



March Best Practices Webinar:

LLM Based Precision Mode & Algo Weight Customization



Introducing our hosts...



Team Introduction



Sammy Schwartz

Senior Experience Strategist



Harshil Goyal

Senior Product Manager





Agenda

- 1. LLM-Based Precision Mode** - Revamped search recall algorithm utilizing the latest generative AI technology for better recall
 - What is LLM Precision Mode?
 - How Does it Work?
 - How to Enable
- 2. Algo Weight Customization** - Greater flexibility to configure weights of different signals in the Bloomreach algorithms
 - What is Algo Weight Customization?
 - How Does it Work?
 - How to Enable
 - Best Practices
 - Conflict Resolution
 - A/B Testing Plan



**Please add your
questions to the chat as
we go along, which we
will be monitoring!**

How Search Works

There are 2 key functions of search: Recall and Ranking



Recall

The process of **determining which products to return** for a given query.



Ranking

The **order products appear in** once they are retrieved.

How Search Works

There are 2 key functions of search: Recall and Ranking



Recall

The process of **determining which products to return** for a given query.



**LLM Based
Precision
Mode**



Ranking

The **order products appear in** once they are retrieved.



**Algo Weight
Customization**

LLM Based Precision Mode





LLM-Based Precision

Drive deeper relevance and higher precision for your recall sets using the latest machine-learning techniques with LLM-Based Precision Mode. LLMs, or Large Language Models, use AI to learn.

Key Benefits


- **Better query coverage** - LLM-based precision mode will address relevance issues for more tail/torso queries, not just for top queries

Use Cases

- The other precision modes available in Bloomreach are not producing the optimal search results for your site

Enablement

- Enable in the dashboard at a global or query level
- Can also enable in the API
- Able to A/B test outside of the dashboard using the API parameter



Please note, this feature is currently available as a free preview until July 2024

LLM-Based Precision

What does LLM mean?

LLM = Large Language Model

- LLM based precision utilizes cutting edge machine learning technology to truly recognize the related product types behind the entire range of head, torso, and tail queries
- Fine tuned an open source LLM with e-commerce data

Why did we develop this precision mode?


- Other precision modes (Category & Product Type Precision) sometimes can be too precise.
- They rely on user behavior (pixel) data, and cannot always effectively optimize for new & long tail queries

How do I know if LLM Based Precision is the right precision mode for me?

- Text Match precision is resulting in a noisy recall set
- Category & Product Type precision are too precise
- You have a large portion of new & long tail queries where you'd want better relevance



How LLM-Based Precision Works





Comparing Precision Modes

Text Match (Default)	Category Precision	Product Type Precision	LLM Based Precision
- Looks to match query keywords and keyword synonyms to keywords in Searchable Fields - Can sometimes pull in noise, but will prioritize products that match all keywords Least Precise	- Filters the result set by keeping only the the products that are part of categories associated with the top 50 products in the result set - Only works for queries that have user behavior (pixel) data - Fall backs to text match Precision varies	- Utilizes user behavior data to identify related product types - Only works for queries that have user behavior (pixel) data - Fall backs to text match Most Precise	- Uses LLM (Large Language Model) and user behavior data (when available) to identify related product types - Works for all queries where user intent can be identified by BR's semantic engine - Fall backs to text match Optimal Precision

Example: Red Shoes

Returns products that have text matching “red” and “Shoes” and their synonyms	Product type “shoes” identified from user query. Dominant Categories (like shoes, dress shoes, loafers) identified based on customer’s categories associated with the initial 50 products in the recall	Product type “shoes” identified from user query. Related product types (like loafers, trainers..) identified through user behavior data, and synonyms)	Product type “shoes” identified from user query. Related product types (like cleats, boots, pumps, loafers, trainers) identified through AI, user behavior data, and synonyms)
---	--	---	---

How To Enable In Dashboard



Enabling LLM Precision Mode

1. Go to the **Algorithm Customizations** tab and click into **Recall Algorithms**
2. View your **Global Recall Algorithms** and click the pencil icon to edit your **Search Recall Precision**
3. Click the dropdown, select **LLM-based Precision**
4. **Preview** your newly selected precision mode and **Save**

Algorithm Customizations

Real-time customer segments

Recall Algorithms

Algorithm Customizations > Recall Algorithms

Recall algorithms

Customize the core Bloomreach algorithm to optimize for your unique business needs. Configure, preview, A/B test, and launch various algorithm settings [Learn more](#)

Global RulesQuery Overrides

Search...

