A COMPUTATIONAL LAYER FOR CPET

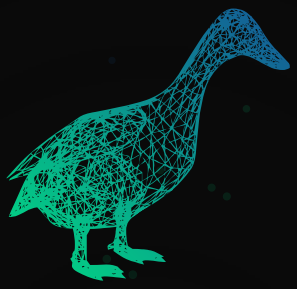

Oxynet

Physiological intelligence for exercise testing.

Structured physiological measurement from cardiopulmonary exercise test signals, consistent across protocols, populations and devices, over one API that people, clinical software and AI agents all call.

THE MANUAL

VERSION 1.8 · SEPTEMBER 2026

PREPARED BY
Andrea Zignoli, MEng, PhD
andrea.zignoli@unitn.it

ABOUT THIS DOCUMENT

What this is, and what it is for

A description of the Oxynet engine written to be read end to end by someone deciding whether to build on it, fund it, or run their laboratory's tests through it.

VERSION	1.8, September 2026
ENGINE DESCRIBED	API v1 at app.oxynet.net, analysis version 2026.08
AUTHORITATIVE SOURCE	This manual is not regenerated on every release. The API reference at **app.oxynet.net/docs** is generated from the running service and wins wherever the two differ.
STATUS	Research software. Not a medical device, and not a diagnosis. See 4.1 on page 29.
CONTACT	Andrea Zignoli, MEng, PhD, **andrea.zignoli@unitn.it**, in a private capacity. See 4.3 on page 31.

> BEFORE CIRCULATING
>
> This document describes capabilities and states their limits. It is not an offer of securities, a solicitation to invest, a clinical claim, or a regulatory submission. Nothing in it has been reviewed by a notified body or a regulator.

Contents

PART ONE

01

What Oxynet is

What it looks like, what is production today, where it sits, and the cycle behind it. Six pages, no physiology required.

The graph is the interface. This is the engine.

PART ONE · 1.1

The mission

CPET data is abundant. Physiological interpretation is still fragmented.

Oxynet is the computational layer that turns cardiopulmonary exercise test signals into structured, reproducible physiological measurements, independent of device, protocol or software.

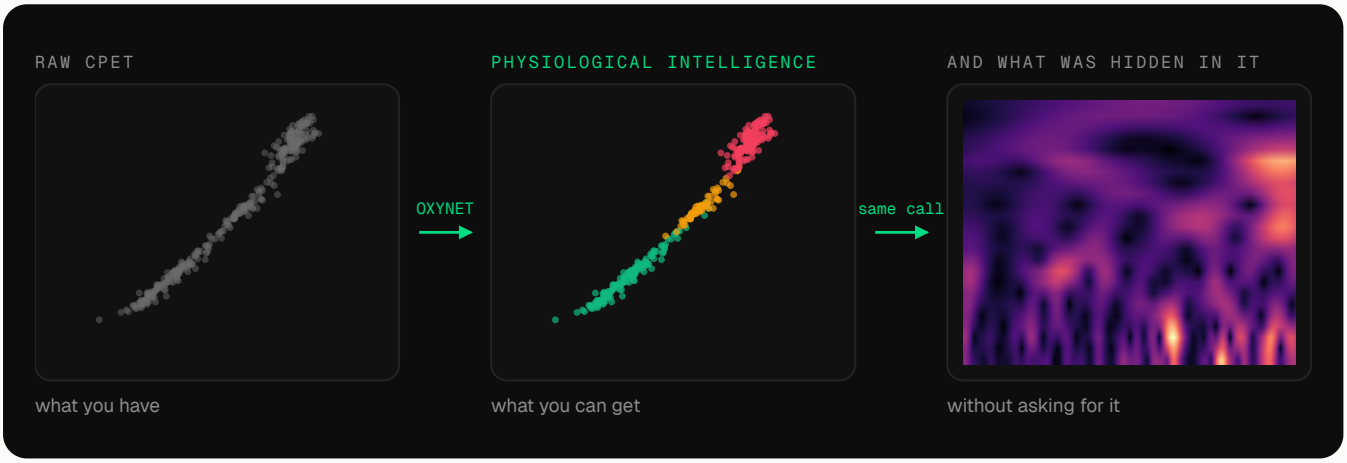

Every point is one breath of one real test. On the left, the file as the cart exported it. In the middle, the same breaths placed into intensity domains with both thresholds located, 97 milliseconds of model time later. On the right, a ventilatory rhythm the file was never asked about.

DEEP LEARNING

Convolutional networks over the breath series, with recurrent, temporal-convolutional and transformer variants benchmarked against them.

SIGNAL PROCESSING

Continuous wavelet analysis for the ventilatory rhythms, cross-channel coherence and phase for whether the trace can be believed.

LEARNED REPRESENTATION

Structure recovered without labels, and self-supervised pretraining in development. Named in full in 2.1 on page 12.

The mission is to make that content readable by any system that needs it, consistently, wherever and on whatever it was recorded.

PART ONE · 1.2

What Oxynet does today

Everything below answers on the live API now, and the note says what is behind each one.

CAPABILITY	STATUS	NOTE
VT1 and VT2 detection	PRODUCTION	Convolutional network over the breath series. Time and $\dot{V}O_2$ axes
Intensity domain per second	PRODUCTION	Three class probabilities per second, from the same network
Signal integrity checks	PRODUCTION	Cross-channel coherence and phase, sampling and clock analysis
Derived CPET quantities	PRODUCTION	Peaks, O_2 pulse, ventilatory equivalent slope profile
Substrate use and FATMAX	PRODUCTION	Indirect calorimetry from gas exchange, load channel optional
Vendor format parsing	PRODUCTION	21 formats, detected automatically
REST API	PRODUCTION	Key per caller, products gated per key
MCP endpoint	AVAILABLE	Native connector for any MCP client
OpenAPI 3.0 schema	AVAILABLE	Client generation, agent actions, integration platforms
Browser application	AVAILABLE	app.oxynet.net, for producing results without building
Python package	AVAILABLE	Open source, local inference
Oscillation analysis	RESEARCH	Continuous wavelet detector. One cohort, transportability untested
Self-supervised pretraining	RESEARCH	Pretraining on unlabelled recordings. In no promoted model
Local and embedded deployment	AVAILABLE	Discussed in 3.5 on page 26

PART ONE · 1.3

Where it sits, and what it is not

Oxynet is a computational layer between the machine that recorded the test and whatever is going to use the result. It replaces neither end.

