Free and Open Source Software in Archaeological Research Processes: an Application to the Study of African Red Slip Ware in Northern Italy

Costa, S.
Dipartimento di Archeologia e Storia delle Arti, University of Siena, Italy
steko@unisi.it

Free and open source software makes a good environment available for routine archaeological work, but what about a new research that starts from scratch with specific aims and requirements? How are software tools created and assembled regarded to the specific aims of a research project? In 2009, we carried out a research about the distribution, circulation, and consumption of African Red Slip Ware in Northern Italy during Late Antiquity.
The entire process of data collection, exploratory and advanced analysis, was performed using free and open source software. Most importantly, we also developed dedicated tools and programs, following the same philosophy and methods of open source software. The collected data and the developed software programs are available on the web under free licenses that allow anyone to use, modify and redistribute them.

Keywords: Late Roman Pottery, Open Source GIS, Open Data

1. Introduction

It is now widely recognized that free and open source software represents a good environment for routine archaeological work, including the recording of excavation data and cultural resources management. On the other hand, there is very little literature about research projects that started from scratch with specific aims and requirements. How are software tools created and assembled with regard to the aims of such a project? This paper means to explore the full potential of free and open source software, presenting a case study about the use, development, and application of free and open source software in the context of archaeological research carried out during the author's Master's thesis at the University of Siena. In 2009, we studied the distribution, circulation, and consumption of African Red Slip Ware in Northern Italy during Late Antiquity.

The entire process of data collection, exploratory and complete data analysis, was performed using free and open source software. Most importantly, we also developed dedicated tools and programs, by doing some interesting experiments that we found worth sharing, all following the same philosophy and methods of open source software communities. The collected data and the developed software programs are available on the web under free licenses that allow anyone to use, modify, and redistribute.

Both source code and raw data can be found in a Mercurial repository at the URL http://bitbucket.org/steko/thesis-app/

We will go through the entire process step by step, starting from the choice of tools and data collection and explaining in detail how the same tools were used to perform the required analyis.

All the work was performed using the Debian GNU/Linux operating system (http://www.debian.org/).

2. African Red Slip Ware in Late Antique Italy

African Red Slip Ware (ARSW) is a fine tableware commonly found at late antique sites. It was produced in *Africa Proconsularis,* and since the 1960s it has been studied in great detail by many scholars (LAMBOGLIA, 1963; HAYES, 1972; HAYES, 1980; CARANDINI, 1981; FULFORD, PEACOCK, 1984; MACKENSEN, 1993; BONIFAY, 2004). Nowadays its typology and chronology are among the most fine-grained for the whole ancient world. Its wide geographical distribution makes it an ideal medium for the study of trade and commerce.

ARSW is by far the most widely found fine pottery of the Late Roman world, particularly in the Western Mediterranean. Its geographic distribution has been studied by many scholars in the past 4 decades, but almost always at a general level of detail, with some notable exceptions (FONTANA, 1991). This research has instead focused on a specific region, in the hope that using a different scale can shed new light on this well-known phenomena of Late Antiquity. ARSW is not even at its *debut* at CAA, since it was the subject of a paper by Hawthorne (HAWTHORNE, 2000), touching some of the topics that we deal with here, like quantification of pottery and social meanings of ceramic objects.

The specific aim of the present research was the definition of specific aspects of complexity that arise from the geographical and social distribution of ARSW in Northern Italy, starting from the beginning of the fifth century AD until the end of the seventh century, when the production of ARSW, or at least the export outside the Carthage region, came to a definitive end. The research was based on published data, relying also on previously published catalogues of archaeological sites (above all TORTORELLA, 1998). The study region was arbitrarily defined based on the current administrative boundaries of Valle d'Aosta, Piemonte, Lombardia, Liguria, Emilia-Romagna, Toscana and Marche . While this choice may appear completely nonsense, it is arbitrary like any other, particularly if we consider that geopolitical and military boundaries were continuously moving during the three centuries that were being studied.

Furthermore, this vast macro-region has the advantage of including a high variability in geographical and natural characterization, since it spans from one coast of Italy to another, and allows us to have a good sample of the territory that was subject to the Lombard conquest. We will see below in more detail why this was a relevant aspect for this research.

