

# Notes on cave surveying and GIS

## Summary

The idea that reduced cave survey data should be easily readable into a Geographic Information System (GIS) platform such as [QGIS](#) is practically a no-brainer, as it can then be integrated with other geographical data such as maps, satellite imagery, digital elevation models, and the like. This is much closer to being achievable than one might think. Here's the contents of a typical [surveyx](#) .3d file as exposed by running `dump3d`:

- survey metadata: title, date, and co-ordinate reference system;
- strings of survey legs with metadata: names, flags (normal, duplicate, splay, surface);
- survey stations with metadata: names, flags (exported, entrance, fixed, surface) and passage cross-sections (LRUD data).

Now compare this to a typical [ESRI shapefile](#), or the [GeoPackage](#) data format from the [Open Geospatial Consortium](#), which are well known containers for GIS vector data. These formats specify:

- a co-ordinate reference system, and possible other metadata;
- geometries comprising points, lines, polylines (line strings), and polygons, with or without *z*-dimension (elevation) data;
- geometry attributes consisting of records of various kinds that are user-configurable.

At this point you are supposed to slap yourself on the head and ask why on earth we haven't been using a GIS-compatible format for storing reduced survey data all along! It's certainly flexible enough to contain all the information normally included in a .3d file.

## Spatial Reference Systems

In order for this to work smoothly, we first have to be on top of our [spatial reference system \(SRS\)](#) in general GIS parlance, or co-ordinate reference system (CRS) in QGIS language. The following notes hopefully contain enough of the truth to be useful. Something closer to the truth can be found in [A guide to coordinate systems in Great Britain](#) by the Ordnance Survey (OS).

An SRS usually comprises:

- a [geodetic datum](#) or Terrestrial Reference System, which is usually an ellipsoid which specifies the overall shape of the earth's surface: [WGS84](#) used in GPS, and [OSGB36](#) used by the Ordnance Survey in the UK are two examples;
- a map projection which is nearly always a [Transverse Mercator](#) projection, such as the [Universal Transverse Mercator \(UTM\)](#) system used in GPS: the map projection tries to flatten *locally* the curved surface of the earth (there

is always some compromise involved here) and a given geodetic datum may have more than one map projection depending which part of the Earth's surface one is interested in (UTM has 60 zones for instance, and all karst in mainland Britain lies in UTM zone 30N);

- a co-ordinate system (typically metric) defined on top of the map projection, usually specifying a 'false origin' so that co-ordinates are always positive.

Given the geodetic datum one can always work with latitudes and longitudes, but these aren't terribly convenient for cave survey data crunching. Also beware that the same point on the earth's surface may have a different latitude and longitude according to the choice of datum: this is known as a datum shift, and a well-known example is the [datum shift between WGS84 and OSGB36](#) that nowadays only shows up in [Magic Map](#).

WGS84 is pretty much universally used nowadays on the internet, for example Google's [Keyhole Markup Language \(KML\)](#) only supports WGS84 latitude and longitude, to upload to [Google Earth](#). Also most GPS devices report latitude and longitude for WGS84, though more often than not you won't see this directly but rather get metric UTM co-ordinates, or metric British National Grid co-ordinates in the UK.

To further add to the confusion, latitude and longitude can be reported in decimal degrees; or degrees, minutes, and seconds (or even degrees and decimal minutes). For example the entrance to Dow Cave is at NGR SD 98378 74300, which translates to (WGS84)  $54^{\circ} 9' 52.2''$  N  $2^{\circ} 1' 34.8''$  W where one decimal place in the seconds corresponds to approximately 3m on the ground, or (WGS84)  $54.16450^{\circ}$  N  $2.02634^{\circ}$  W where five decimal places corresponds approximately to 1m on the ground, or (WGS84) 30U 6 002 262 563 570 for UTM. Online converters between British National Grid references and WGS84 latitudes and longitudes can be found on the internet by searching for 'OSGB36 to WGS84 converter'. To check things, the WGS84 latitude and longitude in decimal degrees can be copied and pasted into Google maps for example, or for that matter directly into the Google search engine.

