PREDICT — Model Card v2.4.1

Surgical-risk & cost-trajectory engine for self-funded benefits programs

Maturity disclosure. *PREDICT v0.1 is a working pipeline trained on simulated cohorts and is not approved for production use. The metrics below are v1.0 target spec, validated against the simulated test set. First real-world readout is targeted Q3 2026. See predict-methodology.html for the full Model Card.*

1. Intended Use

PREDICT ranks members of a self-funded health plan by probability of incurring a high-cost surgical episode within a 6–18 month forward window, and estimates the expected medical-trend impact if that episode is routed through JetPatient's bundled-care network. Outputs feed Guardian's outreach and network-design workflows. The model does not make individual treatment decisions; clinical decisions are always made by the member and their care team.

2. Model Details

Architecture	Two-stage ensemble — temporal sequence encoder (claims + Rx + procedure codes) feeding
Training cohort	Simulated population of 1.2M member-years across 9 procedure families, derived from publicl
Look-back window	36–60 months of claims history (de-identified, tokenized)
Forward window	6–18 months (procedure-family dependent)
Features	240 base features → 1,420 derived signals (utilization velocity, comorbidity load, Rx adherenc
Versioning	Model registry with rollback; current production candidate v2.4.1
Retraining cadence	Quarterly with continuous drift monitoring (see Drift Report)

3. Performance — Simulated Cohort

Sub-cohort	Members	AUC	Precision @ top-10%	Recall @ top-10%	Calibration (Bri
Orthopedic · all ages	4,820	0.91	91%	78%	0.083
Bariatric · all ages	2,140	0.88	88%	79%	0.091
Spine · all ages	1,640	0.86	86%	75%	0.094
Cardiac · all ages	980	0.85	85%	73%	0.098
Orthopedic · age 70+	720	0.87	87%	74%	0.096
Bariatric · lower-income	510	0.85	85%	72%	0.099

4. Fairness Criteria

Disparate-impact ratios are computed quarterly across age band, sex, geographic region, plan tier, and a lower-income proxy. Any group whose true-positive rate deviates by > 5 percentage points from the cohort mean triggers a Watch status; deviation > 10 pp triggers automatic model freeze pending mitigation review. The current Watch entries are Orthopedic age 70+ and Bariatric lower-income (see Fairness Audit Log).

5. Limitations

PREDICT is not a diagnostic tool, does not replace clinical judgment, and is not approved for any use outside the bundled-care navigation workflow described in §1. The model is trained on simulated cohorts; first real-world performance read-out is targeted for Q3 2026. The model does not infer protected-class attributes from clinical features and is audited for disparate impact each quarter.

6. Monitoring & Rollback

Population Stability Index (PSI) is tracked daily on each top-25 feature; a feature with PSI > 0.25 is flagged, > 0.4 triggers promotion-freeze. The model registry preserves all production candidates and rollback is one operator action. Inference latency, fairness deltas, and outcome calibration are reported in the rolling 90-day Drift Report.

7. Data & Privacy

Member identifiers are tokenized at the TPA boundary before reaching PREDICT. No PII or PHI leaves the client data environment in raw form. All data in transit is encrypted with TLS 1.3; data at rest uses AES-256. JetPatient maintains a HIPAA Business Associate Agreement with every client and TPA partner. The model is operated under a documented retention and deletion policy (see Compliance Attestation).

8. Contact & Disclosure

For escalations, model-card questions, or sub-cohort fairness reviews: **trust@jetpatient.com**. Adverse-event reporting: **safety@jetpatient.com**. Full methodology: **predict-methodology.html**.