DATA SERVICES FOR GEOSPATIAL
GROUND TRUTH

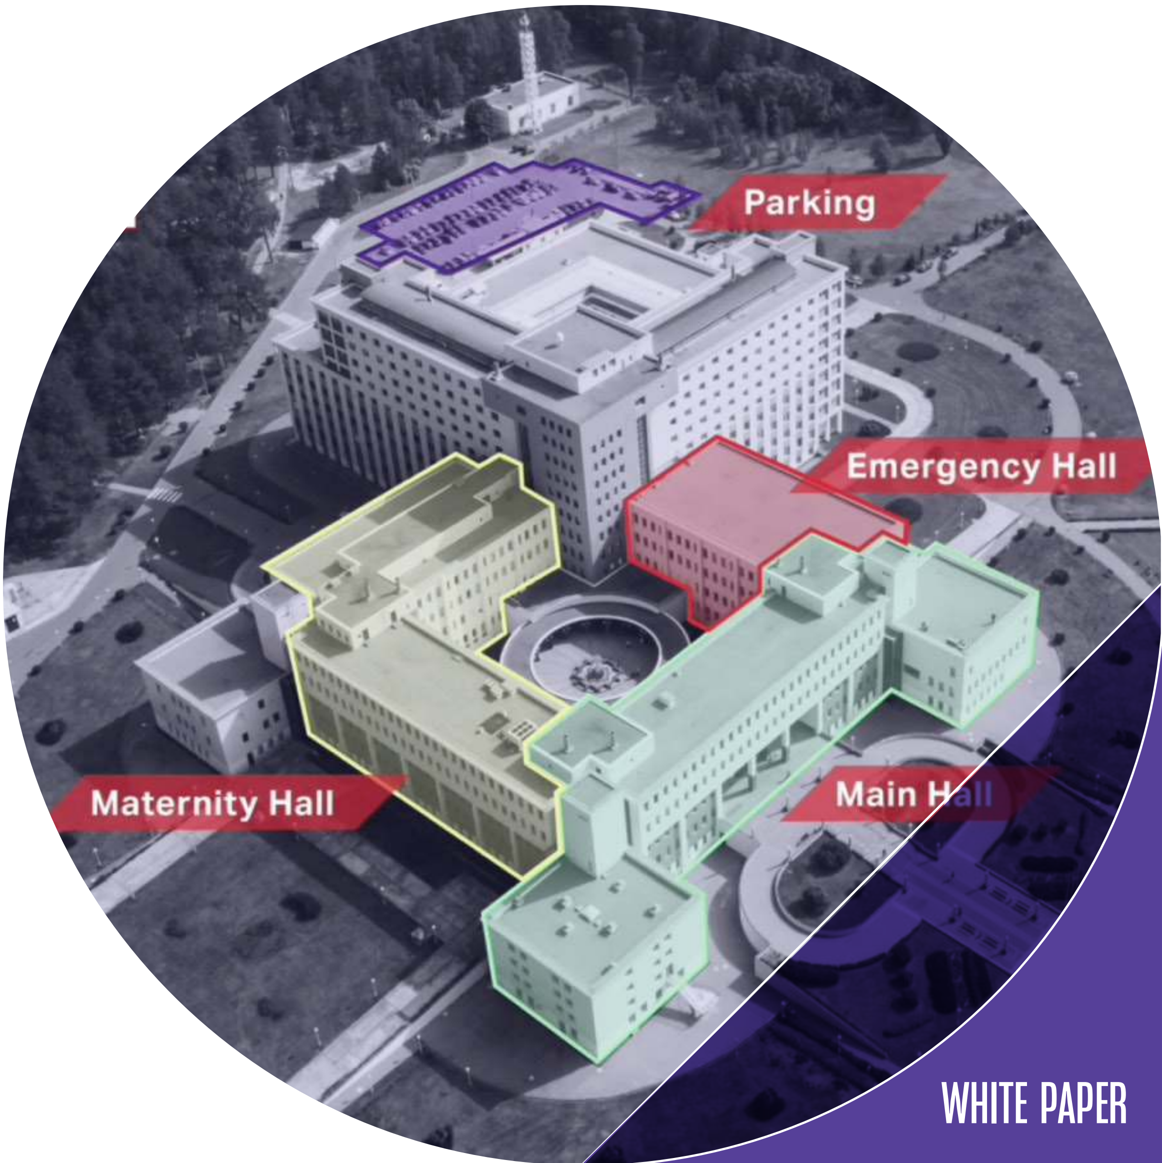

WHITE PAPER

IT WAS A BRISK MORNING IN PARIS

The first flight into geospatial intelligence began on a brisk morning in November 1783 when Jean Frances Pilatre de Rozier – the world's first Aeronaut – and a companion, the Marquis d'Arlandes, rose about 3,000 feet above Paris to gain an entirely new perspective on the City of Lights. With apologies to the 5th Century BC Babylonian scientists who developed civilization's first rudimentary geospatial intelligence tools to map the city and monitor agriculture and other planning functions for the empire, de Rozier and the Marquis were the first humans we know of to gain a truly bird's eye view of civilization in progress.

