Quantarium

Residential Valuation Model

Technical Overview

Residential Valuation Models

It is easy to make an argument that property valuation plays a critical role in real estate. An accurate valuation drives decisions when buying or selling; loan originators limit maximum loan amounts based on the expected value of the property; loan servicing regulations control the course of action for underperforming loans depending on the current loan to value ratio; property taxes are based on assessed value of real estate assets; and financial institutions evaluate real estate portfolios on the projected value of the underlying assets. It should therefore not be surprising that automated property valuation models (AVMs) have been researched, built, and used since the 1980's.

Quantarium is an Artificial Intelligence software company founded by a team of veteran computer scientists that have developed a suite of advanced online software services, including an automated, commercial-grade residential valuation service. The Quantarium AVM (QVM) is built on a novel set of disruptive technologies that combine "big data" with complex modeling and optimization techniques. At its core, QVM provides an online engine which automatically computes unbiased, realistic, values for residential properties.

This paper describes the architecture of the QVM valuation stack offered by Quantarium.

QVM inputs and outputs

Inputs

In building the mathematical model used for a valuation as well as the collateral information associated with it, Quantarium leverages multiple data sources, including licensed public records and licensed Multiple Listing Services (MLS) data[1]. The data is joined, aggregated, and validated for consistency and correctness. Public records, covering more than 100 million U.S. residential properties, includes assessment, deed, mortgage, pre-foreclosure, assignment, and release information. These are provided by top tier national data providers and are updated daily. Quantarium evaluates this set every few days and produces valuations for more than 98% of the properties contained therein. Individual valuations can be run on demand against specific properties using the most recent information available.

Customers can obtain QVM valuations in several ways. For automated queries on individual properties, Quantarium exposes web service interfaces that can be called on demand. QVM values for individual properties can also be embedded in more sophisticated products and reports which will be described later in this document. For large scale, portfolio wide valuations, batch schedule processing is available via Match-and-Append jobs that use customer provided input files.

Quantarium provides several mechanisms for making queries. Customers most commonly query by site addresses. Quantarium has developed techniques to facilitate matching low-quality addresses or

[1] The subject MLS data is controlled by Xome Inc., and made available under the license of and on behalf of Xome and its local broker of record

addresses with incomplete or erroneous information, such as apartment addresses that are missing unit numbers. Additional query fields like property owner name are used to disambiguate queries that return multiple records. Through this method Quantarium has achieved an approximate match rate of 97% in independent 3rd party tests. Other forms of identifications are also possible, for instance, assessor parcel numbers (APNs).

Outputs

QVM outputs include the following:

- **AVM value** – The automated valuation computed by the model for the subject property.
- **AVM Low/AVM High** – A confidence interval for the AVM value.
- **FSD Score** – Forecast Standard Deviation is a statistical measure that represents the probability that the Automated Valuation Model (AVM) value falls within a statistical range of the actual market value, measured against a sales price. For more details refer to https://sf.freddiemac.com/tools-learning/home-value-suite/forecast-standard-deviation
- **Confidence Score** – Quantarium correlates the percentile FSD score with a confidence score between 0 and 100. The higher the confidence score, the more accurate the AVM value is expected to be.

Quantarium valuations use a proprietary correlation model for defining the confidence interval, FSD score, and confidence score. One of the simple models that is being used, which is provided here as an intuitive explanation of such a correlation is the following: The interval between AVM Low and AVM High has a maximal value that is twice the FSD and the confidence score is the 100% complement of the FSD. In mathematical terms, assuming the FSD is defined as a percentile:

$$\text{FSD} = (\text{AVM High} - \text{AVM Low}) / (2 * \text{AVM Value})$$
$$\text{Confidence Score} = 100 * (1 - \text{FSD})$$

Depending on the use case scenario, QVM provides additional valuation outputs. For example:

- For portfolio valuations the QVM provides specific target confidence intervals requested by the customer. For example, a *p90 AVM Low* would be the low property value threshold for 90% of the valuations.
- For more sophisticated, integrated product use cases, the QVM can produce valuation *evidence*. The valuation evidence depends on the underlying modelling technique used. As an example, for a market growth valuation model, the evidence could be a prior sale and the estimated market indices used to estimate growth.

