



Transparency Note: Intelligent Recommendations Service (IR)

Last Updated 9/17/2021

Table of Contents

Transparency Note: Intelligent Recommendations Service (IR).....	1
What is a Transparency Note?	3
Basics of the Intelligent Recommendations service.....	3
Introduction to Intelligent Recommendations.....	3
Key terminology	3
Capabilities.....	4
System Behavior	4
Use cases	5
Intended use cases.....	5
Considerations when choosing other use cases	6
Limitations.....	6
Evaluating Intelligent Recommendations in your applications.....	7
Learn more about responsible AI.....	8
Learn more about the Intelligent Recommendations (IR) service.....	8
Contact us	8
About this document	8

What is a Transparency Note?

An AI (Artificial Intelligence) system includes not only technology, but also the people who will use it, the people who will be affected by it, and the environment in which it is deployed. Creating a system that is fit for its intended purpose requires an understanding of how the technology works, its capabilities, and limitations, and how to achieve the best performance. Microsoft's Transparency Notes are intended to help you understand how our AI technology works, the choices system owners can make that influence system performance and behavior, and the importance of thinking about the complete system, including the technology, the people, and the environment. You can use Transparency Notes when developing or deploying your own system or share them with the people who will use or be affected by your system. Transparency Notes are part of a broader effort at Microsoft to put our AI principles into practice. To find out more, see [Microsoft AI Principles](#).

Basics of the Intelligent Recommendations service

Introduction to Intelligent Recommendations

The Intelligent Recommendations service provides businesses with the ability to produce personalized, machine-learned product suggestions. User behavior inputs, such as transactions, downloads, or use actions, are used to generate recommendations. The service uses a variety of AI and machine learning (ML) algorithms to produce ranked sets of items that improve item discovery for users. These algorithms can predict user interests, connecting users to items that fit their unique taste.

Key terminology

Terminology	Definition
Recommendation List	A list in the context of Recommendations is the output from the service for a given scenario like "Picks for you." The service inputs both user behavior (such as transactions, downloads, page views, etc.) and/or item metadata (such as description, release date, etc.). This input data is then used to derive an ordered list of items from the Catalog.
Catalog	A Catalog refers to the full set of items that are possible to recommend and/or present in the input data. Each list type is tightly coupled with a scenario and has a corresponding algorithm or two algorithms that can serve that scenario.
Seed-item	A seed-item is the item ID of a product that is input into an algorithm to generate a recommendation list related to that product. For example, if a user viewed a handbag, the handbag's item ID would be the seed-item. A "People Also Viewed" algorithm would use the handbag's item ID to generate recommendations of products other users viewed in addition to the handbag.

Capabilities

System Behavior

The Intelligent Recommendations (IR) service relies on ML models built by a customer using defining user interaction signals and item metadata to produce a variety of recommendation lists. Examples of the relevant user interaction signals can include *transactional data*, *downloads*, *session ID* and *tracked interactions (user actions) between users and products/items*. These signals drive the recommendations lists which employ [Collaborative Filtering](#) ML techniques to generate both personalized recommendations ("Picks for you"), as well as other item-based, non-personal, item discovery, such as Trending Charts. Some examples include:

- *Discovery of items based on user interactions (purchase, views, ratings)*. Examples of this discovery type include:
 - "People also like" – People who liked this item also liked.
 - "People also buy" – People who purchased this item also purchased.
 - "Frequently bought together" – People who purchased this item most often purchased it along with.
 - "New" – Items ranked by most recent release date.
 - "Trending" – Items ranked by high sales and recent release date.
 - "Popular" – Items ranked by a selected customer volume input: purchase, download, views, etc.
- *Discovery of items based on a unique user's interactions (purchase, views, ratings)*. An example of discovery type:
 - "Picks for you" – Predicted item suggestions based on the individual user's past interactions (purchases, views, etc.).
- *Discovery of items based on the metadata of all the other items in the catalog*. Examples of this discovery type include:
 - "Shop by similar looks"
 - This works using product images from various angles, such as those often used on e-commerce product detail pages and shows other items with similar product images. The experience works well for areas and industries when the items have an intrinsic aesthetic appeal such as fashion, jewelry, or furniture.
 - "Shop by similar description"
 - This works using natural language processing (NLP) to process textual information such as the description or title. It works well for products with text descriptions, particularly if the defining and distinguishing traits of the product are not captured in product images such as articles, wine, or chocolates.

