DATA SAUDI
بيانات السعودية

وزارة الاقتصاد والتخطيط
MINISTRY OF ECONOMY & PLANNING

VIZBUILDER

USER GUIDE

March 2024

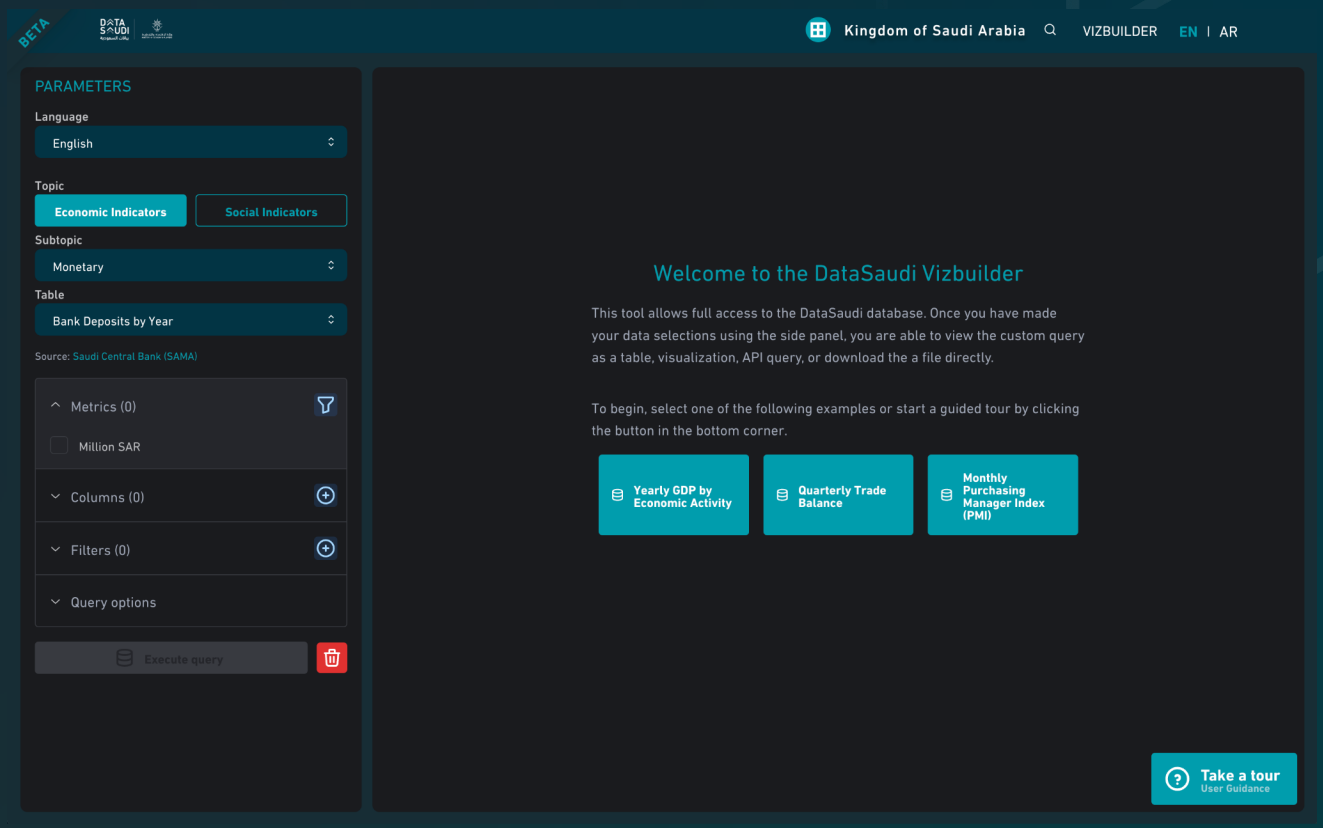

The Vizbuilder is an advanced data exploration tool that allows users to query and visualize all of the available datasets ingested into the DataSaudi database. While the DataSaudi narrative reports function as an introductory guide to selected indicators across various tables, the Vizbuilder enables a transparent view into all of the underlying tables in the database. This allows data savvy users to create custom filters and aggregations of data that may not be currently featured across the various reports.

Building a Query 3
- Data Output Language 4
- Selecting a Dataset 5
- Choosing Metrics 6
- Choosing Columns 7
- Column Filters 8
- Executing the Query 9
- Downloading the Data 9

Viewing Results 10

Data Table 10
- Sorting Data 10
- Searching Data 11
- Filtering Columns 11
- Showing/Hiding Columns 12
- Fullscreen Table 12

Visualizations 13
- Enlarging a Visualization 14
- Downloading a Visualization 14

Raw Response 15
- LogicLayer API URL 15
- Javascript Source 16
- Response Headers 16

Interactive User Guide 17

Starting Examples 17

Building a Query

- The Vizbuilder interface is divided into 2 primary sections:
- **Parameters:** The available parameters for building a data query.
- **Results:** The resulting data from executing the selected parameters (initially displays a welcome screen when no data is selected).

