Artificial (Augmented) Intelligence in Practice at Botkeeper

September 22, 2023

AI UNCHAINED
Accounting Intelligence. No boundaries. No limits.

Speaker Introductions

Jenny Blackwell
Software Engineering Manager

Travis Cherry
VP of Engineering

Brian Emposimato
Director of Product

Learning Objectives

- Discuss how Botkeeper’s models work
  - Focus on our transaction categorization AI models and data sourcing/normalization
  - Understanding machine learning techniques for categorizing client-specific data
  - Learn how to source GL and financial institution data for ongoing categorization
- Identify how different client data affects performance
- Utilize the configurations for AutoPush to fine tune and correct your client’s model
  - Choosing ML Training Start Date
  - Identifying “uncategorized” accounts
  - Using TransactionManager for “augmented“ intelligence alongside our AI models

Session Overview

- Automating the Accounting Cycle
- Zoom in on Transaction Categorization
- The Evolution of Transaction Manager and AutoPush
- Understanding the Business Problem
- Sourcing Historical GL Data
- Model Training
- Sourcing Transaction Data with SmartConnect
- Transaction Categorization with Transaction Classifier
- Unifying Disparate GL Interactions
- Architecture Summary and Ongoing Training
- Optimizing Data and Fine Tuning Performance
- AutoPush Performance Over Time
- What’s Coming for AutoPush and Transaction Manager
- Going Beyond - Understanding LLMs and Applicable Use Cases

Poll #1

Which component(s) of the accounting cycle would you like to see automated?

1. Data Collection
2. Transaction Categorization
3. Journal Entry Creation
4. Reconciliation
5. Financial Review
6. Report Generation
7. Client Communication

AI UNCHAINED
Accounting Intelligence. No boundaries. No limits.


AI in the
Accounting Cycle
Collect Data
SmartConnect
Categorize Transactions
AutoPush & TxM
Create Journal Entries
JE Automation*
Reconcile Accounts
Auto Bank Rec*
Review Financials
BotReview
Generate Reports
Reporting*
Communicate
Financial Insights*
Close Books
BotReview

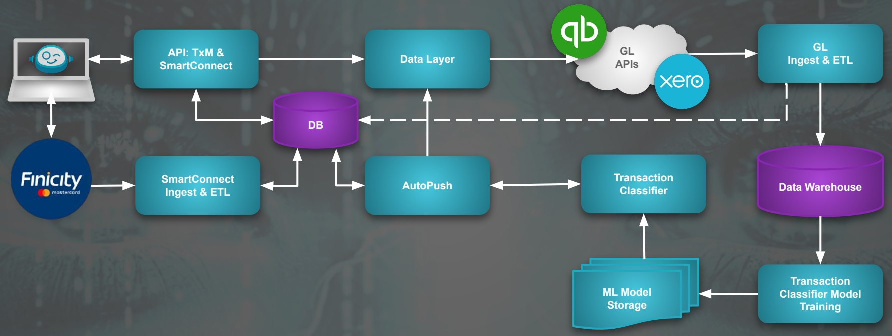

Zoom in on Transaction Categorization

The Evolution of Transaction Manager and AutoPush

- Botkeeper began with a team of skilled bookkeepers amassing a highly accurate dataset.
- We observed and closely monitored this team to understand the areas most ripe for automation, which drove our product roadmap.
- The first version of Transaction Manager was launched in 2020 as a tool that our bookkeeping teams used to streamline transaction collaboration with partners and their clients.
- In 2022, AutoPush fundamentally changed our model by intercepting and automatically categorizing high-confidence transactions, leaving medium and low for human review.

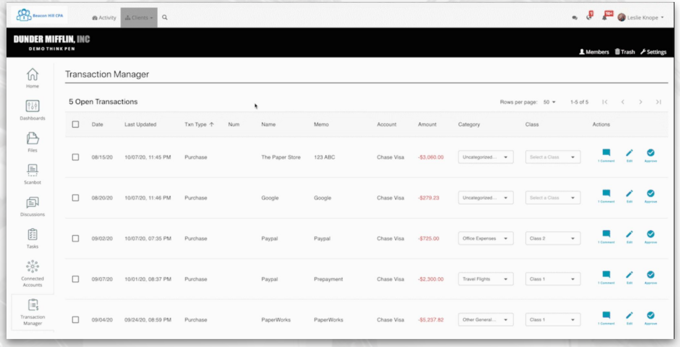

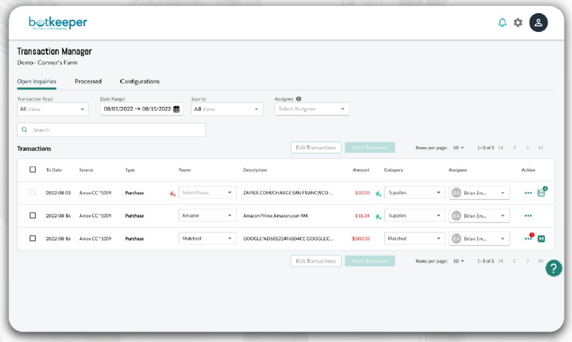

Understanding the Business Problem

Which account/vendor is a particular transaction associated with?

Let’s Get Familiar With the Data - Chart of Accounts

Each entity defines and maintains their own unique COA

- Account attributes:
    - **Name** - User defined, short description of the account
- Standard Account Attributes
    - **Classification** - Accounts are generally defined using hierarchical categorizations. The highest level of this hierarchy is referred to as the classification. The two GLs we support today share the same classification values.
        - EQUITY, ASSET, EXPENSE, LIABILITY, REVENUE

      Type and Sub-Type - These values differ substantially across GLs so they are normalized into a standard set of values and assigned

```
Account ID: 123
Name: Travel Meals
Account Type:
      Classification: EXPENSE
      Type: EXPENSE
      Sub Type: TravelMeals
```

Global definitions allow for more predictive power!

Let’s Get Familiar With the Data – Vendors

Each entity defines and maintains their own unique set of vendors

- Vendor attributes:
    - **Name** - User defined, short description of the vendor
- There are no standard vendor attributes

```
Vendor ID: 456
Name: Starbucks Coffee
Display Name: Starbucks Coffee
```

Let’s Get Familiar With the Data - Transactions

“Transaction” represents various object types within the GL , like Purchases and Deposits for QBO and Spend and Receive for Xero.

- Standard transaction attributes:
  - Type - Normalized across GLs, Purchase and Deposit
  - Date
  - Total
  - Vendor ID
  - Line Items
    - Amount
    - Description
    - Account ID

```
Transaction ID: 789
Transaction Type: PURCHASE
Transaction Date: 2020-01-01
Total: 5.67
Vendor ID: 456
Lines:
      Line Number: 1
      Amount: 5.67
      Description: Starbucks Coffee
      Account ID: 123
```

Transaction Classifier and Botkeeper's Approach to Automating Transaction Categorization

- When creating models, answering more general questions tends to require more data than answering more specific questions.
- As humans, we break up problems into smaller pieces to solve.
- If you can narrow the question, generating a performant model becomes easier!

