

Peer-Review

Poludas, Harshitha. 2026. "FoodByte: A Personalized Food Scoring and Additive Risk Assessment Algorithm for Smarter Eating." *Journal of High School Science* 10 (3): 265–95. <https://doi.org/10.64336/001c.165457>.

Good work. Well thought out, well researched and well presented. Congratulations. My comments have to do with correct statistical model use and reporting, sensitivity/robustness studies and access to raw data and code.

1. Use hierarchical or clustered statistical models that respect repeated measures (products $\times$ profiles)": Right now, your analysis (a simple paired t test) treats every food $\times$ profile combination as completely independent. But the same food is used across multiple profiles, and the same profile is applied to multiple foods—these are linked/repeated measurements. Therefore, you need to use statistics that account for these dependencies (e.g., mixed-effects models) to get more accurate conclusions.

A hierarchical or mixed-effects model separates the variation in your data into fixed effects and random effects: Fixed effects include FoodByte vs Nutri-Score, Effect of food category (snacks vs whole grains), Effect of health profile (A, B, C)

Random effects: These account for the correlated/ repeated measurements, like Each food $\rightarrow$ multiple profiles and Each profile $\rightarrow$ multiple foods. By modeling random effects, you don't overestimate confidence due to repeated measures.

Using the hierarchical or mixed-effect model lets you Test whether FoodByte really differs from Nutri-Score while accounting for repeated measures. See how much variation is due to food vs profile (random effects). Make predictions for new foods or profiles.. Get confidence intervals that are more realistic than simple t-tests ignoring dependencies.

Your statistics teacher can help you with this.

2. When you report these statistics, please make sure you also report effect sizes and confidence intervals.

3. Perform sensitivity analyses on your data. Deliberately change your assumptions, recompute the FoodByte scores, and check whether your main conclusions still hold. In other words, confirm that the result of all your particular weights, thresholds, additive effects is not a fluke for this particular combination. As examples, then you can say "'Scores were highly correlated ($r > 0.9$) across $\pm 50\%$ changes in additive weights.'" or, "Only 6% of foods changed risk category under threshold perturbation." or "Across all aggregation rules tested (simple sum, capped sum, class-weighted sum, max-risk additive only), FoodByte scores remained significantly lower than Nutri-Score for ultra-processed snacks."

Note, that you will need to calculate these differently aggregated sums (simple, capped, class-weighted and max-risk) for this part of your analysis.

Or apply a different penalty for (say) CVD profile and then you can state "Personalization effects were robust across a wide range of profile-weight multipliers."

I am not suggesting you do this for every profile and food; just those that you can justify would cause the most change in scores when perturbed.

You can then show a table or a heat map with Spearman/Pearson correlations, Rank overlap (% same top 10), Mean absolute score change, % of foods changing category.

4. I appreciate that you have included the raw data and code in your submission. However, please also make a public GitHub repository and deposit your data there with a link to it included in the revised manuscript.

I look forward to your revised manuscript.

FoodByte: A Personalized Food Scoring and Additive Risk Assessment Algorithm for Smarter Eating

Reviewer Response

COMMENT #1: Statistical Methods - Hierarchical/Clustered Models Comment:

“Use hierarchical or clustered statistical models that respect repeated measures (products × profiles): Right now, your analysis (a simple paired t test) treats every food × profile combination as completely independent. But the same food is used across multiple profiles, and the same profile is applied to multiple foods—these are linked/repeated measurements.

Therefore, you need to use statistics that account for these dependencies (e.g., mixed-effects models) to get more accurate conclusions. A hierarchical or mixed-effects model separates the variation in your data into fixed effects and random effects: Fixed effects include FoodByte vs Nutri-Score, Effect of food category (snacks vs whole grains), Effect of health profile (A, B, C) Random effects: These account for the correlated/repeated measurements, like Each food → multiple profiles and Each profile → multiple foods. By modeling random effects, you don't overestimate confidence due to repeated measures. Using the hierarchical or mixed-effect model lets you Test whether FoodByte really differs from Nutri-Score while accounting for repeated measures. See how much variation is due to food vs profile (random effects). Make predictions for new foods or profiles. Get confidence intervals that are more realistic than simple t-tests ignoring dependencies. Your statistics teacher can help you with this.”

Response:

I addressed this concern by implementing a linear mixed-effects model that properly accounts for the repeated measures structure in the data.

1. Section 2.3 (Statistical Analysis) - Completely Rewritten:

- a. I replaced the paired t-test methodology with a comprehensive mixed-effects modeling framework. The new text explicitly describes:
 - i. The repeated measures problem (same products evaluated across multiple profiles)
 - ii. Why treating observations as independent would be inappropriate
 - iii. The mixed-effects model structure (random intercept for product, fixed effect for scoring system)
 - iv. How the model separates variance into product-level and profile-level components

v. The calculation and interpretation of the intraclass correlation coefficient (ICC)

“Descriptive statistics were computed to summarize score differences between FoodByte and Nutri-Score, including mean percentage difference, standard deviation, minimum, maximum, and range across all food–profile evaluations. Category-specific and profile-specific summaries were calculated descriptively to characterize structural patterns in score divergence across product types and dietary profiles.

Because each food product was evaluated under multiple dietary profiles, observations were not independent. Treating all food-profile combinations as independent would underestimate uncertainty and inflate statistical significance. To account for this repeated-measures structure, inferential analysis was conducted using a linear mixed-effects model with the score difference (FoodByte minus Nutri-Score) as the outcome and a random intercept for product. This model estimates the overall mean difference between scoring systems while accounting for correlation induced by repeated evaluation of the same product. Results are reported with 95% confidence intervals, and the intraclass correlation coefficient (ICC) was calculated to quantify the proportion of variance attributable to between-product differences.

Standardized effect sizes (Cohen’s d) with 95% confidence intervals were computed separately for each dietary profile to quantify the magnitude of divergence between scoring systems. These effect sizes are reported descriptively to contextualize the practical magnitude of observed differences rather than as independent hypothesis tests.

Pearson correlation coefficients were calculated to describe the relationship between FoodByte and Nutri-Score values. Correlation analyses are presented descriptively and do not assume independence of observations.

To assess robustness to algorithmic specification, sensitivity analyses were conducted by perturbing additive scoring weights ($\pm 50\%$), threshold values ($\pm 20\%$), and profile-specific multipliers for cardiovascular disease penalties (0.5x to 2.0x). For each perturbation, stability of score rankings and category assignments was evaluated using Pearson correlations between perturbed and original scores. The percentage of products maintaining their letter-grade category under threshold perturbations was calculated to assess categorical stability.”