De Rozier's flight lasted barely long enough to enable a few sketches, and the two aerial explorers were captive to the prevailing winds that eventually pulled their balloon to a landing in a field on the outskirts of the city. But the fledgling field of earth observation changed irrevocably that day, even if the two men had little to no idea of what was to come.

Fortunately, this field of knowledge, process and profession all rolled into one has progressed remarkably since then. Geospatial intelligence has become an essential tool for everything from national security, to land use and planning, agriculture, and a host of other commercial and government functions.

As with many broadly adopted commercial and consumer technologies, geospatial data services developed initially through military and government intelligence programs created to provide ground-level surveillance from an eye in the sky perspective. The 1962 Cuban Missile Crisis was grounded in an attempt by the Soviet Union to install nuclear missiles (and launchers) just off the coast of the United States – but the confrontation itself was triggered once surveillance photos taken by U.S. military intelligence U-2 aircraft flying over Cuban territory had been thoroughly analyzed and annotated by U.S. military imagery analysts.

With the proliferation of commercial imaging satellites, starting in the early 2000s, a broad range of private sector industries and public sector regional and national government agencies have access to detailed imagery delivered through high-resolution imaging satellites, drones and even some traditional aircraft. What had been the exclusive province of military and government intelligence agencies is now available to anyone with the budget – and expertise – to convert raw imagery into useful data.


PILÂTRE DE ROZIER
PREMIER VOYAGE AÉRIEN EXÉCUTÉ DANS UNE MONTGOLFIÈRE, PAR PILÂIRE DE ROZIER ET D'ARLANDES, le 21 Nov. 1783

The Open Street Map project, formed in 2004 (openstreetmap.org), enabled what amounts to crowd sourced geospatial data, and, combined with the declassification in the U.S. of GIS data developed by military and intelligence agencies, opened a virtual floodgate of raw data. This has led to a revolution in government programs, such as land use planning – where even state and other regional agencies have access to satellite imagery – as well as a wide range of commercial applications, such as agriculture, commercial and residential insurance, and commercial real estate management and oil storage analysis.

Most tellingly, the industry is nowhere near its peak. Global spending on geospatial infrastructure, services, software and related initiatives is expected to almost triple, from an estimated U.S. $386 billion in 2019 to roughly $1.1 trillion in 2025. That includes spending on Geographic Information System (GIS) spatial analytics, Earth Observation, Scanning (through digital sensors), and Global Navigation Satellite Systems (GNSS).

GNSS (GPS, surveying, and indoor positioning) spending accounts for just about 60 percent of the entire geospatial market, followed by GIS-based initiatives (location intelligence, image analytics, and map content) at approximately 21 percent. Earth Observation via satellites, aerial and drones (just under 17 percent), and Scanning by LiDAR and Radar (roughly 3 percent) make up the rest of the market.

