2011. *A Practical Guide to Building Owl Ontologies Using Protégé 4 and Co-ode Tools Edition 1.3* retrieved march 10, 2013.
- Horton, T., Kroh, A., Ah Yong, S., Bailly, N., Boyko, C.B., Brandão, S.N., Gofas, S., Hooper, J.N.A., Hernandez, F., Holovachov, O., Mees, J., Molodtsova, T.N., Paulay, G., Decock, W., Dekeyser, S., Lanssens, T., Vandepitte, L., Vanhoorne, B., Verfaillie, K., Adlard, R., Adriaens, P., Agatha, S., Ahn, K.J., Akkari, N., Alvarez, B., Anderson, G., Angel, M.V., Arango, C., Artois, T., Atkinson, S., Bank, R., Barber, A., Barbosa, J.P., Bartsch, I., Bellan-Santini, D., Bernot, J., Berta, A., Bezerra, T.N., Bieler, R., Blanco, S., Blasco-Costa, I., Blazewicz, M., Bock, P., Böttger-Schnack, R., Bouchet, P., Boury-Esnault, N., Boxshall, G., Bray, R., Breure, B., Bruce, N.L., Cairns, S., Cárdenas, P., Carstens, E., Chan, B.K., Chan, T.Y., Cheng, L., Churchill, M., Coleman, C.O., Collins, A.G., Corbari, L., Cordeiro, R., Cornils, A., Coste, M., Costello, M.J., Crandall, K.A., Cremonte, F., Cribb, T., Cutmore, S., Dahdouh-Guebas, F., Daly, M., Daneliya, M., Dauvin, J.C., Davie, P., De Broyer, C., De Grave, S., de Mazancourt, V., de Voogd, N.J., Decker, P., Decraemer, W., Defaye, D., d'Hondt, J.L., Dippenaar, S., Dohrmann, M., Dolan, J., Domning, D., Downey, R., Ector, L., Eisendle-Flockner, U., Eitel, M., Encarnação, S.C.d., Enghoff, H., Epler, J., Ewers-Saucedo, C., Faber, M., Feist, S., Figueroa, D., Finn, J., Fišer, C., Fordyce, E., Foster, W., Frank, J.H., Fransen, C., Furuya, H., Galea, H., Garcia-Alvarez, O., Garic, R., Gasca, R., Gaviria-Melo, S., Gerken, S., Gibson, D., Gibson, R., Gil, J., Gittenberger, A., Glasby, C., Glover, A., Gómez-Noguera, S.E., González-Solís, D., Gordon, D., Grabowski, M., Gravili, C., Guerra-García, J.M., Guidetti, R., Guiry, M. D., Hadfield, K.A., Hajdu, E., Hallermann, J., Hayward, B.W., Hendrycks, E., Herbert, D., Herrera Bachiller, A., Ho, J.s., Hodda, M., Hoeg, J., Hoeksema, B., Houart, R., Hughes, L., Hyžný, M., Iñiesta, L.F.M., Iseto, T., Ivanenko, S., Iwataki, M., Janssen, R., Jarms, G., Jaume, D., Jazdzewski, K., Jersabek, C.D., Jóźwiak, P., Kabat, A., Kantor, Y., Karanovic, I., Karthick, B., Kim, Y.H., King, R., Kirk, P.M., Klautau, M., Kocielek, J.P., Köhler, F., Kolb, J., Kotov, A., Kremenetskaia, A., Kristensen, R.M., Kulikovskiy, M., Kullander, S., Lambert, G., Lazarus, D., Le Coze, F., LeCroy, S., Leduc, D., Lefkowitz, E.J., Lemaitre, R., Liu, Y., Lörz, A.N., Lowry, J., Ludwig, T., Lundholm, N., Macpherson, E., Madin, L., Mah, C., Mamo, B., Mamos, T., Manconi, R., Mapstone, G., Marek, P.E., Marshall, B., Marshall, D.J., Martin, P., Mast, R., McFadden, C., McInnes, S.J., Meidla, T., Meland, K., Merrin, K. L., Messing, C., Miljutin, D., Mills, C., Moestrup, Ø., Mokievsky, V., Monniot, F., Mooi, R., Morandini, A.C., Moreira da Rocha, R., Moretzsohn, F., Morrow, C., Mortelmans, J., Mortimer, J., Musco, L., Neubauer, T.A., Neubert, E., Neuhaus, B., Ng, P., Nguyen, A.D., Nielsen, C., Nishikawa, T., Norenburg, J., O'Hara, T., Opreas, D., Osawa, M., Osigus, H.J., Ota, Y., Páll-Gergely, B., Patterson, D., Paxton, H., Peña-Santiago, R., Perrier, V., Perrin, W., Petrescu, I., Picton, B., Pilger, J.F., Piser, A.B., Polhemus, D., Poore, G.C., Potapova, M., Pugh, P., Read, G., Reich, M., Reimer, J.D., Reip, H., Reuscher, M., Reynolds, J.W., Richling, I., Rimet, F., Ríos, P., Rius, M., Rogers, D.C., Rosenberg, G., Rützler, K., Sabbe, K., Saiz-Salinas, J., Sala, S., Santagata, S., Santos, S., Sar, E., Satoh, A., Saucède, T., Schatz, H., Schierwater, B., Schmidt-Rhaesa, A., Schneider, S., Schönborg, C., Schuchert, P., Senna, A.R., Serejo, C., Shaik, S., Shamsi, S., Sharma, J., Shear, W.A., Shenkar, N., Shinn, A., Short, M., Sicinski, J., Sierwald, P., Simmons, E., Sinniger, F., Sivell, D., Sket, B., Smit, H., Smit, N., Smol, N., Souza-Filho, J.F., Spelda, J., Sterrer, W., Stienen, E., Stoev, P., Stöhr, S., Strand, M., Suárez-Morales, E., Summers, M., Suppan, L., Suttle, C., Swalla, B.J., Taiti, S., Tanaka, M., Tandberg, A. H., Tang, D., Tasker, M., Taylor, J., Taylor, J., Tchesunov, A., ten