Gadsden Meridian — Technical Documentation

UK Residential Automated Valuation Model

Version	2.1
Issued	14 June 2026
Model version described	Gadsden Meridian v6.12 (recalibration release of 14 June 2026 on the v6.11 prediction model)
Issuer	Gadsden Valuations Ltd (Company No. 17079555)

This is the technical documentation for underwriters, risk managers, and model-risk teams assessing Gadsden Valuations as an AVM supplier. Each section states what the model does, what the evidence is, and where the live, dated evidence can be checked. Every accuracy figure carries its test cohort and model version, and is generated from the same pipeline as the public accuracy page. This document is reissued, with a new version number and date, when a model or testing change occurs.

1. Model Overview

The Gadsden Meridian AVM produces point estimates of residential property market value for properties in England and Wales, consistent with the RICS definition of Market Value (IVS 104).

For each property, the model returns:

- A point estimate of market value (£) with upper and lower bounds derived from the per-property FSD
- A per-property Forecast Standard Deviation (FSD) and its corresponding Fitch classification band (A–D) — see Section 5
- Per-property feature contributions (SHAP), expressed as pound-value influences on the estimate
- The comparable transactions used as market evidence
- A full audit-trail record (Section 6), downloadable PDF report, and API response (Annex A)

Where the model cannot produce a meaningful estimate, it declines to produce one rather than returning a low-quality value (Section 5.3). Declined requests are not charged.

1.1 Intended Use

The AVM provides model outputs for use within the lender's own valuation and risk processes — as an input reviewed under the lender's policies, alongside or in place of other valuation instruments where the lender's policy permits. It is a valuation model in the IVS 105 sense.

Typical use cases: portfolio revaluation and monitoring, remortgage valuations where physical inspection is not proportionate, an independent cross-check against physical valuations or competing AVM outputs, and desktop-review support.

1.2 Coverage and Exclusions

England and Wales. Four standard residential property types: detached, semi-detached, terraced, and flats/maisonettes. The model does not cover Scotland, Northern Ireland, commercial property, new-build prior to first sale, or non-standard constructions.

Evidence: gadsdenvaluations.com/methodology, /no-value-policy.

---

2. Architecture and Data

2.1 Model

Gadsden Meridian v6.12 is a gradient-boosted decision tree ensemble (LightGBM) using **49 features** across six categories. v6.12 is a recalibration release on the v6.11 prediction model — the LightGBM weights are unchanged from v6.11; what was rebuilt is the FSD/band calibration (Section 5).

CATEGORY	DESCRIPTION	EXAMPLE INPUTS
Spatial	Location, coordinates, geographic context	UPRN-resolved coordinates, ONS geographies, comparable pricing from similar sales
Physical	Property characteristics and construction	Floor area, type, age band, tenure, bedrooms, EPC-derived attributes
Temporal	Transaction timing, market cycle position	Sale date features, prior-sale price adjusted by house price index
Neighbourhood	Area demographics, deprivation, and economics	Census 2021 by LSOA, Indices of Multiple Deprivation, recorded crime rates
Environmental	Flood risk, infrastructure and transport proximity	Environment Agency flood zones, rail station distance
Educational	School quality and proximity	GIAS register, Ofsted ratings

The v6.11 prediction model's principal additions over the v6.x line: comparable pricing as a per-property input (derived from prior sales of similar character — same type, bedrooms and bathrooms — similarity-weighted and growth-adjusted), per-property SHAP contributions, and the FSD confidence model. v6.12 (this release) leaves the prediction model untouched and rebuilds the FSD model on a single-pipeline, density-aware, robust-dispersion basis with stored decile boundaries (Section 5).

The choice of gradient-boosted trees over neural networks is deliberate: on structured tabular data they match or exceed deep-learning performance while remaining fully explainable via SHAP — the model produces no output through black-box techniques alone.

