IBM PAIRS: Scalable big geo-spatial-temporal data and analytics as-a-service

Siyuan Lu*, Hendrik F. Hamann
IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, USA

Table of Contents

Abstract

The rapid growth of geospatial-temporal data from sources like satellites, drones, weather modeling, IoT sensors etc., accumulating at a pace of PetaBytes to Exa-Bytes annually, opens unprecedented opportunities for both science and industrial

* Corresponding Author. Email: lus@us.ibm.com

applications. However, the sheer size and complexity of such data presents significant challenges for conventional geospatial information systems (GIS) which are supported by relational geospatial databases and cloud-based geospatial services based on file systems (manifested as object stores or "cold" tape storages).

To fully exploit the value of geospatial-temporal data, particularly by leveraging the latest advances in machine-learning (ML) and artificial intelligence (AI), a new paradigm for platforms and services is required. Some of the necessary salient features include: (i) scalable cloud-based deployment capable of handling hundreds of PetaBytes of data, (ii) harmonization of data in order to mask the complexity of data (schema, map projection etc.) from end users, (iii) advanced search capabilities of data at a "pixel level" (in contrast to "file level"), and (iv) "in-data" analytics and computation to avoid downloading the mammoth amount of data through the internet.

In this chapter, we review the current trend of the design, implementation, and functionalities of such geospatial-temporal platforms and associated services, focusing on those based upon scalable key-value datastores. IBM PAIRS (Physical Analytics Integrated Data and Repository Services) Geoscope will be used as an example through which we illustrate how the architecture and key-value datastore design supports the aforementioned features and high-performance data ingestion, query, and analytics. The specific implementation of a publicly available PAIRS instance will be presented along with its performance benchmarking.

Furthermore, we review the RESTful API interface of IBM PAIRS. The APIs are minimalistic and designed to provide the end users from different perspectives - data providers, industrial analysts, software developers, data scientists - a smooth experience to seamlessly exploit and use geospatial-temporal data. The API interaction with PAIRS will be illustrated through a few query examples and use cases in extended range weather forecasting and electric utilities. The use cases also highlight how contextual insights can be rapidly gained through a variety of "cross-layer" queries and analytics to reveal relationships/patterns and to predict trends.

Introduction

Traditional geographic information systems connect data with geolocations (e.g., weather, maps etc.). Since their first computerized instantiations in 1960 [1] such systems have been widely used to process and analyze (mostly) static, geo-coded "vector" data (=points, lines and polygons). Traditional GIS is a central technology to geospatial analytics which is a fast growing market projected to reach 86 billion USD by 2023 at a 16.3% CAGR (Compound Annual Growth Rate)[2]. However, GIS is at an inflection point for mainly two reasons:

On the technology side, the backends of traditional GIS are hitting serious scalability limits as a result of the emergence of "mega" big data in the form of geocoded imagery (e.g., from drones and satellites) [3, 4], time-series IoT (Internet of Things) [5, 6], LiDAR (Light Detection and Ranging) [7] or RaDAR (Radio Detection and Ranging) data. By way of example, the European Space Agency (ESA)

produces more than 10 TeraBytes of satellite data in a single day [8]. Ten TeraBytes cannot by handled by most GIS backends. The growth of GIS data generation is expected to grow exponentially, considering the emergence of new platforms for data collections such as drones [7] or nanosatellites [4] or new sensor types such as hyperspectral LiDAR [9].

On the application side, GIS users are now looking more and more to take full advantage of these ever-growing, ubiquitous new data sources leveraging the latest advances of machine-learning and artificial intelligence with the goal to operationalize GIS use cases [10, 11]. Examples of such "geospatial-temporal" use cases are plentiful and cut across different industries ranging from the energy and utility industry (when and where to trim vegetation to avoid costly outages), agribusiness (when and where to buy or sell agricultural commodities), insurance (when and where are the highest risk assets) to governments (when and where to optimally respond to a natural disaster).

Dimensionality Data type	Spatial	Temporal	Spatial and temporal
Vector	Traditional GIS (for planning)		
Raster			
Vector and Raster			Future GIS (for operations)

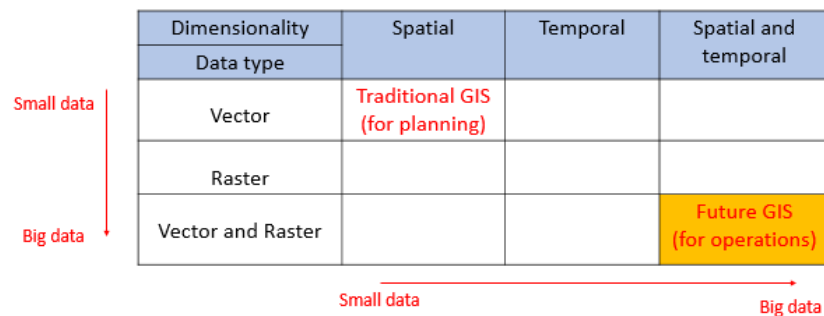Small data — Big data (vertical axis). Small data — Big data (horizontal axis)

Figure 1: Transformation of traditional GIS

Figure 1 illustrates the transformation of traditional GIS from a static, mostly vector-based, planning tool to an operational, real-time technology, which can process all kinds of different data at scale. An example for this is the application of smart meter monitoring (advanced metering infrastructure = AMI) in the context of renewable energy management [12].

The technical challenges for the transformation as depicted in Figure 1 are at least twofold. On a more practical level, scalable integration of data from different data sources is still a major bottleneck, where often more the 90% of all effort is spent on data pre-processing, curation and integration. Most use cases require a combination of different data whether this is raster (or imagery), vector (points, lines, and polygons), or time-series information. It is well known that such data can be highly complex, with hundreds of different formats, resolutions, projections and reference systems.

On a more fundamental level though, while such multi-modal data integration can be very difficult, it is arguably much more challenging to do this at scale. Many of the emerging geospatial-temporal data sets, which users seek exploiting are simply too big to be moved or downloaded in time to be useful for an operational

application. By way of example, the daily 10 TeraBytes from the European Space Agency (at 100 MB/s - read speed of a hard disk drive) takes more than a day to "move" from a storage device to the memory of a processor for subsequent computation.

The facts that (i) many of new emerging geospatial-temporal data sets (LiDAR, RaDAR, imagery, time-series) are too big to be moved and (ii) most use cases require the integration of multiple data sets, leads to the notion of data gravity. Data gravity means that big data tends to attract more data - in the same way a bigger mass attracts a smaller mass – and with that big data attracts more compute and applications. Most traditional GIS "move" data to the application or analytics and thus they are inherently limited in terms of how much can be processed and exploited. To be more specific, the database backend of GIS must become much more powerful to cope with these challenges, where in the future analytics and data must be directly collocated.

The solution to these data gravity challenges involves many technologies. First, given the size of the data and the fact that many users require the same big data sets (such as weather) for their different applications, a shared, often cloud-based system becomes more economical, which can be used remotely as a service. Other key technologies may include HDFS (Hadoop Distributed File System) [13, 14], which allows a scalable distributed storage layer exemplified by key-value stores such as HBase [15, 16] to be combined with a highly parallel processing layer using frameworks such as MapReduce [17, 18]. This in turn enables processing of very large data sets and more importantly pushes analytical tasks "into" queries (or executes these tasks during the queries) avoiding data movement.

By way of comparison, GIS systems even today rely often on relational database systems such as Postgres, mostly for vector data and/or file-based storage for raster data. It is well known that relational databases have difficulty in scaling to data sizes beyond a few tens of TeraBytes. The use of file-based storage comes with other major drawbacks. Often users need to assemble different images thereby dealing with different timestamps, resolutions, map projections etc.. Even in simple cases where a user wants to extract a time series from multiple satellite observations for the same location, one would have to download and open often thousands of files to extract the right information. Ironically, the inability to perform analytical tasks within the data and without downloading often leads to more data, where data providers compute ahead of time more derivatives of the raw data (such as vegetation index from hyperspectral satellite data).

To address the aforementioned challenges recently the IBM PAIRS Geoscope (Physical Analytics Integrated Data Repository & Services) was introduced [19, 20], which unlike most systems does not use relational database systems or file-based storage (object or cold store). PAIRS is based on a distributed, highly parallel, key-value big data system with a big, ready-made catalog of carefully indexed, diverse, and continually updated geospatial-temporal information (of both spatial and/or temporal vector and imagery data) in the cloud, enabling scalable access to

complex queries and machine learning-based analytics and AI to run without the need for downloading data.

PAIRS provides several benefits to the users. Firstly, PAIRS *allows access to PetaBytes of geospatial-temporal data sets at low cost*. That is because many users require the same data sets (e.g., weather, satellite etc.) and analytics capabilities and thus the shared services of PAIRS are much more efficient and cost effective. Second and as we will discuss in more detail below, *PAIRS drastically accelerates the analytics by reducing the time to insights* when retrieving and analyzing geospatial-temporal information - whether PAIRS (i) just provides AI-ready curated data, or (ii) returns results from search and analytics queries involving multiple data sets (by filtering, aggregating, applying mathematical functions etc.) or (iii) provides platform services for custom analytics without downloading the data or (iv) enables clients to integrate their own data thereby allowing them to exploit, analyze or monetize their data along with the PetaBytes of already curated data. Finally, and thirdly, due to the technology's unique scalability, *PAIRS enables users – often for the first time - to scale and operationalize their respective geospatial-temporal use cases.*

PAIRS is not the only technology for geospatial information which leverages a combination of key-value store with a distributed parallel big data system to overcome scalability limits. GeoMesa and GeoWave are two exciting and innovative, open-source research projects using a similar idea [21, 22]. By way of comparison, GeoMesa and GeoWave designs are primarily centered on vector data, while PAIRS complements this capability by focusing on raster data. In addition, PAIRS aims to provide end-to-end functionality from data curation to customizable "in data" analytics which a user can directly use without performing deployment or configuration optimization. Such "as-a-service" nature of PAIRS is reflected in its architecture and user experience as discussed next.

PAIRS architecture overview

Key components of the PAIRS architecture are shown in Fig. 2. In overview, PAIRS has four main components: (i) an *ingestion/data curation engine*, (ii) a massive *distributed compute and data store based on HDFS/HBase*, (iii) an *analytics and data platform*, which enables users via (iv) an *interface* to interact with the system.

