CrowdAI

Company

CrowdA uses cutting edge deep learning techniques, including convolutional neural networks, to automatically detect features from drone, satellite, aerial imagery, and multi-spectral data. CrowdAI has produced some of the most accurate deep learning models (classification, bounding boxes, segmentation, and instance-based segmentation) in the industry and has created partnerships with some of the largest satellite companies in the industry. Our fast, flexible models allow for detection and categorization of numerous features across large geographic regions, including buildings, roads, ships, vegetation, water, and more. CrowdAI currently works with three USG agencies and several commercial partners in insurance, telecommunications, mapping, and oil & gas.

CrowdAI is a team of computer scientists and deep learning specialists from IBM Watson, OpenAI, UC Berkeley, University of Oxford, and Google.

CrowdAI is committed to giving back to the community. CrowdAI lends its technology pro bono for various social good projects specifically issues associated with natural disaster issues. Published research paper on this subject: https://arxiv.org/abs/1812.07033

Differentiators

Other ML vendors who service insurance industry primarily focus on proactive use cases, have less experience on detecting vegetation, are less scalable during time-sensitive events, and use less sources of imagery. CrowdAI differentiates itself from other technologies by:

- **On premise/Cloud deployed** - Customer collects all the data. CrowdAI deploys models in Customer's private clouds. Instant and unlimited analysis of Customer's imagery. Limits latency and data privacy issues. PII data stays internal.
- **Ownership** - Customer owns all the data.
- **Geographic scale** - Classify, detect, and monitor objects over large Areas of Interest
- **Speed** - Classify, detect, and monitor thousands of objects of an entire city in a matter of minutes
- **Customization** - Proprietary models per use case - per customer; trained on customers data and specific object needs.
- **Third party validation** - Won highly competitive bids with US government and publish peer-reviewed white papers
- **Imagery agnostic** - Quickly analyze data regardless of the source of the imagery

- **Maturity** - Deep-learning models have been training for 3.5 years across 125 countries, multiple biomes and different imagery types: 3m Planet (monitoring daily flood change post-disaster), 30-50cm Digital Globe (buildings and vegetation encroachment, swimming pool, roof size, boats), 5.8cm Aerial (roof tile missing, fire hydrant) and 3cm private source (cracks on sidewalks or roofs). CrowdAI creates all the training data. Trains models in 2.5 weeks or less.

How it works

Convolutional Neural Networks

CrowdAI uses the latest deep learning techniques, including convolutional neural networks (CNNs), and proprietary computer vision techniques to identify features such as roads, ships and buildings at scale. Specifically, the CNNs we use identify each pixel as either “object: or “not object” using segmentation approaches discussed in the following sub section.

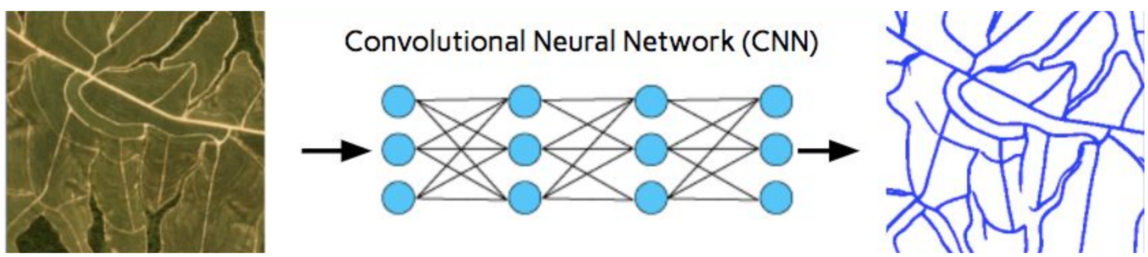

Currently, we use cutting-edge techniques–like fully convolutional encoder-decoder networks with skip connections—to identify and classify target features. One of the key challenges when trying to segment fine-grained objects is the need for relevant context that allows the model to build a consistent definition of the object. The way we solve this is by carefully analyzing the receptive field of the neurons at each layer of the CNN and adding the appropriate amount of model capacity at each depth.