HOW IT USUALLY WORKS

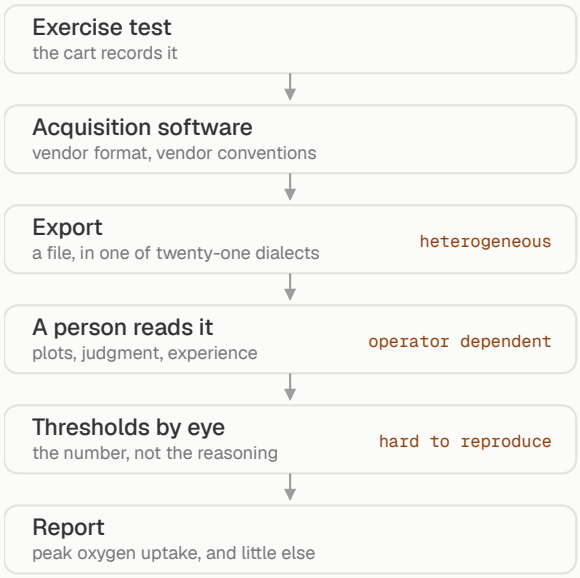

WITH OXYNET IN THE LOOP

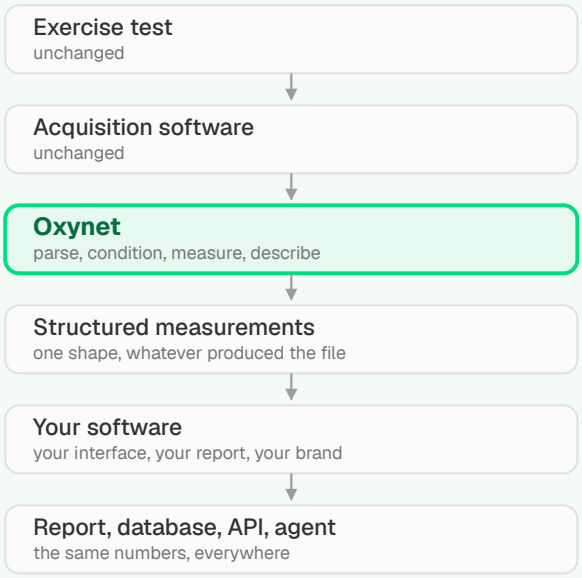

OXYNET IS NOT

An acquisition system. A metabolic cart. A replacement for the software you already ship. A patient-management system. A general-purpose assistant. A competitor for the customer relationship.

OXYNET IS

A device-independent platform for CPET interpretation, reachable over an API, deployable in your cloud or inside your product, that leaves the interface, the report and the clinical judgment where they already are.

PART ONE · 1.4

Built for the AI era

Agents and human operators, in the cloud, on a local network and on the device alike. One engine, reached whichever way the caller works.

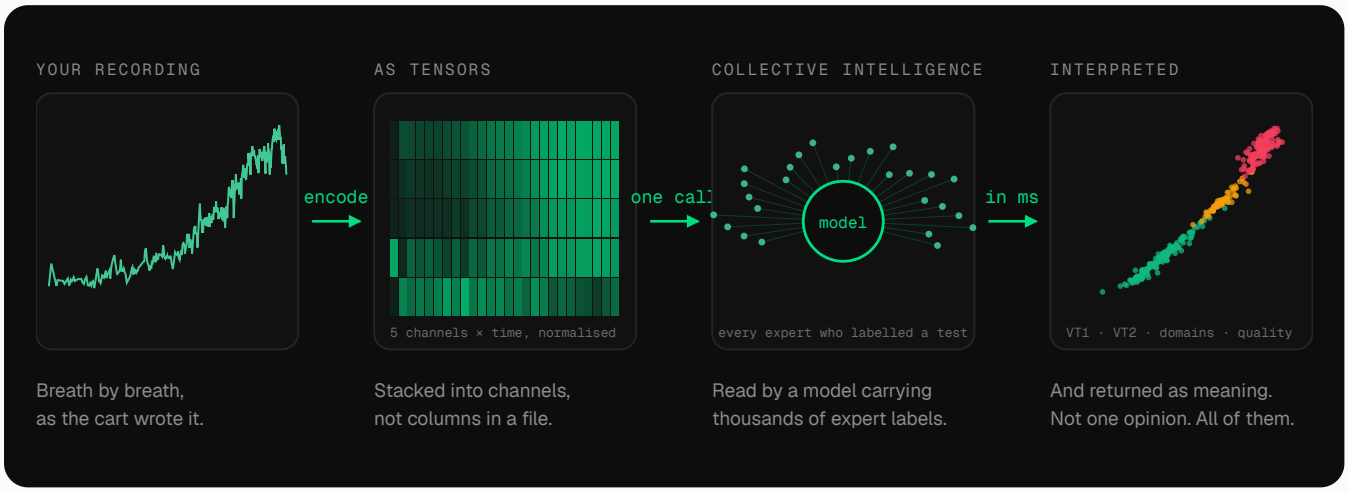

The recording is encoded as tensors, not read as columns. What comes back carries the judgment of every expert who ever labelled a test in the corpus. The result is reproducible: the same recording, on the same model version, returns the same answer next year and in another country.

HUMANS

The browser

A clinician or a researcher opens a test and reads it. No integration, no build.

SOFTWARE

REST and OpenAPI

Your product calls the engine and renders the result in your own interface.

AGENTS

MCP

An assistant calls the physiological models directly, and the signals never enter the conversation.

CLOUD	The hosted API. Nothing to deploy, models always current.
LOCAL	Inside the institution, where recordings must not leave the network.
DEVICE	Small convolutional models that export to on-device runtimes, for interpretation during the test rather than after it.

PART ONE · 1.5

Research, development, deployment

Oxynet is not a set of algorithms written once. It is a platform whose cycle runs continuously, and the cycle is the asset.

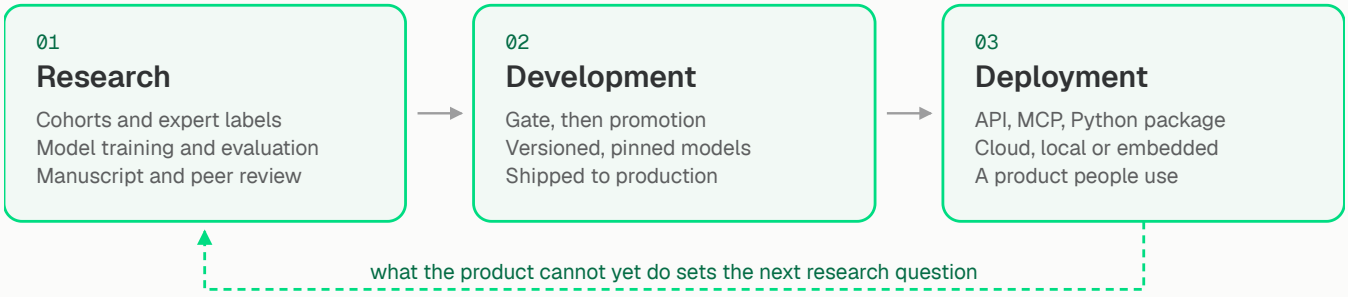

All three stages are Oxynet. The research is not a credential attached to a product; it is where the physiology comes from, and it is maintained rather than cited.

RESEARCH

Cohorts arrive from laboratories that do this work for its own sake. Experts label them. Models are trained and evaluated, and what is learned is written up and put through peer review, which is the part we do not control and the reason to trust it.

DEVELOPMENT

A model that clears the gate on a cohort it never saw is promoted, versioned and pinned, with the evaluation that let it through stored beside the weights.

DEPLOYMENT

It answers on an endpoint, in a package, or inside a partner's product. What it then cannot yet do is the next research question, and the cycle closes.

Most groups doing this work stop after the first stage, not because they cannot go further but because a paper is what the work is for. Oxynet carries the same scientific standard through to something a laboratory can call and a manufacturer can ship, and sends what it learns there back to the beginning.

7 peer-reviewed papers

on the methods and their evaluation, the first in 2019 and the most recent in 2026.

Four continents

Cohorts contributed by laboratories in Europe, North and South America and Oceania.

In production, not in a paper

Every capability marked Production on page 6 is answering on a live endpoint today.