(i) The *ingestion/data curation engine* includes data cleaning, filtering, re-projecting, resampling. It is a highly tuned C++ program compatible with a large variety of geospatial-temporal data in over 200 formats built on top of GDAL/OGR [23] and PDAL [24]. During the ingestion process, all imagery data is remapped onto a set of nested resolution layers and to a common projection and reference system. Details were described previously [19, 20]. The data curation engine can process with today's infrastructure (over 100 servers, ~4,000 cores, ~30 TB of RAM, and ~ 500GB/sec total network switching bandwidth) more than 50 TeraBytes per day. Routinely, PAIRS curates more than 15 TeraBytes per day and has subscribed to

many agencies such as NASA, ECMWF, ESA, NOAA etc. for continuous, near real-time data ingestion.
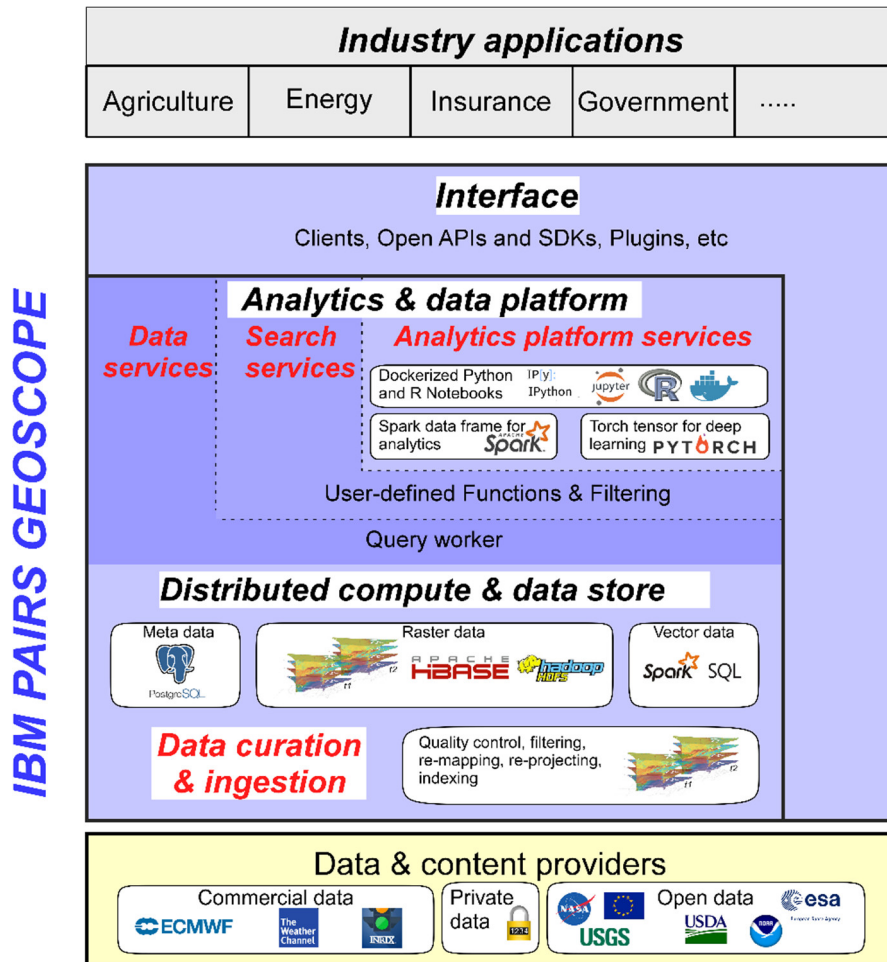

Figure 2: Overview of PAIRS architecture.

All data uploaded are indexed to a *massive distributed compute and data store based on HDFS/HBase*. In this key-value store, every raster data layer is modeled as (x, y, t, θ - value), i.e. value as a function of x (longitude), y (latitude), t (time), and additional dimensions θ. Here θ represents any additional dimensions, other than x, y, t, which are required to uniquely define the value. For example, additional dimensions may include vertical elevation (for 3D atmospheric data) or forecast lead time, Δt (between the issue of a weather forecast and the actual forecasted time)

or a sensor ID. A distinguishing feature of PAIRS is that all data are ready for use without a data staging or preparation step. Unlike many other technologies, PAIRS uses object and cold store *only* for archiving data which have already been ingested into the key-value store. All PAIRS data are organized in layers, where each layer is linked in space and time. Layers can be access controlled (visualization only, read, write, admin) according to the privileges of user groups. In addition to PetaBytes of raster data stored, vector data (discrete points, polygons, typically much smaller in volume) are store in PAIRS in Postgres or a key-value store which can be queried using SPARK SQL.

A MapReduce (M/R) query and built-in analytics engine is at the core of the *analytics and data platform*. It enables data retrieval, filtering, logical joins and complex math of a layer or between different layers. A basic PAIRS query is based on four elements: (i) what (specifies the data layers and additional dimensions if needed), (ii) where (geographical region), (iii) when (time period, maybe different for different data layers), and (iv) post-processing or built-in analytics (aggregation, mathematical computation, filtering). The query syntax is unified for different data layers. The key differences with respect to conventional platforms are the following: (i) A PAIRS query returns physical and logically organized data which is ready for analytics. This contrasts with a conventional platform's "search for data" function, which simply returns the reference to a set of files containing relevant data. The PAIRS data store design ensures that most of the data required by the same processor are co-located on the same cluster, which minimizes the burden of data reorganization. (ii) A PAIRS query takes care of commonly encountered post-processing tasks, such as aggregation and filtering, which often effectively reduce the data returned to the users (compared to the raw data) by over one order of magnitude. More detailed examples will be given below.

The query results are available as files for download, for visualization or processing via OGC web map service (WMS) [25] and web processing service (WPS) [26] which are served from a set of geo-servers [27], or as Pandas data frames or SPARK data frames ready to be used for analytics without downloading. For this, a Docker encapsulated ("dockerized") Python Jupyter Notebook with access to the data frame can be readily "spun" up.

All interactions with PAIRS are available to end users via an *interface* of an open RESTful API and a PAIRS client application, which includes a query GUI (graphical user interface) and Python Jupyter Notebooks. Two screenshots of the PAIRS GUI are shown in Figure 3. A freemium version of this PAIRS client is available at this reference [28]. Further updates of the PAIRS Client will be made including user-enabled uploading and 3D visualization. An initial version of the PAIRS API is available at this reference [29]. For the convenience of Python users, an open sourced PAIRS SDK wrapping API functionalities is available at [30] or from pip or conda Python package management system.
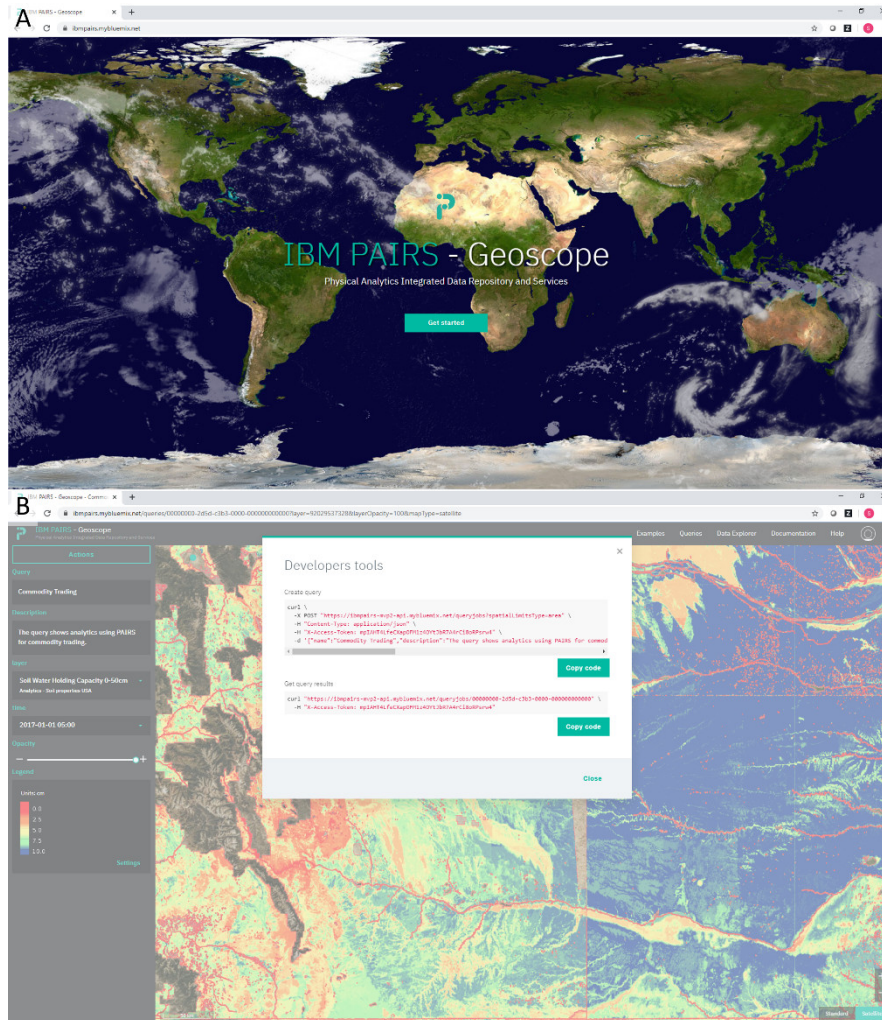
Figure 3: (A) PAIRS Client landing page and (B) screenshot of the PAIRS query UI.

Multiple PAIRS deployment models including SaaS, on premise, or hybrid can co-exist to accommodate clients' focus, e.g. business user, data distributor, application developer. In addition, PAIRS supports multiple data protection schemes to accommodate full data and/or analytics privacy protection, residency requirements, and flexible selective data sharing.

While we have described PAIRS from an architectural point of view it is equally important to understand PAIRS from a users' perspective. From a users' perspective, PAIRS can be a (i) *data curation service*, where users upload their geospatial-

temporal data into PAIRS. The benefit of that is that the user's data becomes query-able with all the already curated PAIRS data. Many users exploit PAIRS as a smart (ii) *data service*, where for example, time-series information from satellite data is being requested at multiple points. On the next level PAIRS enables a (iii) *search or discovery service* for geospatial-temporal data. More specifically, users can query PAIRS to identify or find all locations in a certain geographic area, which meet a couple of requirements. For example, show me all areas in the United States, where the population is larger than 1,000 people per square mile and the temperature will be below freezing in the next five days (for heating energy consumption esti-mation). Finally, a user can fully leverage the different (iv) *analytics platform ser-vice*, which enables users to customize analytics without downloading the infor-mation first. The anticipation of such PAIRS usage patterns from a users' perspective dictates the design of PAIRS key-value store which we detail next.

Key-value store design and performance

At the core of PAIRS is its data store based on HBase on top of Hadoop [15, 16]. For brevity, in the following discussion we focus on the implementation of the raster data store. Interested readers may refer to GeoMesa [21] and GeoWave [22] for the implementation of big vector data store. In HBase all data abstractly can be thought of as being stored as multi-level ordered key-value pairs on a distributed system, which extends over many data nodes (region servers) controlled by a master. Each region server hosts several consecutive key-value pair sections. Such sections are referred to as regions. Using MapReduce (M/R) and SPARK respectively, queries and analytical tasks are executed, which may access multiple regions on different region servers of HBase in parallel, thereby providing excellent scalable perfor-mance [17, 18] . Unlike relational databases, carefully tuned key-value stores are scalable to hundreds of PetaBytes [31].

