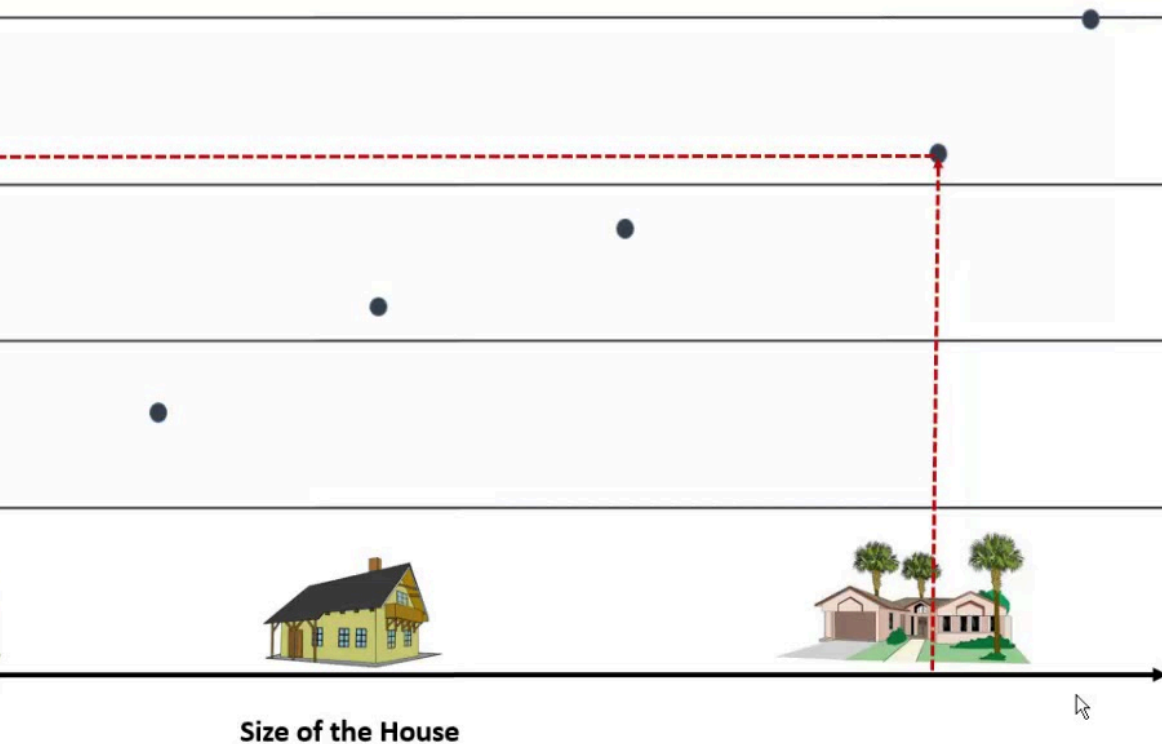


2: Estimating The Price of a House



Predicting House Prices with Linear Regression

Explore how to use Linear Regression to predict house prices based on factors like area, number of bedrooms, and location.

Introduction



Predicting House Prices

Use Linear Regression to predict house prices based on factors like area, number of bedrooms, and location