Note there is a confusing difference between the [Web Mercator](#) often used for online maps, and the standard conformal Mercator projection. For this reason the US DoD has declared the Web Mercator projection 'to be unacceptable for any official use', which is a good enough reason for me to avoid where possible.

In the UK, Ordnance Survey (OS) [British National Grid](#) co-ordinates provide a metric SRS which is convenient for cave survey data. Typically one fixes cave entrances using the numeric part of the national grid reference (NGR). The NGR can be specified in two ways. The most convenient way is to use the OS grid letter system in which a pair of letters specifies a  $100\text{ km} \times 100\text{ km}$  square. Then, within that, a 10-figure national grid reference (NGR) specifies a location to within a square metre. This system (two letters plus 10 figures) is what is usually encountered when using a GPS device set to the British National Grid.

Many datasets in the Cave Registry have entrance fixes specified as 10-fig NGRs, without the grid letters which are assumed known.

Alternatively, and more commonly in GIS, one can use an all-numeric 12-figure NGR in which the leading figures signal the  $100\text{ km} \times 100\text{ km}$  square. For example in the all-numeric scheme the entrance to Dow Cave is at NGR 398378 474300.

In the letter-based system the co-ordinates are often truncated to 8-fig or 6-fig NGRs, to reflect the accuracy of the GPS device for instance (thus 8-fig NGRs are used in the new Northern Caves). In case you forgot your school geography lessons, recall that the correct way to truncate an NGR is to *drop* the least significant figures, not to round to the nearest 10 or 100. This is because an 8-fig (or 6-fig) NGR actually specifies a  $10\text{ m} \times 10\text{ m}$  (or  $100\text{ m} \times 100\text{ m}$ ) *square* and not an approximate position as such. Thus the 6-fig NGR for the Dow Cave entrance is NGR SD 983 743.

To check NGRs in the UK, one can use the ‘Where am I?’ tool in the [Magic Map](#) application. Note that unless explicitly set to use WGS84, Magic Map reports OSGB36 latitude and longitude, which as mentioned is offset from WGS84 by a datum shift of up to 50–100 m. Beware copying and pasting these OSGB36 latitudes and longitudes into Google Maps!

Elsewhere in the world, or for that matter in the UK as well, the UTM system offers a convenient metric SRS for embedding cave survey data. Typically one fixes the entrance co-ordinates as the numeric part of the UTM position, making a note of the UTM grid zone. Online converters from WGS84 latitude and longitude to UTM or back are easily found (thus, for example, the Dow Cave entrance above). Perhaps it’s restating the obvious but if you accidentally paste OSGB36 latitudes and longitudes into a UTM converter, you will likely be out by 50–100 m.

## Georeferencing cave survey data

Back to cave surveying: for most surveys the earth’s surface can be regarded as essentially flat, so one is working in a 3d world with eastings, northings, and altitudes, with the origin of the co-ordinate system chosen at one’s convenience. Perhaps for synoptic maps of very large karst areas, one might be worried about the curvature of the earth’s surface, but for the most part assuming the world is flat should introduce negligible errors, at least in comparison to the errors that typically creep into cave survey projects.

As long as this local cave co-ordinate system can be tied into one of the known geodetic SRS schemes (*i. e.* [georeferenced](#)), then any feature in the cave will have a known position in GIS terms, and can thus be tied into any other georeferenced data such as maps, satellite imagery, digital elevation models, *etc.* Given that most cave surveying is done in metres, it is obviously convenient to tie into an SRS which uses metric co-ordinates, such as UTM or British National Grid. Note that once you’ve tied the dataset into a recognised SRS, any GIS platform

worth its salt will be able to re-project into a different SRS, and will be able to display and combine information from different sources irrespective of the SRS.