2. Section 3.1 (Results) - New Statistical Results:

a. I replaced the paired t-test results with mixed-effects model results:

i. OLD: $t(249) = 13.993$, $p = 3.23 \times 10^{-33}$

ii. NEW: Mean difference: -10.49 points (95% CI: [-13.07, -7.90]), t-statistic: -7.952, $p = 1.84 \times 10^{-15}$, ICC: 0.518

b. Added ICC interpretation (2 new paragraphs):

i. Explains that 51.8% of variance is due to product differences

ii. Explains that 48.2% is due to profile-specific adjustments

iii. Discusses implications for generalization to new products

iv. “A linear mixed-effects model accounting for repeated measures confirmed that the observed differences were highly statistically significant:

- Mean difference: -10.49 points (95% CI: [-13.07, -7.90])

- t-statistic: -7.952

- $p = 1.84 \times 10^{-15}$ ($p < 0.001$)

- Intraclass correlation (ICC): 0.518

These results decisively reject the null hypothesis of equivalence between generic and personalized scoring approaches. The mean difference of -10.49 points indicates that FoodByte systematically assigns lower scores than Nutri-Score when accounting for additive burden and disease-specific risk factors. Even when averaged across all product categories and health profiles, FoodByte produces materially different evaluations than Nutri-Score. The magnitude of divergence substantially exceeds the originally hypothesized 20% discrepancy threshold, confirming that one-size-fits-all

nutritional scoring fails to adequately represent health risks across heterogeneous consumer populations.

The intraclass correlation ($ICC = 0.518$) indicates that 51.8% of the variance in score differences is attributable to systematic differences between products, with the remainder (48.2%) due to profile-specific adjustments and residual variation. This substantial ICC confirms that product identity, particularly the presence of multiple additives and baseline nutritional quality, strongly influences the magnitude of personalization effects. Products with high additive burden show consistently large downgrades across all profiles, while minimally processed foods show smaller, more variable adjustments depending on disease-specific nutritional criteria.

The mixed-effects modeling framework employed in this analysis also provides a principled basis for generalization beyond the observed sample. By modeling product identity as a random effect, the model estimates population-level variance in personalization effects rather than fitting idiosyncratic differences tied only to the 50 evaluated foods. This structure enables inference about how new, previously unseen products would be expected to behave under personalized scoring, conditional on their nutritional composition and additive burden.

Similarly, because health profiles are defined by rule-based threshold and penalty structures rather than learned parameters, FoodByte scores for new or modified profiles can be generated deterministically without refitting the model. Together, these features support extrapolation to novel foods and profiles while maintaining statistically conservative confidence intervals that respect repeated-measure dependencies.”

3. Section 2.4 (Implementation) - Updated Software:

- a. Added statsmodels library for mixed-effects modeling
- b. Updated workflow to reflect new statistical approach

“statsmodels: Linear mixed-effects models accounting for repeated measures structure.”

4. Section 4.1 (Discussion) - Updated Interpretation:

- a. First paragraph now references mixed-effects results instead of paired t-test

“Across 250 product–profile combinations, FoodByte diverged from Nutri-Score by an average of 29.18%, with mixed-effects modeling accounting for repeated measures confirming highly significant differences (mean difference: -10.49 points, 95% CI: [-13.07, -7.90], $p < 0.001$).”

5. Section 5 (Conclusion) - Updated Summary:

a. Now references mixed-effects modeling

“Mixed-effects modeling accounting for repeated measures confirmed highly significant differences (mean: -10.49 points, 95% CI: [-13.07, -7.90], $p < 0.001$), with four of five health profiles demonstrating large effect sizes (Cohen's $|d| > 0.8$). Sensitivity analyses showed robust product rankings ($r > 0.983$) across algorithmic perturbations.”

Why this approach is appropriate:

- Product as a random effect allows generalization to new foods beyond sample
- Properly accounts for clustering/correlation in the data
- Provides more realistic (conservative) confidence intervals
- ICC quantifies how much variation is structural (product-dependent) vs. personalization-driven

COMMENT #2: Effect Sizes and Confidence Intervals Comment:

“When you report these statistics, please make sure you also report effect sizes and confidence intervals.”

Response:

I addressed this concern by computing and reporting Cohen's d effect sizes with 95% confidence intervals for all health profiles.

1. Section 2.3 (Statistical Analysis) - Added Effect Size Methodology:

a. Description of effect size calculation using bootstrap confidence intervals

“Standardized effect sizes (Cohen's d) with 95% confidence intervals were computed separately for each dietary profile to quantify the magnitude of divergence between scoring systems. These effect sizes are reported descriptively to contextualize the practical magnitude of observed differences rather than as independent hypothesis tests.”

2. Section 3.2 (Results) - Restructured Table 1:

a. Transformed Table 1 from a simple 3-column table to a comprehensive 7-column table

OLD Table 1

Continues on next page

Table 1. Mean Percentage Difference by Health Profile

Profile	Mean Diff (%)	Min (%)	Max (%)
Control (Baseline)	11.6	0	83.89
Diabetes + Weight Loss	35	0	97.30
Hypertension + CVD	41.6	0	97.74
IBS + Lactose Intolerance	33.2	0	97.74
CKD + Asthma	24.4	0	97.16

NEW Table 1:

Table 1. Effect Sizes and Statistical Comparisons by Health Profile

Profile	n	Mean Diff (SD)	95% CI	Cohen's d	95% CI (d)	Mean % Diff (SD)	95% CI (%)
Control (Baseline)	50	-4.92 (8.68)	[-7.39, -2.45]	-0.567	[-0.740, -0.414]	11.6 (22.8)	[5.1, 18.1]
Diabetes + Weight Loss	50	-11.80 (11.48)	[-15.07, -8.54]	-1.028	[-1.378, -0.769]	35.0 (33.9)	[25.4, 44.7]
Hypertension + CVD	50	-14.16 (14.83)	[-18.38, -9.95]	-0.955	[-1.332, -0.673]	41.6 (34.6)	[31.8, 51.5]
IBS + Lactose Intolerance	50	-12.64 (10.50)	[-15.62, -9.66]	-1.204	[-1.538, -0.979]	33.2 (33.6)	[23.7, 42.8]
CKD + Asthma	50	-8.91 (11.03)	[-12.04, -5.77]	-0.808	[-1.063, -0.604]	24.4 (28.0)	[16.5, 32.4]

- b. Each profile now reports:
- i. Sample size (n = 50)
 - ii. Mean difference in points with standard deviation
 - iii. 95% confidence interval for mean difference
 - iv. Cohen's d effect size
 - v. 95% confidence interval for Cohen's d (via bootstrap with 10,000 iterations)
 - vi. Mean percentage difference with standard deviation
 - vii. 95% confidence interval for percentage difference

3. Section 3.2 (Results) - Added Effect Size Interpretation:

- a. Complete paragraph interpreting effect sizes

“Effect size analysis revealed that personalization produces substantial practical differences beyond statistical significance. Four of five profiles demonstrated large effect sizes (Cohen's $|d| > 0.8$), indicating that FoodByte's personalized scoring produces meaningfully different evaluations from Nutri-Score for disease-specific populations. The Control profile showed a

medium effect ($d = -0.567$), attributable primarily to additive penalties applied uniformly across all profiles.”

