WHITE PAPER

DESIGN PRINCIPLES

Transaction Categorization

STRANDS

WHITE PAPER

DESIGN PRINCIPLES

Transaction Categorization

STRANDS

Automatic categorization of financial transactions is at the center of any PFM solution. But where do transaction categories come from and how are they automatically assigned? How is success measured and how do we ensure that the system improves over time?

Strands offers this white paper to anyone that wishes to understand the design of an automatic categorization engine.

Key Topics

- **Category Taxonomy Definition**
- **Categorization Rules**
- **Metrics, Monitoring, Maintenance**
- **Categorization Terminology**

How to Categorize Transactions?

Nearly every book, every blog, every advisor that sets out to help you better manage your finances will start the same way: **First, know where your money is going.** If knowledge is power, then today's automated personal finance management (PFM) tools are the engines that deliver it. What was once a tedious task beyond the patience of all but the most dedicated self-improvers is now only a click away for customers having PFM integrated into their online banking. Through interactive pie charts and innovative visualizations users can see if their spending on restaurants, entertainment, housing and the rest are as aligned with their life goals and priorities as they wish. By offering varying levels of granularity with both general and more specific spending and income categories, and by allowing comparisons against spending in these categories among an individual's peer group, PFM gives users actionable insight into their finances. With that they can take control by taking advantage of other PFM features such as creating budgets, setting goals and employing what-if scenarios to see the effects of today's decisions on their future financial health.

It becomes clear that automatic categorization of financial transactions is at the center of any PFM solution, but where do these categories come from and how are they automatically assigned? How is success measured and how do we ensure that the system improves over time?

6 Things to Consider

Define a Category Taxonomy

Before anything, a PFM categorization component requires an initial set of categories. Here we are referring to the fixed set of categories available to all users of PFM from the start. As most PFM applications allow users to add their own custom categories, we don't have to worry about creating a set of categories that will cover all users. Nonetheless, there are **6 things to consider when creating a category taxonomy:**

1

Topology. A taxonomy can be flat or hierarchical. If hierarchical there can be two or more levels. We have found that a hierarchical set of categories with two levels (parent and child categories) offers the best compromise between

simplicity and detail. Customers can get a high-level view of spending by looking at the parent categories while being able to drill down into the details as desired.

2 **Granularity**. Choosing the number of parent categories and, for each, the number of children is again a balancing act between simplicity and detail. Most people find that 10 to12 parents each with no more than 5 or so children works well.

Different granularities for the category 'Food' : simplicity or detail

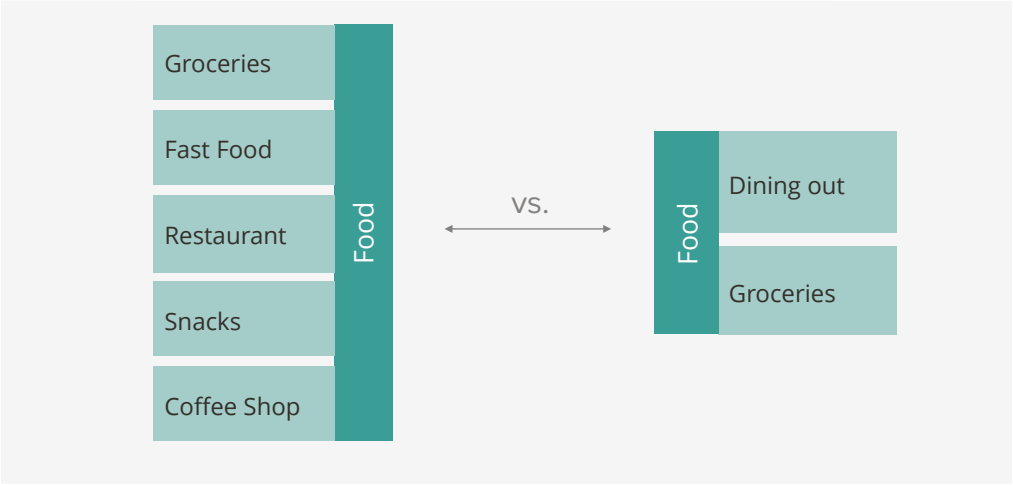

3 **Utility**. Related to granularity is utility. After all, we want to include categories that may be useful and exclude the others. Is a restaurants category enough to cover expenses related to dining out or should there be additional categories such as fast food and coffee shops? As a category may be considered useful if it can generate actionable insights, it's useful to consider use-cases such as:

