



Yahoo! Research Webscope Data Sets

Thank you for your interest in the Yahoo! Research Webscope Program. Below is a catalogue of the data sets available. To request a data set, please send an email to research-data-requests@yahoo-inc.com.

Academic Researchers are required to sign the Webscope Data Sharing Agreement to access the data sets. Researchers will receive the agreement that they will need to complete and sign. Executed agreements should be returned to lucky013@yahoo-inc.com or faxed to 408-349-2270. The researcher will receive a copy of the data on DVD in the mail.

Yahoo! Search Marketing Advertiser Bidding Data, version 1.0

Yahoo! Search Marketing operates Yahoo!'s auction-based platform for selling advertising space next to Yahoo! Search results. Advertisers bid for the right to appear alongside the results of particular search queries. For example, a travel vendor might bid for the right to appear alongside the results of the search query "Las Vegas travel". An advertiser's bid is the price the advertiser is willing to pay whenever a user actually clicks on their ad. Yahoo! Search Marketing auctions are continuous and dynamic.

Advertisers may alter their bids at any time, for example raising their bid for the query "buy flowers" during the week before Valentines Day. This dataset contains the bids over time of all advertisers participating in Yahoo! Search Marketing auctions for the top 1000 search queries during the period from June 15, 2002, to June 14, 2003. Advertisers' identities and query phrases are represented as meaningless anonymous numbers so that no identifying information about advertisers is revealed. The data may be used by economists or other researchers to investigate the behavior of bidders in this unique real-time auction format, responsible for roughly two billion dollars in revenue in 2005 and growing.

Yahoo! Search Marketing Advertiser-Phrase Bipartite Graph, version 1.0

Yahoo! Search Marketing operates Yahoo!'s auction-based platform for selling advertising space next to Yahoo! Search results. Advertisers bid for the right to appear alongside the results of particular search queries. For example, a travel vendor might bid for the right to appear alongside the results of the search query "Las Vegas travel." An advertiser's bid is the price the advertiser is willing to pay whenever a user actually clicks on their ad. Yahoo! Search Marketing auctions are continuous and dynamic: advertisers may alter their bids at any time, for example raising their bid for the query "buy flowers" during the week before Valentines Day.

This is a small completely anonymized graph reflecting the pattern of connectivity between some Yahoo! Search Marketing advertisers and some of the search keyword phrases that they bid on. Both advertisers and keyword phrases are represented as meaningless anonymous numbers so that no identifying information is revealed. The dataset contains 459,678 anonymous phrases ids, 193,582 anonymous advertiser ids, and 2,278,448 edges, representing the act of an advertiser bidding on a phrase.

Yahoo! Music User Ratings of Musical Artists, version 1.0

This dataset represents a snapshot of the Yahoo! Music community's preferences for various musical artists. The dataset contains over ten million ratings of musical artists given by Yahoo! Music users over the course of a one month period sometime prior to March 2004. Users are represented as meaningless anonymous numbers so that no identifying information is revealed. The dataset may be used by researchers to validate recommender systems or collaborative filtering algorithms. The dataset may serve as a testbed for matrix and graph algorithms including PCA and clustering algorithms.

Yahoo! Music User Ratings of Songs with Artist, Album, and Genre Meta Information, version 1.0

This dataset represents a snapshot of the Yahoo! Music community's preferences for various songs. The dataset contains over 717 million ratings of 136 thousand songs given by 1.8 million users of Yahoo! Music services.

The data was collected between 2002 and 2006. Each song in the data set is accompanied by artist, album, and genre attributes. The users, songs, artists, and albums are represented by randomly assigned numeric id's so that no identifying information is revealed. The mapping from genre id's to genre, as well as the genre hierarchy, is given.

Yahoo! Music ratings for User Selected and Randomly Selected songs, version 1.0

This dataset contains ratings for songs collected from two different sources. The first source consists of ratings supplied by users during normal interaction with Yahoo! Music services. The second source consists of ratings for randomly selected songs collected during an online survey conducted by Yahoo! Research. The rating data includes 15,400 users, and 1000 songs. The data contains at least ten ratings collected during normal use of Yahoo! Music services for each user, and exactly ten ratings for randomly selected songs for each of the first 5400 users in the data set. The data set includes approximately 300,000 user supplied ratings, and exactly 54,000 ratings for randomly selected songs. All users and items are represented by randomly assigned numeric identification numbers. In addition, the data set includes responses to seven multiple-choice survey questions regarding rating-behavior for each of the first 5400 users. The survey data and ratings for randomly selected songs were collected between August 22, 2006 and September 7, 2006. The normal-interaction data was collected between 2002 and 2006.

Yahoo! Movies User Ratings of Movies, and Movie Descriptive Content Information, version 1.0

This dataset contains a small sample of the Yahoo! Movies community's preferences for various movies, rated on a scale from A+ to F. Users are represented as meaningless anonymous numbers so that no identifying information is revealed. The dataset also contains a large amount of descriptive information about many movies released prior to November 2003, including cast, crew, synopsis, genre, average ratings, awards, etc. The dataset may be used by researchers to validate recommender systems or collaborative filtering algorithms, including hybrid content and collaborative filtering algorithms. The dataset may serve as a testbed for relational learning and data mining algorithms as well as matrix and graph algorithms including PCA and clustering algorithms.