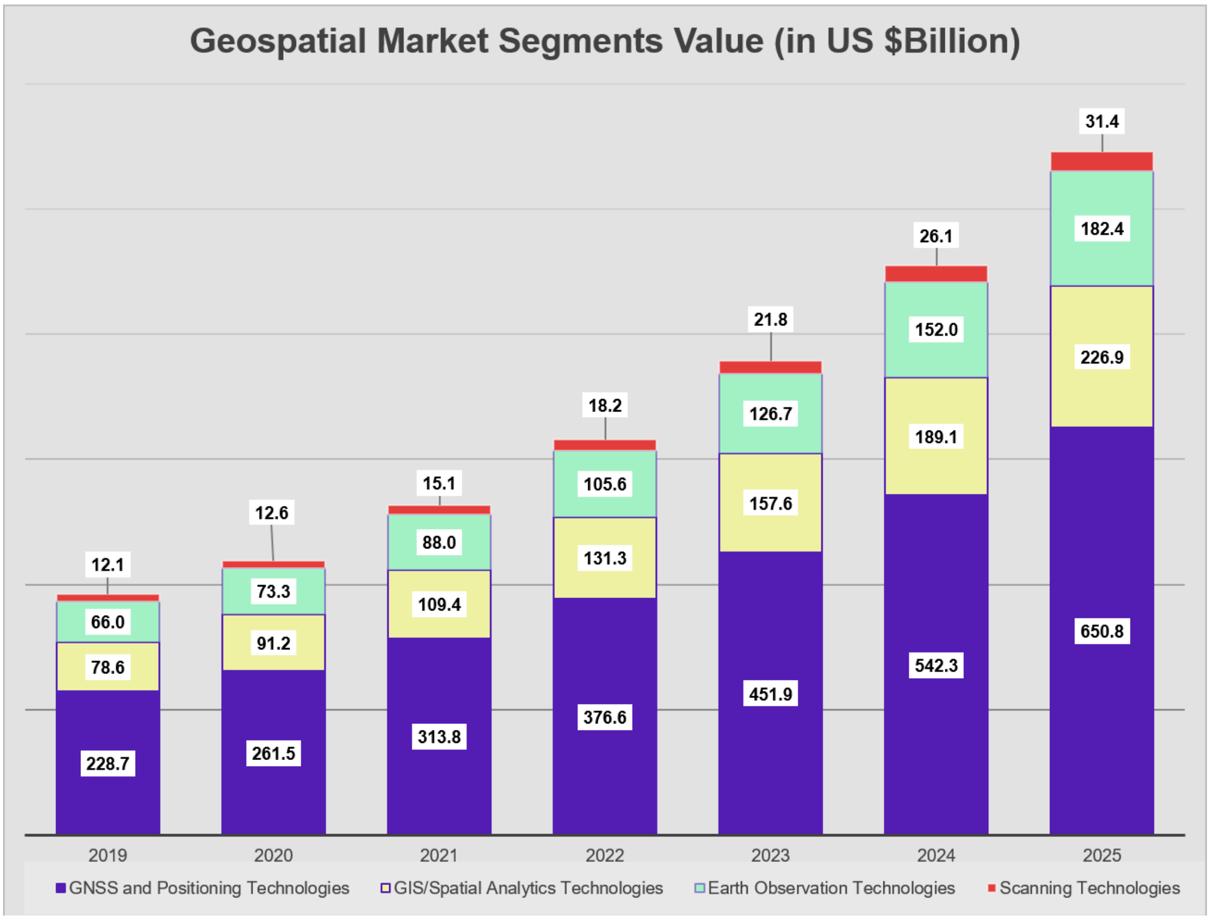

Source: Pitney Bowes

EYE IN THE SKY

Geospatial data evaluation is, at its most basic level, the conversion of unstructured data (images, sensor data and the like) into structured data that can be used for machine learning. Regardless of the specific industry (or public sector) vertical, the process remains constant: taking a raw image full of unstructured data and applying data annotation to break it down into individual components that can be characterized and organized to help train an algorithm. What separates geospatial data from other types of imagery and sensor data is the common visual perspective. All of the data is derived from satellites, drones and other aerial sources that capture everything within a specific geographic area, so the bounding boxes and semantic segmentation tools applied to the process are addressing a common set of parameters. The objects vary widely from one project to another, but the challenge remains the same: building a ground truth that separates and identifies objects within a set of images that are central to a business or government objective.

Geospatial ground truth is a hard-won result, though, particularly when it involves so many different types of landscapes and objects within them. Imagine spending years developing expertise in agriculture, with different crops, fields, and other natural features juxtaposed within a set of images – and then transitioning to maritime analysis, with crowded harbors filled with fleets of different vessels, cranes, storage facilities and the like. The switch from natural to manmade objects, relatively slow-moving developments to vastly more quicker temporal changes, and even differences in overlapping objects all present barriers to efficient and accurate data annotation. That's why geospatial data analysts tend to specialize in specific verticals that enable them to develop expertise through experience – and in the process build institutional knowledge that can be applied to new projects and even new, related use cases.

ONE VIEW, MANY PERSPECTIVES

The use cases for all of this data run a broad gamut of public and private sector activities, among them land use planning, commercial and residential insurance, agriculture, and of course national security. In essence, any initiative that relies wholly or in part on ground truth data – and temporal changes to that ground truth – is either a current or potential customer of geospatial imagery providers and analysts. The implications of this for imagery annotation and analysis effectively limits the vendor field to those companies with both broad and deep experience and expertise. The scope of potential applications has led to a corresponding demand for customized workflow – and project design – to adapt to the seemingly endless number of applications. Unlike other sectors that rely on image and data annotation, such as autonomous vehicle development, geospatial data labeling demands an ability to interpret (and differentiate between sources of) detailed satellite and drone imagery, and to then process the work according to the needs of often very specialized use cases.

