

Accounting for the nutritional risks of calorie restriction yields protein-adequate personalised diet plans

Tazayun Shakeel

tazayunbhat2020@cukashmir.edu.in

Central University of Kashmir

Afaq Khan

Central University of Kashmir

Shagufta khan

childrens hospital bemina srinagar

Article

Keywords: Calorie restriction, Many-objective optimisation, NSGA-III, Diet planning, Nutrient adequacy, Analytic hierarchy process

Posted Date: October 22nd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11160634/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Accounting for the nutritional risks of calorie restriction yields protein-adequate personalised diet plans

Tazayun Shakeel^{1*}, Afaq Alam Khan¹ & Shagufta Khan²

¹Department of Information Technology, School of Engineering & Technology, Central University of Kashmir, Ganderbal 191131, Jammu and Kashmir, India. ²Children's Hospital, Bemina, Srinagar 190018, Jammu and Kashmir, India. *email: tazayunbhat2020@cukashmir.edu.in

Abstract

Eating less than the body needs underlies most weight-loss plans, but a sustained calorie deficit also costs lean muscle, thins hair, lowers mood and leaves people hungry. Diet-planning software rarely models these effects and usually optimises the calorie count alone. We built a diet planner that treats them as goals in their own right. From a person's sex, age, height, weight and activity level the system sets personal energy and nutrient targets, then uses an evolutionary algorithm to search 2,307 foods from the USDA database for a day of eating that supports mood, hair, muscle and fullness while limiting unhealthy nutrients and holding the energy target. Each goal is scored only from nutrients whose supporting evidence survived an audit, which removed fourteen candidate terms. For an active 27-year-old man on an 801 kcal daily deficit, the recommended plan met 101% of an activity-scaled protein target, where Mediterranean, DASH and MIND plans at the same energy reached 71 to 84%, and its micronutrient adequacy (95.4%) was higher than that of all 2,000 random feasible diets tested. All eight test profiles met their protein targets. A nutrition counsellor judged the plans nutritionally sound but often impractical to serve as written.

Keywords Calorie restriction · Many-objective optimisation · NSGA-III · Diet planning · Nutrient adequacy · Analytic hierarchy process

Introduction

Eat fewer calories than the body uses and weight comes off. That idea, the calorie deficit, sits under nearly every weight-loss plan in circulation, and over time it does make the body draw on its fat stores. The reality is messier. Keep a deficit going long enough and some lean muscle goes with the fat [1], hair can thin [2], mood tends to drop [3–5] and hunger becomes persistent [6–8]. Hunger in particular is one of the main reasons people abandon weight-loss

diets [9]. None of this is new. Yet the planners and apps people actually use still aim at the calorie number, sometimes with a rough split between carbohydrate, fat and protein, and stop there.

Excess body weight is among the largest health problems of the present day and carries a higher risk of hypertension and cardiovascular disease [10], so demand for workable deficits is not going away. Cutting intake without looking at what is left on the plate has a price. Smaller portions carry less of every vitamin and mineral, and the shortfall shows up in hair, mood and muscle. Lean mass lost during the deficit takes some resting energy expenditure with it. What a planner rarely asks is whether the meals it proposes protect muscle and hair, support mood and keep the person reasonably full while the deficit holds.

As an optimisation problem, diet planning is old. Stigler posed the least-cost diet meeting a set of nutrient requirements in 1945 [11], and later work added variety, cost, preparation time and preference, moving from linear programming to evolutionary search [12–14].

Turkmenoğlu et al. [14] treated menu planning as a many-objective knapsack problem over 356 cooked dishes, minimising cost and preparation time and maximising preference under two-sided bounds on 23 nutrients, and solved it with NSGA-III. That is the closest precedent to the present work in size and algorithm, though its objectives are economic. MenuGene [13] personalises nutrient targets from a medical record and enforces them at the meal, day and week levels at once. Ahourag et al. [15] minimised glycaemic load alongside cost for patients with diabetes, one of the few cases in which a physiological quantity is an objective rather than a constraint. Sahin and Aytekin-Sahin [16] scored colour, consistency and variety next to nutrient adequacy and found AGE-MOEA and SMS-EMOA ahead of both NSGA variants on their problem. Babalola et al. [17] added seasonal availability, and ACMO-Diet [18] handled an ultra-sparse binary representation with custom variation operators inside U-NSGA-III. Supplementary Table S1 sets these studies side by side.

Read together, the studies show three gaps. The objectives are almost always economic, logistical or hedonic, and nutrition enters only as constraints, so a plan that barely clears its nutrient floors scores the same as one well above them, even though their likely physiological effects differ. No study models what a calorie deficit does to the person eating the plan. Muscle loss, hair thinning, low mood and hunger are absent from every objective set, and these are exactly the ways weight-loss diets are known to fail. Where weight loss is the aim, it is pursued through the calorie figure alone, with no provision for people whose appetite is suppressed by GLP-1 receptor agonists, a group that has grown quickly in clinical practice.

This paper addresses those gaps. We formulate diet planning under a calorie deficit as a six-objective problem. Support for mood, hair, muscle and satiety is maximised over the whole day's intake, and an unhealthy-nutrient composite and the deviation from the energy target are minimised. Almost every scoring constant is derived from the user's own targets, so two people with different profiles face different objective surfaces over the same food table. Within each score, nutrient terms are weighted by an analytic hierarchy process built on an explicit evidence audit, and any term whose own citation reports a null or unfavourable result is removed rather than down-weighted. The hair objective and the way the energy target is set both change when the user takes a GLP-1 receptor agonist. We test the approach on eight profiles covering a 2.7-fold range of energy targets, against 2,000 random feasible diets, against Mediterranean, DASH and MIND plans at matched energy, and through a practitioner's review of the generated plans. That review turned out to be the most informative part of the evaluation, because it shows where nutritional optimisation stops being enough.

Results

Overview of the framework

The system takes a physical profile and returns a five-meal plan for one day, assembled from a real food-composition table (Fig. 1). Steps 1 to 3 prepare the inputs: personal energy and nutrient targets, scoring constants calibrated to those targets, and a filtered pool of candidate foods. Steps 4 to 7 formulate the optimisation problem, step 8 solves it with NSGA-III, and steps 9 and 10 select one plan and evaluate it. Step 11 rebuilds the chosen plan for delivery. It runs after every reported number has been computed and cannot change any result below.




Figure 1. Pipeline of the proposed system, from the user profile to the issued plan. Steps 1 to 10 produce and evaluate the plan. Step 11 rebuilds it for delivery (food substitution, meal shifting and recipe text from a language model) and cannot alter any reported measurement. HV, hypervolume; IGD+, inverted generational distance plus; RDA, recommended dietary allowance; UL, tolerable upper intake level.

The six objectives are four support scores (mood, hair, muscle and satiety), each a weighted sum of bounded saturation terms in nutrients such as protein, iron, zinc and magnesium (Table 1), together with a minimised unhealthy-nutrient composite and a minimised energy-deviation term. Hard constraints cover micronutrient upper limits, energy and sodium bands and per-subgroup portion caps, and soft penalties cover nutrient floors. Methods gives the full formulation.

