

The World Spatiotemporal Analytics and Mapping Project (WSTAMP): Further Progress in Discovering, Exploring, and Mapping Spatiotemporal Patterns Across the World's Largest Open Source Data Sets.

Jesse Piburn^{a,*}, Robert Stewart^a, Aaron Myers^a, Alex Sorokine^a, Eric Axley^b, David Anderson^b, Jordan Burdette^c, Christian Biddle^b, Alexander Hohl^d, Ryan Eberle^e, Jason Kaufman^a, and April Morton^a

^aOak Ridge National Laboratory, 1 Bethel Valley Road, Oak Ridge, TN 37831 - (stewartrn, piburnjo, myersat, sorokina, kaufmanjc, mortonam)@ornl.gov

^bCadre 5, 804 Innovation Drive, Knoxville, TN 37932 - (eric.axley, david.anderson, christian.biddle)@cadre5.com

^cIngenium Professional Services, 136 S. Illinois Ave. Suite 202 Oak Ridge, TN 37830 - burdetteja@ornl.gov

^dUniversity of North Carolina at Charlotte, Department of Geography, 9201 University City Blvd, Charlotte, NC 28223-0001, ahohl@uncc.edu

^eCarson-Newman University Mathematics Department, 1645 Russell Ave., Jefferson City, TN 37760, rseberle@cn.edu

KEY WORDS: Global, Data Mining, Spatio-Temporal, Visualization, Tool, Analytics.

ABSTRACT:

Spatiotemporal (ST) analytics applied to major data sources such as the World Bank and World Health Organization has shown tremendous value in shedding light on the evolution of cultural, health, economic, and geopolitical landscapes on a global level. WSTAMP engages this opportunity by situating analysts, data, and analytics together within a visually rich and computationally rigorous online analysis environment. Since introducing WSTAMP at the First International Workshop on Spatiotemporal Computing, several transformative advances have occurred. Collaboration with human computer interaction experts led to a complete interface redesign that deeply immerses the analyst within a ST context, significantly increases visual and textual content, provides navigational crosswalks for attribute discovery, substantially reduce mouse and keyboard actions, and supports user data uploads. Secondly, the database has been expanded to include over 16,000 attributes, 50 years of time, and 200+ nation states and redesigned to support non-annual, non-national, city, and interaction data. Finally, two new analytics are implemented for analyzing large portfolios of multi-attribute data and measuring the behavioral stability of regions along different dimensions. These advances required substantial new approaches in design, algorithmic innovations, and increased computational efficiency. We report on these advances and inform how others may freely access the tool.

1. INTRODUCTION

Stewart et al. (2015) introduced the World SpatioTemporal Analytics and Mapping Project (WSTAMP) and detailed much of the background information on the WSTAMP project. This paper is meant to serve as an update on the status and developments since this last review, therefore many items detailed in the previous paper are omitted from this one.

Specifically, this paper will introduce recent updates to the WSTAMP ontological schema, both theoretically and technically, briefly highlight two analytical developments, point to user experience enhancements in the WSTAMP online tool, and layout selected next steps for the project. First, a brief outline of what WSTAMP is and what it is meant to accomplish.

1.1 Background & Motivation

WSTAMP seeks to provide contextually situated spatiotemporal understanding on the evolution of socio-cultural, socio-environmental, and geopolitical landscapes on a global level through access to international datasets, tailored spatiotemporal analytics, and ontologically tethered entities that relate data through time and space.

A central component of WSTAMP, the project, is the WSTAMP online tool where these goals are realized through a browser based analytical interface. The WSTAMP tool allows users to engage arbitrary space-time data in an analytical environment, but understanding why we are engaging this data informs how we do it.

The only reason data is at all useful is because we can use it to help inform and guide our understanding of how some process occurred and, when coupled with reflective analysis, attempt to answer the much more difficult question of why. Whether this knowledge is used to clarify the past, operate in the present, or anticipate the future, we are faced with the challenge of trying to make sense of what has been observed.

Understanding of a process, however, does not come from knowing a single fact about it, rather it is achieved from a systematic interconnection of facts, the full meaning of which, is discovered in the implications it has on other processes. If I have 7, is that good? Well it depends. 7 of what? What was it yesterday? What is everyone else's? Who else has 7? Who has more? To answer this question, we need to know more than just the fact that I have 7, we need the interconnection of other facts. We need context. And context is at the center of how WSTAMP is organized.