Sometimes the work involves fairly static data – such as the makeup of an established residential community. But the work more often requires an expertise in temporal changes – everything from the management of a constantly in motion commercial port to the slower rhythms of crop growth and yield predictions.

The land use classification system first developed by the United States Geological Survey in 1976, and implemented through manually annotated aerial photography into the early 1980s, delineates nine so-called Level I broad land use categories, ranging from urban, to agriculture, to forest and even tundra and perennial snow and ice. But those classifications are further broken down into a series of Level II sub-categories. Urban land, for example, can be labeled forest industrial, commercial, industrial, transportation, industrial and commercial complexes, mixed use and other built-up land. And even a sub-category as straightforward as "residential" can be further detailed as single-home, multiple dwelling, residential hotels, mobile home parks and the like.

Depending on the use case, the same USGS data can be used for a richly diverse set of objectives. A proposed oil pipeline project might require a detailed analysis of forested wetlands, bays and estuaries, and even shrub and brush rangeland in order to develop ecological impact statements. The same data might be analyzed with an entirely different set of data targets (and annotation objectives) for wildlife management, drought impact reports or virtually anything else that requires the breakdown and delineation of a set of land use imagery and sensor data.

One use-case example that illustrates both the challenges and parameters of a geospatial intelligence project: monitoring oil supplies through surveillance of oil tanks and temporal changes to stored supplies. In this scenario, the energy company (or government entity) has a staged set of objectives: (1) identify each individual oil tank within a storage facility; (2) develop a system to describe the percent of capacity in use; develop the means to track temporal changes in capacity.

Geospatial data can be retrieved through satellite imagery or airborne surveillance, such as unmanned drones. That imagery is then annotated to describe the field of storage tanks (and individual tanks) through bounding boxes and semantic segmentation. Then each individual tank must be further described with a polygon annotation tool to identify the top, bottom, and even shadows of each tank, and then classified by type (floating roof, fixed roof, etc.), before rendering an adjustable 3D image.

Sometimes the use case is deceptively simple – even if the subsequent work is anything but. Consider the dilemma of a major American big box retailer that needs to forecast retail traffic and resulting financial results for a series of brick and mortar locations around the country. The retailer might use consumer and even customer surveys, historical in-store data, and external economic information to determine patterns that can then be used to build predictive models. But there's a saying that people vote with their feet (or cars) and one of the quickest and most convincing ways to acquire a series of snapshots on a store's performance is to analyze car traffic in the store's parking lot. Feeding a predictive engine, though, requires several steps to move from raw images to analytical data. On a basic level, image annotators need to identify the objects that are cars, trucks, minivans and other vehicles. Which vehicles are likely owned by families, and which ones aren't? How long do shoppers spend in the store (based on parking time)? How does parking lot use vary by time of day, and day of the week?

Regardless of the specific use case (and custom requirements), however, there are common elements that play across the various industry verticals that comprise the sector's customer base. Geospatial data analysts undergo at least two main levels of training.

> Sometimes the use case is deceptively simple – even if the subsequent work is anything but.

First, they need to develop a core competency in analyzing satellite and aerial (drones and planes) imagery, and even use of the more basic tools, such as Google Maps and Earth. Viewing a satellite image, for example, is a meaningless exercise without the ability to break down that image to its component objects – including the expertise, for example, to differentiate a collection of agricultural sheds from a missile launching facility.

The next level of training adds domain expertise for more nuanced annotation work, and often requires a degree of business intelligence capabilities as well as the more technical aspects of annotation. Residential insurance providers, for example, need imagery broken down by property lines, main and secondary structures, roof types and composition, ancillary features, such as in-ground and above ground swimming pools, and the like. A project to mark retail footprints, on the other hand, might require expertise in tagging both merged and acquired stores in different locations and different markets. This type of work often involves intensive searches for the most appropriate raw data and then the creation of a structured custom workflow to build the required data sets.

And because of the sheer variety, and divergence of potential projects, the sector has seen the development of its own specialized annotation tools specific to different industry verticals and use cases. The tools are designed specifically to extract vertical-specific intelligence from raw geospatial data.

