Sawtooth Software

RESEARCH PAPER SERIES

Understanding Item Contribution in TURF Analysis

Bryan Orme
Sawtooth Software, Inc.

3210 N. Canyon Rd., Provo, Utah
+1 801 477 4700
www.sawtoothsoftware.com

Understanding Item Contribution in TURF Analysis

Bryan Orme, Sawtooth Software

April 2025

When seeing results of a TURF (Total Unduplicated Reach and Frequency) analysis, clients often ask: *How much does each item contribute to the success of the portfolio? (i.e., the set or product line of items).* For example, in a set of ice cream flavors {Cherry, Chocolate, Coconut}, which collectively reach 97% of customers, how much does each individual flavor contribute?

Answering this seemingly simple question is more computationally intensive than it may appear, as we explain below.

Review of TURF

(Feel free to skip this section if you're already familiar with TURF.)

TURF analysis helps identify the optimal combination of items—such as product features, SKUs, or flavors—that maximizes reach across a customer base. A classic example involves a grocer with freezer space for only five ice cream flavors. Which five should they stock to ensure that the maximum number of customers find at least one flavor they would buy?

It's not as straightforward as selecting the five most popular flavors overall. A flavor with niche appeal may significantly expand reach even if it isn't among the top five favorites in the general population.

Suppose there are 21 total flavors included in a survey using a 5-point purchase intent scale or anchored MaxDiff. For each respondent, we can determine whether they would purchase (i.e., be *reached by*) each flavor. The challenge is that there are 20,349 possible combinations of five flavors out of 21, calculated as:

$$\binom{21}{5} = \frac{21!}{5!(21-5)!} = 20,349$$

TURF optimization evaluates all these combinations and reports the top sets (portfolios) that achieve the highest reach.

Because several portfolios may have similar reach percentages, researchers often supplement the analysis with frequency counts—i.e., how many items in a portfolio reach each respondent. For example, if a respondent would buy 4 of the 5 flavors in a portfolio, their frequency is 4. Summing these across respondents provides further insight into which portfolios might lead to the most sales.

Why *Order of Entry* Matters

Back to the question of item contribution to a specific portfolio (item lineup): one intuitive way to estimate each item's contribution is by analyzing marginal reach in sequence. For instance:

- How many customers are reached by Cherry alone?
- How many more are reached when Chocolate is added (portfolio: {Cherry, Chocolate})?
- What's the additional reach when Coconut is included (portfolio: {Cherry, Chocolate, Coconut})?

However, this approach is order-dependent. Changing the order of flavor entry—e.g., starting with Coconut instead of Cherry—can yield significantly different results.

To account for this, we use Shapley Value Attribution, which averages the marginal contribution of each item across all possible orderings. This method offers a fair and mathematically grounded way to understand each item's value.

Example Calculation

Let's say the portfolio {Cherry, Chocolate, Coconut} reaches 97.44% of respondents. For three flavors, there are 3! = 6 possible entry orders:

1. Cherry → Chocolate → Coconut
2. Cherry → Coconut → Chocolate
3. Chocolate → Cherry → Coconut
4. Chocolate → Coconut → Cherry
5. Coconut → Cherry → Chocolate
6. Coconut → Chocolate → Cherry

For each ordering:

- We compute the initial contribution of the first flavor (i.e., in a portfolio of size 1).
- Then we calculate the marginal contribution of the second (i.e., in a portfolio of size 2).
- Finally, we compute the marginal contribution of the third (i.e., in a portfolio of size 3).

	Cherry	Chocolate	Coconut
Order 123	59.42%	36.74%	1.28%
Order 132	59.42%	25.24%	12.78%
Order 213	4.79%	91.37%	1.28%
Order 231	0.96%	91.37%	5.11%
Order 312	20.13%	25.24%	52.08%
Order 321	0.96%	44.41%	52.08%
Average:	24.28%	52.40%	20.77%
Sum:			97.44%

From this, we see that **Chocolate** has the highest average contribution to the total reach (52.40%), while **Coconut** has the lowest (20.77%).

Scaling Up: Shapley Analysis for All 21 Flavors

Clients may also ask: *What's the average contribution of each of 21 flavors for all possible portfolios?*

To answer that, let's start with the case of size 3 portfolios. There are 1,140 combinations of size 3 portfolios out of 20 flavors that don't involve item 1. We compare the reach of those 1,140 portfolios that don't include item 1 to the 1,140 matched portfolios that add item 1 (portfolios of size 4). The average of those differences in reach is the Shapley value.

The full Shapley value calculation addresses how much additional reach an item adds to all possible portfolios that don't originally include that item. Sometimes researchers limit the size of the problem by computing Shapley values over a reduced range of portfolios sizes, such as across portfolios of size 3 to 5.

By aggregating and normalizing these values, we derive Shapley values that sum to 100%, representing each flavor's *average contribution across all analyzed portfolio sizes*.

When to Use Which Method

Use TURF when:

- Your goal is to create a lineup of items that maximizes total reach.
- You don't need to isolate the value of individual items.

Use Shapley Value Attribution when:

- You want to understand each item's individual value considering its synergies with others.
- You care about marginal contributions and product synergy.
- You aim to optimize your product line not just for reach, but for strategic insight.

Technical Notes

For **Shapley Value Attribution** (i.e., how much does each item contribute to the reach of a specific portfolio of items), the number of order of entry permutations grows exponentially with the size of the portfolio. With just three items in a specific portfolio, it's just 3! = 6 possible orders of entry to analyze across. But, for a portfolio of nine items, it's 9! = 362,880. Then, within each of those permutations, we need to calculate the incremental reach for the first item, the first two items, etc. The computational burden explodes geometrically.

Our colleague, Dave Lyon (Aurora Market Modeling) pointed out to us that there are order of entry experimental designs that involve computing portfolio reach across a very manageable reduced set of experimentally balanced orders of entry. While estimating Shapley Value Attribution across a small, but carefully drawn, subset of all permutations does not lead to exactly the same result as when computing across all permutations, the results should be extremely close and all that is needed in practice.

For generalized **Shapley Value Calculation** (i.e., how much does each item contribute to all possible portfolios of, say, size 3, 4, and 5), it can become very computationally intensive as well as the total number of items in the study expands. When the number of permutations exceeds a certain threshold, the researcher could use a BIBD design (or near-perfect BIBD design, such as used in Sawtooth's MaxDiff designer) to select a manageable and balanced subset of all possible permutations to simulate. In such designs, each item appears nearly an equal number of times and each item appears with every other item nearly an equal number of times.