Overview of long-baseline accelerator neutrino experiments

Current status and the next generation

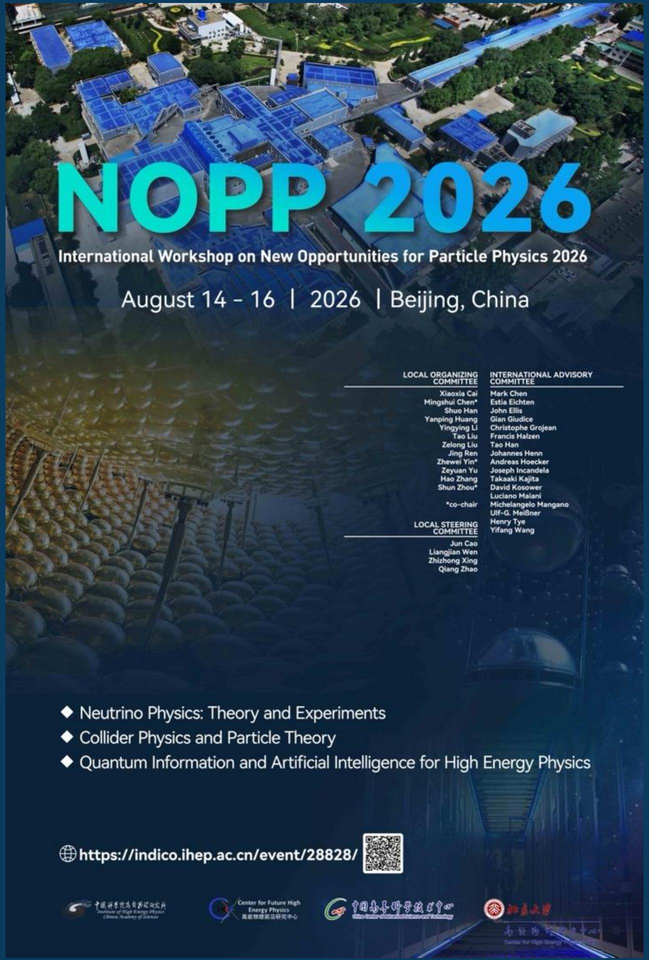

Francesca Di Lodovico

King's College London · Visiting Professor, Tokyo College, University of Tokyo

International Workshop on New Opportunities for Particle Physics 2026 · IHEP · 16 August 2026

What this talk covers

Twenty-five years of accelerator long-baseline physics, and the decade ahead

1. **The technique** — beams, near detectors, and the degeneracy that shapes every design choice
2. **Where we are** — T2K, NOvA, and the first joint fits
3. **What limits us — the interaction model, the flux, and the event count: all three have to improve together**
4. **The next generation — Hyper-Kamiokande and DUNE**: an order of magnitude more events, and systematics built down to match
5. **And alongside — reactor and atmospheric experiments, and the detector and cross-section R&D that feeds all of it**

The state of play in one sentence

CP conservation excluded at 95% C.L. by T2K; $|\Delta m^2_{32}|$ known to under 2%. Still open: the mass ordering, inseparable from δCP — and the θ_{23} octant, since the data sit near maximal mixing.

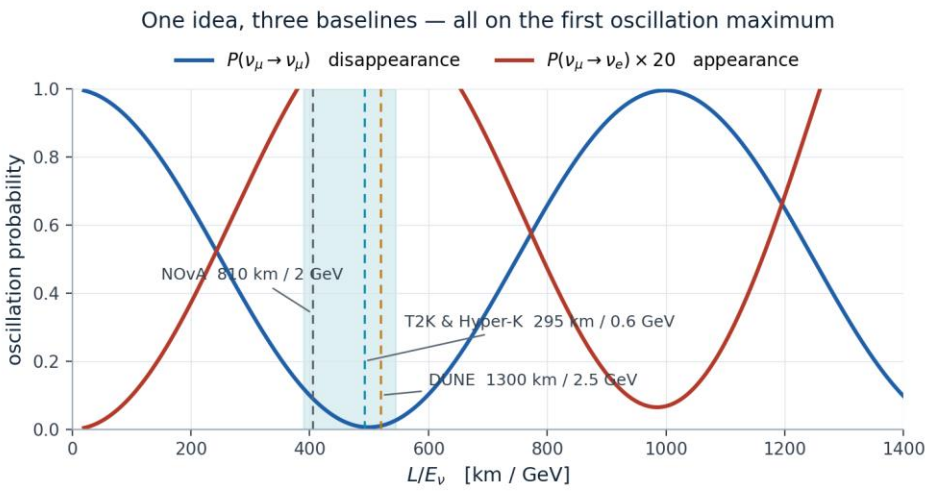

The programme at a glance

Four generations

Experiment	Baseline	Energy	L/⟨E⟩	Far detector	Years	Principal legacy
K2K	250 km	1.3 GeV	190	Water Cherenkov (SK)	1999–2004	First accelerator confirmation of oscillations
MINOS / MINOS+	735 km	3 GeV	245	Magnetised steel–scintillator	2005–2016	Precision Δm^2_{32}; refined θ_{23} presented 2026
OPERA	730 km	17 GeV	43	Emulsion	2008–2012	Direct ν_τ appearance
T2K	295 km	0.6 GeV	490	Water Cherenkov (SK)	2010–	θ_{13}; first indications of CP violation
NOvA	810 km	2 GeV	405	Segmented liquid scintillator	2014–	Most precise single-experiment Δm^2_{32}
Hyper-Kamiokande	295 km	0.6 GeV	490	Water Cherenkov, 188 kt	2028–	5σ CP violation; multi-purpose observatory
DUNE	1300 km	2–3 GeV	520	Liquid argon TPC, >20 kt	2031–	5σ CP violation; mass ordering from matter effects; multi-purpose observatory

L/⟨E⟩ in km/GeV — the first oscillation maximum is at ≈ 500, the second at ≈ 1500, where no experiment sits today.

The questions

And why each one is hard

Question	Observable	Why it is difficult
Is CP violated?	$\nu_\mu \to \nu_e$ vs $\bar{\nu}_\mu \to \bar{\nu}_e$ asymmetry	Matter effects fake the same asymmetry. The formula is degenerate in δCP too: sin δ and cos δ terms let δ and π − δ give the same ν_e rate — only ν vs $\bar{\nu}$ and the spectrum break it.
Mass ordering?	Matter effect, or the reactor spectrum	~10% of the rate at 295 km, ~30–40% at 1300 km. Small by design at 295 km, so the ordering has to come from atmospherics, a reactor, or a longer baseline.
θ_{23} octant?	ν_μ disappearance + ν_e appearance	Disappearance measures $\sin^2 2\theta_{23}$ — degenerate about 45°
Sub-per-cent precision?	Spectral shape	Energy scale, and the nuclear model relating true to reconstructed energy

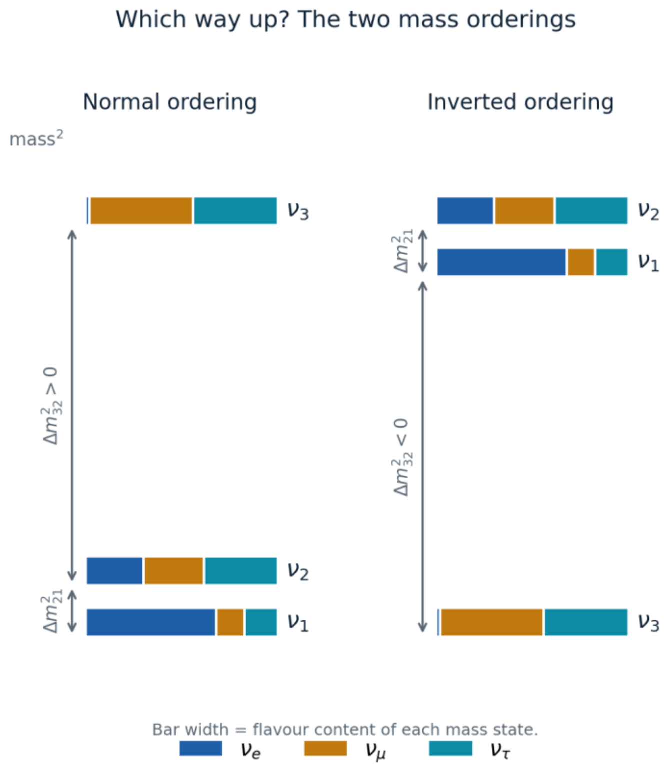

These four are not independent. Every current result quotes the others' assumptions — which is why joint fits, and experiments with different degeneracy directions, matter so much.