Another issue especially persistent in satellite imagery is drastic variability in viewpoint, color distortion and compression artifacts. We solve this issue by synthetically adding these artifacts during the training process to regularize the model and improve the robustness of the performance across many regions. Finally, we use standard normalization techniques like batch-norm to significantly speed up our training process. The figure below is an example of a typical encoder-decoder network architecture we use for building our semantic segmentation models.

Segmentation Approach

For objects like buildings, CrowdAI uses an instance segmentation learning paradigm wherein the model explicitly learns to delimit many instances of the same object. This unlocks useful applications like counting and surface area calculation. We achieve significant improvement on state-of-the-art by incorporating a "human-in-the-loop" for training these models.

A common limitation for instance segmentation models is the ability to delimit and detect small objects. We overcome this by integrating our proprietary semantic segmentation model. We have built a suite of proprietary segmentation models that incorporate and iteratively improve upon lower quality annotations. This line of research inquiry is underexplored in academia, where the datasets are generally fixed and an infrastrastructure that is tightly looped with human annotators in unavailable. Leveraging improvements from this research has allowed us to build models using a significantly lower number of annotated examples for a given performance requirement. In addition to lowering the cost of the annotation process, this also has the potential to unlock deep learning applications for domains where training data is hard to obtain.

Model Architecture Types

U-Net: One of our initial models was inspired by the family of U-Net architectures, where low-level feature maps are combined with higher-level ones, which enables precise localization. This type of network architecture was designed to effectively solve image segmentation problems, particularly in the medical imaging field. U-Net is generally a default choice for many types of segmentation challenges. The encoder of the model consists of a Visual Geometry Group (VSG) network with the addition of batch norm and downsampling layers. The choice of the depth of the network will be informed by careful analysis of the dataset, task, and the receptive field. The number of feature maps throughout the network will need to be determined by key observations so that we can afford the network to lose some representational power in the encoder half, especially since the model has access to low-level features in the decoder half via the skip connection. The decoder is similar to the encoder where—instead of max-pooling—we use deconvolution layers to upsample with a skip connection from the encoder, combining deep representations of the prior decoder layer with more precise spatial representations from the corresponding encoder layer. The final head may consist of a sigmoid activation function.

DeepLab: Our encoder will consist of a ResNet block with a yet-unknown number of convolutional layers with a yet-unknown number of kernels. We have yet to determine the stride and residual blocks and number of convolutional layers and how many parts will divide the whole encoding network. We also have to determine the number of fully convolutional layers for the decoder and the number of kernels. Each of these layers up-samples its input to be double its resolution. The last convolutional layer converts the feature map into scores followed by a sigmoid activation function. Thus, the whole network consists of a yet-unknown number of convolutional layers for the encoder, a yet-unknown number of fully convolutional layers for the decoder, followed by a convolutional layer to output the class labels.

Residual Inception Skip Net: This model is a convolutional encoder-decoder architecture with inception modules instead of standard convolution blocks. Instead, the inception models include asymmetric convolutions. For example, a 3×3 convolution is replaced with a 3×1 convolution, then batch norm followed by 1×3 convolution. This methodology reduces the number of parameters while providing similar performance. All weights are initialized with the He norm and all convolutions are followed by batch norm and then activation. Similar to U-Net architecture, we will also add a skip connection linking identically sized layers between the encoder and the decoder.

Existing Detectors & Capabilities

These are the existing algorithms CrowdAI has built and deployed. Additional features can be considered, but there will need to be a significant business opportunity present to justify development time and resources.

Custom models not mentioned below are available to government and enterprise customers upon request.

