IQANOVA FOUNDATION MODELS FOR MIDD

Models you can audit and trust.

Public-knowledge and proprietary foundation models for model-informed drug development: executable models, graded evidence and a configured LLM layer, delivered as one versioned package.

PK/PD | Pop PK/PD | PBPK/PBBM | QSP | Disease progression

Offered with Pharmetheus, which provides expert review of model outcomes.

One foundation model, four layers.

An IQANOVA Foundation Model is the LLM layer plus all files, optionally delivered on an AI PC that runs it locally.

Models	Executable models (native, SBML, COMBINE OMEX), simulation code and a pinned environment. The EHMN2027 scaffold is supplied separately.
Files	Nineteen folders: evidence ledgers, calibration, verification and validation, analyses, reports, regulatory cross-references and risk registers.
LLM layer	IQANOVA prompt set, exact run history and checksums, a retrieval index over all files, and 22 critic checks before any claim.
Runtime	An AI PC with a local open-weight LLM and the IQANOVA runner, configured and tested before it ships.

Start from public knowledge.
Add your own data when ready.

Tier 1 · Public Knowledge

1–2 MONTHS

- Built only from public literature, labels, regulatory reviews, registries and guidance
- You supply the assets, the context of use and your questions
- No confidential data leaves you, because none is needed
- Public-data work can use a frontier LLM

Tier 2 · Proprietary

2–3 MONTHS

- Adds your confidential drug details, PK/PD results, in vivo and preclinical data
- Detailed questions on your own context of use
- Runs on a local open-weight LLM; no confidential content goes to external APIs
- NDA and data-handling terms first

Durations run from the agreed brief. Pricing is to be discussed and depends on stage, scope and complexity.

Both tiers ship as a versioned package with every file labelled by evidence status.

Everything a reviewer needs to find, rerun and check.

19

Folders, from scope and evidence to regulatory cross-references

22

Critic checks before any claim

4

Status axes on every file: content, evidence, applicability, decision

D0–D10

Diligence stages from target to post-market

REAL

Trial arms, literature values, digitised figures and your supplied data, with provenance

SEMI-SYNTHETIC

Observed inputs plus model-generated parts; the origin of every variable is identified

SYNTHETIC

Virtual patients or systems, with seeds and quality control

Every dataset is labelled by origin. Literature is graded A1 to E, from regulator-confirmed to in-vitro or assumption.

Ask the model the questions that matter.

TIER	EXAMPLE CONTEXT OF USE	QUESTION YOU ASK THE MODEL
1 · Public	First-in-human starting dose for a class of T-cell engagers	Which MABEL-based dose range follows from published data, and which parameters move it most?
1 · Public	Exposure of an approved drug in a special population	Is the labelled dose adjustment consistent with a mechanistic prediction?
1 · Public	Phase 2 design scenarios	How do duration, interval and sample size change power, given public placebo variability?
2 · Proprietary	First-in-human dose for your biologic	How far apart are MABEL-based and NOAEL-based doses, and which extra experiment cuts uncertainty most?
2 · Proprietary	Translation from in vivo preclinical data	What human exposure is predicted to be efficacious, and under which alternative mechanisms does that change?
2 · Proprietary	Review of your existing model	Do rate laws and units check out, and which claims exceed the evidence?

Each context of use is written to a template that can be given to an LLM such as Claude, or to the local model for Tier 2.

Confidential data stays where you put it.

OPTION	PUBLIC DATA	CONFIDENTIAL DATA	FITS
Cloud-assisted	Frontier LLM	None	Tier 1
Hybrid	Frontier LLM	Local model only	Tier 2 with a public evidence build
Fully local	Local model	Local model	Tier 2, offline or air-gapped

Local open-weight models are less capable than frontier models. IQANOVA's critic checks, replay audit and frozen-prediction rule are what keep local runs trustworthy, and local output is benchmarked against frontier output before it is relied on.

Frozen at a date. Updated when the evidence moves.

Every release carries an evidence cut-off date. The prior version stays frozen, and each new release ships with a change log, repeat checks and a new manifest.

Trial success or failure, same or similar target | New class of therapy | Safety finding | Target association | Published patent | New experimental data or drug

1 · TRIGGER	2 · IMPACT REVIEW	3 · NEW VERSION	4 · RE-CHECK
A qualifying event is found	Each source classed new, updated, contradictory, withdrawn or unchanged	Public and/or proprietary data; or a new model	Repeat checks, plus optional expert review, before release

A · Subscribe

- **Evidence Watch:** weekly source checks, monthly impact brief
- **Model Maintenance:** up to one minor checked release a quarter
- **Portfolio:** up to three models, one shared release a month

B · Order an update

- Event-triggered or on request
- Public data, proprietary data, or both
- Proprietary updates run on the local model only

Monitoring starts only under an activated service agreement. Prices are to be discussed.

OPTIONAL WITH PHARMETHEUS

A named expert signs what they checked.

The model is an audited input, never the source of truth. An independent reviewer records name, date and scope before a file is marked REVIEWED.

LEVEL 1

Fidelity check: rate laws, units, parameter provenance, evidence grades

LEVEL 2

Scientific review: assumptions, identifiability, calibration, and whether each answer follows from the evidence

LEVEL 3

Regulatory-readiness review: fit to the context of use and regulatory expectations

A reviewer can raise a HOLD, which stays visible in the package. Each new release is reviewed again before its status changes. Reviews are professional advice, not regulator decisions.

Review is an optional service, quoted separately.

Send your brief. We return a scoped quote.

1 · BRIEF
Assets, context of use, questions, model families, stage, markets

2 · SCOPE AND QUOTE
Outputs, data tier, licence and delivery agreed in writing

3 · BUILD THROUGH GATES
Fidelity, critic and claim checks; HOLDs stay visible

4 · RELEASE
Versioned package, handover and one technical revision

TIER	TARGET	PRICE
Tier 1 · Public	1–2 months	To be discussed
Tier 2 · Proprietary	2–3 months	To be discussed
Update or subscription	Per event, or monthly	To be discussed

REQUEST A MODEL

sales@iqanova.org
iqanova.org/ai-midd

Research use and decision support only; not medical, regulatory, legal or investment advice. The written quote defines price, scope and licence.