The technique

Two design philosophies

	Narrow band, off-axis	Wide band, on-axis
Experiments	K2K, **T2K**, **NOvA**, **Hyper-K**	MINOS / MINOS+, OPERA, **DUNE**
Spectrum	Peaked at the first oscillation maximum	Covers first and second maxima
Advantage	Small feed-down background, well-defined energy	Spectral shape breaks degeneracies within one experiment
Implications	Little information away from the peak	Harder energy reconstruction, more background

Bold = running now or approved.

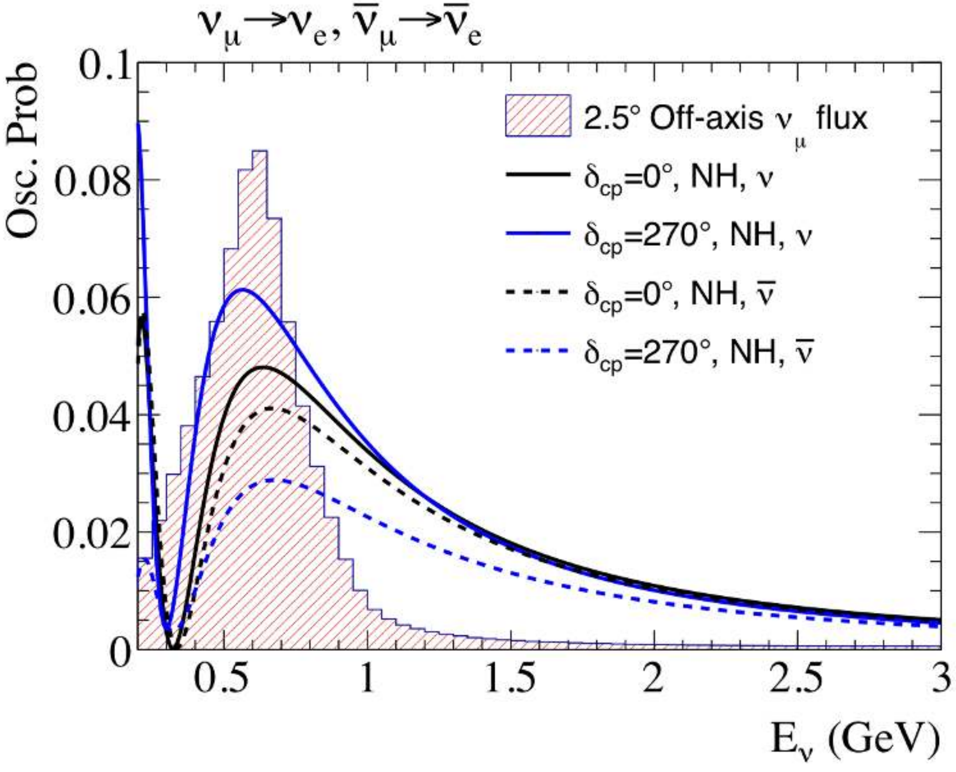

The off-axis technique.

- Both share the same chain: protons on target → horns → decay volume → a mostly-ν_μ beam, with a near and a far detector.
- **And both face the same three-way degeneracy: δCP ↔ mass ordering ↔ θ_{23} octant.**
- **Four ways out: a longer baseline, atmospherics, reactor constraints, joint fits.**

Why off-axis

The same beamline, seen from two angles

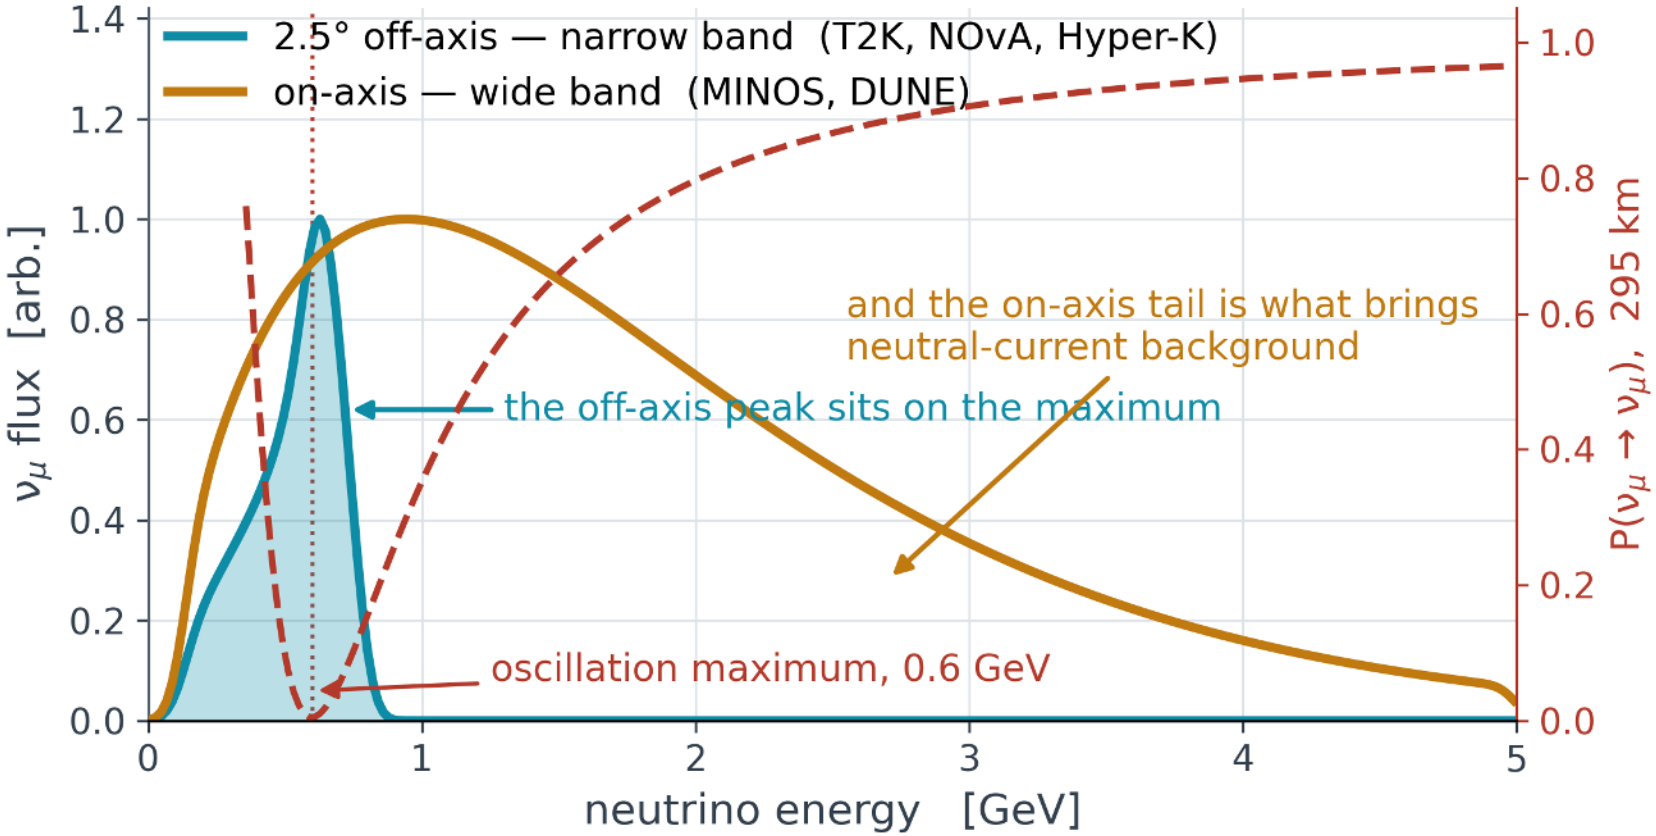

Illustrative, from two-body pion-decay kinematics: $E_\nu = 0.43\, E_\pi / (1 + \gamma^2\theta^2)$, same beamline viewed at two angles.