While details vary, fundamental to the implementation based on key-value store is how to effectively translate or index multi-dimensional geospatial-temporal data (at least 3 dimensions x, y, t) to a one-dimensional key so that optimal and balanced performance of writing/reading is achieved for different types of raster data. The salient design decisions of the PAIRS key-value store are summarized in Table 1 and actual implementation is provided in Table 2. The design decisions are motivated by the anticipated read/write patterns of the geospatial-temporal data en-coded in the key-value store, and importantly, how to efficiently handle raster data at both of the two extreme ends of the spectrum. As depicted in Figure 4, on one end, there are cases with data of very high spatial resolution but only a few timestamps, such as the high-resolution aerial imagery from the US Department of Agriculture (USDA) National Agriculture Imagery Program (NAIP) dataset. The spatial resolution of this dataset is 0.5 to 1 m but data is only available every other year. On the other end, there is data of lower spatial resolution with many timestamps. For example, weather forecasts often come with hourly resolution but are of few kilometers' spatial resolution.

Design Decision	Rationale
Use a fixed coordinate reference system and a fixed set of nested grids.	To enable analytics-ready data at the cost of reprojection errors.
Employ a multi-aspect row key to encode spatial and temporal information.	To enable efficient parallelized data processing for large queries.
Supercells (32x32 pixels) as the value of the key-value store.	To reduce the storage overhead of the keys and the computional overhead of the key-value operations.
Temporal hash in the HBase row keys.	To mitigate the resorting of the HBase key-value stores during ingestion of data layers of high temporal frequency.
Spatial hash in the HBase row keys.	To mitigate HBase region server "hot-spotting" when ingesting data layers of high spatial resolution.
Construct a set of coarse-grained overview layers for each data layers.	To enable (1) effective retrieval of available timestamps at given locations, (2) rapid preview of the results of large queries, and (3) acceleration of queries which require the filtering of data layer values.

Table 1. Summary of PAIRS design decisions.

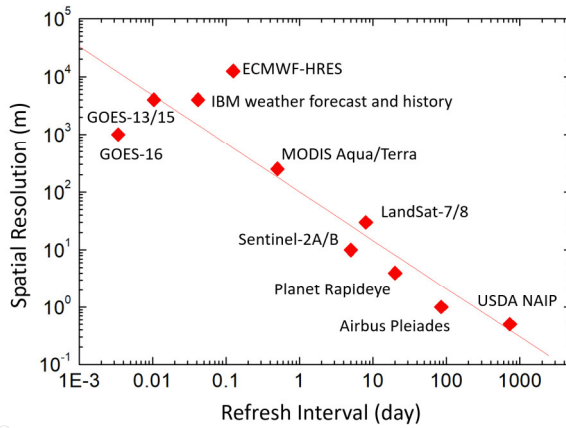|	Key			Value
Row Key [128 bits]	Column		Version	2^N x 2^N pixel super cell
	Column Family	Column Qualifier	Timestamp	
[4 bits reserved] + [4 bits Temporal Hash] + [4 bits Spatial Hash] + [52 bits Spatial] + [16 bits reserved] + [48 bits Temporal Key]	Data Layer	Additional Dimensions	Not used	Typical 32x32 pixels

Table 2: Design of the PAIRS key-value store.

Figure 4: Spatial and temporal resolutions of different raster data sets

Moreover, merely reading the key-value store to retrieve data is not enough. As discussed earlier, the retrieved data must be organized logically and physically in a way that readily enables downstream data analytics (i.e. analytics that proceeds without major reshuffling of the data, which can be a bottleneck in a public cloud environment).

In PAIRS, each dataset (such as a satellite imagery product) is represented as an HBase table (HTable), which conceptually is a hierarchical, three-level key-value store – the three keys being row key, column family and column qualifier. The key-value store is ordered by the three keys with row keys being at the top of the hierarchy. PAIRS employs a key design as shown in Table 2. The highest level HBase row key encodes space and time (i.e. x, y, t). The 2nd level column family encodes a data layer (such as a band of a satellite imagery product). The 3rd level column qualifier encodes any additional dimensions. For example, atmospheric data usually comes in at different attitude. The attitude information is stored as one of the additional dimensions. Weather forecasts may also involve a forecast lead time (e.g. forecast is for 1 day or 10 days ahead), the forecast lead time can be stored as another additional dimension. In the following, we note a few salient features of the raster key-value store design.

To encode the location information of incoming data by spatial keys, a predetermined map projection and spatial resolution are necessary. PAIRS is designed primarily for a cloud hosted data and analytics service for industrial applications. We also anticipate many of its users may come from non-geospatial background. Thus, PAIRS adopted the WGS 84 projection in favor of its simplicity. The inefficiency and non-convergence problem of WGS 84 near the poles are typically of less importance for most conceivable industrial use cases but this can be addressed by alternative projections. A fixed and nested resolution hierarchy (spatial resolutions) as shown in Table 3 is adopted. The grid size reduces by a factor of 2 when resolution increases one level. During ingestion, all data is re-sampled to the next higher resolution level, e.g., data from a satellite with 1.0 m resolution is re-sampled to level 26 (0.89 m at equator). While such implementation increases the data volume, it has the advantage that all information is linked and thus PAIRS can provide very fast "contextual" information (e.g. from different satellites with different resolutions) compared to other systems. Queries including multiple layers of geospatial-temporal data – e.g. "show me all areas in the Middle East where the accumulated precipitation in the last week was lower than 0.2 mm and the population density is larger than 1,000 people/km^2" – are orders of magnitude faster because no re-sampling of the data is required at query time.

Levels	Δθ, Δφ [degree]	Δy Δx (θ=0°)	Δx (θ=40°)
31	2.5e-7	2.78 cm	2.09 cm
30	5e-7	5.56 cm	4.19 cm
29	1e-6	11.1 cm	8.37 cm
...	...	...	...
25	1.6e-5	1.78 m	1.34 m
...	...	...	...
20	5.12e-4	57.0 m	42.9m
...	...	...	...
10	0.524288	58.3 km	43.9 km

Table 3. Shows the PAIRS grid size. The different rows show the grid size in latitude/longitude degree (Δθ,Δφ).

Temporal and spatial hash in row keys : Naively, one might use a key with location and time information and then a z-order to ensure that the data from the same location is stored close on the same part of the disk. However, this creates "hotspots" when reading/writing always hits the same server [32]. To overcome this, PAIRS introduces a special spatial and temporal hash in the beginning of the key to achieve an improved and balanced performance for data layers at the two extreme ends of Figure 4.

First consider a data layer with low spatial resolution but high temporal refresh rate such as the GOES-16 satellite data (5 minutes refreshing interval). Because the main temporal key comes after the spatial key, writing data with a new timestamp to a spatial location means inserting (in contrast to appending) new rows into the HTable. Doing such insertion frequently is computationally quite expensive because within the HTable one must re-sort and re-compact to keep the key-store ordered [33]. To mitigate this problem, we have introduced a four-bit temporal hash. We note that GeoMesa employs a very similar temporal hash. This hash ensures that only a small part of the HTable must be re-sorted and re-compacted as new timestamps get inserted. Tests have shown that this temporal hash improves the data ingestion/curation process by more than 10x.

A different problem is encountered when ingesting and querying a data layer of high spatial resolution but very low temporal refreshing such as the NAIP data (two year refresh rate). In such case, because the tailing temporal key has only a few different values, when querying a polygon area or ingesting a new image tile, one will be effectively reading/ writing a set of continuous keys of HBase which usually are hosted on the same region server. This causes the aforementioned and well-known issue of "region hot-spotting" [32]. To overcome this difficulty, a 4 bits spatial hash is introduced after the temporal hash. This hash ensures that reading or writing of a large spatial area is parallelized on multiple regions to avoid "hot-spotting".

Supercell as values: Moreover, PAIRS uses supercells, which are arrays of 32x32 pixels, as the value of the key-value store. In this way the storage taken by the key (16 Bytes) becomes negligible compared to the value (4 KiloBytes for pixels of 4-Byte float type). Reading/writing each key-value pair then processes 1024 pixels at once, significantly enhancing performance. Our benchmarking showed over 50x improvement compared to one pixel per key-value pair.

PAIRS aims to achieve high performance for both "big" queries of raster data for a "large" area and "small" queries of point locations. Indeed, the profiling of PAIRS queries (Figure 5) indicates a large fraction of queries are the "small" ones for point location, typically 1 to 100 KiloBytes in size. We found empirically that supercells made of 32x32 pixels is a good size as it enables balanced performance. With a 32x32 pixel supercell, the time required to seek a key is already much less than the time to reading/writing each value (involves 1024 pixels, 4 KiloBytes for a pixel of 4 Byte float type). Thus for "big" spatial queries, the effect of further reducing the time to seek keys by increasing of supercell size is only marginal. Other the other hand, for queries of point locations, even though we need to retrieve a supercell (1024 pixels) for a single pixel value, its performance is not significantly degraded either. This is because even with a 32x32 supercell, the time for data retrieval is still about 100 microsecond or less, insignificant compared to the overhead (establishing https connection, logging the query etc.).

The overview layer key-value stores: The discussed key design favors data retrieval from point spatial locations for a period, which matches the preferred user behavior (see Figure 5). In this case, the starting and ending row keys for the point locations can be readily determined. One may simply retrieve all the row keys in between (retrieving all the key-values between a starting and an ending key is called a "scan" operation). In contrast, to retrieve data for an area for a period (or a set of time periods) is more problematic. The reason is that for any given spatial key, one does not necessarily know what temporal keys exist. For example, in satellite imagery, different parts of the earth are imaged at different times. Without prior timestamp information, one will have to either read out all the timestamps possible or first scan each spatial location to know what keys are available. In key-value store operations, such a "scan" operation often takes on more than 1 millisecond, which is much too slow compared to merely retrieving the value for a known key (called a "get" operation, typically or the order of 1 microsecond).
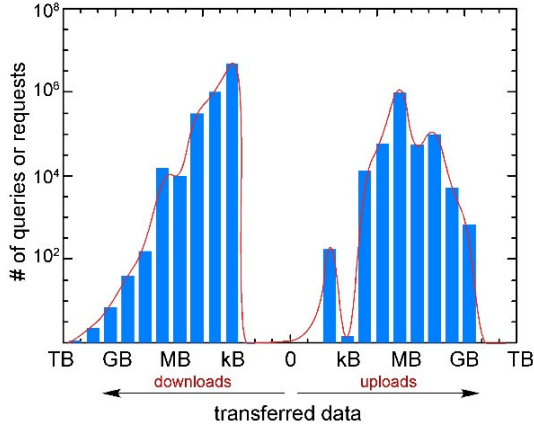
14

Figure 5: PAIRS user behavior sampled from more than 7M requests between 01/01 and 12/31 2017

To overcome such difficulty, the PAIRS innovation includes the introduction of multi-level overview layers as illustrated in Fig. 6. The overview layer uses a similar key-value store structure (Table 2) as the main layer, except that its spatial resolution is coarse-grained by a factor of two per level-up. At selected overview levels, statistics of the supercells are stored. For example, each pixel of the 5th level overview layer stores the mean, min, and max values of the corresponding 32x32 pixel supercell in the main layer (1,024x reduction in the number of pixels). Similarly, the 10th level overview layer stores the mean, min, max values of the corresponding 32x32 supercell in the 5th level overview layer.

