

Technical Disclosure Commons

Defensive Publications Series

July 2020

Refining Shopping Search Queries and Results Based On Inferred Product Purchases

Matthew Sharifi

Victor Carbune

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

Sharifi, Matthew and Carbune, Victor, "Refining Shopping Search Queries and Results Based On Inferred Product Purchases", Technical Disclosure Commons, (July 02, 2020)

https://www.tdcommons.org/dpubs_series/3400



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

Refining Shopping Search Queries and Results Based On Inferred Product Purchases

ABSTRACT

Knowing a user's purchase history can be helpful in refining, rewriting, or correcting a shopping related search query the user enters since users often purchase products that are related to or same as products bought previously. However, the text input in searches for products is often short, ambiguous, or underspecified. Even in cases where a user includes relevant detail in a query, manual formulation of the search terms, often from memory, can result in errors. Current shopping search mechanisms do not help detect and refine underspecified or erroneous queries. This disclosure describes techniques to improve keyword-based product search by appropriate query reformulation and search result annotation. With user permission, the reformulation and annotation is based on the user's inferred product purchase history.

KEYWORDS

- Online shopping
- Product search
- Purchase history
- Shopping history
- Shopping aggregator
- Query refinement
- Query suggestion
- Query correction
- Search result annotation
- Product compatibility

BACKGROUND

Shopping aggregators allow people to use keyword search to look for products across a range of merchants. As with general web search, the text input in searches for products is often short, ambiguous, or underspecified. For instance, users may search for “printer cartridges” without specifying that they are looking for cartridges for their recently purchased printer. Knowing a user’s purchase history can be helpful in refining, rewriting, or correcting the search query the user enters since users often wish to purchase products that are related to or same as products bought previously. Even in cases where a user includes relevant detail in a query, manual formulation of the search terms, often from memory, can result in errors. For example, a user may search for “laser printer cartridges” along with a printer brand and model even though the user owns an inkjet printer of the manufacturer, not a laser printer.

Current shopping aggregators do not provide any mechanisms to help detect and refine underspecified or erroneous queries. As a result, the search results presented by such sites lack relevant cues that can guide users toward the most likely product being sought by the query.

DESCRIPTION

This disclosure describes techniques to improve keyword-based product search by appropriate reformulation of a user’s search queries. With user permission, the reformulation of the user’s query is based on the user’s product purchase history. For instance, if the user searches for “printer cartridges” the query can be modified to include the make, model, and type of the printer that the user previously purchased. Alternatively, with the user’s permission, the user can be offered query suggestions to correct the query to one that includes the printer make, model, and type.

Query suggestions can be presented at the top of the search results to ask if the user meant to issue the more expanded query. In addition, the list of products retrieved in response to the query can be annotated to denote whether the product corresponding to each search result is compatible with the user's prior purchases. For instance, the annotation can include information such as "These cartridges are not compatible with the printer you recently purchased."