Why it matters

- **Narrow band on the oscillation maximum: the energy is known before you reconstruct it.**
- The tail is the problem — a neutral-current π^0 fakes a ν_e event. Cutting the tail removes it at source.
- The cost: one point in L/E. Both experiments recover the rest by moving the NEAR detector off-axis (IWCD, DUNE-PRISM); DUNE's far detector is on-axis.

The experiments in this talk

Four accelerator experiments, three detector technologies

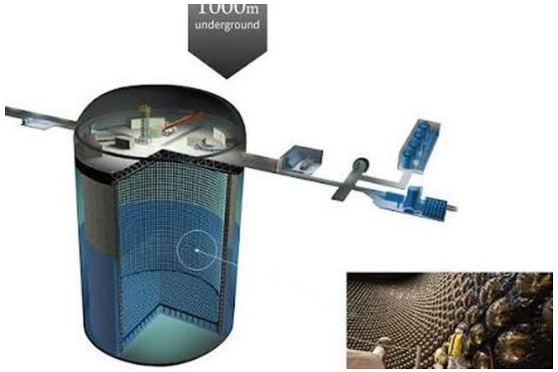

T2K

295 km, 0.6 GeV
50 kt water Cherenkov (Super-K)
running since 2010
target: oxygen

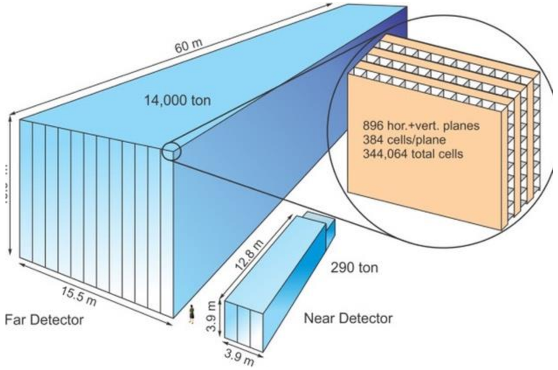

NOvA

810 km, 2 GeV
14 kt segmented scintillator
running since 2014
target: carbon

Hyper-Kamiokande

295 km, 0.6 GeV
188 kt water Cherenkov
first data 2028
target: oxygen

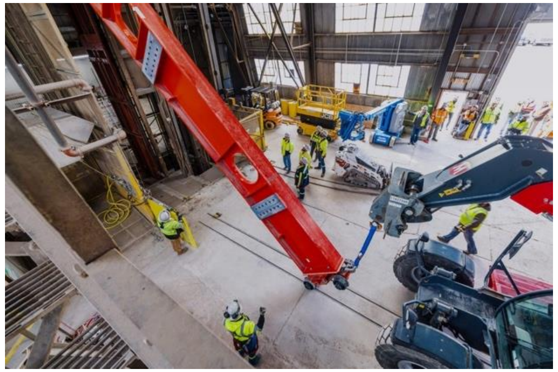

DUNE

1300 km, wide band 0.5–5 GeV
liquid-argon TPC
first beam 2031
target: argon

All four measure the same PMNS parameters with different baselines, energies and target nuclei — which is why their systematic uncertainties do not cancel against each other. JUNO, at 52.5 km from reactors, adds a fifth route in which the matter effect is a small correction rather than the signal; it has its own talk at this workshop.

1

Where we are

T2K, NOvA, and the first joint fits between experiments.

T2K

295 km at 2.5° off-axis, and a beam that keeps improving

900 kW	2.4 × 10^{14}	1.3 MW
stable beam power since January 2026	protons per pulse — a synchrotron world record	target after the summer-2026 upgrade

- 30 GeV protons from the J-PARC Main Ring; Super-Kamiokande at 295 km; 2.5° off-axis giving ⟨Eν⟩ ≈ 0.6 GeV at the first oscillation maximum.
- The Main Ring has already demonstrated 1 MW with fast extraction to the neutrino beamline.
- **ND280 upgrade — SuperFGD, two High-Angle TPCs, six time-of-flight planes — complete since May 2024 and taking data. Full 4π acceptance, ~300 MeV/c proton threshold.**
- **It is not yet in the oscillation fit — the next systematic step.**

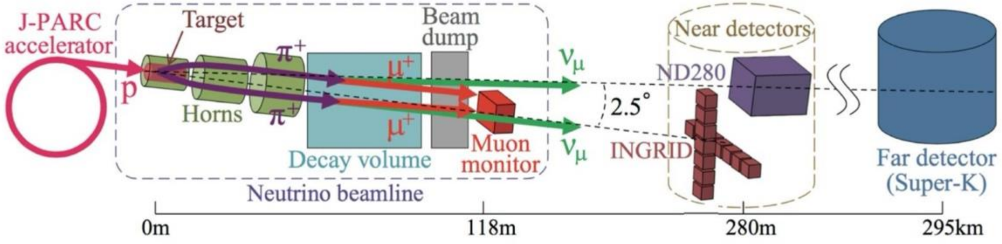

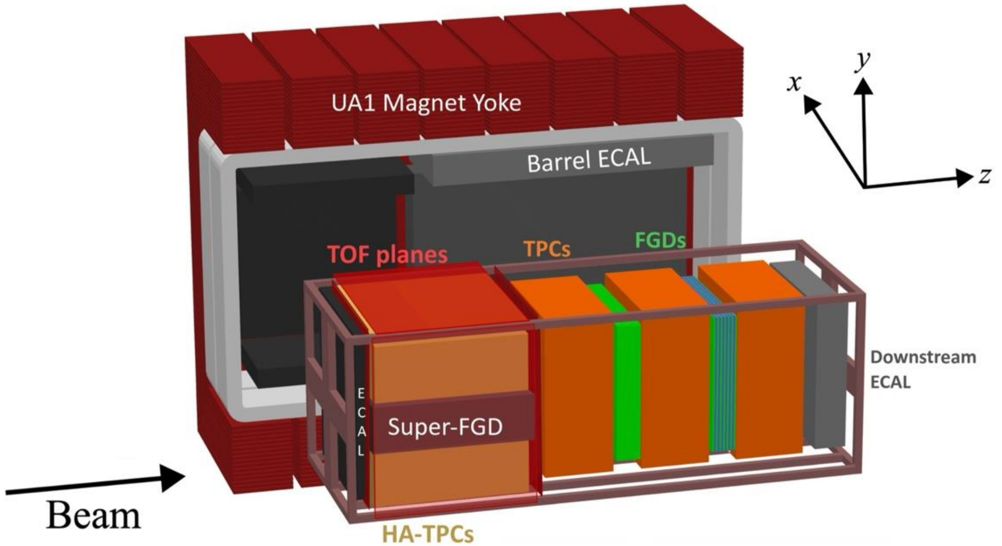

T2K and the upgraded near detector.