TEACHING AN ALGORITHM TO SEE

As with other sectors where imagery is annotated to train an algorithm, ground reference data for machine learning must be specific, singular, bounded and geo specific. Given the preponderance of still imagery rather than video, algorithm training focuses on the identification of relatively static objects – such as roads, oil tanks, crops, and structures of various types – though tracking temporal changes are vital to many applications, such as agriculture, ship building and national security.

Useful ground reference data must embody four key attributes...

1. Include specific geolocation data – the coordinates that place the object in a geographically identifiable location
2. Describe one, singular type of object
3. Include specific geometry that sets one object apart from another
4. Have clear boundaries such that any pixel in the image can be described as within or outside the object's borders

In addition, data annotation must include the required metadata to enable the end user to feed the data through the algorithm. Typically, the metadata will include a date stamp, the geolocation coordinates, classification fields, for example, in agriculture, crop type(s), yield, land cover, etcetera, and any required descriptions. The end goal is applying deep learning methods, such as Convolutional Neural Networks (CNNs), to specifically targeted use cases, where the raw data processing and metadata output can be employed to develop a use-case-specific, finely tuned algorithm.

Korean researchers in 2018, for example, published the results of their research project using drone imagery that had been processed with semantic segmentation to help train an algorithm to detect unregistered buildings in Korean urban areas. By identifying and detailing the buildings in each image, using semantic segmentation tools, they hoped to develop a machine learning system that can compare their building prediction maps against existing local government GIS data to spot illegal building projects and unauthorized land use.

SEMANTIC SEGMENTATION IS GEOSPATIAL KEY

Most agree semantic segmentation is the cornerstone tool of data annotation, a tool based on a labor intensive, multi-step process to accurately break down a satellite image, for example, into its component objects. On a pixel-by-pixel basis, a data analyst will identify and set the boundaries of each object within the image. The segmented objects are then used to train an algorithm to detect it in various other images (and differentiate it from other objects).

In the Korean researchers' initial implementation, they applied semantic segmentation to the drone imagery to generate a building prediction map, and then subtracted existing government GIS data from their prediction maps. They assumed the remaining pixels after that processing were unauthorized buildings. While it remains a work in progress for commercial implementation (the research did hit some initial rough patches due, in no small part, to the use of automated annotation tools with little or no human intervention), it largely succeeded as a proof of concept.

No geospatial market tasks are easy to facilitate. Below are examples of building footprint extraction results processed by open-access publisher MDPI that show how different processes sometimes lead to different results. MDPI analyzed the cities of Las Vegas (A, B), Paris (C, D) Shanghai (E, F) and Khartoum (G, H). Annotations in green denote correctly extracted buildings (True Positives), annotations in red denote other objects extracted as buildings by mistake (False Positives) and annotations in yellow denote ground truth buildings that were not extracted correctly by the processing method (False Negatives).