Additionally, item metadata is also used to enable certain other models and scenarios that use elements such as product photos or product descriptions, as well as ensure that proper filtering of items can be applied as needed on the results.

The algorithm options for several types of lists are provided out-of-the-box, and some recommendation scenarios come with a few controls to adapt their functionality or modify the results being returned. The combination of input and list logic chosen by a customer is largely what determines the output recommendation list. For example, a customer may choose to produce a "People also view" list by using page view information with the "People also" list logic.

The Intelligent Recommendations service also allows for customers to edit the computed lists: specific item(s) can be put at the top, or specific item(s) can be removed from a list, or the list can be fully replaced by specified item(s), all prior to the resultant list being returned from the service. Other types of recommendations provide non-personal product suggestions based on item-to-item relationships (product purchased together, for example), or based on global, aggregated user interactions.

Read more here: [What is Intelligent Recommendations? | Microsoft Docs](#)

Use cases

Intended use cases

Intelligent Recommendations can be used for a variety of discovery scenarios across a multitude of industries. Example use cases for the service include (but are not limited to):

- *Recommending items in an online store:* Customers can use Intelligent Recommendations to recommend items for users who are online shopping. In this scenario, Intelligent Recommendations can recommend visually similar or related items a user might like to purchase. Intelligent Recommendations can also make recommendations based on the user's past activities for personalized or session-based recommendations.
- *Recommending visually similar products based on image data:* Customers can recommend items (clothing, handbags, shoes, jewelry, etc.) using only item images. For example, a user may want to find a similar dress that has a gradient pattern. Intelligent Recommendations can view and compare all item images related to the seed-item and return similar gradient dresses in addition to cross-category suggestions like a gradient-colored handbag or shoes.
- *Recommending movies or multimedia content:* Customers can recommend multimedia content such as movies, music, or songs. Music recommendations might include personalized lists of what users may want to listen to next, ranked genre categories for Most Popular Songs, and more. Similarly, for video, customers can suggest trending media content, personalized watch lists, related videos based on what other users have also viewed/clicked-on, session-based video recommendations, and more.
- *Recommending documentation content:* Customers can recommend related documents based on what users may be reading or based on what users may have recently read. This is also great for products that rely on textual descriptions, which benefit from similarly described items, such as tasting notes of wine.
- *Recommending the next best action:* Customers can recommend repetitive tasks based on what users have been working on most recently to predict what a user may want to work on next, such as troubleshooting steps.
- *Recommending add-on content for video games and applications:* Customers can recommend a personalized list of relevant add-ons (downloadable content and characters, costumes, in-game currency, etc.) based on the purchase history of the user or based on what other users have purchased.

These are just a few of the many ways that Intelligent Recommendations can be applied to numerous discovery verticals. **Read more here:** [Intelligent Recommendations scenarios | Microsoft Docs](#)

Considerations when choosing other use cases

Do not use Intelligent Recommendations for decisions that may have serious adverse impacts

Intelligent Recommendations was not designed or tested to recommend items that require additional considerations related to accuracy, governance, policy, legal, or expert knowledge as these often exist outside the scope of the usage patterns carried out by regular (non-expert) users. Examples of such use cases include medical diagnostics, banking, or financial recommendations, hiring or job placement recommendations, or recommendations related to housing.

Limitations

Note about data types and their implications for the results:

Personalized recommendations

The “Picks for you” personalized recommendations list returns a unique list of items based on the *past activity or interactions of the specific user*. The system was built with the purpose of user interactions powering the models and suggestions to best capture the tastes or interests of individuals. Because of this, **it is suggested to not input user metadata (gender, race, demographic identifiers, etc.) into the model**.

