Implementation of Business Intelligence with Power BI in Agricultural Commodities Traded in the United States

Torrez, David[1]; Calix, Marvin[2]; Leal, Martin[3]
[1-3]Zamorano University, Honduras, david.torrez@est.zamorano.edu, mjcalix@zamorano.edu, mleal@zamorano.edu,

Abstract*– Business intelligence makes it possible to efficiently visualize large volumes of data, thereby supporting a wide range of decision makers. In this study, a business intelligence tool was implemented to analyze the prices and quantities of agricultural commodities traded in the United States using Power BI. The objective was to design an algorithm that, through the My Market News Application Programming Interface (API), downloads three types of reports—terminal market, shipping point, and movement—bringing in information from 2015 onward. An extract, transform and load (ETL) process was established to automatically cleanse and unify the data, enabling seamless updates. Finally, a relational model was built that facilitated the creation of user-friendly interactive dashboards (control panels). The procedure included Python scripts to streamline the downloading, cleansing, and refreshing of the information. In addition, Power Query and DAX were employed to enhance data cleaning and ensure accurate visualization. The end product was a report containing five dashboards: "Prices and Quantity," "Conventional vs. Organic," "Seasonal Analysis," "Quantity Analysis," and "Origin Analysis". Moreover, near real time updates allow users to monitor market trends and support strategic decisions. The application of these visualizations enables producers, buyers, sellers, and analysts to identify trends and seasonality in key commodity variables—such as price, quantity, variety, and origin—thereby greatly facilitating informed decision making.

***Keywords*-- Business Intelligence, Power BI, Python, commodities, United States**

I. Introduction

In 1915, in Hammond, Louisiana, the United States Department of Agriculture (USDA) began issuing reports on certain agricultural commodities, which included information on price, variety, color, and other elements. These reports were produced with the intention of providing accurate information to the markets in question, thereby facilitating decision-making by producers, distributors, sellers, and buyers. Making this information publicly available allowed for fair and balanced competition among all stakeholders [1].

Currently, the Agricultural Marketing Service (AMS), an agency of the USDA, has continued to issue reports on the various agricultural commodities traded in the United States. In addition, they have created several reports that respond to the different needs that the agriculture-related public has encountered [2]. Among all these reports, three stand out for a very broad audience, namely: terminal market, point of shipment, and movement.

The Terminal Market Report provides data from 11 wholesale terminal markets across the United States, detailing each commodity's price, variety, color, origin, and other characteristics. Information is collected through in-person or telephone interviews with buyers and sellers. Participation is voluntary, and strict confidentiality is maintained: no individual business or respondent is identified. This anonymity is achieved by reporting average prices aggregated from a wide range of firms trading the same commodity [3].

On the other hand, there is the Shipping Point Report, which provides information very similar to that of the terminal market, but in this case from different commercial ports throughout the United States. To obtain this information, reporters contact hundreds of carriers, producers, and importers, always maintaining confidentiality with respondents. The data is then analyzed and sometimes corrected before being published in the report called Shipping Point [4].

The third report of interest is the Movement Report, which collects, verifies, analyzes, and reports information on domestic and imported movement. In addition, these reports also contain information similar to that of the terminal market and point of shipment, such as variety, origin, size, and packaging, but they do not include price information. The Movement Report collects some of its information from individual carriers, falling into the categories of boat, air, and land, which provide daily, weekly, and monthly information [5].

The other data source for the movement report is the Automated Commercial Environment (ACE) database, which is maintained by the Department of Homeland Security (DHS) and U.S. Customs and Border Protection (CBP) [5].

Currently, these reports, which were once transmitted via handwritten notes, paper reports, telephone records, and radio broadcasts, are now available to the public through a web portal, a mobile app, and more recently through an initiative called My Market News (MMN) [1].

MMN offers real-time, accurate, and unbiased information through various agricultural commodity reports. The biggest innovation of MMN is the inclusion of an application programming interface (API), which allows for quick and automated access to all these reports through the use of programming languages. Such access has never been so convenient. However, this in turn creates the complication of

having sufficient knowledge of programming and data wrangling to use the API effectively.

This is where business intelligence comes into play, facilitating interactive access to data using different tools, including dashboards [6]. A dashboard is an interactive visual tool that presents large volumes of information in a way that is understandable to the target audience for which it was designed [7]. A dashboard uses graphs, tables, and indicators to convey all this information to the different stakeholders, allowing them to monitor the behavior of the variables of interest, greatly facilitating decision-making.