2.2 Data Sources

SOURCE	SCALE	UPDATE FREQUENCY	LICENCE
HM Land Registry Price Paid	31.3M transactions (full England & Wales record)	Monthly	OGL v3
EPC Register	29M certificates	Monthly	OGL v3
ONS Census 2021	All LSOAs	Decennial	OGL v3
ONS House Price Index	National/regional	Monthly	OGL v3
GIAS School Register	~24,000 schools	Daily	OGL v3
Ofsted Inspection Ratings	All inspected schools	Rolling	OGL v3
Environment Agency	Flood zones 2/3	Annual	OGL v3

Additional open government datasets (deprivation indices, recorded crime, transport accessibility) feed the neighbourhood and environmental feature categories. All data is sourced under open government licensing. No proprietary third-party datasets — the model's operation carries no data-supplier dependency or intellectual-property encumbrance.

The training set comprises 8,401,368 EPC-matched residential sales (2015 onwards) of the four covered property types, filtered from the full 31,312,370-transaction Land Registry Price Paid record and enriched with area data. The set includes Land Registry Category B (additional price paid) records — it is not filtered to arm's-length or market sales only. Missing values are not imputed — missingness is treated as informative, and properties with insufficient critical data are declined rather than valued.

2.3 Training Process

The production model is trained on 8,401,368 completed transactions (January 2015 – June 2025) and tested on the 295,026 transactions that completed after the training cutoff (Section 3). **Training data is current to 1 July 2025** — v6.12 is a recalibration on the v6.11 prediction model and does not refresh the training window; the next data refresh is the scope of v6.13. Excluded from training: transactions without a matched EPC, prices below £10,000 or above £10,000,000, floor areas outside 10–1,000 m², and records that cannot be geocoded. The target variable is log-transformed; the loss function is L1 (median-seeking), which is robust to outlier transactions. Overfitting is controlled by early stopping on a temporal validation set, row and feature subsampling, and minimum leaf sizes — and ultimately tested by the temporal holdout itself, since the model must generalise to a period it has never seen.

2.4 Hyperlocal House Price Index

Temporal adjustment uses a proprietary hyperlocal index rather than the national ONS HPI alone: a postcode-district × property-type stratified quarterly index built from the full 31.3Mtransaction Land Registry record. This is the model's principal mechanism for tracking local market movement between a comparable's sale date and the valuation date.

2.5 Proxies and Derived Inputs

Two derived inputs matter for model-risk review: prior-sale prices are brought to current terms using the index above (an indexation proxy), and floor area / construction attributes derive from EPC records where present (an administrative-register proxy). Index error in thin markets and stale or absent EPCs are among the drivers of the segments addressed in the v6.12 development cycle (Section 7).

Evidence: /methodology, /changelog (per-version feature history).

---

3. Validation

Validation is by the following means:

1. **Internal:** bulk testing against known sales (Land Registry) on every model release, results of which are published and versioned.
2. **External:** clients' own testing against known sales (their own recent data), which prevents any possibility of the test being gamed (since we have never seen those sales).

In a future iteration, it is intended to enlist third-party validation by Fitch.

3.1 Walk-Forward Out-of-Sample Testing

All published accuracy figures come from strict walk-forward testing: the model is trained on transactions up to a cutoff date and tested exclusively on later transactions it has never seen. The current published test is the **H2-2025 census**: 295,026 Land Registry transactions completing between July and December 2025. This is a census of qualifying completions in the window, not a sample.

3.2 Benchmark: Actual Sale Prices

All accuracy metrics are measured against **Land Registry completion prices** — the amount actually paid in completed transactions. This is a materially harder benchmark than measuring against surveyor opinion, which is itself an estimate, and which in published research is anchored by the known sale price in most lending contexts. Published PE10 figures from providers benchmarking against surveyor valuations are not comparable with the figures in this document; comparisons should be made on the benchmark, not the headline number.

3.3 Cadence and Republication

Bulk tests are re-run and republished with every model release, and the published accuracy page is generated from the same results database as this document — the figures here and at /accuracy have a single source. The model is retrained quarterly as new Land Registry data accumulates, with ad hoc retraining if monitoring indicates drift; every release carries a dated changelog entry with its own bulk-test results.

Evidence: /accuracy, /audit-trail, /methodology.

