



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

Predictive Routing Help

Lift Estimation Report

12/15/2025

Contents

- 1 Lift Estimation Report
 - 1.1 Lift Estimation End-to-End Example
 - 1.2 How to Generate a Lift Estimation Report
 - 1.3 Reading a Lift Estimation Report
 - 1.4 Lift Estimation Report In-Depth
 - 1.5 Calculating Lift for Metrics You Want to Minimize (releases earlier than 9.0.012.01 only)

Lift Estimation Report

The Lift Estimation report generates an estimate of what lift you might be able to achieve in the value of the target metric optimized compared with your existing routing. It functions as a sort of testing of the waters, providing a broad-stroke outline of the sort of results you might expect with a full implementation of Genesys Predictive Routing (GPR). In contrast with model evaluation metrics, such as Area Under the Curve (AUC) or Mean Absolute Error, which reflect the accuracy of a GPR Model in predicting the outcome of an interaction + agent combination, Lift Estimation attempts to estimate the improvement in the target metric you ultimately care about.

- For an in-depth discussion of how GPR handles metrics, see [Understanding Score Expressions](#).

For example, in a First Contact Resolution maximization use case, the AUC of the Model indicates how well the Model predicts the number of resolved and unresolved outcomes. The Lift Estimation report directly estimates the amount of improvement you might expect in FCR when the predictive Model is deployed in production.

The key challenge with an estimate is that it is based on historical data. As a result, we only know the actual outcome of an handled by a specific agent. We do not know *for certain* what the outcome would have been if the exact same interaction was handled by the agent GPR selected as optimal. The underlying algorithm *estimates* the outcome for the GPR-selected agent based on historical interactions where the agent actually selected matches the one GPR would have selected.

Keep in mind the following points regarding the results of the analysis. The Lift Estimation report:

- is based on historical data from your contact center.
- assumes an agent-surplus situation.
- assumes your baseline or existing routing selects agents randomly without using any predictive model.
- includes only agents found in the Dataset; agents found only in the Agent Profile are not included.
- does not factor in handle time for the interactions
- depends on the quality of the Model you use to create it. See [View Model quality and agent coverage reports](#) for information on how to assess Model quality.

Important

- See [Lift Estimation Report In-Depth](#) for a detailed discussion of how the Lift Estimation report works.
- See [Lift Estimation vs A/B Testing](#) for an explanation of how these reports differ and the advantages of each.
- See [Best Practices and Troubleshooting](#) to ensure that you are setting up the report as effectively as possible.
- See [Calculating Lift for Metrics You Want to Minimize](#) (releases earlier than 9.0.012.01

only) for how generate estimates for metrics, such as AHT and number of escalations, where a lower value is better.

Lift Estimation End-to-End Example

James' business has been doing reasonably well, but customer satisfaction numbers have stagnated and even seem to be decreasing. James knows his business needs to pick up the pace. Genesys Predictive Routing (GPR) looks like a great idea. But can it make the difference he's looking for?

James' business has a lot of data available, about agents and customers, and about how the interactions turned out. He deploys GPR in a test environment and creates an initial Dataset, drawing especially on interaction data from Genesys Info Mart, agent data from Configuration Server, customer data from the CRM system, and results from customer satisfaction surveys (CSAT). He combines this data into a CSV file, and imports it into GPR.

For these initial trial runs of GPR, James does not create separate Agent Profile and Customer Profile schemas. He can get an adequate estimate from a single Dataset.

James know that he has CSAT results for only a small percentage of the total interactions. To make use of more of his rich data stockpile, he decides he will also run a Lift Estimation based on first-contact resolution (FCR)—all interactions where customers did not call back for the same reason within a set time. So James opens his GPR web interface and creates two Predictors, one for CSAT, one for FCR. For each, he uses the default full-feature set Model, which is created automatically when the Predictor is created. At this point, he's ready to consider all factors as possible drivers for improvement.