Hove, H., ter Poorten, J.J., Thomas, J.D., Thuesen, E.V., Thurston, M., Thuy, B., Timi, J.T., Timm, T., Todaro, A., Turon, X., Tyler, S., Uetz, P., Uribe-Palomino, J., Utevsky, S., Vacelet, J., Vachard, D., Vader, W., Väinölä, R., Van de Vijver, B., van der Meij, S.E., van Haaren, T., van Soest, R.W., Vanreusel, A., Venekey, V., Vinarski, M., Vonk, R., Vos, C., Walker-Smith, G., Walter, T.C., Watling, L., Wayland, M., Wesener, T., Wetzell, C.E., Whipps, C., White, K., Wieneke, U., Williams, D.M., Williams, G., Wilson, R., Witkowski, A., Witkowski, J., Wyatt, N., Wylezich, C., Xu, K., Zanol, J., Zeidler, W., Zhao, Z., 2019. World register of marine species (worms). <http://www.marinespecies.org>. <http://www.marinespecies.org>. Accessed: 2019-09-11.
- IUCN, 2019. The iucn red list of threatened species. <http://www.iucnredlist.org>.
- Keil, Petr, Chase, Jonathan M., 2019. Global patterns and drivers of tree diversity integrated across a continuum of spatial grains. *Nat. Ecol. Evol.* 3 (3), 390. <https://doi.org/10.1038/s41559-019-0799-0>.
- Kreft, Holger, Walter, Jetz, 2007. Global patterns and determinants of vascular plant diversity. *Proc. Natl. Acad. Sci. Unit. States Am.* <https://doi.org/10.1073/pnas.0608361104>. <https://www.pnas.org/content/104/14/5925>.
- Lemmen, Christiaan, Van Oosterom, Peter, Bennett, Rohan, 2015. The land administration domain model. *Land Use Pol.* 49, 535–545. <https://doi.org/10.1016/j.landusepol.2015.01.014>. <https://www.sciencedirect.com/science/article/pii/S0264837715000174?via%3Dihub>.
- Lowndes, Julia S Stewart, Best, Benjamin D., Scarborough, Courtney, C Afflerbach, Jamie, Frazier, Melanie R., O'Hara, Casey C., Jiang, Ning, Halpern, Benjamin S., 2017. Our path to better science in less time using open data science tools. *Nat. Ecol. Evol.* 1 (6), 1–7. <https://doi.org/10.1038/s41559-017-0160>.
- Otto, Ilona M., Biewald, Anne, Coumou, Dim, Feulner, Georg, Köhler, Claudia, Nocke, Thomas, Anders, Blok, Albert, Gröber, Sabine Selchow, Tyfield, David, et al., 2015. Socio-economic data for global environmental change research. *Nat. Clim. Change* 5 (6), 503. <https://doi.org/10.1038/nclimate2593>. <https://www.nature.com/articles/nclimate2593.pdf>.
- Pebesma, Edzer, 2018. Simple features for R: standardized support for spatial vector data. *The R Journal* 10 (1), 439–446. <https://doi.org/10.32614/rj-2018-009> doi: 10.32614/RJ-2018-009.
- Petersen, Anne Helby, Ekström, Claus Thorn, 2019. *datamaid: your assistant for documenting supervised data quality screening in r*. *J. Stat. Software* 90 (1), 1–38.
- The Plant List, 2013. The plant list. <http://www.theplantlist.org>. Version 1.1.
- van den Brink, Linda, Janssen, Paul, Quak, Wilko, Stoter, Jantien, 2017. Towards a high level of semantic harmonisation in the geospatial domain. *Comput. Environ. Urban Syst.* 62, 233–242. <https://doi.org/10.1016/j.compenvurbysys.2016.12.002>.
- Wickham, Hadley, 2014. Tidy data. *J. Stat. Software* 59 (10), 1–23. <https://doi.org/10.18637/jss.v059.i10>. <https://www.jstatsoft.org/article/view/v059i10>.
- Wieczorek, John, Bloom, David, Guralnick, Robert, Blum, Stan, Döring, Markus, Giovanni, Renato, Robertson, Tim, Vieglais, David, 2012. Darwin core: an evolving community-developed biodiversity data standard. *PLoS One* 7 (1). <https://doi.org/10.1371/journal.pone.0029715>. <https://journals.plos.org/plosone/article/file?id=10.1371/journal.pone.0029715&type=printable>.