Table 1. AHP weights of the four maximised support scores. Each column sums to 1. A dash means the term is not part of that score.

Term	Mood	Hair	Muscle	Satiety
Magnesium	0.219	—	0.104	—
Zinc	0.199	0.212	—	—
Vitamin D	0.150	0.136	—	—
Omega-3 (ALA, EPA+DHA)	0.147	—	0.104	—
Iron	0.143	0.352	0.156	—
Selenium	0.078	—	—	—
Vitamin B-6	0.064	—	—	—
Protein	—	0.300	0.265	0.307
Carbohydrate	—	—	0.140	—
Potassium	—	—	0.116	—

Term	Mood	Hair	Muscle	Satiety
Leucine	–	–	0.116	–
Energy density (decreasing)	–	–	–	0.492
Fibre	–	–	–	0.201
Number of terms	7	4	7	3

The results below come from a primary profile, P1: an active 27-year-old man (74 kg, 171 cm) with a maintenance energy of 2,854 kcal/day and a deficit of 801 kcal/day. That gives an energy target of 2,053 kcal and keeps all six objectives active. The search used a population of 210 over 12,000 generations with 21 independent seeds and 2,307 candidate foods (Supplementary Table S3).

Convergence and the Pareto front

Figure 2a plots the hypervolume of each seed against generation. The mean rises steeply until about generation 500 and slowly after that. Early on, the search is carrying an unstructured population into the feasible region; later, gains come from refining the trade-off surface rather than finding it. The seeds end far apart, with final values in Fig. 2a between 0.136 and 1.283. Measured against the pooled front, final hypervolume is 0.139 ± 0.223 (mean $\pm$ s.d.; median 0.025, interquartile range 0.001 to 0.171), a coefficient of variation of 160%, and IGD+ is 0.751 ± 0.488 . The distribution is strongly skewed: most seeds contribute little and a few contribute most of the front, and no single seed approximates the pooled result closely (median IGD+ 0.696). Pooling the fronts of all 21 seeds is therefore how the front is obtained, not a refinement, and a single-seed run would misrepresent what the method can produce. Each trace is the hypervolume of a cumulative archive of the non-dominated plans found so far, so a mean that is still rising at generation 12,000 means new non-dominated plans are still being found. The archive is capped at 300 points, and the small dips visible in a few individual traces occur when that cap prunes points from the archive.




Figure 2. Convergence of the search for profile P1 over the 12,000-generation budget. (a) Hypervolume of the cumulative archive (capped at 300 points) against generation for each of the 21 seeds (grey), with the mean (red). Small dips in individual traces come from pruning at the cap. (b) Per-generation hypervolume improvement rate, raw (grey) and smoothed over a 15-generation window (red). Hypervolume uses a reference point of 1.1 in every normalised objective. The values plotted use the figure's own normalisation and are not on the same scale as the seed statistics measured against the pooled front, which are reported in the text.

For choosing a budget, the improvement rate in Fig. 2b is more useful. The smoothed per-generation gain peaks at roughly 5×10^{-4} near generation 200, falls below 2×10^{-4} by about generation 500, and then stays low with occasional bumps when a single seed makes a step gain. It stays near zero only after about generation 7,000. The 12,000-generation budget is therefore a judgement about how long to wait for the next step, not a point at which the search has provably finished. The feasible share of the population rises from 0% at the first sample to 100% at the last, and every seed ends fully feasible, so the constraints are not what limits convergence in this run.

The pooled front holds 135 non-dominated plans (Fig. 3). Plans that score highest on mood are not the ones that score highest on muscle, and between the satiety and unhealthy-nutrient axes the lines fan out, so those two objectives trade against each other instead of moving together. Those trade-offs are what justify solving the six objectives jointly. One result from the front concerns the objective set itself. Across the 135 plans, mood and hair correlate at +0.92, above the 0.9 level at which the framework flags a pair for merging. The two scores share iron, zinc and vitamin D in their nutrient bases. The other correlations are the ones the formulation intends: muscle against satiety at -0.56 , and the unhealthy composite almost independent of mood at $+0.17$.




Figure 3. Pooled Pareto front for profile P1 in parallel coordinates. Each of the 135 grey lines is one non-dominated plan and the recommended plan is shown in red. Arrows show the direction of improvement: the four support scores are maximised and the unhealthy-nutrient score minimised. The energy axis shows each plan's total energy in kcal; the objective minimised on it is the deviation from the 2,053 kcal target. Labelled values give the range of each axis across the front.

Comparison with random feasible diets

A Pareto front is optimal by construction, so the more useful question is what the search gains over the feasible region it explores. Table 2 sets the recommended plan against 2,000 random feasible diets, each matched to the 2,053 kcal target and checked against every tolerable upper limit. The plan beats every random diet on mood, hair and muscle support and on nutrient adequacy. Two rows are less clean. On satiety it scores 0.492 against a random best of 0.493, and on the unhealthy-nutrient composite it sits at the 91st percentile, the best

random diet (0.189) being better than any plan on the front. That is not a failure of the search. A random diet can happen to be low in saturated fat with no regard for the four support scores, whereas every plan on the front must be non-dominated across all six objectives.

Table 2. The recommended plan for P1 against 2,000 random feasible diets. All random diets are energy-matched to 2,053 kcal and within every tolerable upper limit. Percentile is the share of random diets the recommended plan beats. The unhealthy-nutrient score is minimised, so lower is better on that row.

Metric	Recommended	Random mean	Random best	Percentile
Mood support	0.711	0.461	0.620	100%
Hair support	0.674	0.519	0.639	100%
Muscle support	0.629	0.477	0.601	100%
Satiety support	0.492	0.371	0.493	100%
Unhealthy nutrients	0.295	0.388	0.189	91%
Mean Adequacy Ratio	95.4%	60.4%	81.3%	100%

Adequacy is where the gap is widest. The recommended plan reaches a Mean Adequacy Ratio (MAR) of 95.4%, against a random mean of 60.4% and a random best of 81.3%. Over the whole front the median MAR is 95.44%, against a random median of 60.96%. A two-sided Mann–Whitney U test ($n = 135$ front plans, $n = 2,000$ random diets) gives $U = 270,000$ and $P = 1.8 \times 10^{-84}$, with a rank-biserial correlation of 1.000. That is complete separation: every plan on the front has a higher MAR than every random diet. Across the front, MAR has a mean of 94.83% and ranges from 86.08% to 97.62%.

Comparison with Mediterranean, DASH and MIND patterns

The recommended plan was compared with three reference diets built from the published patterns and scored with the same code and the same personalised constants (Table 3). Protein separates the four plans most sharply. Against an activity-scaled floor of 133.2 g/day the optimised plan supplies 135.2 g, while the Mediterranean, DASH and MIND plans reach 71%, 84% and 78% of it. None of them was designed around a floor scaled to body weight and activity, and at an 801 kcal deficit the shortfall is the difference between roughly 1.3 to 1.5 g/kg and the 1.8 g/kg targeted here.

Table 3. Nutrient delivery of the optimised plan and three reference patterns for P1 at 2,053 kcal. Macronutrients are given in grams with the percentage of the deficit target in

parentheses. Micronutrients are given as a percentage of the RDA or AI floor. Med., Mediterranean; MAR, Mean Adequacy Ratio.