The easiest way to georeference cave survey data, with a modern survey distribution, is to **\*fix** cave entrances with appropriate co-ordinates and make judicious use of the **\*cs** commands (for co-ordinate system): use a plain **\*cs** command to specify the input SRS that the entrance co-ordinates are given in, and a **\*cs out** command to specify what the output SRS should be. In the UK for instance one can use this to convert between the OS grid letter system and the all-numeric scheme.

The cave survey data used in the examples below is included in the repository under the **DowProv** directory. It is for the [Dow Cave - Providence Pot system](#) (Great Whernside, Wharfedale, UK), and is essentially a snapshot of the data held in the [Cave Registry Data Archive](#). Note that the **.svx** files have [unix-style line endings](#) so on Windows you might have to use something like [Notepad++](#) to look at them. The processed data is **DowProv.3d**, generated using **survex 1.2.32**.

Back to georeferencing, the cave-specific file **DowCave.svx** (for example) contains

```
*begin DowCave
*export entrance
...
*entrance entrance
*fix entrance 98378 74300 334
*equat entrance dow1.1
...
```

and the master file **DowProv.svx** contains

```
*cs OSGB:SD
*cs out EPSG:7405
...
*begin DowProv
*include DowCave
...
```

(obviously this is only one of many possible ways to add the metadata into the **survex** files).

Thus the file **DowCave.svx** contains a **\*fix** which specifies the entrance location as a 10-fig NGR SD 98378 74300, without the SD part. The easting and northing here (and elevation [ODN](#)) were obtained by field work. Then the file **DowProv.svx** specifies input SRS is the SD square, and asks that the reduced data should be exported using the all-numeric British National Grid scheme, here codified with [European Petroleum Survey Group \(EPSG\)](#) code **EPSG:7405**. Using **EPSG** numbers avoids potential misunderstanding when importing into a GIS platform, for example in **QGIS** one can find the exact exported SRS easily enough by searching on the **EPSG** number.

If you check the processed survey in `aven`, or run `3dtopos` on the `.3d` file, the processed entrance co-ordinates are now indeed

```
(398378.00, 474300.00, 334.00 ) dowprov.dowcave.entrance
```

Whilst this may seem like a crazily over-the-top way to add a ‘3’ and ‘4’ to the entrance co-ordinates, it is actually very simple to implement: one only needs to add two lines (the `*cs` and `*cs out` commands) to the `survex` file. The benefit is that it is robust, clean, and unambiguous. Moreover, the output SRS is included as metadata in the `.3d` file; thus with `dump3d` one sees

```
CS +init=epsg:7405 +no_defs
```

(this is in fact a [PROJ.4](#) string which species the map projection, and can be directly pushed to a GIS application).

As a slightly less trivial example, one can ask for the reduced survey data to be re-projected as UTM co-ordinates. This can be done almost totally trivially by replacing the previous `*cs out` command with `*cs out EPSG:3042` which specifies the output SRS is (WGS84) UTM zone 30N (this includes zone 30U). If we now reduce the data with `cavern` and check with `3dtopos` we find the Dow Cave entrance has magically moved to

```
(563570.22, 6002262.20, 384.57 ) dowprov.dowcave.entrance
```

and the exported SRS from `dump3d` is

```
CS +init=epsg:3042 +no_defs
```

As expected, the entrance location in UTM is the same as obtained by converting the original NGR first to WGS84 latitude and longitude, then to UTM, using the online converters. Note that in re-projecting to UTM, we also get a vertical datum shift.