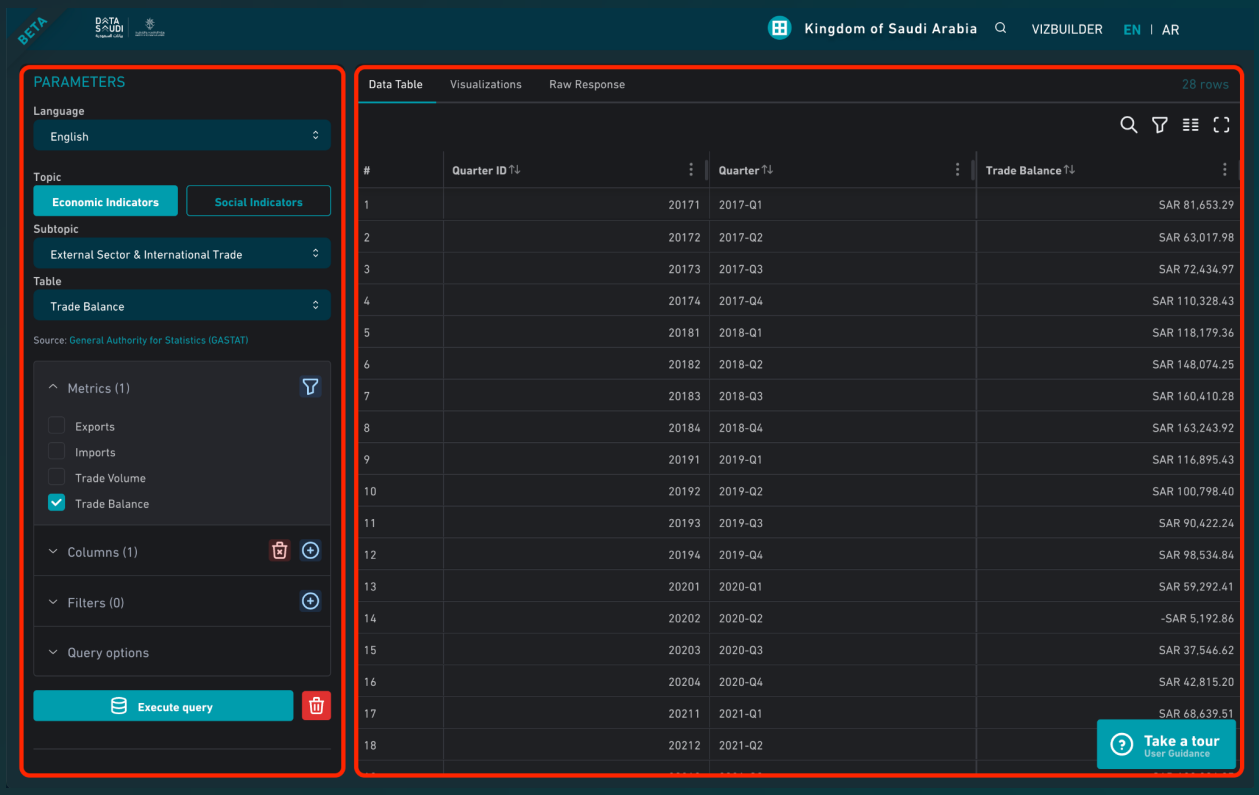

A query showing user selected Parameters (left) and the resulting Data Table (right).

Each query has many required and non-required parameters available, and the following section outlines each option in detail.

Data Output Language

DataSaudi makes every dataset available in both Arabic and English. This allows users viewing the site in one language to query the data in the other language. For example, a user exploring the site and interface in Arabic is able to load data with English names by selecting “English” from this dropdown.

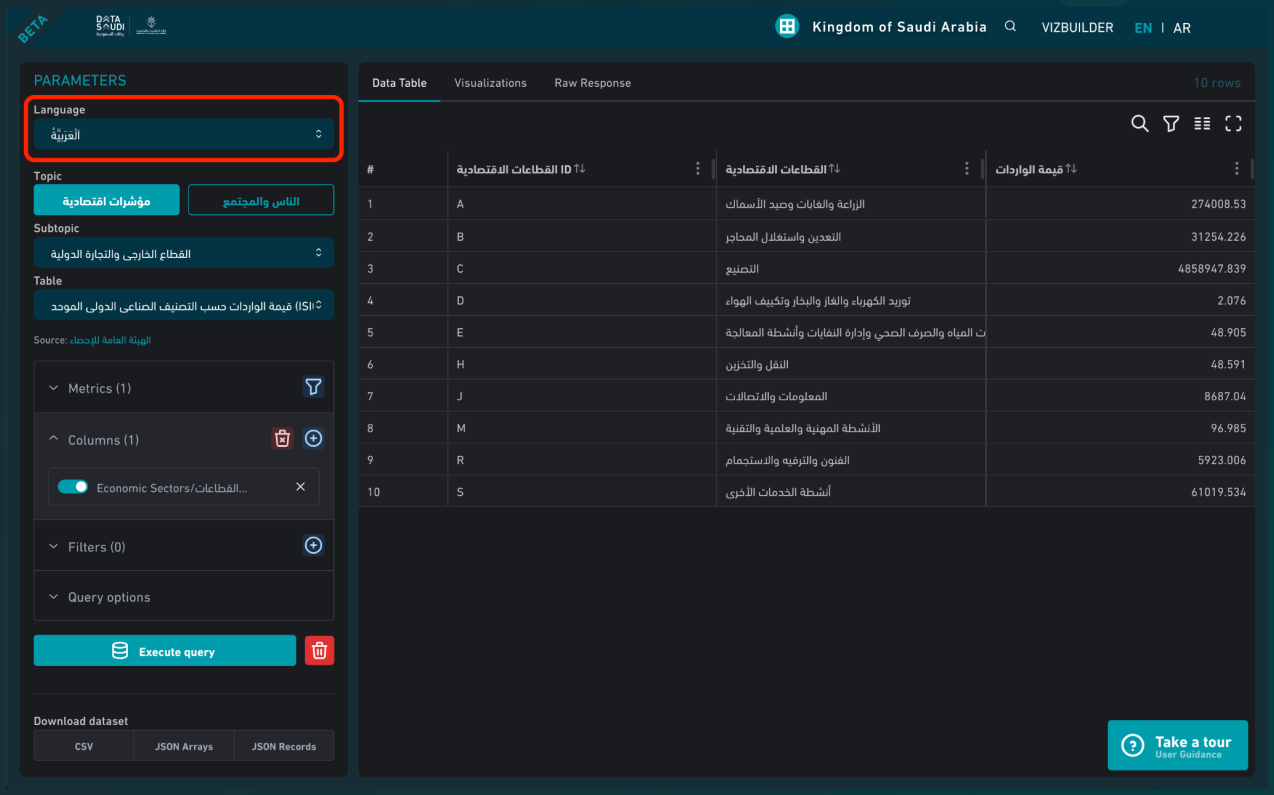

A query showing Arabic data being explored in the English side of DataSaudi.

Selecting a Dataset

Datasets are organized by the same groupings and labels used in the report navigation. As an example, the Gross Domestic Product dataset is located under the “Economic Indicators” topic within the “Real Sector Indicators” subtopic (seen below):

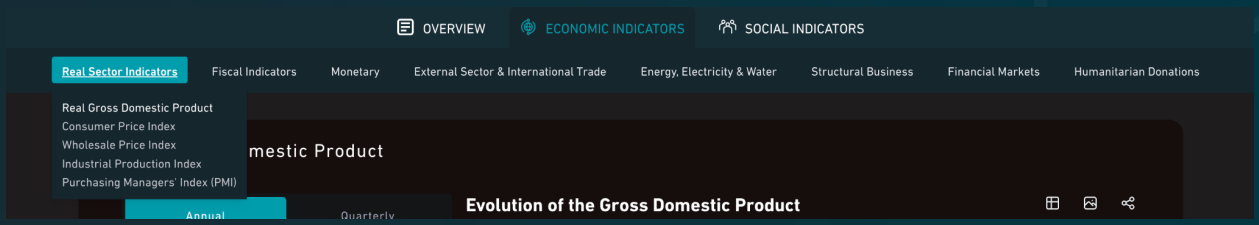

Gross Domestic Product as found within the Kingdom of Saudi Arabia report.