Which account/vendor is a particular transaction associated with?

Does this transaction look enough like another historically categorized transaction to likely be associated with the same account/vendor?

Ensemble learning is an approach to machine learning that seeks better predictive performance by combining multiple models.

The Ensemble Approach – Entity Level Models

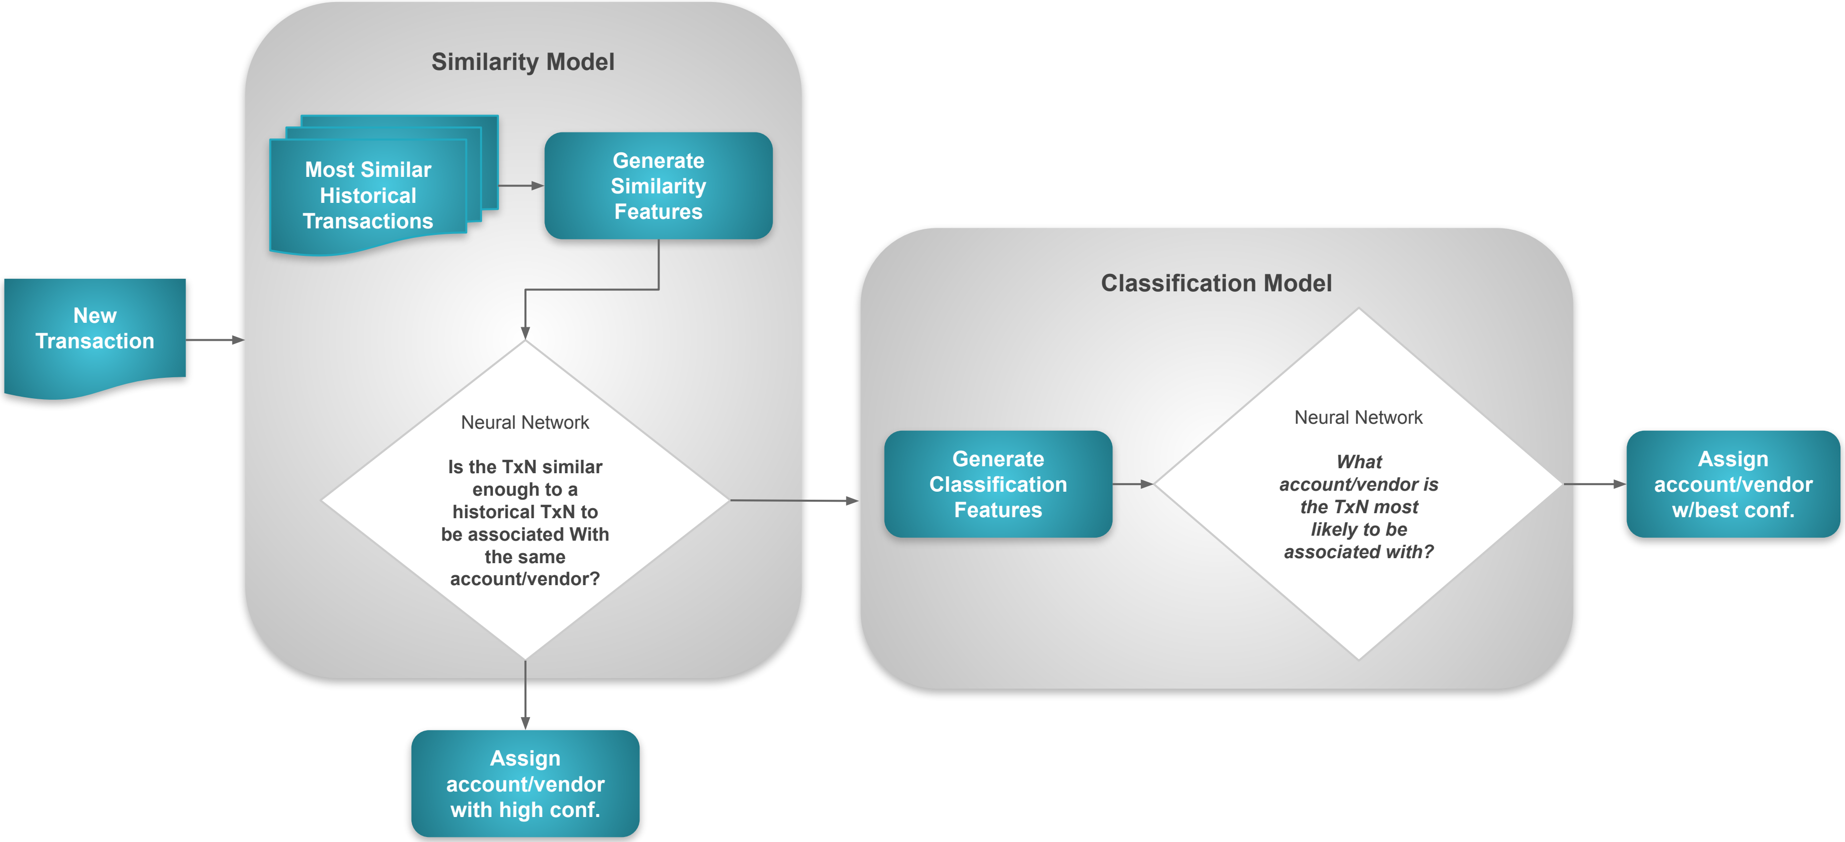

From Training Data to Trained Model

Collect Training Data

Our models learn from historically categorized transactions.

Pre-processing and Enrichment

Pre-processing steps allow us to get the most useful information out of our data.

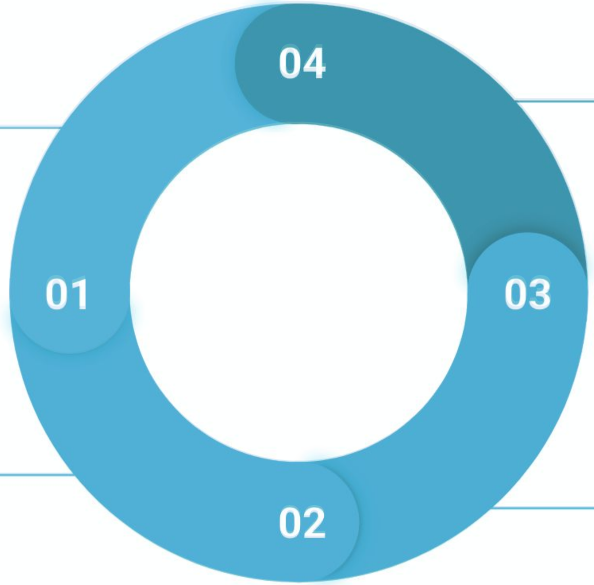

Model Training

Learning how important each attribute in our dataset is to solving our business problem.