To retrieve data for an area for a time period (or a set of time periods), one first scans the overview layer to obtain the timestamps available for the area. This enables one to pre-calculate all the row keys needed to retrieve in the main layer, leading to much faster data retrieval using the "get" operation (instead of the "scan" operation).

Moreover, the overview layer also brings the added benefit that it enables rapid overview visualization of the data layers and accelerated data filtering (e.g. "get data where temperature is below freezing"), because if the filtering condition can be ascertained by the overview layer, the retrieval of unnecessary data layers can be completely skipped.

PAIRS queries of point locations are usually served in real-time (with hundreds of milliseconds response times). In contrast, for queries of areas, Fig. 7 is a useful way to characterize their performance. In Fig. 7 the time for retrieving data is plotted as a function of data size. More complex queries would apparently change the curve. Since writing output to disk is the most expensive part of a query, query time can reduce substantially if a query involves data reduction (see examples in the next section).

From Fig. 7 one observes that if little data is retrieved, the performance is limited by *latency*, which is determined by the overhead of logging, authorization of a query, and queueing for the availability of resources before starting a MapReduce job. As the processing time increases with query size, latency becomes negligible.

Within an *optimal query size range*, the time is only weakly dependent on the size because the number of parallel mappers (in MapReduce) processing the query increases with its size. The slope in Figure 7 in this regime characterizes the *scalability*. For a typical user, the PAIRS system has an optimal query size range from 0.5 to 500 GB. The scalability slope is ~ 0.2 because the number of mappers scales sublinearly with the query size for a typical user. Beyond the optimal query size, the number of mappers reaches an upper bound, and the query time scales about linearly with size. The slope of the linear relation defines the max query *speed*, which is currently ~400 MB/s for a typical user. The maximum *throughput* is the largest possible query size limited by the memory available to hold the result for in-memory analytics, which is ~ 2 TB for a typical user.
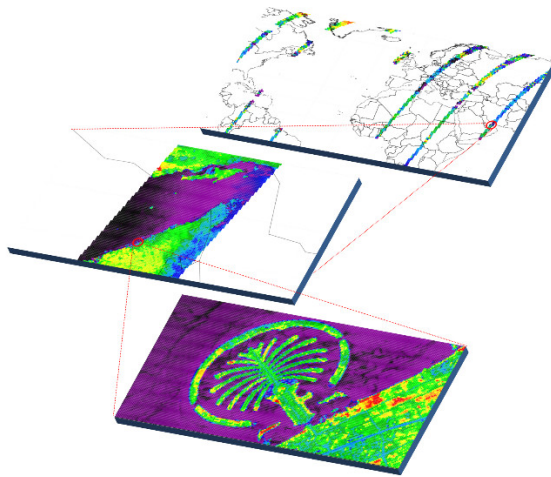

Figure 6: Landsat-8 (NIR band) illustrating the relation of main layer (bottom), 5th level overview layer (middle) and 10th level overview layer (top).
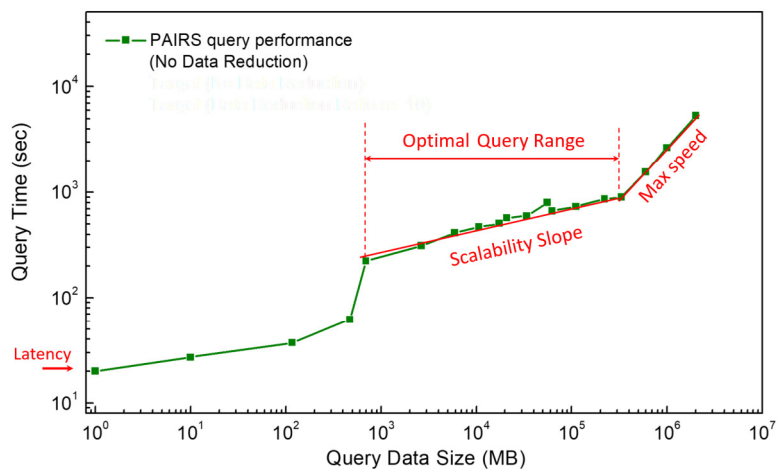

Figure. 7: PAIRS query times as a function of query data size.

PAIRS user experience

The implantation and performance of the PAIRS key-value store discussed above enables a paradigm-shifting user experience for many geospatial-temporal use cases when compared to a conventional system. Indeed, the conventional usage of geospatial-temporal data at scale is convoluted. For example, a user intends to run analytics on satellite imagery using a conventional system such as the US Geological Survey (USGS) EarthExplorer [34]. One first selects an area-of-interest (AOI) and a time range to obtain a list of image files of the relevant satellite tiles. Then these image files are downloaded and processed for the user's task at hand. Suppose one is to obtain a time series of near-infrared surface reflectance from ESA's Sentinel-2 for satellite for a particular region over many years. In such a scenario, often hundreds of tiles will need to be downloaded and opened/sought to extract the pixel(s) corresponding to the region of interest. Such a task gets increasingly complicated as more data sources become involved – e.g. we need temperature in addition to surface reflectance and/or we want to use data from other satellites such as LandSat or MODIS.

The PAIRS design emphasizes that the platform relieves the user from performing such data reorganization and provides a simple and unified experience regardless of the details of the original data. The PAIRS API and GUI capabilities are detailed in its documentation [29]. In this section we illustrate how the PAIRS provides a new user experience. For generality, we discuss the PAIRS query examples using the native PAIRS REST API. Often the open sourced PAIRS Python SDK (short PAW = PAIRS Geoscope RESTful API Wrapper) [30] which wraps the native PAIRS API offers more convenient interaction with PAIRS.

Data service

The simplest service which PAIRS offers is the *data service*. For example, to obtain the time series of surface reflectance as well as temperature, one simply sends a POST request to PAIRS with a JSON payload (query_json). The sample Python code snippet is:

```python
import pandas as pd, requests
pairs_auth = ('<username>', '<password>') # your username and password here.
query_json = {
    "layers" : [
        {
            "id" : "49361" # near IR surface reflectance, Sentinel-2 L2 band 8
        },
        {
            "id" : "49257" # 2m temperature, TWC gCOD hourly weather
        }
    ],
    "spatial" : {"type" : "point",  "coordinates" : ["41.213", "-73.798"]},
    "temporal" : {"intervals" : [
        {"start" : "2017-01-01T00:00:00Z", "end" : "2019-10-31T00:00:00Z"}
    ]}
}
```

```
# send a POST request containing query_json to PAIRS API endpoint.
api_response = requests.post(
    'https://pairs.res.ibm.com/v2/query', auth = pairs_auth, json = query_json
)
# convert the response json into a pandas dataframe
pairs_data = pd.DataFrame(api_response.json()['data'])
```

The query above requests about 3 years of near-infrared surface reflectance (Sentinel-2 band 8, PAIRS data layer id=49361) and 2 m temperature (global weather history hourly temperature from TWC (The Weather Company, an IBM business), PAIRS data layer id=49257) for a location somewhere in New York with the coordinates 41.213/-73.798 degree (latitude/longitude). PAIRS responds with a JSON (api_response) with about 100 records of surface reflectance and about 24,000 records of hourly temperature typically within a few hundreds of milliseconds. The last line of the code snippet above converts the JSON into a Pandas data frame for downstream analytics. Beyond point location query, a user may specify a query JSON with the spatial part representing a bounding box or a multi-polygon. In such case PAIRS returns either a set of geotiffs or CSVs (latitude, longitude, timestamp, value) for the queried area.

As the GUI counterpart to such query capability, a user can in real-time visualize the content of a data layer in PAIRS GUI as show in Figure 8. This includes (1) picking a timestamp to visualize a data layer as a color map and (2) picking a location and time range to visualize a time-series.
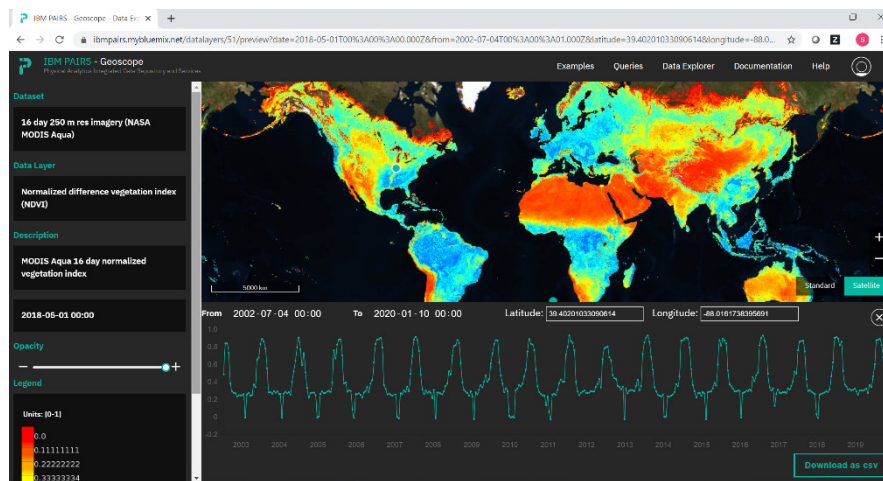

Figure 8. PAIRS GUI (Geoscope) enables the real-time visualization of the content of data layers in PAIRS. This example shows global NDVI from MODIS Aqua on 2015-11-25 and a twenty-year time series at latitude 37.3°/longitude -89.1°.

The performance of the query is the benefit of the key-value store design of PAIRS. Take the surface reflectance and temperature timeseries query above for

example, given the key design shown in Table 2, the location and time range specified in the query can be directly translated to a set of starting and ending keys for scanning the HBase to retrieve data. In contrast, in a conventional file-based system, one would be forced to open and seek tens of thousands of files (each for different timestamps) to retrieve data, which limits performance. It may also be obvious to the readers that the visualization of the data layer content on the GUI relies on the construction of the overview layers (Figure 6). Thus, at a given zoom level, the PAIRS overview layers can supply data to the GUI at the appropriately coarse-grained resolution.