Nutrient	Target	NSGA-III	Med.	DASH	MIND
<i>Macronutrients</i>					
Protein (g)	133.2	135.2 (101%)	95.0 (71%)	111.5 (84%)	103.7 (78%)
Carbohydrate (g)	241.5	241.6 (100%)	229.6 (95%)	278.0 (115%)	225.4 (93%)
Fat (g)	61.6	65.7 (107%)	90.8 (147%)	61.6 (100%)	88.4 (144%)
Fibre (g)	28.7	32.7 (114%)	40.0 (139%)	46.9 (163%)	49.0 (171%)
<i>Micronutrients (% of floor)</i>					
Magnesium (mg)	400	126%	133%	153%	165%
Zinc (mg)	11	326%	107%	118%	106%
Vitamin C (mg)	90	109%	367%	280%	336%
Vitamin D (µg)	15	136%	8%	28%	1%
Vitamin E (mg)	15	76%	98%	132%	157%
Selenium (µg)	55	548%	216%	203%	188%
Vitamin A (µg)	900	112%	88%	261%	162%
Folate (µg)	400	115%	180%	147%	184%
Iron (mg)	8	402%	204%	184%	222%
Vitamin B-6 (mg)	1.3	201%	228%	227%	228%
Potassium (mg)	3,400	143%	115%	142%	117%
Calcium (mg)	1,000	100%	72%	116%	75%
Niacin (mg)	16	194%	129%	178%	216%
Thiamin (mg)	1.2	116%	144%	146%	157%
Riboflavin (mg)	1.3	130%	162%	167%	147%
Vitamin B-12 (µg)	2.4	2,345%	239%	143%	148%
Copper (mg)	0.9	504%	256%	204%	221%
Phosphorus (mg)	700	330%	232%	273%	223%
MAR	—	95.4%	92.6%	96.0%	93.1%
Floors met	—	18/21	14/18	17/18	16/18

The micronutrient picture is closer. The optimised plan reaches a MAR of 95.4%, ahead of MIND (93.1%) and the Mediterranean pattern (92.6%) but just behind DASH (96.0%). Only 36.3% of plans on the front exceed DASH, so the typical plan from this method is slightly less complete in micronutrients than a pattern designed by an expert committee. Against MIND and the Mediterranean pattern the ordering reverses, with 86.7% and 97.8% of front plans ahead. Vitamin E is the plan's weakest point, at 76% of its floor, a consequence of the fat-composition problem discussed below. This comparison covers P1 only, because the three reference plans were built and scored once at that profile's 2,053 kcal target.

Consistency across profiles

To check whether the method behaves consistently across the people it is meant to serve, six further standard-deficit profiles and two GLP-1 profiles were generated. They cover both sexes, ages 23 to 63, weights 58 to 95 kg and all three activity categories, giving energy targets from 1,168 to 3,105 kcal (Table 4).

Table 4. Results across eight profiles. P1 to P6 use a user-specified deficit and G1 and G2 are GLP-1 profiles. Hair is inactive (–) where the deficit is below 800 kcal/day and the person is not on a GLP-1 drug. P1 and G1 ran for 12,000 generations and the other six for 3,000. MAR and sodium are as issued after the rebuild stage. Musc., muscle; Sat., satiety; Unh., unhealthy-nutrient composite; Prot., protein as % of target; Floors, share of soft minimums met.

ID	Target (kcal)	Mood	Hair	Musc.	Sat.	Unh.	Prot. (%)	MAR (%)	Floors (%)
P1	2,053	0.711	0.674	0.629	0.492	0.295	101	95.4	85.7
P2	1,585	0.613	–	0.564	0.518	0.352	103	91.7	82.9
P3	2,110	0.707	–	0.615	0.476	0.268	102	94.8	88.6
P4	1,701	0.616	–	0.601	0.479	0.307	100	90.5	77.1
P5	1,168	0.551	–	0.549	0.494	0.299	102	79.5	57.1
P6	3,105	0.718	0.646	0.645	0.460	0.289	105	95.2	77.1
G1	1,200	0.591	0.526	0.530	0.484	0.279	101	80.7	57.1
G2	1,500	0.633	0.624	0.594	0.494	0.319	100	85.0	65.7

The protein floor binds in every case, and all eight plans clear it, from 100% at the tightest targets to 105% at P6. That is the most consistent result in the table and it bears directly on the problem the work set out to address. Nutrient adequacy follows the energy budget and falls away sharply at the bottom. MAR ranges from 79.5% to 95.4% and tracks the target almost monotonically, because a per-kilogram protein floor claims a share of energy that does not shrink as the target falls. On a small budget the macronutrient floors take nearly everything and micronutrient adequacy is what the search gives up.

Sodium does not follow the energy target, but the highest loads have a common cause: heavy use of canned and preserved proteins, which carry both the mineral density the support scores reward and the sodium the composite penalises. All eight plans stay inside the 2,300 mg ceiling. The conditional hair objective behaves as specified, active for P1, P6, G1 and G2 and returning zero elsewhere.

The constraint system failed once, at the largest budget. On P6 the fibre target of 43.5 g exceeded the absolute 40 g ceiling, and the band logic clamped the ceiling to the target, leaving a single admissible value that a third of the seeds could not reach. The framework did not detect this and printed its standard note that constraints were not limiting. A band whose width falls to zero should be reported as an infeasible configuration before the search starts. This is a flaw in how the constraints are constructed, not in NSGA-III.

Practitioner review

The comparisons so far measure the plan against reference patterns and reference intakes. Neither says whether a practitioner would hand it to a client. Eight plans were therefore reviewed by a nutrition counsellor (co-author S.K.) and scored out of ten on five criteria (Table 5).

Table 5. Practitioner scores for the eight reviewed plans, out of ten. One reviewer applied the same five criteria to every plan.

Criterion	P1	P2	P3	P4	P5	P6	G1	G2
Nutritional adequacy	9	8	9	7	3	7	6	6
Practicality	6	7	8	6	6	6	7	5
Safety	7	8	8	7	3	8	5	6
Realism as a deliverable plan	4	6	7	5	3	4	5	4
Overall acceptability	6.5	7.2	8	6.2	3.7	6.2	5.75	5.2

Two criticisms came up for every plan. Fat is rewarded by mass and never by composition, which leaves vitamin E and linoleic acid short at every energy level and in both drug branches. And no constraint anywhere in the formulation concerns how a plate reads as a meal, so the acceptability problems differ from plan to plan but are never absent. Realism is the one criterion that does not follow the energy budget. P6, with the largest target, scores 4, the same as P1 and G2. Adequacy can be bought with energy. Acceptability cannot, because nothing in the objective set rewards it. Within the objectives defined here the optimiser does what it was built to do, and every criticism the reviewer raised falls outside them. These scores come with a clear caveat: they are one practitioner's judgement on a rubric that is not a validated instrument, no inter-rater agreement was measured, and the reviewer is a co-author, so the evaluation is not independent of the model.