4. New Figure 2 - Effect Sizes Visualization:

- a. Entirely new two-panel figure showing:
 - i. Panel A: Cohen's d values with 95% confidence intervals and reference lines for small/medium/large effects
 - ii. Panel B: Mean percentage differences with 95% confidence intervals and 20% hypothesis threshold line




Figure caption:

“Figure 2. Effect sizes and percentage differences with 95% confidence intervals across health profiles. The first plot demonstrates Cohen's d values with bootstrap confidence intervals (10,000 iterations). The reference lines indicate small ($d = 0.2$), medium ($d = 0.5$), and large ($d = 0.8$) effect thresholds. Four of five profiles demonstrate large effects ($|d| > 0.8$), with IBS + Lactose Intolerance showing the largest effect ($d = -1.204$).”

I inserted this figure before the former figure 2, shifting them all down. The former figure 2 is now figure 3.

5. Section 2.4 (Implementation) - Updated Workflow:

- a. Workflow item 3.b
- b. Figure 2 listed in visualization workflow

Exact text:

“b. Figure 2: Effect sizes (Cohen's d) with 95% confidence intervals and percentage differences across health profiles.”

6. Section 2.4 (Implementation) - Updated Libraries:

- a. Statistical Analysis item 4.c
- b. Bootstrap methodology for effect sizes

“c. Effect sizes: Computed Cohen's d with bootstrap confidence intervals (10,000 iterations) for each health profile.”

7. Section 4.1 (Discussion) - New Paragraph on Effect Sizes:

a. Entire new second paragraph discussing practical significance

“Effect size analysis demonstrated that these differences represent substantial practical significance beyond statistical significance. Four of five health profiles exhibited large standardized effects (Cohen's $|d| > 0.8$), with IBS + Lactose Intolerance showing the largest effect ($d = -1.204$, 95% CI: $[-1.538, -0.979]$). Even the Control profile, which received only additive penalties without disease-specific adjustments, demonstrated a medium effect ($d = -0.567$, 95% CI: $[-0.740, -0.414]$).”

Key Results:

- Control: $d = -0.567$ (medium effect)
- Diabetes + WL: $d = -1.028$ (large effect)
- HTN + CVD: $d = -0.955$ (large effect)
- IBS + Lactose: $d = -1.204$ (large effect)
- CKD + Asthma: $d = -0.808$ (large effect)

All confidence intervals reported using bootstrap with 10,000 iterations.

Summary: Effect sizes are now reported in Table 1, visualized in Figure 2, interpreted in Results section 3.2, and discussed in Discussion section 4.1.

COMMENT #3: Sensitivity Analyses Comment:

“Perform sensitivity analyses on your data. Deliberately change your assumptions, recompute the FoodByte scores, and check whether your main conclusions still hold. In other words, confirm that the result of all your particular weights, thresholds, additive effects is not a fluke for this particular combination. As examples, then you can say 'Scores were highly correlated ($r > 0.9$) across $\pm 50\%$ changes in additive weights.' or, 'Only 6% of foods changed risk category under threshold perturbation.' or 'Across all aggregation rules tested (simple sum, capped sum, class-weighted sum, max-risk additive only), FoodByte scores remained significantly lower than Nutri-Score for ultra-processed snacks.' Note, that you will need to calculate these differently aggregated sums (simple, capped, class-weighted and max-risk) for this part of your analysis. Or apply a different penalty for (say) CVD profile and then you can state 'Personalization effects are robust across a wide range of profile-weight multipliers.' I am not suggesting you do this for every profile and food; just those that you can justify would cause the most change in scores when perturbed. You can then show a table or a heat map with Spearman/Pearson correlations, Rank overlap (% same top 10), Mean absolute score change, % of foods changing category.”

Response:

I have substantially addressed this concern by conducting three comprehensive sensitivity analyses that test robustness to parameter perturbations. I provide detailed justification for why I did not test alternative aggregation methods, as the existing perturbations already stress-test the effective behavior of the aggregation mechanism.

New Section 3.7: Sensitivity and Robustness Analyses (~1,500 words)

I added a complete new results section with four subsections addressing different types of sensitivity:

3.7.1 Additive Weight Perturbation

- What I tested: Scaled all additive risk scores by $\pm 50\%$ and recalculated FoodByte scores
- Results reported:
 - -50% additive weights: mean $r = 0.987$ (range: $0.983\text{--}0.992$)
 - $+50\%$ additive weights: mean $r = 0.993$ (range: $0.990\text{--}0.996$)
 - All correlations exceeded $r = 0.983$

“Additive risk scores represent relative ordinal rankings based on scientific evidence strength. To assess sensitivity to these assignments, I recalculated all FoodByte scores with additive penalties scaled by factors of $0.5\times$ (50% reduction) and $1.5\times$ (50% increase). Pearson correlations between perturbed and original scores remained very high across all five profiles... All correlations exceeded $r = 0.983$, indicating that product rankings remain highly stable even under substantial changes to additive penalty magnitudes.”

Conclusion: Product rankings are robust to large changes in additive weights.

3.7.2 Threshold Perturbation

- What I tested: Perturbed all disease-specific thresholds (sodium limits, sugar limits, etc.) by $\pm 20\%$
- What I measured: Changes in letter-grade categories
- Results reported:
 - 90.4% of products maintained original letter grade
 - 9.6% changed by one category (e.g., $C \rightarrow D$)
 - 0% changed by more than one category

“Disease-specific thresholds (e.g., sodium limits for hypertension, sugar limits for diabetes) reflect clinical guideline targets but may require adjustment based on individual circumstances. To test categorical stability, I perturbed all thresholds by $\pm 20\%$ and assessed changes in letter-grade categories... 90.4% of products maintained their original letter grade, 9.6% changed by one category (e.g., $C \rightarrow D$ or $D \rightarrow C$), [and] No products changed by more than one category.”

