

West South-Central Division Single-Family Housing Starts Using Satellite Imagery Technical Documentation

1. Introduction

The U.S. Census Bureau has been exploring the feasibility of using satellite imagery to measure new residential construction activity for the Survey of Construction (SOC). Estimates for the West South-Central (WSC) Division single-family housing starts utilizing satellite data are being published as an experimental data product¹ to serve as proof of concept for satellite imagery as a data collection method for SOC.

For permit issuing places in the SOC sample, samples of new residential building permits are selected monthly and data collection is accomplished by field representative interviews or site visits. With the new data collection method, satellite images for a subset of places in the SOC sample for the WSC division are collected monthly, allowing for the observation of all new residential construction in the selected places rather than a sample of new residential construction projects. The starts data obtained from the satellite image are then combined with the SOC survey data for sampled places not collected by satellite in the WSC division to compute a division level estimate of monthly housing starts.

2. SOC background

2.1 Sample Design

The Survey of Construction sample design consists of three stages: (1) a subsample of the 2004 Current Population Survey (CPS) primary sampling units (PSUs), which are land areas (groups of counties, towns or townships within a state) that represent the entire United States; (2) selection of permit/non-permit areas; and (3) a monthly selection of permits. The permit areas correspond to a building permit office (BPO). Currently, no changes to the sample design are planned to accommodate the introduction of satellite data collection. See the [SOC methodology page](#) for additional details on the SOC sample design.

2.2 Data Collection

The major innovation for the experimental data product is the use of satellite imagery as a data collection method for the entire area covered by a BPO for a subset of BPOs in the current SOC sample. By observing the entire area covered by the BPO, the final stage of sample selection, the permit sampling stage, is eliminated, expanding the sample to include all permits from the selected BPO. For the initial satellite implementation, the selected BPOs are self-representing places from self-representing PSUs so their combined first and second stage sampling weights are equal to one.

¹ The Census Bureau has reviewed this product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data (Disclosure Review Board (DRB) approval number: CBDRB-FY25-0086).

3. Scope

3.1 Geographic

The current set of BPOs that are collected by satellite are self-representing BPOs from self-representing PSUs in the current SOC sample for the West South-Central Division (Division 7). Given the limited use of satellite data collection at this time, the estimate is at the division level for this division only. This provides a more geographically granular tabulation than the current production estimate which goes down only to the region level.

3.2 Type of Construction

The estimate is limited to housing starts. Although there are many data items that are collected for SOC the focus for this product is housing starts because they are easily observable from satellite imagery and match the SOC definition of a housing start. Estimates for other construction stages will require further development and is an area of future research.

3.3 Residential building type

This product is limited to single-family homes because each observation corresponds to one housing unit and the estimate can be obtained by simply counting the number of observed starts.

Furthermore, single-family homes are typically built in areas with other single-family homes. These single-family residential neighborhoods are easily identified on satellite imagery and are used in the building type model described in section 5.1 to limit the counting of starts to areas that are likely to contain single-family homes. However, multifamily buildings may be hard to distinguish from an office building or hospital at the start of construction. Estimating multifamily housing starts from satellite imagery is another area for future research.

4. Image Collection

Areas of interest (AOIs) for satellite data collection are selected from the existing SOC sample of permit issuing places, focusing on areas with significant single-family home construction activity. These AOIs include incorporated places (cities, towns, and villages), census-designated places (CDPs), counties, and unincorporated areas. **Places** refer to specific, bounded areas that are either incorporated under state law with their own local government or, in the case of CDPs, unincorporated but identified for statistical purposes. **Counties**, on the other hand, are larger administrative units that encompass multiple places as well as unincorporated areas. To define the exact boundaries of these AOIs, Census Tiger/Line Shapefiles are used. County geometries are derived from the U.S. County file, while geometries for places come from state-specific place files. **Unincorporated areas** of a county are determined by combining county geometries with incorporated places, identified through a spatial join using BPS data and place Tiger/Line Shapefiles, leaving the unincorporated portion as the new AOI for that county's unincorporated area.