- *Is someone likely to want to make separate budgets for fast food vs. non-fast food restaurants?*
- *Is someone likely to take a different action after seeing that 80% of dining-out expenses are in restaurants and 20% in coffee shops vs. a 20-80 split in the other direction?*

4 **Business Intelligence**. Categories are useful both for the end user and for the financial institution to analyze their customers' financial behavior. Expenses related to mortgage, rent, insurance, and credit cards are just some of the spending patterns that can be interesting for personalized marketing and overall business strategies. A category taxonomy should include categories of interest to Business Intelligence departments as well.

5 **Grouping**. Sometimes the same categories can be grouped in different ways. Insurance is a good example. You can imagine an insurance parent category with children such as home insurance, car insurance and health insurance. On the other hand, one could equally imagine a home insurance category in a group headed by a home parent for all home related expense, car insurance under an auto parent and health insurance in yet another group with other medical expenses. Once again, use cases help make such decisions.

Some concepts can be sliced horizontally or vertically

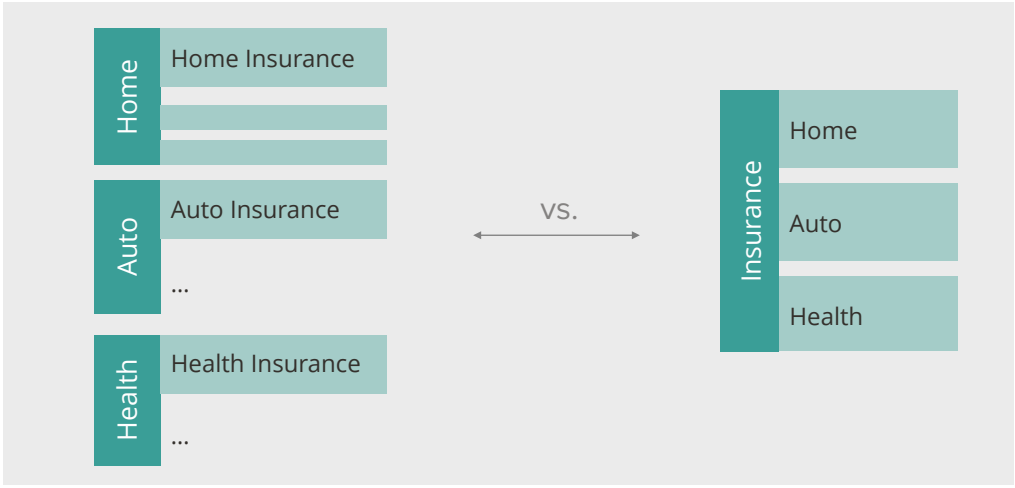

6 **Regionalisms**. Categories that make sense in one part of the world may be less useful in another. In a place where student loans are the norm, a student loan payment category is an important part of being able to plan discretionary spending. Conversely, in a place where a college education is provided at little cost, such a category would not be missed if not present.

Categorization Rules

Strands automatic categorization relies on a set of rules. For each financial transaction, we look for an applicable rule that says something on the order of "transactions like this should get that category". Of course that begs the question, what does it mean for a transaction to be "like this"? Furthermore we need to understand, where do these rules come from and how are they created?

Define Transaction Metadata

Financial transactions accessed by a PFM solution carry various data about the transaction, or metadata. As bank customers ourselves we know there is always the transaction amount and some narrative text describing what the transaction was for. But there is generally much more than that. This metadata is what forms the basis for creating the rules. Thus, an initial task before creating a rule set is to identify available metadata that may be useful for categorization.

Typical metadata we find in financial institutions include:

- **Transaction Code:** Transactions almost always carry some sort of code that classifies it from the bank's point of view. Such codes typically indicate transaction types such as ATM withdrawals, interest payments, bank fees or service charges.