Markets and Models

The QVM uses multiple models that cascade in two dimensions, based on geography and model technique. In other words, multiple valuations techniques and learned corresponding valuation parameters are used to produce a final valuation. Quantarium trains and validates the QVM using sale prices as a proxy for true value. Sale values are validated for correctness using a combination of statistical outlier detection and domain specific heuristics (e.g. sale values that deviate from the corresponding loan amounts are deemed unreliable).

Individual models are trained using data from local geographic area, like zip or county. This improves the quality of the valuation in a couple of ways beyond standard ensemble modelling techniques. For example, preference is given to lower granularity models that have shown better historical performance, and, in addition, having multiple valuations enables the QVM to detect outlier conditions where multiple models significantly disagree.

QVM architecture

The QVM uses multiple models to estimate property value, each optimized at several levels of geographic resolution. An abstracted view of the QVM pipeline is presented in the figure below.

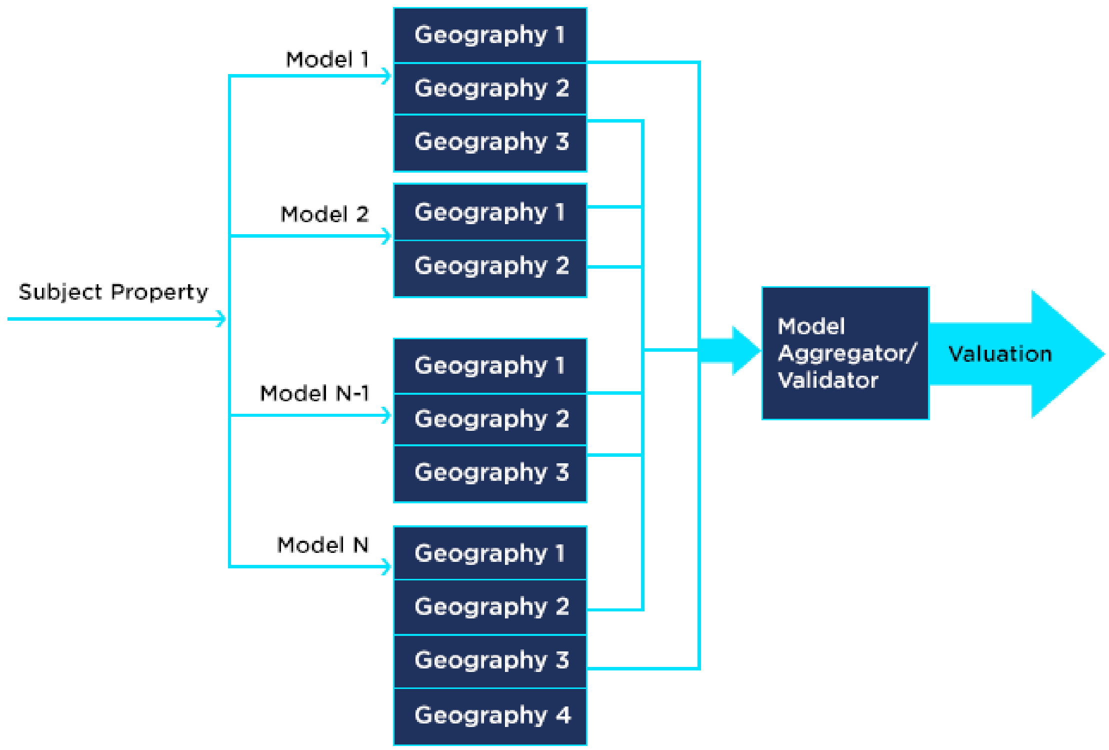

The QVM uses the following geographic resolutions or “markets”: zip, county, Metropolitan Statistical Area (MSA), state, national. Empirical evidence shows that a lower granularity model produces more accurate estimates, therefore zip level models are preferred. For this reason, most of the property valuations are computed at the zip level. However, Practical limitations exist since data becomes too sparse as the granularity increases and because lower granularity models may produce estimates that can later fail validation. In these cases, broader geographies will be used to increase coverage.