In this work, an automated business intelligence tool was created that allows information to be updated in near real time, consisting of interactive dashboards (aggregations of information presented in different types of intuitive graphs) with information from the three types of reports of interest issued by My Market News. The information only includes data from 2015 onwards and is automatically updated, covering data up to a date close to the user's query. The dashboards only allow for exploratory analysis of the data, but do not allow for predictive or prescriptive analysis.

The objectives of this study were as follows:

•1 Create an algorithm that allows a virtual connection with the My Market News API to automatically download information about fruits and vegetables sold in the United States.

•2 Establish an automated data manipulation process to facilitate the continuous updating of databases and make them available for proper data modeling.

•3 Develop a relational model between the three databases for the proper integration of information to facilitate the creation of interactive graphics.

•4 Create interactive dashboards that display information graphically to monitor the behavior of fruit and vegetable prices and quantities, helping users make informed decisions by identifying trends and seasonality.

II. Methodology

A. Obtaining a Valid API Key

The information in the reports was obtained thanks to the MMN API. An account was created, which is required in order to obtain an API key, which was used to download the various reports.

B. Obtaining a Valid API Key

With the API key in hand and a Python code, it was possible to check all the reports available in the MMN database. Once this information was checked, it was saved in a CSV file for quick exploratory analysis. This file contained 1,015 available reports and showed the different identifiers for each one, as well as the last time it was updated and a brief description of the information it contained.

C. Selection of Reports of Interest

The pattern of names corresponding to terminal market, point of shipment, and movement reports was identified. Once the reports corresponding to each type of interest were identified, those that had been updated at least once in 2024 were filtered out, with the intention of having reports that continue to be updated. Once this was done, an Excel file with three sheets was obtained, each containing the information from the three types of reports of interest, respectively.

D. Downloading Reports

In order to download each report, it must first be read in the Python environment in JSON format. Once read in this format, it can be saved to different types of files. For this project, we decided to use the comma-separated values (CSV) format, as it greatly reduces storage space. The tasks of reading and downloading were achieved by defining two functions in the Python environment.

Next, a script was created that would allow for year-by-year downloads. This script was only responsible for downloading one year's worth of information in six-month sets, in order to generate a five-second pause between sets and avoid saturating the API. Additionally, the function identifies whether the year being downloaded is the current year, making a slight change so that it downloads from January 1st to the date of execution.

Finally, the main function was created, which, using the previously defined functions, would be able to perform the complete download of the three different types of reports. Three parameters were assigned to this function: start year, end year, and report type. In this way, it was only necessary to execute this function three times to obtain all the raw information of interest.

E. Data Transformation and Updating in Python

Data transformation was performed in Python and Power Query. In the Python environment, only columns that were not relevant to this study were removed, several of which contained very little information or were simply comments that were difficult to represent in graphs. After deleting the columns from the individual reports of a report type, we took the opportunity to unify them. Once this process was complete, there were three large files corresponding to each report type.

To finalize the use of Python, we created a code that allowed us to continue updating the unified files. This code efficiently compiles the previous steps, which also includes the removal of irrelevant columns. This allows the reports to be updated from the last day they contain to the current day when the code is compiled. Once the script was finished, the task scheduler was used to run the code daily, requiring only that the computer of choice be turned on.

It was necessary to link the files to OneDrive to facilitate database updates. To do this, all that had to be done was to define the data source with the OneDrive URLs in Power BI.

With the available resources, this was the most effective method for frequently updating the information.

F. *Data Transformation in Power Query*

Once the source of the files in Power BI had been defined, we proceeded to clean them up in Power Query. In this environment, we simply replaced some values in certain columns and created a custom column in the point of shipment and terminal market databases with M code. In addition, filters were applied to display only the information of interest on the dashboard.

G. *Creation of Lookup Tables*

After completing the cleanup, we began creating lookup tables, which are necessary elements for generating visualizations with information shared between databases. We evaluated which columns the databases shared and proceeded to create the corresponding lookup tables with DAX code. To create these tables, we extracted the unique values from the selected column in the three databases and used them to generate a new table.

H. *Data Modeling*

With the lookup tables ready, data modeling could proceed. Carefully, relationships were established between the three main databases and the newly created lookup tables. Only one-to-many relationships were used, paying special attention to cardinality. An erroneous relationship would prevent the information from being displayed correctly.