Once the geometries are established, they are sent to an image vendor for augmentation to meet satellite tasking requirements. The vendor simplifies the AOIs by reducing the number of vertices to ensure the total for the project remains below a defined limit. Additionally, if an AOI is too small to meet tasking requirements, it may be combined with a larger nearby AOI. In some cases, the vendor may remove areas such as airports, state and federal parks, military bases, and bodies of water, where residential construction is not possible. We review the modified AOIs with the vendor to confirm that any changes are appropriate. These augmentations do not affect the original AOIs used

for data reporting. Finally, the vendor and image provider assess whether the areas can be captured during the required timeframe before tasking begins.

Ideally, satellite imagery would be captured on the first day of the month following the reference month. However, weather conditions can render an AOI or portions of an AOI unobservable by satellite on that exact day. A longer image collection window allows for a greater chance of obtaining usable satellite imagery. The current satellite tasking window has been set at the 25th of the reference month to the 5th of the following month.

5. Measuring housing starts from satellite images

Single-family housing starts are identified from satellite images using a series of Convolutional Neural Network (CNN) models and post-processing steps. To identify construction stages present in an image, we use a U-Net Attention model that integrates an attention mechanism in its encoder-decoder structure. This enhances the model's focus on relevant spatial features and is especially beneficial for distinguishing construction stages within cluttered satellite images. The model's prediction mask output is a grayscale Geo-TIF image file which consists of different grayscale colors for each pixel representing different categories of interest. Next a blob search process extracts clustered or connected pixels of the same classification, referred to as blobs. This provides a clear outcome of the property's stage of construction without the need for third party parcel data. Each blob represents one of the following: a pool, a roof, a particular stage of construction of a building, land, or vegetation. Figure 1 provides an example of a satellite image and the corresponding construction stage prediction mask.



Figure 1. Satellite image and construction stage prediction mask

5.1 Exclusions

To reduce noise in the blob search output, several filtering processes are applied to remove start blobs that are likely to be false positives. This includes overlaying start blobs with building footprint polygons available from Bing Maps and excluding start blobs that overlap with existing buildings. Additionally, start blobs are overlayed with known roads using the 2023 Tiger/Line Road Data to exclude starts that overlap with roads.

A second CNN model is applied to the satellite images to classify areas by building type (single-family detached, single-family attached, multi-family, and non-residential). When identifying building type among a large area, images usually contain too many features so a DeepLabV3+ with ResNet50 model architecture is employed. This combination proved effective at handling complex urban features by grouping different small features. Starts that fall into areas that are classified as multi-family or non-residential are excluded from the satellite estimates for single-family starts.

5.2 Business classification and additional exclusions

Specific business rules are applied to interpret construction progress by analyzing transitions between predictions from the previous and current months. These rules help determine the construction status for each property site based on the logic defined in Table 1.

For example, if a site was predicted to be in the LAND phase in the previous month and the current month shows EXCAVATION, the business class START is assigned, indicating that construction has progressed from land preparation to active development. Similarly, a transition from EXCAVATION to FOUNDATION would indicate further progress, thus the business class IN PROGRESS would be assigned.

The previous month's blobs are matched with the current month's blobs using geo-metadata and matching algorithms, accounting for any size differences between the blobs across months. Once the business rules are applied, predictions are aggregated across the defined place. This aggregation provides an estimate of the number of construction starts, within a specific area.

The model may predict some noise because of cloud cover, poor visibility, or limited sunlight during winter months. These cases can produce discrepancies in what should be observed. A post process step is taken to adjust these mis-matched areas accordingly. For example, if an area has been predicted to have a roof for 3 months, the likelihood of it becoming vegetation is highly unlikely. Therefore, we modify the result to be roof to roof which takes the MISCLASSIFIED BACKWARDS to an EXIST category.