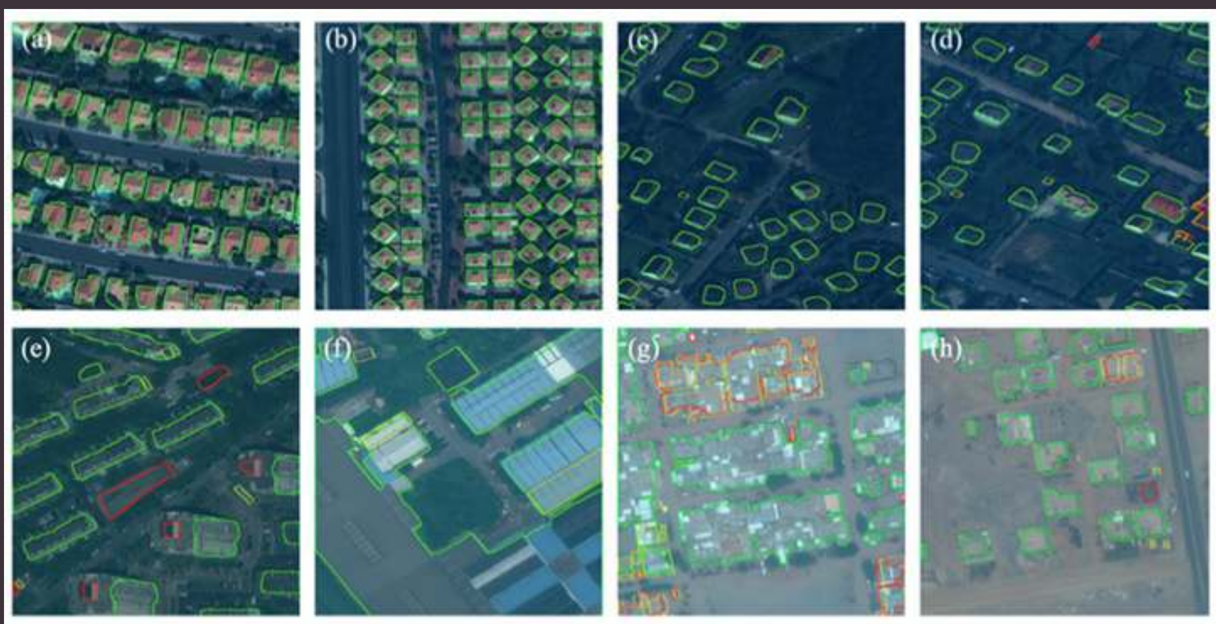

BUILDING ARTIFICIAL INTELLIGENCE THROUGH ANALYSIS

All of this presupposes good quality imagery that eliminates the "garbage in, garbage out" rule of data analysis. One of the more crucial roles of a geospatial imagery annotator is (1) determining the suitability of the imagery to the project's objectives and (2) identifying holes in the imagery library that might limit the ability to train the algorithm. While some, typically repetitive data analysis can be automated, the most challenging data annotation projects require an experienced set of human eyes. That approach holds true regardless of whether the imagery involves a mixed-use community in the Midwest United States or a potential military hotspot in the Middle East.

In 2017, speaking at a geo-intelligence symposium, the director of the National Geospatial Agency, Robert Cardillo, noted that the agency planned to automate as much as 75 percent of its repetitive imagery analysis – primarily so its analysts could devote their time and expertise to the 25 percent of critical imagery that requires a skilled and experienced set of eyes. "To manually exploit the imagery we will have over the next 20 years, we would need eight million imagery analysts," he said. "Even now – every day – in just one combat theater, with a single sensor, we collect the data equivalent of three NFL seasons in high definition. We intend to automate 75 percent of the repetitive tasks our analysts perform so they have more time to analyze that last play and more accurately anticipate the next one. And then they can look much harder at our toughest problems – the 25 percent that require the most attention."

SEPARATING ONE DATA ANNOTATION PROVIDER FROM ANOTHER

The almost endless variety of geospatial intelligence projects argues against anything resembling a cookie cutter approach to data analysis – and places a valuable premium on experience, training, and the institutional knowledge of a company that specializes in data annotation. By developing a cross-disciplinary expertise built on multiple use cases and different machine learning scenarios, a top data annotation provider develops the agility required to tackle even the most challenging assignments.

That in itself explains the growth in human capital devoted to geospatial image annotation in the past few years – both in terms of specialized teams that can build vertically aligned expertise over time (developing institutional knowledge that can be applied to similar projects with overlapping requirements), and the individual annotation experts who understand the difference between useful raw imagery for a given project and imagery best left on the cutting room floor. Moreover, as with any successful collaboration, the data annotation firm and client must be able to take the specific data requirements of the project and develop robust guidelines the annotation solutions team and client use to support QA goals. Those guidelines become "live" documents that are updated through a feedback loop that is used to improve quality and accuracy metrics.

NEXT GENERATION 9-1-1 TO BE NEXT BIG THING

Few initiatives illustrate the opportunities and challenges inherent in the use of geospatial data as well as the move to next generation 9-1-1 services (otherwise known as NG9-1-1). As emergency response systems transition from switched network landline infrastructure to Internet Protocol communications, the organizations tasked with managing emergency response systems are envisioning a blank canvas to be filled by entirely new categories of rich data, where (almost) anything is possible.

Anyone designing an emergency response system from scratch in 2020 would have a common set of basic requirements:

- Precisely determining the location of the emergency – not just as a point on a map, but complete with sub-address data including location within a structure, the nature of that structure and the nature of the surrounding topography that may affect the ability of first responders to act effectively.
- Calculating that location relative to the various 9-1-1 call centers to which telcos route emergency calls – and then the optimal routing for first responders, taking into account the ground truth of each potential route.

The underlying opportunities inherent in those system requirements: over time, provide first responders with an incredibly rich data set that could include a 3D model of house on fire – complete with any barriers, access points, bedroom locations, and building materials – optimal routing to the destination and even live updates.

The challenges: developing and / or gaining access to the geospatial imagery (complete with annotation) and integrating that data within the NG9-1-1 infrastructure. And doing so with a common metadata structure that can be implemented across multiple constituencies, service areas and first responder organizations.