I. *Creation of Measures*

Measures were created to respond to needs that were identified when creating certain charts. Most were recalculations of information that was already available; however, the Power BI environment works much better when most of the key data is handled in the form of measures. DAX (Data Analysis Expression) code was used to perform this task effectively.

J. *Creation of Interactive Dashboards*

Six pages were created, including the home page, which serves as an introduction for users of the tool. The other five pages contain the five interactive dashboards, which are: prices & quantities, conventional vs. organic, seasonal analysis, quantity analysis, and analysis by origin. The following graphs were used: line graph, area graph, 100% stacked area graph, and ring graph. In addition, the Power BI component called small multiples was used to provide better visualizations.

Furthermore, data segmenters were added to all pages, as these elements are essential for the Dashboard to fulfill its function of applying multiple filters, as many categorical variables have been configured as advanced filters. The segmenters used allowed the following information to be filtered: years, commodity, variety, packaging, and origin. In addition, a home button was added to allow access to the Dashboard's main page at any time.

III. RESULTS

The various one-to-many connections were made using the previously created lookup tables. Great care was taken to make the connections with the correct columns and only in the corresponding databases. Close attention was also paid to the cardinality of these relationships. Any failure in the aforementioned factors would render the model useless when it came to obtaining effective visualizations. The model can be seen in Fig. 1.

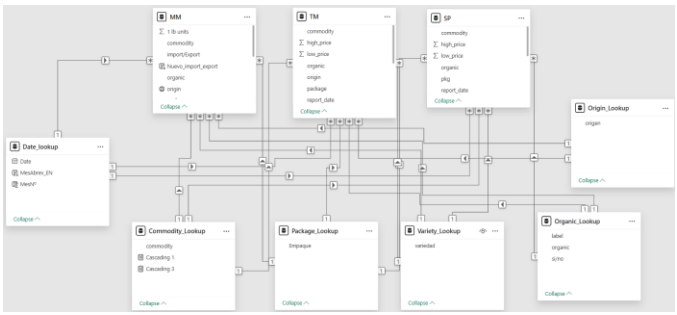

Fig. 1 Database modeling.

The first dashboard was deliberately crafted to present, briefly, the essential insights from all three report types. By combining high-information-density visuals with succinct summary metrics, it enables users to grasp the key aspects of their chosen commodity in just seconds. The final layout is shown in Fig. 2.

The main visual is a dual-axis line chart showing monthly average prices and cumulative quantities. The remaining visualizations and components were designed to complement the primary price-and-quantity chart. A donut chart was created to display total quantities, segmented by the shares of imported, exported, and domestic product. In addition, a small-multiples view was developed to show quantities by origin, enabling users to discern differences across sources. Given the large number of origins for certain commodities, small multiples were deemed the most effective approach for revealing location-specific patterns and trends.

To complete the page, a series of data slicers was implemented to enable deeper, more granular analysis of every variable available in the databases. The slicers provide filters for year, commodity, organic status, packaging, variety, origin, terminal market, port, and movement. Leveraging Bookmarks, each slicer not only filters all visuals on the page but also cross-filters the other slicers, thereby preventing the display of empty charts to the user.

Fig. 2 Dashboard of prices & quantities

We deemed it essential to dedicate an entire page to the differences between organic and conventional commodities. This section highlights the contrast in quantities and prices for organic products, as shown in Fig. 3.

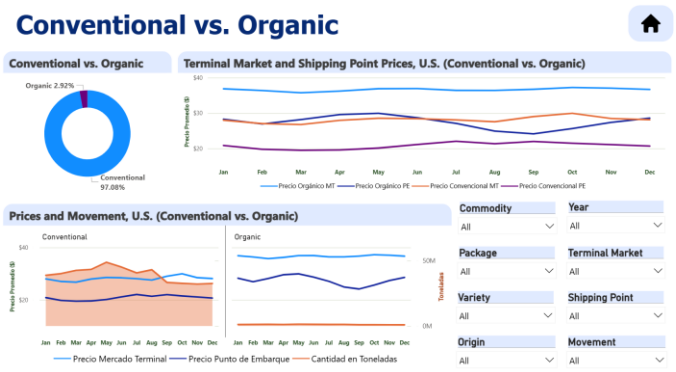

Fig. 3 Dashboard of conventional vs. organic

We recognized the need to analyze intra-year trends. To that end, three separate visuals—one for each report type—were created. Fig. 4 shows the full page, which combines these charts with the complete set of data slicers. Designed for historical and trend-focused analysis of each commodity, the page relies heavily on the year slicer: by selecting specific years, analysts can quickly plot and compare periods of interest, making differences immediately apparent.