PART ONE · 1.6

Why build on Oxynet

Nobody has to change what they already do. The layer sits underneath all of it.

FOR DEVICE MANUFACTURERS

Computational interpretation, without building a physiology stack

Add advanced interpretation to the product you already ship, with no change to acquisition hardware and no physiology team to hire. The numbers stay comparable across an installed base that spans a dozen markets.

FOR CLINICAL SOFTWARE

Raw exports become structured, reproducible outputs

Turn whatever the cart wrote into the same shape every time, in your own interface and your own report. One integration covers twenty-one vendor formats.

FOR LABORATORIES AND RESEARCHERS

Heterogeneous data, processed consistently

Pool cohorts across devices, protocols and years without pooling the readers with them. A method section a reviewer can reproduce, and a result that does not move when the rota does.

FOR AI SYSTEMS

Physiological computation an agent can call

An assistant reaches the validated models directly instead of estimating a threshold from numbers in its context. The agent orchestrates and explains. It does not do the physiology.

The four capabilities behind all of this are shown on one real recording over the next ten pages, each with the payload the API returned and the caveat that travels with it. They begin on page 12.

PART TWO

02

What the engine measures

Four capabilities on one real recording, each with the payload the API returned and the caveat that travels with it.

The graph is the interface. This is the engine.

PART TWO · 2.1

The technology

Signal processing, machine learning and deep learning, applied to a signal that has had comparatively little of it.

Conditioning, spectral analysis, phase

SIGNAL PROCESSING

Sampling-aware per-channel smoothing, resampling to a common grid, continuous wavelet analysis over the oscillation band, cross-channel coherence and phase. Deterministic arithmetic, and where signal integrity is decided.

Convolutional networks over the breath series

DEEP LEARNING

The promoted threshold models are CNNs over the detrended series, trained per cohort. Recurrent (LSTM and GRU), temporal-convolutional and transformer variants are trained and benchmarked against them; what ships is whatever clears the gate on a held-out cohort.

Structure without labels

REPRESENTATION LEARNING

An unsupervised representation learned from the recordings themselves, in which the intensity domains emerge as regions rather than lines. It is how a threshold acquires an extent instead of being an instant.

Not every channel, every time

ABLATION

Models are trained and evaluated with channels deliberately withheld, so what each one actually contributes is measured rather than assumed. A model that needs five channels when three carry the signal is a model that cannot run on half the exports in the field, and the registry records what each one requires.

In development, not deployed

SELF-SUPERVISED LEARNING

Pretraining on unlabelled recordings before fine-tuning on the labelled ones. The corpus is far larger than its labelled subset, so this is where the next gain is expected. It is in no promoted model today and nothing here rests on it.

BUILT FOR THE AGENTIC ERA

CPET interpretation has mostly lived inside desktop software. An assistant asked to read a test today will estimate a threshold from the numbers in front of it.

Through MCP it calls a model that was evaluated against expert labelling instead, and reports what came back, refusals included. See 3.4 on page 25.

NAMED, VERSIONED, PINNED

Models carry names rather than architecture strings. Same name at a new version means the same cohort, improved; a new name means a different cohort or family.

The manifest wins: the version pinned server-side is served whatever a client asks for, and the response says which one answered.

DETERMINISTIC BY DESIGN

The same recording returns the same numbers, on the same model version, indefinitely. There is no sampling, no temperature, and no run-to-run variation.

PART TWO · 2.2 · WORKED EXAMPLE

Whether the numbers can be believed

Before any physiology is reported, the engine establishes whether the recording can support it, in terms that describe the file rather than the patient.

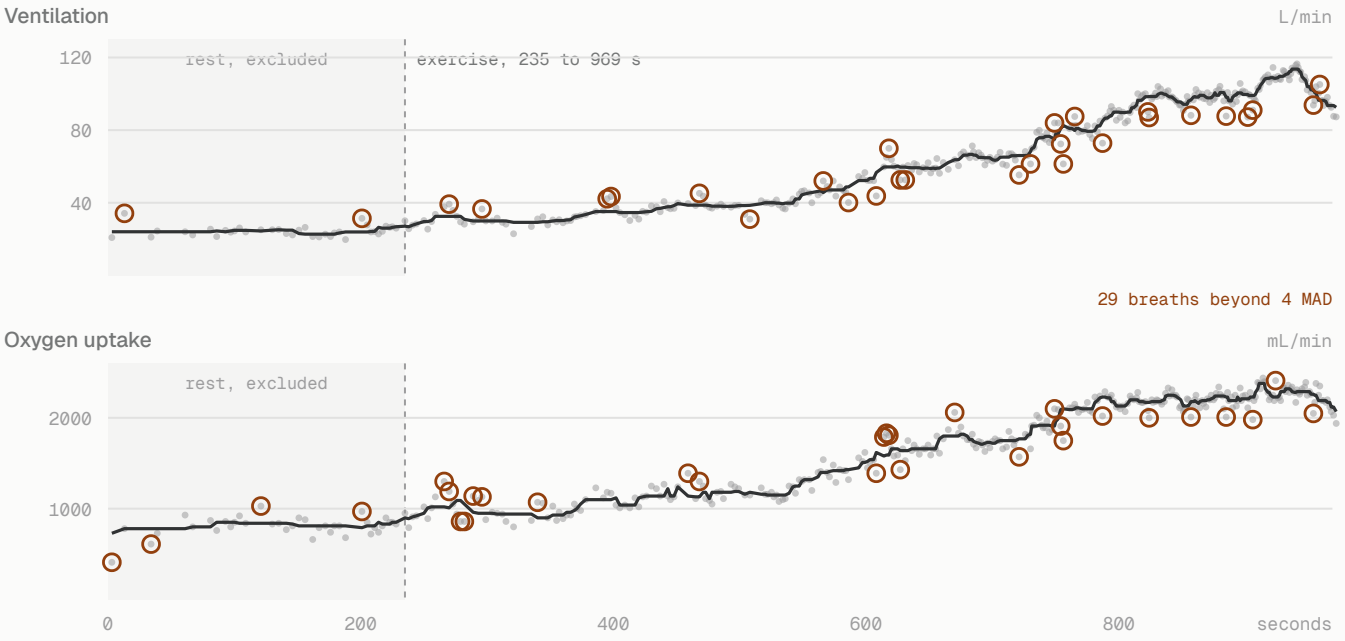

Two raw breath series exactly as the cart wrote them, with the conditioned trace running through each and the breaths a robust rule flags as outliers ringed. The shaded span at the start is the resting phase the engine excluded, 37 samples of it, before any window was taken. No intensity colouring here on purpose: domains are physiology, and nothing physiological should be claimed until the trace itself has been believed.

GAS AGREEMENT

Do the channels behave as the physiology requires? A cart with a drifting analyser fails here while every individual trace still looks plausible.

SAMPLING ADEQUACY

Which oscillation periods can this file resolve at all? Coarse sampling narrows the searchable band; it does not invalidate the recording, and the narrowing is reported.

CLOCK INTEGRITY

Many carts restart the clock at each phase transition, which makes the raw duration meaningless. A repair is always reported, never silently applied.