Discussion

The main finding is simple. When the side effects of a calorie deficit are written into the objectives, the resulting plans hold protein where it needs to be across a 2.7-fold range of energy targets, and they reach micronutrient adequacy that random feasible diets never approach. Three widely recommended patterns, scaled to the same energy, fall 16 to 29% short on protein for an active adult in a deficit. That gap reflects purpose as much as quality. The Mediterranean, DASH and MIND patterns were designed for long-term health in a general population, not for protecting lean mass during weight loss, and the comparison should be read that way.

Some of what the optimiser leads on is circular. The support scores are its own objective functions, so its advantage on them is expected and says little. The independent evidence is the nutrient comparison: protein, micronutrients against published floors, and harmful nutrients against published limits. On those measures the method matches or beats the reference patterns for P1, with DASH still slightly ahead on MAR.

The objective set is also less six-dimensional than it looks. Mood and hair correlate at +0.92 on P1's front, so its effective dimensionality is closer to five. The merge test that should catch this is applied per run, and it reported the overlap inconsistently, flagging mood against hair on one profile and mood against muscle on another. That suggests the overlap comes from the shared nutrient bases rather than from any single run. Some objectives also barely vary across the front: muscle support spans only 0.038 and satiety 0.042, against 0.088 for mood. An objective that hardly varies over the non-dominated set does little to separate plans, whatever weight the selection rule gives it. Merging or re-basing the overlapping scores is an obvious next step.

The gap between what the optimiser measures and what a practitioner would issue is, we think, the most useful result of the study. The plans are nutritionally efficient, yet they scored between 3 and 7 out of 10 for realism as deliverable documents. A diet nobody follows delivers none of the benefits the objectives measure. Cost, availability, preparation effort and familiarity are therefore not refinements to add later. They are conditions for the approach being usable at all.

Two further limitations sit inside the model. The fat objective rewards mass rather than composition, so essential fatty acid and vitamin E floors are missed even when total fat is on

target. And the constraint construction can, at a large energy target, clamp a two-sided band to a single value without flagging the configuration as infeasible, as the P6 run showed.

The data impose a limit of their own. Missing nutrient values in the merged table are set to zero, and some nutrients are sparsely measured in the source: biotin, iodine and molybdenum are missing for more than 99% of foods, and vitamin D, EPA and DHA for about three-quarters. A food's true content of these is often unknown rather than nil. Treating unknown as zero makes any score or adequacy figure that relies on these nutrients conservative, and for the sparsest of them unreliable, which is why the micronutrient comparison is restricted to well-measured nutrients.

The evaluation has boundaries too. The practitioner review rests on a single reviewer using an unvalidated rubric. The reference-pattern comparison was run for one profile, so it shows the method can outperform three published patterns for that profile, not that it does so in general. The cross-profile evidence relies on internal measures (MAR, floor coverage and objective scores), and because six of the eight profiles ran for 3,000 generations rather than 12,000, the difference in budget is confounded with the energy target. Finally, the search algorithm itself was not compared with alternatives such as AGE-MOEA or SMS-EMOA, which outperformed NSGA-III on a related menu-planning problem [16].

Future work follows from these results. Adding cost, familiarity and plate structure to the objective set, replacing the fat-mass term with one sensitive to fat composition, and checking constraint bands for zero width before the search begins would close most of the gap between a plan that is nutritionally optimal and one a client will actually eat. Validation with independent dietitians, and eventually with people following the plans, would test whether the support scores translate into the outcomes they are named after.

Methods

Personalised energy and nutrient targets

The system collects sex, age, weight, height, activity level, a flag for GLP-1 receptor agonist or other weight-loss medication and, for women, pregnancy or lactation status. From these it derives maintenance energy, an energy target, macronutrient targets and a full set of micronutrient floors and ceilings. Every scoring constant in the objective functions is then calibrated against those targets.

Maintenance energy and the deficit

Maintenance energy for adults is the Mifflin–St Jeor resting energy expenditure [19] multiplied by the physical activity level of the FAO/WHO/UNU categories [20]: 1.40 sedentary, 1.55 light, 1.70 active and 1.90 vigorous. For children and adolescents the system uses the age- and sex-specific Estimated Energy Requirement equations instead, and pregnancy and lactation add their published increments. A deficit is applied only to non-pregnant, non-lactating adults aged 19 or over. For everyone else the energy target equals maintenance energy and the macronutrient targets are the ordinary life-stage values. This boundary is deliberate, since intentional energy restriction is not appropriate in those groups.

Macronutrient targets

Where a deficit applies, the energy target is maintenance energy minus the deficit, and macronutrient targets are set in sequence so that protein and fat are protected and carbohydrate absorbs what remains. Protein comes first, at an activity-scaled 1.6, 1.8 or 2.0 g/kg of body weight for a light, active or vigorous lifestyle. This exceeds the 0.8 g/kg RDA because higher protein intake spares lean mass during energy restriction [21,22], and it stays inside the protein acceptable macronutrient distribution range (AMDR). Fat is set next, at 27% of the target energy, as a floor for hormone synthesis and absorption of the fat-soluble vitamins [23,24]. Carbohydrate receives the energy left after the protein and fat floors but is never allowed below 130 g/day, the RDA that covers the brain's glucose requirement [23]. Fibre is set at 14 g per 1,000 kcal of target energy [25].

Two of these figures were revised during development. An earlier protein ladder of 1.6, 2.0 and 2.2 g/kg suited a resistance-trained athlete in a hard cut rather than a moderately active adult: protein ran near 30% of energy while fat sat on its floor, and the linoleic acid and vitamin E minimums became unreachable. The fat floor was raised from 20% to 27% at the same time, since every plan had come out at 20 to 21% fat with linoleic acid near 5% of its target. The search had been right to fail; the targets were wrong. Each plan is now also checked against acceptable ranges of 10 to 40% protein, 15 to 40% fat and 20 to 65% carbohydrate, which are reported but not enforced.

Micronutrients

Micronutrient RDA or AI floors and tolerable upper intake levels (ULs) are taken from the Dietary Reference Intake tables for the person's life stage and sex, covering 23 tracked micronutrients, 16 of which carry a UL for a healthy adult. ULs are applied at 90% of the

published value, so that a plan sitting on its ceiling still leaves room for intake the model does not see. The 90% margin is a modelling choice made for this work, not a published figure.

Food database and candidate pool

The food-composition table is built from USDA FoodData Central [26] by merging the Foundation Foods and SR Legacy data types. Energy is harmonised to kilocalories, using the Atwater specific-factor value where available and converting kilojoule entries at 4.184 kJ per kcal. Rows with no energy value are removed and missing nutrient cells are set to zero. The result is a table of 7,964 foods described by energy and 39 nutrient fields (macronutrients, fibre, sugars, added sugars, leucine, the individual fatty-acid classes, sodium and 23 micronutrients). Setting missing values to zero keeps every food evaluable but treats an unmeasured nutrient as absent, which is discussed as a limitation.