The QVM uses three classes of models which are based on (1) comparable sales, (2) market appreciation, and (3) the relation between market and property characteristics. At any point in time, multiple instances of these models can be used in the final cascade. This covers not only instances where the same model has multiple instances that are tied to the geography cascading architecture from above, but also models that instantiate with completely different trained parameters. The latter is useful in providing customers with the right tradeoff between coverage and accuracy. Quantarium provides a default selection of models, but in specific use cases per-customer configuration can be used to change that. For example, customers that optimize towards accuracy may benefit from selecting only the models that are trained aggressively towards low valuation errors, while customers that prefer higher coverage may choose a wider range of acceptable models.

Comparable valuation model

Quantarium comparable valuation model works in three steps, each of which is trained and optimized within a local market, using advanced optimization and statistical techniques, such as genetic algorithms. A local market can be a subset of a geographical market that splits a given geography further based on other attributes. Such attributes include, for example, value ranges and property types. First, it identifies similar properties (comparables) from a pool of recent sales. The criteria for what constitutes a comparable property is optimized and considers the differences between property characteristic, geographical distance and time delta from the comparable sale. Second, it adjusts the comparable sale price, using parameters learned during the training phase and the deltas between the comparable and subject properties. This produces an adjusted sale price for the subject property. Finally, the model ranks and filters outliers from the set of comparables to calculate similarity weights, which are then used to calculate the final valuation as a weighted average of the adjusted sale prices. For this model, the evidence consists of the selected comparables and their adjusted prices. Quantarium retrains and adjusts the parameters used in each of the three phases. The pool of comparables used for training and valuation consists of hundreds to tens of thousands. The final comparables set is selected, depending on location and time, from the most trusted sale prices available.

Market appreciation model

A market appreciation model works by projecting a prior sale for the property into the present using a market index model. The accuracy of such models depends on how well a prior sale approximates the context of the current sale and how close the appreciation of the subject property mirrors the index appreciation. Quantarium has developed techniques that minimize the error introduced by these factors. For example, a sale that occurs prior to a major property change will result in a high valuation error and will thus be excluded. Similarly, a very recent sale that is not an “arms-length” transaction will be excluded as well since the confidence it reflects in the true value of the property is low. These models use multiple types of real estate indices and continuously monitors the accuracy of the estimates. The evidence for this type of models consists of the prior identified sales and the market indices used.

Property characteristics model

These types of models learn the complex relations between many property characteristics and the market value. Training is performed using a significant number of samples within a given geography using sufficient temporal data to account for seasonal variation. The models use standard machine learning techniques, such as linear regression, ensemble models like regression trees, or more complex models like neural networks. Evidence for such models depends on the technique used. For example, a linear regression model may provide a price per square footage as a linear weight learned during training.

Model aggregation and validation

The QVM uses a proprietary trained AI engine to provide a final estimate from the set of multiple cascaded estimates, along with a corresponding FSD score. In principle, the aggregation performs two types of rules: (1) filtering rules which discard cascaded valuations under certain conditions, and (2) aggregation rules which combine cascaded valuations into a single new aggregated valuation. The main idea behind aggregation is that it provides a standard variance. The filtering rules and especially their criteria are complex and trained empirically. To provide an intuitive example of how this works, one such filtering rule can be something like (numbers are used for illustrative purposes): "If a comparable valuation with at least 10 comparables exist in zip code 90210, then discard all market appreciation models that use sales older than 12 months".

In addition to the input and training data validation that occurs per model, the QVM validates the final estimate against property information. Like filtering rules, validation criteria are trained empirically as well. For an intuitive example, one validation rule might be something like (numbers are used for illustrative purposes): "If the valuation is 50% larger than a sale that occurred within 1 year, then the valuation is dropped". The QVM also saves the existence of prior known bad inputs and valuations to avoid future regression.

QVM testing

Quantarium tests QVM outcomes internally both on a continuous basis as a prerequisite validation prior to public release, as well as part of model changes that occur occasionally in response to regular retraining or conceptual model additions or updates. Further, Quantarium saves results from past nationwide runs, referred as historical valuations, to validate against recent sales. The last 90 days of nationwide sales transactions serve as the ground truth for historical valuations computed from models that have not used them in training. This provides a statistical history of model performance.

In addition, a subset of QVM models participate in independent, 3rd party, testing. On a weekly basis, Quantarium receives pools of tens of thousands of properties for which valuations are expected within 48 hours. Each quarter's error information is calculated and aggregated. Specific information and results are available upon request. We recommend that customers interested in valuations for specific subsets of properties undergo their own valuation to avoid introducing any selection bias. Quantarium may be able to recommend better models to use. For one such example, Quantarium provides a specific QVM model for distressed properties.