---

4. Accuracy Results — H2-2025, Meridian v6.12

All tables below: H2-2025 out-of-sample census, 295,026 transactions, model v6.12, benchmark = Land Registry completion prices. **v6.12 is a recalibration release**: point valuations are byte-identical to v6.11, so the overall PE-family and MdAPE in 4.1 are unchanged from v6.11. The band split and per-band statistics in 4.2 are the v6.12 honest re-allocation produced by the rebuilt FSD calibration (Section 5).

Note on bias. From v6.12 onwards, bias is reported as the **median signed percentage error** — the robust measure adopted alongside FSD per Option A (EAA ESSVM convention). v6.11 figures were the mean signed percentage error. This is a measurement-definition change; it can move the number, and occasionally flip its sign, **without any change to the model's predictions**.

4.1 Headline

METRIC	VALUE
Median Absolute % Error (MdAPE)	**6.69%**
Within ±5% (PE5)	**39.6%**
Within ±10% (PE10)	**65.4%**
Within ±15% (PE15)	**79.8%**
Within ±20% (PE20)	**87.6%**
Bias (median signed)	**−0.79%**
Test cohort	295,026 (Jul–Dec 2025 completions)

4.2 By Fitch Classification Band

Confidence-segmented accuracy is the view a lender's band policy consumes. Bands are assigned per property from the FSD (Section 5), following Fitch's published RMBS classification framework; the table reports how each band performed out of sample under v6.12.

BAND	FSD RANGE	N	% OF TEST	PE10	MDAPE	BIAS
A	≤ 0.05	0	0.0%	—	—	—
B	≤ 0.10	195,369	66.2%	**70.3%**	5.93%	−0.66%
C	≤ 0.20	99,077	33.6%	56.1%	8.56%	−1.17%
D	> 0.20	580	0.2%	35.0%	14.30%	+1.58%

Under robust dispersion no cell falls at or below the 0.05 threshold, so no valuation is currently issued in Band A. Band B expanded under v6.12 from 48.7% to 66.2% of the test at unchanged PE10 (~70%): the recalibration recognises more reliable cells than the raw-σ v6.11 calibration did, while the precision of the band itself is held constant. Band D tightened from 9.4% (27,693 properties) to 0.2% (580) and now identifies the genuinely uncertain residual rather than acting as a catch-all; a lender's own policy can decline, discount, or re-route Band D on the band alone.

4.3 By Property Type

TYPE	MDAPE	PE10	PE15	PE20	BIAS	N
Semi-Detached	6.21%	68.3%	82.2%	89.6%	−0.82%	87,931
Detached	6.59%	66.6%	81.5%	89.2%	−0.65%	63,698
Terraced	6.66%	65.2%	79.2%	86.9%	−0.63%	94,834
Flat	7.90%	59.4%	74.4%	83.2%	−1.28%	48,563

4.4 By Price Band

PRICE BAND	MDAPE	PE10	BIAS	N
£300K–£500K	5.76%	72.0%	−1.33%	85,292
£150K–£300K	6.26%	68.2%	−0.71%	117,964
£500K–£1M	6.81%	65.5%	−2.32%	41,628
Under £150K	10.75%	47.5%	+3.76%	44,215
Over £1M	10.12%	49.5%	−5.31%	5,927

4.5 By Region

REGION	MDAPE	PE10	BIAS	N
East of England	5.98%	69.9%	−0.54%	30,478
South East	5.98%	69.8%	−0.48%	45,148
East Midlands	6.11%	68.4%	−0.51%	21,132
South West	6.36%	67.5%	−0.65%	33,037
London	6.68%	66.0%	−0.41%	37,463
West Midlands	6.69%	65.7%	−1.35%	25,190
Yorkshire	7.34%	62.1%	−1.50%	31,967
Wales	7.42%	61.3%	−0.97%	14,770
North West	7.59%	60.7%	−0.89%	41,147
North East	8.28%	56.9%	−1.57%	14,595

Read together with the band table in 4.2, the segmentation in 4.3–4.5 is the model's operating envelope: performance is strongest for houses in the £150K–£500K range in southern and eastern England. The under-£150K and over-£1M segments and the weaker northern regions are the targets of the v6.13 training-window refresh (Section 7).