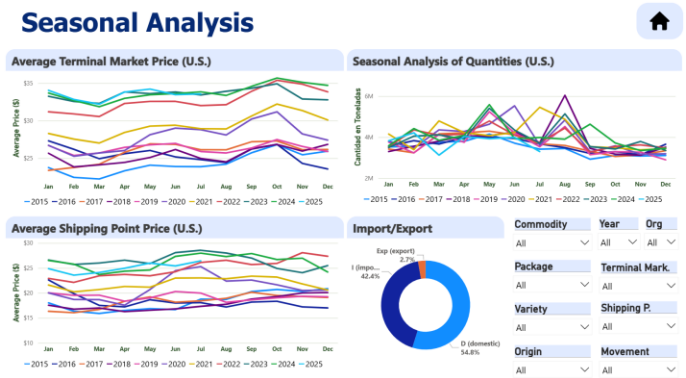

Fig. 4 Dashboard of seasonal analysis

Fig. 5 presents the final version of this page, which is intended to help analysts examine quantity dynamics in isolation. The charts work well together, enabling users to view both percentage and absolute figures and thereby offering a broad yet straightforward perspective.

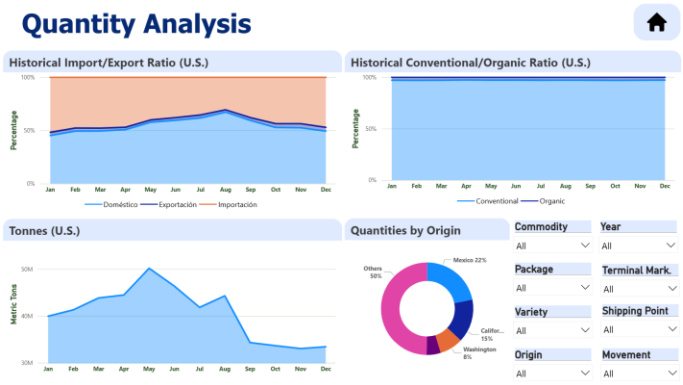

Fig. 5 Dashboard of quantity analysis

The fifth and final dashboard was designed for users primarily interested in the origin of the various commodities. We successfully created a section that provides clear exploratory insights into each product's provenance. The finished layout is shown in Fig. 6.

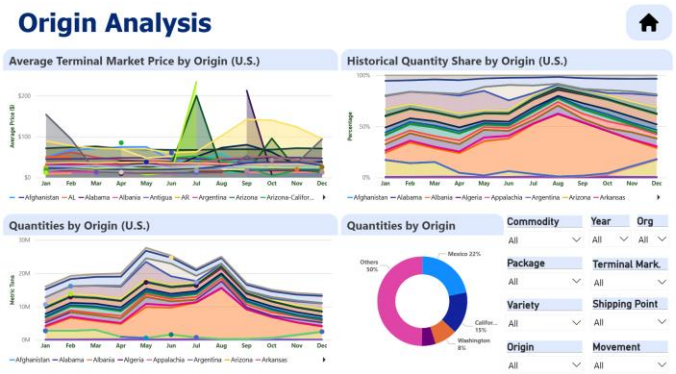

Fig. 6 Dashboard of origin analysis

To enhance the user experience, a landing page was created to streamline navigation across the entire dashboard. Its purpose is to present—succinctly and intuitively—the key information and features covered by the tool. The final layout of this landing page is shown in Fig. 7.

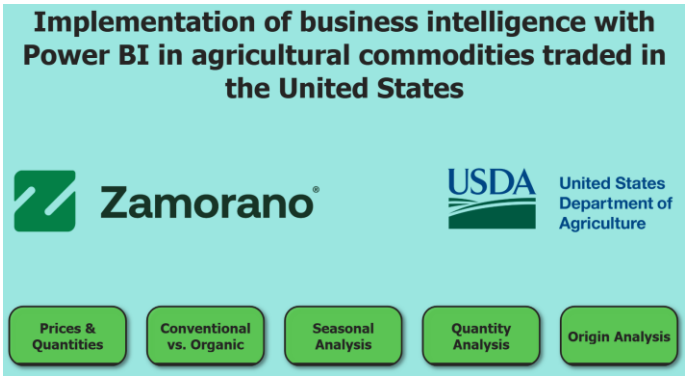

Fig. 7 Home page

The page displays the study's title—serving simultaneously as the page header—and features the logos of Zamorano and the United States Department of Agriculture (USDA). At the bottom, a navigation bar allows users to jump directly to every section of the report.