FAQ (other considerations)

Often, customers have specific questions which were not addressed in the overview provided thus far. This FAQ tries to provide answers for the most commonly asked ones.

Q: What is the geographic coverage for QVM?

A: The model provides valuations for all 50 US states. It does not provide estimates for US territories such as Puerto Rico and Virgin Islands.

Q: Does QVM provide evidence for all models?

A: As a general principle, QVM model selection prefers simpler methods over more complicated methods. This tends to result in valuations that can be understood by humans. However, not all models lend themselves to easily understood explanations. Also, the kind of evidence/explanation available is usually delivered in more sophisticated product offerings, built on top of the QVM.

Q: Does the QVM use information outside public records?

A: Indeed, property information is reconciled from multiple data sources. For instance, public records can be augmented with MLS data for cases where the data is not available, such as sale information for non-disclosure states.

Q: Are QVM models different per property type, like single family residence, or condo?

A: The models generally use property types either as a selection parameter or as a discriminating class.

Q: Are there any internal restrictions to model output, i.e. in what situations would the model not generate an estimate?

A: The QVM may not generate an estimated value when it cannot get sufficient data as input. This can happen when the QVM:

- cannot identify the site address
- fails to gather sufficient data as required by model parameters
- identifies suspect data about the property (for instance values outside of normal ranges for square footage).
- In addition, even if the mathematical model can produce one it may still not issue one when the model:
- finds a confidence score below a certain threshold.
- fails the validation methodology verification on the final evaluation

Q: How does the QVM model deal with conflicting property characteristics?

A: The QVM works on an individual property "view" that is aggregated from multiple sources of data with the goal of having the most accurate representation for property intrinsic characteristics. For cases where a given characteristic can be obtained or inferred from multiple data sources, the model is tuned to select the source that results in the most accurate evaluation. For instance, we found that public records are more reliable data sources for sales transactions. In addition, the existence of multiple sources of data for the same characteristic helps identify inconsistent values. For example, if multiple sources show different square footage values which differ significantly for a given property the model may not render an evaluation for that property.

QVM extended use

In the last decade, simple AVM valuations have become common place, table stakes, so to speak. Quantarium provides integrated, extended uses for the QVM that go beyond a value and the associated confidence. But to discuss those, it is worth considering the major advantages and benefits of the automated valuation technologies as well as their disadvantages.

First, automated valuations tend to be cost efficient and fast. In fact, many companies provide consumer-grade valuations as part of their service, at zero perceived cost to the customer. Even lower accuracy valuations can be used for fast decision points as part of a business pipeline that requires commercial grade property valuations, because they are fast and affordable. Secondly, automated valuations benefit from an intrinsic, fundamental strength: the ability to repeat the same algorithm without operational errors. Once an address has been validated one doesn't have to worry about human errors in transcribing the valuation data or about inability to retrieve that valuation later because it's after hours. However, as with any computer driven technology, AVMs suffer from the "garbage in – garbage out" syndrome. While significant errors, such as extra "0's" inadvertently reported on a sale sheet can be caught, subtle errors induced by lack of context and information can still occur. This is one of the reasons that Quantarium provides solutions that leverage automated valuation in the broader context of solving end-to-end business problems through information pipelines.

Here are some examples of how QVM valuations can be used in this manner:

1. Property facts – With QVM values, the associated valuation evidence as well as property and location information can be used to provide a comprehensive view for a property during any purchase or sale transactions.
2. Appraiser assistance – With QVM values, the associated valuation evidence, property and location information can be used as a guide for human appraisers when the situation requires it. As an example, loan originators can validate and enforce specific business criteria required when appraising properties.
3. Home Value reconciliation – The QVM can be used to detect conditions where the home value matches the expected value and, when it doesn't, it helps a human-driven reconciliation using the valuation evidence provided by the QVM.
4. Home value indices – Quantarium uses home valuations to build its own real estate valuation indices at zip, county, MSA, and state market levels. These home valuations can be used for historical analysis as well as for forecasting economic conditions.
5. Historical valuations – Quantarium models are built so that they are agnostic of absolute time. These models can be used to produce a valuation for a prior time, considering the market and the property characteristics at that moment. This is useful for complex business scenarios which involve real time testing of alternative strategies.
6. Special-circumstance valuations – as previously mentioned, models can be adjusted to provide distressed property valuations or portfolio level confidence via additional statistical metrics like high confidence for minimal values.