Now James starts running Lift Estimation reports, starting with the default numbers of samples and simulations. He starts with no Group By setting, then experiments with grouping by different parameters, such as agent groups, customer types, customer intents, agent tenure, and so on.

To get a view-at-a-glance compilation of the results of his Lift Estimation experiments, James creates a table similar to the following one.


Metric	Group By	Model Name	Report Name	Model Quality	No GPR	GPR		Uplift		# Simulations	# Samples	Tuning Threshold
5%	10%	5%	10%	5%	10%							
FCR	VAG-1	FCR_Uplift	FCR_Uplift	AUC: 0.98 R2S: 0.98 MAE: 8-10.92	0.92	0.96	0.97	50%	62.5%	100	10,000	50
CSAT	VAG-1	CSAT_Uplift	CSAT_Uplift	AUC: 0.85 R2S: 0.85 MAE: 8-12.08	0.80	0.86	0.86	7.5%	7.5%	100	10,000	50

As James views the results from these reports, he starts to breath more easily. He's seeing potential lift. Of course, the more agent availability, the better GPR looks. That's only natural; the more agents available, the more likely GPR can pick the absolutely ideal matchup. But even a real-life availability levels, ones he's used to from his years of experience, GPR can get interactions to agents who are

better at specific kinds of issues. He can improve his company's customer service, even without a single change to training or other factors.

And with the Feature Analysis report, he can target factors that will most effectively drive improvement, enabling his company to allocate resources where it really matters. And by integrating out of the box with Genesys Reporting, he can run that A/B testing report and see a head to head matchup between his current routing and Predictive Routing. He's not going to bet against Genesys on that one.

How to Generate a Lift Estimation Report



The screenshot shows a web form for generating a Lift Estimation report. The form is titled 'Report' and has a dropdown menu set to 'Lift Estimation'. Below this, there are several input fields: 'Report title*' with the value 'Lift Estimation', 'Predictor Name*' with the value 'TEST1', 'Predictor Model*' with a dropdown set to 'm2', 'Number of Simulations*' with the value '100', and 'Number of Samples*' with the value '100'. There is also a 'Group By' dropdown set to 'None Selected'. At the bottom right of the form is a blue button labeled 'Run Analysis'.

To create a **Lift Estimation** analysis report, start on the **Predictors** window. The report uses the metric and parameters set in the Predictor you select on the **Trend** tab or **Details** tab of the **Predictors** window.

Important

- The Lift Estimation report does not support composite Predictors.

- To configure this report correctly, see the [Lift Estimation Best Practice Recommendations](#), below.
- Predictive Routing supports report generation that includes up to 250 features (columns).
- For all reports, mandatory fields are marked with an asterisk.

To configure the report, complete the following fields:

1. Click **Analysis**, located on the right side of the top navigation bar.
2. Select **Lift Estimation** from the **Report** drop-down menu.
3. Select a Model from the drop-down list. This Model can be active or inactive, but it must have been trained.
4. In the **Number of Simulations** field, enter a figure for how many times to select subsets of agents; each subset of agents is used in one simulation and GPR calculates an average across the specified number of simulations to account for the randomness in subset selection. The size of these subsets is determined by the agent availability figure calculated. Predictive Routing calculates lift for all agent availability percentages, from 0.01, when only 1% of total agents are available, to 1, when 100% of agents are available. For example, 50% availability uses random selections of half the total agents; 25% availability uses simulations that each pick a random 25% of the total number of agents.
By default, the number of simulations is set to 100, meaning 100 subsets of agents are selected by random sampling. The field accepts any value larger than 0 and less than or equal to 500.
5. In the **Number of Samples** field, enter a value for how many samples to use from the test set. By default, this is set to 100. This value should be approximately 30 times the number of agents (that is, an average of 30 samples per agent). Note that the number of samples cannot exceed the number of records in the test section of the Dataset. If you enter a larger number, the Lift Estimate report runs only the available number of samples in the test portion.
6. In the **Group By** field, enter a parameter to use to group interactions for estimation. Only agents who handled interactions of the type specified in the group-by value are used in the estimation for that group.
When you specify a high-cardinality feature as a grouping parameter, the top 20 values of that feature, as determined by interaction volume, are extracted and a report is generated for each one of those groups.
7. If you want to use values other than the default top 20 group values for high-cardinality features, select the **Advanced** check box, and then choose up to 20 group values against which you want to run the Lift Estimation report. By default, the 20 feature values with the highest interaction volume are available for selection in the **Group Values** field when **Advanced** mode is on.
If there are no interactions in the test part of the Dataset for a selected Group By feature, that feature does not appear in the generated report.
8. Click **Run Analysis**.
The result appears on the **Reports** tab.