Conclusion: Categorical recommendations are robust to moderate threshold adjustments.

3.7.3 Profile-Specific Sensitivity: Cardiovascular Disease

- What I tested: Varied CVD profile penalty multipliers from 0.5× to 2.0× (representing conservative to aggressive adjustments)
- New Table 3 added: Shows mean scores and correlations for each multiplier
- Results reported:
 - All correlations $r > 0.978$
 - Product rankings highly stable across entire range

Table 3:

Table 3. Cardiovascular Disease Profile Sensitivity to Penalty Multipliers

Multiplier	Mean Score	Correlation
0.50x	46.87	0.983
0.75x	43.33	0.997
1.00x	39.79	1.000
1.25x	36.25	0.998
1.50x	32.71	0.992
2.00x	25.63	0.979

“To illustrate profile-specific robustness, I examined the HTN + CVD profile under varying penalty multipliers from 0.5x to 2.0x, representing conservative to aggressive disease-specific adjustments. Correlations between original and perturbed scores remained above $r = 0.978$ for all realistic multipliers... Product rankings remained highly stable (Pearson $r > 0.90$) across the entire tested range.”

Conclusion: CVD profile is robust to uncertainty in penalty severity.

3.7.4 Sensitivity Summary

- New Figure 5 added: Comprehensive sensitivity heatmap
- What it shows: Pearson correlations for all 35 perturbation-profile combinations
- Results: All $r \geq 0.983$, most exceed $r = 0.99$




Figure 5:

Figure 5 caption:

“Sensitivity analysis heatmap showing Pearson correlations between original FoodByte scores and perturbed scores under various algorithmic modifications. Columns represent perturbation types: Original (baseline), Additive -50%/+50% (scaled additive penalties), Threshold -20%/+20% (adjusted disease-specific thresholds), and CVD 0.5x/2.0x (cardiovascular penalty multipliers, HTN+CVD profile only). All correlations exceed $r = 0.983$ (green cells), indicating high stability of product rankings across perturbations.”

“Figure 5 presents a comprehensive heatmap of correlations between original and perturbed FoodByte scores across all tested perturbations and health profiles. All 35 perturbation-profile combinations yielded correlations $r \geq 0.983$, with the vast majority exceeding $r = 0.99$.”

3.7.5 Aggregation Rule Considerations

- What I address: Why I did NOT test alternative aggregation methods
- Justification:

“An additional potential source of model sensitivity lies in the specific rule used to aggregate multiple additive penalties (e.g., simple summation, capped summation, max-risk-only, or class-weighted aggregation). In principle, these alternatives could be evaluated by re-executing the full FoodByte algorithm under each aggregation scheme.

In the present study, I did not enumerate alternative aggregation rules for two methodological reasons. First, the additive-weight, threshold, and profile-multiplier perturbations already stress-test the effective behavior of the aggregation mechanism across wide and extreme ranges. These perturbations directly alter the magnitude and relative contribution of additive penalties, capturing the primary dimensions along which alternative aggregation rules would differ. Second, introducing simulated aggregation schemes without full algorithmic re-execution would require heuristic assumptions that risk introducing artifacts rather than strengthening inference. The uniformly high rank correlations observed across all perturbations (minimum $r = 0.983$), together with minimal category instability under threshold shifts, indicate that FoodByte's conclusions are driven by the presence and co-occurrence of additives and disease-specific sensitivity rather than by the precise mathematical form of aggregation.”

Why this is defensible:

1. The three sensitivity tests already indirectly test aggregation behavior
2. Simulating aggregation without full re-execution would introduce methodological artifacts
3. The high stability ($r > 0.983$) across all existing perturbations demonstrates robustness

Updates to Other Sections:

1. Section 2.3 (Methods) - Added Sensitivity Methodology:

- a. Last paragraph

“To assess robustness to algorithmic specification, sensitivity analyses were conducted by perturbing additive scoring weights ($\pm 50\%$), threshold values ($\pm 20\%$), and profile-specific multipliers for cardiovascular disease penalties (0.5x to 2.0x). For each perturbation, stability of score rankings and category assignments was evaluated using Pearson correlations between perturbed and original scores. The percentage of products maintaining their letter-grade category under threshold perturbations was calculated to assess categorical stability.”

2. Section 2.4 (Implementation) - Updated Workflow:

- a. Items 3.e and 4.e

“e. Figure 5: Sensitivity analysis heatmap showing correlations under algorithmic perturbations.”

“e. Sensitivity analyses: Tested robustness under perturbations to additive weights, disease-specific thresholds, and profile-specific penalty multipliers.”

3. Section 4.1 (Discussion) - New Paragraph on Sensitivity:

a. Paragraph 3

“Sensitivity analyses confirmed the robustness of these findings across algorithmic variations. Product rankings remained highly stable ($r > 0.983$) under $\pm 50\%$ changes to additive weights, with 90.4% of products maintaining their letter-grade category under $\pm 20\%$ threshold perturbations. Cardiovascular disease penalties could vary from 0.5x to 2.0x of baseline values while maintaining rank correlations above $r = 0.978$. This robustness indicates that FoodByte's divergence from Nutri-Score reflects systematic differences in additive presence and disease-specific risk profiles rather than arbitrary parameter calibration.”

Summary of Sensitivity Tests:

1. Additive weight perturbation ($\pm 50\%$) - Correlations $r > 0.983$
2. Threshold perturbation ($\pm 20\%$) - 90.4% maintain category
3. CVD profile multipliers (0.5-2.0x) - Correlations $r > 0.978$
4. Aggregation methods - Not tested; methodological justification provided in Section 3.7.5

I conducted 3 comprehensive sensitivity analyses and provided scientific justification for why a 4th test (aggregation methods) was not necessary given the existing evidence of robustness.

COMMENT #4: GitHub Repository Comment:

“I appreciate that you have included the raw data and code in your submission. However, please also make a public GitHub repository and deposit your data there with a link to it included in the revised manuscript.”

Response:

I have fully addressed this concern by creating a public GitHub repository and including the link in the revised manuscript.

1. Section 2.4 (Implementation) - Added GitHub Link:

- a. Added direct link to public GitHub repository

“Complete data and analysis code are publicly available

<https://github.com/harshithapoludas/FoodByte.git> to enable reproduction of all reported results.”

Reproducibility: Any researcher can clone the repository and reproduce the entire analysis using the provided code and data.