What the API returned

POST /v1/cpet/{id}/compute {"metrics": [...]}


```
{
  "gas_quality": {
    "verdict": "confirmed",
    "n_checks": 2,
    "n_passed": 2,
    "vco2_coherence": 0.9829,
    "vco2_phase_error_deg": 6.57
  },
  "sampling_adequacy": {
    "raw_interval_s": 2,
    "period_floor_s": 25,
    "period_ceiling_s": 140,
    "band_narrowed": false
  },
  "clock_integrity": {
    "raw_span_s": 969,
    "repair_applied": false
  }
}
```


Phase recognition, and the error it carries

A recording is not all exercise. It contains rest, a warm-up, the test and a recovery, and most measurements are meaningless taken across the wrong one. The engine annotates the phases rather than cropping to them, so a consumer can see what was excluded.

The boundaries are estimated and the estimate carries a mean error of about thirty-four seconds. That number is published rather than hidden, and it propagates: peak oxygen uptake is taken over the whole record instead of the detected exercise phase, because the peak sits at the boundary where that error bites.

These describe the recording, not the patient. A file can fail every check and come from a healthy subject. **An absent check is not a passed check:** `n_checks` reports how many were runnable.

The figure opposite is the interface. This is the engine.

PART TWO · 2.3 · WORKED EXAMPLE

Where VT1 and VT2 sit, and whether they are there at all

The entry product, and the part most often read differently by two competent people looking at the same test.

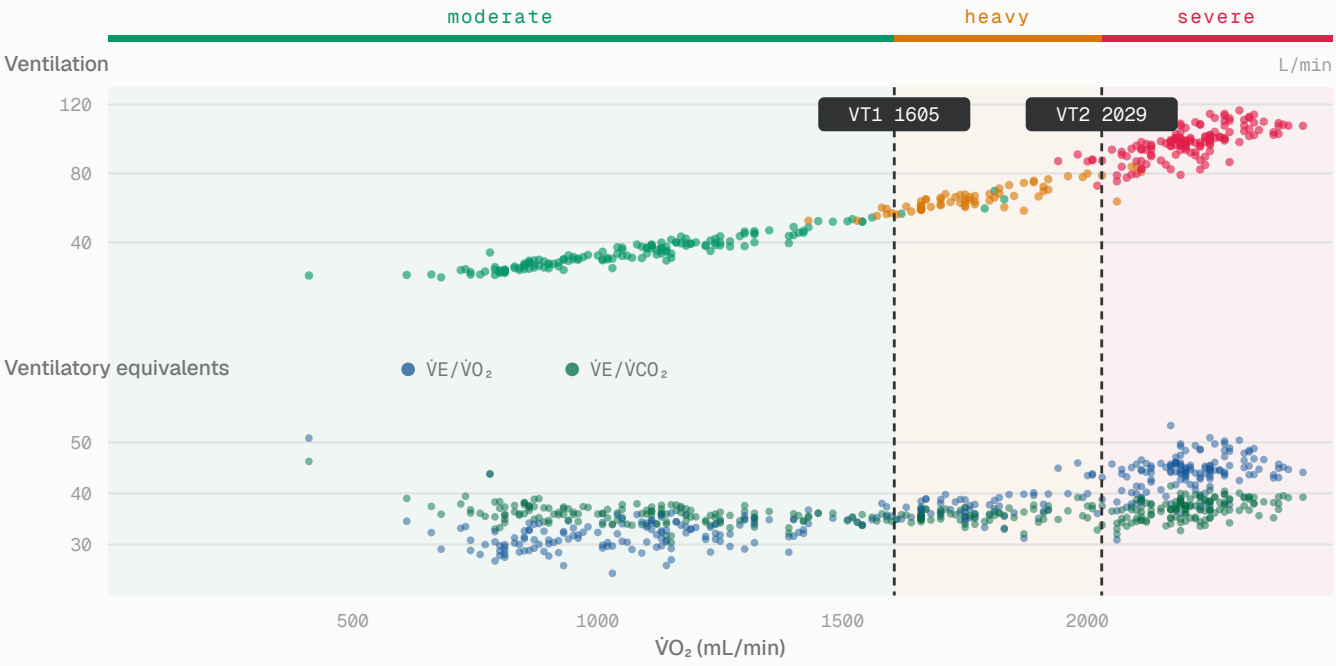

Both panels against oxygen uptake, which is the axis these are read on, breath by breath and unsmoothed. Above, ventilation. Below, the ventilatory equivalents, where the textbook criteria live: VT1 sits at the nadir of the equivalent for oxygen while the one for carbon dioxide is still flat, and VT2 where the carbon dioxide equivalent turns up as well. The equivalents come from the file's own columns and the dashed lines are what the model returned, so whether the two agree is visible rather than asserted.

What the network actually emits

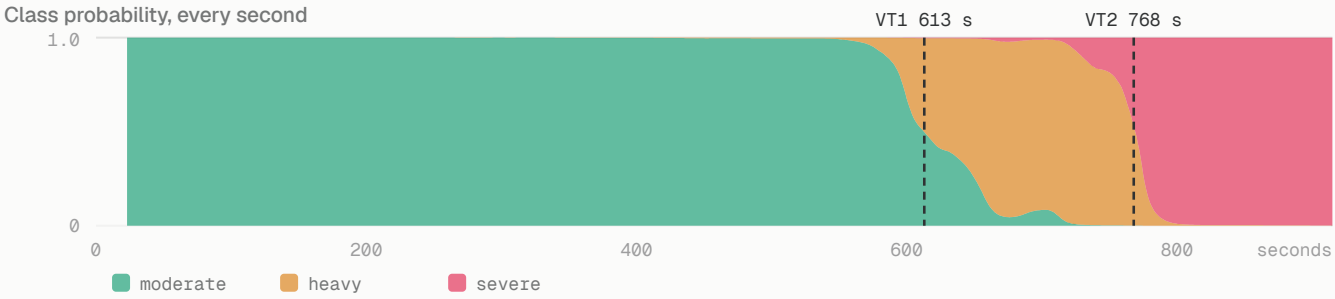

Three class probabilities every second, stacked, straight from the model. The thresholds are a detector run over these curves rather than a separate model, which is why a subject who never leaves the moderate domain simply never produces a crossing. Taken by calling the same model on the breath records directly: that path conditions the file slightly differently from the upload path used opposite, so it places the thresholds at 613 and 768 seconds against 619 and 763. Same weights, different conditioning, and both are reported rather than reconciled.

What the API returned

POST /v1/cpet/{id}/analyze {"analyses": ["vt"]}


```
{
  "analysis": "vt",
  "status": "ok",
  "findings": {
    "vt1_time_s": 619,
    "vt2_time_s": 763,
    "vt1_vo2_ml_min": 1605.2,
    "vt2_vo2_ml_min": 2028.9
  },
  "quality": "good",
  "provenance": {
    "model": "gandalf",
    "model_version": "v0.1.0",
    "analysis_version": "2026.08",
    "latency_ms": 97.41
  }
}
```


Against the label the file carried

This export arrived with thresholds already stored in it, set by the clinician who ran the test. They were not shown to the model.

	STORED LABEL	OXYNET	Δ
VT1	1,540	1,605	+65
VT2	1,940	2,029	+89

Oxygen uptake in mL/min. One test is an illustration, not an evaluation: the promotion gate is what a model is held to.