Influence

Settings ⓘ

Search Recall Precision : Text Match

Facet Precision : Standard

Query Relaxation : Product Type

Spell Correct : Term Frequency

Account homeoasis.bloomreach.com

CancelPreviewSave

Algorithm Customizations > Recall Algorithms

Recall algorithms

Customize the core Bloomreach algorithm to optimize for your unique business needs. Configure, preview, A/B test, and launch various algorithm settings [Learn more](#)

Search Recall Precision

Removes noise from the recall set automatically without manual intervention.

Text Match

Category Precision

Product Type Precision

Text Match

LLM-based Precision

Product Type

Facet Precision

Declutters the facet noise from the irrelevant products in the recall set without impacting the actual product recall itself.

Standard

Spell Correct

Spellcorrect reduces the number of null result queries due to spelling errors.

Term Frequency



Algo Weight Customization

Algo Weight Customization

Gain full control over the ranking algorithms powering your search and category pages with the ability to select signals and assign their weights for ranking on all or specific search and category pages.


The default Bloomreach algorithm weighs most heavily the performance of a product for a given query or category to determine its ranking. With Algo Weight Customization, customers can create rules to give more weight to sitewide performance in determining ranking.

Key Benefits

- More customized algorithms that utilize the key signals important to your business
- Greater control over the algorithms that power your ranking results directly in the dashboard

Use Cases

- Define weights for signals (sitewide views, conversions, ATC, and revenue) for select search or category pages.
- Replace default ranking algorithm with a customized algorithm for select search or category pages



Algo Weight Customization

How Does it Work?

- Just like any numeric or product description ranking rule, create rules to boost by sitewide Performance Attributes
 - Sitewide Revenue**
 - Sitewide Conversions**
 - Sitewide ATC**
 - Sitewide Views**
- Use Performance Attribute boosts in conjunction with the default BR algorithm or override part or all of the BR algorithm to customize ranking
 - Use Global Settings (selected by default):** If there are Global Rules with performance attribute boosts, inherit those settings and apply new Performance Attribute rule in combination
 - Use Default Ranking:** Use default BR algo in combination with the new Performance Attribute rules (disregarding any Global Performance Attribute rules)
 - Disable Performance Attributes:** The weighting of Performance Attributes in the default BR Algo will not be considered, only the Performance Attribute weighting defined in the new Rule. Relevance will still be considered.
 - Disable All Ranking Attributes:** The default BR Algo will be disabled, this query will only be impacted by the Performance Attribute rules, or other ranking rules, no other signals will be considered. Relevance will not be considered. **[Do Not Use without testing first]**

Step 1 → Step 2 → Step 3 Next

Performance Attributes

Select the performance attributes below to boost linearly (larger the value, stronger the impact).

- ☐ Sitewide ATC
- ☐ Sitewide conversion
- ☐ Sitewide revenue
- ☐ Sitewide view

Step 1 → Step 2 → Step 3 Done

Optionally disable the default BR ranking algorithm partially or fully. It will prevent double-counting of certain attributes in the default ranking. If disabled fully, the ranking will be determined solely based on performance attributes configured in the previous step.

CUSTOMIZE RANKING ALGORITHM

- ☒ Use global settings ⓘ
- ☐ Use default ranking ⓘ
- ☐ Disable performance attributes ⓘ
- ☐ Disable all ranking attributes ⓘ

Example

Use Case

- I want my Sale PLP to give greater weight to products that are generating the most revenue across the entire site, not just the products that are selling well on the Sale PLP

Steps to Enable

- Add a new Category Ranking Rule
- Create a New Attribute Rule and Select 'Performance Attributes'
- Add Attribute rule for Sitewide Revenue and boost at desired strength
- View estimated impact
- Use Preview to review impact of rule

Product Grid

Category: Audience: ALL (Default) Duration: All Time

Product Attribute Insights

Create New Attribute Rule

Step 1 → Step 2

Choose Attribute Type

Numeric Attributes

Product Description Attributes

Performance Attributes

Step 1 → Step 2 → Step 3

Next

Performance Attributes

Select the performance attributes below to boost linearly (larger the value, stronger the impact).

☐ Sitewide ATC

☐ Sitewide conversion

☐ Sitewide revenue

☐ Sitewide view

Step 1 → Step 2 → Step 3

Done

Optionally disable the default BR ranking algorithm partially or fully. It will prevent double-counting of certain attributes in the default ranking. If disabled fully, the ranking will be determined solely based on performance attributes configured in the previous step.