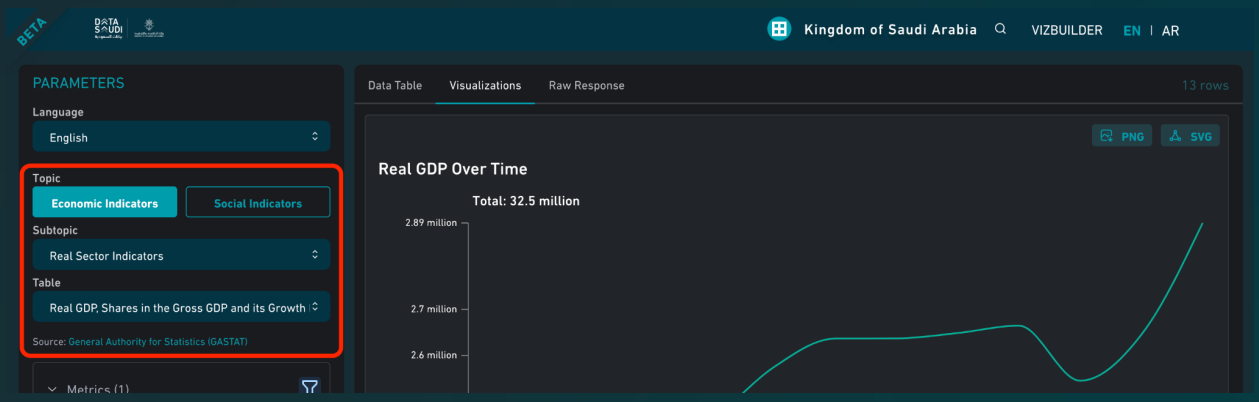

Gross Domestic Product as found within the Vizbuilder, using the same groupings.

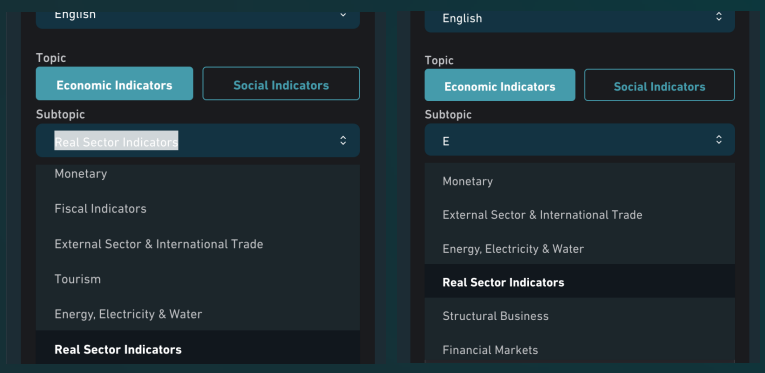

Using search inside of the Subtopic dropdown.

Choosing Metrics

Note: Users must select at least one Metric to make a query.

Once a specific data table has been selected, users are presented with all of the available numerical measures in the "Metrics" panel. Metrics selected here will be presented as columns in the resulting data table.

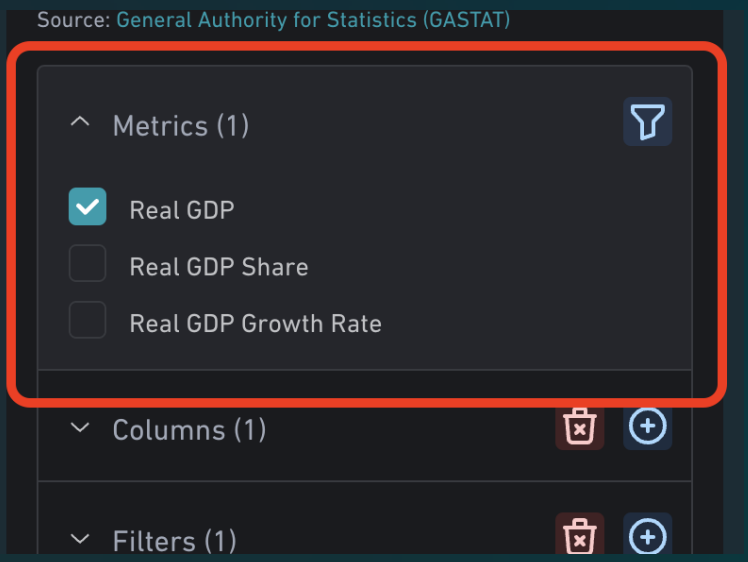

Selecting the first of 3 available Metrics.

The DataSaudi database was constructed using a Relational OLAP (ROLAP) architecture, which allows storing data at its more granular level and allowing the query interface to perform data aggregations on demand. This enables users to create custom queries that may not be available in the source data by aggregating metric values based on the user's selected Columns (see the "Choosing Columns" of this document).

Choosing Columns

Note: Users must select at least one Column to make a query.

Each data table in DataSaudi can have many different non-numeric Columns used to organize and filter the resulting data table. The more columns that a user selects, the deeper and more granular the resulting data will be.

Additionally, each Column has the potential to have multiple different levels within itself, like “Year” and “Quarter” inside of a “Time” dimension or a “Geography” dimension that contains both “Continent” and “Country”. Only 1 level of a specific dimension can be selected for a query, and the depth of the level determines how the data will be aggregated (for example, selecting a “Continent” level will aggregate the numeric values for all Countries within each Continent).

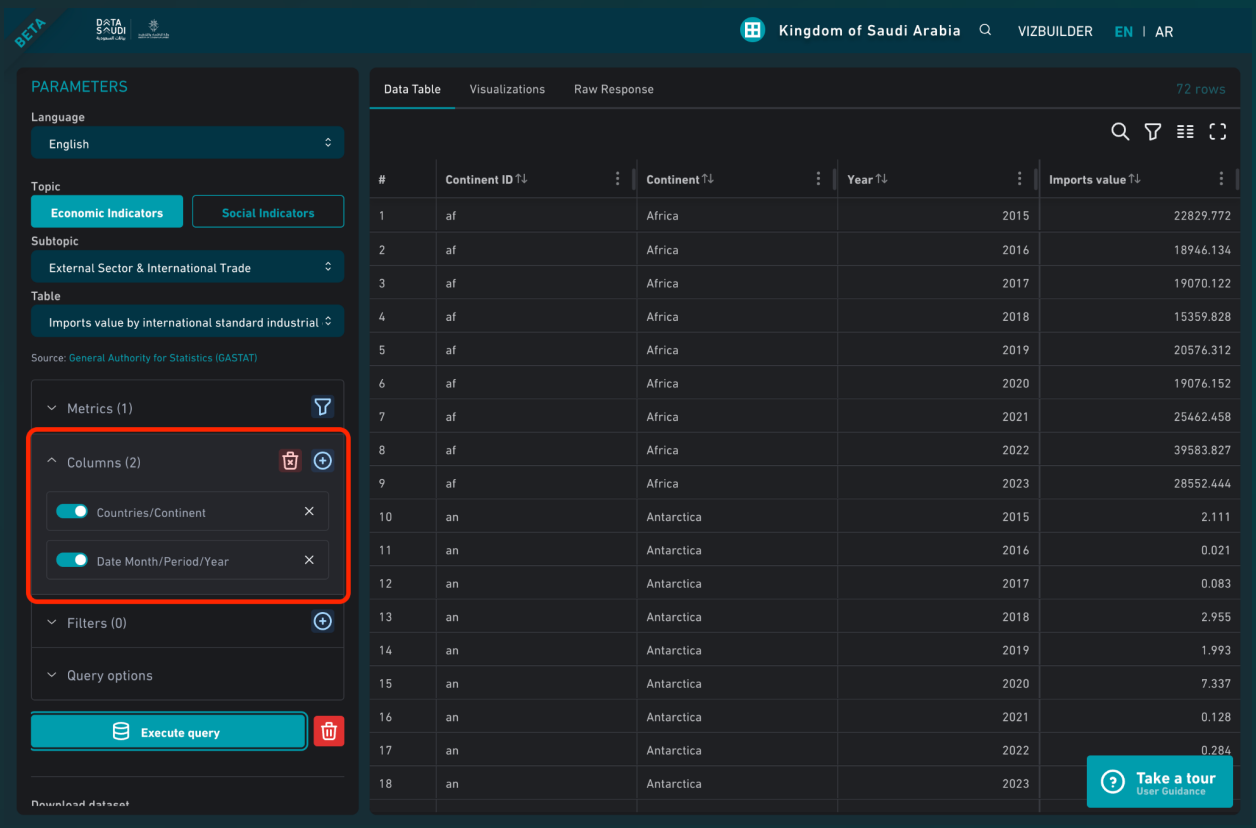

Continent values being aggregated for the Imports value metric.