PROMOTION GATE

No model reaches the registry without clearing, on a held-out cohort:

$\dot{V}O_2$ bias ≤ 120 mL/min
$\dot{V}O_2$ correlation ≥ 0.80
time bias ≤ 60 s
time correlation ≥ 0.80

A threshold that is not there

In a symptom-limited clinical cohort a subject may stop before a threshold is crossed, and a model that always returns two numbers will invent one. The current generation answers the presence question separately from the location question. In a development cohort of 103 clinical tests, 22 had no threshold at all; reporting them as located would have been the more confident output and the wrong one.

No result carries a confidence score. Nothing in Oxynet is calibrated against clinical outcomes, so a percentage would be a number without a meaning. The `quality` field describes the recording, not the certainty of the answer.

The figure opposite is the interface. This is the engine.

PART TWO · 2.4 · WORKED EXAMPLE

Oscillation, measured rather than declared BETA

Conventionally reported as present or absent. What the engine returns is the rhythm itself, with the evidence that it is real and the reason it was not graded.

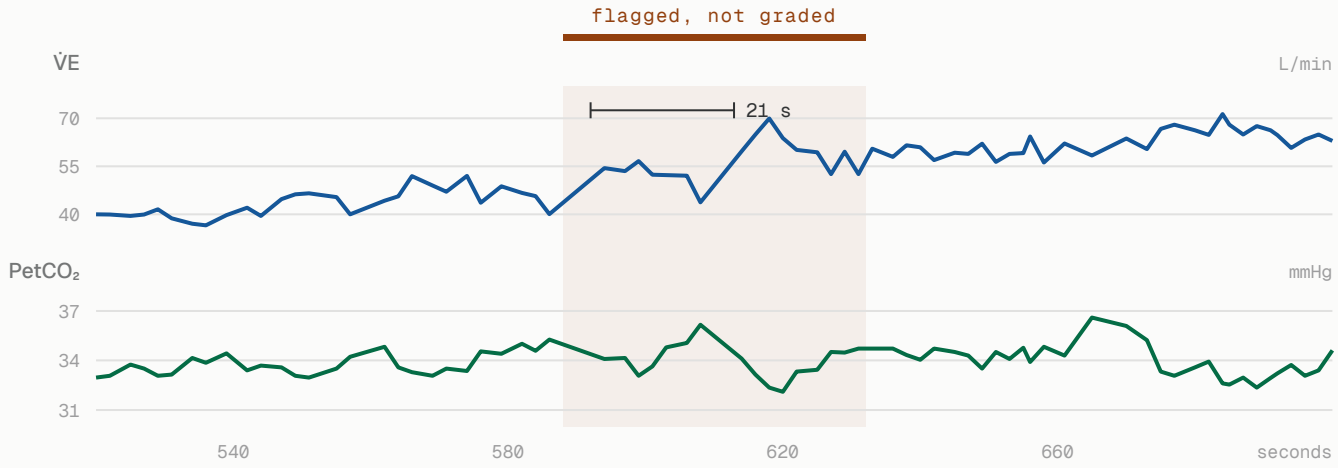

The episode in the raw traces first: ventilation cycling while end-tidal carbon dioxide moves against it, 164 degrees out of phase, which is what a disturbance in the ventilatory control loop does and what a movement artefact does not.

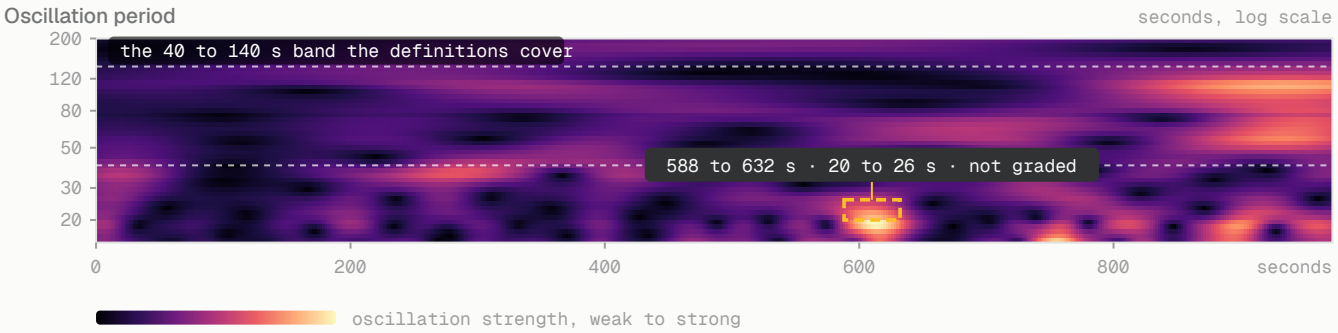

The same view the Oxynet application draws: a wavelet map of ventilation, time against oscillation period, brighter where the rhythm is stronger, with the detected episode boxed and dashed because it was flagged and not graded. The box, its period band and everything quantitative about it came back from the API. The map behind it is computed from the same ventilation trace for this figure, because the detector returns the episode list rather than the raster.

PERIOD	AMPLITUDE	CLARITY	DRIVEN BY
21 s	**17 %**	**6.7×**	**Rate**
band 20 to 26 s	of ventilation, 8.6 L/min	above this recording's own background	76 % breathing frequency

What the API returned

POST /v1/cpet/{id}/analyze {"analyses": ["eov"]}


```
{
  "analysis": "eov",
  "grade": "none",
  "findings": {
    "period_peak_s": 20.87,
    "amplitude_pct_ve": 16.78,
    "clarity": 6.65,
    "co2_coupling": "tight antiphase",
    "phase_deg": -164.4,
    "corroborated": true,
    "graded": false
  }
}
```


Not graded, and not normal either

The published definitions cover periods between roughly forty and one hundred and forty seconds. This rhythm ran at twenty-one, so it falls outside them and the engine declines to grade it. The burden is returned as zero, which is the correct answer to the question that was asked.

It also reports the rhythm, because a classifier answering only "no" would have discarded a corroborated oscillation seven times above background in a patient's recording. Which fact matters is the clinician's call, not ours.

> THE CAVEAT THAT TRAVELS WITH THIS
>
> Oscillation analysis is beta. It was developed on a single heart-failure cohort of 44 patients and has not been tested for transportability to a second population. It is offered for research use on that basis, and finding that second cohort is one of the collaborations sought in 4.4 on page 32.

PART TWO · 2.5 · WORKED EXAMPLE

Which fuel, at which intensity

Fat and carbohydrate oxidation read from gas exchange by indirect calorimetry, so it runs on any test with oxygen uptake and carbon dioxide output, with or without an ergometer channel.