Leveraging Data

Gather relevant data on house prices, area, number of bedrooms, and location

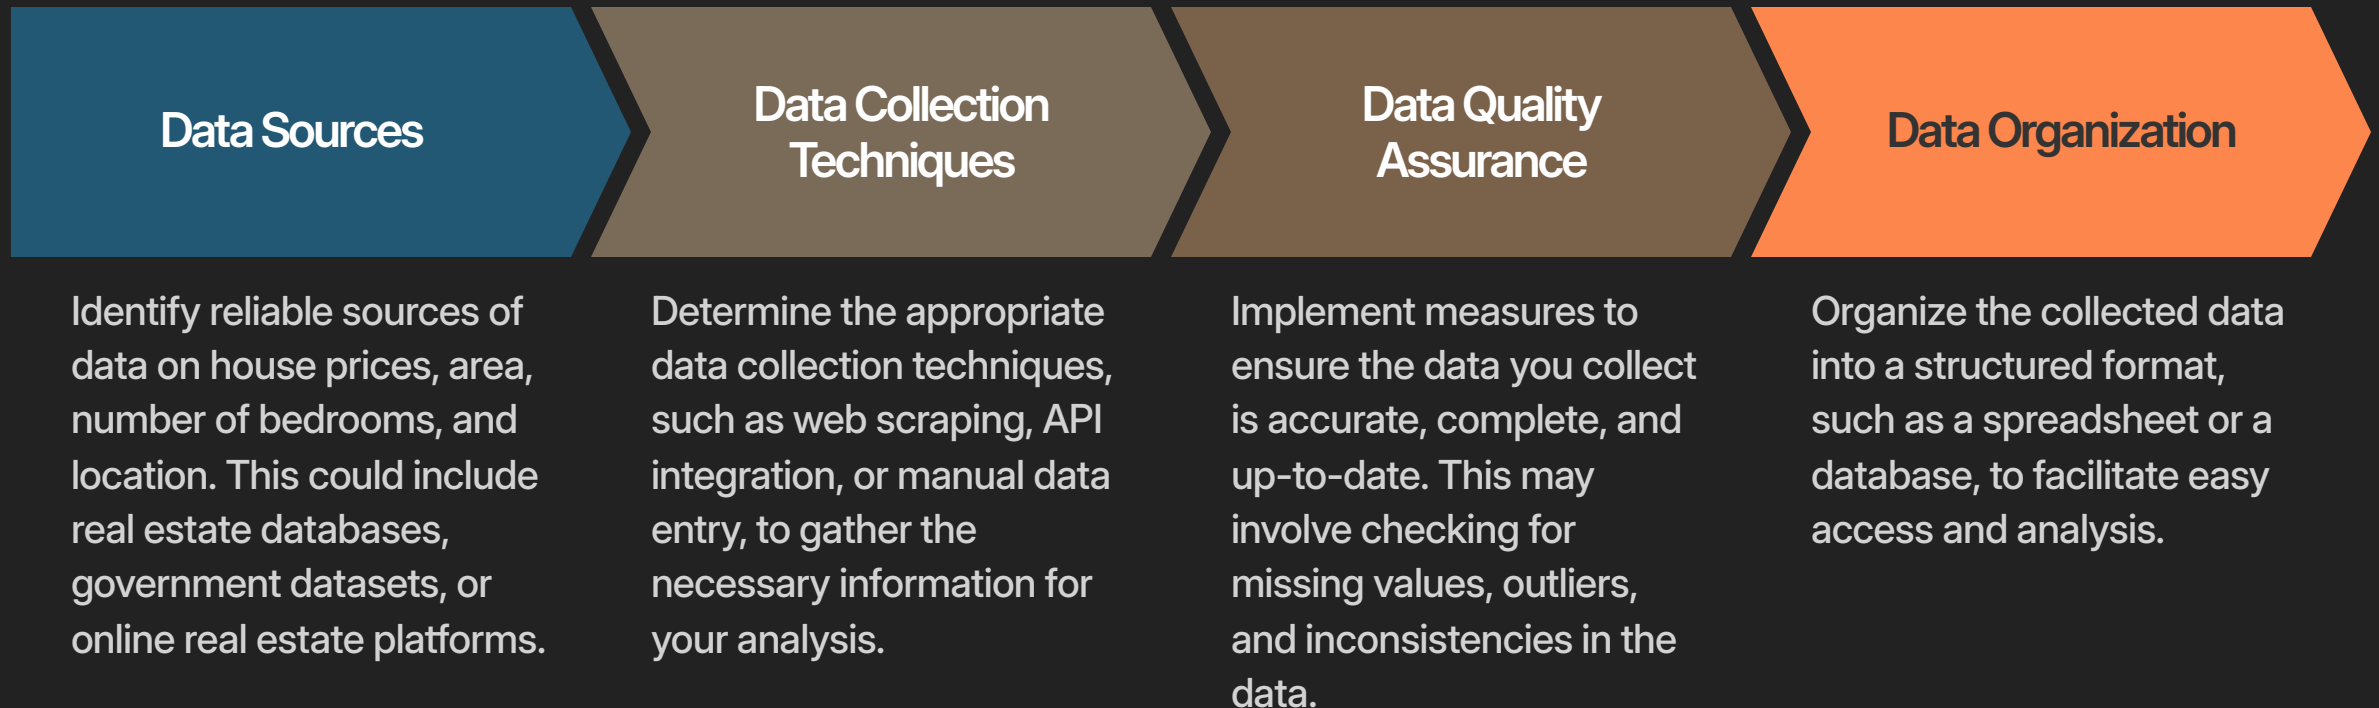


Data Preprocessing

Clean and prepare the data for model training, including handling missing values and normalizing features

By exploring the use of Linear Regression, we can leverage data on house attributes to accurately predict house prices and support decision-making in the real estate market.