As our ability to collect data outpaces our ability to find, sort, organize, and make sense of it, contextual information is being disregarded as the data itself is becoming an obstacle to understanding the process it is measuring. The goal of WSTAMP is to ground all facts in the interconnection of others, to ground all answers in the context that gives meaning to them.

This is a serious challenge particularly when accessing the empirical view of data science and the technological tools of computer science. These alone are not sufficient to address a comprehensive view of not only how I see the world and its history but how others interpret, organize, and act on the world as well. For this we must call on the domains such as geography,

*Corresponding author

history, political science, and philosophy; treating data and computer science as useful tools to help us find an answer, not as the answers themselves.

2. DATA

In our introductory paper (Stewart et al., 2015) we present two primary challenges in developing a spatiotemporal capability that is tightly connected to long-horizon geographic histories. Briefly reviewing that work, we divide this challenge into two separate but interconnected endeavors: 1) how do we settle on, store, and represent the collection of geographic entities (e.g. nations, territories, regions) and 2) how to elegantly capture and record a variety of things that happen (e.g. merger, split) or values (e.g. birth rate, trade) from multiple disparate sources. Our solution was to envision the world as comprised of a set of entities and things that happen to those entities are events. For example, Sudan is a world entity that over time produced measures of GDP, birth rate, and population. Sudan also split into Sudan and South Sudan, however the entity that is Sudan, although changed, persists. A new annual report for GDP, for birth rate, and for national succession are all considered ontologically as events for our purposes. Indeed, we demonstrate through ontologically rich structures how entities relate to each other of a long history of splits, mergers, etc. In this way, the WSTAMP data structure is organized around only two principles: entities and events.

2.1 Entities

At the time of the last update, all entities in WSTAMP were limited to the national level. This was in part due to project scope and the availability of data but there's no ontological reason why entities that are recognized as nations are more important than non-national entities. Continuing our efforts, we have now expanded our tracking to other, non-national entities, as well. Most prominent in this category are entities such as admin1 and admin2 sub-national administrative units, but this also includes entities that are themselves defined as collections of other entities such as, Environmental Protection Agency (EPA) regions being defined by U.S. States, as well as, entities that are not traditionally represented by the notion of regions or areas, such as rivers or weather stations.

With the expansion of our domain of discourse it is useful to describe entities in a more abstract manner. In general, an entity is any physical, social, or theoretical construct that exists now or in the past. These include traditional geographic entities such as admin0, admin1, and admin2 regions but also massless entities such as Facebook users or specialized entities such as a meteorological weather station. These entities are governed by the concept of perspectives. Perspectives offer a world view on what entities exist and how they are related to one another.

In WSTAMP we derive perspectives, and therefore each perspective's member entities, from the sources from which we collect data. An entity *must* exist in at least one perspective but *can* exist in several, for example, Nigeria exists in both the United Nations and World Health Organization perspectives. Furthermore, if a source does not collect data on an entity, it is considered to not exist in that perspective.

For example, the World Bank has a collection of entities for which it publishes data. While the World Bank is certainly aware of the U.S. State of Nevada, since the World Bank does not publish data for Nevada it is considered to not exist in World Bank perspective.

Some entities imply a hierarchical relationship, Admin2s divide completely Admin1s, which themselves divide completely Admin0s, while others are multi-hierarchical, Tennessee is a child of the US and also of EPA region 4. Still others have fuzzy or partial relationships as in Tennessee is partly a member of the Tennessee Valley. We respond to this set of varied relationships with the concept of membership. By relaxing the notion of hierarchy, we arrive at a much more flexible and consistent way to organize, group, and relate entities in the world, throughout time and across perspectives.

Furthermore, introducing the notion of membership to all entities, given that a perspectives' entities can be related to at least one other already existing group of entities, it is possible to propagate the membership structure across all entities and perspectives and infer the membership across entities without explicitly stating their relationship. For example, the EPA regional perspective says nothing about cities or counties, but is defined by U.S. States. Since we know that Trousdale County is a member of the state of Tennessee, we therefore know, without it being manually defined, that Trousdale County is in EPA Region 4. The same can be done for cities, rivers, Facebook users or any other entity. In this simple example, it seems trivial but the capability of doing this through time and across multiple perspectives represents a significant capability.