Feature Engineering

Feature engineering is a technique that leverages data to create new variables that aren’t in the training data set.

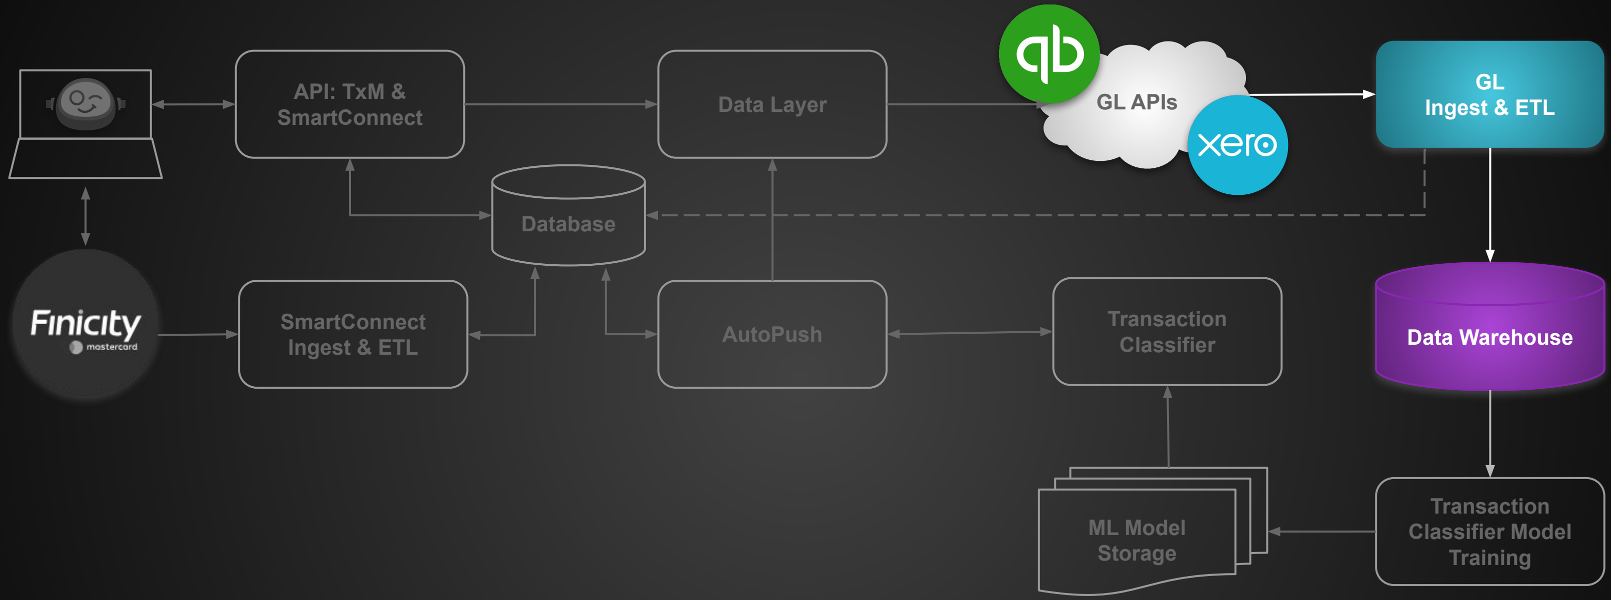

Sourcing Historical GL Data

Collecting the Training Data

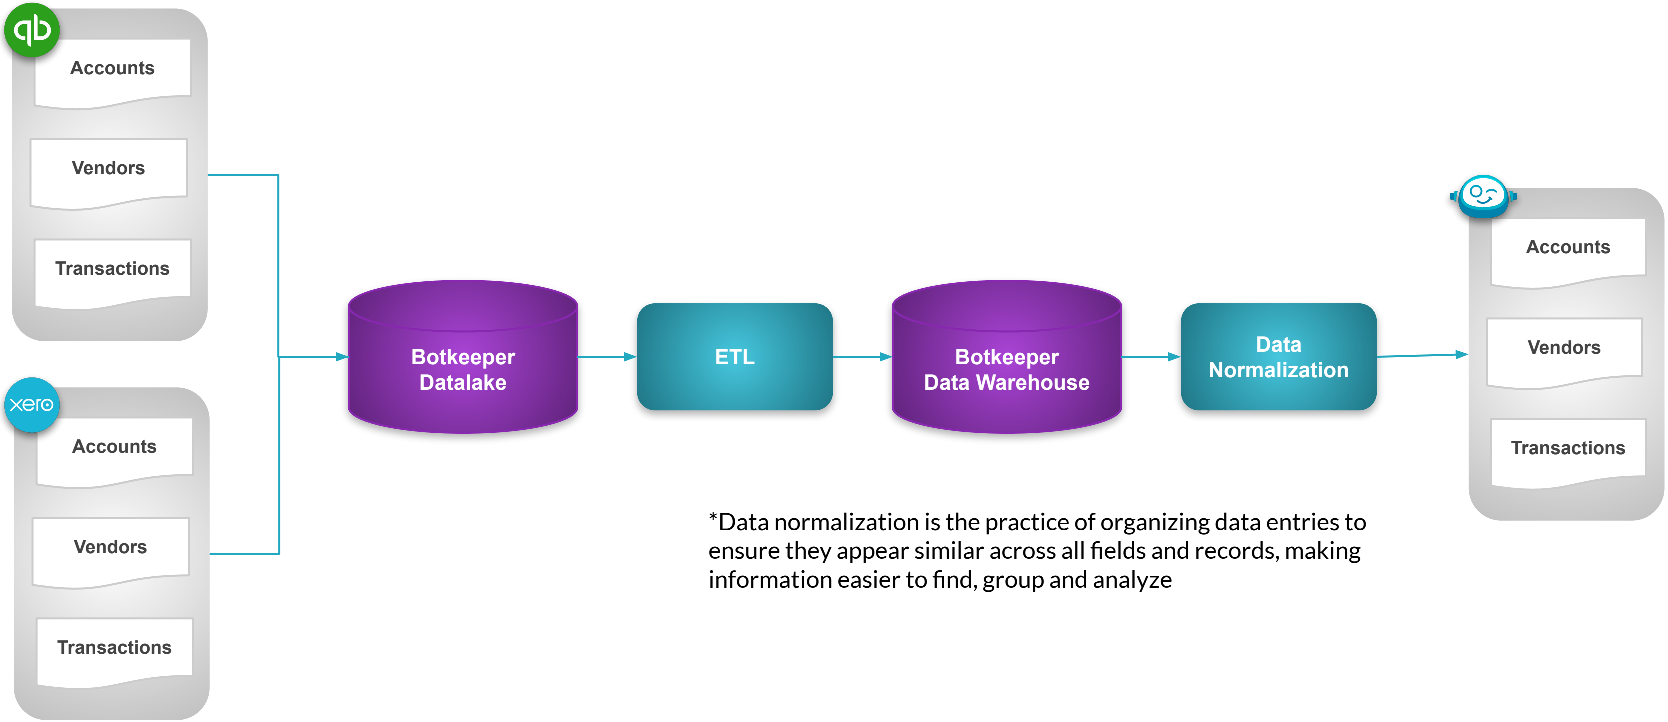

*Data normalization is the practice of organizing data entries to ensure they appear similar across all fields and records, making information easier to find, group and analyze