The service performance measurements review, which is focused on scoring the overall performance of the recommendation models, showed higher relevance in the model’s results when using user interaction data instead of identifiable user metadata. In other words, the Intelligent Recommendations models produced relevant results for users without needing to know targeted information about the users that would otherwise categorize or identify them based on demographic identifiers. Results were checked for individual results in objective ways by using standard techniques to measure accuracy using a training subset, running predictions, and then measuring accuracy and precision to a test subset which were actual prior observations. Subjective measures with user feedback have also shown a distinct alignment towards individual actual taste, instead of popularity or demographic based assumptions.

Additionally, since the Intelligent Recommendations models rely heavily on inputs regarding how users *interact* with an experience (clicks, session id, purchase, downloads, views, etc.) it means that users without any past purchase, or interaction history, may have no or a limited number of product suggestions until they have some accumulated activity.

Empty or reduced number of recommendations results

In general, with Recommendations lists that use Collaborative Filtering AI-ML (Artificial Intelligence and Machine Learning) techniques (such as “People also like,” “Picks for you”) having sufficient data about prior observations is needed to generate relevant predictions. When there are few signals for items seen thus far, as can happen with very new or not immensely popular items, then those items may return short, irrelevant, or even zero results for “People also like” and would not show up in the “People also like” lists for other items, or “Picks for you” for users. When a given user has little to no past user interactions, they would also have empty, or less personalized results, until more signals are available.

Visual recommendations require multiple images

Visually enriched recommendations only require images. This means recommendations for a specific item can be quite different in many ways from the seed item (original item that is receiving recommendations), even if images are similar. For accurate results, the system needs enough specific images of customers' items. The system also requires that some of the items have more than one image per item. Customers with a catalog of items with aesthetic appeal with a normal expectation that product images represent the qualities that appeal to users, can run our visual recommendations algorithms to produce great similarity results, without the need for additional labelling or annotation tasks. Limitations for this type of recommendation are that items with no images or badly representative images can lower the relevance of recommendations results.

Textual recommendations

Textually enriched recommendations only need the textual data (title and description) of an item to start returning results. Textual data is used to power the "Shop by similar text," using NLP algorithms.

The Intelligent Recommendation service's NLP deep learning algorithm is a fine-grained approach for learning language models and scoring similarities that are hyper-focused on catalogs for exceptional text-based item recommendations. Customers with a catalog of items with the normal expectations of providing item content (e.g., title, description, document text) can run our textual recommendations algorithms to produce great similarity results, without the need for additional labelling or annotation tasks. Limitations for this type of recommendation are that items with no descriptions, poorly written or badly representative text can lower the relevance of recommendations results.

Evaluating Intelligent Recommendations in your applications

The performance of Intelligent Recommendations will vary depending on the real-world uses and conditions in which people use them. True usage by end-users can be a combination of many elements in addition to the recommendations algorithms, including, for example, the user interface placement and title of a recommendation result on page, product quality or pricing, user base overall preferences, and business or global trends.

To ensure optimal performance in their scenarios, customers should conduct their own evaluations of the solutions they implement using Intelligent Recommendations. Customers should generally follow an evaluation process that 1) uses some internal stakeholders to evaluate results, 2) uses A/B experimentation to rollout Intelligent Recommendations to end-users, 3) incorporates KPI (Key Performance Indicators) and metrics monitoring when deploying the service in experiences for the first time, and 4) tests and tweaks the Intelligent Recommendations configuration including the surrounding experiences such as user interface placement or business processes.

Learn more about responsible AI

[Microsoft AI principles](#)

[Microsoft responsible AI resources](#)

[Microsoft Azure Learning courses on responsible AI](#)

Learn more about the Intelligent Recommendations (IR) service

[What is Intelligent Recommendations? | Microsoft Docs](#)

[Intelligent Recommendations architecture | Microsoft Docs](#)

[Intelligent Recommendations scenarios | Microsoft Docs](#)

[Intelligent Recommendations REST API reference? - Industry Intelligent-recommendations Rest API | Microsoft Docs](#)

[Data contract reference | Microsoft Docs](#)

Contact us

[Give us feedback on this document](#)

About this document

© 2021 Microsoft Corporation. All rights reserved. This document is provided "as-is" and for informational purposes only. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it. Some examples are for illustration only and are fictitious. No real association is intended or inferred.

Published: September 17, 2021

Last updated: September 17, 2021