Changing seasons are an additional source of noise in the model predictions. Houses with trees nearby may be classified as VEGETATION when leaves are present but classified as ROOF after the leaves fall. This results in START COMPLETION categorization, so named because the project moves from a preconstruction stage to a completed roof in a single month. The vast majority of blobs identified as START COMPLETION are noise. However, while the progression from a preconstruction stage (BACKGROUND, LAND, or VEGETATION) to a roof within a month is uncommon, it does occur, most often in areas with many construction projects occurring close together, such as a new neighborhood, and so cannot be completely ignored. Consequently, START COMPLETION blobs that lie within the same Geohash (see Section 5.3 below for more details on Geohashes) as a blob categorized as START are included in the starts estimate while all others are excluded.

Table 1. Business classification rules based on construction stage predictions.

	Current Month Prediction							
Previous Prediction	Background	Land	Vegetation	Excavation	Foundation	Framing	Unfinished Roof	Roof
Background	Unknown	Land	Land	Start	Start	Start	Start	Start Completion
Land	Land	Land	Land	Start	Start	Start	Start	Start Completion
Vegetation	Land	Land	Land	Start	Start	Start	Start	Start Completion
Excavation	Image Error	Land	Land	In Progress	In Progress	In Progress	In Progress	Completion
Foundation	Image Error	Misclassified Backwards	Misclassified Backwards	In Progress	In Progress	In Progress	In Progress	Completion
Framing	Image Error	Misclassified Backwards	Misclassified Backwards	In Progress	In Progress	In Progress	In Progress	Completion
Unfinished Roof	Image Error	Misclassified Backwards	Misclassified Backwards	In Progress	In Progress	In Progress	In Progress	Completion
Roof	Image Error	Misclassified Backwards	Misclassified Backwards	Start Reconstruction	Start Reconstruction	Start Reconstruction	Rehab	Exist

5.3 Imputation

As with other data collection methods, there can be missing data in the satellite imagery, most often due to cloud cover. The entire AOI could be missing, or more frequently, only part of an AOI is missing. Furthermore, there could be consecutive months of missing data for a given area.

New residential construction is not typically evenly distributed across an AOI, which could be as large as an entire city or county. It is typically localized to a small number of areas with many new homes being built close together. When a partial AOI is observed, rather than applying a missing data mechanism at the AOI level, we take a localized approach using Geohashes. Geohash is a public domain geocoding system to define geographic areas in a hierarchical spatial structure. The number of characters in the hash represents the size, a 6-character Geohash is larger in area compared to 7-character Geohash. For imputation, we divide an AOI into its component 7-character Geohashes (153m x 153m). At this more granular geographic level, both the assumption of evenly distribution construction activity in the current month and an assumed correlation in activity in consecutive months are more plausible than at the AOI level.

When we encounter missing data for an entire Geohash or part of a Geohash, we draw from the most recent observed data for that Geohash to develop an imputed value for the missing data. When less than 50% of a Geohash is observed in the current month, the Geohash's starts value is set to the maximum of the Geohash starts observed in the current month and the Geohash's starts from the prior (or last observed) month. When at least 50% of the Geohash is observed in the current month, the current month's starts are adjusted by dividing the observed start by the proportion of the Geohash that is observed. For example, if 75% of the Geohash was observed and 3 starts were identified in the observed area, the final starts value for the Geohash would be calculated as $3/0.75 = 4$ starts.

5.4 Capture window adjustment

Capture window adjustment is used to adjust the post-imputed counts for the difference in days between image dates. Due to the long image capture window, the number of days between observations could vary between 20 days to 42 days. We adjust the estimated starts for a given area by the ratio of the number of days in the target month to the number of days between observations.

For example, if the image for May was taken on June 1st and the image for June was taken on July 4th, there would be 33 days between images when the June estimates were calculated and there are only 30 days in June. Therefore, the June estimate would be multiplied by $30/33$ to correct for the length of time between observations.