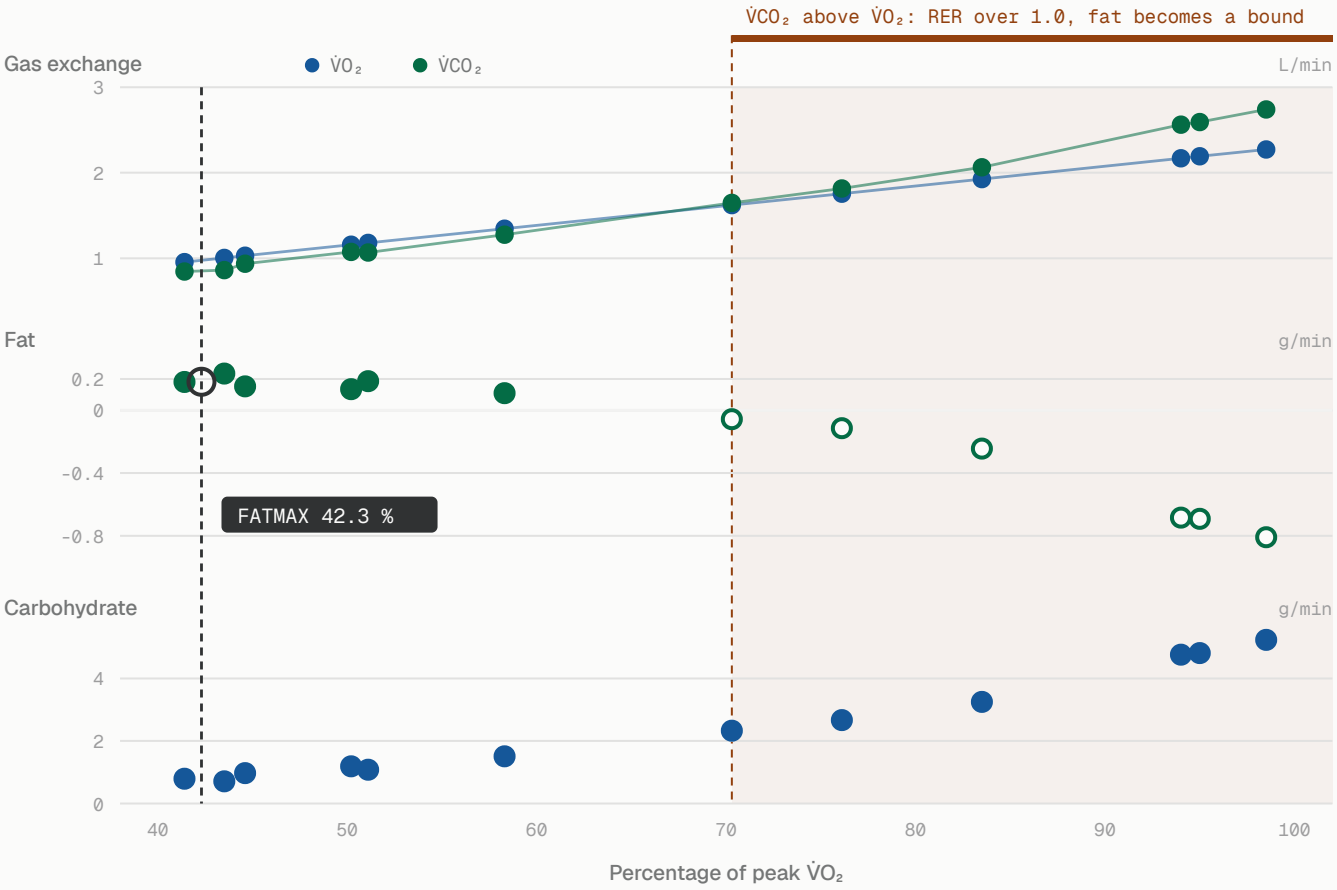

Twelve 60-second windows from the reference recording, each a measurement rather than a fitted point. FATMAX at 42.3 per cent of peak oxygen uptake is the turning point of a fitted cubic, not a measured sample. Beyond RER 1.0 the shaded region begins: there the exhaled carbon dioxide includes a non-metabolic part, the stoichiometry stops holding, and the fat figure is an upper bound that reads low and can read negative. Open markers are those bounds, drawn because hiding them would make the curve look better than the measurement is.

FATMAX

42.3 %

of peak oxygen uptake

PEAK FAT RATE

0.182 g/min

34.1 % of energy

WINDOWS

12

60 s clock bins

BOUNDS, NOT RATES

6 of 12

windows above RER 1.0

What the API returned

POST /v1/cpet/{id}/analyze {"analyses": ["substrate"]}


```
{
  "analysis": "substrate",
  "findings": {
    "fatmax": { "found": true, "pct_vo2peak": 42.3,
                "fat_g_min": 0.182, "load": null },
    "crossover": {
      "found": false,
      "reason": "Carbohydrate was already the larger fuel
                 at the lowest intensity measured."
    },
    "efficiency": {
      "available": false,
      "reason": "Needs a work rate in watts. This test
                 recorded no external load."
    }
  }
}
```


Two refusals, and why they are the point

The crossover was not found, and the engine says why: carbohydrate was already the larger fuel at the lowest intensity this protocol reached, so the crossover lies below anything the test covered. That is a fact about the protocol, not a missing value, and a system returning the lowest measured point as the crossover would be wrong rather than incomplete.

Gross efficiency was refused outright. It is mechanical work over metabolic energy and this export carries no work rate. Watts could be back-calculated from oxygen uptake, but only by assuming the efficiency the calculation is meant to measure, so the engine declines rather than returning a circular number.

One more caveat rides in the response and belongs in any report built on it: these windows are clock bins on a ramp, not steady-state stages. Nothing settled, so FATMAX reads high and is not comparable with a value from a graded test. The engine reports `steady_state: false` rather than leaving the reader to notice.

Method: fat and carbohydrate rates from Frayn (1983) and Jeukendrup and Wallis (2001), assuming negligible protein oxidation; energy at 9.75 and 4.18 kcal per gram. FATMAX is the peak of a cubic through the origin fitted after Achten and Jeukendrup, a fitted turning point that carries no confidence.

The figures in this part are the interface. The payloads are the engine.

PART THREE

03

How it is integrated

What the engine needs, what it returns, and the three architectures it can be deployed as.

PART THREE · 3.1

Hand Oxynet the file, work with the handle

One architectural rule shapes the whole API, and everything else follows from it.

The caller posts the vendor file exactly as the cart exported it and receives a handle. Every later operation names the handle, and the signals never pass back through the caller unless a picture is being drawn.

The parsing stays where the knowledge is

Twenty-one vendor formats, each with its own column names, unit conventions and clock quirks. A caller that normalises the file first has to reimplement all of it, and will lose information doing so.

The measurement runs at full resolution

Every analysis runs on the complete record, never on a decimated copy, so a narrow feature cannot fall between two kept samples.

A cohort becomes possible

A 360 KB export is roughly ninety thousand tokens inlined into a conversation. Through a handle, an assistant works through a hundred recordings holding only the structured results.

The lifecycle is four calls, whether the caller is a hospital integration, a script or an assistant.

```
get_capabilities()             # what may this key do
create_upload()                # -> a one-shot upload URL
$ curl -F "file=@test.csv" URL # -> cpet_id
analyze_cpet(id, ["vt","eov"])
compute_metrics(id, ["vo2max"])
get_cpet_series(id, ["VE"])    # only to draw a picture
```

ASK FIRST, DO NOT DISCOVER

Analyses are sold separately, so a key holding thresholds does not necessarily hold oscillation.

`get_capabilities` reports what the key may do, which models it can reach and what the limits are. Cheaper than handling a refusal.

EPHEMERAL BY DEFAULT

Uploaded bytes are parsed and discarded, never stored. The parsed record is deleted after 24 hours unless the caller asks for longer, up to 168.

PART THREE · 3.2

What the engine needs

The three analyses have different appetites. A file that supports thresholds does not necessarily support oscillation, and the API says which before anything is attempted.