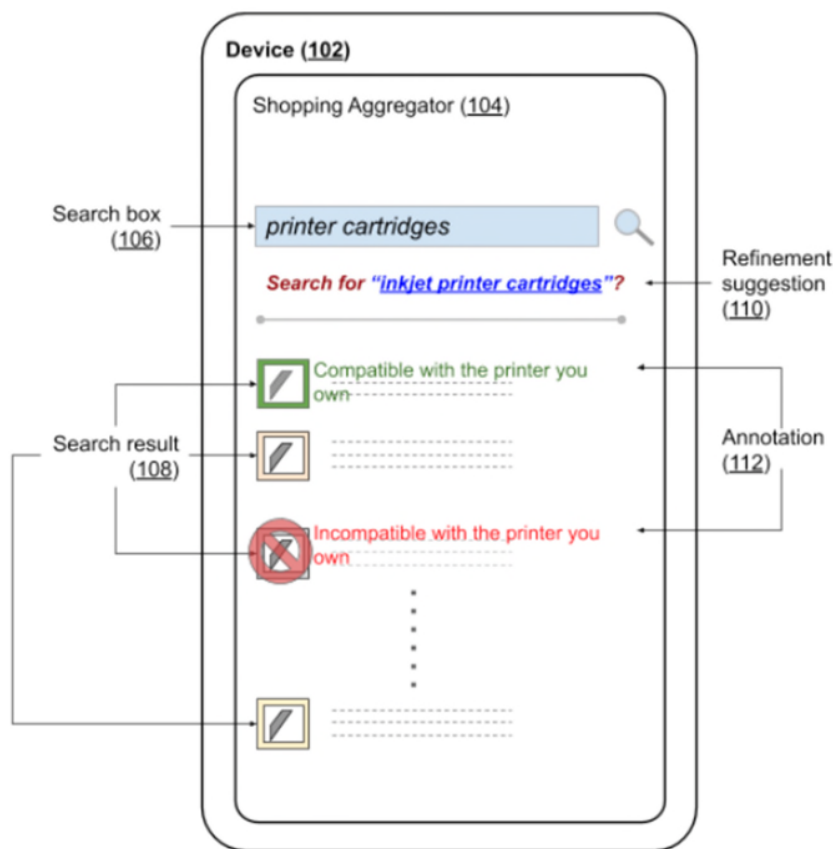


Fig. 1: Query refinement and result annotation based on products owned by a user

Fig. 1 shows an operational implementation of the techniques described in this disclosure. A user searches for "printer cartridges" using the search box (106) of a shopping aggregator (104) accessed via a device (102). Based on the user's prior purchase of an inkjet printer, the user is offered a refinement suggestion (110) to make the query more specific. In addition, the

list of products listed in the search results (108) is marked with annotations (112) that indicate whether the item is compatible with the user's previously purchased inkjet printer.

A person can purchase products from a variety of sources such as online stores, offline stores, shopping aggregator referrals, etc. The techniques described in this disclosure rely on inferring a complete picture of a user's purchases across the various sources with the user's permission in one or more of the following ways:

- **Interaction with a shopping aggregator site:** If the user permits, conversion detection mechanisms of a shopping aggregator can explicitly record when a user purchases a product from a merchant after clicking on the corresponding link within the search results presented by the shopping aggregator.
- **Analysis of web search terms:** Examining the terms in a user's web search queries can indicate whether the user owns a particular product. For instance, a query that includes a specific printer make and model can imply ownership if it contains the term "calibration" or can indicate lack of ownership if issued with the term "reviews." Similarly, queries for items related to a product can imply ownership of the product. For example, searches for specific printer ink cartridges suggest that the user owns the corresponding printer. With user permission, a trained machine learning classifier is employed to determine whether a given query implies product ownership. If the model indicates product ownership, the query is further analyzed to extract the name of the product(s) that the user likely owns.
- **Installed apps:** If the user permits, apps installed or searched for by the user can be checked for those that are connected to specific products, thus signifying ownership of the corresponding items. For instance, printer drivers installed on a device can be indicative of

ownership of specific printer models, apps for smart home devices can indicate that the user has the specific device in their home, etc.

- **Content within specific apps:** With user permission, product ownership can be inferred from the content of various apps, such as email, shopping, etc. If the user permits, an on-device trained machine learning classifier can process the content within relevant apps to detect relevant information about product purchases. For instance, an email receipt for the purchase of a product is an indicator of ownership.

With user permission, inferences from each of the above mechanisms can be aggregated across one or more queries to determine likely product ownership with higher confidence.

To incorporate product ownership inferences in the user's shopping search, the user's query is first associated with a product or a product type. Such associating can be done by analyzing the query terms and/or examining the results returned in response to the query. For instance, a query for "printer cartridges" typically returns results that mention a variety of printer models, thus signifying that it is related to printers.

After associating a query with a product or a product type, the inferred product ownership information can be used to take appropriate actions as permitted by the user. In cases where the ownership of a product is inferred with high confidence, the original query can be reissued by automatically inserting relevant information about the product, such as make, model, product type, etc. When it is less certain that the user owns a product, addition of the product information to the query can be offered as a refinement suggestion. When the user's query contradicts inferred product ownership, the query refinement suggestion can be worded to request the user to double check whether the user meant to inquire about the owned product instead. For instance, if

the user searches for “cartridges for a laser printer” then the user can be asked: “Did you mean ‘cartridges for an inkjet printer?’”