Search or Discovery service

More sophisticated than the data service, PAIRS enables a user to push spatial and/or temporal aggregation, filtering, math computation, and basic geospatial transformation into a query, which we refer to as the *search or discovery service*. Let us use the following query as an example. The example computes the summer 2018 max temperature for all the corn fields averaged for all states in the contiguous US (CONUS).

```
query_json = {
    "layers" : [
        {
            "id" : "92", # PRISM daily maximum temperature
            "temporal" : {"intervals" : [
                {"start" : "2018-06-01T00:00:00Z", "end" : "2018-08-31T00:00:00Z"}
            ]},
            "aggregation" : "Max"
        },
        {
            "id" : "111", # USDA cropscape
            "temporal" : {"intervals" : [{
                "start" : "2018-01-01T00:00:00Z", "end" : "2018-12-31T00:00:00Z"}
            ]},
            "aggregation" : "Mean", # collapse crop type in the temporal range into a single value
            # filter out spatial area for which crop type equals 1 (corn per USDA designation)
            "filter": {"expression" : "EQ 1"}
        }
    ],
    "spatial" : {
        "type" : "poly",
        "aoi" : "24", # polygon of Contiguous US
        # list of polygon id for 48 states
        "aggregation" : {"aoi": [121, 123, 124, … , 130, 131, 133, 134, … , 171]}
    },
    # "temporal" below is irrelevant as it is overridden by "temporal" within the data layers above
    "temporal" : {"intervals" : [
        {"start" : "2018-06-01T00:00:00Z", "end" : "2018-09-30T00:00:00Z "}
    ]}
}
```

There is quite a lot going on in the example. To begin we are requesting data for spatial area CONUS ("aoi" : "24") and for two data layers: daily maximum temperature from PRISM dataset ("id" : "92"), and crop type from USDA cropscape dataset ("id" : "111"). The spatial resolution of PRISM is ∼ 4 km (PAIRS level 14), while crop type is ∼ 30 m (PAIRS level 21). PAIRS automatically samples the temperature data to match the crop type which is the highest resolution data layer of this query. Moreover, for each of the two data layers we use a different temporal range. We

requested temporal range 06/01/2018 to 09/30/2018 for temperature and requested PAIRS to apply max aggregation for the time period ("aggregation" : "Max") to obtain the highest temperature during the time period for each pixel. Separately for crop type, the temporal is 01/01/2018 to 12/31/2018. Temporal aggregation ("aggregation" : "Mean") is applied to collapse crop type within the temporal range into a single value. Finally a filter ("filter": {"expression" : "EQ 1"}) is applied, which request PAIRS to retrieve data for only spatial areas with crop type =1 (i.e. corn per US department of agriculture convention). Finally, a spatial aggregation is specified with a list of polygon id's for the 48 contiguous US states ("spatial": {"type": "poly", "aoi": "24","aggregation": {"aoi": [121, 123, 124, 125, … … 171]}}). PAIRS thus spatially aggregates the corn field temperature by states and provides an output file in the default CSV format.

A salient character of such a query is that all the computations are performed in parallel in the PAIRS cluster. While there is a large amount of data being processed, a user merely retrieves a CSV with around 3,000 rows in which each row contains the 2018 spatially averaged summer max temperature of corn fields for one state.

We can now take such query capability to answer some less trivial questions. Say we are interested in the impact of global warming on agriculture, thus would like to know in which part of the croplands in the northern hemisphere have seen a substantial summer daily maximum temperature (Tmax) rise of over 1.5 °C in the last forty years. The sample query is below.

```
query_json = {
    "layers" : [
        {
            # data layer id 49188 is daily maximum temperature (Tmax) at 2 m above surface
            # virtual layer "Y2018" is the mean summer Tmax of Jun to Aug 2018. Same below.
            "alias" : "Y2018","temporal" : {"intervals" : [{"start" : "2018-06-01",\
            "end" : "2018-08-31"}]}, "id" : 49188, "aggregation" : "Mean",
            "output" : "false" # PAIRS does not write output for this layer
        },
        {
            # virtual layer "Y2017" is the mean summer Tmax of 2017.
            "alias" : "Y2017","temporal" : {"intervals" : [{"start" : "2017-06-01",\
            "end" : "2017-08-31"}]}, "id" : 49188, "aggregation" : "Mean", "output" : "false"
        },

        … … # not showing "Y1980" to "Y2016" due to space limitation

        {
            # virtual layer "Y1979" is the mean summer Tmax of 1979.
            "alias" : "Y1979","temporal" : {"intervals" : [{"start" : "1979-06-01",\
            "end" : "1979-08-31"}]}, "id" : 49188, "aggregation" : "Mean", "output" : "false"
        },
        {
            # the mean summer Tmax difference between 2009 to 2018 and 1979 to 1998
            "alias" : "TempDiff",
            "expression" : " ($Y2018 + $Y2017 + $Y2016 + $Y2015 + $Y2014 \
                        + $Y2013 + $Y2012 + $Y2011 + $Y2010 + $Y2009)/10 \
                        - ($Y1988 + $Y1987 + $Y1986 + $Y1985 + $Y1984 \
                        + $Y1983 + $Y1982 + $Y1981 + $Y1980 + $Y1979)/10",
            "filter" : {"expression" : "GT 1.5"} # filter out pixels of value greater than 1.5
        },
        {
            "alias" : "crop_fraction",
            "temporal" : {"intervals": [{"snapshot" : "2017-01-01"}]},
            "id" : 49307, # crop fraction at 250 m resolution, survey of timestamp 2017-01-01
            "aggregation" : "Mean",
            "filter" : {"expression" : "GT 0.5"} #  filter where crop dominance > 50%
```

```
            }
        ],
        "spatial" : {
            # bbox north hemisphere, latitude 0 to 80 deg north and longitude -179.9 to 179.9 deg east
            "type" : "square", 'coordinates': [0,-179.9, 80, 179.9]
        },
        # "temporal" below is overridden by "temporal" within the data layers above
        "temporal" : {"intervals" : [
            {"start" : "1976-01-01", "end" : "2018-12-31"}
        ]}
}
```

In this example, we are requesting data for north hemisphere between latitude -0 to 80 degree north and longitude -179.9 to 179.9 degree east (defined by 'coordinates': [0, -179.9, 80, 179.9]). A number of user-defined "intermediate" layers are created by "mean" aggregation of Tmax (PAIRS data layer id=49188). For example layer "Y2018"

```
        {
            # data layer id 49188 is daily maximum temperature (Tmax) at 2 m above surface
            # virtual layer "Y2018" is the mean summer Tmax of Jun to Aug 2018. Same below.
            "alias" : "Y2018","temporal" : [{"start" : "2018-06-01",\
            "end" : "2018-08-31"}]}, "id" : 49188, "aggregation" : "Mean",
            "output" : "false" # PAIRS does not write output for this layer
        },
```

represents the mean Tmax in summer (June to August) 2018. Note that "output" : "false" instructs PAIRS to not write output, thus the intermediate layer stays only in memory.

Based on the intermediate layers, a user defined function (UDF)

```
        {
            # the mean summer Tmax difference between 2009 to 2018 and 1979 to 1998
            "alias" : "TempDiff",
            "expression" : " ($Y2018 + $Y2017 + $Y2016 + $Y2015 + $Y2014 \
                    + $Y2013 + $Y2012 + $Y2011 + $Y2010 + $Y2009)/10 \
                    - ($Y1988 + $Y1987 + $Y1986 + $Y1985 + $Y1984 \
                    + $Y1983 + $Y1982 + $Y1981 + $Y1980 + $Y1979)/10",
            "filter" : {"expression" : "GT 1.5"} # filter out pixels of value greater than 1.5
        },
```

computes the mean summer daily maximum temperature difference between 2009 to 2018 and 1979 to 1988. The temperature difference is subsequently filtered by "filter": {"expression" : "GT 1.5"}, i.e. selecting the pixels of temperature rise over 1.5 °C. Moreover, a filter using crop fraction (data layer id 49307) selects the pixels in which the crop coverage percentage is over 50%.

The result of the query is shown in Figure 9 below, which concludes that Europe croplands have the most notable summer daily maximum temperature rise in the last forty years.

The query above involves over 1,800 timestamps of temperature and 6e9 spatial grid points (northern hemisphere at around 250 m meter resolution, PAIRS level 18 for crop fraction). This combination represents ~ 1e13 spatial-temporal grid points. Such a task, starting from raw data gathering may conventionally take a data scientist from days to weeks to complete. In contrast, it took a single query and around 60 seconds to execute on PAIRS, showcasing the processing power and user experience on the PAIRS platform. While the UDF employed in the query is simplistic involving only arithmetic operations, PAIRS UDF supports common mathematical and logical operators and functions. More sophisticated analytics which can be

pushed into UDFs includes regression models and decision tree. Some examples can be found in this reference [35].

In addition to the functionalities discussed above, a PAIRS query may also include common geospatial processing including coarse graining, contouring, as well as customized functions such as delineating trees from satellite images (see section below on vegetation management) etc.. The list of built-in functionalities is continually evolving. For the latest refer to documentation [29].
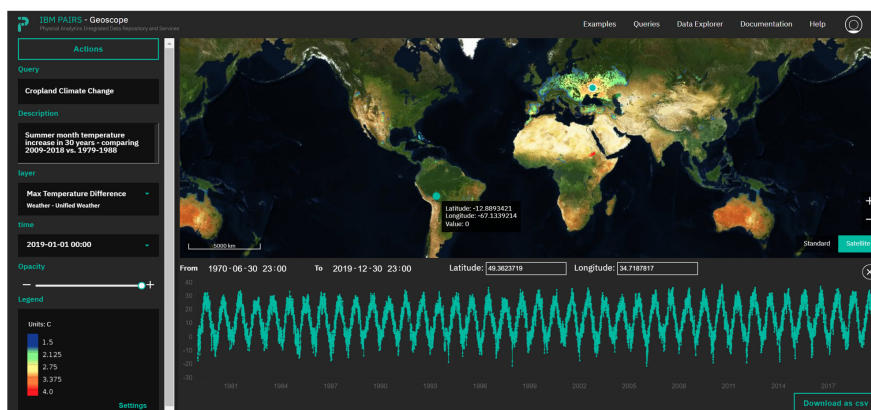

Figure 9. A PAIRS query showing northern hemisphere croplands in which mean summer daily maximum temperature has risen over 1.5 °C comparing 1979 to 1988 and 2008 to 2018. Note that Europe stands out as the most affected region.