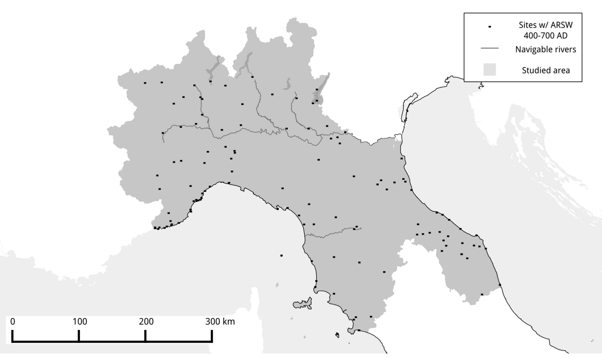

Figure 1: *Map of the region under study and sites.*

3. Choosing the tools

Python is a high-level programming language (http://www.python.org/), it is available for all operating systems and already comes installed on most GNU/Linux distributions and on Mac OS X. Python supports several programming paradigms but it is mainly based on object-oriented programming. Its syntax is simple, particularly when compared with other languages like C++ or Java, and has a gentle learning curve. Furthermore, Python is an interpreted language, and one can thus learn the basics using the interactive interpreter, without any need to compile source code beforehand. These features, together with the large number of third-party libraries developed worldwide by thousands of people, make Python an ideal "swiss knife" for the development of archaeological software, for anything ranging from data management to web publishing and quantitative analysis.

GeoDjango is a popular extension to Django, an open source web framework designed for high-traffic dynamic websites (http://djangoproject.com/), written in Python. GeoDjango is available in the main Django distribution under the django.contrib.gis module. Django uses a structure called Model-Template-View (MTV), in which data models (Model), data handling (View) and visualization (Template), are separately implemented and can be freely assembled. GeoDjango enables the integration of geospatial data into the development of web applications, through its Object-Relational Mapper (ORM). Using the ORM means to avoid using SQL directly, instead data models are defined declaratively in the Python language. Using Django, models were built in a direct and concise way, while data was stored in PostgreSQL/PostGIS underneath (http://www.postgresql.org/ and http://www.postgis.org/). In other words, Django offers an alternative way to interact with a spatially-enabled relational database, but data is stored in a traditional fashion and other programs can safely access the same database for data visualization (an example is explained below).

4. Building a data model during the collection

The data model is relatively simple, and is conceived to hold most information available from published data. A severe problem was encountered while designing a model for quantitative pottery data, because existing international standards (ARCELIN and TUFFREAU-LIBRE, 1998) are often ignored by many authors, and are a recent acquisition. As a consequence, quantitative data has been scarcely used for analysis, but has been nevertheless recorded with the same detail as they were published by the original authors.

The data model keeps the two concepts of "excavation" and "site" separated. This separation allows for more than one excavation at the same site (in several campaigns, or different areas of a large settlement), like in urban archaeology. The atomic unit is the "sherd count", which is a quantification of how many vessels of the same type and ware were found in a single stratigraphic context (in some cases, in aggregate form as a site phase or period). Lots of published find reports bring unclear figures like "some sherds" or "presence of", in these cases the recorded information is merely that a specific type was present in a context. Every data class (*Site, Excavation, Context, etc*) was described as a child class of the base *Model* class, following common object-oriented programming practice. Where needed, constraints of unique values (name of site) or combination of values were added (SherdCount is unique for the *Ware, Assemblage* and *Type* attributes taken together). All