Lift Estimation Best Practices and Troubleshooting

Important

If you are troubleshooting and encounter problems not covered here, contact Genesys Customer Care for additional support.

Number of Simulations

- This value should be higher for larger numbers of agents. 100 is an appropriate value for most environments. Increase this value to the maximum, 500, if you are scoring for large number of agents (more than 5000).

The Group By Parameter

- If you are using the **Group By** functionality, the Lift Estimation dataset is constructed with latest **Number of Samples** rows for each unique value in the **Group By** feature selected. The rows are drawn from the Predictor Dataset, using rows that fit the criteria. If there are insufficient records for a certain **Group By** value, that set is underrepresented.
For example, if **Group By** is set to Queue, which has two values, VQ1 and VQ2, and **Number of Samples** is set to 1000, GPR constructs two Lift Estimate Datasets (VQ1_Dataset, VQ2_Dataset). Each Dataset should be analyzed based on the most recent 1,000 samples from the corresponding queue. However, if there are only 200 rows in the Predictor Dataset where Queue=VQ2, the VQ2_Dataset would be short of 800 records compared with the Dataset for VQ_1. As a result, the results for VQ_2 are probably less accurate than those for VQ_1, and do not provide equally strong results.
- To group a Lift Estimation report by queue, you must include the desired queue as a Predictor (context) feature.

Remove Erroneously-Collected Target Variables

Data collection errors can distort predictions, especially with respect to the value of a target variable. Before uploading your data to GPR, remove any target variables that contain clearly erroneous values and that do not represent the actual outcome of the interaction.

Examples of values that should be removed:

- A handling time of 0 seconds. GPR processes zeros like any other number. However, a handling time of 0 is not a viable real-world value, which means that this value is generated by some data collection issue.
- A handling time of 10+ hours for an immediate-response interaction type, such as voice. No real-world customer phone call lasts for ten hours, indicating some upstream error that resulted in the erroneous value.
- A customer satisfaction score of 99 when the normal range is 1-10. Some statisticians encode missing values with an arbitrary number, such as 99. Any scores falling outside the valid range should be removed.

These examples are not exhaustive. Part of the process of creating an effective Dataset is to remove values at your discretion when you see that they do not represent a realistic outcome for an interaction.

Lift Estimation Report Not Available for Selection

- Check that you are on the **Predictors Trend** tab or **Details** tab.
- Set the date range to a period that has data.

The Lift Estimation Reports is Empty


- Verify that you have data for the selected period and that you have created and trained your Predictor.

A "No data available for your chosen parameters for analysis type lift_estimation" Error Appears

- The selected Group By filter might not yield enough data.
 - The trained Model is old and the Predictor has been updated since it was created and trained.
To resolve this issue, regenerate the Predictor, add all of the Agent and Customer features, and retrain/reactivate a new Model.
-

Reading a Lift Estimation Report

The Lift Estimation graph shows two lines. The blue line represents your baseline routing. The black line represents the results of using Predictive Routing. The distance between the two lines represents the expected lift. The diagram shows the estimated outcome for various levels of agent availability.



To produce the estimate, the Lift Estimation analysis takes the specified number of interactions, identifies target agent pools from the data based on the day of interaction and group-by value, if any, and then runs multiple simulations for differing levels of availability. The restriction of target agents to those in agent groups active on the day of the interaction is to ensure that GPR has the exact same choices as what the existing routing had, and thus to have a level-playing field.