Yahoo! Altavista Web Page Hyperlink Connectivity Graph, circa 2002

This dataset contains URLs and hyperlinks for over 1.4 billion public web pages indexed by the Yahoo! Altavista search engine in 2002. The dataset encodes the graph or map of links among web pages, where nodes in the graph are URLs. The Yahoo! Altavista web graph is an example of a large real-world graph. The dataset may serve as a testbed for matrix, graph, clustering, data mining, and machine learning algorithms.

Yahoo! Buzz Game Transactions with Buzz Scores, version 1.0

The Yahoo!/O'Reilly Tech Buzz Game tests the theory that a free electronic market can predict trends in technology. In the game, users buy and sell fantasy "stocks" in various technologies. The prices of the stocks fluctuate according to supply and demand using a market mechanism invented at Yahoo! called a dynamic pari-mutuel market. Weekly dividends are paid out to stockholders based on the current "buzz score", or percentage of Yahoo! searches associated with each technology. This dataset shows all the transactions over the course of the Tech Buzz Game, divided into two periods. The first period spans April 1, 2005 to July 31, 2005, after which all markets were closed and cashed out according to final buzz scores at that time. The second period spans August 22, 2005 to the present, at the beginning of which new markets were established and the market reset to a starting point. Traders are represented as meaningless anonymous numbers so that no identifying information is revealed. In addition, the dataset contains buzz score data showing the daily percentages of search volume for each stock. Researchers may use the data to test the predictive value of the market or to test market behavioral theories.

Yahoo! Groups User-Group Membership Bipartite Graph, version 1.0

Millions of communities and groups use Yahoo! Groups as a meeting place and forum to discuss mutual interests on nearly any topic. This dataset contains a sample of the "membership graph" of Yahoo! Groups, where both users and groups are represented as meaningless anonymous numbers so that no identifying information is revealed. Users and groups are nodes in the membership graph, with edges indicating that a user is a member of a group. The dataset consists only of the anonymous bipartite membership graph, and does not contain any information about users, groups, or discussions. The Yahoo! Groups membership graph is an example of a large real-world power-law graph. The dataset may serve as a testbed for matrix and graph algorithms including PCA and graph clustering algorithms, as well as machine learning algorithms.

Yahoo! N-Grams, version 2.0

This data set contains n-grams (contiguous words of length n), $n = 1$ to 5, extracted from a corpus of 14.6 million documents (126 million unique sentences, 3.4 billion running words) crawled from over 12000 news-oriented sites. The documents were published on these sites between February 2006 and December 2006. The dataset does not contain the documents themselves, but only the n-grams which occur at least twice. It provides statistics such as frequency of occurrence, number and entropy of different left (right) single-token contexts of each n-gram. This dataset may be used by researchers to build statistical language models for speech or handwriting recognition or machine translation.

Yahoo! Network Flows Data, version 1.0

Yahoo! network flows data contains communication patterns between end users in the large Internet and Yahoo servers. A netflow record includes timestamp, source IP address, destination IP address, source port, destination port, protocol, number of packets, and number of bytes transferred from the source to the destination. The record does not include the content of the data communication.

Each Netflow data file consists of sampled netflow records exported from routers in 15-minute

intervals. The dataset includes netflow data files collected from three border routers in October 11 2007. All IP addresses in the dataset are anonymized using a random permutation algorithm. **Note:** due to the size of this dataset, we can provide five dual layer DVDs or ten 4.7 GB DVDs or you can send a mini portable hard drive that we'll burn the data onto and return to you.

Metadata Extracted from Publicly Available Web Pages, version 1.0

The dataset contains about 100 million triples of RDF data obtained by extracting metadata from publicly available webpages. Three forms of embedded metadata are extracted: microformats (hCard, hCalendar and hReview), RDFa metadata and RDF documents linked to webpages. All metadata extracted from a webpage is converted to RDF. The data is made available in the WARC format, version 0.9.

The dataset may serve as a testbed for research in scalability in the Semantic Web area and also for developing methods to deal with metadata that is incomplete, erroneous or biased in some way.

Yahoo! Semantically Annotated Snapshot of the English Wikipedia, version 1.0

This SW1 dataset contains a snapshot of the English Wikipedia dated from 2006-11-04 processed with a number of public-available NLP tools. In order to build SW1, we started from the XML-ized wikipedia dump distributed by the University of Amsterdam. This snapshot of the English wikipedia contains 1,490,688 entries (excluding redirects). First, the text is extracted from the XML entry and split into sentences using simple heuristics. Then we ran several syntactic and semantic NLP taggers on it and collected their output. Raw Data (Multitag format)

The multitag format contains all the Wikipedia text plus all the semantic tags. All other data files can be reconstructed from this. A multitag file contains several Wikipedia entries. The Wikipedia snapshot was cut into 3000 multitag files each containing roughly 500 entries.

Yahoo! Answers Manners Questions, version 1.0

Yahoo! Answers is a web site where people post questions and answers, all of which are public to any web user willing to browse or download them. The data we have collected is a subset of the Yahoo! Answers corpus from a 10/25/2007 dump. It is a small subset of the questions, selected for their linguistic properties (for example they all start with "how {to|do|did|does|can|would|could|should}"). Additionally, we removed questions and answers of obvious low quality, i.e., we kept only questions and answers that have at least four words, out of which at least one is a noun and at least one is a verb. The final subset contains 142,627 questions and their answers.

In addition of question and answer text, the corpus contains a small amount of meta data, i.e., which answer was selected as the best answer, and the category and sub-category that was assigned to this question. No personal information is included in the corpus. The question URIs were replaced with locally-generated identifiers.

This dataset may be used by researchers to learn and validate answer extraction models for manner questions. An example of such work was published by Surdeanu et al. (2008).