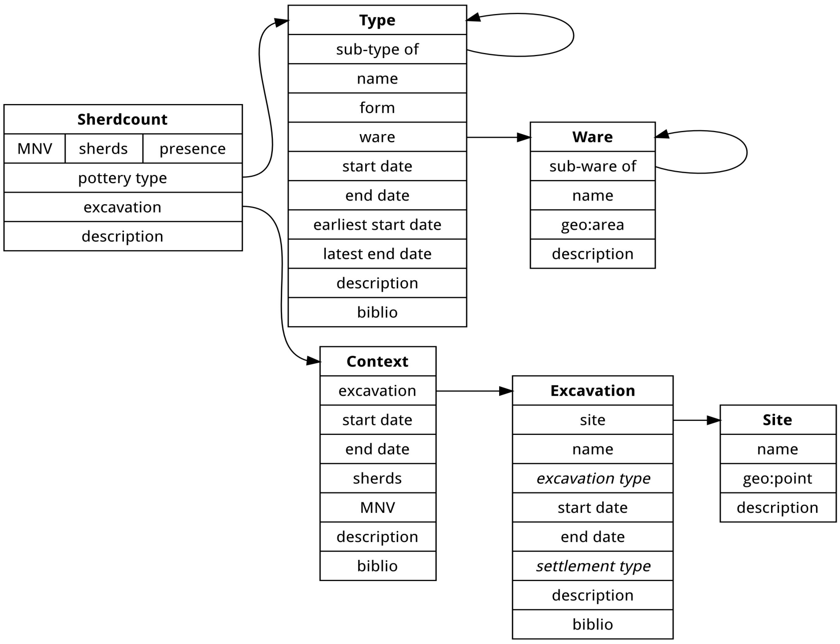

Figure 2: *Data model. Attributes prefixed with "geo" are geospatial. Attributes in italic have a closed list of choices.*

variable and attribute names were defined in English, following Python best practices. In theory, the software could be easily adopted by an international team thanks to its clear and basic structure.

Looking at the interoperability of this data model, we will argue here that it is quite similar to the ArchVocabsemantic ontology (http://archvocab.net/ - see ISAKSEN, 2009) first introduced at CAA in 2009. The ArchVocab Excavation ontology has a hierarchical structure, even if it mandates geolocation to external vocabularies or services. The *SherdCount* object would correspond to the *Find* class in ArchVocab, while other classes are close matches. A mapping to ArchVocab is thus certainly possible, even though vocabularies for the various ARSW typologies that are missing.

Ware and Type objects can be recorded as sub-types of other objects of the same class. This feature is useful for recording variants of a general ceramic type (e.g. Hayes 104A and 104B are sub-types of type Hayes 104), not for the sake of precision but to allow for data with a different level of detail. Lots of sub-types have a more fine-grained dating if compared to general types, or are known to be products of different workshops. In some cases, being unable to specify a sub-type might mean getting a very large timespan, like for example the Hayes 91 flanged bowl, that was manufactured for more than 3 centuries in several variants. Sub-types were used also to bring into a common and consistent framework the different typologies that are used for the classification of African Red Slip Ware. The most used typologies in the study area are those of Hayes, Lamboglia, and the *Atlante delle forme ceramiche*, but also the more recent works by Fulford, Mackensen, and Bonifay are sometimes used for specific forms.

A satisfying data model can be effectively developed only after collecting data, following an iterative process. Django is quite flexible at this respect, and the South application can handle "migrations" (upgrades) to data models, like when adding a new column or changing the type of an existing column (http://south.aeracode.org/). Adding a new model is supported directly by the Django core utilities.

5. Collecting and recording data

Django was chosen primarily for its Object-Relational Mapper and the flexibility that derives from the use of Python in every context. However, there are several additional advantages that come "for free". The most notable piece is the auto-generated administration interface. With less than 10 lines of code, a complete data entry environment is available, based on the data models. This interface is an easy-to-use web application, where each field is restricted to the allowed values, including foreign keys and lists of choices. Unique constraints are enforced. GeoDjango adds geospatial functionality with an OpenLayers map for entering point, linear, and polygonal features on the background OpenStreetMap data (http://www.openlayers.org/ and http://www.openstreetmap.org/). All data were recorded using this interface on *localhost* through Django native development webserver.

Creating such an interface from scratch is a boring, repetitive operation, and nevertheless it seems quite common for archaeologists developing database applications to spend time on it. We strongly encourage the adoption of dedicated tools for this purpose. Together with the extensive use of ORMs, this approach can improve a lot the development process and the mainteinance of code.