- Buildings
    - Detection - buildings of all type on 30-50cm imagery, large buildings only beyond that; can detect precise georeferenced building footprints on Airbus and (some) Planet imagery; building centroids are available
    - Classification - categorize into building type (commercial, residential, industrial) based on size and context
    - Change detection - can determine building damage after disasters (e.g. hurricanes, wildfires); can determine where new buildings are built between two time periods
- Roads
    - Detection - roads/paths of all types on 30-50cm imagery, most roads on 3m imagery (first company to prove out roads on Planet's 3m daily imagery)
        - Additional detection options related to roads: sidewalks, bike lanes, crosswalks, pedestrian plazas
    - Classification - can classify roads by surface type (paved, unpaved) and width; can count lanes on 30cm imagery
    - Change detection - determine where new roads are built across a city, region, state, or entire country between two time periods

- Ships
    - Detection - identify ships of all sizes on 30-50cm imagery, large ocean-going vessels on 1.5-3m imagery.
    - Classification - can determine ship size and orientation/directionality
- Well Pads
    - Detection - identify oil/gas exploration & production well pads at scale across entire basins on 3m imagery or better
    - Classification - determine pad activity state (e.g. “no activity” or “possible fracking”)
    - Change detection - identify where new pads are built and what activity status has changed across entire regions on a weekly basis
- Vegetation
    - Change detection - determine areas of vegetation growth/destruction on 3-5m+ imagery
- Vehicles
    - Detection - identify and count all vehicles on 30cm satellite and all aerial/drone imagery. Can create density maps.
    - Classification - determine vehicle type (sedan, SUV, truck, etc.) on sub-30cm aerial and drone imagery/video.
    - Tracking - in full-motion video, can track individual vehicles throughout frames
    - Change detection - Helps understand “patterns of life” by determining where vehicles of different types tend to be in a sequence of images/videos over time (e.g. “How many passenger vehicles tend to be in this parking lot?”)
- Water/Land Cover
    - Change detection - determine changes to land and water cover from 3-5m+ imagery

Geographic Range

CrowdAI's models are trained to be globally flexible and can be fine-tuned on imagery from commercial and proprietary satellite, manned aerial, and unmanned aerial sources. The geographic range is relatively unlimited, as satellite constellations image the landmass of the earth. There is one important caveat: most areas of the open ocean are not imaged regularly, and therefore are not typically available.

Global Reach of Current Customer Base

Grab - Southeast Asia
NTT Data - Japan and Australia
US Government - Global
Intel - Uganda

Resolution Quality

Different objects have different resolution quality based on available imagery sources. All of these are commercially available over major populated areas. In addition, post-CAT, imagery is often sourced from 25cm/pixel aerial data. CrowdAI has strong partnerships with all the major third-party imagery providers including (but not limited to) Digital Globe, Airbus, Nearmap, NOAA, and Planet.

Object	Source	Quality
Shingles	Nearmap	~5.8cm/pixel
Commercial Buildings	Digital Globe, Planet	30-50cm/pixel and 3m/pixel
Residential Buildings	Digital Globe	30-50cm/pixel
Roads	Digital Globe, Planet	30-50cm/pixel and 3m/pixel
Roofs	Digital Globe	30-50cm/pixel
Vegetation	Digital Globe	30-50cm/pixel
Vehicles / Boats	Digital Globe	30-50cm/pixel
Water	Digital Globe	30-50cm/pixel
Pavement	Digital Globe	30-50cm/pixel
Trampoline	Nearmap Commercial	~5.8cm/pixel
Fire Hydrant	Nearmap Commercial	~5.8cm/pixel
Solar Panels	Nearmap Commercial/ Digital Globe	~5.8cm/pixel and 30-50cm/pixel

Proven Success in Insurance

- One of the largest American property insurance companies based in New York. The carrier's Chief Architect sponsored the project and a team of data scientists with expertise in AI put CrowdAI through an extreme vetting process while they evaluated several other options - including building in-house. Their business driver for quoting is roof type, roof condition, and roof square footage. Use Case: Underwriting. Metrics: Speed and Accuracy
- European multinational insurance company. Project with NatCAT risk team to analyze imagery from two previous CAT events. Use Case: Claims. Metrics: Reduction in time to close a claim and its impact on CSAT and reduction in loss adjustment expense.