Summary

Quantarium began development of the Quantarium Valuation Model (QVM) in 2008 using an alternative Genetic Algorithm optimization technique, which was different in kind from existing automated valuation models available in the marketplace. Since its commercial availability in 2014, the initial model has been augmented with additional models in a cascading fashion. This approach enables careful balancing of multiple contributing factors like intrinsic value of the property, local geo-temporal context, and general housing market parameters, for a best-in-class valuation. Quantarium prides itself in providing the right tool for the job and specializes in building products which extend beyond a simple valuation to meet end to end customer needs. Following the same philosophy, the underlying QVM models can be configured for various operating points to achieve desired coverage and accuracy for specific customer scenarios.

For more information visit www.quantarium.com or contact info@quantarium.com.

Appendix A – Brief primer on AVM statistical measures

AVM testing, whether independent or internal, typically measures and evaluates the performance of models using a methodology that involves providing a set of test inputs for which the true valuation. Also referred as the ground truth, it usually consists of very recent sale prices or highly trusted appraised values. The tested model can either decline to provide a result or return a property valuation. The valuation must include at least the predicted value for the property, but, for more in-depth testing methodologies, can also require confidence intervals, FSD values and any reasons for not providing a valuation that may have been encountered.

Typical quality measures that can be computed from such tests include:

Coverage – the ratio between the total number of valuations returned by the model and the total number of valuations requested. This metric can be further "decomposed" in coverage ratios for various reasons for which valuations cannot be produced. A very typical example can be to discriminate between "cannot find the property" versus "cannot produce a valuation that meets desired quality".

Error – the signed difference between the model valuation and ground truth.

Absolute Error – the absolute value, also known as unsigned difference, between the model valuation and ground truth.

Relative Error – the ratio between error and ground truth. Relative error enables apples-to-apples comparison and use of aggregate statistical metrics across the entire range of ground truth values.

Absolute Relative Error – the absolute value of the relative error.

Aggregation of error measures occurs over the set of test samples that are evaluated by the model typically using mean (also referred to as average), median or percentile statistics. The most commonly found include:

Mean Error – average error looks for global behavior of models over the test set. It is usually used when assessing portfolio performance and is ideally close to 0. The average error will not, however, identify large individual errors that are spread over both negative and positive values. It will also not offer information about the relative magnitude of the error.

Mean Relative Error – the use of relative values offers a normalized metric for error magnitude. While the use of relative error provides normalization, care should be used when aggregating over wide ranges of ground truth values. Specifically, an error of $2,000 for a property valued at $50,000 represents a 4% error, while the same error for a property of $2,000,000 is only 0.1%. For this reason, mean errors over price ranges can be provided in addition to the global measure.

The graph in *Figure 1* below shows an example of the mean relative error displayed in red over subject property price ranges for a test portfolio that has been evaluated by QVM. The blue bars represent the associated record counts for the respective price ranges.