Every piece of information can be traced back to its source through common bibliographic references, given that the research was based on published data. The format used for bibliographic references is BibTeX, usually associated with the LaTeX typesetting environment (http://en.wikipedia.org/wiki/BibTeX). The choice of BibTeX might appear inconvenient, observing how many archaeological databases included a bibliographic database inside. However, we had no need to perform bibliographic queries, instead only to visualize citations for each site or context. Building a dedicated module would have been totally superfluous, moreover if we consider how many different kinds of publications exist. Bibliographic references were formatted using the Pybliographer programming library (http://pybliographer.org/). The BibTeX parsing engine was integrated into the application using an extension to the templates used by Django, producing a properly formatted citation in HTML format. Apart from visualizing citations in the web application, BibTeX was used in parallel in the LaTeX environment for advanced typesetting of the dissertation (http://www.latex-project.org/). As an example, a listing of all recorded sites with extensive bibliographic references was automatically created through a Python script querying the database.

Data entry should not be a blind operation, and immediate feedback was a requirement for assessing the progress of data collection and upgrading the data models as needed. For these two reasons, we created a simple exploratory environment, using Django templates and views for a simple, yet complete, web application. This environment includes summary reports such as the lists of all sites and ceramic types, but also detail views for each excavation context, quantitative data for each site, and the theoretical timespan of ARSW imports at each site. This timespan has been used in some further analysis. Dynamic maps appear in the detail pages of each ceramic type, allowing to get a clear preliminary picture.

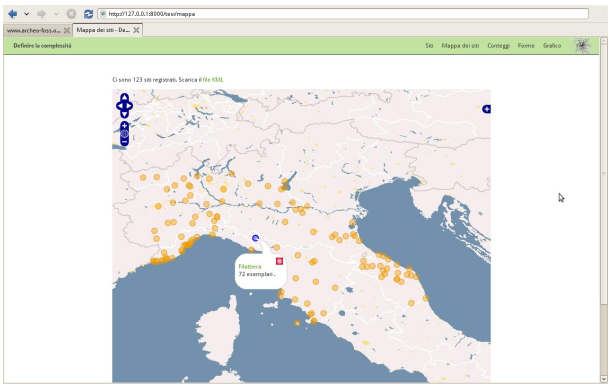

Figure 3: *Dynamic map of sites. There are links from each site to the summary page.*

GeoDjango native capabilities for handling standard geospatial data formats like KML, GML, and GeoJSON has made it easy to integrate with the entire suite of OSGeo software (like OpenLayers and QuantumGIS, http://www.qgis.org/) for advanced cartographic visualization and layout of high-quality maps for printing. For the latter, we used QuantumGIS (QGIS). QGIS is a desktop GIS software, easy to use but with advanced features, capable of reading data stright from PostGIS or Spatialite.

What really makes the difference in adopting Django instead of a custom in-house framework is the pervasive use of Python and of the several programming libraries. The popular Matplotlib was used for the advanced creation of data plots, and other parts of the code were re-used from software we had developed previously (http://matplotlib.sf.net/ Matplotlib is based on the NumPy numerical programming library, that has an interface similar to MATLAB, but is completely free software). Re-use of code is a key aspect of good programming practice, and we found that Python consistently encourages modularity.

The web application is obviously a prototype, yet a working one, that was actually used to follow the development of research and show in-progress results to other people.

Integration with Matplotlib was particularly interesting with regard to real-time generation of plots obtained by crossing existing knowledge about dating of ceramic types with the collected quantitative data. These plots can be used as a framework for visualizing other attributes like the vessel function (individual or collective dish, plate, bowl) and origin. Scatterplots aren't very common for representing chronological information, because they result in a more abstract visualization than other possible plots. However, we found this abstraction useful because it helps focusing on chronology rather than using typology strictly as a classification tool.

Summary pages about single sites include a cumulative plot made with the "weighted means" technique. This kind of visualization has been used mainly for the study of ARSW production, and for summary studies on data resulting from archaeological survey (FENTRESS and PERKINS, 1988; FENTRESS *et al.*, 2004; TERRENATO and RICCI, 1998; ZANINI, 1996). It is rare, even though more frequent recently, to see the application of this technique for the study of single sites. The code implementing this method was re-used from a previously developed software, written in 2007, for the same purpose in a more limited framework. This source code is available as a small stand-alone application at the URL http://bitbucket.org/steko/weighted-means/ and it can read simple quantitative data in CSV format.