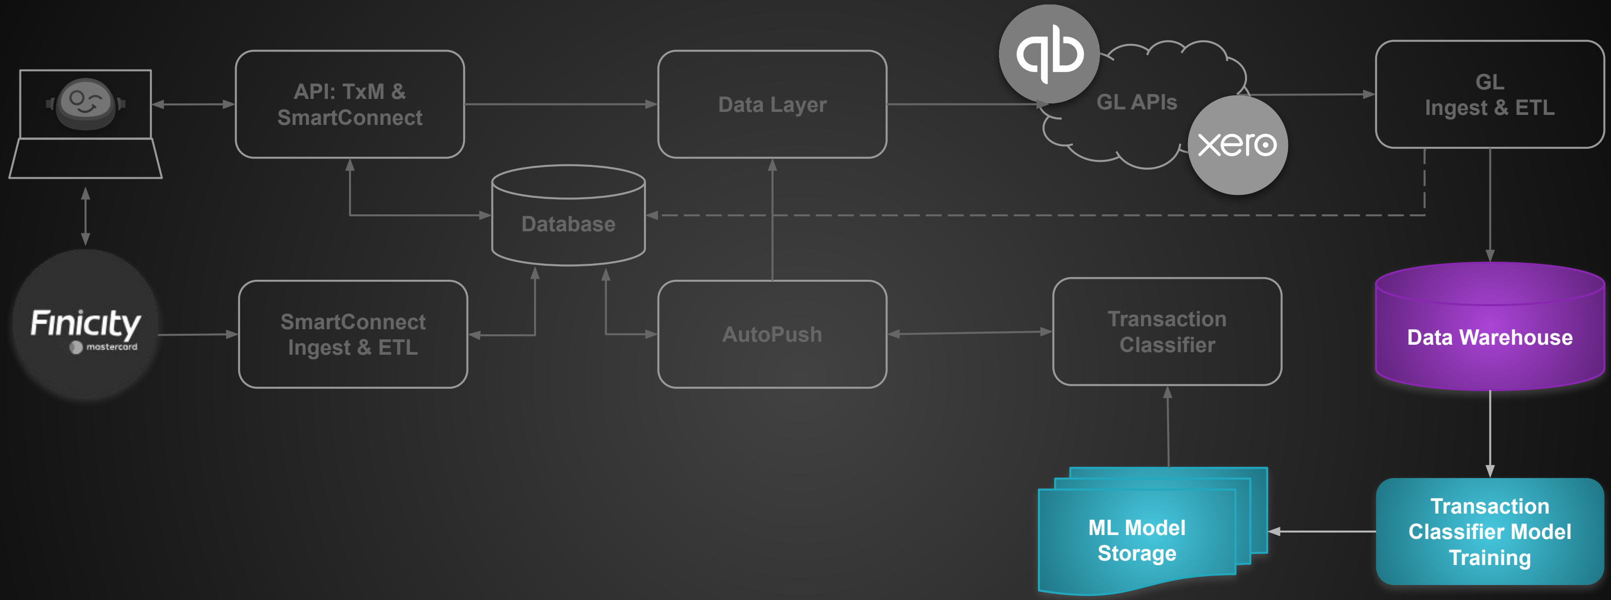

Model Training

Pre-processing and Enrichment

- **Pre-processing** steps allow us to get the most useful information out of our data.
    - Example:
        - Convert line item descriptions to lower case
        - Replace or remove some common patterns found in transaction line item descriptions
            - Names of months, dates, times, punctuation

          Book Club September becomes **book club MONTH**

- **Enrichment** steps allow us to add any derived attributes that are needed
    - Example:
        - A normalized version of the amount value (monetary values are non-normal)

Feature Engineering

- Feature engineering is a technique that leverages data to create new variables that aren't in the training data set.
- The goal of each feature is to capture a predictive attribute about a record.
- When all of these attributes are analyzed together, the unique patterns that they form should correlate with a particular value you are looking to predict.
- The features we have available to us help to inform what type of model we should employ. Once a model is chosen, features can be further refined.

- Examples:
  - Normalized edit distance between descriptions
  - Normalized difference between transaction amounts
  - Word sequencing

Edit Distance

An edit distance is a method used to quantify how dissimilar two strings are, measured by counting how many operations are needed to transform one into the other. There are several defined ways to determine edit distance, each allowing for different types of operations.

Levenshtein distance between **kitten** and **sitting** is 3**:**

1. **k**itten → **s**itten
2. sitt**e**n→ sitt**i**n
3. sittin → sittin**g**

Normalized Difference Between Amounts

In our case, we take the absolute difference between transaction amounts and divide it by the mean of the two amounts.

Transaction 1: 9.56
Transaction 2: 5.34

Absolute difference: abs(9.56 - 5.34) = 4.22
Mean: (9.56 + 5.34)/2 = 7.45
Normalized difference: 4.22/7.45 = 56.64%

Word Sequencing

Representing a given string, like a transaction description, in a way that allows for comparison to an entire corpus of words used across all transaction descriptions for an entity. Each word in the corpus is assigned a number greater than 0, where 0 is reserved unknown or no words, and the string is then translated into the corresponding sequence of numbers. Something like word sequencing is necessary, because models work on numerical values.

Description: starbucks coffee

Corpus:

- starbucks: 1
- coffee: 2
- expense: 3
- food: 4
- phone: 5

Word Sequence: [1,2,0,0,0]

Model Training

- Both the similarity and classification models used within our ensemble are neural networks, and trained at the entity level.
- Neural networks are supervised models - take an input, make a prediction, compare to the expected output, and make adjustments accordingly, iterating until the output is optimal.
- Neural networks are made up of interconnected nodes or neurons in a layered structure that resembles the human brain.
- Neurons are simple processing units that receive input, perform calculations on that input, and produce an output
- Once trained, model attributes are stored, ready to make predictions for newly introduced transactions.