CHANNEL	THRESHOLDS	OSCILLATION	SUBSTRATE	NOTES
Time s	●	●	●	Seconds from the start of the recording. Does not need a uniform grid
$\dot{V}O_2$ mL/min	●	●	●	The axis thresholds and substrate use are both reported against
$\dot{V}CO_2$ mL/min	●	●	●	Without it there is no substrate analysis at all
$\dot{V}E$ L/min	●	●	◐	The channel the oscillation is measured in
$PetO_2$ mmHg	●	◐	·	Half of the end-tidal closure check
$PetCO_2$ mmHg	●	●	·	Without it an oscillation cannot be corroborated, and the analysis says so rather than guessing
Heart rate bpm	○	○	○	Needed for O_2 pulse, and for nothing else
Breathing frequency min^{-1}	○	◐	·	Separates a rhythm driven by rate from one driven by volume
Work rate W	○	○	◐	Without it there is no gross efficiency, and FATMAX is an intensity rather than a wattage

● required ◐ preferred ○ optional · not used

Sampling

Breath by breath or interval averaged, both accepted, and records do not have to sit on a uniform grid: the engine resamples before inference and reports the interval it received.

The raw interval sets which oscillation periods can be resolved at all, so a coarse export narrows the searchable band. The narrowing is reported rather than being allowed to look like an absence.

What is never inferred

A channel that was not recorded is reported as absent, with the reason, and nothing is substituted for it.

Work rate is the case that matters commercially: watts can be back-calculated from oxygen uptake only by assuming the efficiency such a calculation is meant to measure. The engine refuses instead.

PART THREE · 3.3

REST

The surface a clinical system or a data pipeline integrates against. Authentication is an API key in a header; models and analyses are granted per key.

`GET /v1/capabilities`

What this key may do: licensed analyses, reachable models, the metric registry, limits and the retention policy.

`POST /v1/cpet`

Upload a recording. The format is detected; the response says what was parsed and which channels are absent, with a reason for each.

`POST /v1/cpet/{id}/analyze`

Run models. One envelope per analysis, and one that cannot run does not stop the others.

`POST /v1/cpet/{id}/compute`

Derived quantities from the metric registry: peaks, the slope profile, classical landmarks, input-quality checks.

`GET /v1/cpet/{id}/series`

A decimated slice of the signals, for plotting only. Never for computing.

The envelope

Every analysis returns the same shape, so a consumer writes one parser and not three.

```
{
  "status":     "ok" | "not_analysable",
  "findings":   { ... } | null,
  "quality":    "good" | "acceptable"
                | "poor" | "unusable",
  "notes":      [ "plain-language caveats" ],
  "provenance": { "model", "model_version",
                  "analysis_version",
                  "latency_ms" }
}
```

TWO FIELDS PEOPLE MISREAD

`quality` describes the INPUT recording. It is not a confidence in the answer, and there is no field that is.

`notes` is not decoration. It carries the caveats that must travel with the numbers into any report built on them.

MACHINE-READABLE

OpenAPI 3.0 at **/v1/openapi.json**, importable into a client generator, an agent action or an integration platform.

Plain-text integration guide at **/llms.txt**, so an agent reads the documentation in one request.

PROVENANCE ON EVERY RESULT

Which model, which version, which analysis version, and how long it took. A result can always be traced to the thing that produced it.

PART THREE · 3.4

MCP, and what an assistant may do

A native connector for Claude, Gemini CLI and any MCP client. The assistant orchestrates and explains. It does not do the physiology.

REPRODUCED SESSION · ABRIDGED

YOU

Analyse the test in ~/exports/R5-T3.json. Is there oscillatory ventilation?

TOOL

```
create_upload()  →  $ curl -F "file=@R5-T3.json"
```

cpet_id: cpet_7485262aa3f7 · exercise_threshold_app (detected)
353 breaths · 969 s · Work: absent, and nothing is inferred from $\dot{V}O_2$

TOOL

```
analyze_cpet(cpet_id, ["vt", "eov"])
```

vt → VT1 619 s / 1605 mL/min · VT2 763 s / 2029 mL/min · quality good
eov → grade "none" · 1 event flagged, not graded · gas quality confirmed

AGENT

Not by the graded definition, and that is not the whole answer. Oxynet found one rhythm, from 9.8 to 10.5 minutes: breathing cycled every 21 seconds, swinging 17 % of ventilation, seven times above this recording's own background, with end-tidal CO_2 moving in antiphase to corroborate it. That period is faster than the 40 to 140 seconds the EOV definitions cover, so it is flagged and not graded.

Oscillation analysis is beta and was developed on a single heart-failure cohort, so treat this as a research finding. There is no work-rate channel in this export, so nothing here is read against a work rate.

THE RULE

The file goes from disk to Oxynet and only the summary comes back. The signals never travel through the conversation, which is what makes a cohort tractable at all.

ONE LINE TO CONNECT

```
claude mcp add --transport http oxynet https://app.oxynet.net/oxynet-mcp --header "X-API-Key: ..."
```

WHY THIS MATTERS

A general-purpose assistant asked to read a CPET will estimate a threshold from the numbers in front of it. Through MCP it calls the model that was evaluated against expert labelling, and reports what came back.

A REFUSAL IS INFORMATION

Where a recording cannot support an analysis the API says so and says why, and the assistant is instructed to report that rather than work around it. So is declining to attach a confidence to a result that has none.

PART THREE · 3.5

Cloud, local, embedded

The same engine, deployed three ways. Which one fits is a product, regulatory and infrastructure decision rather than a technical one, and it is a conversation rather than a menu.

A Hosted API

```
your software
↓ HTTPS
Oxynet, hosted
↓ JSON
your software
```

For: the fastest route to a working integration, central model updates, nothing to deploy or keep current

Against: recordings leave the local environment, and connectivity becomes a dependency

Live today. This is what the worked examples in Part Two were run against.

B Local engine

```
your software
↓
Oxynet, on your
infrastructure
↓
your software
```

For: recordings never leave the institution, low latency, suits restricted clinical environments

Against: you own deployment and update management, and models have to be shipped to you

The Python package already performs local inference. A packaged server deployment is scoped per partner.

C Embedded

```
your product
┌──────────┐
│ Oxynet   │
│ engine   │
└──────────┘
```

For: interpretation inside the product, available during the test rather than after it, no network at all

Against: the tightest coupling, and the longest conversation about versioning and support

Models are small convolutional networks and export to on-device runtimes. Real-time and edge deployment is a direction the architecture supports, and a joint scoping exercise rather than a shipping product.

The engine is the same in all three. What changes is where it runs and who operates it, so a partnership can start on the hosted API and move inward without the interpretation changing underneath the customer.

PART THREE · 3.6

Limits, retention and governance

What an integrator needs before a pilot, and what a data protection officer needs before signing one.

Operational limits

	VALUE
Upload size	25 MB
Retention, default	24 h
Retention, maximum	168 h
Upload ticket life	5 min
Record length	40 to 4,800 samples
Formats read	21, auto-detected

Access

An API key per caller, in a header. Analyses and models are granted per key, so a key that holds thresholds does not hold oscillation unless it was sold.

The upload URL carries its own short-lived credential, so the shell command that posts a file handles no API key at all.

Data handling, stated exactly

Uploaded bytes are parsed and discarded. They are never written to storage. What persists is the parsed record, deleted after 24 hours unless the caller asks for longer.

Nothing a customer uploads enters the training corpus. Research cohorts that do are pseudonymised at the point they are emitted, before they reach it.