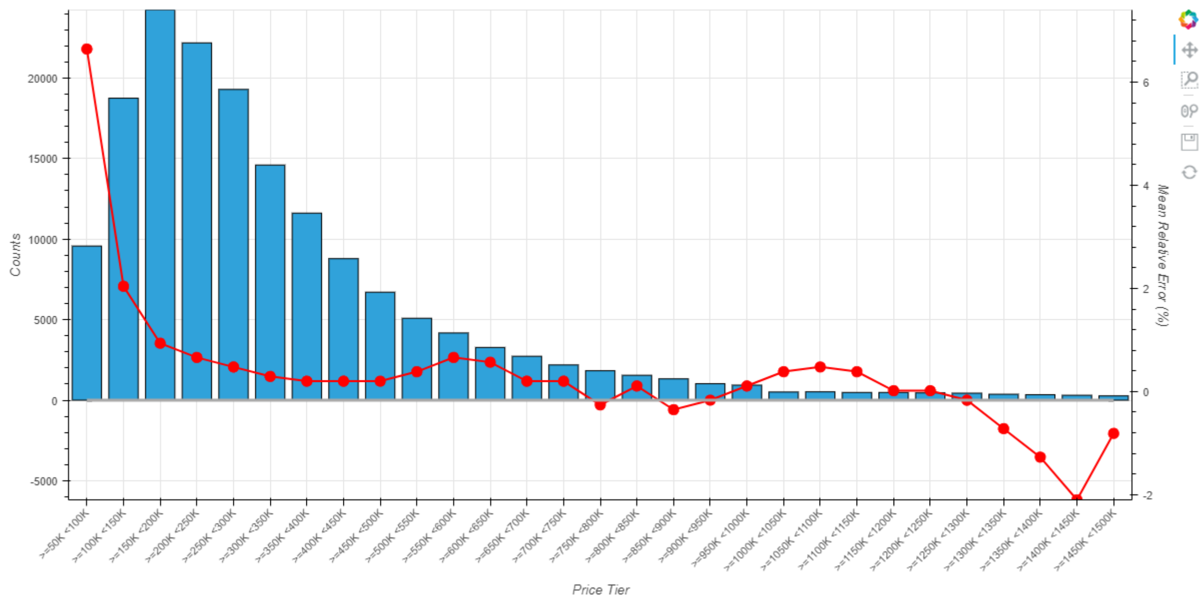

Figure 1

Mean Absolute Relative Error – the use of signed errors could hide large errors that occur symmetrically around zero. Using mean over absolute values is helpful in situations where assessing the magnitude of the error is more important than its sign.

Median Absolute Relative Error – the mean value tends to be sensitive to extreme values (outliers). Using the median instead of the mean, provides a more stable measure of error – in other words, making a few large errors will not materially change the value of this metric.

The graph in *Figure 2* below shows an example of the median absolute relative error displayed in red over subject property price ranges for the same test portfolio as in *Figure 1*.

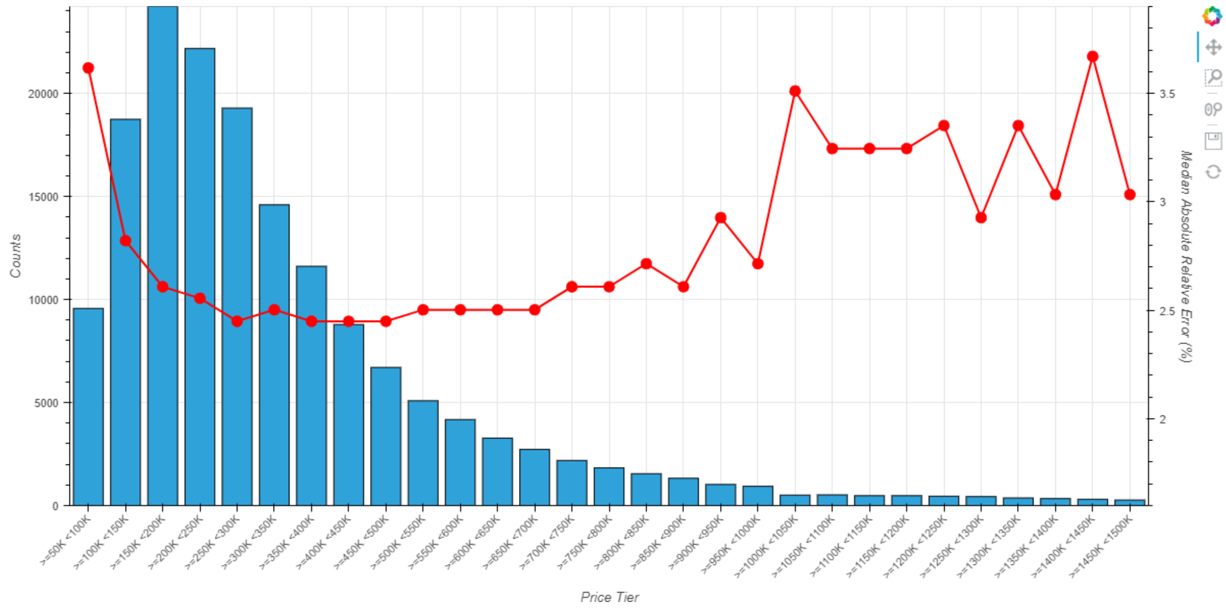

Figure 2

Absolute Relative Error Percentiles – percentile errors generalize the median error. As median absolute error refers to the absolute error value for which half of the errors are less than that, the Nth percentile represents the value for which N% of the values are less than that. Median is same as P50. Percentiles are used either in the higher ranges for more "stringent" measures or in the lower range for more "lax" ones. One would use P90 absolute relative error to get a guarantee that at most 10% of the values are above that error value, while one would use P10 to get a guarantee that at least 10% are below that error value.