SUMMARY OF CHANGES

1. Statistical Methods (Section 2.3):

- a. Completely rewritten to describe mixed-effects modeling
- b. Added effect size calculation methodology
- c. Added sensitivity analysis methodology
- d. Word count: Increased from ~200 to ~500 words

2. Results - Statistical Analyses (Section 3.1):

- a. Replaced paired t-test with mixed-effects model results
- b. Added ICC = 0.518 with detailed interpretation
- c. Added discussion of generalization to new products
- d. Added: ~300 words

3. Results - Effect Sizes (Section 3.2):

- a. Completely restructured Table 1 (3 columns → 7 columns)
- b. Added Figure 2 (effect sizes with CIs)
- c. Added interpretation of effect sizes
- d. Added: ~200 words, 1 new figure, enhanced table
- 4. Results - Sensitivity Analyses (Section 3.7):**
 - a. Entirely new section: ~1,500 words
 - b. Added Table 3 (CVD multipliers)
 - c. Added Figure 5 (sensitivity heatmap)
 - d. Four subsections (3.7.1-3.7.4 plus aggregation justification 3.7.5)
- 5. Discussion (Section 4.1):**
 - a. Added paragraph on effect sizes
 - b. Added paragraph on sensitivity analyses
 - c. Updated paragraph 1 to reference mixed-effects results
 - d. Added: ~250 words
- 6. Implementation (Section 2.4):**
 - a. Added statsmodels library
 - b. Updated workflow to include new figures
 - c. Added GitHub repository link
 - d. Added: ~100 words
- 7. Conclusion (Section 5):**
 - a. Updated to reference mixed-effects modeling and effect sizes
 - b. Modified: ~50 words

New Content Summary:

- **New sections:** 1 (Section 3.7, ~1,500 words)
- **New figures:** 2 (Figure 2: Effect sizes, Figure 5: Sensitivity heatmap)
- **New tables:** 1 (Table 3: CVD multipliers)
- **Enhanced tables:** 1 (Table 1: now 7 columns instead of 3)
- **Total new/revised content:** ~2,900 words
- **Figures renumbered:** Old Figures 2→3, Old Figure 3→4

ADDITIONAL IMPROVEMENTS

Beyond addressing the reviewer's specific comments, I made the following improvements:

- 1. Added Interpretive Text Below Figures:**
 - a. Figure 2 (Effect sizes): Added paragraph explaining why disease-specific profiles show larger effects
 - b. Figure 5 (Sensitivity): Added paragraph explaining what high correlations mean for validity
- 2. Methodological Transparency:**
 - a. Section 3.7.5 honestly acknowledges limitation (aggregation methods not tested)
 - b. Provides scientific justification for this decision
 - c. Demonstrates mature scientific reasoning

CONCLUSION

The revised manuscript now includes:

More rigorous statistical methods accounting for repeated measures

- Comprehensive reporting of effect sizes and confidence intervals
- Extensive sensitivity analyses demonstrating robustness

- Full reproducibility via public GitHub repository

-

Thank you for addressing my comments. The manuscript is much improved. However, some inconsistencies and unaddressed issues remain. Please address these.

1. You have modeled only product as a random effect "...using a linear mixed-effects model with the score difference as the outcome and a random intercept for product." But this is a crossed design; each food has multiple profiles and each profile has multiple foods. You have included random intercept for product and fixed effect for scoring system. But you have not included random intercept for profile. Profiles are not independent replications — they are systematic rule structures. If profiles influence score differences systematically, failing to include them as a random effect underestimates uncertainty, inflates t-statistic and overstates generalizability. Right now, your model accounts for clustering by product, but not clustering by profile.

2. Please tone down overreach and fix language. For example [1] "Results decisively reject the null hypothesis of equivalence" requires TOST or equivalent. [2] You interpret: 51.8% product variance, 48.2% profile-specific adjustments. That is only true if profile is explicitly modeled. If profile is not in the random structure, then 48.2% is just residual variance — not necessarily attributable to profile adjustments. [3] replace "personalization produces substantial practical differences" with "The magnitude of divergence suggests practical relevance." Tone down promotional claims. [4] You state "This structure enables inference about how new, previously unseen products would behave." That is only partially true. You can generalize to: similar product distribution But: You did not sample products randomly. Sample size is 50 foods. Food selection method matters. Unless you used random sampling from a population frame, generalization claims should be toned down. [5] You compare FoodByte vs Nutri-Score. But: Nutri-Score was not designed to be personalized. You are not comparing against clinical outcomes. There is no validation against health endpoints. So the study shows divergence — not superiority. Remove language implying failure of Nutri-Score. [6] You mention: "relative ordinal rankings based on scientific evidence strength". These weights are expert-defined, not empirically estimated. That's fine — but should be labeled as rule-based heuristic weights. [7] "Correlation analyses are presented descriptively and do not assume independence." That is incorrect. Pearson correlation assumes independent observations. If repeated measures exist, correlation CIs are inflated. Please state explicitly that they are descriptive only without inferential testing.

3. Sample Size and Power: 50 products, 5 profiles, 250 observations. ICC 0.518 suggests high clustering. Effective sample size may be closer to ~50–100, not 250. Please report: Model convergence diagnostics, Residual diagnostics or Normality checks

-

-

FoodByte: A Personalized Food Scoring and Additive Risk Assessment Algorithm for Smarter Eating

-

Reviewer Response

-

Author: Harshitha Poludas Date: 07/07/2026

-

-

-

-

COMMENT #1: Crossed Mixed-Effects Model / Random Intercept for Profile

Comment:

You have modeled only product as a random effect "...using a linear mixed-effects model with the score difference as the outcome and a random intercept for product." But this is a crossed design; each food has multiple profiles and each profile has multiple foods. You have included random intercept for product and fixed effect for scoring system. But you have not included random intercept for profile. Profiles are not independent replications — they are

systematic rule structures. If profiles influence score differences systematically, failing to include them as a random effect underestimates uncertainty, inflates t-statistic and overstates generalizability. Right now, your model accounts for clustering by product, but not clustering by profile.

-
-

• **Response:**

-

• To address this, I replaced the product-only mixed-effects model with a **cross-classified linear mixed-effects model** including random intercepts for both product and health profile.

• **Section 2.3: Statistical Analysis — Revised Model Specification:**

-

• I revised Section 2.3 to explicitly describe the crossed data structure and the need to model clustering by both product and profile. The revised Methods now state that all 250 product–profile observations should not be treated as independent and that the inferential model includes random intercepts for both clustering dimensions.