A quick note on boundaries. The geometric representation of an entity is itself an attribute of the entity, not an intrinsic quality of the entity and needs to be considered as such. Boundaries that represent an entity, need to be considered separately from the entity itself. There can be several potential geometric representations for one entity arising from data driven reasons such as the resolution the geometry was drawn at, a city can be represented by a point or a polygon with varying degrees of detail, to more nuanced reasons such as conflicting definitions of what the entity itself is, to disputed or undefined boundaries between entities.

2.2 Events

Entities are interesting to us because they have events, in fact they exist only because an event as occurred. For our purposes, we distinguish between two types of events, existential and observational. Existential events are events that alter in some way the existence of the entity itself, such as the unification of Yemen in 1990, the dissolution of Yugoslavia from 1989 - 1992, or the independence of South Sudan, and simultaneous continuance of a pre-existing but now altered Sudan, in 2011. Observational events however are events that are produced by the entity itself or by the entity's associated qualities that are relevant for some duration in time, such as the people living within the entity, its geography, or its economic production. Examples of observational events are measures such as Italy's population count in 2015, the humidity in Tokyo, Japan on December 1 2016 at 6:24 PM, or unemployment rate in Sumner County, Tennessee for June, 2017.

Observational events are predicated on existential ones. How could an entity have a GDP if that entity doesn't exist? Existential events therefore limit the potential duration of any observational event. Many observational events however are measured at much shorter intervals than the entire existence of an entity. The United States produces estimates of unemployment rate at annual, quarterly, and monthly relevancies, and although too burdensome to track has some *true* value of instantaneous unemployment at this exact moment (assuming the entity that is the United States, still exists at the time you are reading this). This highlights an

important adaptation for the WTSAMP schema. All observational events have an associated duration for which they are relevant.

Annual estimates of birth rate are annual “events” that have a duration, a relevancy of 1 year. Some events happen on longer or shorter cycles, happen sporadically, and have durations that vary considerably. Furthermore, some events are estimates of a continuous progress while other events occur and are instantaneously over. For example, the unemployment rate for the United States is continuously occurring and our monthly, quarterly or yearly measures of this process are events which estimate the its value over a particular duration and due to its continuous nature can be estimated on a regular cadence. On the other hand, a terrorist attack has no reliable cadence and can be instantaneous in its duration.

Events are not restricted to being produced by one entity, in fact some of the most transformative, and devastating, events throughout history have been events that were born through the interaction of multiple entities. WSTAMP handles these events by including a “from entity” and “to entity” with every observational event. For example, the total value of exports from Bolivia to Brazil for 1999 would have a “from entity” of Bolivia and a “to entity” of Brazil. Handling observational events relating only to just one entity, such as birth rate, is simply handled by having the from and to entity being the same.

Just as entities can be related to one another through the idea of membership, we to can use the same framework, to represent the relationship observational events have to other observational events through perspectives and collections of entities. The Allied invasion of Normandy in 1944 was an event with its own duration and involved entities. That event was itself apart of another event known as World War II with its own duration and involved entities.

All observational events have the potential to have member events, but as in the unemployment rate example, it is the function of importance and burden that determines whether it is explicitly tracked.

2.3 Database Implementation

To track entities and events as described above, several changes to the WSTAMP database schema needed to occur. Including how entities are uniquely identified, how they are mapped between perspectives, and how non-national entities are populated.

2.3.1 Tracking Entities and Events

One of the goals on new identifier system is to trace the same entities between database versions. Most other datasets use simple integer identifiers, but this makes maintenance tasks very hard and may cause lots of errors. Some other datasets (like Correlates of War) use abbreviations or numerical codes, but that causes problems when the entity disappears or changes its name. For example, a few years ago, the updated version of the database retroactively renamed a couple of countries, making the two versions incompatible.

Until version 8 of the database schema, we did not have strict conventions for identifiers of entities and events. Previously when only considering national level entities, in most cases, we used a full country name that remains unique within a single perspective. There are multiple problems with that

approach. First, these identifiers sometimes are very long, which wastes space, error prone, and difficult to maintain. Second, they are rather arbitrary, because typically an entity has multiple names and multiple spellings of the same name. Third, there are some entities with very similar names. A separate system of identifiers for the entities was used by the front facing WSTAMP online tool inside the baskets. These identifiers were not visible from the backend, which lead to broken logic of the baskets after database updates.