Analytics platform service

Finally, it is anticipated that query (*discovery service*) by itself may not be able to perform all the analytics a user may want to. In such cases, it is expected that a PAIRS query would have substantially reduced the amount of data via aggregation, filtering etc. as discussed earlier. The last mile of customized analytics beyond query capability is handled by *analytics platform service*. A user may request a Python Jupyter Notebook on the PAIRS cluster or an IBM Watson Studio Notebook which contains the query result as data frame(s) to be launched. In API mode, a user makes the request using the id of a completed query and gets in the response a unique URL for the notebook. In the GUI, a user clicks on the "generate Jupyter Notebook" button in the "Actions" menu as shown in the screenshot (Figure 10). For resource management and access control purposes the notebook is dockerized [36]. In the Jupyter Notebook, the user can take advantage of all the latest modules of Python including PyTorch for deep learning and SPARK for scalable processing. For privileged users, a big query result will be returned as a SPARK data frame instead of a usual Pandas data frame. A SPARK data frame is distributed throughout

the memory of the PAIRS cluster when possible. Using PySpark, the data frame may be accessed from the Jupyter Notebook through a set of RESTful APIs orchestrated by an Apache Livy server.
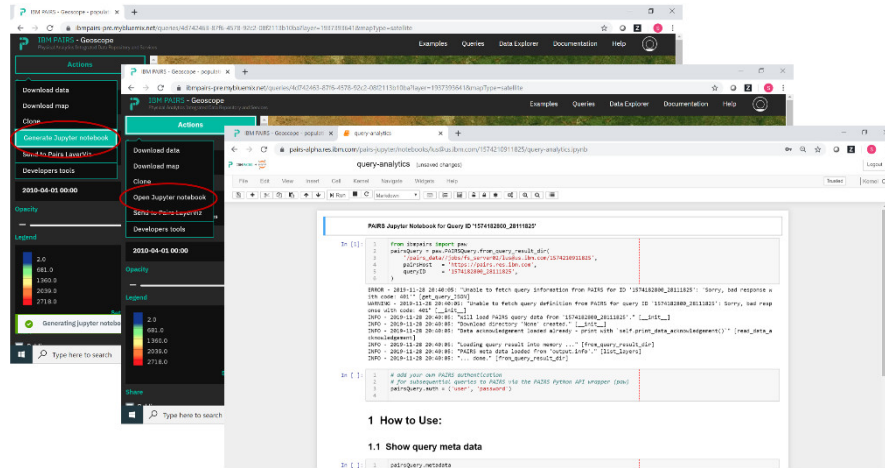

Figure 10. GUI interface by which a user may spin up a The Jupyter Notebook on the PAIRS cluster from a query result. The Jupyter Notebook is preloaded with the query result as dataframe(s).

.

Selected industry applications

Following the discussion of PAIRS architecture and user experience above, we present next a couple of selected geospatial-temporal use cases whose solutions were developed using PAIRS in the past few years. Geospatial-temporal use cases are, generally speaking "What-When-and-Where" type of applications and are naturally plentiful cutting across multiple industries and sectors (Table 4), such as government (how, when and where to respond to a disaster such as a hurricane?), retail (what, when and where to promote a product?), finance (when and where to buy and sell what commodities), agriculture (when and where to apply the right amount of fertigation), or energy (when, where and how much renewable energy will be generated?). While sometimes such PAIRS enabled applications are described as being "on top" of PAIRS, we note that this notion is misleading. It is better to refer to such applications as ones "within" PAIRS as they fully exploit the in-data computation capabilities as discussed in the previous sections.
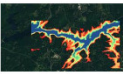
Industry	Finance	Utility	Agriculture	Insurance	Retail	Public
Example	Commodity trading	Vegetation management	Decision support for agriculture	Develop risk models (Flood, Fire)	Supply change	Disaster response
Example Queries / Questions	How much crop is planted? How much corn will be produced in Argentina?	Where do trees infringe on utility assets? When to schedule tree trimming?	What is the best crop to plant? When to apply fertilizers, pesticides?	Where are my assets at risk? At what time of the year is the risk the highest?	Where, when and what to ship? Where should I promote a product?	Who is being impacted by a hurricane? What is the best emergency response?
PAIRS Data Layers	Weather, land class, soil, satellite	Weather, satellite, power line data	Weather, land class, soil, satellites	Climate, vegetation, traffic, census	Weather, socio-economic data, store locations	Weather, climate, satellites, map, socio-economic data
Example	**Early Crop Recognition**	**Tree Identification**	**Precision Agriculture**	**Flooding Risk**	**Optimal store locations**	**Emergency Response**

Table 4: Exemplary PAIRS industrial use cases

PAIRS enabled improvements in weather forecasting:

One of the useful applications of PAIRS is weather forecasting, which is an old field of science (i.e., meteorology) but still very actively researched. For one, weather impacts literally every aspect of our lives and the economy, and for another, weather is highly complex with many aspects of the underlying physical phenomena not quite understood. It is thus well known that not every weather model provides the same forecasts and accuracy. By way of example, this becomes very evident during extreme weather events such as hurricanes, where multiple weather models can project very different pathways [37]. The difference in the forecasts is naturally most pronounced for longer term forecasts (going beyond 10 days) and thus we focus in the following on such long-term forecasts. Specifically, we discuss briefly how PAIRS can be used to improve the accuracy of such forecasts using its scalable big data processing capabilities by leveraging state-of-the-art machine-learning techniques. While the discussion will be focused on weather, it should be noted that the general framework, presented below, is applicable to "consolidate" between different forecasts or prediction modalities, which is a common challenge. For example, the presented framework could be used to consolidate the information received from different IoT sensor systems, which measure similar or related but not agreeing parameters.

	Weather station			Satellite	Model		
Source	RAWS	NOAA ISD	NOAA SurfRad	GPS-RO (Cosmic I)	NOAA CFS v2	ECMWF Extended	JMA Extended
Type	Vector	Vector	Vector	Vector	Raster	Raster	Raster
Coverage	US	Global	US	Global	Global	Global	Practically Global
Spatial Resolution	Point ~2,200 sites	Point ~30,000 sites	Point 7 sites	Point ~ 1000 ROs per day	0.5 deg	0.4 deg	1 deg

Temporal Resolution	< 1 h	1 h	1 min	~ 1 min per ROs	6 h	6 h	24 h
Forecasting Horizon	NA	NA	NA	NA	0 – 6 months	0 – 46 days	0 – 30 days
Forecasting Issuance	NA	NA	NA	NA	4 times per day	Twice per week	Once per week
Ensemble Forecast	NA	NA	NA	NA	4 Members	51 Members	50 Members
Estimated Date size	~4 MB/day	~ 30 MB/day	~2 MB/day	~9 MB/day	~1 TB/day	~ 400 GB/day	~1 GB/day

Table 5: Selected data sets in PAIRS including long-term weather models and weather station data.

Table 5 shows a selected list of data sets available in PAIRS, which are relevant for improving the accuracy of long-term forecasting. This includes weather station data from RAWS (=Remote Automatic Weather Stations) [38, 39], the ISD (=Integrated Surface Database) from NOAA (National Oceanic and Atmospheric Administration) [40, 41], and NOAA's Surface Radiation network (SurfRad) [42, 43]. PAIRS also includes data from GPS-RO (=Radio Occultation) [44, 45], which is a technique for measuring atmospheric parameters from space. There are also outputs from several extended or long range weather forecast models available, including NOAA's CFS v2 (=Climate Forecasting System)[46], as well as extended range forecasts from ECMWF (=European Centre for Medium-Range Weather Forecasts) [47] and from JMA (=Japanese Meteorology Agency) [48]. Note that Table 5 is only a rough estimate about the total amount of data content one may be able to retrieve from the data sources from a user's perspective. It is not about the internal complexity and data processing necessary to produce the user accessible outputs. For example, the ECMWF extended range forecasts are for 46 days ahead at 6 hourly resolution (185 timestamps) and for a 0.4 degree global grid (globally ~4e5 grid points). On the order of 100 parameters and/or pressure levels and 51 ensemble members are available from the forecasts. The forecasts are issued twice a week. Assuming parameter values are stored as four Bytes floating-point numbers, we estimate that, phenomenologically, the daily data content is around $185 \times 4e5 \times 100 \times 51 \times 2 / 7 \times 4 \sim 400$ GB/day.

Table 5 highlights the complexity of the data integration. For example, the forecast models not only differ in the underlying physics and assumptions which are used to generate them but also provide data at different spatial and temporal resolutions and cover different forecasting horizons. By virtue of PAIRS, the output from all these different models are "automatically" harmonized, integrated and spatially linked. Note that although many relevant datasets, such as ECMWF weather reanalysis etc. are omitted in Table 5, the amount of data listed already amounts to over 500 TeraBytes annually. Clearly to exploit all that data within a reasonable processing time, a scalable big data platform such as PAIRS is required.

As for any machine-learning task, the analytics includes two steps (training and deployment). In the training, first, an error analysis is performed to identify the most important features, where historical forecasts are compared with actual

measurements from high quality weather station. Because PAIRS allows quick and scalable access to data, this can be followed by a very granular functional analysis of variance (FANOVA) [49], which identifies 0^{th}, 1^{st} and 2^{nd} order errors of the forecasted parameters (such as temperature, precipitation rate etc.). By way of example, the 0^{th} order of a temperature forecast is just a bias, while the 1^{st} order error depends on one feature and the 2^{nd} order error on two features and so on. Such features can include other forecasted parameters. An example of a 2^{nd} order error of a 30-day temperature forecast from NOAA CFSv2 member 1 is shown in Figure 11, where it is compared to a class I weather station from the SurfRad network (here for Bondville, IL, 40.05192°N, 88.37309°W). In this example, the 2^{nd} order error is a function of wind speed and solar irradiance. As shown in Figure 11, FANOVA reveals that for this specific location the NOAA CFSv2 model on average *overpredicts* the temperature if the forecasted wind speed and solar irradiance are high. However, the same model *underpredicts* the temperature if the wind speed and solar irradiance are low. Clearly and as shown in Figure 11, different regimes of weather categories can be identified. We note that the 2^{nd} order error does not only depend on the forecast location, but it is also a function of forecast horizon (e.g., 30 days ahead vs 60 days ahead), and forecast parameter (e.g., temperature, precipitation). While we only show in Figure 11 four weather categories, many others are identified in FANOVA for different features, forecast horizons, locations etc.

Next we train an individual machine learning (ML) for each weather situation, forecast horizon, location and forecast parameter. Best results have been achieved using ensemble learning methods for regression (random forest) although other ML methods show good results as well. One may even run multiple ML models in parallel and adopt a multi-expert learning system which dynamically picks the most accurate ML method based on recent performance. The key, however, is to have on-demand access to data frames of training label (ground truth) and a large number of features so that the important parameters for different weather situations and different ML methods can be selected.