Column Filters

The “Filters” panel allows users to only display the members of a Column that they are interested in. For example, a user could filter a “Year” column by only the most recent year, or a “Sector” column by only the pertinent sector(s) for their research.

Once a filter has been added for a specific column, clicking on the name of the column inside of the Filter interface will present the user with a pop-up modal window showing all of the available members of that column. If users want to see data for all of the available members of a column, there is no need to apply a filter to that column.

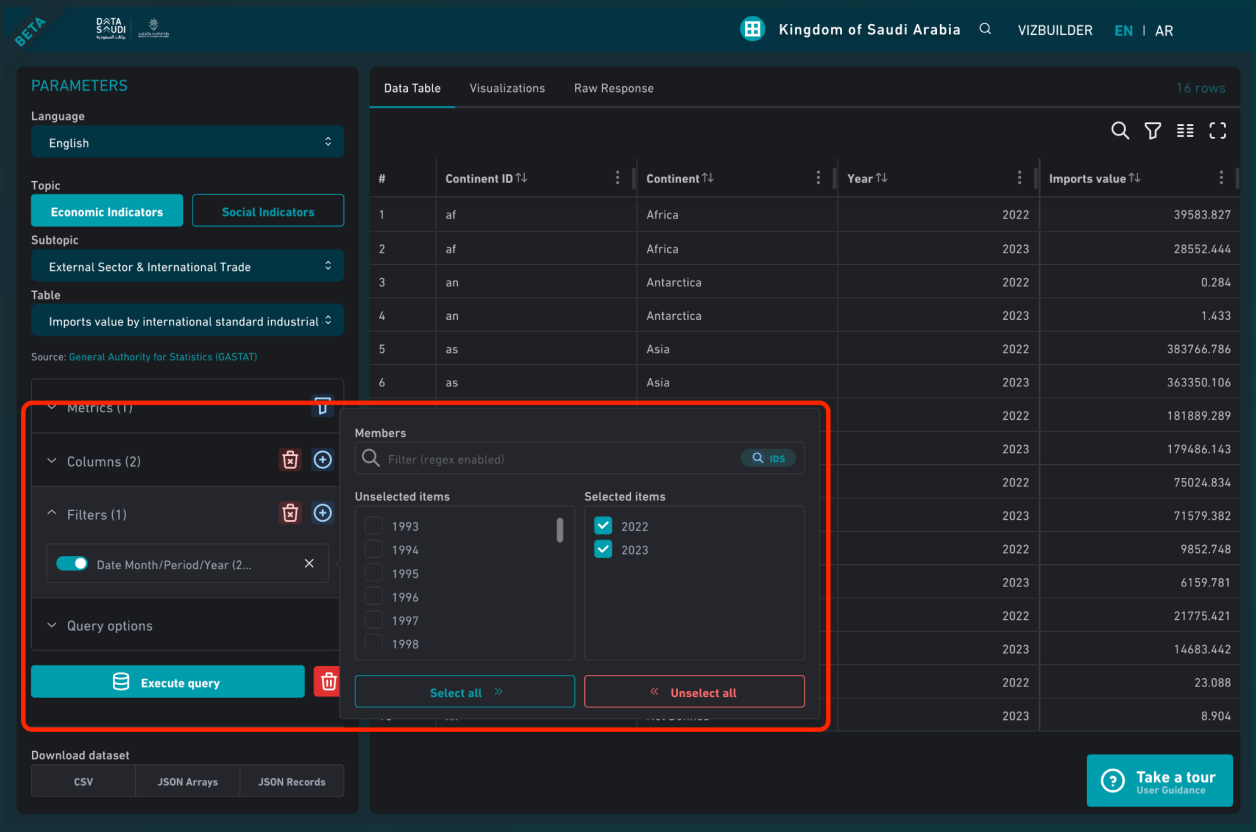

A “Year” column being filtered to show only 2022 and 2023.

Executing the Query

Once the desired data table, metrics, columns, and filters have been selected, users need to press the “Execute query” button to request the data from the database. Once the server responds with the custom data query, its contents are displayed in the tabbed interface next to the query Parameters, replacing the welcome text (or any previously queried data).