We have examined multiple approaches of how identifiers can be created. For countries, we cannot use commonly used identifiers such as ISO, ANSI, FIPS or USPS codes, because there are numerous entities (historical and other) that do not have such codes. As an alternative to ISO codes, we have investigated several alternative approaches:

1. One method is the use of standard phonetic abbreviation algorithms like Soundex and Metaphone, but they cannot be used here without significant modification. In most cases, country names are represented as rather long phrases that are not abbreviated well by these algorithms. There are also many similar country and attribute names that will be abbreviated to the same representation.
2. Another method is to abbreviate the names by the first letters of significant words and then check for overlap with competing, existing identifiers.

Eventually, we decided on human-readable mnemonic abbreviations for identifiers to simplify and improve robustness of the database maintenance tasks. Examples of such identifiers include FRA for France, NZ for New Zealand, and so forth. Our identifiers are similar to commonly used ISO or ANSI codes but not the same. In this approach, identifiers should be short but not necessarily very short. They should be understandable and help to avoid collisions between similar names and common errors, and they should be easily typed on a standard US keyboard using only 7-bit ASCII.

Generation of the identifiers now is mostly automatic and they are loosely based on ISO 3-letter country codes. During the generation process, identifiers are checked manually and are editable by the operator. In the current version of the database, all world entities are identified using the new format of the identifiers at the national level. Each perspective has its own set of identifiers that are matched together.

2.3.2 Matching Between Perspectives

All perspectives that exist in the database were matched to a single WSTAMP perspective (matches are stored in the equivalents table). For each entity, we found a corresponding entity in the WSTAMP perspective and specified a time-period when this match is valid. The matches were performed using one of the following methods:

1. Matching the country names and synonyms and manually verifying each match
2. For the entities that were not found to match via their names, we have used a purely manual approach using agreement among multiple source of information, such as CIA World Factbook notes, the web site of the Office of the Historian of the U.S. State Department, and Wikipedia.

- For some entities, we have used geographic overlap as a match indicator for specific periods of time as supplementary information.

The matches are persistent and will be preserved between database updates. They can be examined along with historical information using internal web interface for the backend database.

2.3.3 Populating Non-national Entities

In cases of non-national entities, we encode the identifier the same as national ones if there is no hierarchical nature to them. For example, the Rio Grande watershed observed from a USGS dataset would be encoded USGS:RG. For entities that do have a hierarchical nature to them, this information is included in the identifier. This hierarchy is included because at the level below national there are several examples of entities that use the same names, in the U.S for example there are 31 Washington Counties, 26 Jefferson Counties, and 25 Franklin Counties. To handle this situation, we prepend the perspective's identifier to the entity identifier to avoid confusion of the entities belonging to different perspectives and then continue this down to the name of the entity that we are identifying. For example, The U.S. Census Bureau's perspective on the entity of Washington County, TN would have the following identifier: CEN:US.TN.WASH.

The biggest problem with the use of mnemonics in identifiers is that they do not always reflect the state of the world at the time when they are used compared to the time when they are created; this is why they are not exposed to the end user. For example, a province may change its country. If, in 2016, we create an identifier for Saar (part of Germany) as DE.SAAR, then this identifier has to remain the same even for the period of 1945-1956 when Saar was a part of France. This happens because identifiers track the actual entity while its history is tracked in a different part of the database. This could create confusion for the end users if outwardly available, but we anticipate that database administrators can easily handle this situation.

This creates the scaffolding for tracking non-national entities. The remaining challenge is resolving external sources of information about non-national entities and identifying reliable sources for recording their long term ontological evolution. We identify the following open issues for effective handling of non-national entities:

- Entity identity and mode of existence
- Spatial region and type of individuation (e.g., vague or disputed boundaries)
- Temporal region(s)/context
- Relations to other entities
- Uncertainty and incomplete information
- Multiple perspectives on the associations of non-national and nation-level entities

At present, the LandScan Global product includes sub-national entities (ADMIN 1) that are adopted for WSTAMP. LandScan collections, however, are only available back to 2012. One potential source for continuing this ontological connection further back is the Global Administrative Unit Layers (GAUL) produced by the Food and Agricultural Organization.

3. ANALYTICS

Collecting data and tracking the evolution of entities are a necessary but not sufficient endeavour for contextually situated spatiotemporal understanding. In the face of this, WSTAMP is focused on developing analytic methodologies that help facilitate the exploration, analysis, and understanding of the entities and events for which data is collected.