That dual set of challenges and opportunities also creates an opening for public and private sector organizations that already collect, annotate, and organize geospatial data for other purposes. NG9-1-1 infrastructure developers are already eying ways to integrate annotated imagery in call response centers – but gaining access to that data is a significant challenge and will likely remain the case for the next few years at least.

This will be a walk first and then run transition.

The typical 9-1-1 call center, where dispatchers take incoming distress calls and route the emergency to the correct first responders (police, ambulance, and / or firefighters), is based on communications architecture that goes back decades, if not longer. The system is designed to coordinate response based on (1) the location of the emergency, (2) the emergency response service sector for that location, and (3) the type of emergency.

That infrastructure, though, has largely failed to keep pace with advances in technology as well as changing personal habits – most notably the widespread transition from landlines to cell phone use. Most 9-1-1 call centers, for example, base a cell phone's location on the nearest cellphone tower rather than precise latitude / longitude, to say nothing of altitude / floor location.

In a world where telephone calls – and various forms of rich data – increasingly travel over IP networks – and where basic telephony is but one option for communicating, diagnosing and managing responses to emergency situations, it was inevitable that emergency response systems evolve to keep up.

Over the past five to six years, government agencies in the U.S., Canada, and the European Union have been developing the regulatory framework for the transition to NG9-1-1, and the emergency response ecosystems have been grappling with all of the questions and challenges inherent in the transition.

"We're preparing for the impact of GIS on (next generation) 9-1-1," says Robert Darts, Senior Service Delivery Manager, Fire Technology and GIS, at E-Comm 911, the emergency services center operator for much of British Columbia, Canada.

Darts, who co-chairs two of the Canadian standardization bodies for NG9-1-1, notes that much if not all of the Canadian standards related to GIS will likely closely match the work already in progress under the auspices of the National Emergency Number Association (NENA), the U.S. standards setting body for emergency response infrastructure.

The new NENA Standard for NG9-1-1 GIS Data Model specifies three types of spatial data for its initial implementation...

- Points – Discrete locations such as address points, premise locations, and hydrants.
- Lines – Linear features, such as roads, rivers, and railways.
- Polygons – Geographic coverage areas such as (public safety answering points i.e.- 9-1-1 call center) boundaries as well as Emergency Service Boundaries, and cities.

Just as importantly, it establishes a common set of metadata attributes for the exchange and use of GIS data, as well as a roadmap for rich data sets based on multiple GIS data file formats. The standard also leverages existing GIS data specifications, including the World Geodetic System 1984 (WGS84) as well as the geodetic parameters specified by the European Petroleum Survey Group for both 2D and 3D geometries.

The scope and detail of NG 9-1-1 metadata requirements can be seen in the following NENA table describing the boundary of a neighborhood, subdivision, or commercial area for a an emergency response call center operator.

Descriptive name	Field Name	M/C/O	Type	Field Width
Discrepancy Agency ID	DiscrpAgID	M	P	75
Date Updated	DateUpdate	M	D	-
Effective Date	Effective	O	D	-
Expiration Date	Expire	O	D	-
Neighborhood NENA Globally Unique ID	NbrhdNGUID	M	P	254
Country	Country	M	P	2
State	State	M	P	2
County	County	M	P	75
Additional Code	AddCode	C	P	6
Incorporated Municipality	IncMuni	M	P	100
Unincorporated Community	UnincComm	C	E	100
Neighborhood Community	NbrhdComm	M	E	100

Source: National Emergency Number Association

In the NENA table above, each row of data contains a descriptive name, a field name, which is used when exchanging GIS data, a Mandatory (M), Conditional (C) or Optional (O) status, the types of data that are used within the data field, and the field widths which represent guidelines for interoperability. This table, and configuration, also includes the neighborhood name -- helps resolve locations of similar sounding street names, and is critically improtant to telecommunicators.

The planned transition to NG9-1-1, expected to take several more years, will include coordination with local and regional governments in the U.S. and Canada, including what will amount to a gap analysis of available GIS data.

As Darts notes, local and / or state / provincial governments will need to aggregate entire libraries of new data for emergency response systems. “Some smaller governments might not have all the data,” he says. Regardless, the days of emergency response call centers relying solely on raw Google Map and Google Earth data are numbered.

PRIVATE, PUBLIC SECTOR GEOSPATIAL USE CASES

The multitude of use cases for geospatial applications can be seen running a broad gamut of public and private sector activities, including land use planning, commercial and residential insurance, agriculture, national security, oil and gas exploration, and retail activity. The data is gathered through satellites, drones and other aerial sources that capture everything within a specific geographic area, and the data annotation required varies depending on the final use case associated with the engagement. At iMerit, workflows and project design are tailored to adapt to the seemingly endless stream of geospatial applications.