Data Collection



Data Preprocessing

Handle Missing Values

Identify and address any missing values in the dataset using techniques like imputation or data removal.

Feature Normalization

Normalize the feature values to a common scale, such as using standardization or min-max scaling, to ensure no feature dominates the model training.

Categorical Feature Encoding

Convert any categorical features (e.g., location) into a numerical representation that the model can understand, such as using one-hot encoding.

Data Splitting

Divide the dataset into training and testing sets to evaluate the model's performance on unseen data.

Feature Engineering

Create new features or transform existing ones to improve the model's ability to capture the underlying relationships in the data.

Model Training



The diagram illustrates the sequential steps of model training. It features four horizontal arrows of increasing length, each representing a step. The arrows are colored in a gradient from blue to orange. The first arrow is blue and labeled 'Train-Test Split'. The second is light orange and labeled 'Hyperparameter Tuning'. The third is a darker orange and labeled 'Model Convergence'. The fourth is the darkest orange and labeled 'Training Time'. Each arrow starts from the left edge of the frame and points to the right, with its length increasing from top to bottom.

Train-Test Split

Hyperparameter Tuning

Model Convergence

Training Time

Model Evaluation

Assess the accuracy and performance of the Linear Regression model using the following metrics:

1%

R-squared

0%

Mean Absolute Error

0%

Root Mean Squared Error

1%

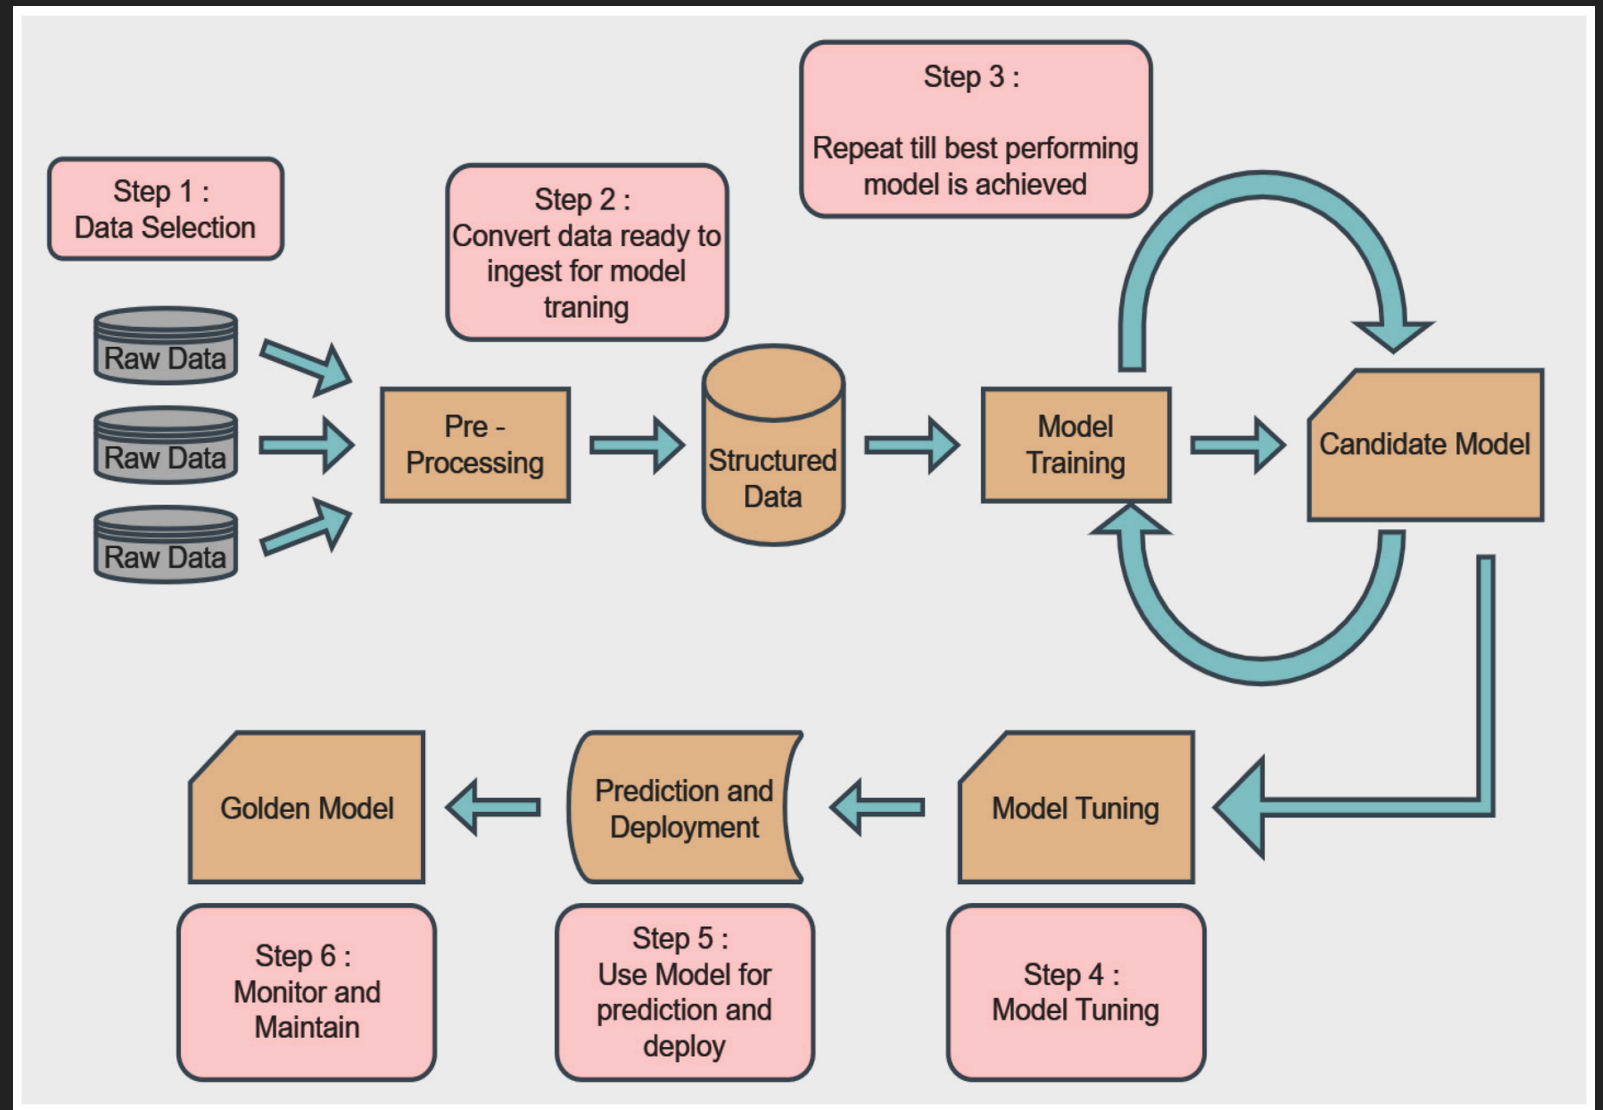
Cross-Validation Score

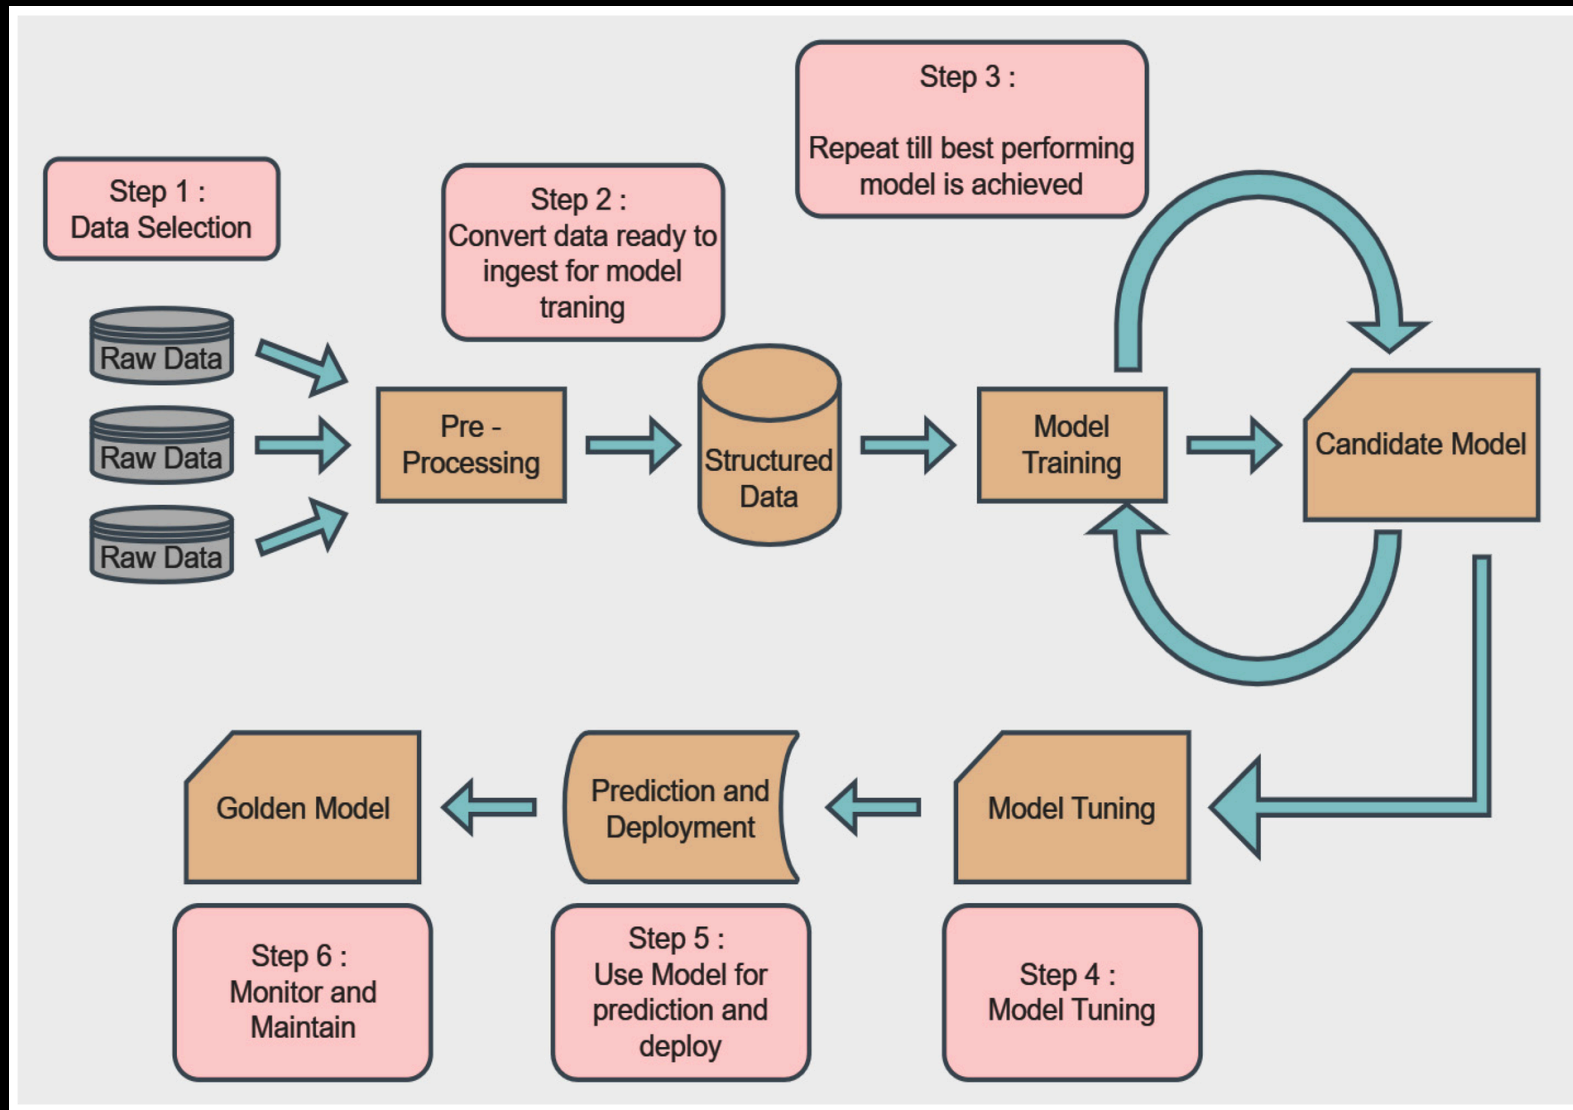
Model Deployment



Real-World Application

The Linear Regression model can be invaluable for real estate professionals, homebuyers, and policymakers. By applying this model to factors like property size, number of bedrooms, and location, they can generate accurate price predictions to inform buying, selling, and investment decisions. This data-driven approach can help identify undervalued properties, guide pricing strategies, and support policy decisions aimed at improving housing affordability.





The Linear Regression model has demonstrated its power in predicting house prices by leveraging factors like area, number of bedrooms, and location. By implementing the steps outlined in this presentation, we can train, evaluate, and deploy the model to generate accurate predictions, ultimately supporting real-world decision-making in the housing market.