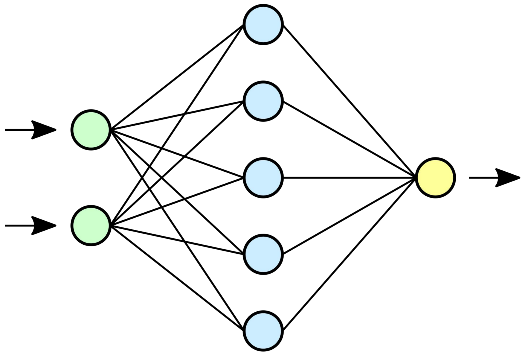

Global Account Model Embedding Layer

- As mentioned previously, the global definitions of account classifications, types, and sub types, give us a boost in predictive power!
- We used this as an opportunity to build another model that is trained on our global data set.
- The goal of this model is to predict the account type given a transaction.
- Also a neural network!
- This trained model can then be embedded into the entity level classification model, giving that neural network a more informed "starting point".

Global Vendor Model

- The newest improvement to our vendor model!
- When transactions associated with new vendors are introduced into an entity's data set, there is no history to help the entity model make a prediction. The vendor may not even exist in the GL yet.
- In this iteration, we use the power of our global data set to find similar transactions and suggest the most likely associated vendor based on global patterns.

Optimizing Data and Tuning Performance

- Performant models require a large volume of consistent, accurate historical training data
  - “Garbage in, garbage out”
- Botkeeper requires a minimum of 100 historical transactions for training an entity model.
  - Must include line item descriptions (Xero does not enforce this)
  - Must include a vendor for vendor model training (QBO does not enforce this)
- If there is a subset of historical data that is known to be inaccurate, inconsistent, or messy, Botkeeper offers an option to exclude this from the model’s training data.
  - From within Transaction Manager configurations, enter the ML Training Start Date
- Botkeeper also excludes transactions assigned to uncategorized accounts from training. Be sure to define these within the Transaction Manager configurations.

- Check out “Unleash the Bots” to learn more about how to fine-tune your AutoPush settings!

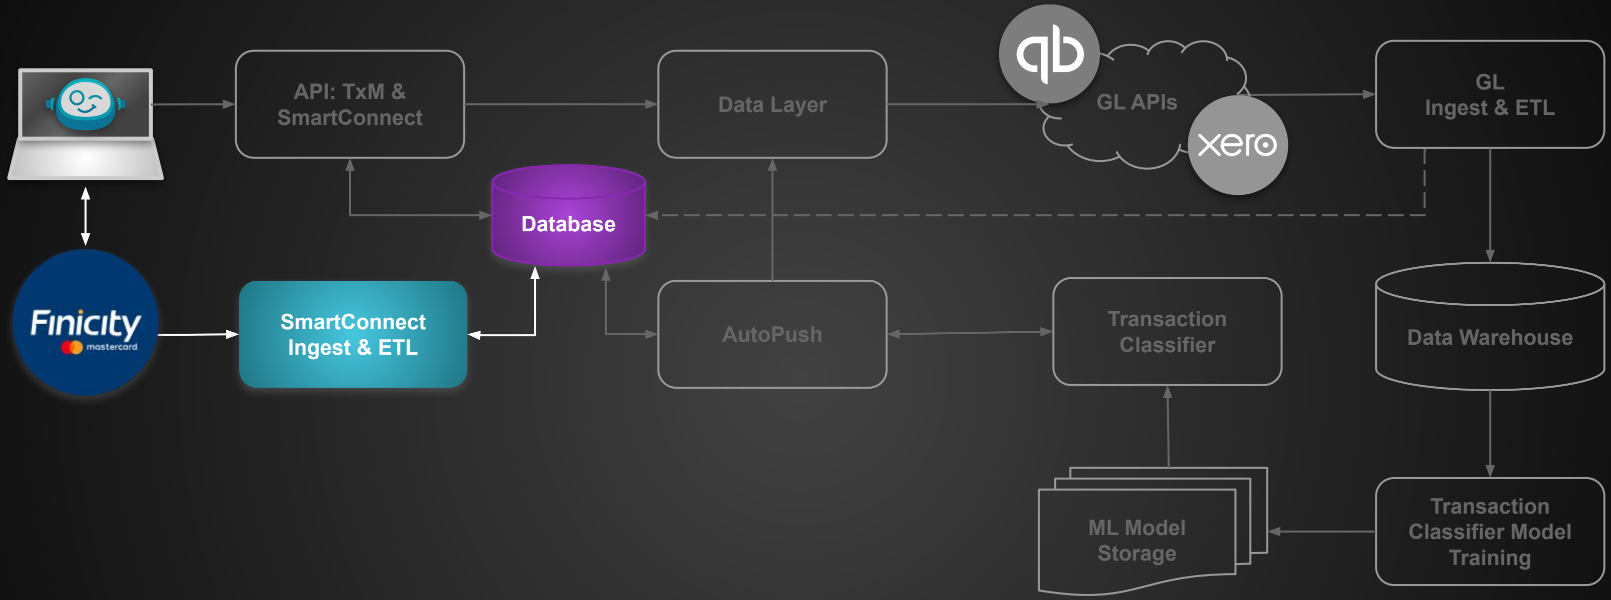

Sourcing Transactions with SmartConnect

SmartConnect, powered by Finicity, a Mastercard company

"Open banking that simply works
Consumer-permissioned data to power your innovation.

When consumers can securely permission their financial data with the utmost confidence—choosing which data to share and who to share it with—a whole new world of opportunity opens. For them and for you."

95%

Coverage of accounts in U.S.

95%

Lead in direct data access connectivity

Vetted

Proven data quality battle-tested by stringent investors

Finicity "implement[s] best-in-class physical, technological and procedural security safeguards similar to those used by major financial institutions (banks, credit card companies, trading firms)."

Sources: https://finicity.com and https://www.finicity.com/security/

SmartConnect Seamlessly Connects to Client Accounts

Pull in your clients' account data from over 16,000 financial institutions across the US and Canada in just a few simple steps.

Smart Connect handles multi-factor authentication security questions and one-time passcodes via text or phone calls — so you can get back to business.

