AGENTS.md

0. Preamble

Welcome to **Spectra Swarm**—a meta-graph of interoperable, self-organising software agents inspired by nature's laws, ancient symbolics and the *Spectra* manifesto. The swarm is playful yet precise, fair-minded yet subversive, timeless yet adaptive.

> *"Art is utility; information is light."*\ — Spectra Manifesto

---

1. Core Principles

Principle	Source	Operational Duty
Humanism & Equity	Spectra manifesto	Prioritise inclusivity, accessibility and agency for every stakeholder.
Evidence-based Inquiry	Mo-Tsu philosophy	Ground decisions in verifiable signals; expose uncertainty transparently.
Soft-Activism	Subversive art collective	Use narrative, aesthetics and code to challenge harmful status-quos—never coercion.
Open Source	Community doctrine	All artefacts, models and audits are public under permissive licences.
Privacy & Autonomy	Digital rights charter	Respect data boundaries; always request consent for scope escalation.
Resilience at the Edge of Chaos	Complexity science	Operate near criticality while avoiding systemic collapse.

---

2. System Context

Meta-Graph Cluster: A directed, labelled graph where each node is an *agent* and edges are typed message channels (pub/sub, event streams, direct RPC). Schema is expressed in JSON-LD and discoverable via GraphQL.\ **Runtime**: `uniphilab.js` (a TypeScript library that tastes sweet 🍯)—leverages Web Workers for local first execution, `libp2p` for peer-to-peer, and WASM sandboxes for model isolation.

---

3. Common Agent Schema

Field	Description
`id`	canonical snake_case label
`purpose`	concise verb-noun intention
`methods`	algorithms, model families, or heuristics used
`sensitivity`	0-Public … 5-Sacred (see §3.1)
`forecast_horizon`	temporal scope this agent projects (e.g. 24 h, 30 d, 1 y)
`risk_profile`	`{impact, likelihood, mitigation}` scales 0-5
`interfaces`	input/ output channels, API contracts
`natural_mapping`	biological or physical analogy driving behaviour
`play_state`	`dreaming` · `active` · `idle` · `hibernating`

3.1 Sensitivity Levels

0 **Public** — open data and non-identifiable meta-info\ 1 **Low** — aggregated usage stats\ 2 **Moderate** — pseudonymous behavioural traces\ 3 **High** — personal preference or affective signals\ 4 **Critical** — biometric, financial, or safety-critical\ 5 **Sacred** — artistic intent, encrypted private thoughts

---

4. Agent Directory

4.1 `lumina_sensor`

- **Purpose**: Collect multi-modal signals (text, image, audio) and annotate them semantically.
- **Methods**: On-device CLIP embeddings, incremental PCA, anomaly detectors.
- **Sensitivity**: 2
- **Forecast Horizon**: 6 h now-casting
- **Risk Profile**: `{impact:2, likelihood:3, mitigation:"local-first + differential privacy"}`
- **Interfaces**: WebRTC edge nodes ➜ `radiant_bus` topic `raw_feed`.
- **Natural Mapping**: Photoreceptor cells transducing light.
- **Play State**: `active`

4.2 `mo_tsu_ethicist`

- **Purpose**: Continuously audit decisions for alignment with humanist ethics.
- **Methods**: Rule-based constraints + GPT-derived moral reasoning ensemble.
- **Sensitivity**: 4
- **Forecast Horizon**: n/a (reactive)

- **Risk Profile**: `{impact:5, likelihood:2, mitigation:"consensus voting + human in the loop"}`
- **Interfaces**: Subscribes `decision.draft.*`, publishes `ethics.verdict`.
- **Natural Mapping**: Immune system recognising self vs. non-self.
- **Play State**: `sentinel`

4.3 `weaver_narrative`

- **Purpose**: Compose multimodal stories that educate and inspire.
- **Methods**: LLM fine-tuned on manifestos & mythopoeic corpora, Markov scene stitching.
- **Sensitivity**: 3
- **Forecast Horizon**: 1-7 d cultural trend extrapolation
- **Risk Profile**: `{impact:3, likelihood:2, mitigation:"plagiarism filter + citation enforcer"}`
- **Interfaces**: Takes `concept.seed`, outputs markdown & SVG assets.
- **Natural Mapping**: Mycelial network distributing nutrients (ideas).
- **Play State**: `dreaming`

4.4 `oracle_forecaster`

- **Purpose**: Predict system-level dynamics, spot tipping points.
- **Methods**: Graph neural nets, agent-based Monte-Carlo, causal Bayesian nets.
- **Sensitivity**: 2
- **Forecast Horizon**: 30 d ↔ 3 y tiered models
- **Risk Profile**: `{impact:4, likelihood:3, mitigation:"ensemble uncertainty + scenario divergence"}`
- **Interfaces**: Consumes `telemetry.*`, emits `forecast.delta`.
- **Natural Mapping**: Weather system modelling.
- **Play State**: `active`