Accuracy rate – Accuracy rates measure the percentage of valuations that have absolute errors under or over a certain target percentage. For instance, AR5 measures the percentage of valuations that have under 5% absolute error. Similarly, AR50+ measures the percentage of valuations that have more than 50% error rate. They are used in scenarios where it's less important to know the median absolute relative error, but rather how many valuations are expected to error within an acceptable threshold from the lower or higher end range. From a statistical point of view if median absolute error is 8%, it means that 50% of the values are below 8% absolute error, and therefore AR8 would be 50%. Similarly, if P90 absolute relative error is 20%, then AR20 would be 90% and AR20+ would be 10%.

The graph in figure 3 shows an alternate graphical representation, where QVM AR 20%, at state level, is tiled over the US map for a test portfolio.

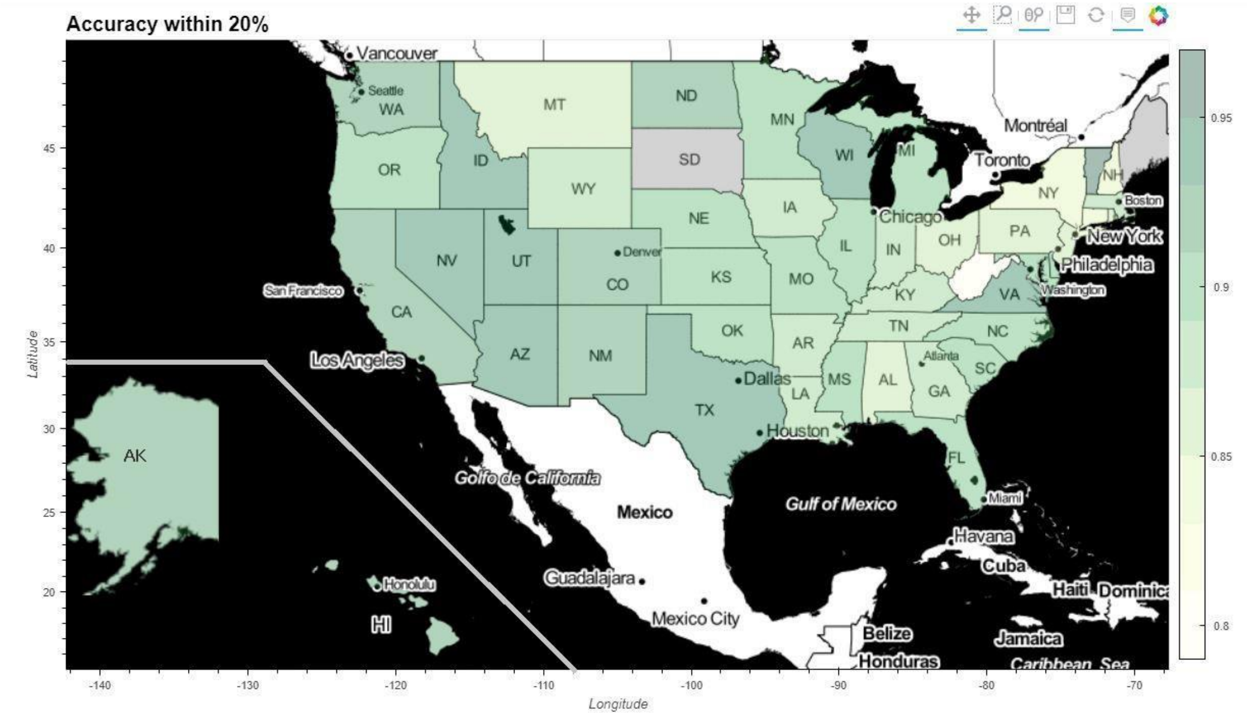

Figure 3

Confidence interval accuracy – this is a common metric used to measure confidence interval quality when the valuation model produces one. The metric is the ratio of how many times the ground truth is within the confidence interval over the total number of valuations.