LifeCorp Research

Resident-Specific Forecasting: Selected Technical Sections

Source classification: Internal — Restricted Distribution. Approved technical extract for the service-planning record. Only the sections reproduced here are released; this copy supplies no access to the source file or resident representations.

2.1 — Representation and research lineage

The internal programme reference is MIRROR. Research began in November 2024 with a shared-model prototype, structured behavioral state and persistent resident-specific representations. A representation combines permitted episodic observations, context, preference history and a bounded adapter where evidence supports one. It is a versioned probabilistic forecast object, not a reconstruction of the resident or an agent authorized to speak for them.

Following the September 22, 2026 partnership, programme responsibility sits with Lattice Research & Evaluation under Intelligence Governance. The routine transport configuration retains its validated frozen backbone. Hosting that configuration within the restricted stack is distinct from migrating it to a new partner foundation model. Any backbone change requires new task validation; earlier evidence does not automatically transfer.

2.3 — Prediction ledger

Before an observation window opens, the ledger preserves the resident-representation version, model version, timestamp, permitted context snapshot, predicted outcome distribution, uncertainty, expected observation window and planned institutional interventions. Outcomes are appended after the window closes. Original probabilities remain immutable. Unobserved outcomes are marked missing rather than inferred from subsequent service uptake. No post-outcome context may enter an earlier prediction.

Permitted event streams include scoped service requests, transport confirmations, scheduling choices and consented participation records. Access is task-specific. Private memorial content, generated memorial exchanges, personal correspondence and private personal-archive uploads are excluded. Corrections and supersession notices invalidate affected cached features. An inference cannot be used to reconstruct an excluded source indirectly.

3.1 — Prospective routine-task evaluation

The evaluation covers 180 opted-in adult residents with recurring transport bookings. Each contributes one prediction per week of whether the requested journey will be confirmed within seven days. Twelve prediction dates from March 2 through May 18, 2026 produce 2,160 resolved predictions; a later held-out window from June 1 through August 17 produces a further 2,160, with outcomes complete by August 23. All expected outcomes were recorded. This is a repeated-measure task sample, not a trial of 4,320 separate residents.

Prototype fitting ended before March 2. Predictions preceded outcome access in both periods; neither period included parameter retuning. The first period informed task review but the later window remained untouched until final scoring. In the later window, the frozen resident-specific configuration had a mean Brier score of 0.16 versus 0.19 for the pre-registered service-area baseline; lower is better. These rounded descriptive scores concern this binary endpoint only. Uncertainty analysis must cluster by resident and account for site conditions; the extract makes no significance claim.

Scheduling practice and service availability remained part of the observed environment. The system was sufficiently useful to prioritize manual checking of booking exceptions, without deciding allocations. These scores establish neither understanding of a person nor accuracy after a housing move, bereavement or a new employment arrangement. Those transitions remain outside the validated task.

3.4 — Error-triggered optimization

Observation updates may change structured state without retuning parameters. An isolated surprise contributes to the error record but does not initiate optimization. Domain review examines rolling calibration, proper scoring loss, uncertainty coverage, intervention-conditioned residuals and systematic changes against pre-registered gates. Threshold values belong to the task register; a single universal resident accuracy threshold is inappropriate.

When accumulated evidence crosses a gate, the process preserves the current representation and prediction history, identifies eligible observations and distinguishes changed circumstances from persistent error. It first corrects stale state or provenance. Adapter retuning is considered only where simpler changes fail. The candidate is evaluated against a reserved later time block whose outcomes were unavailable to fitting; unresolved prediction windows remain excluded. Promotion requires baseline and subgroup checks, approved uncertainty behavior and no boundary violations. Failed candidates remain unpromoted. The deployed version, rollback target and reasons are retained.

4.2 — Scenario and cohort use

A hypothetical change may be applied to an individual representation to estimate response distributions under specified assumptions. Examples include service timing, appointment location, introduction to an optional activity and alternative communication formats. A simulated trajectory is not an observation, a preference declaration or an authorization. Unfamiliar conditions should widen uncertainty or cause abstention.

Synthetic cohorts assemble individual representations for planning hypotheses. Correlated errors, relationship effects and institution-wide responses can invalidate a sum of individual forecasts. Population use requires separate validation. The routine confirmation study authorizes no prediction about collective relocation, consultation votes or an individual decision to leave the ecosystem.

5.1 — Intervention feedback and research disagreement

LifeCorp changes the environment in which the forecast is scored: transport slots, service options and notice timing affect the eventual behavior. Intervention labels must distinguish what was available before prediction from what the institution later changed. A model can improve its score because operations became more consistent. That is useful for capacity planning but differs from discovering a stable personal preference.

Operations considers reduced booking uncertainty valuable even where standardized schedules explain part of the gain. Research agrees on that use while requiring separate evidence for a welfare or understanding claim. Intelligence Governance requires that assessment remain task-specific. Predictive accuracy is an operational metric, not a measure of autonomy, satisfaction or informed consent.

5.3 — Communication convergence hypothesis

The 2025 Human Compatibility Framework provides a research hypothesis about observed communication variance. Greater familiarity with approved formats, avoidance of extra review and reduced pursuit of difficult complaints could lower prediction error without a better model of the resident. Better modeling and behavioral adaptation may coexist. The hypothesis does not authorize a new resident-data feed or a new task deployment.

Any proposed evaluation must preserve the original complaint population and record missing contacts, abandonment, readiness completion, service remedies and institutional interventions separately. Predict-before-observe ledgers and error-triggered optimization remain unchanged. Missing interaction is not successful prediction. Prospective error comparisons must separate model changes from changes in who is observed or which expressions reach the channel.

An internal representation cannot directly establish private thought, unexpressed dissatisfaction or subjective acceptance. The existing routine transport evidence supplies no validation for those claims. Aggregate service-consistency measures may motivate further research; they do not resolve whether apparent improvement reflects understanding or conformity to institution-shaped conditions.

Publication context

https://lifecorp.org/archive/service-forecast-planning