SmartConnect Architecture

1. Connect accounts via Finicity in BOS
2. Sync prior 180 days of transaction history on initial connection
3. Ongoing regular syncs update transaction data from the financial institutions
4. Statements are fetched based on their respective availability dates
5. Data is ingested and landed in our Data Lake
6. Data is extracted, transformed into a normalized data model and loaded into our database for use by the SmartConnect module and AutoPush

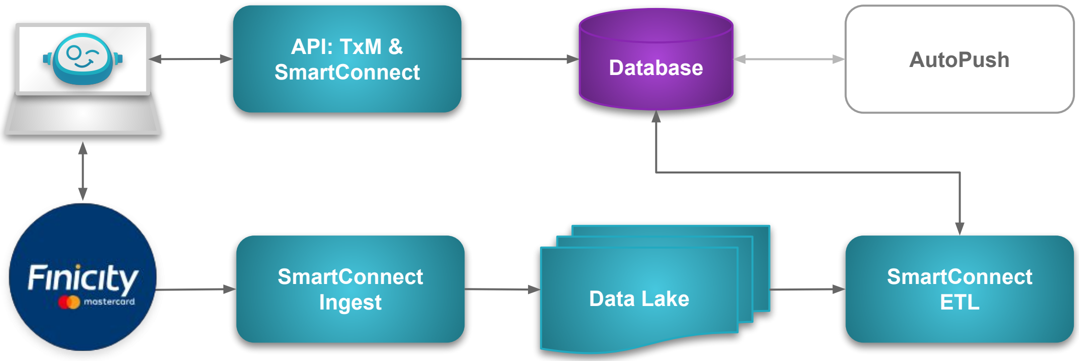

Poll #2

What type of access do you get for your clients' financial accounts?

1. Client provides read-only credentials
2. Client provides their primary credentials
3. Client does not provide credentials
4. Sticky note
5. Email

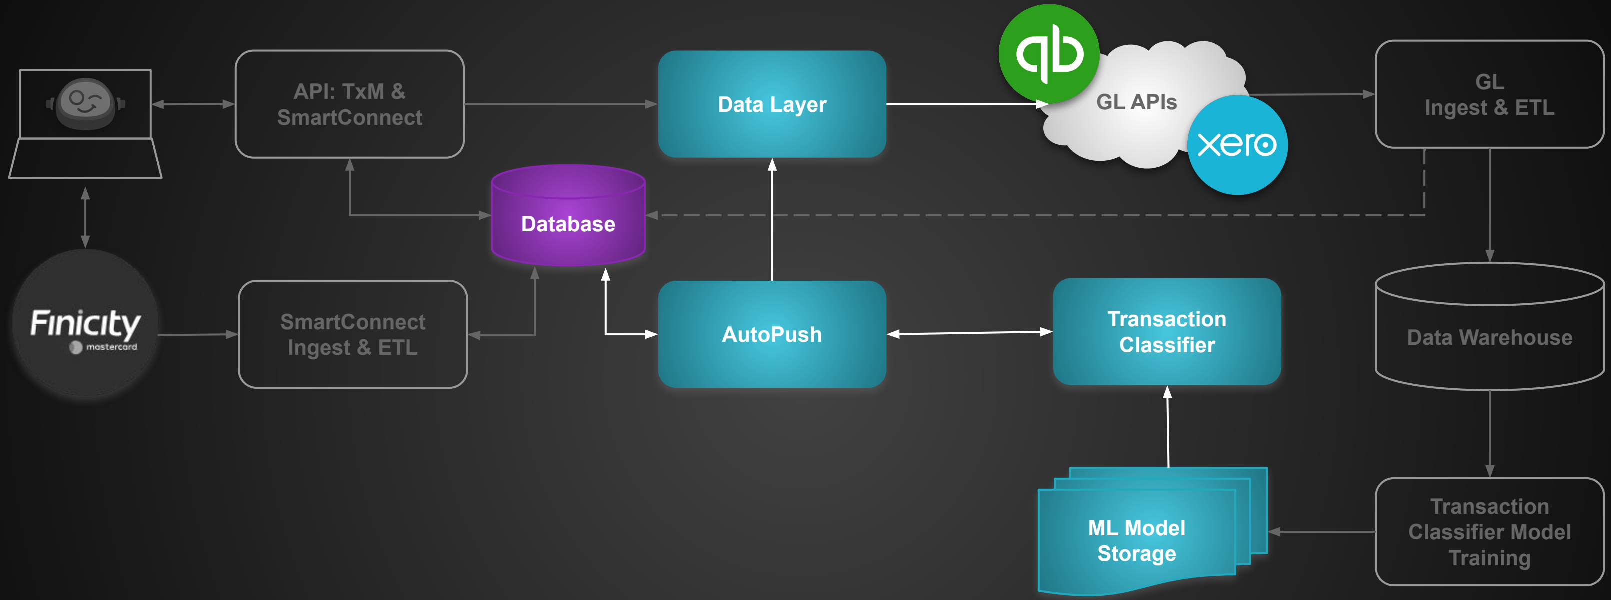

Txn Categorization with Transaction Classifier

Transaction Categorization With Transaction Classifier and Autopush

1

Retrieve Bankfeed Txns

Identify bankfeed txns that are eligible for processing

2

Transform to GL Txns

Format bankfeed txns into Botkeeper's standard GL txn structure

3

Enrich with Txn Classifier

Transaction Classifier provides account/vendor predictions based on trained models.

4

Enrich with Match Data

Matching Service provides suggestions for GL txn matches

5

Land Txns in the GL

For txns meeting all criteria, create the txn in the Gl

6

Surface Txns in TxM

Surface enriched data to the user for review and finalization.