- The y-axis shows the target metric used in the Predictor. If you want to maximize your target metric, as for FCR, for example, the higher up the y-axis, the better the routing performance. If you want to minimize your target metric, such as Handle Time, the lower on the y-axis, the better the result.
- The x-axis represents the agent availability factor, which is part of the simulation. The agent availability could range from 0.01, when only 1% of total agents are available, to 1, when 100% of agents are available. The range on the x-axis can vary depending on the total number of agents under consideration.

You can identify the potential lift in your environment by checking the difference between the blue and black lines at the point on the x-axis that corresponds to your average agent availability.

Important

If the target metric you are evaluating is a Boolean, the Boolean values are converted to the equivalent numeric values, with false = 0 and true = 1. Values between 0 and 1 can be interpreted as the percent chance of the metric being true.

If you select a **Group By** value when defining the report parameters, target candidate agents for scoring and lift estimation are defined within the specified group, and the tabs above the chart show Lift Estimation reports for each group. The **Aggregate** tab displays the weighted average for each availability based on the number of interactions for each group. The shape of the lift estimation curve depends on the following two factors:

- The variation in scores across agents for each interaction. If many agents have similar scores that are close to the maximum, an increase in agent availability does not significantly impact the lift, and the line extending across the range of availability percentages does not increase or decrease as the availability changes.
- The accuracy of the Model. The position of the data points on the curve relative to the base line is determined by the reliability of the predicted scores against the true outcome.

Important

In some cases, overfitting of the underlying Models might produce negative lift in scenarios with high agent availability. While this may look counterintuitive, it is a result of negative correlation between the predicted and the actual scores. To reduce overfitting, ensure that individual agent Models (created for DISJOINT and HYBRID Model configurations) are trained with enough samples and/or reduce the number of features included in the Model.

Customizing the Display

- You can toggle between the default graph view and a tabular view of the same data points. The rows highlighted in green show positive lift.
- To save your report outcome in CSV format, click the **Export** button. The report downloads immediately.

Lift Estimation

Graph/Table Toggle

Created: Oct 16, 2019 3:05 PM
Author: Y

Analysis Type: Lift Estimation (maximization) | Predictor: PREDICTOR | Predictor Model: Full Feature set | Group By: Interaction_Agent_SkillLevelForCall | Number Of Samples: 100 | Number Of Simulations: | Delete | Export

Interaction_Agent_SkillLevelForCall:Aggregate | Interaction_Agent_SkillLevelForCall:5 | Interaction_Agent_SkillLevelForCall:4 | Interaction_Agent_SkillLevelForCall:6 | Interaction_Agent_SkillLevelForCall:7 | Interaction_Agent_SkillLevelForCall:3 | Interaction_Agent_SkillLevelForCall:2 | Interaction_Agent_SkillLevelForCall:8 | Interaction_Agent_SkillLevelForCall:9 | Interaction_Agent_SkillLevelForCall:1

Agent Availability	Target Metric	Lift
Existing Routing	0.07	-
0.01	0	0%
0.02	0	0%
0.05	0	0%
0.10	0	0%
0.25	0.05	0.17%
0.50	0.08	1.40%
0.75	0.11	0.07%

Lift Estimation Report In-Depth

The Lift Estimation report does a *hindsight analysis* to estimate the lift in a specified KPI if the interactions in the test dataset had been routed using the selected predictive Model instead of the baseline or existing routing.

To recreate the original routing scenarios as closely as possible, GPR limits the choice of agents to those who handled calls on a given day as found in the test data collected through the existing baseline routing and further assumes the agent-specific features are stable for the given day.

Agent Availability

Agent availability is defined as the fraction of the total number of agents who handled calls on the day of the interaction within the group selected in the **Group By** parameter (if used).