Evidence: every table above is generated from the same bulk-test results record that drives /accuracy; downloadable PDF reports of the same results are linked from that page.

5. Confidence Model — Per-Property FSD, Fitch Bands, and Value Range

5.1 Forecast Standard Deviation (Option A / EAA ESSVM convention)

Every valuation carries a per-property **Forecast Standard Deviation (FSD)** — under v6.12 onwards, the **robust per-cell dispersion** of percentage errors for properties with that profile, measured as IQR / 1.349 of signed percentage errors. This is the Option-A precision figure adopted alongside the EAA ESSVM convention: FSD is dispersion only; per-cell bias is reported as a separate figure (Section 4) for disclosure, not folded into FSD or the range.

FSDs are assigned from a granular calibration cell — keyed on confidence decile × property type × price band × evidence density — built and stored at calibration time. The serving code reads stored decile boundaries and looks up the cell directly; calibration-time and serving-time scoring share a single code path by construction. **Decile agreement (calibration vs serving) = 100% on the H2-2025 census** (v6.12 build).

FSD-segmented accuracy is republished with every bulk test so the calibration can be checked against outcomes (Section 4.2).

5.2 Mapping to the Fitch Band Schema

For lender convenience the FSD is mapped to the classification bands published in Fitch Ratings' RMBS criteria for AVM-derived valuations:

BAND	FSD
A	≤ 0.05
B	≤ 0.10
C	≤ 0.20
D	> 0.20

This is a mapping to Fitch's published framework, stated so that our confidence metric arrives in units a securitisation or credit-risk team already uses. It implies no Fitch review, recognition, or approval of Gadsden Valuations. FSD is the standard metric Fitch's published AVM treatment consumes, which is why we express confidence in it natively rather than in a proprietary scale. (The API additionally returns a 0–7 confidence level on the European AVM Alliance's common scale, derived from the same FSD.)

The FSD and band are returned on every valuation — API, dashboard, and PDF — so a lender can apply its own band policy per property with no further computation. Gadsden Valuations does not prescribe haircuts; any treatment applied to a band is the lender's own policy decision.

5.3 Value Range — FSD-to-Coverage Mapping (ESSVM §5.3)

The European AVM Alliance ESSVM 3rd Edition §5.3 requires that providers state a clear mapping between their confidence measure (FSD) and the empirical coverage of the published valuation range. The Gadsden Meridian mapping under v6.12 is as follows.

Published range convention. The published 90% valuation range is centred on the predicted value at ± **z = 1.645** × FSD — the standard normal-curve convention the EAA ESSVM range schema uses. The same convention is used by the principal UK AVM incumbents; a lender's model-risk team can read GV's range alongside its other providers without reconvention.

Measured empirical coverage on the H2-2025 census (v6.12). Because residential property valuation errors are fat-tailed, the empirical coverage of a robust-σ ± 1.645 interval is below the theoretical normal-curve 90%:

BAND	MEAN FSD	EMPIRICAL COVERAGE @ Z=1.645	Z FOR EMPIRICAL 90%
B	0.0873	82.6%	2.20
C	0.1270	82.9%	2.20
D	0.2176	85.2%	2.02
Overall	—	**82.7%**	**2.20**

Two things to note. First, the implied-90% empirical coverage is **~83% uniform across bands** — the v6.12 robust-σ measure means the same thing in every band; the only gap from a literal 90% interval is the fat tail, captured by a single global multiplier (z ≈ 2.20 across all bands). Under v6.11's raw-σ calibration the per-band empirical coverages were widely uneven (Band B 78%, Band C 87%, Band D 96%) and FSD's meaning drifted band to band. v6.12 fixes that.

Second, evidence density is now a calibration dimension. Where comparable evidence is sparse (the 20th comparable beyond 1.5 km of the subject), the calibration assigns a higher mean FSD (0.117 vs 0.100 for the dense bucket), and the assigned FSD tracks realised robust σ to within 1 percentage point in each density bucket. The decline gate (Section 5.4) builds on this honest density signal.