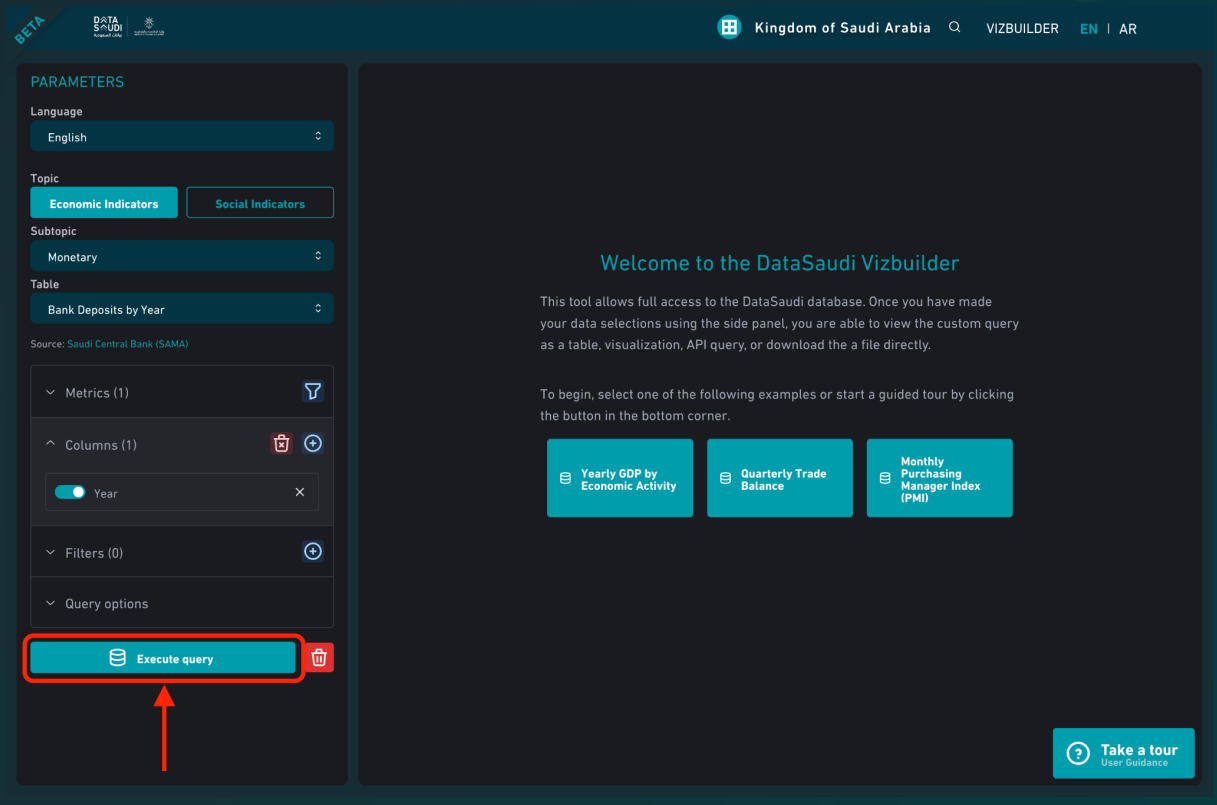

The “Execute query” button ready to be clicked after making parameter selections.

Downloading the Data

At any point after a query has been executed, users can immediately download the data in both CSV and JSON formats. These buttons are located just below the Execute query button at the bottom of the Parameters panel.

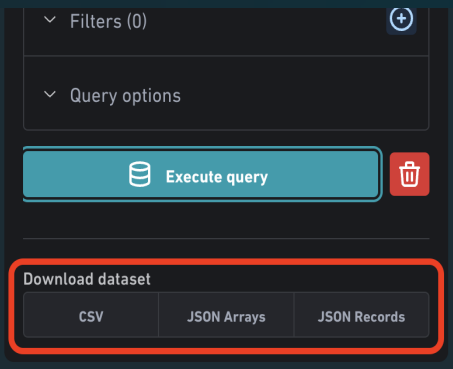

Download dataset buttons placement.

Viewing Results

Data Table

After executing a query, the results are initially displayed as a large data table.

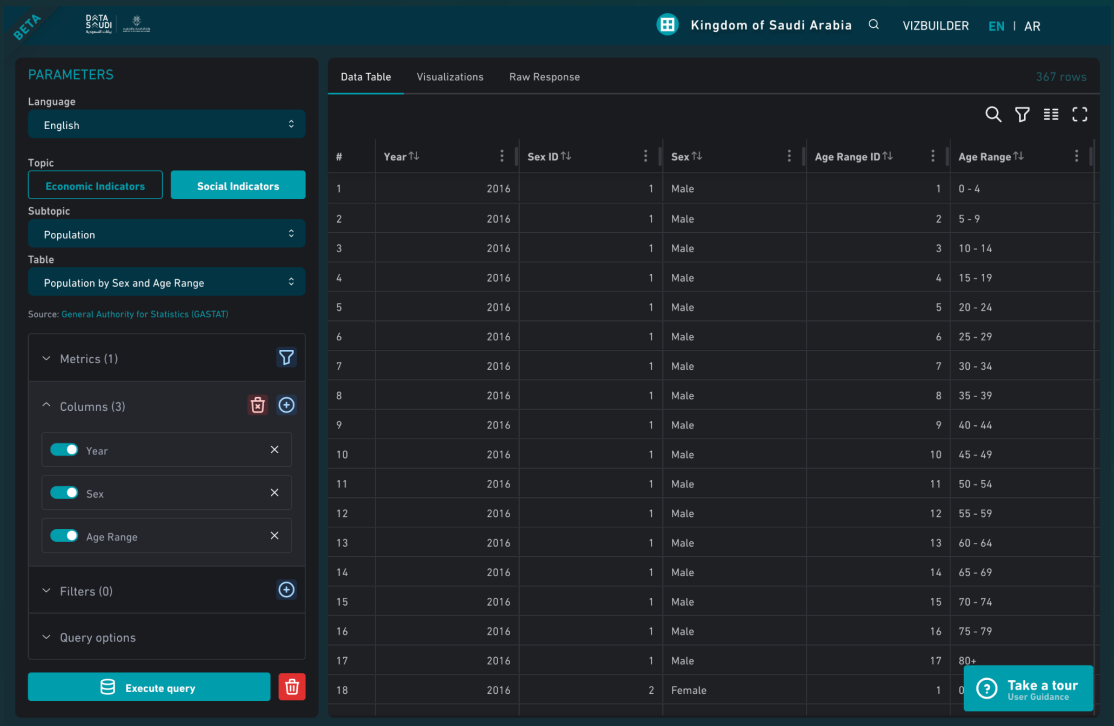

A data table result after pressing the “Execute query” button.

Sorting Data

Clicking on a column header allows users to sort the data in both ascending and descending order, as seen in this screenshot:

	Sex ID	Sex	Age Range ID	Age Range	Population
2020	3	Total	24	Total	35013414
2019	3	Total	24	Total	34218169
2021	3	Total	24	Total	34110821
2018	3	Total	24	Total	33413660
2017	3	Total	24	Total	32612846
2022	3	Total	24	Total	32175224
2016	3	Total	24	Total	31742308
2022	3	Total	7	30 - 34	3892832

Sorting the Data Table by descending “Population”.

Searching Data

Clicking the search icon, the first icon in the top corner of the Data Table, allows users to search the returned data for a specific query, filtering the results live as the user types.

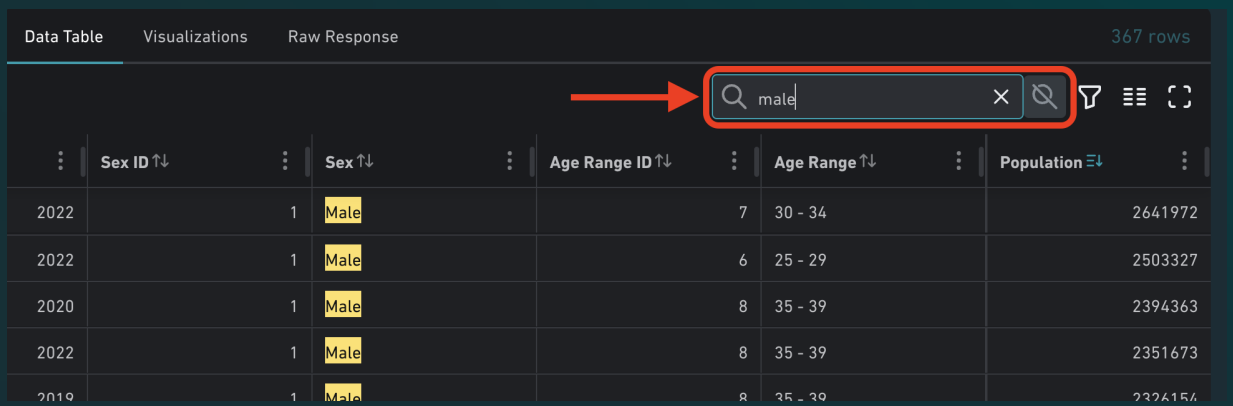

	Sex ID	Sex	Age Range ID	Age Range	Population
2022	1	Male	7	30 - 34	2641972
2022	1	Male	6	25 - 29	2503327
2020	1	Male	8	35 - 39	2394363
2022	1	Male	8	35 - 39	2351673
2019	1	Male	8	35 - 39	2326154

Searching the Data Table for a specific string (“male”).