-

• The revised model is: $\text{Score difference}_{ij} = \beta_0 + u_{\text{product},i} + v_{\text{profile},j} + \varepsilon_{ij}$ where **u_product,i** represents the random intercept for product *i*, **v_profile,j** represents the random intercept for profile *j*, and ε_{ij} represents residual variation.

-

-

• The methods now state “Because each food product was evaluated under each of five health profiles, the resulting observations had a cross-classified repeated-measures structure: observations could be correlated because they shared the same product and because they shared the same profile rule structure.

-

• Treating all 250 product–profile combinations as independent would underestimate uncertainty. Therefore, inferential analysis was conducted using a cross-classified linear mixed-effects model with score difference, defined as FoodByte score minus Nutri-Score, as the outcome. The model included a fixed intercept estimating the overall mean score difference and random intercepts for both product and health profile.”

• **Section 2.4: Implementation — Updated Software Description:**

-

• I updated the implementation workflow to reflect the revised crossed model. The statistical analysis workflow now states:

• - “Inferential statistics: Cross-classified linear mixed-effects model with score difference as the outcome and random intercepts for both product and health profile.”

-

• The implementation section also identifies statsmodels as the library used for linear mixed-effects modeling.

-

-

• **Section 3.1: Results — New Cross-Classified Mixed Model Results:**

-

• I replaced the old product-only model results with the new cross-classified model results. The revised results are:

- Mean score difference: **-10.49 points**
- 95% CI: **[-14.50, -6.48]**
- Model-based p-value: **2.88×10^{-7}**
- Product variance: **75.73**, representing **52.6%** of modeled variance
- Profile variance: **12.21**, representing **8.5%** of modeled variance
- Residual variance: **56.02**, representing **38.9%** of modeled variance

- Model convergence: **Yes**
-
- The manuscript now states that product identity was the largest modeled source of clustering, while health-profile structure also contributed measurable variance and was retained in the random-effects structure. I also added the explicit caution that variance components are model-based clustering estimates rather than causal attribution.
- This revision removes the earlier unsupported interpretation that the non-product variance was “profile-specific adjustment” when profile was not explicitly modeled. Profile is now included in the random-effects structure, and the results are reported as product-level, profile-level, and residual variance components.
-
- **COMMENT #2: Effect Sizes and Confidence Intervals Comment:**
- 1. “Please tone down overreach and fix language. For example [1] “Results decisively reject the null hypothesis of equivalence” requires TOST or equivalent. [2] You interpret: 51.8% product variance, 48.2% profile-specific adjustments. That is only true if profile is explicitly modeled. If profile is not in the random structure, then 48.2% is just residual variance — not necessarily attributable to profile adjustments.
- [3] replace “personalization produces substantial practical differences” with “The magnitude of divergence suggests practical relevance.” Tone down promotional claims. [4] You state “This structure enables inference about how new, previously unseen products would behave.” That is only partially true. You can generalize to: similar product distribution But: You did not sample products randomly. Sample size is 50 foods. Food selection method matters. Unless you used random sampling from a population frame, generalization claims should be toned down. [5] You compare FoodByte vs Nutri-Score. But: Nutri-Score was not designed to be personalized. You are not comparing against clinical outcomes. There is no validation against health endpoints. So the study shows divergence — not superiority. Remove language implying failure of Nutri-Score. [6] You mention: “relative ordinal rankings based on scientific evidence strength”. These weights are expert-defined, not empirically estimated. That’s fine — but should be labeled as rule-based heuristic weights. [7] “Correlation analyses are presented descriptively and do not assume independence.” That is incorrect. Pearson correlation assumes independent observations. If repeated measures exist, correlation CIs are inflated. Please state explicitly that they are descriptive only without inferential testing.”
-
- **Response:**
-
- The revised manuscript now frames the study as showing **algorithmic score divergence** within a purposively selected product sample, rather than as evidence that FoodByte is clinically superior to Nutri-Score or that Nutri-Score fails.
- **Removal of Equivalence Language**
-
- I removed the phrase “Results decisively reject the null hypothesis of equivalence.”
- I also removed or revised language implying formal equivalence or non-equivalence testing. The manuscript no longer uses “decisively reject” or claims equivalence testing, because no TOST or equivalent procedure was performed.
- The revised wording in Section 3.1 states:
- “These results indicate a consistent divergence between FoodByte and Nutri-Score after accounting for clustering by both food product and health-profile rule structure.”
-
- **Variance Component Interpretation Corrected**
-
- The earlier product-only interpretation of variance was removed. The manuscript now reports variance components from the corrected cross-classified model: “Variance-component

estimates indicated that product-level clustering accounted for 52.6% of modeled variance, profile-level clustering accounted for 8.5%, and residual variation accounted for 38.9%.”

- The manuscript also clarifies: “These variance components should be interpreted as model-based clustering estimates rather than causal attribution.”

-

- **Practical Relevance Language Revised**

-

- I replaced stronger language such as “personalization produces substantial practical differences” with more cautious wording like “The magnitude of divergence suggests practical relevance within the evaluated sample.”

- The revised Section 3.2 states:

- “Effect size analysis showed that the magnitude of divergence suggests practical relevance within the evaluated sample. Four of five profiles had large descriptive standardized differences (Cohen’s $|d| > 0.8$), while the Control profile had a medium descriptive effect size.”

- The Discussion now also states that the effect sizes should not be interpreted as evidence of clinical outcome improvement.

- **Generalizability Claims Toned Down**

-

- I removed the earlier statement that the model enables inference about “new, previously unseen products.” Because the 50 foods were purposively selected rather than randomly sampled, the manuscript now limits interpretation to the evaluated sample and similar product/profile structures.

- The revised Methods state:

- “Because the products were purposively selected and the five profiles were prespecified rule-based scenarios rather than randomly sampled from defined population frames, generalization is limited to similar packaged-food products and comparable health-profile structures.”

- The Abstract and Discussion now also explicitly state that the findings should not be interpreted as population-wide generalizability or clinical superiority.

- **FoodByte vs Nutri-Score Framed as Divergence, Not Superiority**

-

- I revised language throughout the manuscript to clarify that Nutri-Score was not designed to perform personalized additive-risk or disease-specific scoring. I removed language implying that Nutri-Score “fails,” “misrates,” or is clinically inferior.

- For example, Section 3.4 now states “Nutri-Score does not include additive-risk scoring within its nutrient-profiling framework. FoodByte adds an additive-focused layer to the evaluation, allowing additive presence to influence scores in a way that is not part of Nutri-Score’s intended design.”