In theory, comparing the curve for a single site or region with the general one should be able to assess the relative trend of ARSW imports through time. It is much less common to compare a single site with the region where it lies, probably due to a lack of enough detailed and homogeneous data.

This set of tools is flexible and allows for fast development. A summary view of the 123 recorded sites shows a total number of 3126 vessels, and both the chronological and geographical framework were clear at the general and detailed scale. We were able to early detect some phenomena like the marked absence of ARSW in some regions, even during the peak of its diffusion, the pervasive presence of some ceramic types (like the Hayes 61 B dish) and the more outstanding contexts.

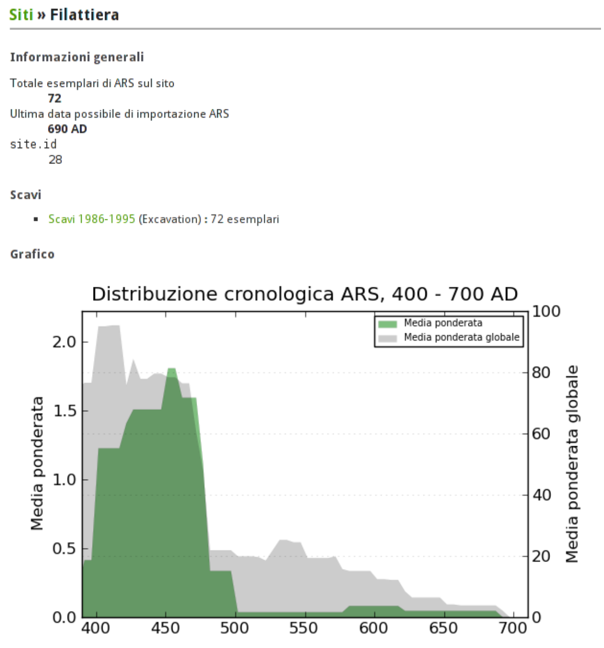

Figure 4: *Screenshot of the web application. The summary page for the Byzantine military settlement in Filattiera shows.*

6. From database to analysis

Studying the distribution of African Red Slip Ware as a socio-economic phenomenon requires a more in-depth analysis of the different aspects that contribute to its the understanding. Descriptive plots are just the first step of an iterative process, and it isn't really acceptable to stop at this point. Otherwise, at the end of the story we will find out that we had just added some "archaeological numbers" on a fixed background of historical knowledge. ARSW has already been studied through advanced quantitative methods, but that was twenty years ago when the number of available sites was considerably smaller than now. In the early 90s Sergio Fontana has published a pioneering application of multivariate statistics to this same topic, using the published data available at that time (FONTANA, 1991). The method adopted by Fontana was applied to the data we had collected. As it was expected, the results are meaningful only for those periods that have at least a certain amount of published archaeological contexts.. Nowadays, the general picture is characterized by a large number of published excavation and find reports, so we can shift our focus from a general view to regional studies.

During data analysis we continuously used the R statistical language for the assessment of the different interpretative models known from the existing literature (http://www.r-project.org/). R is a GNU software and it is largely compatible with the S language, being its most popular implementation. We found that Python and R made up a powerful, automated environment to query, extract and analyse data using some scripts that we wrote for these specific tasks. All scripts allow fine-tuning of the required parameters along with high-quality output plots. Parametrization is helpful to analyze the same data under varying conditions. Scripts allow for strict repeatability of analysis and care was taken in documenting their code for both personal and external use, and they are available under the GNU GPLv3 license at a separate repository than the main application, found at the URL http://bitbucket.org/steko/thesis-stats/. For this analytical work, we found that using an advanced programming environment was important, and we used GNU Emacs (version 23) together with the "Emacs Speaks Statistics" extension, for using the R interactive interpreter and writing code in both R and Python (http://www.gnu.org/software/emacs/ and http://ess.r-project.org/).