- **Merchant Category Code (MCC):** Credit and debit card transactions almost always carry an MCC, a four-digit number assigned to a business by credit card companies. This code indicates the type of business performed by the merchant at a granularity often sufficient for directly assigning a PFM category.

- **Point-of-Sale (POS) ID:** The POS ID identifies the actual point-of-sale terminal used to process a card transaction. This can help pinpoint a more specific category than just knowing the transaction took place at a large department store, for example.

- **Direct Debit Originator:** Direct debit transactions usually carry an ID or account number identifying the party that is permitted to remove money from the customer's account on a one-time or scheduled basis.

- **Bill pay / transfers counter party:** When a customer arranges to pay an online bill or transfer funds there is often an ID or account number identifying the payment's beneficiary.

- **Clearing and settlement codes:** As credits and debits between financial institutions flow through the various clearing houses there is often a code attached that is useful for identifying transactions such as payroll deposits, mortgage, insurance and tax payments, to name a few.

While the above may be common, the list is by no means exhaustive. Useful metadata may include internal proprietary codes as well as regional and country-specific classifications.

Rule Creation

Once the available metadata has been identified, one can create rules combining one or more of their possible values. Often rules can be very simple, based on a single data value:

- *Transactions with Merchant Category Code 5411 get category 'Groceries'.*

- *Transactions with Transaction Code T04 get category 'Interest Payment'.*

More complicated rules can be fashioned by combining multiple values:

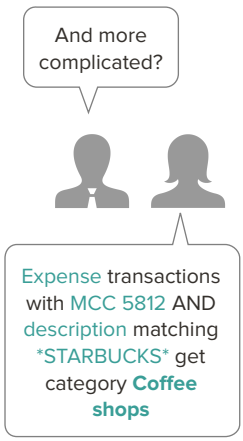

- *Transactions with Merchant Category Code 5812 and Description matching *STARBUCKS* get category 'Coffee Shops'.*

Note also that we can specify the type of matching to perform for each type of metadata. Matches can be exact or pattern-based using wild cards ("*" in the example, above). Pattern matches are especially useful for transaction descriptions, which often contain variable information such as dates, accounts, and point-of-sale details.

3 Basic Priority Principles

Rule Application

Having created a set of rules, categorization is performed by taking each transaction in turn and, for each, finding a rule with matching values and assigning that rule's category to the transaction. But what if there is more than one matching rule and they don't agree on the category to assign? For example, a single transaction could easily match both of these rules:

- *Transactions with Beneficiary Account 912239 get category 'Clothing'.*
- *Transactions with Transaction Code POS get category 'Shopping'.*

As a practical matter, we simply order the rules and the first applicable rule wins. But the ordering task is made much easier when guided by **3 basic principles for favoring some rules over others:**

1. **Precision:** Normally we prefer the rule that gives the more precise category. 'Clothing' is a more precise indication of how money is being spent than the more general category 'Shopping'.

2. **Certainty:** Often the better rule is the one that has a greater chance of giving the correct category. A rule with a Merchant ID known to sell jewelry is better than a rule that looks for DIAMOND in the transaction description. (Imagine the full description may be DIAMOND LAKE HARDWARE EMPORIUM).

3. **Specificity:** Sometimes the information used by one rule is a superset of the other. A rule that includes more metadata is more specific, and it is natural to favor the more specific rule. For example, a rule based on transaction type, beneficiary account and description is better than one with the same values but based only on transaction type and beneficiary account.

Metrics, Monitoring, Maintenance

As with any endeavor we need a way to measure how we are progressing, a way to be sure we are staying on track and, most important perhaps, a way to leverage data as it accumulates to constantly be improving the system over time.

There are two key metrics to look at with respect to categorization:

- **Coverage** is the percentage of all transactions assigned a category after applying the rules.
- **Accuracy** is the percentage of categorized transactions that are correctly categorized.

~87%

Average initial coverage

The first of these, coverage, is straightforward to calculate and can generally be improved by simply making more rules. For an initial rollout, a percentage in the high 80's is a reasonable goal.