- Section 4.3.1 now frames the comparison as scope differences relative to Nutri-Score, rather than as failure of Nutri-Score. The revised text states “This divergence should not be interpreted as evidence that Nutri-Score misrates products, because Nutri-Score was not designed to perform individualized clinical or additive-risk assessment.”

- The Conclusion now states “These findings should be interpreted as evidence of score divergence and practical relevance within the evaluated product sample, not as proof of clinical superiority or validation against health outcomes.”

- **Additive Weights Relabeled as Rule-Based Heuristic Weights**

-

- I revised the additive scoring description to clarify that additive risk scores are expert-defined, rule-based heuristic weights, not empirically estimated clinical risk coefficients.

- The revised Section 2.1.4 states “Food additives were assigned expert-defined, rule-based heuristic risk scores on an ordinal 1–10 scale. These scores were informed by published toxicological, epidemiological, regulatory, and clinical evidence, but they were not empirically

estimated from clinical outcome data, exposure-response modeling, or population-level validation.”

- The revised Conclusion states “Its additive-risk scores are literature-informed and rule-based rather than empirically estimated clinical risk coefficients.”

- **Pearson Correlation Language Corrected**

-

- I corrected the Methods to explicitly state that Pearson correlation assumes independent observations and that the reported correlations are descriptive only.

- The revised Methods state “Pearson correlation coefficients and R^2 values were calculated to describe the pooled relationship between FoodByte and Nutri-Score values. Because Pearson correlation assumes independent observations and each product appears repeatedly across profiles, these correlations were treated as descriptive summaries only. No inferential p-values or confidence intervals for correlation coefficients were interpreted.”

- The Correlation Analysis section now also states that no inferential or independence-based interpretation is attached to the correlation

- **Additional Tone and Wording Revisions Throughout**

-

- Beyond the reviewer’s specific examples, I made additional revisions throughout the Abstract, Introduction, Results, Discussion, and Conclusion to remove promotional or causal language. I replaced words and phrases such as “proves,” “confirms,” “failure,” “clinically superior,” “misleading endorsements,” “life-threatening implications,” and “improves health outcomes” with more cautious phrasing such as:

- “suggests”
- “is consistent with”
- “in this sample”
- “descriptive”
- “algorithmic divergence”
- “rule-based”
- “literature-informed”
- “requires empirical validation”
- “not evidence of clinical superiority”

-

-

- I also revised the Doritos case study and Discussion to make clear that the case study illustrates differences in scoring rules, not clinical validation or evidence that one system predicts health outcomes better than the other.

-

- **COMMENT #3: Sample Size, Power, Clustering, and Diagnostics Comment:**

- “Sample Size and Power: 50 products, 5 profiles, 250 observations. ICC 0.518 suggests high

- clustering.. Effective sample size may be closer to ~50–100, not 250. Please report: Model convergence diagnostics, Residual diagnostics or Normality checks”

-

- **Response:**

-

- I revised the Methods, Results, and Limitations to explicitly describe the clustered repeated-measures design and to report approximate effective sample size calculations. I also added model convergence diagnostics, residual diagnostics, a Shapiro-Wilk normality check, and a new diagnostic figure.

-

- **Section 2.3: Effective Sample Size Added to Methods**

-

- I added a paragraph explaining that the 250 rows are repeated product–profile evaluations, not independent product samples. I added approximate design-effect calculations using the product-level and profile-level variance components from the mixed-effects model.
- The revised Methods state “Because the 250 observations were repeated product–profile evaluations rather than independent product samples, the nominal observation count was not treated as the effective independent sample size. To contextualize clustering, an approximate design-effect calculation was performed using the product-level variance proportion from the mixed-effects model. With five profile evaluations per product and product-level clustering accounting for 52.6% of modeled variance, the approximate design effect was $1 + (5 - 1)(0.526)$
- $= 3.10$, corresponding to an approximate effective sample size of $250 / 3.10 \approx 81$ independent observations. Including both product- and profile-level variance components yields a more conservative approximation of $250 / [1 + 4(0.526 + 0.085)] \approx 73$.”
- The manuscript also clarifies that these are descriptive approximations rather than exact power estimates for a crossed repeated-measures design.

• Effective Sample Size Reported in Results

- I added the design-effect and effective sample size results to Section 3.1:

- Product-only design effect: **3.104**
- Approximate effective sample size, product-only: **80.5**
- Product + profile design effect: **3.443**
- Approximate effective sample size, product + profile: **72.6**

- The revised Results state “Although the dataset contained 250 product–profile observations, these observations were clustered by both product and profile and should not be interpreted as 250 independent samples. Using the product-level variance proportion as an approximate ICC-like clustering estimate, the design effect was 3.104, giving an approximate effective sample size of 80.5. A more conservative approximation incorporating both product- and profile-level variance components yielded a design effect of 3.443 and an approximate effective sample size of 72.6.”

• Model Convergence and Residual Diagnostics Added

- I added model diagnostics to Section 3.1. The revised manuscript reports:

- Model convergence: **successful**
- Residual mean: approximately **0.0000**
- Residual SD: **6.7524**
- Shapiro-Wilk W: **0.991**
- Shapiro-Wilk p-value: **0.136**

- The manuscript states “Model convergence and residual diagnostics were assessed for the cross-classified mixed-effects model. The model converged successfully. Residual diagnostics were generally acceptable: the residual Q-Q plot showed approximate normality with mild tail deviations, and the residuals-versus-fitted plot showed residuals centered approximately around zero, although several outlying observations and some visible structure were present. The Shapiro-Wilk test did not indicate strong evidence against residual normality ($W = 0.991$, $p = 0.136$).”

• New Figure 1: Model Diagnostics

- I added a new main-text diagnostic figure as Figure 1. The figure includes:
- Panel A: residual Q-Q plot

- Panel B: residuals versus fitted values plot
-
- The caption states that the model included random intercepts for both product and health profile and reports the Shapiro-Wilk diagnostic result.
- This figure was inserted before the previous Figure 1, so all subsequent figures were renumbered.

- **Section 4.5: Limitations Expanded**
-

- I revised the Sample Size limitation to explicitly state that the effective information content is lower than the nominal 250 observations.
- The revised limitation states “The dataset included 50 unique food products evaluated across 5 prespecified health profiles, yielding 250 product–profile observations. However, these observations are not equivalent to 250 independent product samples because each product was repeatedly evaluated across profiles and each profile was repeatedly applied across products. The mixed-effects model indicated substantial clustering, with product-level variance accounting for 52.6% of modeled variance and profile-level variance accounting for 8.5%. Approximate design-effect calculations suggested an effective sample size of roughly 73–81 independent observations, despite the nominal 250 product–profile rows. Therefore, statistical precision and generalizability should be interpreted cautiously.”
-
-

- **ADDITIONAL REVISIONS MADE FOR INTERNAL CONSISTENCY**
-

- In addition to the three main reviewer concerns, I made several smaller revisions throughout the manuscript to improve clarity, consistency, and reviewer-safe wording.