- To simulate 100% agent availability, for each interaction, GPR scores all agents who are part of the agent pool, which includes all agents working on the relevant day and belonging to the group selected in the group-by value. GPR picks the agent from that pool with the maximum score.
- To simulate 50% agent availability GPR randomly selects half the agents from the daily agent pool, scoring them, and picking the maximum-scored agent from that half.

To ensure that the lift produced by routing with GPR in a reduced agent-availability scenario does not happen by chance, GPR repeats the Lift Estimation analysis on a random selection of agents many times (100 runs, for example), and averages the estimate across the runs.

Important

100% agent availability is unrealistic in any real-world setting. To focus on meaningful results, disregard the section of the generated report showing agent availabilities

higher than 50%, or higher than you would normally expect to see.

Adjusting the availability of agents shows how different operating constraints can yield different outcomes. On average, agents should have approximately 30 samples for good-quality analysis results.

Lift Estimation vs A/B Testing

The Lift Estimation report is a useful, efficient way to get a reasonably quick assessment of how well a Model performs and the expected lift in a specified KPI over baseline routing. However, it does not replace A/B testing. The following table highlights key differences between the two types of reports.

Important

The A/B Testing report is generated from data written to the Genesys Info Mart database.

- For an explanation of how to create and view the A/B Testing report, see [Predictive Routing A/B Testing Report](#) in the *Genesys Customer Experience Insights User's Guide*.
- For instructions on how to set up Genesys Reporting to store GPR data, see [Integrate with Genesys Reporting](#) in the *Predictive Routing Deployment and Operations Guide*.

Lift Estimation	A/B Testing
Offline hindsight analysis using historical interaction data	Online testing in the live production environment; considered as the ultimate test
Estimates the scope for KPI improvement under various assumptions	Determines the real improvement in KPI with minimal or no assumptions
Being offline analysis, no real-time processing resources involved, which might slow down performance	A poor Model or incorrect routing strategy could potentially have detrimental effects
Assumes that the baseline routing selects agents randomly	No assumption on baseline agent selection
Lift is estimated only for agent-surplus scenarios	Can measure performance for both caller-surplus and agent-surplus scenarios
Likely agent availability to be calculated from past data prior to determining the applicable estimate	Lift is calculated based on real agent availability during the period of assessment, hence more accurate
Like any other statistical estimate, it is associated with error which could be higher when assumptions are violated	When ensured Control and Target routing methods operate in similar conditions, the calculated lift is likely to be unbiased and reliable.

Calculating Lift for Metrics You Want to Minimize (releases earlier than 9.0.012.01 only)

For releases earlier than 9.0.012.01, creating a Lift Estimation report for a metric where a lower value is preferable requires additional configuration. The following procedure uses AHT and EscalationFlag as the example metrics. AHT is a straightforward example of a metric where a lower outcome value is better. EscalationFlag is a Boolean metric, where 0 = no escalation, which is the preferred result.

To generate values for the Lift Estimation report, use the following procedure:

1. Create a new outcome column in your dataset that stores the *negative* of the original outcome. For example, use the following expression types to transform your metrics:
 - AHT: $\text{New_AHT} = -1 * \text{AHT}$
 - EscalationFlag (where 1 is escalated and 0 is not): $\text{New_EscalationFlag} = 1 - \text{EscalationFlag}$
2. Run a Feature Analysis report on the *original* outcome.
3. Build a Predictor using the key features identified in the Feature Analysis report and built on the new (modified) metric. For example, instead of selecting AHT as your metric, select New_AHT.
4. Follow the standard process of training your Predictor and building a Model using the new Predictor.
5. Run a Lift Estimation report from the **Predictors** tab with the desired Predictor selected.

The Lift Estimation report now correctly displays the expected improvement in the metric value.

Important

For all purposes *other than* running a Lift Estimation report--that is, Feature Analysis, Agent Variance, Predictors and Models--use the *original* metric. The Feature Analysis report, Predictor, and Model created using the new modified metric should be used *only* in the process of creating a Lift Estimation report.