This mapping is restated on the public /accuracy page so lenders can verify it against the live data on the same day.

5.4 Decline ("No-Value") Policy

The model declines to value rather than guessing. A valuation is declined when the available comparable market evidence is insufficient or critical property data is missing. Declines return a machine-readable reason code and a plain-language description (Annex A.4), and are not charged.

For model-risk purposes the decline policy and the band schema together function as pre-packaged use restrictions: the model self-identifies the cases where it should not be relied on, before any lender-side policy is applied.

Evidence: /no-value-policy; sample reports across bands B, C, D and a no-value response at /downloads.

6. Explainability and Audit Trail

6.1 Per-Property SHAP Contributions

Every valuation includes per-property SHAP feature contributions — the top contributors to that specific estimate, expressed as positive or negative pound values — in the API response (`feature_contributions`) and as the "What Drives This Valuation" section of the PDF report. This is per-property attribution, not a global importance chart: a reviewer can see why this property received this value.

6.2 Comparable Evidence

Each valuation is accompanied by the comparable transactions used as market evidence — recent sales of the same property type, nearest first, with address, sale date, price, floor area, and distance. This gives a reviewer the same evidence a valuer would seek independently.

6.3 Audit Trail

Each valuation permanently records: the exact model version that produced it; date and time; a unique valuation reference; the full feature vector (every input the model saw); the output — point estimate, range, FSD, and band; the comparables selected as evidence; and the requesting user and channel (web, API, bulk). A valuation is therefore reproducible and reviewable after the fact, including across model versions.

6.4 Sample Evidence

The published sample portfolio includes contrasting cases — Band B, Band C, Band D, and a declined (no-value) response — so reviewers can see the full range of model output.

Evidence: /audit-trail, /downloads, Annex A (API response schema).

---

7. Structural Limitations and the v6.13 Development Cycle

7.1 Structural Limitations (any AVM)

The model values from recorded data and cannot observe: internal condition and renovation status (the principal limitation of any AVM); bespoke improvements not yet in EPC or planning records; legal encumbrances beyond Land Registry data; micro-location factors (view, aspect, noise); and the 2–4 month Land Registry registration lag, which the hyperlocal index partially compensates. The per-property FSD is the mechanism by which data thinness surfaces at the individual-valuation level rather than only in aggregate statistics.

7.2 v6.13 Improvement Agenda

The next cycle (v6.13) is a training-window refresh — the principal data update v6.12 deferred — and addresses the residual segment weaknesses visible in Section 4. In priority order:

1. **Sub-£150K pricing** — the segment shows the strongest residual bias (+3.8% median, Section 4.4) and the weakest PE10 (47.5%); the training refresh and feature work on this segment is the cycle's principal objective.

2. **£1M+ pricing** — the £1M+ segment carries a −5.3% bias and 49.5% PE10; comparable evidence thins at the top of the market.
3. **Sparse-market comparable evidence** — refinement of the comparable-selection rules so they adapt fully to local market density and property character in rural and thin markets. The v6.12 evidence-density dimension surfaces these cases honestly in FSD; v6.13 targets the underlying prediction quality.
4. **Flat-specific inputs** — lease and block-level factors are the gap behind the flat segment's results (Section 4.3, PE10 59.4%).

Each release is published with its own bulk-test results, so the effect of the cycle is measurable on the same tables as Section 4.

---

8. Test Us Against Your Own Book

We offer a free retrospective backtest against your completed transactions — your addresses and completion dates in, our valuations out, scored on your numbers. The same walk-forward discipline applies, and since we have never seen your completion data, the test cannot be gamed. This produces validation evidence on your portfolio's composition rather than the national census.

Evidence: /for-lenders.