Filtering Columns

Clicking the filter icon, the second icon in the top corner of the Data Table, toggles a filter interface for each column that appears underneath each individual column header. This interface allows for multiple types of String and Number filtering, such as “starts with” and “great than/less than”, which are applied to the data as the user types.

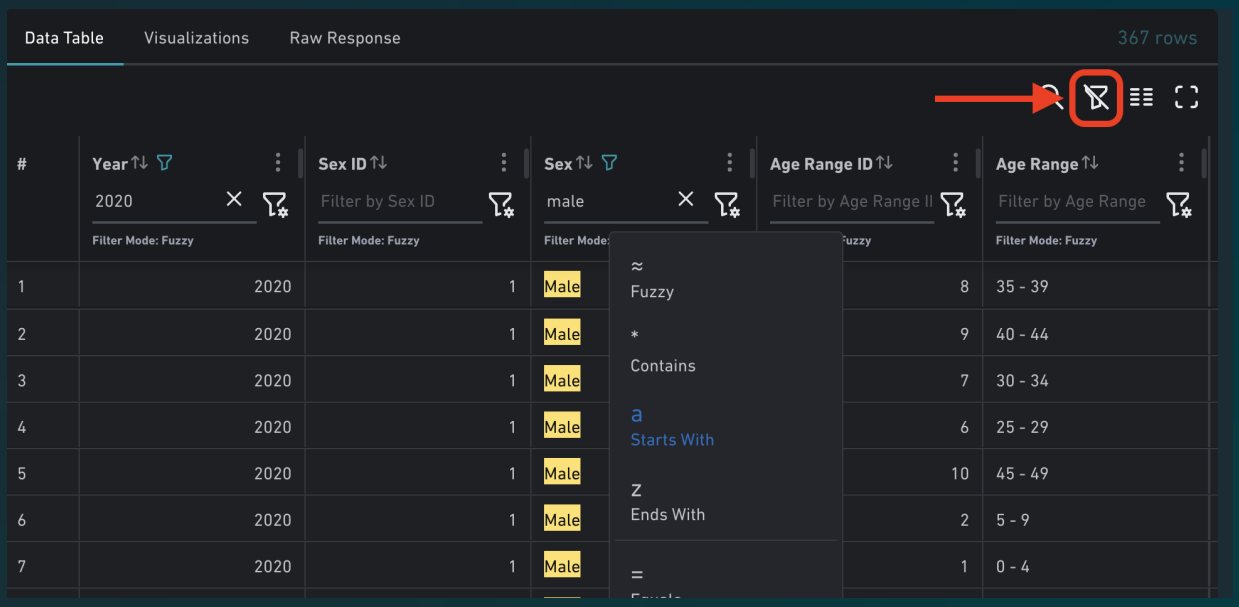

Data Table results being filtered by “Fuzzy” matching and “Starts With” on two separate columns.

Showing/Hiding Columns

The third icon in the top corner of the Data Table allows the user to edit which specific columns are returned with the data. This is commonly used to hide ID columns which may not be pertinent to a user's specific need.

#	Year	Sex ID	Sex	Age Range	Population
1	2020	3	Total	Total	
2	2019	3	Total	Total	
3	2021	3	Total	Total	
4	2018	3	Total	Total	
5	2017	3	Total	Total	
6	2022	3	Total	Total	
7	2016	3	Total	Total	31742308
8	2022	3	Total	30 - 34	3892832

Hiding the "Age Range ID" column using the Show/Hide interface.

Fullscreen Table

The final and forth icon in the top corner of the Data Table will enable fullscreen mode for the table, allowing users to explore large datasets with as much screen real estate as possible.

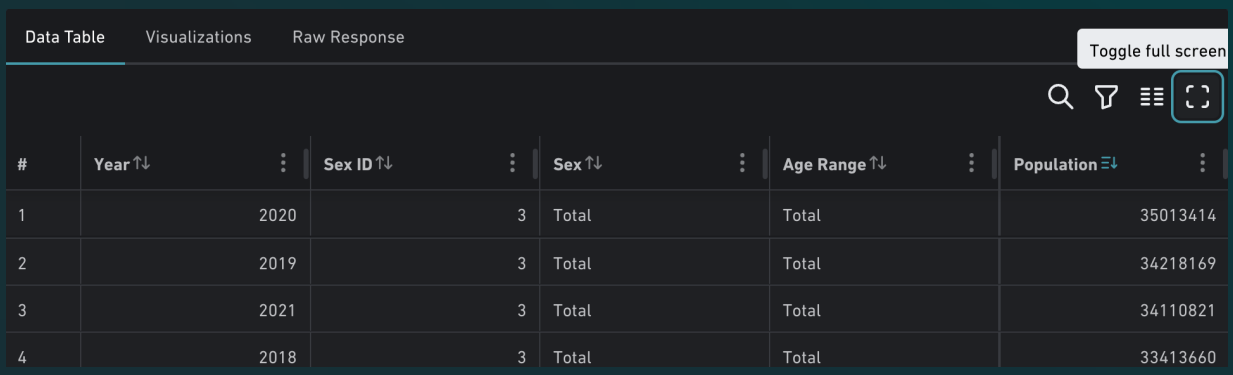

#	Year	Sex ID	Sex	Age Range	Population
1	2020	3	Total	Total	35013414
2	2019	3	Total	Total	34218169
3	2021	3	Total	Total	34110821
4	2018	3	Total	Total	33413660

The full screen toggle button about to be clicked.

Visualizations

The second tab in the results interface will take the currently selected data query and produce as many visualizations as possible. For complex/large queries, this often involves looking at each individual Metric across each individual Column, transforming the data into various slices which make the most sense for comprehensible visualizations.

The Visualizations tab showing small multiple visualizations based on a user query.

Enlarging a Visualization

By default, the Visualizations appear as small tiles, showing the full breadth of possibilities and combinations. When finding a visualization of interest, users can click the “Enlarge” button in the top corner of that visualization to zoom in on that single visualization, often showing more detail, labeling, and a timeline.