T2K at Neutrino 2026

Run 1–11 "v2" — the same protons, a substantially new analysis

- **Exposure 21 × 10^{20} POT (ν) and 16 × 10^{20} POT ($\bar{\nu}$); 5.6 × 10^{21} POT collected to date.**
- Updates: new ND280 backward-going and high-angle selections; multi-ring events in the SK ν_e sample; and a rebuilt interaction model — new spectral function for oxygen, retuned 2p2h and pion production.
- Uncertainty on the predicted far-detector event rate, before **→ after the near-detector fit: 1-Ring μ 19.2% → 3.7%; 1-Ring e 20.3% → 5.7%.**

95% C.L. exclusion of CP conservation — 90% in 2024	**78 → 81%** posterior probability for NO; IO rejected at 93% C.L.	**1.9%** total error on Δm^2_{32} at fixed ordering	**67%** posterior probability for the upper octant — both allowed

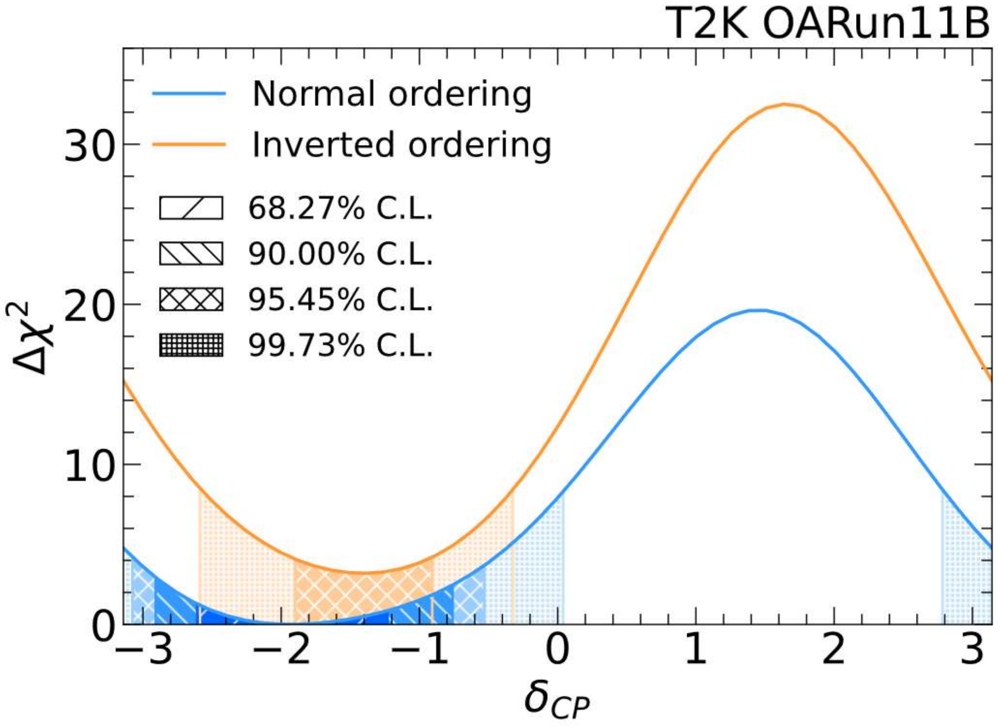

T2K Run 1–11 v2, preliminary.

NOvA — ten years

The most precise single-experiment $|\Delta m^2_{32}|$

- 810 km, 14 mrad off-axis, ⟨Eν⟩ ≈ 2 GeV; 14 kt segmented liquid scintillator.
- 26.61 / 12.50 × 10^{20} POT (ν / ν̄): 181 ν_e, 384 ν_μ, 32 $\bar{\nu}_e$, 106 $\bar{\nu}_\mu$ candidates.
- **Δm^2_{32} = 2.431 (+0.036 / −0.034) × 10^{-3} eV² (NO).**
- **$\sin^2\theta_{23}$ = 0.55 (+0.02 / −0.06) — close to maximal, so no octant preference. δCP = 0.93 (+0.21 / −0.89) π.**
- **Ordering: normal preferred** — Bayes factor 2.2 (69% posterior) with no reactor constraint, 3.3 (77%) with the 1D Daya Bay constraint, 6.6 (87%) with the 2D Δm^2_{32}–$\sin^2 2\theta_{13}$ one.
- NuMI runs to ≈2027, then the shutdown for LBNF and PIP-II.

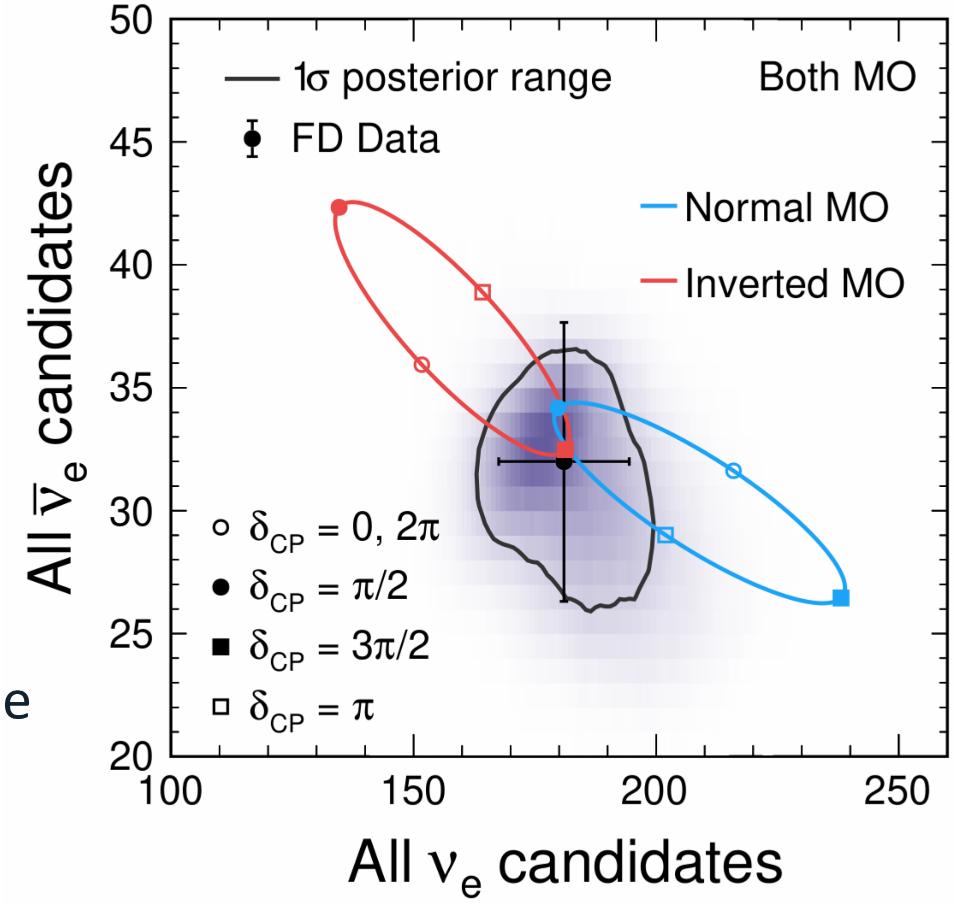

Bi-event plot: the data sit where the two ordering ellipses cross — the least discriminating point.