Proven Success Outside of Insurance

When disasters strike, one of the largest telecommunications companies in the United States used to deploy field teams to determine the cost of damage, estimate clean up time, and allocate additional on-the-ground resources. This process would sometimes take weeks because of the need to meticulously comb NOAA satellite images. The longer it took to clean up after a storm, the longer it would take to turn on electricity, and the more calls they would inbound at their support center from unsatisfied customers. After Hurricane Michael, the company used CrowdAI to detect, classify, and monitor thousands of objects in a matter of minutes. The technology was able to classify objects by category and severity of damage which made it easier to estimate costs and determine the skill sets needed for their field team. More accurate and quicker data reduced operational costs which helped turn on electricity quicker and increase satisfaction of their customers.

https://blog.crowdai.com/assessing-damage-after-the-storm/

Example of deliverable to telco customer:

Hurricane Harvey was one of the costliest storms on record that left hundreds of miles of roads under water. A large location intelligence company needed to measure the extent of the damage to provide better data to first-responders in the Houston area. When previous storms occurred, it took days for the Company to manually comb over areas of interest with the use of publically available imagery. With CrowdAI, in a matter of minutes, they were able to analyze the before and after imagery to detect and classify flooding and the extent of damage from the flooring on roads. This saved the company significant operational expenses and increased the speed of aid to those in need.

Example of deliverables to Customer:

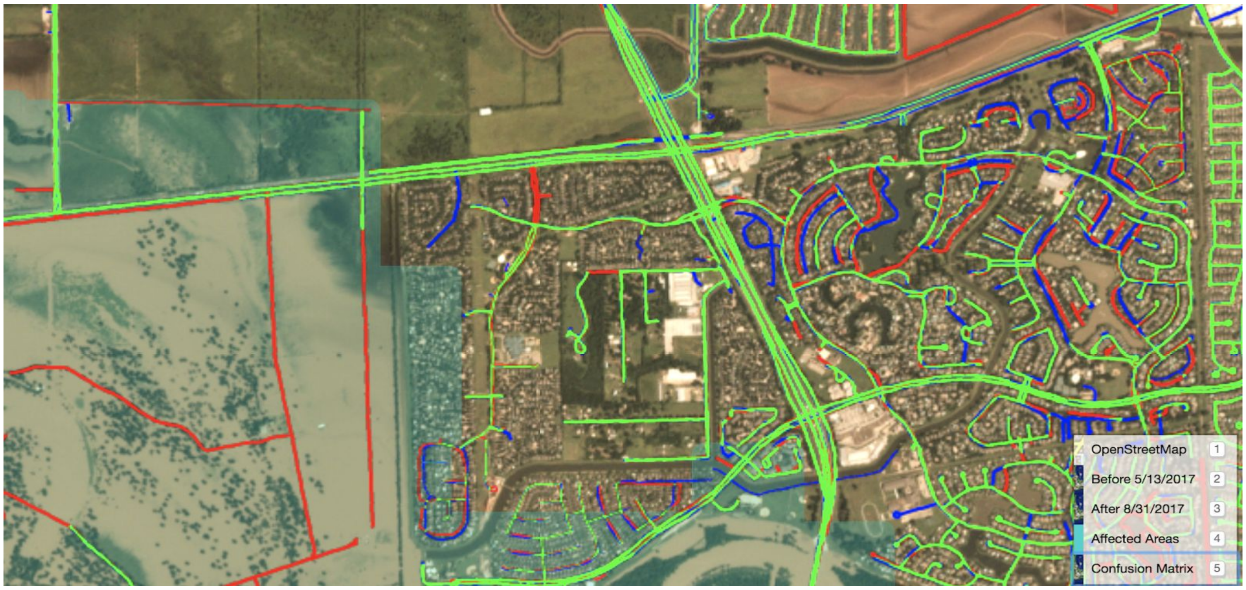

CrowdAI Points of Contacts

Executive Sponsor - Devaki Raj, CEO - Devaki@crowdai.com
Business Lead - Richard Purcell, Head of Commercial Sales - Richard@crowdai.com
Technical Lead - Nic Borensztein, CTO - Nic@crowdai.com