For another example, the CUCC Austria data set which comes as sample data with the `survex` distribution can be georeferenced by adding the following to the top of the `all.svx` file:

```
*cs custom "+proj=tmerc +lat_0=0 +lon_0=13d20 +k=1
+x_0=0 +y_0=-5200000 +ellps=bessel
+towgs84=577.326,90.129,463.919,5.137,1.474,5.297,2.4232
+units=m +no_defs"
*cs out EPSG:31255
```

The first 4 lines are all one line in the real file, and this specifies the custom SRS in which the co-ordinates of the surface fixed points in the Austria data set are specified. The second line determines the output SRS. This doesn’t really matter too much as long as the SRS can be recognised by the GIS platform: this example uses the MGI / Austria Gauss-Krüger (GK) Central SRS (EPSG:31255), where the *only* difference compared to custom SRS is in the `+y_0` false origin. Another output SRS could be EPSG:3045 which is (WGS84) UTM zone 33N.

I've gone into these examples in some detail as the `survex` documentation on the `*cs` command is rather spartan.

As a further benefit, providing the `survex` data files include correctly formatted `*date` commands (as the Dow-Providence dataset does), the `*cs` commands make `survex` aware of the geodetic SRS and magnetic declination corrections can be automatically added. This is another reason one might want to 'do things properly' with `*cs` commands. The `DowProv.svx` master file thus also contains the lines (the first two are just comments)

```
; Mag dec calculated for SD 97480 72608 alt 225m
; Dowbergill Bridge, just above Kettlewell
*declination auto 97480 72608 225
```

This correctly applies the magnetic declination using the [International Geomagnetic Reference Field \(IGRF\)](#) model, calculated at the specified location in the input SRS, and applied to *all* the included survey files, in this case taking into account the range of dates which spans some 30 years.

## GIS import methods

### Using DXF

Once the dataset has been georeferenced as just described, a DXF file can be generated using the `survex cad3d` tool, or exporting as such from `aven`. This DXF file can then be directly imported into QGIS. At present this route does not neatly capture *z*-dimension (elevation) data.

### DXF conversion using GDAL

This route requires command-line access to the [GDAL utilities](#).

From the DXF file, the centreline can be extracted by running

```
ogr2ogr -f "ESRI Shapefile" DowProv_centreline.shp DowProv.dxf \
  -where "Layer='CentreLine'" -a_srs EPSG:7405
```

We take the opportunity here to add an SRS to match that used in the georeferenced survey data. The resulting shapefile can then be imported in QGIS, and this route *does* preserve *z*-dimension (elevation) data.

Similarly the stations with labels (and elevations) can be extracted by running

```
ogr2ogr -f "ESRI Shapefile" DowProv_stations.shp DowProv.dxf \
  -where "Layer='Labels'" -a_srs EPSG:7405
```

### QGIS import plugins (preferred route)

Features of these plugins include:

- no dependencies, natively reads binary `.3d` files ;

- import stations and legs with full metadata ;
- features carry *z*-dimension (elevation) data ;
- create passage walls, cross-sections, and polygons from LRUD data ;
- CRS can be set from a PROJ.4 string embedded in `.3d` file ;
- save results into a GeoPackage (`.gpkg`) shapefile.

More details and installation instructions can be found in the [QGIS2](#) and [QGIS3](#) plugin repositories. Additional scripts can be found in the `extra` directories.

## Georeferencing images, maps, and old surveys

Georeferencing here refers to assigning a co-ordinate system to an image or map, or a scanned hard copy of a survey. The actual steps require identifying so-called Ground Control Points (GCPs), which are identifiable features on the map for which actual co-ordinates are known. One way to do this is to use the [GDAL Georeferencer](#) plugin in QGIS. Georeferencing surveys may be easier if there is more than one entrance and the positions are known, or there is already a surface grid. If there is only one entrance then tracing a centerline in Inkscape and using the `survex` output tool as described [here](#) may help.

## Copying

These notes are licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License \(CC BY-SA 4.0\)](#).

These notes are distributed ‘as is’, in the hope that they will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of FITNESS FOR A PARTICULAR PURPOSE.

## Copyright

Copyright © (2017–19) Patrick B Warren.