- **Category Counts Corrected**

- I corrected the product-category distribution to match the current dataset and Table 2. The manuscript now uses the six-category distribution reflected in the analysis output:

- Snacks: $n = 23$
- Beverages: $n = 5$
- Breakfast: $n = 5$
- Condiments: $n = 5$
- Dairy: $n = 7$
- Grains: $n = 5$

- I removed the inconsistent earlier category distribution that listed frozen meals and different category counts.

- **Figure Numbering Updated**

- Because the new diagnostic figure was added as Figure 1, all subsequent figures were renumbered:

- Figure 1: Mixed-model diagnostic plots
- Figure 2: Distribution of score differences by health profile
- Figure 3: Effect sizes and percentage differences
- Figure 4: Category-specific score differences
- Figure 5: FoodByte vs Nutri-Score scatter plot
- Figure 6: Sensitivity analysis heatmap
- The Implementation workflow and in-text references were updated accordingly.

- **Equation and Notation Fixes**

- I corrected the mixed-effects equation notation to use random-effect terms $u_{\text{product},i}$ and $v_{\text{profile},j}$ rather than incorrect symbols. I also fixed the final score equation by removing an extra parenthesis:

- $\text{final score} = \max(1, \min(100, \text{personalized score}))$
- **Sensitivity Analysis Wording Revised**

- The sensitivity analysis language was revised to avoid treating Pearson correlations as inferential tests. The manuscript now describes them as descriptive similarity measures and emphasizes that sensitivity analyses support internal stability of the scoring outputs, not clinical validation.

- For example, Section 3.7.1 now states “These descriptive correlations suggest that FoodByte score patterns were internally stable under the tested additive-weight perturbations. This analysis indicates that the observed score patterns were not highly sensitive to the exact additive penalty scaling used in this study. However, these sensitivity results do not validate the additive weights against clinical outcomes or establish that the weights are empirically calibrated risk estimates.”

- **Case Study Interpretation Toned Down**

- I revised the Nacho Cheese Doritos case study to remove claims of clinical superiority, “life-threatening implications,” or real-world health outcome effects. The revised case study states that Doritos illustrates how FoodByte’s rule-based framework can produce different scores from Nutri-Score, but not that either system has been clinically validated as predicting health outcomes.

- **Conclusion Revised**

- The Conclusion was substantially revised to reinforce the corrected framing. It now states that the study shows score divergence and internal robustness within the evaluated sample, not clinical superiority or health outcome validation.

- The revised conclusion states:

- “Overall, FoodByte represents a rule-based prototype for additive-aware and profile-specific food scoring. Future research should prioritize expanded product sampling, empirical validation against consumer behavior and clinical outcomes, improved ingredient detection, additive concentration data, alternative aggregation rules, and testing with diverse users.”

-

- **SUMMARY OF CHANGES**

-

- Replaced product-only mixed-effects model with a cross-classified mixed-effects model

- Added random intercepts for both product and health profile

- Corrected mixed-model equation and notation

- Added variance components for product, profile, and residual variation

- Removed unsupported ICC/profile interpretations from the old model

- Updated Section 3.1 with corrected model results:

- Mean difference: **-10.49 points**

- 95% CI: **[-14.50, -6.48]**

- Model-based p-value: **2.88×10^{-7}**

- Product variance: **52.6%**

- Profile variance: **8.5%**

- Residual variance: **38.9%**

- Added effective sample size estimates:

- Approximate effective n: **80.5** using product-level clustering

- Approximate effective n: **72.6** using product + profile clustering

- Added model convergence and residual diagnostics

- Added new Figure 1 showing residual Q-Q and residuals-versus-fitted plots

- Removed overclaiming and promotional phrasing

- Removed equivalence-test language

- Removed claims of FoodByte superiority or Nutri-Score failure

- Clarified that Nutri-Score was not designed for personalization or additive-risk scoring

- Clarified that the study shows score divergence, not health outcome validation

- Relabeled additive weights as literature-informed, rule-based heuristic weights

- Clarified Pearson correlations are descriptive only

- Retained sensitivity analyses but revised interpretation

- Pearson correlations are now described as descriptive stability measures
- Added caution that sensitivity analyses do not validate clinical outcomes or empirically calibrate additive weights
- Expanded sample-size limitation
- Added effective sample size discussion
- Added limitations on purposive sampling, profile simulation, additive concentration data, heuristic scoring, behavioral validation, and clinical outcome validation

-
-

• **CONCLUSION**

-

- The revised manuscript now directly addresses all reviewer concerns by:
- Modeling the crossed repeated-measures structure using random intercepts for both product and health profile.
- Reporting variance components, approximate effective sample size, convergence diagnostics, residual diagnostics, and normality checks.
- Removing overclaiming and reframing the findings as descriptive score divergence rather than clinical superiority.
- Clarifying that additive weights are rule-based heuristic weights, not empirically estimated clinical risk coefficients.
- Treating Pearson correlations and sensitivity correlations as descriptive summaries only.
- Limiting generalizability because the product sample was purposively selected and no clinical outcomes were measured.

-

- Thank you for addressing my comments. Accepted. I also recommend a ‘future research’ or ‘perspectives’ paragraph be included:

“Future research should move beyond demonstrating algorithmic divergence and evaluate whether personalized food-scoring systems provide measurable advantages over existing nutrient-profiling frameworks such as Nutri-Score. This could be accomplished through prospective studies comparing FoodByte recommendations with clinician assessments, registered dietitian evaluations, established dietary guidelines, or patient-centered outcomes. Disease-specific validation in populations with diabetes, hypertension, chronic kidney disease, irritable bowel syndrome, or food allergies would determine whether profile-specific scoring improves dietary decision-making or better predicts clinically relevant biomarkers. In addition, future work should investigate whether incorporating empirical exposure data, additive concentrations, longitudinal dietary records, and machine learning–derived risk estimates improves upon the current rule-based heuristic framework. Ultimately, the scientific value of personalized food scoring will depend not only on its ability to generate scores that differ from generic systems, but on demonstrating that those differences produce more accurate, clinically meaningful, and actionable nutritional recommendations.”