Joint fit: T2K + Super-Kamiokande atmospheric

Using the same far detector twice

- SK I–V, 3244.4 days of atmospherics with T2K's 19.7 / 16.3 × 10^{20} POT.
- **Matter-resonant ν_e appearance at 2–10 GeV constrains ordering and octant at once — the degeneracy that limits T2K alone.**
- Common interaction model and correlated detector systematics across the two.
- **CP conservation (J_CP = 0) excluded at 1.9σ;** P(NO) = 0.90. J_CP is the Jarlskog invariant — zero if CP is conserved.
- **Best fit (NO): δCP = −1.76 (+0.73 / −0.95); $\sin^2\theta_{23}$ = 0.468 (+0.106 / −0.025); Δm^2_{32} = 2.520 (+0.048 / −0.058) × 10^{-3} eV^2.**

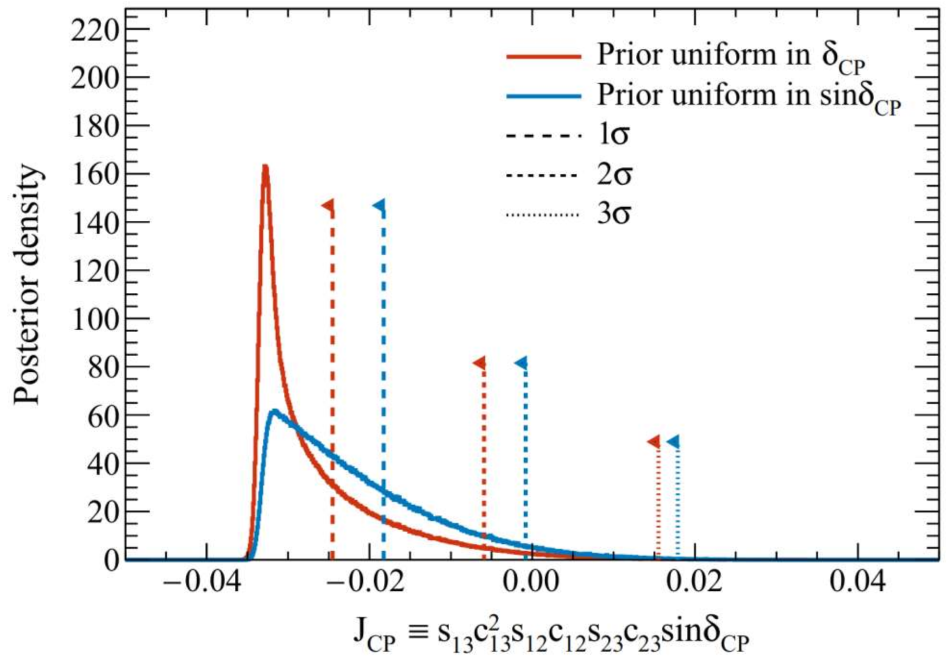

First joint T2K + Super-K oscillation analysis.

Joint fit: T2K + NOvA

The combination can move the answer

- **Δm^2_{32} = 2.43 (+0.04 / −0.03) × 10^{-3} eV^2 (NO) — the smallest experimental uncertainty on $|\Delta m^2_{32}|$ to date.**
- $\sin^2\theta_{23}$ = 0.56 (+0.03 / −0.05); upper-octant Bayes factor 3.5 with the reactor constraint.
- δCP: 3σ credible interval [−1.38π, 0.30π] in NO — δCP ≈ +π/2 lies outside it.
- **Mass ordering: the combination mildly prefers inverted — Bayes factor 1.3 with the reactor constraint, 2.5 without. Neither is statistically significant.**

Each experiment alone prefers normal ordering; the other parameters agree better under inverted, so the joint fit flips — not significantly.

Joint fit: T2K + NOvA — the result

$\sin^2\vartheta_{23}$ against δCP — normal ordering (blue, a) and inverted (orange, b). The combination mildly prefers inverted: Bayes factor 1.3 with the reactor constraint, 2.5 without.