A typical training involves querying over 3 years of historical training data from around 100,000 point locations globally, which means around 1 TeraBytes have to be processed for each forecast variable - a nontrivial task. For this purpose, a specific data assembly module is used. The module uses an XML file as the template to construct complex PAIRS queries, manages those queries, and reorganizes the query results into training or forecasting data frames. As noted in Table 5, NOAA CFS forecasts are issued four times per day, while ECMWF and JMA extended range forecasts are issued twice and once a week, respectively. Thus, one of the roles of the data assembly module is to pick different lead times (Δt between the issue of a forecast and the actual forecasted time) for the different forecasting models in the PAIRS query so that the latest forecasts of the models are selected.
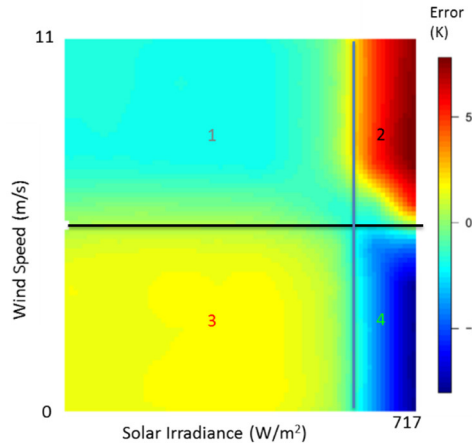
Figure 11. 2nd order error of a 30-day ahead temperature NOAA forecast for Bondville, IL, USA (here as a function of wind speed and solar irradiance)

In the deployment step, after data assembly and classifying the respective weather categorization, we apply the specific trained machine-learning model for this case, for each forecast parameter, location, and forecast horizon if applicable. Figure 12 shows an example of such a forecast, which nicely illustrates the power of the approach and how a scalable platform such as PAIRS can help to develop such fine-grained ML models. We show in Figure 12 the best long-term forecast from the 4 members of NOAA's CFSv2 model (in red) for a location with high quality weather station at Pennsylvania State University (40,72012°N, 77.93085°W). The temperature measurements are shown in blue. Because forecasts were available every 6 hours the comparison between forecast and measurement is performed for the same time interval. As becomes evident the NOAA forecast provides moderate accuracy. For comparison we have plotted forecasts, where we used machine-learning without categorization (green) and with categorization (blue). While the machine-learnt forecasts without categorization show improvements over NOAA's forecast it tends to be "biased towards the mean", which is a common pitfall of certain machine-learning approaches. Clearly, as shown in Figure 12, PAIRS big data capabilities, which enable specific machine-learning for each weather category, this problem can be mitigated and overall the best accuracy can be achieved. For reference in this plot we show four corresponding weather categories (labeled from 1 to 4).

While the data in Figure 12 shows just a snapshot for a single location and forecast parameter (i.e., temperature) we show in Figure 13 the mean absolute error (MAE) results for 7 locations in the US with class 1 weather stations (Bondville, IL (40.05192°N, 88.37309°W), Boulder, CO (40.12498°N, 105.23680°W), Desert Rock, NV (36.62373°N, 116.01947°W), Fort Peck, MT (48.30783°N, 105.10170°W), Goodwin Creek, MS (34.2547°N, 89.8729°W), Penn State, PA (40.72012°N, 77.93085°W), Sioux Fall, SD (43.73403°N, 96.62328°W)) for wind speed and temperature, respectively in comparison with the four NOAA CFS

member models (for the duration from 06/20/15 – 09/20/15). Figure 13 shows improvements in MAE of 30% over the best NOAA model.
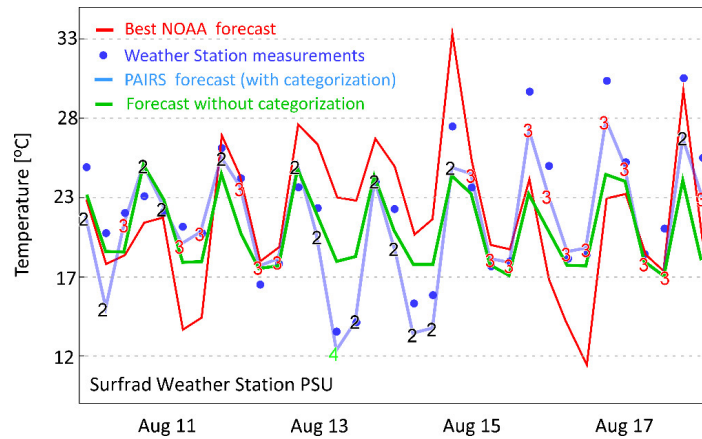

Figure 12. Comparison of a 30 day ahead temperature forecasts from the "best" NOAA forecasts and a PAIRS forecast (with and without categorization) with measurements from a SurfRad weather station at Pennsylvania State University (year 2015)
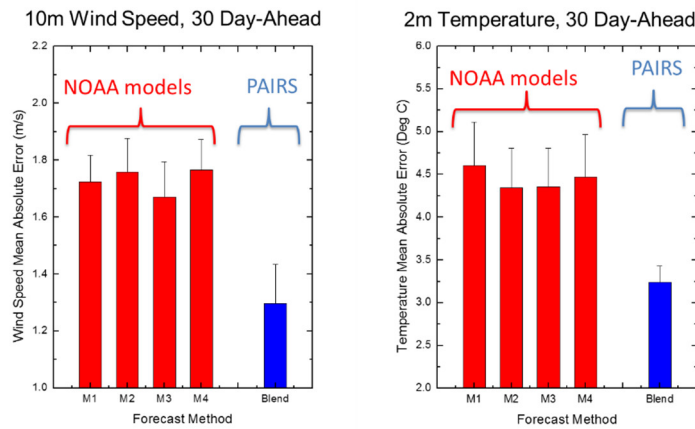

Figure 13. Mean absolute error (MAE) of the PAIRS model for wind speed and temperature for seven locations in the US with class 1 weather stations (Surfrad) in comparison with the 4 NOAA forecast models (for the duration from 06/20/15 – 09/20/15).

Vegetation management

Besides weather forecasting, PAIRS recently attracted many applications in the electric utility industry. In many parts of the world electrical powerlines are above ground. For example, in the US alone, the electrical grid includes 200,000 miles of transmission and 5.5 million miles of distribution lines and almost all of it is above ground, where it can come in contact with vegetation. However, this is a safety hazard and can cause major outages or even spark wildfires. Consequently, many electric utilities have complex vegetation management programs, which include regular pruning, brush removal, herbicide applications, etc., to prevent the vegetation from interfering with the utility assets and overgrowing of the conductors. Naturally such line clearance programs are not only difficult and complex but most importantly very costly with vegetation management being the largest preventive maintenance expense for many utilities [50]. By way of example, it has been reported that San Diego Gas & Electric must trim more than 450,000 trees regularly [51].

It is well known that tree growth can vary tremendously by species and other environmental conditions such as weather, soil etc. [52]. However, the lack of monitoring capabilities, which could provide actionable insights in where and when vegetation management is required, leaves many utilities with no other choice than regular maintenance schedules for their programs, which is naturally non-optimal and adds to the already very large cost.

In the following we discuss briefly how a vegetation management solution has been developed using PAIRS big data processing capabilities. Table 6 shows selective data sets which are available in PAIRS and are relevant to understand and monitor the progression of vegetation. As in the previous example, multiple very large data sources with different data types, resolutions etc. must be integrated.

	Other	Satellite		Weather		Soil	
Source	LiDAR	NAIP	Sentinel II	NAM	GFS	SURGO	Soilgrid
Type	Vector	Raster	Raster	Raster	Raster	Vector	Raster
Coverage	Local	US	Global	US	Global	US	Global
Spatial Resolution	<0.1 m	0.5 m	10 m	5 km	0.25 deg	Point	250 m
Temporal Resolution	NA	2 y	Weekly	1 h	3 h	NA	NA
Forecasting Horizon	NA	NA	NA	0-60 h	6-192 h	NA	NA
Estimated Data size	NA	~80TB/year	~12TB/day	~ 0.6TB/day	~ 1TB/day	~ 1 TB	~ 2 TB

Table 6: Different data sets in PAIRS relevant for vegetation management

Key to the vegetation management solution in PAIRS is the combination of high and low-resolution (spatial) hyperspectral aerial/satellite imagery. The high-resolution imagery, for example, the NAIP dataset in US at 0.5-1 m spatial resolution [53], enables the computation of a vegetation base layer. In some cases, the base layer can also be derived from LiDAR (light detection and ranging) data sets, which can

be easily processed in PAIRS [54]. Both LiDAR and/or NAIP data sets can be used to estimate and delineate the canopy size as illustrated in Figure 14. The computation of the vegetation from NAIP is based on the normalized difference vegetation index (NDVI), which is the normalized difference between the red and near infrared band of the imagery. In some cases, some additional data layers (i.e., land use, OpenStreet map etc.) can be leveraged to improve the tree identification process.

The difference between different types of vegetation can be inferred from a time series of remote observations of the vegetation index, which is often not available at such high spatial resolution. However, the Sentinel-2 satellite from the European space agency (ESA) provides NDVI data at 10 m spatial resolution every 5 days if clouds are not interfering [55]. An alternative data source, which is not listed in Table 6, but available in PAIRS is the LandSat-8 dataset from the USGS [56]. Figure 15 shows a time series of NDVI data for two different vegetation types. In combination with ground truth data, which can also be obtained from LiDAR scans, such information can then be used to identify the vegetation type and tree species as applicable.

Combining tree type identification with consecutive high-resolution imagery or LiDAR scans, one can further estimate the tree growth using the canopy size. The basic relationship between tree canopy size and tree growth is shown in Figure 16 [52] for different weather conditions.
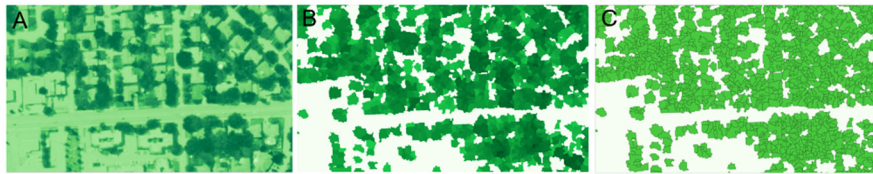

Figure 14 shows the basic process of tree delineation using the normalized differential vegetation index (NDVI): (A) raw NDVI, (B) result of smart thresholding, (C) vectorization to obtain the outline of tree canopies.
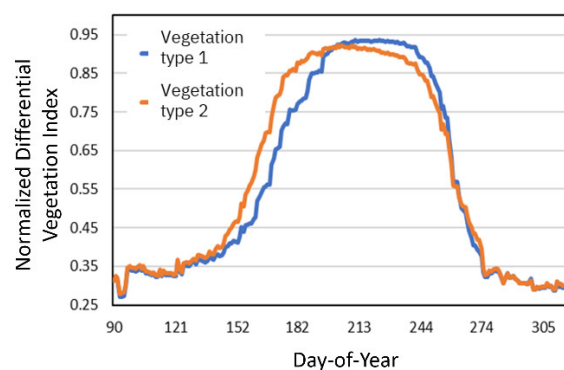

Figure 15: Time-series of normalized difference vegetation index (NDVI) for two different vegetation types.
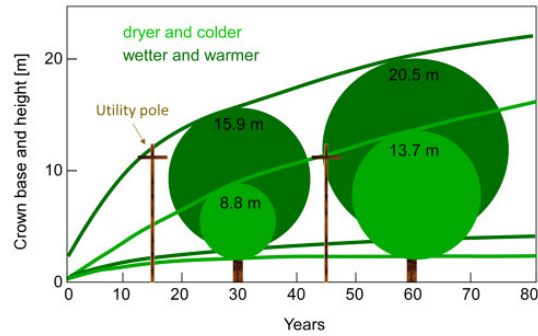
Figure 16 shows the relationship between tree height and canopy size for different environmental conditions for a green ash tree (rendered from [52]).