6. Estimation

Currently, the SOC housing starts estimates are computed as the sum of the Survey of Use Permits (SUP) portion calculated from permit level data with a separate ratio estimator and the Non-Permit (NP) portion calculated as the weighted sum of observed new construction from sampled non-permit areas. See the "Compilation of Data" section of the [SOC technical documentation](#).

However, the estimates obtained from the satellite imagery are at a BPO level and cannot be used with our current estimator which aggregates data at the permit level. Therefore, the WSC single-family starts estimate for month t ($\hat{S}_{SOC,t}$) is a hybrid estimate combining the place level estimates obtained from the satellite imagery ($S_{SAT,t}$) with a modified version of the currently used separate

ratio estimator for the non-satellite SUP places ($\hat{S}_{SUP_NS,t}$) and the NP portion of the estimate ($\hat{S}_{NP,t}$) which remains unchanged.

$$\hat{S}_{SOC,t} = S_{SAT,t} + \hat{S}_{SUP_NS,t} \hat{S}_{NP,t}$$

The $\hat{S}_{SUP_NS,t}$ is the sum of 14 ratio estimates consisting of monthly estimates for the most recent twelve months and two additional ratio estimates that collapse months 13-18 into a single estimate and months 19-60 into the final estimate. Each ratio estimate is of the form $\hat{G}_{\tilde{m},t} \frac{Y_{\tilde{m}}}{\hat{H}_{\tilde{m}}}$, where $\hat{G}_{\tilde{m},t}$ is the SUP division estimate of units authorized in period \tilde{m} and started in month t , $\hat{H}_{\tilde{m}}$ is the SUP division estimate of units authorized in period \tilde{m} , and $Y_{\tilde{m}}$ is the Building Permits Survey (BPS) value of units authorized in the division in period \tilde{m} and \tilde{m} can be a single month or a collection of months.

The modification introduced to the current SUP estimator for this WSC division estimate is the addition of a PSU calibration factor so that the weighted value of sampled permits for a given month from the non-satellite places in each sampled PSU equals the total number of permits for that month from non-satellite places in that PSU from the BPS. See Appendix A for additional details about the estimation process.

7. Reliability of data

7.1 Sampling Error

Sampling error for the non-satellite component is estimated using the modified half-sample method currently used for SOC. The satellite portion of the estimate has no sampling error because all the selected satellite places are self-representing and satellite data collection is attempted for all in-scope areas for the selected places.

7.2 Nonsampling Error

Nonsampling error includes all sources of error, other than sampling error, that contribute to the total error of an estimate. Both the satellite and non-satellite portion of the WSC single-family housing starts estimate are subject to nonsampling errors. These include coverage errors, collection errors, measurement errors, imputation errors, errors due to nonresponse, and other processing errors. Although nonsampling error is not directly measured, quality control procedures are applied throughout the satellite and non-satellite data processes to minimize nonsampling errors.

7.3 Nonresponse

The current nonresponse adjustment method for SOC housing starts accounts for late reported starts that are expected to be received up to twelve months late by multiplying the SUP housing starts estimate by a nonresponse undercoverage adjustment factor (NUAF) computed at the region level. The WSC single-family housing starts estimate uses the same methodology but applied at the division level for the non-satellite component of the estimate, using only non-satellite sampled places in the computation of the NUAFs.

8. Disclosure Avoidance

The current SOC disclosure avoidance method of cell suppression will continue to be used, if necessary, for the WSC single-family starts estimates. The Census Bureau has reviewed this product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data (Disclosure Review Board (DRB) approval number: CBDRB-FY25-0086).

Appendices

A. Computer vision models

A.1 General structure

The **encoder-decoder architecture** is highly beneficial for tasks like semantic segmentation because it effectively balances feature extraction and spatial resolution restoration. In this architecture, the **encoder** focuses on extracting high-level, abstract features from the input image, such as shapes, patterns, and semantic information, by progressively downsampling the spatial dimensions and applying deep convolutional layers. Then the **decoder** restores the spatial details, ensuring fine-grained predictions by upsampling the features back to the original resolution. This combination allows the model to capture global context and preserve local details simultaneously.