Two analytics that were recently developed and implemented in the WSTAMP tool focus on exploratory analysis and helping a user find and sort through large amounts of spatiotemporal trends. The first, called Find Signature Trends, groups similar trend behaviors together into clusters and the second, Attribute Stability Index, was developed to help find erratic or unpredictable behaviors.

3.1 Find Signature Trends

The Find Signature Trends analytic uses a non-linear time series data mining method, known as Dynamic Time Warping, to calculate a distance measure between all pairs of trends in your analysis. Hierarchical clustering is then used to sort the trends into groups based on their temporal similarities. From these groups, an average behavior is computed as the representative trend for that group.

Trying to hypothesize about the spatio-temporal relationships among hundreds or thousands of locations each with potentially hundreds or thousands of time series attributes can be an extremely difficult endeavour when one looks through each location's temporal trends individually, as one would do when looking through successive time series charts. The Find Signature Trends analytic allows users to take numerous individual trends and organize them into groups of signature behaviors, thereby providing an understanding of the emergent trends across the areas of interest (Figure 1).

This analytic was designed to allow users to take a macroscopic approach to understanding spatio-temporal behaviors. By grouping numerous trends into a much smaller comprehensible amount of signature trends a user can quickly get an understanding of the overall types of temporal behaviors before taking a deeper dive into the specific trends for each location and attribute. This broad to narrow approach is visually facilitated by first displaying the signature trends and allows users to deep dive into each signature trend to see the actual trends that were grouped together. The map displays a true spatio-temporal understanding of the trends, because the colours on the map do not represent a single value; rather, they represent a temporal behavior. Displaying these trends on the map allows users to see the spatial distribution of these temporal behaviors.

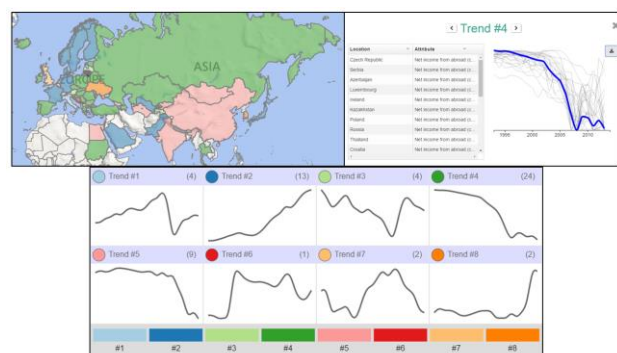


Figure 1. Find Signature Trends

3.2 Attribute Stability Index

A second analytic that was developed and implemented is called the Attribute Stability Index. Identifying erratic or unstable trends is a common question of inquiry. Several well-known methodologies exist for finding erratic time series however these advances rely on having numerous temporal observations. This creates a challenge when attempting to apply these methodologies to trends with much fewer temporal observations such as for spatio-temporal socio-cultural data, where a typical trend of interest might only consist of 20-30 annual observations. Attribute Stability Index is the result of trying to address this need.

To identify unstable trends, we must first define unstable. Instability, by our operational definition, is marked by two characteristics, 1) how widely varying the values are and 2) how predictable that variance is from one observation to the next. That is to say, using variance alone does not suffice, a widely varying but perfectly predictable behavior isn't of concern. Additionally, values for a trend may be unpredictable from one observation to the next but if the resulting values are all within a small window of potential values then it also isn't marked by instability. It is the interaction between these two characteristics that we are interested in identifying. A visual example of these two characteristics are shown in Figure 2

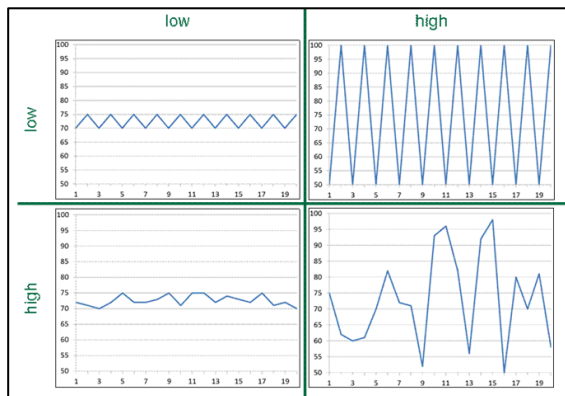


Figure 2. Examples of the interaction of variance and unpredictability

An approximate entropy based methodology is used to characterize each trend with a stability curve and summarized value. An example of this methodology is shown below with the actual trends shown in Figure 3 and the resulting Attribute Stability curves shown in Figure 4. Estimating the integral for each attribute stability curve results in the attribute stability index value. With the larger the value indicating the more erratic and unstable behaviors.