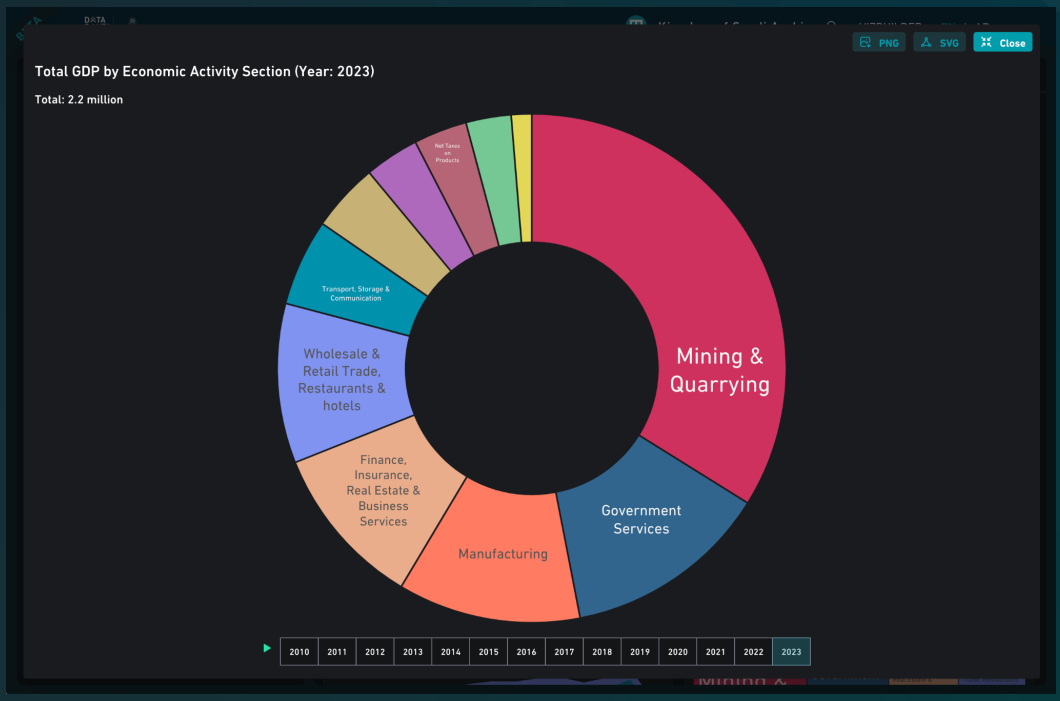

An enlarged Pie Chart appearing in a pop-up modal above the interface.

Downloading a Visualization

When viewing an enlarged visualization, users have the option to download the current visualization as either a PNG or SVG image by using the buttons in the top corner.

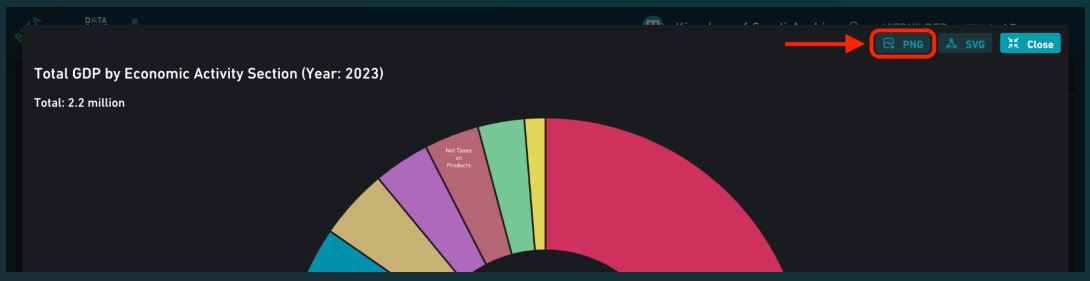

The location of the “Download as a PNG” button.

Raw Response

The final tab in the results panel is the API tab, which shows a few different ways to interact with the data in external/developer applications.

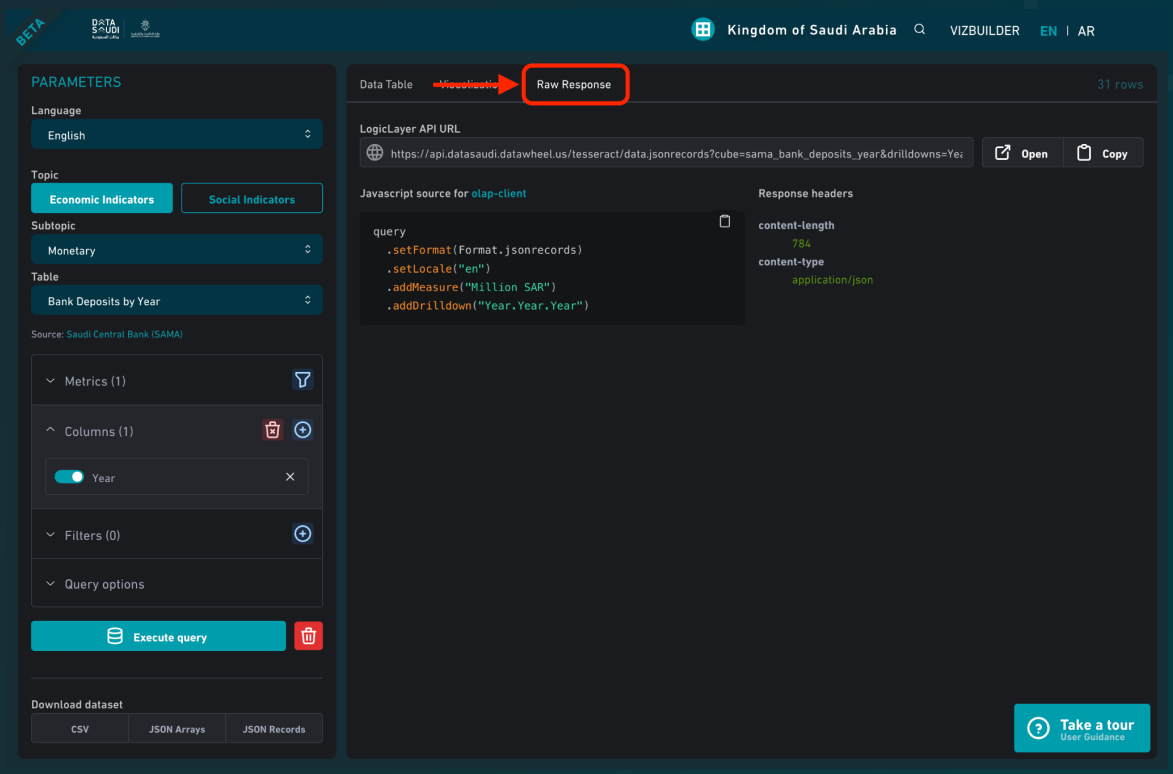

The Raw Response tab for an example query.

LogicLayer API URL

Every data query is made using a URL containing all of the query parameters. This URL can be queried from any application to retrieve data as JSON results. The buttons to the side of the URL can be used to Open the URL in a new browser window, or Copy the URL to the user's clipboard.

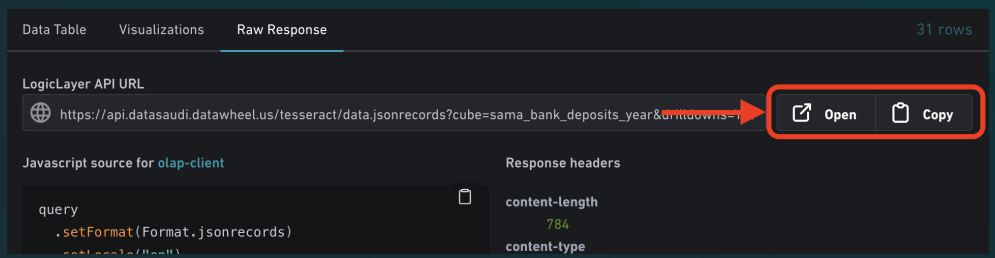

The "Open" and "Copy" buttons for the LogicLayer API URL.