One of the key advantages of this architecture is its ability to learn features at multiple scales. The encoder compresses the image, enabling the model to capture large receptive fields and extract global information, while the decoder combines this with local features to produce detailed segmentation outputs. This multi-scale learning is particularly useful for handling objects of varying sizes within an image. Additionally, skip connections or feature concatenation between the encoder and decoder facilitate the retention of fine-grained spatial details, resulting in sharper boundaries and more accurate predictions.

The encoder-decoder structure is also computationally efficient. By reducing the spatial dimensions in the encoder stage, the model can focus on learning complex semantic patterns while requiring less computational power. The decoder then uses this compact representation to reconstruct meaningful spatial information. Moreover, the flexibility of this architecture allows for easy adaptation to different tasks, such as semantic segmentation, medical image analysis, or instance segmentation, by integrating various backbone networks (e.g., ResNet or MobileNet) as encoders, depending on the task requirements and computational constraints.

Another significant benefit is its robustness to noise and small datasets. The bottleneck structure in the encoder forces the model to learn essential features, which helps reduce overfitting on noisy data. Furthermore, pretrained encoders can be leveraged to transfer knowledge from large datasets, enabling efficient training even with limited labeled data. This architecture can also incorporate attention mechanisms, enhancing its ability to focus on the most important regions of the image, thus improving segmentation quality.

Overall, the encoder-decoder architecture provides an elegant and effective solution for semantic segmentation by combining global context, local detail preservation, computational efficiency, and flexibility, making it a popular choice for dense prediction tasks.