Each food description is passed through a rule-based classifier that uses word-boundary regular expressions and gives precedence to the first comma-separated segment of the description, so a food is classified by what it is rather than by an ingredient or flavour it mentions. Each row receives one of eight food types and one of 61 finer subgroups, 57 of which carry a meal role, which lets a user restrict the pool to foods they will actually eat. Seven of the types occur in the final pool: vegetarian (50.3%), non-vegetarian (35.1%), dairy (8.5%), mixed (2.6%), fats and oils (1.7%), fat (1.0%) and protein (0.7%). Six filters are then applied: infant and toddler products; rows confirmed by audit to carry corrupted nutrient values; foods whose subgroup has no meal-role mapping; a locale filter inferred from the pool; a processing filter above NOVA category 3 [27] together with a condiment filter; and a safety and practicality filter that removes raw meats, analytic components and branded or restaurant entries. For the reported runs this leaves 2,307 candidate foods in 47 subgroups. The safety filter matters most: it removes 706 raw-meat rows, and applying it before the candidate set is fixed means an unsafe or unservable food can never be selected.

Decision variables and structural rules

Let n be the number of candidate foods. A candidate plan is a vector of serving sizes, one per food,

$$\mathbf{x} = (x_1, \dots, x_n)^T \quad (1)$$

where x_i is the amount of food i in units of 100 g. With a_{ij} the amount of nutrient j in 100 g of food i and e_i its energy per 100 g, the daily totals are linear in $\mathbf{x}$:

$$N_j(\mathbf{x}) = \sum_i a_{ij} x_i, E(\mathbf{x}) = \sum_i e_i x_i, W(\mathbf{x}) = 100 \sum_i x_i, \rho(\mathbf{x}) = E(\mathbf{x})/W(\mathbf{x}) \quad (2)$$

where W is the total food mass in grams and ρ the energy density of the day's intake. All objectives are functions of these plan-level totals and never of individual foods, so every reported outcome describes the whole day. Three structural rules are enforced by a repair operator applied after every genetic operation, not by penalty terms. The structural feasible set is

$$\Omega = \{\mathbf{x}: x_i \in \{0\} \cup [\ell_i, u_i], \quad m \leq |\{i: x_i > 0\}| \leq M_c, \quad x_i > 0 \Rightarrow 100x_i \in \delta\mathbb{Z}\} \quad (3)$$

where $[\ell_i, u_i]$ are the per-food serving bounds, $(m, M_c) = (5, 20)$ are the cardinality bounds and δ is the serving grid (5 g by default). Semi-continuity means a food is either absent or present at no less than its minimum serving. Cardinality keeps between 5 and 20 foods active; when the upper bound is exceeded, foods are retained by meal role in round-robin order. Discretisation snaps active servings to the grid. The repair operator also scales active servings by a common factor towards the energy target, then re-clips and re-snaps them, so no structurally infeasible plan is ever evaluated and the search spends its budget on nutrition rather than on rediscovering the shape of a diet plan.

Objective functions

Both scoring forms in the model are the Hill equation [28] at different values of one parameter. With intake v , a reference quantity K and Hill coefficient n ,

$$h(v, K, n) = \frac{(v/K)^n}{1 + (v/K)^n} \in [0, 1] \quad (4)$$

The shape of the response is set by n alone, and K fixes the point at which the term reaches one half. At $n = 1$ this reduces to the Michaelis–Menten form [29], used for every beneficial nutrient term, with K set to the user's own adequacy target:

$$s(v, K) = h(v, K, 1) = \frac{v}{K + v} \quad (5)$$

Intake below target contributes steeply and intake beyond it gives diminishing returns, so the optimiser is rewarded for reaching adequacy and gains little from exceeding it. At $n = 2$ the same equation becomes sigmoid, and this is the penalty applied to every harmful nutrient:

$$p(v, K) = h(v, K, 2) = \frac{(v/K)^2}{1 + (v/K)^2} \quad (6)$$

It equals 0.5 at the limit $v = K$ and stays bounded as v grows. The squared numerator flattens the curve near the origin, so a small amount of a harmful nutrient is tolerated, and steepens it

as the limit approaches, so pressure is applied before the limit is crossed rather than only after. Each of the mood, hair and muscle scores is a convex combination of saturation terms whose AHP weights [30] sum to one:

$$S_{\bullet}(\mathbf{x}) = \sum_k w_{\bullet,k} s(N_k(\mathbf{x}), K_k) \quad (7)$$

where $\bullet$ stands for mood, hair or muscle. The terms, in descending order of weight, are: for mood, magnesium, zinc, vitamin D, omega-3, iron, selenium and vitamin B-6; for hair, iron, protein, zinc and vitamin D; for muscle, protein, iron, carbohydrate, potassium, leucine, magnesium and omega-3 (Table 1). Satiety is the exception, because its energy-density term is decreasing:

$$S_{\text{Sat}}(\mathbf{x}) = 0.492 \frac{K_{\rho}}{K_{\rho} + \rho(\mathbf{x})} + 0.307 s(N_{\text{prot}}, K_{\text{prot}}) + 0.201 s(N_{\text{fib}}, K_{\text{fib}}) \quad (8)$$

with $K_{\rho} = 1.0$ kcal/g. Energy density carries almost half of the satiety score. Of the three surviving terms it is the only one supported by randomised trials with objectively measured intake, and a meta-analysis of 38 such trials found that lowering energy density reduced intake by about 223 kcal while portion size was held constant [31].

Two objectives are minimised. The unhealthy-nutrient composite combines four sigmoid penalties,

$$S_{\text{Unh}}(\mathbf{x}) = 0.274 p(N_{\text{Na}}, K_{\text{Na}}) + 0.248 p(N_{\text{sat}}, K_{\text{O}}) + 0.248 p(N_{\text{trans}}, K_{\text{M}}) + 0.230 p(N_{\text{sug}}, K_{\text{N}}) \quad (9)$$

where $K_{\text{M}} = 0.01E/9$, $K_{\text{O}} = 0.05E/9$ and $K_{\text{N}} = 0.10E/4$ follow the WHO limits at the person's energy target E [32], and $K_{\text{Na}} = 2,000$ mg is fixed. The saturated-fat constant uses 5% of energy rather than the WHO limit of 10% [33], so optimisation pressure starts well below the guideline ceiling. As a result, the composite compares plans with one another and is not a measure of compliance with guidance. Energy deviation uses the same bounded family, with r the fractional deviation from target and β the half-width of the energy band:

$$S_{\text{Edev}}(\mathbf{x}) = \frac{r}{r + \beta/3} \quad (10)$$

so a plan sitting exactly on the band edge ($r = \beta$) scores 0.75 rather than saturating at 1.

Evidence audit and weights

Weights come from a two-criterion analytic hierarchy process [30]. Each candidate term is scored for effect size and for evidence quality against a documented rubric, the two scores are

multiplied, and the resulting ratios are snapped to Saaty's integer scale to build the pairwise comparison matrix. The procedure runs separately within each objective, so weights are comparable inside a score but not across scores. One rule was applied throughout: a term whose own citation reports a null or unfavourable primary result is removed, not down-weighted. A zero weight and a small weight are different claims, and the second cannot be defended when the evidence points the other way. Fourteen terms were removed on this basis (Supplementary Table S2). The hair score shrank the most, from nine terms to four, and it carries a caveat: no term in it now rests on an intervention trial, so the score measures nutritional support, not predicted hair growth.