IV. Discussion

The business-intelligence tool proved highly successful, delivering valuable insights through the five dashboards. These views allow users to identify seasonal trends in quantity, price, and origin; precisely gauge the price differential between organic and conventional production; detect volume concentrations across commodities; and uncover many other insights as they interact with the platform.

Armed with this information, a skilled analyst can mitigate logistical risks, capitalize on market windows when prices are more favorable, and optimize purchase-contract timing by anticipating shifts in supply. Users can also compare price differences by origin, production method, variety, packaging type, and location within the United States, providing supplementary input for broader studies. All these capabilities are delivered with updates that are practically real-time.

In addition to descriptive analytics, the tool's standardised ETL pipeline and integrated relational schema create a robust foundation for more advanced modelling. Because price, volume, origin and production-method dimensions are already normalised within a single data model, these tables can be exported directly to statistical or machine-learning frameworks for time-series forecasting, scenario simulation, or optimisation studies.

Although the current implementation is limited to exploration dashboards, the same Python scripts that orchestrate the daily refresh could be extended to generate predictive features and feed automated forecasts back into Power BI, thereby transforming the platform from a passive reporting service into a proactive early-warning system.

The platform's near-real-time updates also open avenues for policy monitoring and sustainability assessment. Regulators and trade associations can track shifts in import, export and domestic supply as they occur, rapidly detecting anomalies such as sudden volume surges at a particular port or price distortions in each region. At the same time, the dashboards' ability to juxtapose organic and conventional streams, or compare origins and packaging types, supplies evidence for targeted certification programmes and localized marketing campaigns. By enabling granular comparison across origin, production method and variety, the tool supports transparency initiatives aimed at fair pricing, helps logistics planners anticipate bottlenecks, and provides researchers with a rich empirical base for studying how supply-chain interventions ripple through U.S. commodity markets.

V. Conclusions

This study demonstrates the feasibility of integrating the entire data life cycle from My Market News into a single business-intelligence tool. First, an automated algorithm was designed and validated to connect continuously to the official API, downloading price and volume data for agricultural commodities traded in the United States. This approach eliminates manual processing and ensures that the databases remain up to date.

Building on this infrastructure, a fully automated ETL pipeline was implemented to cleanse, transform, and structure the information, allowing each new extraction to be incorporated seamlessly into the analytical model. This workflow not only reduces refresh times but also preserves data quality—an essential prerequisite for rigorous analysis.

The resulting relational model links the three primary sources (terminal market, shipping point, and movement reports), enabling the creation of measures and filters that simultaneously exploit variables such as price, quantity, origin, organic category, and product presentation. This architecture supports a Power BI report comprising five interactive dashboards that reveal seasonal patterns, price differentials between conventional and organic production, and volume concentrations by geographic origin.

Collectively, the solution provides producers, buyers, and analysts with a holistic, near-real-time view of the market, enabling them to anticipate supply fluctuations, optimize purchase contracts, and mitigate logistical risks. The project therefore not only meets its technical objectives but also underscores the strategic value of automation and advanced data visualization for decision-making in the agri-food sector.

VI. Recommendations

Migrate the current data from OneDrive files to a cloud-based relational database such as Azure SQL, SQL Server, or MySQL. Centralizing the data in a managed RDBMS will accelerate daily refreshes, enhance concurrency, and deliver more reliable query performance. Ensure that the migration team has the requisite database-administration and programming expertise to design the schema, configure security, and automate incremental loads.

Deploy the Python update script on a serverless cloud execution service—e.g., Azure Functions, AWS Lambda, or

Google Cloud Run—to achieve a higher degree of automation. Scheduling the script in the cloud removes the need for a physical machine to power each day, lowers operational overhead, and improves overall reliability.

ACKNOWLEDGMENT

I would like to express my deepest gratitude to Eng. Christian Torrez, whose support and expertise were instrumental in the development and completion of this business-intelligence tool.