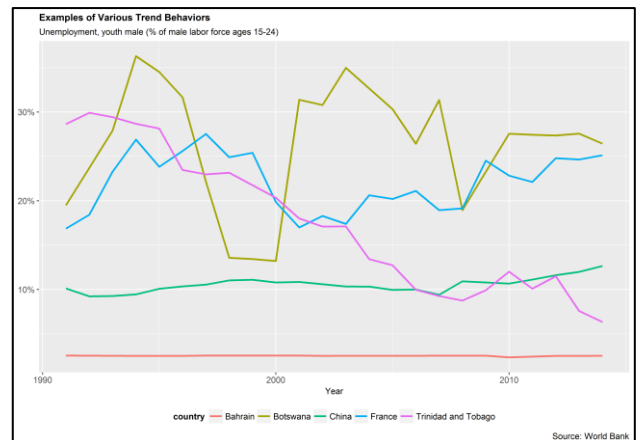


Figure 3. Examples of Various Trend Behaviors

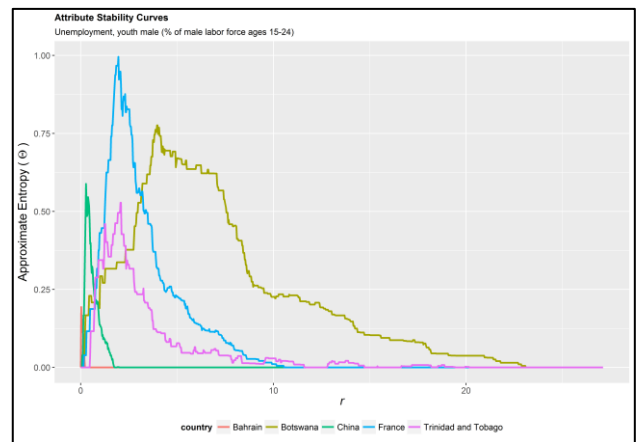


Figure 4. Attribute Stability Curves of Examples Trends

For more information on the attribute stability analytic please see, Piburn, Stewart, Morton (2017).

4. USER INTERFACE

Even with useful analytics to leverage in the task of data exploration, the efficiency and effectiveness of this pursuit is largely influenced by the way a user is presented with performing the task. This is doubly true for browser based analytic environments such as the WSTAMP online tool. Several substantial improvements to the user experience and multiple new features have been added to the tool. Below we discuss a selected few.

4.1 Persistent Context

Heavy use of the previous version of the tool demonstrated that the process of selecting geography, selecting attributes, and exploring trends was highly iterative and interactive. Specifically, the selection of countries, attributes, and time scales relies on the ability to rapidly see the completeness over the cube as well as first order analytics such as trend lines, summary statistics, and outliers. In the previous version of the tool, this required constant movement between tabs (select geography, select attribute, explore, analyse), visually losing important content in previous tabs. This disruption was exacerbated by the presence of pop-up windows requiring user input for every execution of an analytic. Figure 5 shows the previous tabbed version.

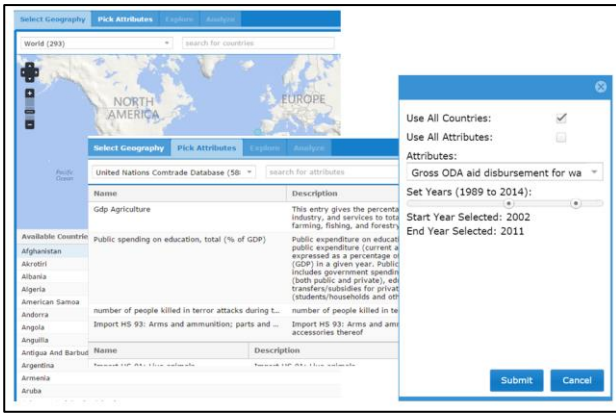


Figure 5. Previous WSTAMP tool interface

The revamped interface situates these first three steps in common virtual real estate, tremendously reducing disruptive windows and creating an environment where the interplay and cross pollination among these three elements is always present. Figure 6 shows the presence of space, time, attribute, and exploratory analytics all in a single persistent space.