Conditional objectives under appetite suppression

The hair score is computed only when the calorie deficit is at least 800 kcal/day or the person takes a GLP-1 receptor agonist; otherwise it returns zero. Rapid weight loss is a documented trigger of telogen effluvium [2], and appetite suppression by GLP-1 receptor agonists produces rapid loss whatever deficit the user reports, which is why the drug flag activates the objective on its own. Pharmacovigilance analyses have reported a disproportionality signal for alopecia under semaglutide and tirzepatide [34], but that evidence is observational and its authors attribute the effect largely to the weight loss itself. The gate is therefore justified by the rate of loss, with the drug flag used as a proxy for it. The 800 kcal/day figure is a modelling parameter, not a clinical cut-off. Using the conventional approximation of 7,700 kcal per kilogram [35] it corresponds to about 0.73 kg of weight loss per week.

For GLP-1 users the energy target is a sex-specific fixed value (1,200 kcal for women, 1,500 kcal for men [36]) rather than a user-specified deficit. Two safety reports accompany it. Any target below the sex-specific threshold is flagged as needing supervision, and any target below 800 kcal is identified as a formal very-low-calorie diet [37]. Neither blocks plan generation, because a low target may be appropriate for a patient whose intake is already limited by the drug.

Constraint handling

Writing the maximised scores in minimisation form and adding the soft penalty $\Pi(\mathbf{x})$ to every component gives the constrained many-objective problem

$$\min_{\mathbf{x} \in \Omega} \mathbf{F}(\mathbf{x}) = (-S_{\text{Mood}}, -S_{\text{Hair}}, -S_{\text{Musc}}, -S_{\text{Sat}}, S_{\text{Unh}}, S_{\text{Edev}})^T + \Pi(\mathbf{x}) \cdot \mathbf{1} \quad (11)$$

subject to hard constraints in the standard form $g(\mathbf{x}) \leq 0$: micronutrient upper limits at a 90% margin, $N_j(\mathbf{x}) - 0.9U_j \leq 0$; a two-sided energy band $E(1 - \beta) \leq E(\mathbf{x}) \leq E(1 + \beta)$ with $\beta = 0.02$; a two-sided sodium band of 1,500 to 2,300 mg; per-subgroup gram caps; and per-subgroup limits on repeated foods. A UL marks the point beyond which risk starts to rise, so a plan sitting exactly on it leaves no room for the difference between a modelled plan and a real day's eating, and the 90% margin provides that room. Trans fat, saturated fat and added sugar are left out of the hard set and handled through the unhealthy-nutrient objective, so a plan that exceeds a limit is penalised in proportion rather than declared infeasible. Sodium is both scored and bounded, because its chronic-disease ceiling is a clinical limit rather than a preference.

Requirements that would make the problem infeasible if enforced strictly are treated as soft shortfalls. Let T be the set of soft minimums: the macronutrient targets, the micronutrient RDA floors and five subgroup gram floors (vegetables 350 g, fruit 200 g, whole grains 48 g, legumes 35 g and healthy fats 15 g). Then

$$\Pi(\mathbf{x}) = \lambda \frac{1}{|T|} \sum_{t \in T} \max\left(0, \frac{T_t - N_t(\mathbf{x})}{T_t}\right) + \lambda_p \pi_{\text{prot}} + \lambda_m \pi_{\text{macro}} + \lambda_f \pi_{\text{div}} \quad (12)$$

with $\lambda = 0.5$, $\lambda_p = 1.0$, $\lambda_m = 0.4$ and $\lambda_f = 0.6$, where π_{prot} is the fractional protein shortfall, π_{macro} the mean shortfall of the remaining macronutrient and subgroup targets, and π_{div} the shortfall in distinct fat and dairy foods against a minimum of three. The three extra terms are additive rather than folded into the mean, so raising pressure on protein, macronutrients and diversity does not weaken the pull of the micronutrient floors. The 38 per-subgroup gram ceilings and the requirement of at least three distinct fat or dairy foods were set in consultation with a practising dietitian (see Acknowledgements), since the Dietary Reference Intakes bound nutrients, not grams of a food subgroup, and no published table supplies these values. They serve two purposes. One is realism: a plan that meets every nutrient target with 400 g of cheese satisfies the mathematics and nobody's dinner. The other is protection against exploitation of the objective functions, since several nutrients can be obtained cheaply from a single dense food.

Optimisation with NSGA-III

The problem is solved with NSGA-III [38] as implemented in pymoo 0.6.2 [39] under Python 3.13.5, using Riesz s-energy reference directions [40] (210 points, matched to a population of 210), simulated binary crossover [41] with distribution index 30 and probability 1.0,

polynomial mutation [42] with distribution index 40, duplicate elimination and the repair operator described above. NSGA-III was chosen because selection based on Pareto dominance alone weakens as the number of objectives grows: with six objectives most of a population is mutually non-dominated and crowding distance can no longer discriminate between individuals. NSGA-III replaces crowding with association to a fixed set of reference directions, which preserves spread in higher dimensions [38].

The per-variable mutation probability follows from the sparsity of the representation. Polynomial mutation conventionally uses one over the number of variables, which works when every variable carries information. Here at most 20 of 2,307 foods can be active, so about 99.1% of the vector is zero and a mutation lands on an active food in only about one of every 115 individuals. Setting the rate to one over the cardinality bound gives about one mutation per plan among the foods it actually contains. This single change accounted for most of the observed improvement in front size and hypervolume.

Plan selection and meal assignment

The fronts from all seeds are pooled and reduced to one global non-dominated set. One plan is recommended from this set by a multi-criteria decision rule: each objective is min–max normalised across the front, and the plan k^* with the smallest weighted distance to the ideal point is selected,

$$k^* = \operatorname{argmin}_k \sum_q w_q \frac{F_q(k) - \min_l F_q(l)}{\max_l F_q(l) - \min_l F_q(l)}, \quad \sum_q w_q = 1 \quad (13)$$

where $F_q(k)$ is the value of objective q for plan k and w_q are the selection weights, equal by default. The selected plan is split into five meals (breakfast, morning snack, lunch, afternoon snack and dinner). Five slots were chosen because the International Society of Sports Nutrition recommends protein doses of 20 to 40 g spread every three to four hours across the day, over four to six meals [43]. Each food is mapped to one of four meal roles through its subgroup and placed whole, largest serving first, into that role's slots. A serving is divided between slots only when a role has fewer distinct foods than mandated slots.

Post-optimisation rebuild and use of a language model

A rebuild stage runs after every core output has been written to disk, and it cannot abort the run. It does three things. Foods still flagged as unsafe or impractical are substituted from the same candidate table, matched on subgroup and energy so that the nutrient profile is

preserved. Foods placed in an unsuitable meal role are moved, provided both meals stay inside their calorie bands. Finally, preparation text is generated for each meal from its ingredient list using a large language model through the Google Gemini API [44]. The API was called without pinning a model version, so the service used whichever Gemini model it offered at the time of each call. If no API key is supplied, recipes fall back to a templated form. Using a language model inside a research pipeline has a reproducibility cost, so it is stated plainly here. Everything before this stage is deterministic, and the rebuild is deterministic in everything measurable (food choices, portions and meal assignments), but the recipe wording is not, both because generation is stochastic and because the underlying model can change between runs. Generated text is checked against the ingredient list and rejected if it introduces anything new, and no quantity reported in Results passes through this stage. Supplementary Fig. S1 shows the layered system architecture in which these stages run.