Javascript Source

Additionally, the OLAP database can be queried directly from a Node/JavaScript environment using the @datawheel/olap-client open source library. The Raw Response tab shows the appropriate code needed to use this programmatic interface.

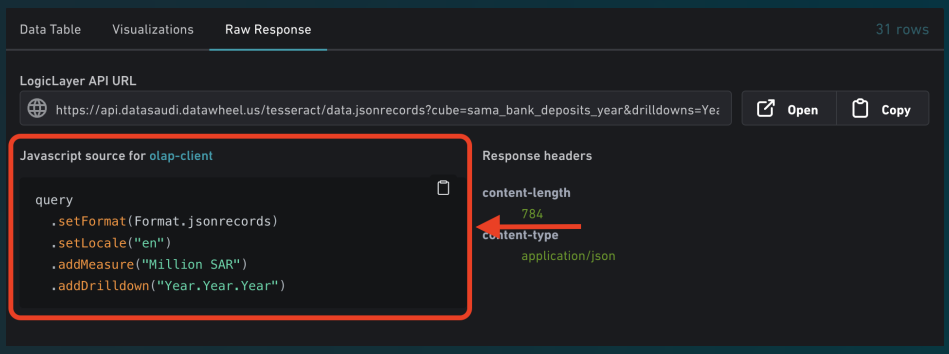

The Javascript Source for an example query.

Response Headers

And finally, the last part of the Raw Response tab shows the headers received from the API response, which can be helpful when debugging and integrating data into external platforms.

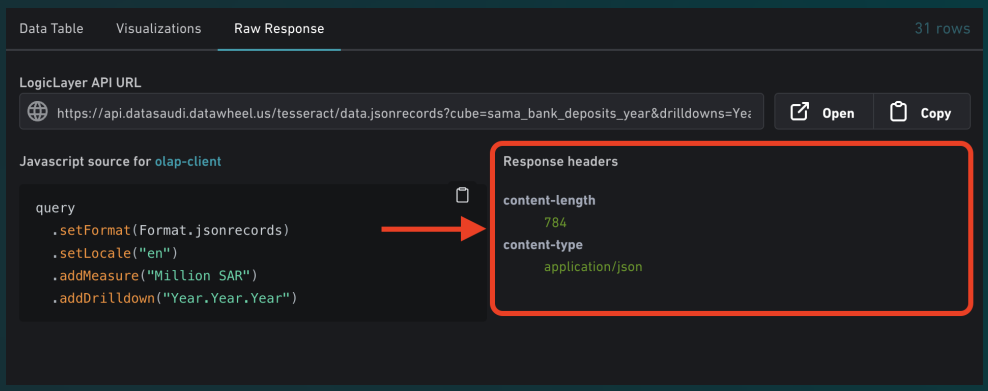

The Response Headers for an example query.

Interactive User Guide

Many of the interface descriptions outlined in this document are also available on the live site by clicking the “Take a Tour” button in the bottom right corner of the window. This will launch a guided tour highlighting the different elements of the interface using an example starting query.

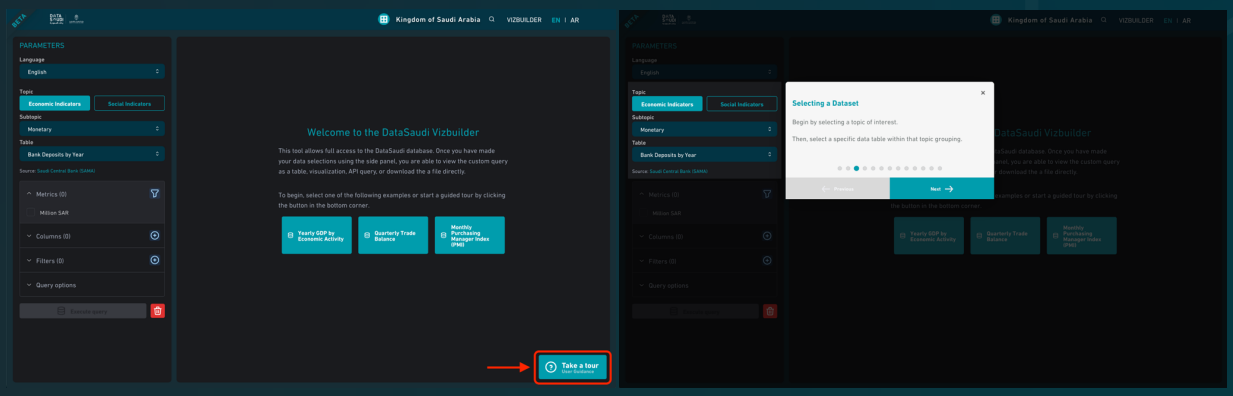

Button location to begin tour and example tour step.

Starting Examples

In addition to the interactive tour, the starting screen also displays example queries to help the user get started. Clicking any of these buttons will populate the interface with the selected dataset, allowing users to jump right into using the interface without having to manually choose a dataset and options.

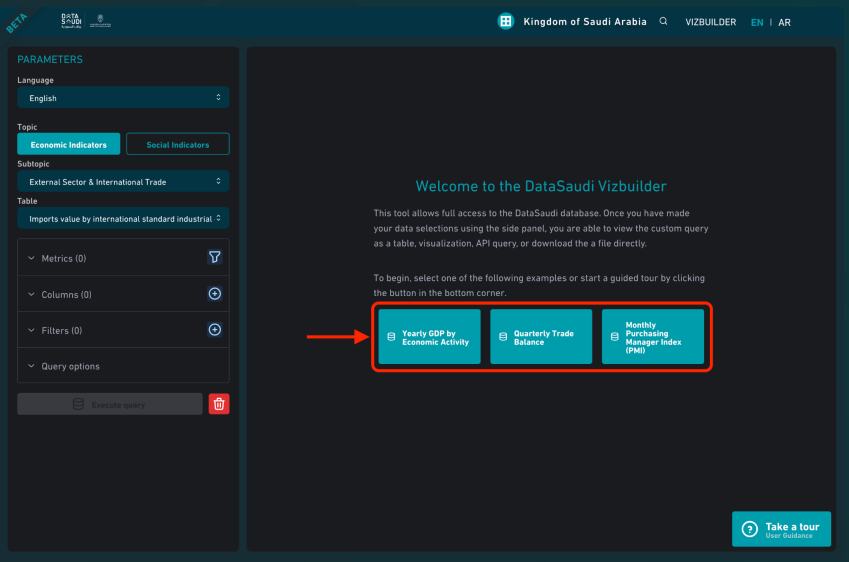

Button location of example queries.