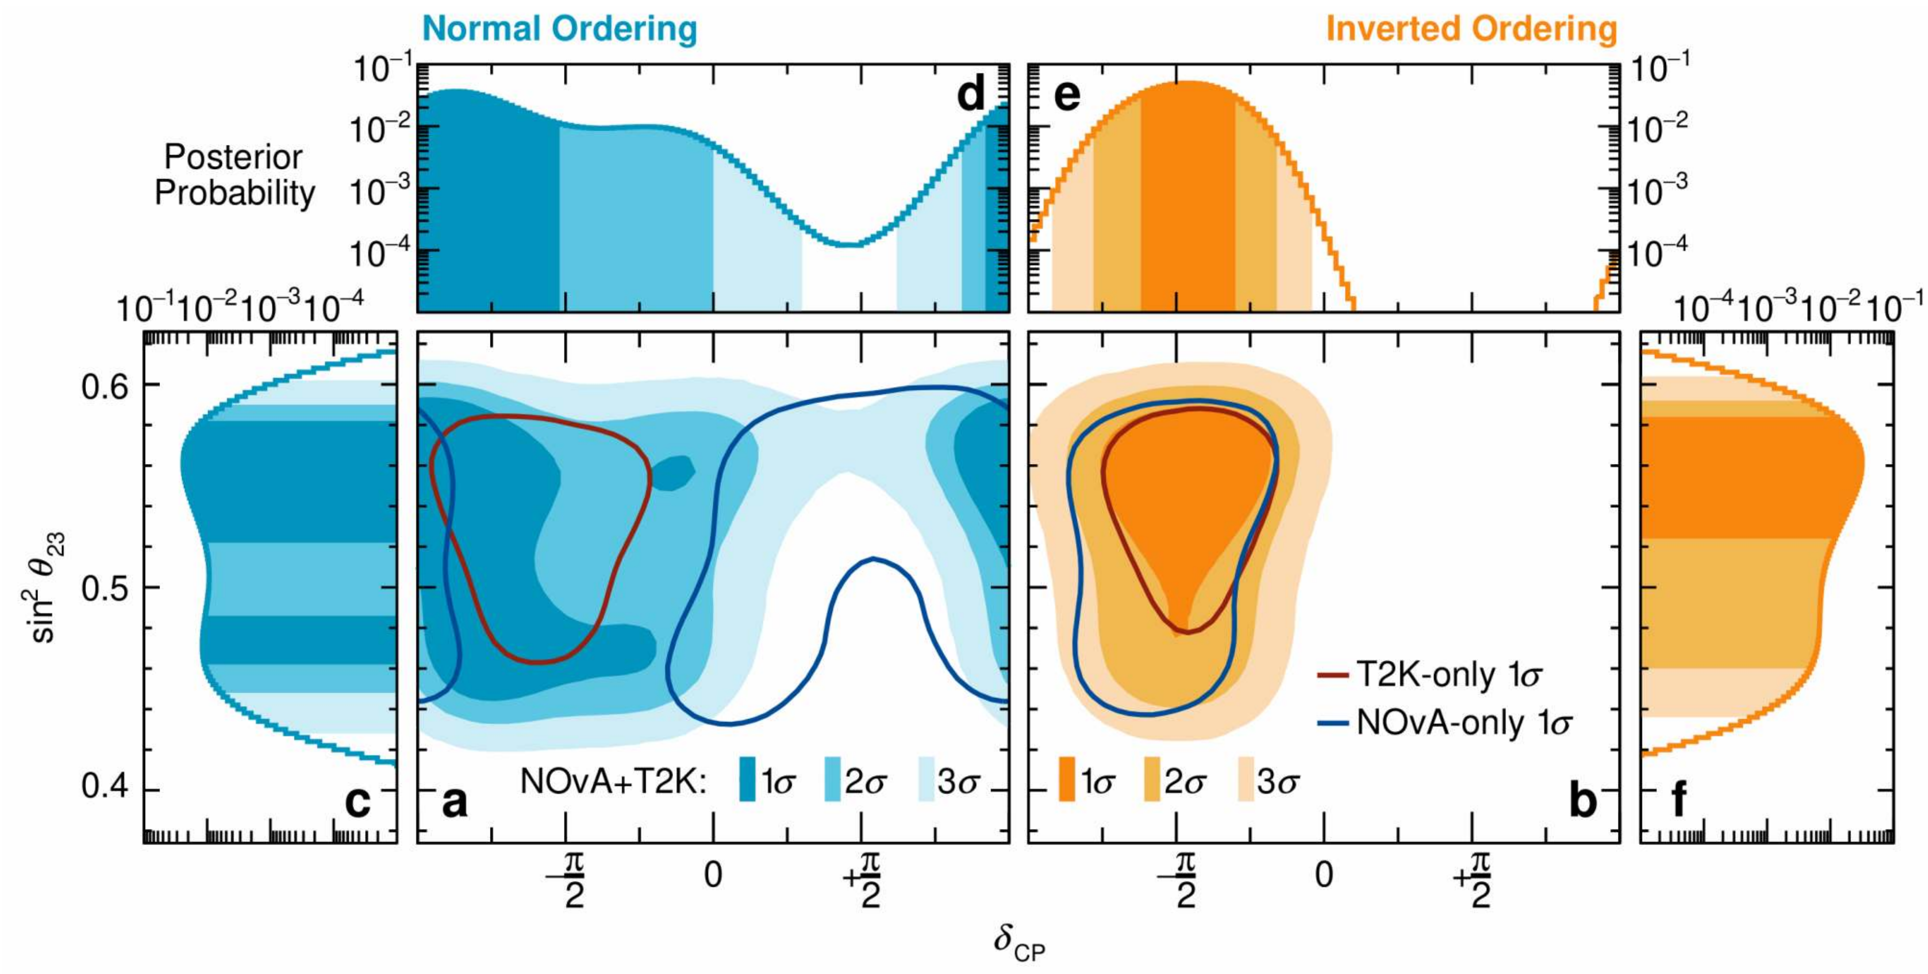

The wider context

Reactors, atmospherics, and the global fit

- Super-K and IceCube/ORCA add an independent ordering handle.
- Global fit: $\Delta\chi^2(\text{IO} - \text{NO}) = 1.5$–$5.9$ depending on which atmospheric data are included — at best $\lesssim 2.4\sigma$.
- **δCP is conditional on the ordering: assume normal and the best fit sits near CP conservation, which stays allowed; assume inverted and the data need δCP ≈ 3π/2, so CP conservation is excluded at ~4σ.**
- *NuFIT 6.1 predates the JUNO data*

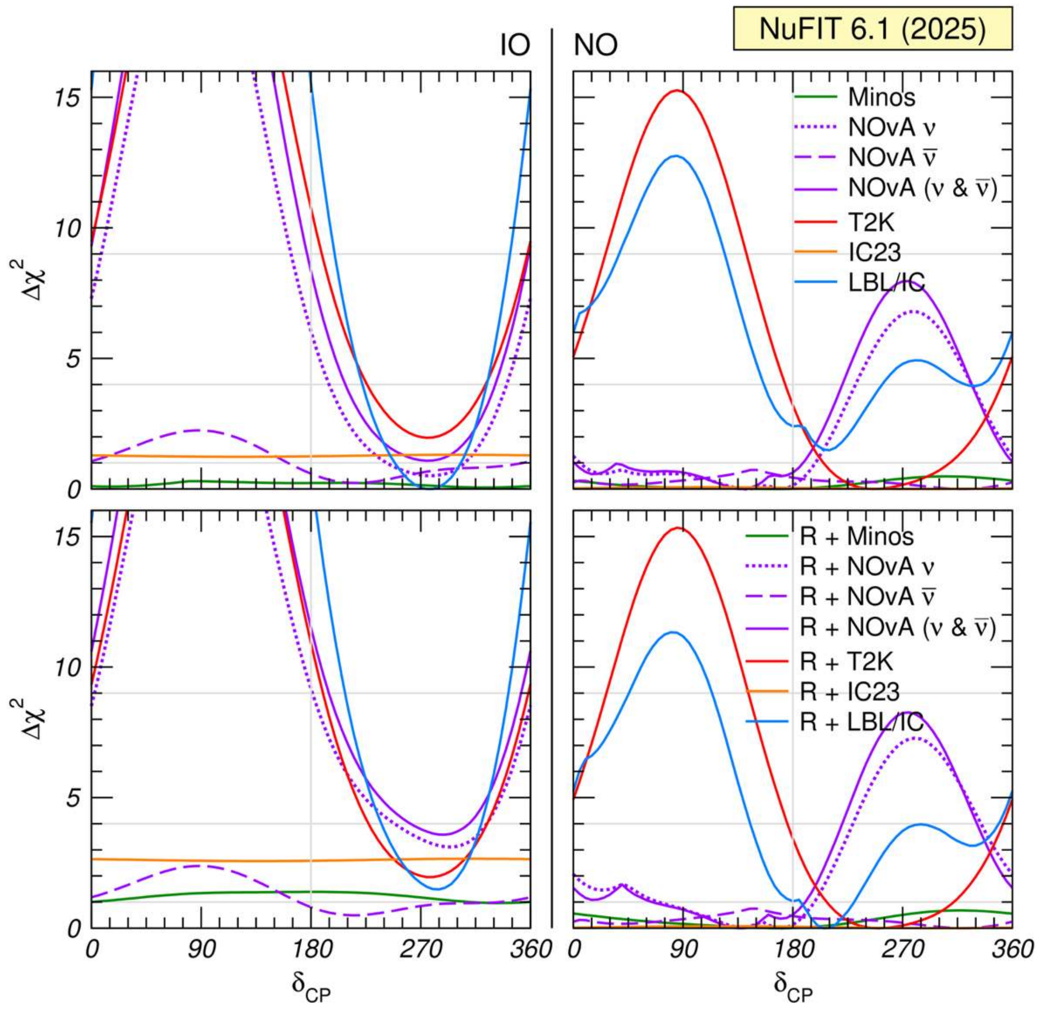

NuFIT 6.1: $\Delta\chi^2$ vs δCP per dataset. Left IO, right NO; bottom row adds the reactor ϑ_{13} constraint (R). Best fit = where a curve touches zero.

Two cautionary notes

Where the field is not yet self-consistent

MINOS+ refined θ_{23}

- A re-analysis of 23.761×10^{20} POT with FLUKA flux reweighting, MEC added to the model, and an improved kNN hadronic-shower energy estimator gives, with the reactor constraint: $\sin^2\theta_{23} = 0.37 \pm 0.03$, $\Delta m^2_{32} = (2.47 \pm 0.08) \times 10^{-3}$ eV² (NH). Maximal mixing excluded at 3σ — the excluded window is 0.44–0.59.
- **The published MINOS+ result (PRL 125, 131802) was $\sin^2\theta_{23} = 0.43$ (+0.20 / −0.04).** The central value barely moved; the upper error shrank by a factor of seven.
- **T2K finds 0.56, NOvA 0.55**. Preliminary and unpublished: the paper is in preparation, and the result wants reproducing with an independent energy estimator before anyone redraws their picture of θ_{23}.