In addition to query refinement, the inferred product ownership can be employed to annotate individual search results. The annotations can indicate compatibility between products in the results list and those likely owned by the user. The compatibility can be derived from examining relevant product information, such as title, description, metadata, etc. Compatibility matching can be performed via the same mechanism used to extract the product or product type mentioned in user queries. Matches and mismatches can then be appropriately highlighted in the shopping search results, e.g., as shown in Fig. 1. For instance, a product in the results list can be marked with annotations such as “These cartridges are compatible with your inkjet printer.”

Each product inferred to be owned by the user can be associated with a score indicating the likelihood of the accuracy of the inference based on one or more of the various factors described above. With user permission, the score can be refined based on additional factors such as purchase date. Further, it can be checked if the user still owns the product by employing similar mechanisms as above to determine if the user sold or otherwise disposed of the product.

With user permission, the techniques described above can support queries related to any product type. For example, queries regarding car accessories can be construed as referring to a car owned by the user, searches for console games can be taken to indicate interest in those that can run on the user’s particular gaming console, etc. Further, the techniques can handle queries that may be applicable to multiple products. For instance, a query for a “filter” can be refined to apply to a vacuum cleaner or an air purifier, depending on what the user owns.

If the user permits, the described techniques can also offer relevant explanations pertaining to the query refinement suggestions and product annotations. The explanations can be

based on the underlying factors. For instance, such an explanation can state: “You may need inkjet printer cartridges because you purchased an inkjet printer last year” or “You are being shown filters for a vacuum cleaner because you previously searched for vacuum cleaners.”

The product ownership classifier can employ any suitable trained machine learning model, such as a neural network. The model can be trained on known positive and negative query examples, obtained with user permission. Positive training examples can be obtained from cases where product ownership is known with certainty, such as via merchant information about conversions of clicks into subsequent purchases. Further, users can be offered the ability to inspect the inferred product purchases and indicate any incorrect inferences. If the user permits, such user input can provide additional training information for the machine learning model.

Similarly, trained machine learning models can be used for other components, such as extracting product information from queries and matching product compatibility for search result annotation. The threshold values for confidence used in the operation of the various trained machine learning models can be set by the developers and/or specified by the user and/or determined dynamically at runtime.

When the user has an account with a service provider such as a shopping aggregator, the various pieces of information used to infer product ownership along with the derived inferences can be stored in such an account, if permitted by the user. In case the entity that provides the shopping aggregator service also provides other services (e.g., web search, app search, etc.), user permission can be obtained to utilize such data to infer product purchases.

The techniques described in this disclosure can be implemented by any entity that supports keyword-based searches for product purchases, such as shopping aggregators that search for products across merchants, merchants that partner with third-party sellers and other

affiliate partners, etc. In addition, an application programming interface (API) can be provided to access the inferred purchases such that it can be used by third parties, if such third-party access is permitted by the user. For instance, if the user permits, inferred purchases accessed via the API can be used to automatically select appropriate compatibility filters on the site or app for a third-party store.

Implementation of the techniques described in this disclosure can improve the relevance of searches related to shopping and enhance the user experience (UX) of the product selection and purchase process. In addition, the techniques can help users avoid non-compatible or unneeded purchases, thus saving the time, effort, and costs of product return and replacement.

Further to the descriptions above, a user may be provided with controls allowing the user to make an election as to both if and when systems, programs or features described herein may enable collection of user information (e.g., information about a user's purchases, searches performed, a user's preferences), and if the user is sent content or communications from a server. In addition, certain data may be treated in one or more ways before it is stored or used, so that personally identifiable information is removed. For example, a user's identity may be treated so that no personally identifiable information can be determined for the user, or a user's geographic location may be generalized where location information is obtained (such as to a city, ZIP code, or state level), so that a particular location of a user cannot be determined. Thus, the user may have control over what information is collected about the user, how that information is used, and what information is provided to the user.

CONCLUSION

This disclosure describes techniques to improve keyword-based product search by appropriate query reformulation and search result annotation. The reformulation and annotation

is based on the user's inferred product purchase history. The techniques described in this disclosure can infer purchases across various shopping sources. The inferred product ownership information can be used for query refinement, suggestion, or correction based on associating the user's query associated with a product or a product type. Further, the inferred product ownership can be employed to annotate individual search results to indicate compatibility with previously purchased products. The described techniques can support queries related to any product type and can be implemented by any entity that includes keyword searches for product purchases. Implementation of the techniques can improve the relevance of shopping related searches and enhance the user experience (UX) of the product selection and purchase process.