Notwithstanding the complexity of data models that are often developed for archaeological purposes, statistical analysis always requires having data in a tabular, flat format. This is a basic assumption, but sometimes it is not recognized. The Python standard programming library includes a dedicated module for reading and creating files in the common Comma-Separated Values format (CSV). CSV was used to export in a tabular form all data that needed to undergo some kind of analysis, like: at how many sites ARSW is found over time, how many vessels and ceramic types are found at each site, how far are sites from the coast or from navigable rivers. R was used as a scripting analytical language, creating short dedicated programs for the study of single aspects. Each program outputs one or more plots as required. We found fascinating to conceive software programming as the construction of a logical and mathematical sequence of routines.

While the archaeological results from this research are not strictly the topic of this paper, we would like to show some of them, to allow a better understanding of the technical discussion above.

Figure 5 shows an epiphenomenon of geopolitical history, namely the difference between those areas that were conquered by the Lombards and those that remained under direct or indirect control of the Byzantine Empire. The importance of the Empire for the prolonged wealthiness of regions such as Liguria and the Exarchate, not to mention Southern Italy altogether, has been stressed by some scholars (ZANINI, 2001), and regarded as secondary by others. For the region under study, we found that the variation in the density of sites where ARSW was present is markedly different between the two influence areas until the mid seventh century AD. The sharp drop at that time goes together with the Byzantine retreat from *Liguria Maritima* and their general loss of positions (and interest) in the most far away regions. We remark that instead of the absolute number of sites where ARSW was found we analysed density, to take into account the ever-changing boundaries between the two opponents. In this situation, absolute numbers were not comparable so we used cartography from available literature (ZANINI, 1998) to create a digital map of the moving Byzantine *limes* from 568 AD to the end of the seventh century.

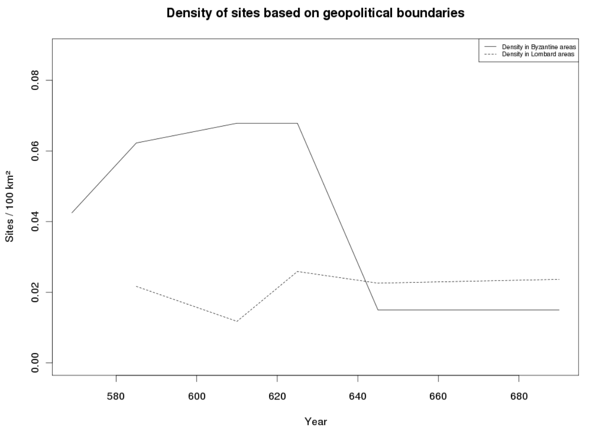

Figure 5: *Density of sites based on geopolitical boundaries.*

We regret to remark that there is a general lack of geographic data about the ancient world, and this is particularly true for Late Antiquity. Unfortunately a large number of scholars end up doing the same work for themselves, and there is little interest in sharing geographic data such as political and administrative boundaries.

For figure 6 we adopted the same method as above, based on site densities, applied to the analysis of geographic factors such as distance from the coast and from navigable rivers. For this purpose, we created a digital map of the late antique coastline based on data from the OpenStreetMap project, combined with data from geological literature for some areas that are known to have undergone significant change in coastline from the Roman times onwards. These areas include the Po delta, the area around Ravenna and the whole coast of Tuscany from Luni to Orbetello. Starting from this coastline, we created a 10 km *buffer* around it (arbitrarily chosen), to select a sub-area where site density over time was calculated. The same steps were then undertaken for rivers, too. The final result is made by three different chronological distributions for coastal regions, areas reached by river, and inland regions. The plot shows the arithmetic ratio between these three density curves and we can spot a remarkable change at the end of the fifth century. Since that time, being near the coast was much more significant than before regarded to the availability of imported goods, of which ARSW is certainly a significant proxy. Also inland waterways continued to play a minor role.