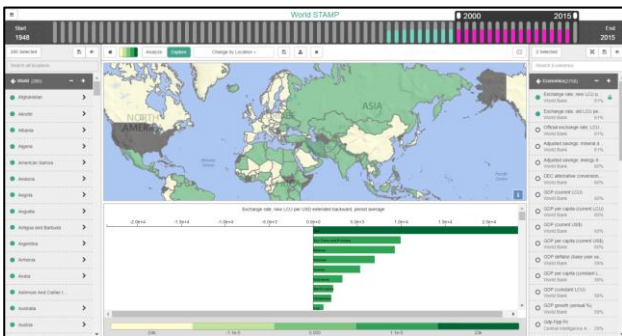


Figure 6. Current WSTAMP tool interface

This iterative process is sped up significantly by the tool not only being interactive, but reactive to the user. WSTAMP now automatically runs analytics for a user's selections every time a selection is altered. When a user has adjusted their currently selected time, locations, or attributes, by the time their attention has returned to the analytic real estate on the screen, the analytics have already been updated.

To enable this responsiveness, a significant amount of effort was focused on performance improvements to allow the interface to rapidly respond to user actions. Including, analytic optimization, novel solutions for data caching, and infrastructure performance tuning.

Another significant improvement to aid in searching available data was the development of an advanced search feature. This feature allows the user a more significant amount of screen space to search for relevant attributes. In this view, shown in Figure 7, users have access to a robust filter set and the ability to explore attribute lists with greater details.

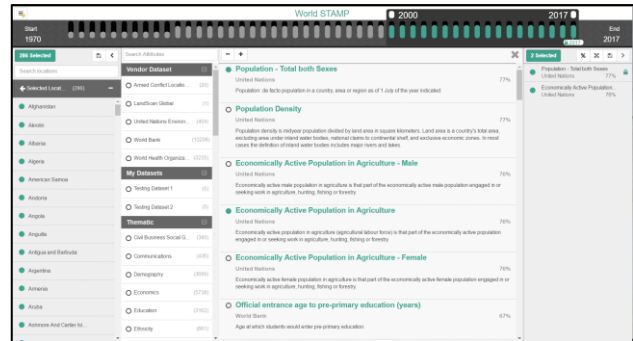


Figure 7. Expanded attribute search feature

Additionally, the “analyze tab” from the previous version was reimagined in a similar way. The full graph-map-table structure was retained but with easy access to geography and attribute panels on the left and right of the tool for quick country and attribute selection during the analysis (Figure 8). Furthermore, all panels are resizable and can expand to full screen.

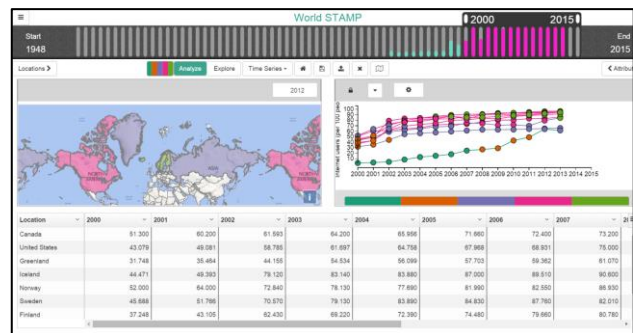


Figure 8. WSTAMP 'Analyze' interface

4.2 Country History

In order to provide additional context to users, the ability to explore an entity's existential evolution has been revealed through the interface. By clicking on the entity in the Geography panel, a window shows the historic information about its evolution (Figure 9). This information is pulled from the existential events that were discussed earlier. Further work will continue to provide contextual events in this interface that are not strictly existential, such as political elections, natural disasters, and other relevant contextual information.

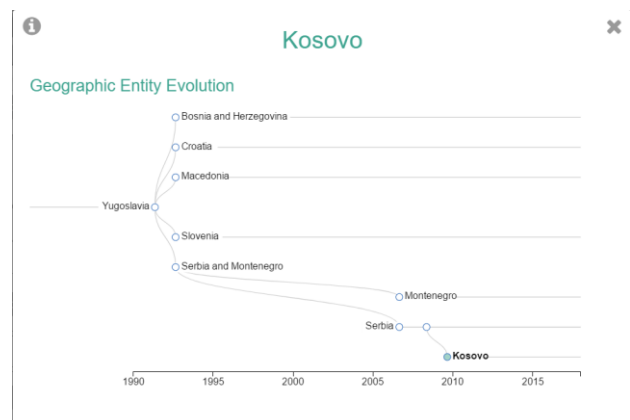


Figure 9. Geographic Entity Evolution

4.3 Attribute Details

When exploring attributes, a user now has additional information about the attribute of interest. By clicking on an attribute name in the attribute panel, a window is displayed (Figure 10) that will give a detailed overview of that attribute, including description, summary statistics, and completeness. A list of related attributes is also provided along with baskets where the attribute has been previously used.