---

9. What We Supply to Your Model-Risk and Vendor Review

Providing AVM outputs to lenders is not a regulated activity. The model-risk frameworks a lender may operate under — such as PRA SS1/23, which applies to UK banks with internal model approval — bind the lender, not the vendor. What we maintain is the material those frameworks expect an in-scope lender to obtain from a third-party model vendor:

- **Validation evidence**: dated, versioned out-of-sample results against a stated benchmark, segmented by type, price, region, and confidence band (Section 4), republished with every release.
- **Methodology documentation**: this document and the published methodology page; deeper feature-level documentation is available to organisations conducting formal vendor due diligence.
- **Data and assumption transparency**: sources, licences, update cadences, proxies, and the no-imputation policy (Section 2).
- **Per-property confidence and restriction logic**: FSD, band, and decline policy (Section 5) — the basis for lender-side policy treatment, overrides, and use restrictions.
- **Explainability and auditability**: per-property SHAP and a complete, reproducible audit record (Section 6).
- **Change governance**: a public, dated version history (the changelog); bulk-test results published per release.

- **Independent verification route**: the free backtest (Section 8), which puts validation evidence on the lender's own outcomes rather than our chosen cohort.

One dated fact on the lender's side of the rulebook: from 1 January 2027, PRA Basel 3.1 rules expressly permit a suitably robust statistical method as the valuation at origination (PRA Rulebook, Art. 124D(8)(a); the PRA has stated this may include AVMs — PS9/24). Whether and how to use AVM output remains a matter for each lender's own policy.

Evidence: /compliance, /for-lenders, /vendor-qualification.

Annex A — API Documentation

The API returns the same outputs as the dashboard and PDF reports, from the same model and audit pipeline. The machine-readable specification is published at `https://gadsdenvaluations.com/api/v1/openapi.json`.

A.1 Authentication

All endpoints require an API key (`gak_…`) supplied as a Bearer token. Keys are created and revoked from the account dashboard or via the key-management endpoints.

A.2 Endpoints

METHOD	PATH	PURPOSE
POST	`/api/v1/valuations`	Request a valuation (by UPRN, property ID, or address + postcode)
GET	`/api/v1/valuations/{id}`	Retrieve a previously issued valuation by reference
POST	`/api/v1/valuations/batch`	Batch valuation (list of property IDs)
GET	`/api/v1/account`	Account status and remaining valuations
GET/POST/DELETE	`/api/v1/keys`	API key management

A.3 Valuation Response

A successful valuation returns:

FIELD	CONTENT
`predicted_value`	Point estimate of market value (£)
`valuation_range`	Lower and upper bounds derived from the FSD
`fsd`, `fsd_band`, `fitch_haircut_pct`	Per-property FSD, its Fitch band (A–D), and the haircut fraction Fitch's published framework associates with that band (reference only)
`confidence`	The same confidence detail, including the EAA 0–7 level
`property`	The subject property as the model saw it: type, bedrooms, floor area, construction period, EPC rating, last sale
`comparables`	The comparable sales used as evidence: address, date, price, type, floor area, distance
`feature_contributions`	Per-property SHAP contributions in pound values
`model_version`	The Gadsden Meridian release that produced the valuation — the same version label shown on /accuracy, /changelog, and the site footer. Advances with each release, whether retrain or recalibration.
`timestamp`, `processing_time_ms`	Issue time and latency

A rental estimate (estimated monthly/annual rent, gross yield, and rental comparables) is included where rental evidence is available.

A.4 No-Value Response

Where the model declines to value, the API returns `status: "no_value"` with a machine-readable `reason` code, a plain-language `reason_description`, and a `missing_data` list identifying the absent inputs. No-value responses are not charged.

A.5 Batch and Backtesting

The batch endpoint accepts a list of properties and returns per-property results in the same schema. Valuations can be requested as at a historical date (`as_of_date`), which is the mechanism behind the retrospective backtest (Section 8): the model uses only data available before that date.

Contact: sales@gadsdenvaluations.com | gadsdenvaluations.com

Figures are from the H2-2025 bulk test of Meridian v6.12 and are superseded by any later bulk test published at gadsdenvaluations.com/accuracy. Use of AVM output within a regulated lending process is governed by the lender's own policies.