Evaluation protocol

The primary results come from one full run for profile P1 (Supplementary Table S3), with all six objectives active. Solution quality is reported per seed as hypervolume [45] with a reference point of 1.1 in every normalised dimension, and as IGD+ against the pooled reference front, together with the feasible percentage of the final population. Convergence is sampled every 25 generations. At each sample the current non-dominated set is merged into a cumulative archive, capped at 300 points to keep the six-objective hypervolume computation affordable, and the hypervolume of that archive is recorded. Each seed starts from the base seed plus its index, so a run is reproducible from its configuration. A complete nutrient panel is produced so the recommended plan can be checked against every floor and ceiling.

Two comparisons follow. The first is against 2,000 random feasible diets, each energy-matched to the 2,053 kcal target and screened against every UL. The second is against Mediterranean [46], DASH [47] and MIND [48] plans built from the published patterns, converted to grams using standard household measures, matched to rows of the same USDA table and scaled to deliver exactly 2,053 kcal. Seven further profiles (P2 to P6, G1 and G2) were run to test consistency; P1 and G1 used the full 12,000-generation budget and the other six used 3,000 generations. Finally, eight generated plans were reviewed by a nutrition counsellor and scored out of ten on nutritional adequacy, practicality, safety, realism as a deliverable plan and overall acceptability.

All physiological profiles are hypothetical. No human participants were involved, and the practitioner review assessed generated documents, not people.

Statistical analysis

MAR is the mean of the nutrient adequacy ratios, each ratio being intake divided by its floor and truncated at one, expressed as a percentage. The truncation stops a nutrient eaten far above its floor from masking a deficiency elsewhere. Unlike the classical MAR, which uses a small set of nutrients with an RDA, the mean here runs over the whole soft-floor list, so it is stricter and not directly comparable with published values. A nutrient is scored and counted only if at least 60% of the candidate pool carries a value for it, and nutrients below that threshold are reported as unknown rather than deficient. Front plans and random diets were compared on MAR with a two-sided Mann–Whitney U test in SciPy 1.18.0, chosen because the MAR distributions are bounded and not assumed to be normal. With separation this complete the effect size is the more informative statistic, so the rank-biserial correlation is reported alongside P ($\alpha = 0.05$). Correlations between objectives were computed across the 135 non-dominated plans of the pooled P1 front. Percentiles in Table 2 give the share of random diets that the recommended plan beats.

Use of AI-assisted tools

Apart from the recipe-text generation described above, which is part of the method, AI-assisted copy-editing was used to improve the language and readability of the manuscript. The authors reviewed all edited text and take full responsibility for the content.

Data availability

The food-composition data are openly available from USDA FoodData Central (<https://fdc.nal.usda.gov>), from which the Foundation Foods and SR Legacy data types were drawn. The merged and cleaned food table, the derived data supporting the findings of this study and the source code are available from the corresponding author on reasonable request.

References

1. Weiss, E. P., Jordan, R. C., Frese, E. M., Albert, S. G. & Villareal, D. T. Effects of weight loss on lean mass, strength, bone, and aerobic capacity. *Med. Sci. Sports Exerc.* **49**, 206–217 (2017).

- 566 2. Kang, D. H., Kwon, S. H., Sim, W. Y. & Lew, B. L. Telogen effluvium associated with
567 weight loss: a single center retrospective study. *Ann. Dermatol.* **36**, 384–388 (2024).
- 568 3. Jackson, S. E., Steptoe, A., Beeken, R. J., Kivimaki, M. & Wardle, J. Psychological
569 changes following weight loss in overweight and obese adults: a prospective cohort
570 study. *PLoS ONE* **9**, e104552 (2014).
- 571 4. Karl, J. P. et al. Transient decrements in mood during energy deficit are independent of
572 dietary protein-to-carbohydrate ratio. *Physiol. Behav.* **139**, 524–531 (2015).
- 573 5. Menniti, G. et al. Mental health consequences of dietary restriction: increased
574 depressive symptoms in biological men and populations with elevated BMI. *BMJ Nutr.*
575 *Prev. Health* **8**, e001167 (2025).
- 576 6. Most, J. & Redman, L. M. Impact of calorie restriction on energy metabolism in
577 humans. *Exp. Gerontol.* **133**, 110875 (2020).
- 578 7. Sumithran, P. et al. Long-term persistence of hormonal adaptations to weight loss. *N.*
579 *Engl. J. Med.* **365**, 1597–1604 (2011).
- 580 8. Nymo, S. et al. Investigation of the long-term sustainability of changes in appetite after
581 weight loss. *Int. J. Obes.* **42**, 1489–1499 (2018).
- 582 9. Randle, M. et al. A systematic review of ecological momentary assessment studies of
583 appetite and affect in the experience of temptations and lapses during weight loss
584 dieting. *Obes. Rev.* **24**, e13596 (2023).
- 585 10. GBD 2015 Obesity Collaborators. Health effects of overweight and obesity in 195
586 countries over 25 years. *N. Engl. J. Med.* **377**, 13–27 (2017).
- 587 11. Stigler, G. J. The cost of subsistence. *J. Farm Econ.* **27**, 303–314 (1945).
- 588 12. Moreira, R. P., Wanner, E., Martins, F. V. & Sarubbi, J. F. An evolutionary mono-
589 objective approach for solving the menu planning problem. In *2018 IEEE Congress on*
590 *Evolutionary Computation (CEC)* 1–8 (IEEE, 2018).
- 591 13. Gaál, B., Vassányi, I. & Kozmann, G. A novel artificial intelligence method for weekly
592 dietary menu planning. *Methods Inf. Med.* **44**, 655–664 (2005).
- 593 14. Türkmenoğlu, Ç., Etaner-Uyar, A. Ş. & Kiraz, B. Recommending healthy meal plans
594 by optimising nature-inspired many-objective diet problem. *Health Inform. J.* **27**,
595 1460458220976719 (2021).
- 596 15. Ahourag, A., El Moutaouakil, K., Cheggour, M., Chellak, S. & Baizri, H.
597 Multiobjective optimization to optimal Moroccan diet using genetic algorithm. *Int. J.*
598 *Eng. Model.* **36**, 67–79 (2023).
- 599 16. Sahin, O. & Aytakin-Sahin, G. Open-source multi-objective optimization software for
600 menu planning. *Expert Syst. Appl.* **252**, 124213 (2024).
- 601 17. Babalola, A., Ojokoh, B., Adetunmbi, A. & Taiwo, O. Non-dominated sorting genetic
602 algorithm-II based optimization for healthy meals. In *2024 International Conference on*
603 *Science, Engineering and Business for Driving Sustainable Development Goals*
604 *(SEB4SDG)* 1–5 (IEEE, 2024).