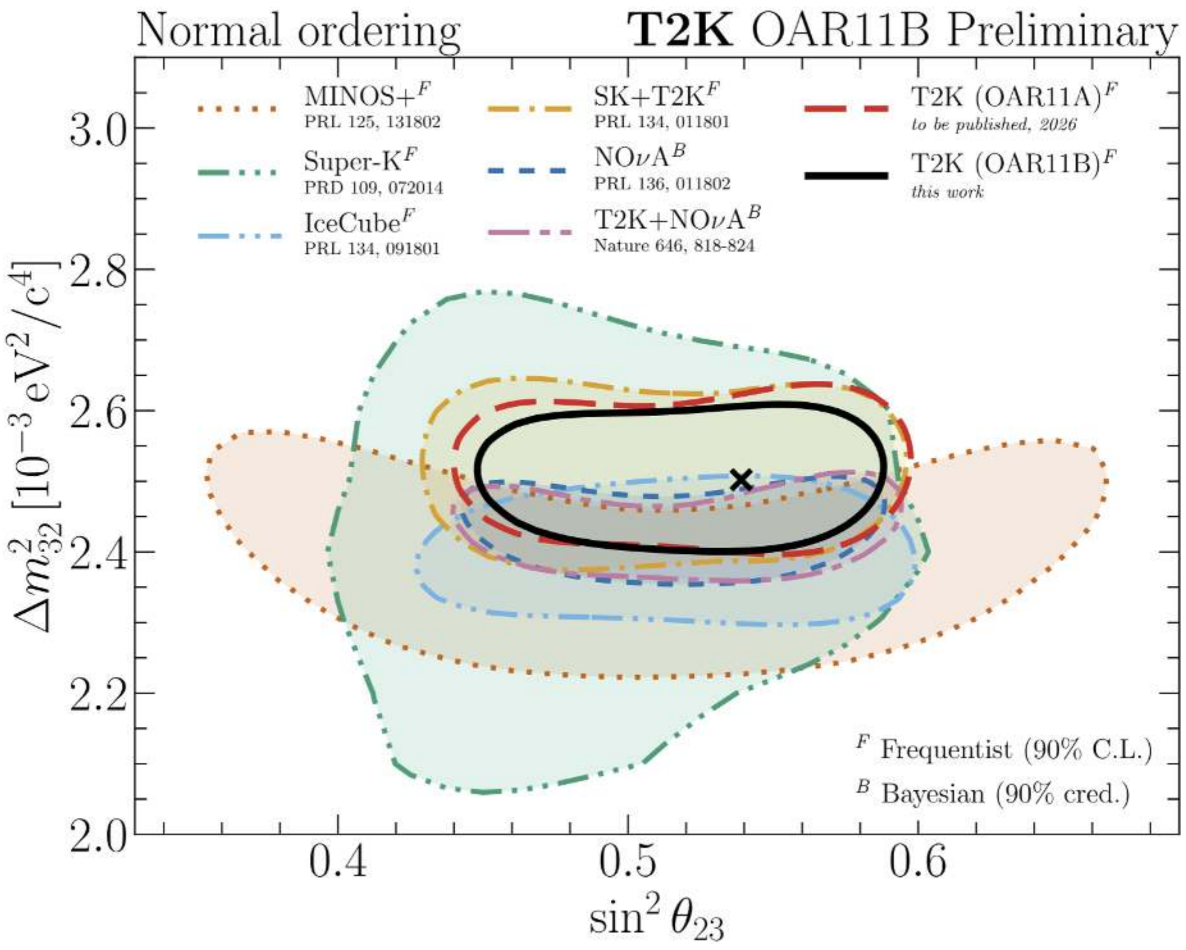

MINOS+ shown with T2K, NOvA, Super-K and IceCube.

What actually limits us

How we can improve for the next generation long baseline neutrino detectors

1. **Neutrino interaction modelling** — now dominant in the appearance channels. Nuclear effects, 2p2h, final-state interactions, and the ν_e/ν_μ cross-section ratio.
2. **Flux** — hadroproduction (NA61/SHINE) and beamline alignment.
3. **Near-to-far extrapolation** — a different target, different acceptance and a different energy dependence between the two detectors.
4. **Energy scale** — what limits Δm^2_{32} once the statistics are large.

The next generation's answer

- Near detectors with the same target and technology as the far detector, spanning off-axis angles — the far spectrum measured, not modelled. That is IWCD, and DUNE-PRISM.
- Plus an order of magnitude more events. Statistics and systematics have to come down together.

2

The next generation

Hyper-Kamiokande from 2028, DUNE from 2031 — and JUNO already running.

Hyper-Kamiokande

The detector and where construction stands

- **68 m diameter × 71 m high; 0.26 Mt total, 188 kt fiducial — 8.4 × Super-Kamiokande; ~600 m rock overburden in the Tochibora mine.**
- 20,000 × 50 cm Box & Line PMTs (≈2× Super-K photon-detection efficiency) plus ~800 multi-PMT modules; 20% photocoverage; ~3,600 outer-detector PMTs.
- 295 km from J-PARC at 2.5° off-axis — deliberately the same configuration as T2K, with the beam going to 1.3 MW.
- **Status: cavern excavation complete 31 July 2025; ~17,000 of 20,000 PMTs delivered by October 2025; IWCD groundbreaking 4 November 2025; water filling 2027–28; first data 2028.**
- ~650 researchers, 106 institutes, 23 countries and regions.

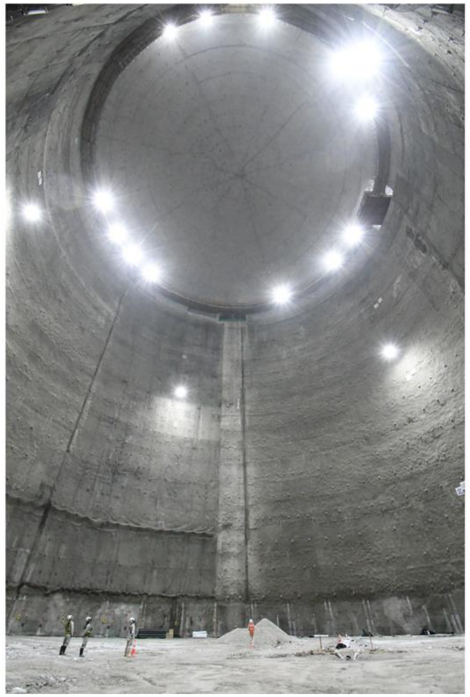

Hyper-Kamiokande, 600 m underground.

Hyper-K near detectors

Why the intermediate detector is the interesting part

- The upgraded **ND280** — already running for T2K — is inherited by Hyper-K.
- **IWCD: a water Cherenkov detector 8.8 m across and 12.7 m tall with 360 multi-PMT modules, in a 50 m deep pit ~870 m from the target.**
- **It moves vertically, spanning a range of off-axis angles.** Each position sees a different spectrum, and linear combinations of those spectra reproduce any spectrum you want — including the oscillated one at the far detector.
- Same target nucleus, same technology, same reconstruction as the far detector. That removes the interaction-model systematic rather than merely constraining it.

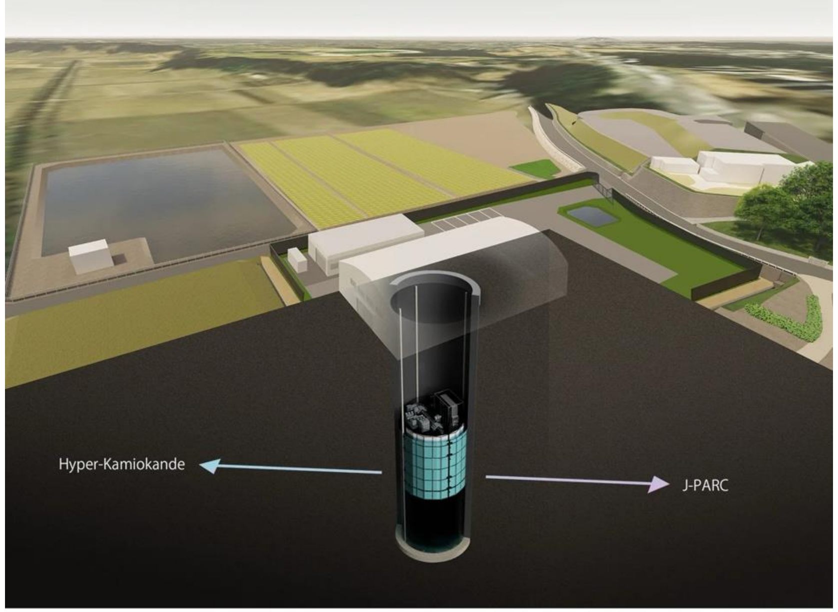

IWCD: synthesising the far-detector flux.