A.2 Construction Stage Model

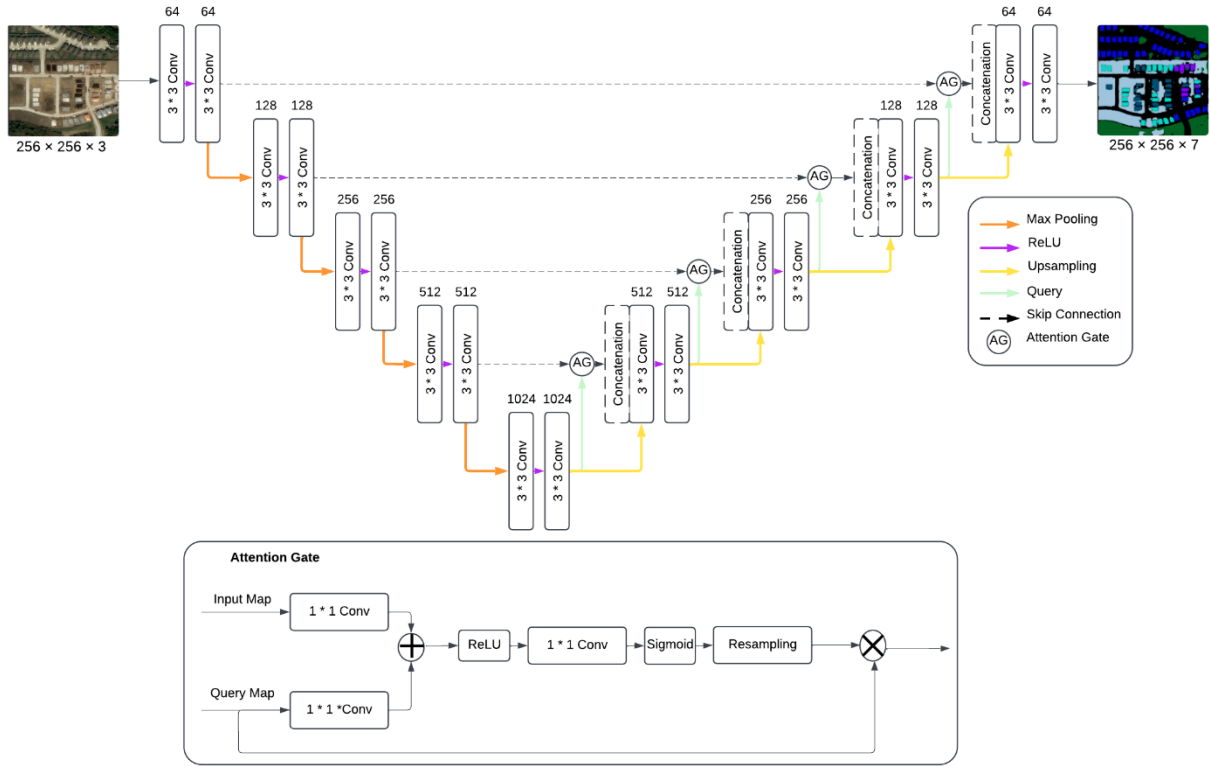


Figure 2. Construction stage model using Attention U-Net architecture.

This model diagram represents an **Attention U-Net architecture** designed for semantic segmentation tasks. The architecture takes an input image of dimensions 256x256x3 and processes it through an encoder-decoder structure, enhanced by attention gates to focus on relevant regions. The encoder consists of convolutional layers (3x3 convolutions with ReLU activation) and max-pooling layers (2x2) to progressively downsample the input and extract hierarchical features. Skip connections between the encoder and decoder preserve spatial details by directly transferring features from the encoder to corresponding layers in the decoder.

The decoder upsamples the feature maps using transpose convolution (2x2) and refines them with convolutional layers, progressively reconstructing the segmentation map to match the input resolution. Attention gates (AG) are introduced along the skip connections to emphasize the most relevant features while suppressing background noise. These gates dynamically learn to focus on key regions by multiplying incoming feature maps with attention weights computed based on contextual importance. At the output, a 1x1 convolution layer maps the refined features into a segmentation mask of 256x256x7, where each pixel corresponds to one of the seven output classes. The use of attention gates ensures precise segmentation by enhancing the model's focus on critical regions while reducing interference from irrelevant parts of the image, making this architecture particularly effective for tasks requiring fine-grained and accurate pixel-wise predictions. For additional details on Attention U-Net see Jing (2019).