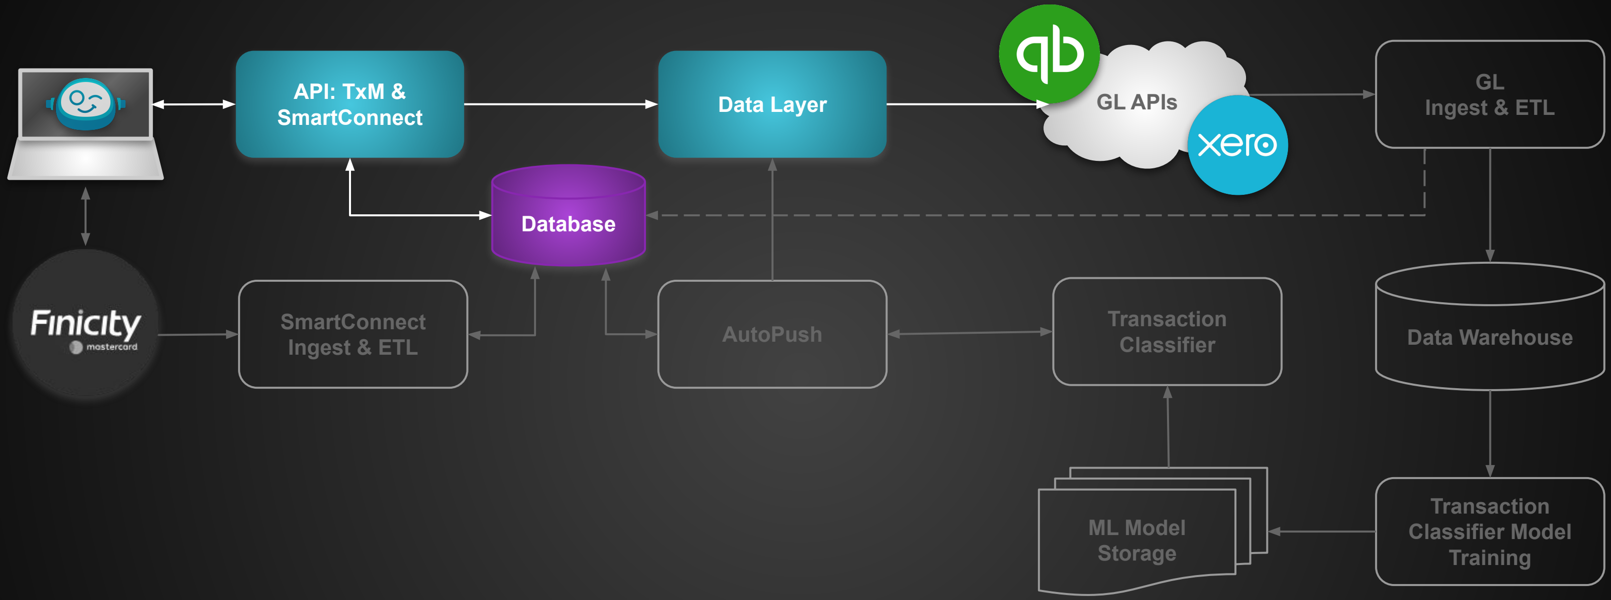

Unifying Disparate GL Interactions

Transaction Manager

- All transactions are visible and available for review on TxM, with the high confidence landing in the “Processed” tab all lower confidence (<98%) on the “Needs Review” tab.
- Inquiries can be assigned from Botkeeper to the Firm to a Client as needed for insight.
- Manage your clients on a single pane of glass - TxM offers a unified user experience across GL’s - currently QBO & Xero with more on the way!
- The models continue to learn as manual adjustments are made, further improving accuracy over time.

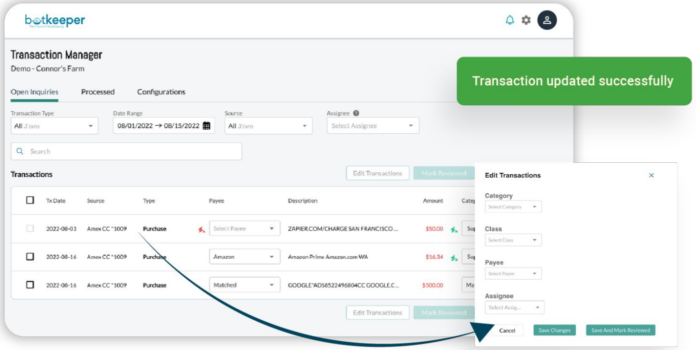

Architecture Summary and Ongoing Training

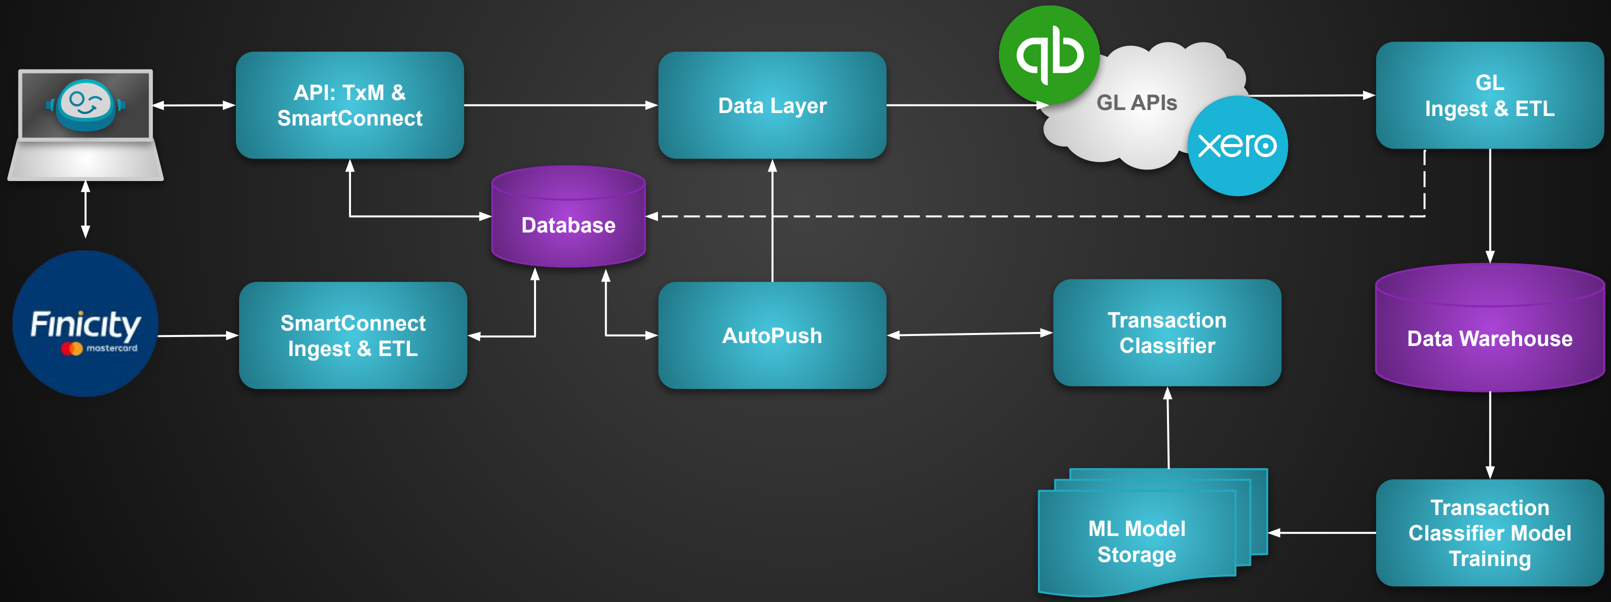

AI & ML improve efficiency but never fully eliminate the need for humans. AutoPush can eliminate over half of transaction categorization work and streamlines the rest.

60%

Average “AutoPushed” Transactions

But why can’t we automate 100%?