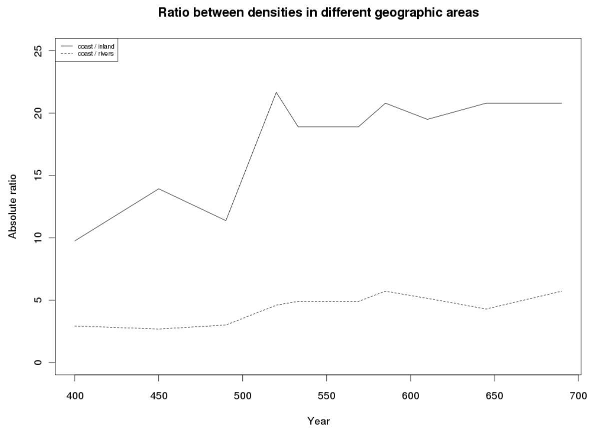

Figure 6: *Ratio between densities in coastal areas vs inland.*

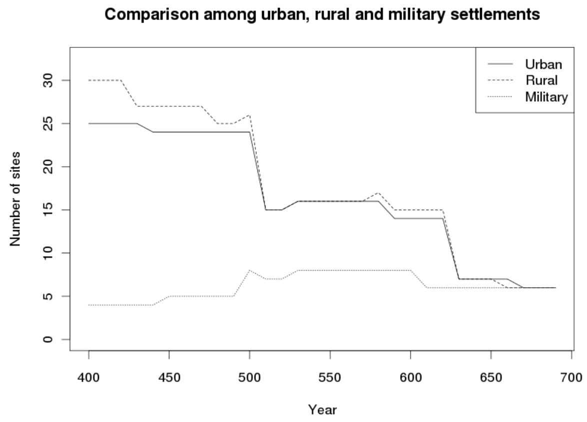

Figure 7: *Comparison between urban, rural and military settlements. Number of sites with ARSW over time for each settlement type.*

A third interesting phenomenon is represented by the decline in imports at both urban and rural settlements, while on the other hand military settlements keep getting a fair amount of ARSW, particularly those under Byzantine control like Sant'Antonino di Perti. A plot is shown in figure 7. We would like to remark also the strong correlation between urban and rural settlements (Spearman's $R^2 = 0.98$). Correlation can be read from a socioeconomic point of view, and it seems to mean that urban sites continued to act as central places for their regions. On the other hand, we ought to look at the absolute numbers and understand that when they become very low, the significance of this correlation is certainly unacceptable. This last issue is particularly strong if we consider that the relationship between cities and countryside should be analyzed at a regional scale, and not at a general level.

However, we recognize that the regional approach would require a much larger database including not only other classes of ceramic finds, but also data from survey projects and a more sound assessment of social contexts. Unfortunately, the work discussed here is arguably the single largest *published* database about late antique and early medieval pottery in Northern Italy.

7. Tracking the research process

Aside from the main research process, other tools were used for tracking changes to source code. However, there are two interesting points here that deserve to be mentioned.

Using code to describe the steps of the process means that we can look at it anatomically.

Once we accept that research is "materialized" as source code, the next step is finding out that tracking changes to source code means intrinsically tracking the research process itself. Version control systems are standard tools for software developers, but in this case we can use them also to allow the process becoming explicit in its evolution, and even to compare it at different times. We used Mercurial, a complete and fast distributed version control system that we already adopted for other software projects (http://mercurial.selenic.com/). Mercurial works without a centralized server (similarly to git, while CVS and Subversion need one), it has a gentle learning curve and has a friendly graphical interface for new users. Using it from the command line is however the most powerful and versatile way to keep under version control any software program, regardless of the programming language that is used.

This same concept can be easily extended to an entire research team, as happens with a wiki (ZANINI and COSTA, 2006). Having not only data, but a whole research process digitally encoded makes it possible to track this process and to record the multivocality that is implicit in a collective work.

Conclusions

This research, briefly summarized here, regarding to its technical part, actually reached some interesting archaeological results for the understanding of Late Roman Italy and of its trade patterns, socioeconomic structure, and how political changes affected them. Most importantly, for the scope of this conference, we demonstrated that, provided a sufficient level of technical skill with high-level, object-oriented programming languages, free and open source software can be not just *used* but *produced* to fulfill the needs of a specific aim. Choosing effective tools is key, and the interoperability among all of them is equally fundamental. Moreover, using and developing free and open source software basically means that we're building open systems, where data is stored but not locked.