CUSTOMIZE RANKING ALGORITHM

☒ Use global settings ⓘ

☐ Use default ranking ⓘ

☐ Disable performance attributes ⓘ

☐ Disable all ranking attributes ⓘ

Product Attribute Insights

Add Additional Attribute Rule

Est. impact for top 20 products

Sitewide revenue \$2,470,068.47 ↗

Performance Attribute Rules

Sitewide revenue

Total revenue generated by a product across the site over the past 30 days.

Operation Boost

Strength 25.0

Algo Weight Customization Best Practices

Best Practices

- **Create an A/B testing plan to find the optimal weights of Sitewide Signal boosts**
 - Keep in mind - what results are you looking to achieve? What does success look like? Determine your key KPIs for testing.
- **The best boosting strategy will likely be one that uses a combination of Sitewide Signals**
 - If you just boost on Sitewide Revenue - may see some very high priced items at the top of the result set
 - If you just boost on Sitewide CVR - may see Sale/Clearance items at the top of the result set
- **Keep in mind the potential impact to relevance - boosting on Sitewide Performance could have a negative impact on relevance, especially for queries with noise, so always test out your boosting strategy**
 - Ensure your site relevance is in a good state before adding Algo Weight Rules - consider a Searchable Field Audit with the Experience Strategy Group to find and fix sources of noise
 - Your site will likely benefit from different global rules for Search vs. Category, since Relevance isn't an issue with Category

Algo Weight Customization Conflict Resolution

Conflict Resolution

What happens if I create two rules that are in conflict with one another?

The algorithm will try to respect all rules when possible, but when there is direct conflict it will follow a standard order of operations to determine which rule to use. Conflict resolution will **always prioritize specificity** - so more local rules (like query/category specific rules) will take precedence when in conflict with a global rule

Scenario:

I have applied...

- "disable performance attributes" as a Global Search ranking rule
- "use default ranking" as a Global ranking rule
- "use global settings" as a local ranking rule for the categories "sofas" and "side tables".

What eventually will apply to the categories "sofas" and "side tables"?

Conflict resolution always prioritizes specificity and hence in this case local ranking rules will be applied for "sofas" and "side tables" i.e. "use global settings" will apply.

Now, since there are 2 global rules applied(Global search ranking rule and global ranking rule), again specificity will be given importance and Global search ranking rule will be applied i.e. "use default ranking".

A/B Testing Program Plan with the *Experience Strategy Group*

Program Plan Pillars of Success



Alignment on Success Metrics

The primary measure of success for optimizing products based on newness or discount % may not be conversion rate. For example, **the primary goal for driving sales of discounted products may be reaching a sell through target**



Analyze Impact Across the Site

Bloomreach can support in providing logs for a **full audit of test performance.** For example,, a boost on conversions may benefit certain departments and be a detriment to others.



Small changes to weightings can have a big impact

Test and retest a boost on the same attribute to determine the optimal weighting in combination with the BR algo. **Other customers are continuously testing varying boost levels for the same attribute**



Hypothesis Library

A library of activities that we can explore to learn more about our consumers who leverage search in the product discovery journey.

The Problem

What are we looking to solve?

What Learning Objectives Does This Support?

DELIVER BEST RETURN	IMPROVE SEARCH CONVERSIONS
IMPROVE SEARCH BOUNCE RATE	PERSONALIZATION
ELEVATE THE USER EXPERIENCE	IMPROVE STOCK/INVENTORY EXPERIENCE
IMPROVE ADD TO CART RATE	BOOST/BURY ACTIVITIES

Hypothesis

The hypothesis forms the backbone for the test. Through testing we will prove or disprove our assumptions

Learning Question:

Methodology

- Control: Default Algo:
- Test Bucket A:
- Test Bucket B:

Measurement

Primary KPI: What are the most important KPIs to solve this problem

Secondary KPI: What are metrics not as important to measure but are related and worth tracking

Other **measures of success:**

Dependencies/Risk

Anything happening that could have a positive or negative impact on this test



**These features are just
the beginning!**

**Look out for more AI &
Customization features to
come in 2024**

Q&A

Thank you



Let's keep the conversation going!

1. Join the **Digital Merchandising Community** on LinkedIn
2. Invite a friend..OR two to join the community and the events
3. If you are interested in being apart of a future panel, we would love to have you join one!