Industry	Annotation Type	Use Case
Transportation	Classification	Classification of the ship building stage to identify the time frame it takes to build a ship in specific docks
Transportation	Point Annotation	Annotation of the head and tail of ships in open waters in strategic regions
Transportation	Point Annotation	Classification of an oil tank based on visual features - floating roof top, fixed roof top, cleanliness, etc.
Transportation	2D Polygon Annotation & Classification	Annotation of military aircraft using bounding boxes and classification of the type of aircraft
Transportation	Semantic Segmentation - Change Detection	Semantic segmentation of parking lot fill percentages to quantify the traffic in certain facilities
Transportation	Polyline Annotation	Annotation of polygons to track railways
Transportation	Classification	Annotation of points on railcars

Industry	Annotation Type	Use Case
Global Oil & Gas	Point Annotation	Identification of each individual oil tank using a point on a region, using Google Maps
Global Oil & Gas	2D Polygon Annotation	Annotation of the top, bottom, and shadows of an oil tank using a circle annotation tool
Global Oil & Gas	Classification	Classification of an oil tank based on visual features - floating roof top, fixed roof top, cleanliness, etc.
Global Oil & Gas	3D Polygon Annotation	Annotation of an oil tank in 3D using a 3D annotation tool
Global Oil & Gas	Content Classification	Creation of metadata on individual oil tanks as well as refineries/terminals to capture the volume of oil, number of oil tanks, type of oil, etc.
Global Oil & Gas	2D Polygon Annotation	Annotation of oil pipelines
Retail	Content Classification	Creation, completion, and verification of a complete list of stores, associated with an entity (e.g. a big box retailer)
Retail	Point Annotation	Identification of entrances of manufacturers, distributors, brick and mortar stores of publicly traded companies using points
Retail	2D Polygon Annotation	Annotation of polygons around the building footprint of manufacturers, distributors, brick and mortar stores of publicly traded companies
Retail	2D Polygon Annotation	Annotation of polygons around parking lots of manufacturers, distributors, brick and mortar stores of publicly traded companies

Industry	Annotation Type	Use Case
Energy	Classification	Classification of the status of a power plant as operational or non-operational based on gas being emitted
Land Use	Image Classification	Classification of images as cloudy, corrupt, or clean to filter which image can be used for training an image classification model
Land Use	Semantic Segmentation	Semantic segmentation on the features of an image (for example buildings, roads, vegetation, desert roads); level of detail ranges from general to very specific
Land Use	2D Polygon Annotation	Annotation of polygons around features in an image (for example intersections, parks, farm land, etc.)
Land Use	Semantic Segmentation - Change Detection	Semantic segmentation on a large area of various geographies and ensuring temporal consistency while annotating over time
Natural Resources	2D Polygon Annotation/ Semantic Segmentation	Semantic segmentation or polygon annotation around piles of coal or iron ore to quantify the volume of a particular resource, and classification of each pile by height level so that the estimated volume of each annotated pile of natural resource can be quantified
Natural Resources	Content Classification	Creation of a global list of iron ore/coal terminals
Natural Resources	2D Polygon Annotation	Annotation of the parts of the mines that are dug up

Industry	Annotation Type	Use Case
Residential Real Estate	2D Polygon Annotation	Annotation of polygons around each individual single family home from an aerial perspective
Residential Real Estate	Image Classification	Classification of annotations on residential homes by building stage, such as foundation, framing, and roofed
Commercial Real Estate	Content classification	Creation of a complete list of various information on REITs, such as grade level, number of stores, store location, capacity, and size
Insurance	2D Polygon Annotation	Annotation of polygons around features of a residential building (e.g. windows, doors, garage, swimming pool, etc.) to assess the insurance rate
Insurance	Image Classification	Classification of annotations of features of residential buildings
Agriculture	2D Polygon Annotation	Annotation of palm plantations or individual palm oil trees and their differentiation from natural forest
Agriculture	Classification	Classification of palm oil annotations by different maturity levels

iMerit

For information on iMerit's Geospatial solutions contact Mallory Marcone in the Americas and Soo Yang in other parts of the world.

SOO YANG
Solutions Architect
soo@imerit.net

MALLORY MARCONE
Solutions Architect
mallory@imerit.net

iMerit

DATA SERVICES FOR GEOSPATIAL GROUND TRUTH

imerit.net