Accuracy, though, has to be considered. After all, we can achieve 100% coverage by assigning a random category to every transaction—an extreme example taken to the point of absurdity, but it does make the point. Compared to coverage, unfortunately, it is decidedly more difficult to measure accuracy. It is often up to one or more humans looking at samples of categorized transactions and judging whether the assigned categories are correct. Even worse, it is not always clear what constitutes a "correct" category. Distinctions such as fast food vs. restaurant can be subjective. The most important criteria in the end is that the PFM user maintains confidence in the system. Automatically assigned categories need to make enough sense that the system is perceived as intelligent and consistent.

Improving the System

A reporting system can help make sure categorization is performing as expected. Reports showing the coverage are produced for the most recently loaded transactions and over recent time periods such as the latest week, month, and quarter. If coverage is declining one needs to look for newly popular spending habits and the corresponding metadata that are not yet considered by the rules.

Improvements through data analysis

By making the reporting system interactive, categorization managers can explore the PFM data to improve the system with new and/or improved rules. To improve coverage, reporting should show the most common values among selected metadata within a set of uncategorized transactions. For example, by exploring the data we might find that the most common Originator Account ID among uncategorized direct debit transactions is 312178, and that adding a new rule for that value could materially improve the coverage.

Improvements through collaborative intelligence

As the PFM system allows users to manually set and change categories for their transactions, reports based on community and individual user behavior are also used to improve both coverage and accuracy. When categorizing we keep a record of the rule used to categorize each transaction. Thus, we can report on the most common rules for which users have modified the automatically assigned category.

Such a modification doesn't necessarily indicate a problem with the rule—a user who shopped at a general grocery store only to buy pet supplies might change the category accordingly even while still considering the originally assigned category to be accurate from the system's point of view. On the other hand, such rules could have suffered a clerical error during their creation or a misunderstanding about how PFM users perceive the classification of the corresponding vendor. When there is a consensus in how users are changing such categories, the system can recommend how the rule might best be modified. Similarly, seeing how users categorize initially uncategorized transactions make it easy to suggest new rules that could be added to the system.

Conclusion

Automatic categorization of financial transactions is a cornerstone of any personal finance management system. It relies on a manageable taxonomy of categories that in turn provides actionable insights for both the institution offering the PFM and their customers. A set of rules for assigning categories based on available transaction metadata can be efficiently applied with conflicts resolved by choosing the rule that makes the best use of the available information. An interactive reporting mechanism allows the system to be monitored and improved over time.

STRANDS

Categorization Terminology

PFM	Personal Financial Management, refers to the managing of money for individuals and families. Expense tracking, budgeting and planning are all key aspects of PFM.
Categorization	Process by which account operations are classified according to the purpose of the transaction.
PFM Taxonomy	Pre-determined set of spending and income categories to be assigned to account operations during categorization.
Custom Category	Spending or income category not included in the taxonomy that can be created by a PFM system user to uniquely categorize account operations according to his or her personal requirements.
Metadata	In the general sense metadata is data that describes content. For example: the number of words and printing date of this paper. In the PFM domain, an example includes the transaction date of an account operation.
Pattern Match	A pattern match is a way of matching text characters. It occurs when certain, but not all, characters correspond with a sample. This in contrast to exact matching, where all characters need to match. For example, ‘BURGER’, matches ‘BURGER KING’ and ‘HAMBURGER’ in a pattern match .
Process Metric	Statistical measurement used to monitor the success or performance of a given task.
Collective Intelligence	Learning technique that analyses the behavior of individuals in a group to improve the overall performance of a system.

STRANDS

DESIGN PRINCIPLES

ABOUT STRANDS

Strands develops innovative software that enables banks to offer personalized digital banking experiences. The Strands Finance Suite includes a comprehensive set of white-label solutions such as **PFM**, **BFM**, **CLO** & **REC**. In 2008, Strands revolutionized online banking by deploying the first PFM in Europe.

From our Barcelona HQ and offices in San Francisco, Miami, Madrid & Buenos Aires, we serve top-tier banks including Barclays, Deutsche Bank, BBVA, BNP Paribas, Bank of Montreal (BMO), & PostFinance.

BARCLAYS

Deutsche Bank

BBVA

AIB

PostFinance

BNP PARIBAS

BMO