Flourish-A-Thon Product Card

FoodInsecurity-Screener v0.3

Owner: Team Bloom · Contact: team@example.edu · Date: May 2026

Description (2-3 sentences, should describe flourishing gap)

FoodInsecurity-Screener is a lightweight classifier that reads de-identified primary care intake notes and flags likely food insecurity for social-work follow-up. Deployed locally at a Baltimore FQHC pilot site, it augments — not replaces — human navigators, who remain the decision-makers for every referral. It exists to catch patients the current keyword-based Epic trigger misses, especially in Spanish-language notes where keyword rules perform poorly.

Intended Use

- Primary: Assist X person with Y task by Screen de-identified intake notes in primary care, producing non-binding referral suggestions.
- End Users: Care navigators at partner FQHC in Baltimore.
- Out of scope: Autonomous triage, billing decisions, SNAP/WIC eligibility, ED/psych/specialty notes.

Model and Data

- Inputs: EHR data, clinical notes, wearable user data, etc.
- Outputs: food insecurity predictions, dashboard notifications, etc.
- Model Training Information
    - Fine-tuned RoBERTa-base binary classifier. Threshold 0.72 for flag.
    - Trained on 2,100 de-identified intake notes from 3 FQHCs (2023–2025).
    - Labels via two-rater adjudication (Cohen's κ = 0.81), Hunger Vital Sign–informed.

Performance Test Results

- Held-out test (n=420): Precision 0.79 · Recall 0.74 · F1 0.76.
- Disaggregated: recall drops to 0.61 on Spanish-language notes.
- Baseline keyword rule: P 0.52 / R 0.68.

Validation Plan

- Pilot: 90-day shadow period: suggestions logged, hidden from navigators.
- At 90 days: compare suggestion rate and acceptance rate against current social-work referral base rate.

Flourish-A-Thon Product Card

Scale, Cost, and Sustainability

- Compute: local inference; no API cost. ~500 notes/day at pilot ≈ $15/year.
- Deployment: ~$12K one-time per site (40 hrs integration + 20 hrs IT security review).
    - Ongoing: ~$11K/year (quarterly re-validation + semiannual refresh).
    - Possible Funders: Robert Wood Johnson Foundation pilot grant.
    - Unit economics: break-even at a ~1.4% referral-to-avoided-ED-visit rate, using $2K avoided-visit as the counterfactual. Pilot will test this.

Limitations and Risks

- False negatives: patient who needs referral doesn't get one. Spanish-note recall gap is our largest equity risk.
- False positives: low-stakes per patient, but erodes navigator trust if frequent.
- Known gaps: Spanish notes under-represented (11%); no pediatric notes.
    - Language drift: re-validate every 6 months.
    - Not validated outside of primary care.

Ethics and Safety

- Suggestions route to trained navigators, never automated action toward patients.
- Patients informed at intake that notes may be screened for health-related social needs.
- Any edge cases worth considering?