Finally, we show in Figure 17 the results from such an analysis with the delineated vegetation and tree height in the vicinity of a power line. While the model is simplistic, initial validations have shown that this model can provide between 80-90% accuracy [54].
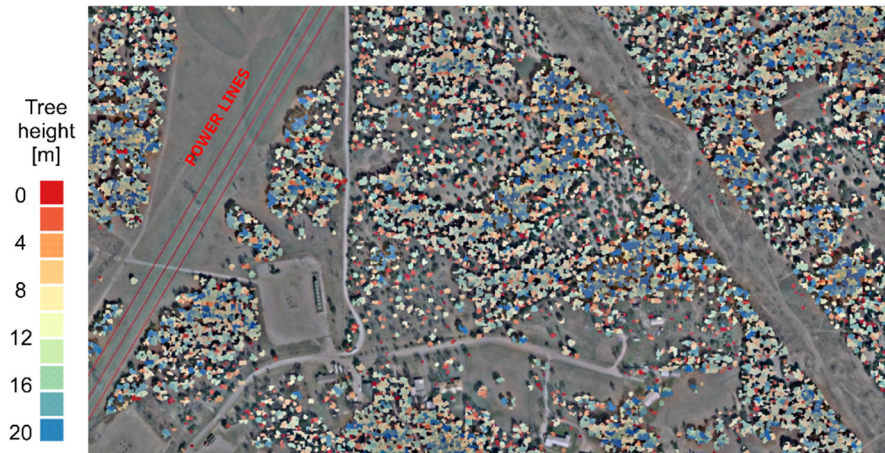

Figure 17: Tree distribution around power lines after processing high resolution aerial and LiDAR data.

Conclusion and PAIRS Resources

To conclude, IBM Physical Analytics Integrated Data Repository and Services (PAIRS) is a big data and analytics service platform designed to support complex

industrial applications which require analytics of a wide range of geospatial-temporal data. PAIRS features highly scalable (PetaBytes scale) storage of curated data and processing close to the data for advanced analytics and offers a unified user interface and user experience independent of the source of such data. It substantially reduces users' data management burden and, in many use cases, enables the users to drastically accelerate "data-to-sights" through its ability to compare, combine, filter, sort and display multiple large data sets simultaneously for correlation discovery and change detection.

Some useful PAIRS resources are provided below.

- Freemium GUI: https://ibmpairs.mybluemix.net
- API tutorial and reference: https://pairs.res.ibm.com/tutorial/
- Data documentation: https://ibmpairs.mybluemix.net/data-explorer
- Python SDK and sample Jupyter Notebooks (open source): https://github.com/ibm/ibmpairs
- PAIRS demos: https://www.youtube.com/watch?v=MlPhTKE189s
- PAIRS Introduction: https://www.youtube.com/watch?v=Nxwi6x0ObT0

References

1. Clarke, K.C., *Advances in geographic information systems.* Computers, environment and urban systems, 1986. **10**(3-4): p. 175-184.
2. *Geospatial Analytics Market worth 86.32 Billion USD by 2023* https://www.marketsandmarkets.com/PressReleases/geospatial-analytics.asp.
3. Tang, L. and G. Shao, *Drone remote sensing for forestry research and practices.* Journal of Forestry Research, 2015. **26**(4): p. 791-797.
4. Bouwmeester, J. and J. Guo, *Survey of worldwide pico-and nanosatellite missions, distributions and subsystem technology.* Acta Astronautica, 2010. **67**(7-8): p. 854-862.
5. Gubbi, J., et al., *Internet of Things (IoT): A vision, architectural elements, and future directions.* Future generation computer systems, 2013. **29**(7): p. 1645-1660.
6. Weber, R.M., *Internet of Things Becomes Next Big Thing.* Journal of Financial Service Professionals, 2016. **70**(6).
7. Dubayah, R.O. and J.B. Drake, *Lidar remote sensing for forestry.* Journal of Forestry, 2000. **98**(6): p. 44-46.

8. Petiteville, I., *personal communication.*

9. Hakala, T., et al., *Full waveform hyperspectral LiDAR for terrestrial laser scanning.* Optics express, 2012. **20**(7): p. 7119-7127.

10. Chen, J., et al., *Exploratory data analysis of activity diary data: a space–time GIS approach.* Journal of Transport Geography, 2011. **19**(3): p. 394-404.

11. Yuan, M., *Challenges and critical issues for temporal GIS research and technologies*, in *Handbook of Research on Geoinformatics.* 2009, IGI Global. p. 144-153.

12. Resch, B., et al., *GIS-based planning and modeling for renewable energy: Challenges and future research avenues.* ISPRS International Journal of Geo-Information, 2014. **3**(2): p. 662-692.

13. Lam, C., *Hadoop in action.* 2010: Manning Publications Co.

14. Zikopoulos, P. and C. Eaton, *Understanding big data: Analytics for enterprise class hadoop and streaming data.* 2011: McGraw-Hill Osborne Media.

15. Dimiduk, N., et al., *HBase in action.* 2013: Manning Shelter Island.

16. Harter, T., et al. *Analysis of {HDFS} Under HBase: A Facebook Messages Case Study.* in *Proceedings of the 12th {USENIX} Conference on File and Storage Technologies ({FAST} 14).* 2014.

17. Dean, J. and S. Ghemawat, *MapReduce: simplified data processing on large clusters.* Communications of the ACM, 2008. **51**(1): p. 107-113.

18. Ekanayake, J., S. Pallickara, and G. Fox. *Mapreduce for data intensive scientific analyses.* in *2008 IEEE Fourth International Conference on eScience.* 2008. IEEE.

19. Klein, L.J., et al. *PAIRS: A scalable geo-spatial data analytics platform.* in *2015 IEEE International Conference on Big Data (Big Data).* 2015. IEEE.

20. Lu, S., et al. *IBM PAIRS curated big data service for accelerated geospatial data analytics and discovery.* in *2016 IEEE International Conference on Big Data (Big Data).* 2016. IEEE.

21. Hughes, J.N., et al. *Geomesa: a distributed architecture for spatio-temporal fusion.* in *Geospatial Informatics, Fusion, and Motion Video Analytics V.* 2015. International Society for Optics and Photonics.

22. Whitby, M.A., R. Fecher, and C. Bennight. *Geowave: Utilizing distributed key-value stores for multidimensional data.* in *International Symposium on Spatial and Temporal Databases.* 2017. Springer.

23. Warmerdam, F., *The geospatial data abstraction library*, in *Open source approaches in spatial data handling.* 2008, Springer. p. 87-104.

24. Contributors, P., *PDAL Point Data Abstraction Library.* 2018.

25. Consortium, O.G., *Web Map Service.*

26. Consortium, O.G., *Web Processing Service.*

27. Henderson, C., *Mastering GeoServer.* 2014: Packt Publishing Ltd.

28. *IBM PAIRS freemium https://ibmpairs.mybluemix.net.*

29.   *IBM   PAIRS   tutorial   and   reference   manual   https://pairs.res.ibm.com/tutorial/.*

30.   *IBM PAIRS SDK https://github.com/ibm/ibmpairs.*

31.   DeCandia, G., et al. *Dynamo: amazon's highly available key-value store.* in *ACM SIGOPS operating systems review.* 2007. ACM.

32.   Vohra, D., *Defining the Row Keys*, in *Apache HBase Primer*. 2016, Springer. p. 117-119.

33.   Harter, T., et al. *Analysis of HDFS under HBase: a facebook messages case study*. in *FAST*. 2014.

34.   *USGS EarthExplorer, https://earthexplorer.usgs.gov/.*

35.   *IBM   PAIRS   application   examples. https://github.com/IBM/ibmpairs/tree/master/examples.*

36.   Merkel, D., *Docker: lightweight linux containers for consistent development and deployment.* Linux Journal, 2014. **2014**(239): p. 2.

37.   Brennan, M.J. and S.J. Majumdar, *An examination of model track forecast errors for Hurricane Ike (2008) in the Gulf of Mexico.* Weather and Forecasting, 2011. **26**(6): p. 848-867.

38.   Horel, J.D. and X. Dong, *An evaluation of the distribution of Remote Automated Weather Stations (RAWS).* Journal of Applied Meteorology and Climatology, 2010. **49**(7): p. 1563-1578.

39.   *National Interagency Fire Center (NIFC) Remote Automated Weather Stations (RAWS) https://raws.nifc.gov/.*

40.   Smith, A., N. Lott, and R. Vose, *The integrated surface database: Recent developments and partnerships.* Bulletin of the American Meteorological Society, 2011. **92**(6): p. 704-708.

41.   *NOAA ISD https://www.ncdc.noaa.gov/isd.*

42.   Augustine, J.A., et al., *An aerosol optical depth climatology for NOAA's national surface radiation budget network (SURFRAD).* Journal of Geophysical Research: Atmospheres, 2008. **113**(D11).

43.   *NOAA Surfrad https://www.esrl.noaa.gov/gmd/grad/surfrad/index.html.*

44.   Kuo, Y.H., et al., *Inversion and Error Estimation of GPS Radio Occultation Data.* 気象集誌. 第2輯, 2004. **82**(1B): p. 507-531.

45.   *UCAR COSMIC https://www.cosmic.ucar.edu/.*

46.   Saha, S., et al., *The NCEP climate forecast system.* Journal of Climate, 2006. **19**(15): p. 3483-3517.

47.   *ECMWF   ENS   Extended   forecast https://www.ecmwf.int/en/forecasts/datasets/set-vi.*

48.   *Japan Meteorological Agency, 2013: Outline of the Operational Numerical Weather Prediction at the Japan Meteorological Agency. Appendix to WMO Tech. Progress Rep. on the Global Data-Processing and Forecasting System and Numerical Weather Prediction, 188 pp. [Available   online   at   http://www.jma.go.jp/jma/jma-eng/jma-center/nwp/outline2013-nwp/index.htm.]*

49. Hooker, G., *Generalized functional anova diagnostics for high-dimensional functions of dependent variables.* Journal of Computational and Graphical Statistics, 2007. **16**(3): p. 709-732.

50. Hollenbaugh, R. and B. Champagne, *Utility vegetation management: the key driver of system reliability.* 2006.

51. Rodgers, D., *Trimming the Cost of Tree Trimming.* T&D World Magazine, 2014.

52. McPherson, E.G. and P.J. Peper, *Urban tree growth modeling.* Journal of Arboriculture & Urban Forestry. 38 (5): 175-183, 2012. **38**(5): p. 175-183.

53. *USDA NAIP (National Agriculture Imagery Program) https://www.fsa.usda.gov/programs-and-services/aerial-photography/imagery-programs/naip-imagery/.*

54. Klein, L.J., *N-dimensional geospatial data and analytics for critical infrastructure risk assessment.* IEEE Big Data Conference, 2019. **in press**.

55. *ESA Sentinel-2 https://sentinel.esa.int/web/sentinel/missions/sentinel-2.*

56. *USGS LandSat https://www.usgs.gov/land-resources/nli/landsat.*