A.3 Building Type Model

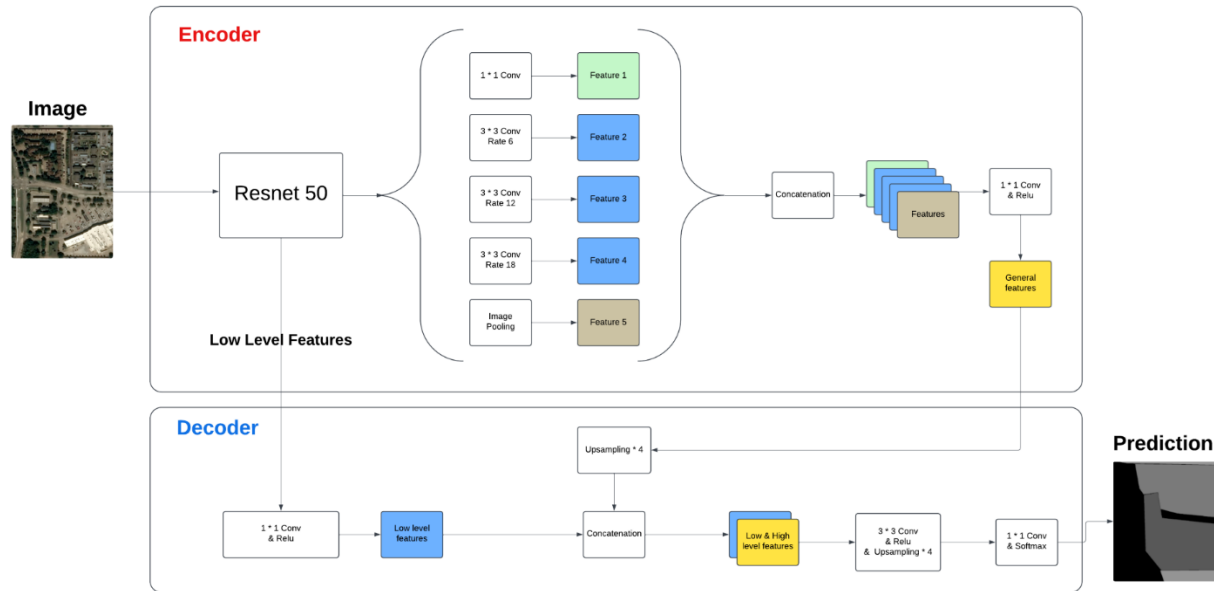


Figure 3. Building type model using a semantic segmentation architecture.

The model diagram represents a semantic segmentation architecture designed to process an input image, a satellite image in our case, through an encoder-decoder framework. The process begins with the input image being fed into the encoder, which uses a ResNet50 backbone for feature extraction. The encoder captures hierarchical feature representations, progressing from low-level details like edges and textures to high-level semantic features. To enhance multi-scale feature learning, the architecture incorporates an Atrous Spatial Pyramid Pooling (ASPP) module within the encoder. This module consists of 1x1 convolutions for point-wise feature extraction, 3x3 convolutions with varying dilation rates (6, 12, and 18) to capture features at different spatial scales, and global image pooling to encode context across the entire image. The outputs of these operations are concatenated and passed through a 1x1 convolution to ensure compact and meaningful feature representations.

The decoder complements the encoder by restoring spatial resolution and refining predictions. Low-level features from earlier layers of the encoder are integrated into the decoder through skip connections, ensuring that fine-grained spatial details are preserved. These low-level features are first processed using a 1x1 convolution to reduce their channel dimensions, making them compatible with the high-level features from the ASPP module. The decoder upsamples the high-level features by a factor of four, combining them with the low-level features through concatenation. Additional refinement is achieved through a series of 3x3 convolutions and further upsampling, which ensures that the output segmentation mask matches the resolution of the original input image.

Finally, the model outputs a pixel-wise segmentation mask, which predicts the class of each pixel in the input image, such as building, road, or vegetation. This architecture combines global and local context, enabling accurate and high-resolution predictions. Its design leverages the strengths of both the encoder for abstract feature extraction and the decoder for spatial reconstruction, making it well-suited for complex semantic segmentation tasks. For additional details on semantic segmentation see Kou et al. (2022).

B. Estimation details

PSU calibration factor

Divide the SUP sample of places (BPOs) into those selected to be collected by satellite (Q_S) and those that will not (Q_N). All calculations using SOC survey data detailed below only use records coming from places in Q_N .

Let w_j be the final weight for permit j (product of CPS sampling, SOC PSU sampling, place sampling, and permit sampling weights) and w_k^{PSU} be the sampling weight for PSU k (product of CPS sampling, SOC PSU sampling). Thus, the permit-place sampling weight for permit j from place p in PSU k is $w_j^{permit-place} = w_j / w_k^{PSU}$.

Define the number of building permits authorized in month m from place p recorded in BPS as $y_{p,m}$ and let $F_p = 0$ if place p is a satellite collected place and 1 otherwise. Also, let $A_{j,m} = 1$ if permit j was authorized in month m and 0 otherwise. The PSU calibration factor for PSU k in month m ($w_{PSU_k,m}^{cal}$) is the ratio of the BPS value of permits authorized in month m for non-satellite places in PSU k ($Y_{PSU_k,m} = \sum_{p \in PSU_k} F_p y_{p,m}$) to the SOC estimate of permits authorized in month m for non-satellite places in PSU k ($H_{PSU_k,m} = \sum_{j \in Q_N \cap PSU_k} A_{j,m} w_j^{permit-place}$)

$$w_{PSU_k,m}^{cal} = \frac{Y_{PSU_k,m}}{H_{PSU_k,m}}$$

The final weight, w_j , for permit j belonging to PSU k , authorized in month m is multiplied by the corresponding PSU calibration factor to define the calibrated final weight (w_j^*) that will be used for the non-satellite portion of the estimate.