Future developments of this research include re-evaluating the data model to allow interoperability with other existing archaeological databases, also using semantic web technologies.

References

ARCELIN, P. and TUFFREAU-LIBRE, M. (edd.), La quantification des céramiques, conditions et protocole. Actes Table-ronde Centre Arch. Mont-Beuvray, 7-9 avril 1998 (Bibracte 1998).

CARANDINI, A. (ed.), 1981. *Atlante delle forme ceramiche I. Ceramica fine romana del bacino mediterraneo*. Istituto per l'Enciclopedia Italiana, Roma.

BONIFAY M. 2004. Études sur la céramique romaine tardive d'Afrique. Archaeopress, Oxford.

FENTRESS, E. and PERKINS, Ph., 198 1988. Counting African Red Slip Ware. In: Mastino, A. (ed.) *L'Africa Romana*. Atti del V convegno di studio, Sassari 11-13 dicembre 1987, pp:205-214.

FENTRESS E. *et. al*, 2004. Accounting for ARS: Fineware and Sites in Sicily and Africa. In: Alcock, S. E., Cherry, J. F. (ed.) *Side-by-side survey: comparative regional studies in the Mediterranean world*, pp:147:162. Oxbow Books, Oxford.

FONTANA S., 1991. Analisi comparata delle attestazioni della ceramica africana nel V sec. d.C.: un'indagine preliminare. *Archeologia e Calcolatori 2*, pp:109-121.

FULFORD M.G. and PEACOCK D.P.S. 1984. *Excavations at Carthage. The British Mission, I, 2. The Avenue du President Habib Bourguiba, Salambbo: the pottery and other ceramic objects from the site*. Sheffield.

HAWTHORNE, J., 2000. Vessel volume as a factor in ceramic quantification: the case of African Red Slip Ware. In: Lockyear, K., Sly, T. J. T., Mihăilescu-Bîrliba, V. (edd.) *CAA* 96. Computer Applications and Quantitative Methods in Archaeology, pp: 19-23.

HAYES J. W. 1972. *Late Roman Pottery*. Ann Arbor, Michigan.

HAYES J. W. 1980. *A supplement to Late Roman Pottery*. Ann Arbor, Michigan.

ISAKSEN, L., MARTINEZ, K., GIBBINS, N., EARL, G. and KEAY, S. 2009. *Linking Archaeological Data*. Available at http://eprints.ecs.soton.ac.uk/18240/ [Accessed 25 May 2010].

LAMBOGLIA N. 1963. Nuove osservazioni sulla "terra sigillata chiara" II. Tipi C, lucente e D. *Rivista di Studi Liguri 29*, pp:145–212.

MACKENSEN M. 1993. Die spätantiken sigillata– und lampentöpferein von El Mahrine (Nordtunesien). Munchen.

TERRENATO, N. and RICCI, M., 1998. I residui nella stratificazione urbana. Metodi di quantificazione e implicazioni per l'interpretazione delle sequenze: un caso di studio dalle pendici settentrionali del Palatino". In: Guidobaldi, F., Pavolini, C., Pergola, P. (edd.) *I materiali residui,* pp: 89-104 . École Française de Rome:Roma.

TORTORELLA, S., 1998. La ceramica africana in Italia. Problemi di circolazione e distribuzione . In Saguì, L. (ed.) Ceramica in Italia VI-VII secolo. Atti del convegno in onore di John Hayes, pp:42-70. Edizioni del Giglio, Firenze.

ZANINI, E., 1996. Ricontando la terra sigillata africana. *Archeologia Medievale 23*, pp:677-688.

ZANINI, E., 1998. *Le Italie bizantine*. Edipuglia, Bari.

ZANINI, E. and COSTA, S., 2006. Organizzare il processo conoscitivo nell'indagine archeologica: riflessioni metodologiche ed esperimenti digitali. *Archeologia e Calcolatori 17*, pp:241-264.