REFERENCES

[1] J. Main, *USDA Market News on MARS/MMN (Market Analysis and Reporting Services/MyMarketNews)*, 2023. [Online]. Available: https://www.zoomgov.com/rec/play/L0gUt8HQ1ZFExHJDZjAZjDi2SJJi_FDPTQKHYWl9_-aJOWxC4ifRqIQO4ClbRE-yt_BYQUfShZWp-j_v.VYGRPp9XqTk1jy9d?eagerLoadZvaPages=&accessLevel=meeting&canPlayFromShare=true&from=share_recording_detail&startTime=1683046839000&componentName=rec-play&originRequestUrl=

[2] J. Okoniewski, *Introduction to Market News- Shared screen with gallery view*, 2023. [Online]. Available: https://www.zoomgov.com/rec/play/q9BUHgRIdswvKWbySefSB0X6uX3M5lcvBJ_w-F1AxnGQwLLmyHZaqnqvdURojTUrXx0c1oKXAufzoUMQ.u-zNrA51IHqhuQ4I?eagerLoadZvaPages=&accessLevel=meeting&canPlayFromShare=true&from=share_recording_detail&continueMode=true&componentName=rec-play&originRequestUrl=https%3A%2F%2Fwww.zoomgov.com%2Frec%2Fshare%2F-uj2o59xSx9SQiRHioL91_X3szRLxwGeS6biBMG-amP_KeM8G-Qv_DRpKwtzQPrD.qpN2QmwTmdZ7gjYZ

[3] A. Swan, *Price Reporting: Wholesale Markets, Shipping Points, Retail, Organic Data Initiative, Hemp*, 2023. [Online]. Available: https://www.zoomgov.com/rec/play/Sm5E9upWwuLtnDVBmCLwLHRYwvfrEErDExjnM6tFsrTqp4EqQpPc4xsw43ENpF3PRvwlp8Fpd1HHk_kR.gGnlAuPS4CKyARMn?eagerLoadZvaPages=&accessLevel=meeting&canPlayFromShare=true&from=share_recording_detail&continueMode=true&componentName=rec-play&originRequestUrl=https%3A%2F%2Fwww.zoomgov.com%2Frec%2Fshare%2FHJBzbDMyN1oOndz6H1jBIkcP0wcbbBxw6cdTvALcWJuALgNviF4RYJMknhCwx4so.cYLdMXGC6cLz9cKD

[4] A. Zaki, *Price Reporting: Wholesale Markets, Shipping Points, Retail, Organic Data baInitiative, Hemp*, 2023. [Online]. Available: https://www.zoomgov.com/rec/play/Sm5E9upWwuLtnDVBmCLwLHRYwvfrEErDExjnM6tFsrTqp4EqQpPc4xsw43ENpF3PRvwlp8Fpd1HHk_kR.gGnlAuPS4CKyARMn?eagerLoadZvaPages=&accessLevel=meeting&canPlayFromShare=true&from=share_recording_detail&continueMode=true&componentName=rec-play&originRequestUrl=https%3A%2F%2Fwww.zoomgov.com%2Frec%2Fshare%2FHJBzbDMyN1oOndz6H1jBIkcP0wcbbBxw6cdTvALcWJuALgNviF4RYJMknhCwx4so.cYLdMXGC6cLz9cKD

[5] P. Willkie, *USDA Market News on Movement/Transportation: Shipments, Truck Rates, Imports, HTS codes*, 2023. [Online]. Available: https://www.zoomgov.com/rec/play/efCw1qWh6n9rn9-h0wA2qXe0q7LvJpEEIIqvUCCH2RS1t3AU2uOT6xQ_KSoF998Ujj5U6YUK6PNV_tET.xnCI-dDLDAKFnp5y?eagerLoadZvaPages=&accessLevel=meeting&canPlayFromShare=true&from=share_recording_detail&continueMode=true&componentName=rec-play&originRequestUrl=https%3A%2F%2Fwww.zoomgov.com%2Frec%2Fshare%2F1vowxFeefeNmmGAodkwWWNDwcWHItGSoBd0z_U8SjeH_ocm4mOOnpQjnm1muW3Ua.5Sg6HEzulOOmbrQj

[6] R. Sharda, D. Delen, and E. Turban, *Business intelligence, analytics, and data science: a managerial perspective,* 4th ed. New York: Pearson, 2018. [Online]. Available: https://opac.atmaluhur.ac.id/uploaded_files/temporary/DigitalCollection/ZTRjYzc5ODA3OTQ0NTJlMTAxODI5MDlhM2QxMTcwM2E2NTE5Yzc4Mw==.pdf

[7] B. Bach *et al.,* "Dashboard Design Patterns," *IEEE transactions on visualization and computer graphics*, vol. 29, no. 1, pp. 342–352, 2022, doi: 10.48550/arXiv.2205.00757.