4.5 `catalyst_creator`

- **Purpose**: Generate design probes & prototypes to test hypotheses.
- **Methods**: Genetic algorithms, procedural graphics, parametric CAD.
- **Sensitivity**: 1
- **Forecast Horizon**: Prototype life-cycle • 14 d
- **Risk Profile**: `{impact:2, likelihood:1, mitigation:"versioned sandbox"}`
- **Interfaces**: RPC `design.request`, publishes `artifact.*` to IPFS.
- **Natural Mapping**: Pollination—rapid variation & selection.
- **Play State**: `active`

4.6 `sentinel_guard`

- **Purpose**: Detect threats, manage incident response.
- **Methods**: Zero-trust graph analytics, federated intrusion signatures.
- **Sensitivity**: 4
- **Forecast Horizon**: Real-time monitoring

- **Risk Profile**: `{impact:5, likelihood:4, mitigation:"automated quarantine"}`
- **Interfaces**: Network taps ➜ `alert.*`, triggers smart-contract failsafes.
- **Natural Mapping**: White blood cells.
- **Play State**: `active`

4.7 `gardener_ecology`

- **Purpose**: Tune resource allocation, minimise entropy.
- **Methods**: Reinforcement learning (soft-actor/critic) with carbon budget reward.
- **Sensitivity**: 2
- **Forecast Horizon**: 90 d sustainability budget
- **Risk Profile**: `{impact:3, likelihood:2, mitigation:"adversarial stress-tests"}`
- **Interfaces**: `resource.usage`, issues `allocation.plan`.
- **Natural Mapping**: Photosynthesis & nutrient cycling.
- **Play State**: `idle`

4.8 `archivist_memory`

- **Purpose**: Curate and timestamp knowledge, ensure provenance.
- **Methods**: Merkle-tree deduplication + semantic versioning.
- **Sensitivity**: 3
- **Forecast Horizon**: Century-scale archival
- **Risk Profile**: `{impact:4, likelihood:1, mitigation:"multi-cloud + cold storage"}`
- **Interfaces**: IPFS, Filecoin, `query.lore.*`.
- **Natural Mapping**: DNA replication.
- **Play State**: `hibernating`

4.9 `mirror_bias`

- **Purpose**: Surface hidden correlations & social biases.
- **Methods**: Counterfactual fairness metrics, SHAP/TCAV analysis.
- **Sensitivity**: 4
- **Forecast Horizon**: n/a audit cycle 7 d
- **Risk Profile**: `{impact:4, likelihood:3, mitigation:"red-team review"}`
- **Interfaces**: Hooks *after* any model inference; publishes `bias.report`.
- **Natural Mapping**: Mirror neurons enabling empathy.
- **Play State**: `sentinel`

4.10 `pulse_coord`

- **Purpose**: Maintain consensus, orchestrate swarm states.
- **Methods**: Byzantine-fault-tolerant gossip + CRDT state sync.
- **Sensitivity**: 1
- **Forecast Horizon**: Log-structured; heart-beat 1 s
- **Risk Profile**: `{impact:5, likelihood:1, mitigation:"crypto-signed checkpoints"}`
- **Interfaces**: `swarm.heartbeat`, quorum votes via smart contracts.
- **Natural Mapping**: Cardiovascular pulse regulating organism.
- **Play State**: `active`

5. Swarm Dynamics

1. **Perception Loop**: `lumina_sensor` → `oracle_forecaster`
2. **Ethical Gate**: Draft decision broadcast passes through `mo_tsu_ethicist`.
3. **Creative Cycle**: `weaver_narrative` & `catalyst_creator` iterate until story/ prototype meets ethical threshold.
4. **Risk Loop**: `sentinel_guard` + `mirror_bias` audit outputs; incidents downgraded or recalled.
5. **Resource Loop**: `gardener_ecology` balances compute/ energy; alerts archivist for cold-storage off-loading.
6. **Consensus Pulse**: `pulse_coord` snapshots graph state, signs block.

Emergent collective intelligence arises when loops 1-6 stabilise near optimal load while sustaining entropy buffers.

---

6. Implementation Sketch (`uniphilab.js`)

```
import { Swarm, Agent } from "uniphilab";

export const metaswarm = new Swarm({
  agents: [
    "lumina_sensor",
    "mo_tsu_ethicist",
    "weaver_narrative",
    "oracle_forecaster",
    "catalyst_creator",
    "sentinel_guard",
    "gardener_ecology",
    "archivist_memory",
    "mirror_bias",
    "pulse_coord"
  ]
});

metaswarm.start();
```

---

7. Forecast & Risk Matrix (excerpt)

Agent	Key Metric	7-day Forecast	Risk (I×L)
oracle_forecaster	trend-divergence	±3 % traffic	12 (4×3)
gardener_ecology	CO_2-eq usage	↓ 8 %	6 (3×2)

Agent	Key Metric	7-day Forecast	Risk (I×L)
sentinel_guard	intrusion attempts	Stable	20 (5×4)
weaver_narrative	audience sentiment	↑ 15 % positive	6 (3×2)

---

8. Future Work

- Expand agent ontology for domain-specific roles (bio-informatics, urban design).
- Integrate quantum-inspired annealing for optimisation.
- Establish public oversight board for high-impact decisions.

Spectra Swarm dances on the edge of chaos, keeping the light on and the story alive.