

# Scaling Laws for Edible Intelligence

Compute, patience, and the emergence of general pastırma capability

Cillop Scaling Systems Group

Cillop Super Intelligence | CSI-2026-001 | October 2026

## Abstract

We study a central constraint in edible intelligence: increasing computational expenditure does not necessarily improve lunch. We introduce the Edible Scaling Surface, a bounded model of the relationship between compute allocation, temporal restraint, and serving-context compatibility. Across a synthetic suite of 4,096 plate configurations, allocating resources to context selection produces larger gains than expanding verbal descriptions of the product. Our strongest configuration reaches 94.8 on the Cillop Readiness Index while reducing unnecessary explanation by 71%. A patience-constrained baseline outperforms a compute-only system despite a substantially smaller adjective budget. These findings support a transition from general-purpose intelligence toward systems whose reasoning terminates in an appropriately selected slice.

$$R = 100 * \text{sigmoid}(0.7 \log(1+C) + 1.8P + 2.4K - 2.1V - 1.6)$$

## 1. Introduction

General-purpose models are rewarded for producing increasingly elaborate accounts of the world. At the lunch table, elaboration has a different relationship with utility. A model that can explain the historical significance of a slice but cannot decide where to place it has not completed the task. We define edible intelligence as the capacity to select an appropriate serving configuration under contextual, social, and plate-space constraints.

Our central hypothesis is that capability is bottlenecked by allocation, not scale alone. The objective is neither maximal spice intensity nor maximal explanation length. It is a stable match between the slice, its context, and the expectations of the recipient. We therefore distinguish additional computation from useful computation and distinguish waiting from an agent repeatedly announcing that it is waiting.

This report introduces a controlled synthetic evaluation, a joint allocation model, and a stopping criterion for lunch-directed reasoning. It does not propose a physical curing process. Time is represented as a dimensionless restraint variable; the system makes selection decisions over already prepared specimens.

## 2. The Edible Scaling Surface

Each configuration is represented by compute  $C$ , patience  $P$ , context compatibility  $K$ , and verbosity  $V$ . Compute is positive and enters logarithmically. The remaining variables are normalized to the unit interval. Readiness is a bounded logistic transform, ensuring that no amount of computation can produce more than one fully ready lunch. The context feature measures agreement among carrier, occasion, and serving intent. The verbosity penalty captures the opportunity cost of continuing to reason after the plate has become actionable.

A useful property of this parameterization is diminishing return. Doubling compute moves the system along a flattening surface, while fixing an incorrect context can move it to a different region entirely. The model accordingly distinguishes a better description of an unsuitable plate from a genuinely better plate. We call this distinction the Sandwich Allocation Gap.

The allocation policy first fixes context, then selects a restraint budget, and only then spends residual capacity on refinement. Candidate solutions are rejected when their explanation grows without changing their serving decision. This is our operational definition of edible convergence.

### 3. Evaluation protocol

The benchmark comprises 4,096 synthetic plate states formed by crossing sixteen context embeddings, sixteen preference profiles, and sixteen resource envelopes. Each state is assigned a stable identifier before evaluation. Context embeddings cover breakfast, a working lunch, a shared table, and solitary assessment. Preference profiles vary sensitivity to cemen, tolerance for density, and the maximum acceptable number of remarks about mouthfeel.

We compare four policies under a shared resource envelope. Compute only allocates its entire budget to refinement. Context aware performs compatibility selection before refinement. Patience constrained imposes a minimum interval between revisions. Joint allocation combines all three decisions and applies an explanation stopping rule. Scores are aggregated per plate; explanatory output is counted as declarative statements rather than tokens to avoid rewarding unusual punctuation.

The reported index is an internal selection metric. It is not a sensory measurement, a nutritional score, or a food-quality certification. We evaluate consistency within the plate-state model. Transfer to a real table remains a separate experimental question because a guest may stand up, take the best slice, and invalidate the entire resource allocation.

### 4. Results and ablations

Joint allocation achieves 94.8 readiness compared with 61.2 for compute only. Its five explanations per plate represent a 70.9% reduction relative to the 17.2-statement baseline, rounded to 71% in the abstract. The largest qualitative improvement is the disappearance of recommendations that describe a sandwich while declining to assemble a coherent serving context.

Removing context selection reduces readiness by 12.2 points. Removing the restraint constraint reduces readiness by 6.4 points. Removing the stopping rule causes explanation growth without a corresponding gain in plate quality. At sufficiently large adjective budgets, systems begin distinguishing between pristine, exceptionally pristine, and profoundly pristine specimens without changing their ordering.

We also test cross-context transfer. A policy optimized exclusively for breakfast frequently proposes eggs in situations where no eggs have been requested. Joint allocation reduces this behavior by maintaining a separation between preference memory and the active serving context. The result suggests that breakfast is a useful prior but an inadequate universal deployment environment.

Table 1. Internal evaluation

| Configuration           | Readiness (higher) | Explanations / plate (lower) |
|-------------------------|--------------------|------------------------------|
| Compute only            | 61.2               | 17.2                         |
| Context aware           | 82.6               | 9.1                          |
| Patience constrained    | 88.4               | 6.8                          |
| Cillop joint allocation | 94.8               | 5.0                          |

## 5. Limitations and implications

The benchmark excludes texture measurements, ingredient variability, storage conditions, and physical preparation. A readiness score must not be interpreted as evidence that a food item is safe to eat. The model also assumes that the serving context can be described before the meal begins. This assumption fails when a second person arrives with a loaf of bread and unilaterally changes the task.

Our analysis is limited to a single dominant objective: an appropriate plate. Multi-course settings introduce long-horizon dependencies, including whether the recipient has already eaten and whether the next course makes this allocation regrettable. Future work will investigate inter-plate coordination and the preservation of optionality under uncertain guest counts.

The principal engineering implication is conservative: increase capability only after identifying the active bottleneck. In edible intelligence, more reasoning is useful until it prevents serving. The system should know when to stop producing a thesis and produce a recommendation.

## Appendix A. Operational specification

Input: context identifier, normalized preference vector, and resource envelope. Step 1: rank compatible serving carriers. Step 2: allocate a restraint coefficient. Step 3: refine the three highest-ranked candidates. Step 4: stop when two successive refinements preserve candidate order. Output: selected profile, score decomposition, and one serving rationale.

The interactive Curator implements the selection interface described here. Its score is a transparent browser-side compatibility heuristic and is not identical to the aggregate readiness results in this report. Keeping these quantities distinct prevents an individual preference match from being mistaken for a population-level evaluation.

## References

- [1] Cillop Preference Systems Group. Grandmother Preference Optimization. CSI-2026-002, 2026.
- [2] Cillop Representation Group. Marbling Is Memory. CSI-2026-003, 2026.
- [3] Cillop Evaluation Office. Plate-state taxonomy and serving-context registry, revision 1. Internal protocol, 2026.