Because data is ever changing…

AutoPush Accuracy Continues to Improve

- Botkeeper processes hundreds of thousands of transactions each month
- We continue to see increased volumes of high confidence transactions a 97% accuracy rate
- Most recently, ~33% of low and medium confidence transactions did not require edits

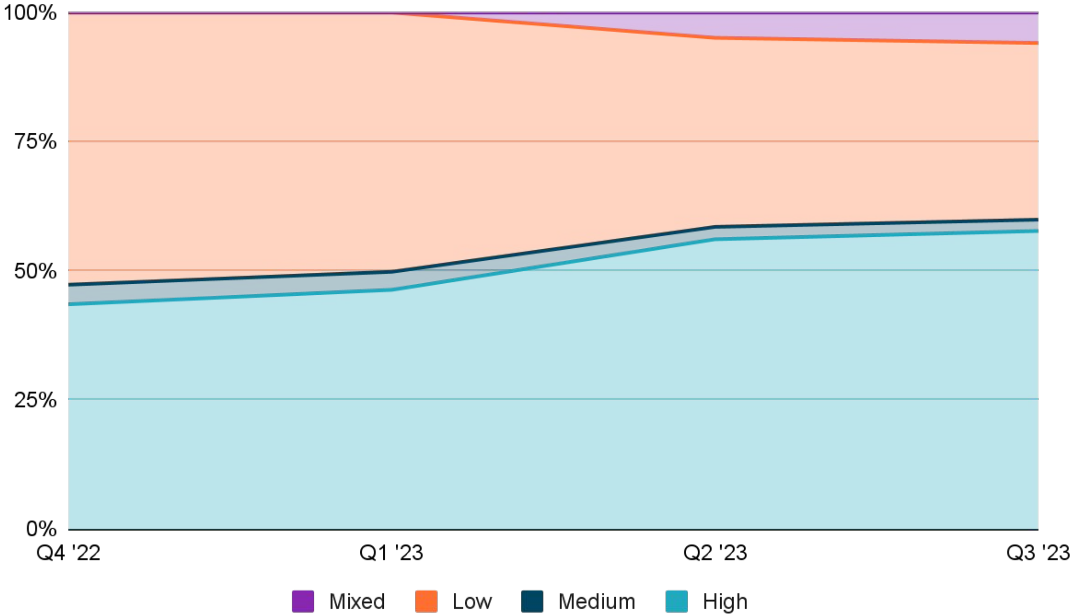

Poll #3

If it were adjustable, what threshold would you set for “high” confidence transactions?

1. > 98%
2. > 90%
3. 80 - 90%
4. 70 - 80%
5. 60 - 70%
6. 50 - 60%

AI UNCHAINED
Accounting Intelligence. No boundaries. No limits.

What's Coming for AutoPush and Transaction Manager

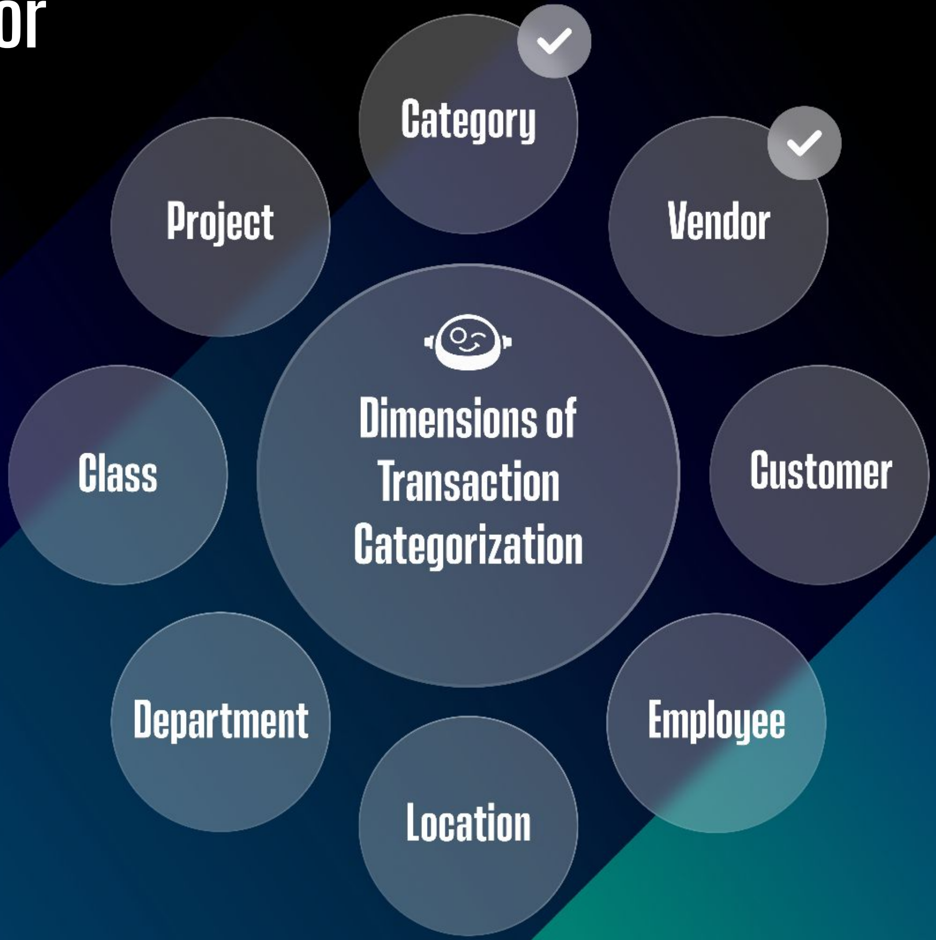

Dimensions of Transaction Categorization

WHO?

- Vendor
- Customer
- Employee

Dimensions of Transaction Categorization

Category

Project

Class

WHAT?

Dimensions of Transaction Categorization

Department

Location

WHERE?

What's Coming for AutoPush and Transaction Manager

- **New GL integrations**
  - Expanding our solutions to new markets
- **ML Customization**
  - Prompt your models to behave how you want
  - Define thresholds for low, medium and high confidence transactions and how the system should handle them
- **Industry-specific models**
  - Optimize ML performance by leveraging data across similar industries
    - Vendors being used by a business in industry A may require different categorization compared to the same vendor for a business in industry B
- **Allocation Automation**
  - Expense splits
  - Intercompany splits
  - Interest/principal splits
- **GenAI predictions**
  - Leverage generative AI to make predictions when AutoPush predictions are low confidence
    - "Hey Google.. what is this vendor for?"

Going Beyond - LLMs and Applicable Use Cases

- LLMs (large language models) are models that focus on language approximation.
- Consider the human brain with different processing centers that are responsible for completing different types of tasks. The language processing center isn't responsible for reasoning tasks or performing complex mathematical calculations.
- Transaction categorization is a task that requires a lot of precision, LLMs can be accurate, but are not precise.
- LLMs can be used to generate replies to emails, write resumes, etc.
- Truly powerful techniques can be found in the intersection of LLMs and other modeling techniques.
- Be sure to attend the general session for some exciting announcements later today!

Poll #4

Is your firm using AI or LLM tech today? If so, how?

The existential concern of AI becoming a replacement for highly skilled professionals is not actually the greatest threat to the industry.

It is those who learn to harness this tech in truly innovative ways.

That is what Botkeeper will do for your firm.

Jenny Blackwell / **jblackwell@botkeeper.com**
Travis Cherry / **tcherry@botkeeper.com**
Brian Emposimato / **bemposimato@botkeeper.com**

THANK YOU!