



Is Your Data AI-Ready?

The strength of your AI depends on the quality of your data. Making sure your data is properly prepared is the first step to unlocking better results and more value from AI.

When your data is AI-ready, everything runs smoother. You'll get more accurate responses with fewer hallucinations and less frustrating back-and-forth. At Domo, we've guided hundreds of teams through the process of getting their data ready for AI, and now, we're sharing our insights. Our checklist covers the key steps for programmatically cleansing your data—with a few helpful hints about how Domo.AI can help ease the workload.

Below, you'll find essential tips for cleaning your data—making sure it's polished and primed for any specific use case you envision. These tips will get you 80 percent of the way—the final 20 percent will depend on your AI application.



Labeling		
Check	Goal	Example
<input type="checkbox"/>	Data in columns must match column labels.	Column with revenue data should include "revenue" in the column name.
<input type="checkbox"/>	Data column labels are unique (do not match any other column names) and concise/short, yet sufficiently intuitive for understanding what the data in the column represents.	<i>Emails_Opened_Last_Year</i> is a concise label for a column containing data on emails opened in the last year.
<input type="checkbox"/>	Data column descriptions are included when necessary or helpful, particularly if intending to use generative AI to explore the data set.	Following specific, consistent naming conventions. For example, if a column is labeled <i>Revenue</i> , but only includes the revenue from certain product lines, this should be noted in the description (or better yet, indicated in the column label).
<input type="checkbox"/>	Columns are appropriately described in the data set when values in a column are not intuitive.	"K-product sales" is an example of a nonintuitive data value (what is a K-product?).

Row considerations		
Check	Goal	Example
<input type="checkbox"/>	There is a unique identifier/primary key for each row, meaning a column or columns that contain values that uniquely identify each row in a data set. The unique identifier column(s) should be the first column(s) in the data set.	<i>EMPLID</i> , a column containing the unique ID number for each employee at a company, is usually the unique identifier/primary key in a data set containing employee data.
<input type="checkbox"/>	Note in the data set description about what the data granularity is (what does a row represent?) and which columns in the data set are unique identifier(s).	Example data set description: <i>Each row corresponds to an employee in the company. EMPLID is the unique identifier.</i>
<input type="checkbox"/>	Rows/observations are reasonable to include in data set for analysis and others are removed.	Remove rows for "admin", "test", "internal", etc.
<input type="checkbox"/>	Duplicate rows are identified, flagged, and managed programmatically.*	Rows that have the same value across all columns.
<input type="checkbox"/>	Rows with duplicate identifiers are identified, flagged, and managed programmatically.*	Rows that may not have the same value across all columns but that have the same unique identifier.

Within-column accuracy & consistency

Check	Goal	Example
<input type="checkbox"/>	Data type is appropriate given the values in a column. Columns are reasonable to include in a data set for analysis and to remove noise, especially for chatbot AI. Being intentional is the key.	Values in a column that are numeric are stored as a fixed decimal, floating decimal, or integer data type; text values are stored as a string data type; dates and times are stored as date or time-stamp data type.
<input type="checkbox"/>	Data formatting is consistent.	States are represented with the correct two-letter abbreviations. Phone numbers formatted consistently: either 10 digits with no specific characters, or 10 digits with dashes, or 10 digits with parentheses for area code. Just be consistent.
<input type="checkbox"/>	Range of data values in each column is appropriate and accurate.	All data values for a given column fall within an expected and specified range; e.g., a column with data on people's ages should generally not include any negative values or values greater than 110.
<input type="checkbox"/>	Outliers are identified, flagged, and managed programmatically for each column.*	Unusually high and low data points in each column that look suspicious must be investigated and appropriately managed.
<input type="checkbox"/>	Categorical variables are clean and categories are grouped following conceptual/business/statistical logic.	Categories are combined based on logic.
<input type="checkbox"/>	Number of buckets for categorical variables are reasonable for analysis.	Are there too many buckets considering the amount of data? Can smaller buckets be combined? Do you have at least 25 observations in each bucket after all segmentation?
<input type="checkbox"/>	Bucket size for categorical variables is reasonable for analysis.	Buckets that describe small percentages of the rows in the data set should be managed.
<input type="checkbox"/>	Distribution of data in each column is understood, and necessary data transformations should be completed.	Normal, skewed, other, etc.
<input type="checkbox"/>	Missing data is coded consistently within each column.	Placeholder values for missing data, "N/A", empty cells, or other methods for identifying missing data should be applied consistently.
<input type="checkbox"/>	Missing data is identified, flagged, and managed programmatically (many data science algorithms automatically exclude any row with missing data in at least one cell; missing data generally needs to be managed to avoid this).*	Although we may have certain features, how well are these features defined? Is the majority of data defined or missing? Can we compute missing data?
<input type="checkbox"/>	Provide formula used for calculation of the metric, and verify that the implementation matches the calculation.	If profit is revenue minus COGS, verify that it is implemented in a Beast Mode.

Between-column accuracy & consistency

Check	Goal	Example
<input type="checkbox"/>	Appropriate granularity is applied or labeled, such that the same unit of analysis is expressed in each column of the data set.	Hourly, daily, weekly, monthly, etc.
<input type="checkbox"/>	Appropriate aggregations are applied or labeled, such that the same unit of analysis is expressed in each column of the data set.	Uniques, counts, averages, etc.
<input type="checkbox"/>	Duplicate columns are identified, flagged, and managed programmatically.*	Are they truly duplicate columns?
<input type="checkbox"/>	Constants are identified, flagged, and managed programmatically.*	Columns with only one value are generally not useful in analyses.
<input type="checkbox"/>	Other data and calculation errors are identified, flagged, and managed programmatically.*	Beast Modes, ETLs, etc.
<input type="checkbox"/>	Missing data has been explored and patterns of missing data across columns is understood.	Do certain rows have missing data consistently across the same features?

Joins

Check	Goal	Example
<input type="checkbox"/>	Row count of data sets must be validated both before and after joins.	If a table is left joined, it should only include up to the number of rows as in the left table.
<input type="checkbox"/>	Join keys must be cardinal.	If you are joining a customer and orders data set, does each customer ID truly correspond to just one customer in the orders data set? All keys in one dataset have a match in other dataset. Make sure there are no orphans IDs, or there is a reasons for not having a match.

*The language of "identified, flagged, and managed programmatically" is vague, but the goal is context-specific.