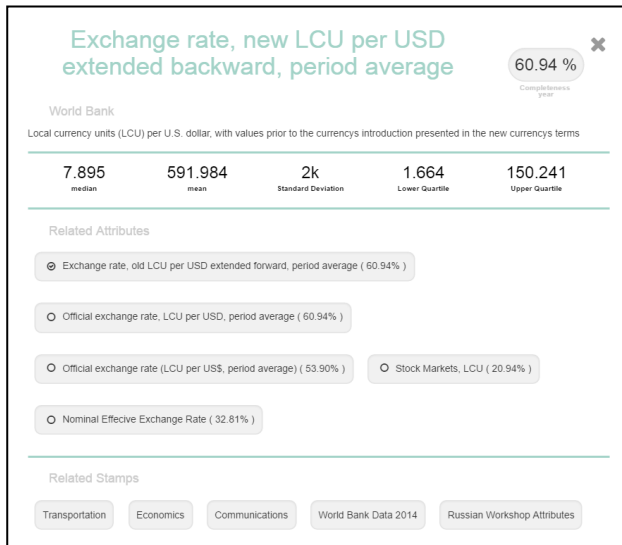


Figure 10. Attribute information interface

4.4 User Uploaded Datasets

WSTAMP now has the capability for users to upload their own datasets and seamlessly integrate their data along with the other sources WSTAMP provides, such as the World Bank and the World Health Organization as well as other user uploaded datasets. Once uploaded the full suite of analytics and visuals are available for use.

5. NEXT STEPS

WSTAMP as a project, not just the online tool, is taking deliberate steps towards becoming a full spatiotemporal and geo-historical information and science system. This is not simply an engineering task. It is not just a matter of collecting data and providing some statistical statements about them, that is only the means to an end, it is about doing this with an intent to further our knowledge of the world and the processes which take place on it. To accomplish this, the current prevailing spatiotemporal epistemologies and geographic data science methodologies cannot be adopted without first critiquing their place within the discipline.

Beginning with this end in mind, we focus our efforts on the process by which we will achieve this goal. WSTAMP is currently undertaking the development of a multi-perspective global administrative boundary ontology, focusing on ingesting datasets from sources such as the Armed Conflict Location and Event Data Project (ACLED) and the Demographic Health Surveys Program (DHS), exploring graph theory based analytics designed specifically for spatiotemporal data, continuing to engage Human Computer Interaction (HCI) experts to improve user experience, and continuing to contribute to the broader discipline through peer-review and academic engagement.

6. CONCLUSION

In this paper, we have updated the status of the World SpatioTemporal Analytics and Mapping Project (WSTAMP). WSTAMP as a system seeks to provide contextually situated spatiotemporal understanding through access to global datasets, tailored spatiotemporal analytics, and ontologically tethered entities. We provide details on the evolution of how WSTAMP tracks and organizes entities and events, include overviews of newly developed analytics, and update the progress of the WSTAMP online tool. Identified next steps include expanding the types of entities included in the WSTAMP schema, researching novel analytic techniques, and expanding the functionality of the WSTAMP tool. Finally, as contextual critique is essential to the understanding of the data in WSTAMP, further work must continue in critiquing the context of WSTAMP itself.

ACKNOWLEDGEMENTS

This manuscript has been authored by employees of UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the U.S. Department of Energy. Accordingly, the United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

REFERENCES

- Piburn, J, Stewart, R, & Morton, A (2017). An Approximate Entropy Based Approach for Quantifying Stability in Spatio-Temporal Data with Limited Temporal Observations. *21st International Conference on GeoComputation*. Leeds, U.K. (accepted)
- Stewart, R, Piburn, J, Sorokine, A, Myers, A, Moehl, J, & White, D (2015). World Spatiotemporal Analytics and Mapping Project (WSTAMP): Discovering, Exploring, and Mapping Spatiotemporal Patterns across the World's Largest Open Source Data Sets. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2(4), 95.