Option B on the previous page exists for environments where identifiable data must not leave the institution at all.

These are the technical characteristics of the system. Whether a given deployment satisfies a given jurisdiction is a compliance assessment, and one no vendor document can perform on your behalf.

This manual is a snapshot. The API reference at **app.oxynet.net/docs** is generated from the running service and is authoritative wherever the two differ, and **app.oxynet.net/v1/openapi.json** is the machine-readable form of the same thing.

PART FOUR

04

Status, evidence and next steps

The limits, stated plainly, and a concrete way to start.

PART FOUR · 4.1

What Oxynet is not

Stated plainly, because a capability list read by someone deciding what to build on is only useful next to its limits.

× **Not a medical device.**
Research software. No CE mark, no FDA clearance, not registered as software as a medical device in any jurisdiction, and not to be the basis of a diagnosis or a treatment decision.

× **Not a diagnosis.**
The engine returns physiological measurements. Interpreting them for a patient is the clinician's act, and the clinician remains the responsible party.

× **Not calibrated against outcomes.**
No output carries a confidence score, because there is nothing to calibrate one against. Agreement with expert labelling is what has been measured, and it is not the same claim.

Capabilities have different evidence maturity

Not everything here stands on the same ground, and a partner needs to know which is which before they build. The three states are used consistently throughout this document.

PRODUCTION

Answering on the live API, with the evaluation that let each model ship stored beside it. Threshold detection carries the longest record.

AVAILABLE

Shipped and usable, and the surface rather than the science: the MCP endpoint, the schema, the package, the deployment options.

RESEARCH

Real, running, and not yet transportable. Oscillation analysis sits here on one cohort. It says so wherever it appears.

Current status

Oxynet is operated by Andrea Zignoli as an independent technology platform. The computational engine, the models, the software interfaces and the research record are established and in use. Commercial incorporation and partnership structures are under development, and this document is not an offer of securities or a solicitation to invest.

PART FOUR · 4.2

The evidence record

What has been shown, how a model is held to it, and what has not been shown. The third column is the one worth reading.

VALIDATED

Threshold detection against expert labelling, across cohorts differing in population, protocol, ergometer and metabolimeter.

7 peer-reviewed papers on the methodology and its evaluation.

Parts of the corpus carry several independent raters per test, so agreement with the engine can be read against agreement between two humans.

OPERATIONAL

Everything marked Production on page 6 is answering on the live API now.

Every promoted model ships with the evaluation that let it through the gate, so a result can be re-checked without re-running inference.

Deterministic inference: the same recording returns the same numbers on the same model version.

RESEARCH

Oscillation analysis, developed on one heart-failure cohort of 44 patients. Transportability to a second population is unknown.

Self-supervised pretraining, in development and in no promoted model.

Nothing is calibrated against a clinical outcome. No result predicts mortality, hospitalisation or response to treatment.

The landmarks corroborate, they do not adjudicate

The engine can also compute the classical criteria: the V-slope and respiratory compensation break points, the ventilatory equivalent nadir, the end-tidal carbon dioxide peak. They are textbook rules evaluated on the file's own signal. They carry no confidence and they never override a model.

On the reference recording the automated V-slope fit returned slopes of 1.30 and 1.42 either side of its break point. Slopes that barely differ mean no break point was found, and the response reports the fit so a reader can see that rather than being handed a number to trust.

PROMOTION GATE

Evaluated on a cohort it never saw, on both axes, and all four must clear:

$\dot{V}O_2$ bias ≤ 120 mL/min
$\dot{V}O_2$ correlation ≥ 0.80
time bias ≤ 60 s
time correlation ≥ 0.80

The publication list is maintained at **www.oxynet.net**, and the open-source Python package at **pypi.org/project/pyoxynet** covers local inference and synthetic data generation.

PART FOUR · 4.3

References

Seven peer-reviewed journal articles behind the methods described here, one preprint, and the writing that explains them.

PEER REVIEWED

01 **AI-Driven Analysis of CPET to Identify Gas Exchange and Ventilatory Thresholds**
Sports Medicine · 2026
pubmed.ncbi.nlm.nih.gov/41784915

02 **AI for CPET Interpretation**
Biomedical Signal Processing and Control · 2023
www.sciencedirect.com/science/article/abs/p…

03 **Regression, Generation, and Explanation**
Sensors (MDPI) · 2023
doi.org/10.3390/s23020826

04 **Crowdsourcing and CNN for Intensity Domain Determination**
European Journal of Sport Science · 2021
doi.org/10.1080/17461391.2020.1866081

05 **LSTM Networks for VO_2 Estimation**
PLOS ONE · 2020
journals.plos.org/plosone/article?id=10.137…

06 **AI Technologies in Exercise Data Processing**
Sport Sciences for Health · 2019
link.springer.com/article/10.1007%2Fs11332-…

07 **LSTM for Intensity Domain Estimation**
European Journal of Sport Science · 2019
doi.org/10.1080/17461391.2019.1587523

PREPRINT

Conditional GANs for Synthetic CPET Data
www.overleaf.com/read/fcmwscvyhtfq

ALSO WRITTEN

Oxynet: A Collective Intelligence Approach
www.linkedin.com/pulse/oxynet-collective-in…

AI in CPET Data Interpretation
andreazignoli.github.io/blog-post-5

Automatic Interpretation of CPET with Deep Learning
medium.com/@andrea.zignoli/automatic-interp…

Generating Realistic CPET Data with Python
medium.com/@andrea.zignoli/automatic-genera…

The list is maintained at **www.oxynet.net** and this page is generated from it, so the two cannot disagree. Blog and Medium pieces are listed because they explain the work, not as evidence for it.

PART FOUR · 4.4

Routes to market, and the next step

Four ways to build on the platform. They are the same engine reached at different depths, not four different products.

01 **Integrate**
Call Oxynet from software you already ship. Your interface, your reporting, your customer relationship, our physiology, kept current.

02 **Embed**
Make it part of the product itself, on your infrastructure or on the device, for interpretation during the test rather than after it.

03 **Validate**
Run cohorts through the engine to establish new capabilities and the evidence behind them. Oscillation analysis needs a second population before it can leave research, and that is a collaboration rather than a purchase.

04 **Partner**
A broader commercial relationship around distribution, licensing or joint product development.

Commercial and licensing structure follows from the route, the deployment and the scope of rights, so it is deliberately not fixed here. Investment conversations are welcome and are a separate one.

START HERE

Send one representative file.

01 **One dataset**
You send a representative export. We run it and return the structured result.

02 **Feasibility note**
What worked, what the format needed, what the engine refused and why.

03 **Integration prototype**
Against the hosted API, in a sandbox, on your side.

04 **Workflow and validation**
Your cohort, your acceptance criteria, your regulatory position.

05 **Deployment**
Hosted, local or embedded, per 3.5.

Phase one needs one file and no commitment.

COMMERCIAL ENQUIRIES

Andrea Zignoli, MEng, PhD

andrea.zignoli@unitn.it

Responsible for the commercial use of these tools. Licensing, pilots, integration and partnership all come to this address.

Acknowledgments: the Exercise Thresholds application, and the laboratories whose cohorts made this work possible. Oxynet Manual v1.8, September 2026. The API reference at app.oxynet.net/docs is generated from the running service and is authoritative wherever it differs from this document.