18. Zhao, L. ACMO-Diet: accelerated constrained multi-objective optimization for diet recommendation. In *2025 IEEE 7th International Conference on Communications, Information System and Computer Engineering (CISCE)* 1249–1254 (IEEE, 2025).
19. Mifflin, M. D. et al. A new predictive equation for resting energy expenditure in healthy individuals. *Am. J. Clin. Nutr.* **51**, 241–247 (1990).
20. FAO/WHO/UNU Expert Consultation. *Human Energy Requirements: Report of a Joint FAO/WHO/UNU Expert Consultation, Rome, 17–24 October 2001*. FAO Food and Nutrition Technical Report Series No. 1 (FAO, 2004).
21. Leidy, H. J. et al. The role of protein in weight loss and maintenance. *Am. J. Clin. Nutr.* **101**, 1320S–1329S (2015).
22. Longland, T. M., Oikawa, S. Y., Mitchell, C. J., Devries, M. C. & Phillips, S. M. Higher compared with lower dietary protein during an energy deficit combined with intense exercise promotes greater lean mass gain and fat mass loss: a randomized trial. *Am. J. Clin. Nutr.* **103**, 738–746 (2016).
23. Institute of Medicine. *Dietary Reference Intakes for Energy, Carbohydrate, Fiber, Fat, Fatty Acids, Cholesterol, Protein, and Amino Acids* (National Academies Press, 2005). <https://doi.org/10.17226/10490>
24. World Health Organization. *Total Fat Intake for the Prevention of Unhealthy Weight Gain in Adults and Children: WHO Guideline* (WHO, 2023).
25. Dahl, W. J. & Stewart, M. L. Position of the Academy of Nutrition and Dietetics: health implications of dietary fiber. *J. Acad. Nutr. Diet.* **115**, 1861–1870 (2015).
26. U.S. Department of Agriculture, Agricultural Research Service. FoodData Central. <https://fdc.nal.usda.gov> (2026). Accessed 21 September 2026.
27. Monteiro, C. A. et al. Ultra-processed foods: what they are and how to identify them. *Public Health Nutr.* **22**, 936–941 (2019).
28. Hill, A. V. The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. *J. Physiol.* **40**, iv–vii (1910).
29. Michaelis, L., Menten, M. L., Johnson, K. A. & Goody, R. S. The original Michaelis constant: translation of the 1913 Michaelis–Menten paper. *Biochemistry* **50**, 8264–8269 (2011).
30. Saaty, T. L. A scaling method for priorities in hierarchical structures. *J. Math. Psychol.* **15**, 234–281 (1977).
31. Klos, B. et al. Impact of energy density on energy intake in children and adults: a systematic review and meta-analysis of randomized controlled trials. *Eur. J. Nutr.* **62**, 1059–1076 (2023).
32. World Health Organization Regional Office for the Eastern Mediterranean. Limit fat, salt and sugar intake. <https://www.emro.who.int/nutrition/reduce-fat-salt-and-sugar-intake/> (n.d.). Accessed 21 September 2026.
33. World Health Organization. *Saturated Fatty Acid and Trans-Fatty Acid Intake for Adults and Children: WHO Guideline* (WHO, 2023).

- 645 34. Godfrey, H., Leibovit-Reiben, Z., Jedlowski, P. & Thiede, R. Alopecia associated with
646 the use of semaglutide and tirzepatide: a disproportionality analysis using the FDA
647 Adverse Event Reporting System (FAERS) from 2022 to 2023. *J. Eur. Acad. Dermatol.*
648 *Venereol.* **39**, e153–e154 (2025).
- 649 35. Wishnofsky, M. Caloric equivalents of gained or lost weight. *Am. J. Clin. Nutr.* **6**, 542–
650 546 (1958).
- 651 36. Jensen, M. D. et al. 2013 AHA/ACC/TOS guideline for the management of overweight
652 and obesity in adults. *Circulation* **129**, S102–S138 (2014).
- 653 37. National Task Force on the Prevention and Treatment of Obesity, National Institutes of
654 Health. Very low-calorie diets. *JAMA* **270**, 967–974 (1993).
- 655 38. Deb, K. & Jain, H. An evolutionary many-objective optimization algorithm using
656 reference-point-based nondominated sorting approach, part I: solving problems with
657 box constraints. *IEEE Trans. Evol. Comput.* **18**, 577–601 (2014).
- 658 39. Blank, J. & Deb, K. pymoo: multi-objective optimization in Python. *IEEE Access* **8**,
659 89497–89509 (2020).
- 660 40. Blank, J., Deb, K., Dhebar, Y., Bandaru, S. & Seada, H. Generating well-spaced points
661 on a unit simplex for evolutionary many-objective optimization. *IEEE Trans. Evol.*
662 *Comput.* **25**, 48–60 (2021).
- 663 41. Deb, K. & Agrawal, R. B. Simulated binary crossover for continuous search space.
664 *Complex Syst.* **9**, 115–148 (1995).
- 665 42. Deb, K. & Goyal, M. A combined genetic adaptive search (GeneAS) for engineering
666 design. *Comput. Sci. Inform.* **26**, 30–45 (1996).
- 667 43. Jäger, R. et al. International Society of Sports Nutrition position stand: protein and
668 exercise. *J. Int. Soc. Sports Nutr.* **14**, 20 (2017).
- 669 44. Google. Gemini API. *Google AI for Developers* <https://ai.google.dev/gemini-api/docs>
670 (2026). Accessed 17 September 2026.
- 671 45. Bader, J. & Zitzler, E. A hypervolume-based optimizer for high-dimensional objective
672 spaces. In *New Developments in Multiple Objective and Goal Programming*, Lecture
673 Notes in Economics and Mathematical Systems Vol. 638, 35–54 (Springer, 2010).
- 674 46. Cleveland Clinic. Mediterranean diet.
675 <https://my.clevelandclinic.org/health/articles/16037-mediterranean-diet> (2022).
676 Accessed 21 September 2026.
- 677 47. National Heart, Lung, and Blood Institute. DASH eating plan.
678 <https://www.nhlbi.nih.gov/health/dash-eating-plan> (2021). Accessed 21 September
679 2026.
- 680 48. Harvard T.H. Chan School of Public Health. Diet review: MIND diet. *The Nutrition*
681 *Source* <https://nutritionsource.hsph.harvard.edu/healthy-weight/diet-reviews/mind-diet/>
682 (2022). Accessed 21 September 2026.

683 **Acknowledgements**

684 The authors thank Rizala Sidiqi, Dietitian and Clinical Nutritionist at Imperial Polyclinic,
685 Srinagar, for providing the per-subgroup portion caps used in the constraint set, and the
686 Department of Information Technology, Central University of Kashmir, for computational
687 support.

688 **Funding**

689 No funding was received for conducting this study.

690 **Author contributions**

691 T.S. conceived the study and carried out the methodology, software development, formal
692 analysis, investigation, data curation and visualisation, and wrote the original draft. A.A.K.
693 supervised the work, contributed to conceptualisation and reviewed and edited the
694 manuscript. S.K. carried out the practitioner review of the generated plans. All authors
695 reviewed the manuscript.

696 **Competing interests**

697 The authors declare no competing interests.

698 **Additional information**

699 **Supplementary information** accompanies this paper as a separate file (Supplementary
700 Tables S1–S3 and Supplementary Fig. S1).

701 **Correspondence** and requests for materials should be addressed to T.S.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [ShakeelSciRepSupplementaryInformation.docx](#)