$$w_j^* = w_j w_{PSU_k,m}^{cal}$$

Estimator

The non-satellite estimate of housing starts uses the separate ratio estimator currently used for SUP estimation. It is the sum of 14 ratio estimates consisting of monthly estimates for the most recent twelve months and two additional ratio estimates that collapse months 13-18 into a single estimate and months 19-60 into the final estimate. Each ratio estimate is of the form $\hat{G}_{\tilde{m},t} \frac{Y_{\tilde{m}}}{\hat{H}_{\tilde{m}}}$, where $\hat{G}_{\tilde{m},t}$ is the SUP division estimate of units authorized in period \tilde{m} and started in month t , $\hat{H}_{\tilde{m}}$ is the SUP division estimate of units authorized in period \tilde{m} , and $Y_{\tilde{m}}$ is the BPS value of units authorized in the division in period \tilde{m} and \tilde{m} can be a single month or a collection of months.

Define $I_{j,t} = 1$ if permit j was a start in month t , and 0 otherwise. Then the non-satellite SUP estimate of units authorized in month m and started in month t is

$$\hat{G}_{m,t} = \sum_{j \in Q_N} w_j^* A_{j,m} I_{j,t}$$

The non-satellite SUP estimate of the total number of permits authorized in month m is

$$\hat{H}_m = \sum_{j \in Q_N} w_j^* A_{j,m}$$

Since all satellite collected places are in the BPS, we can remove their data from the BPS total. Currently, all satellite collected places are self-representing places from self-representing PSUs, so their SOC place level sampling weights are one. Therefore, we can subtract their unweighted building permits for each month, m , from the BPS total number of permits for that month (Y_m). Define the adjusted BPS total as

$$\tilde{Y}_m = Y_m - \sum_p (1 - F_p) y_{p,m}$$

The total starts estimate for the satellite collected places (all with sampling weight=1) is simply the sum of the place level estimates obtained from the satellite imagery for month t .

$$S_{SAT,t} = \sum_{p \in Q_S} S_{SAT,p,t}$$

The estimator for total single family starts for the division is given below.

$$\hat{S}_{SOC,t} = \hat{S}_{SUP,t} + \hat{S}_{NP,t} + S_{SAT,t}$$

$$\hat{S}_{SUP,t} = \left[\sum_{m=1}^{12} \hat{G}_{m,t} \frac{\tilde{Y}_m}{\hat{H}_m} \right] + \left[\sum_{m=13}^{18} \hat{G}_{m,t} \left[\frac{\sum_{m=13}^{18} \tilde{Y}_m}{\sum_{m=13}^{18} \hat{H}_m} \right] \right] + \left[\sum_{m=19}^{60} \hat{G}_{m,t} \left[\frac{\sum_{m=19}^{60} \tilde{Y}_m}{\sum_{m=19}^{60} \hat{H}_m} \right] \right]$$

where $\hat{S}_{NP,t} = \sum_j w_j I_{j,t}$

References:

Jing, Hong. (2019, December 8). Biomedical Image Segmentation – Attention U-Net. *Jinglescode*. <https://jinglescode.github.io/2019/12/08/biomedical-image-segmentation-u-net-attention/>

Kou, Lei & Sysyn, Mykola & Fischer, Szabolcs & Liu, Jianxing & Nabochenko, Olga. (2022). Optical Rail Surface Crack Detection Method Based on Semantic Segmentation Replacement for Magnetic Particle Inspection. *Sensors*. 22. 8214. <https://doi.org/10.3390/s22218214>.