Hyper-K sensitivity

27×10^{21} POT over ten years, 1:3 neutrino to antineutrino

< 3 yr	**~6 yr**	**> 60%**	**6°–20°**
to 5σ CPV at δCP = −π/2, ordering known	with no external ordering constraint at all	of δCP at 5σ after 10 years — ~80% at 3σ	1σ resolution on δCP

- Expected samples: ~10,000 ν_μ, ~10,000 $\bar{\nu}_\mu$, ~2,000 ν_e and ~800 $\bar{\nu}_e$ — roughly T2K's entire-programme ν_e yield every year.
- $\sin^2\theta_{23}$ to 0.4–2%; Δm^2_{32} to ~0.4%; $\sin^2\theta_{13}$ improved by 14–23% over the reactors alone.
- **Mass ordering from the atmospheric sample: 3.8σ to 6.2σ in about six years, depending on $\sin^2\theta_{23}$.**
- Two systematic scenarios throughout: today's T2K errors, and the improved model.

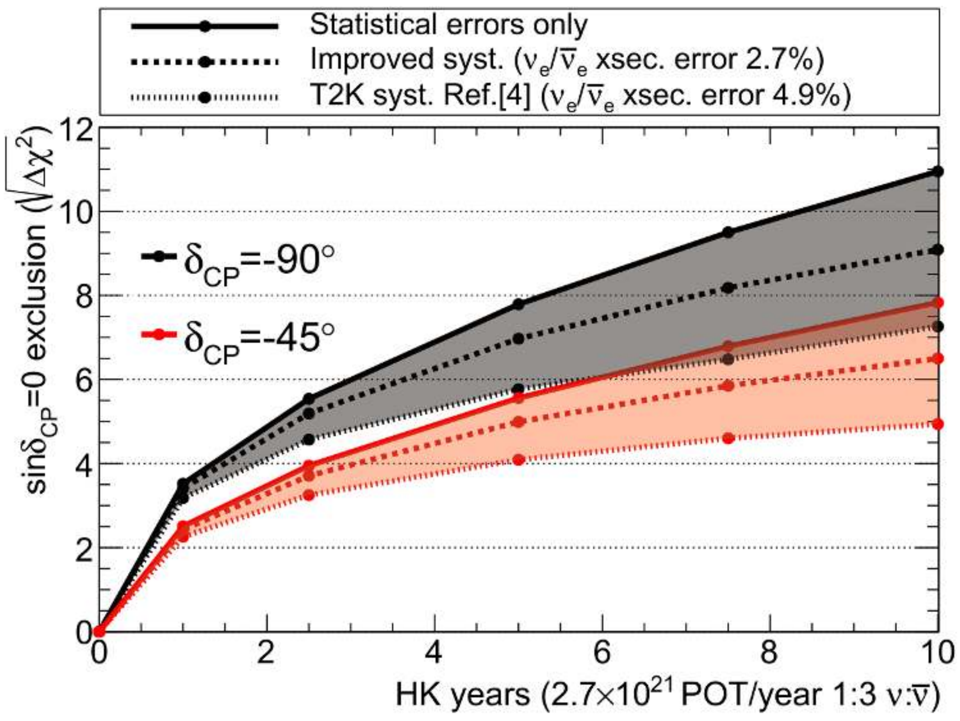

Hyper-K CP-violation discovery reach.

DUNE

1300 km, wide band, liquid argon

- Fermilab to SURF, South Dakota: 1300 km, 1478 m underground — the longest baseline in the field.
- **On-axis wide-band beam, ⟨Eν⟩ ≈ 2–3 GeV, covering the first and second oscillation maxima at once.**
- **Four far-detector caverns: three liquid-argon TPC modules of 17 kt each (>10 kt fiducial), and a fourth "module of opportunity" whose technology is deliberately not yet fixed.**
- At 1300 km the matter effect is tens of per cent — large enough that the mass ordering comes out of the beam measurement itself, but it must then be disentangled from δCP.
- Excavation of the far-site caverns is complete; the first cryostat is being assembled underground now.

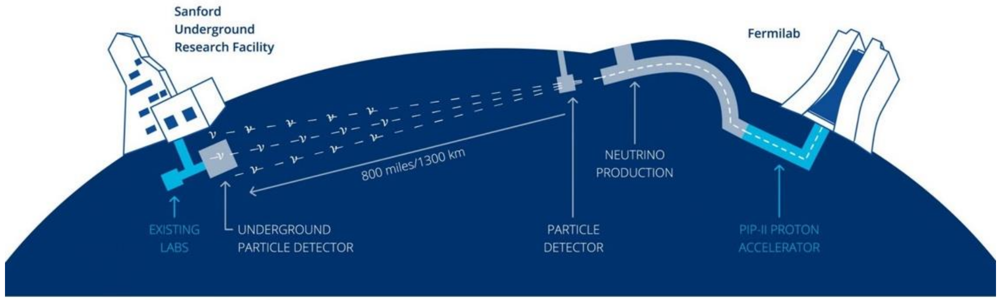

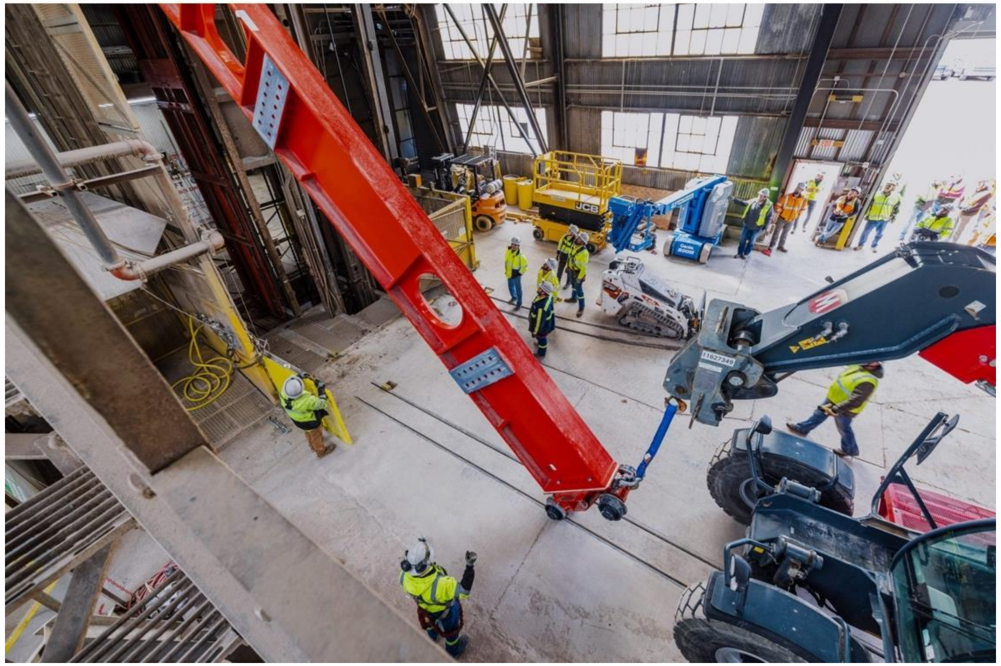

DUNE construction at SURF, 2026.

DUNE detectors and construction

Two phases, and a near detector that moves

- **Phase I:** two far-detector modules — the first with horizontal drift, the second with vertical drift — and a 1.2 MW beam.
- **Phase II:** the third and fourth modules, an upgraded near detector, and a beam above 2 MW.
- **Near detector:** ND-LAr with a temporary muon spectrometer, plus SAND on-axis to watch the beam. ND-GAr, a high-pressure gas TPC, comes in Phase II.
- **DUNE-PRISM: ND-LAr and its muon spectrometer move up to 28.5 m off-axis — part of Phase I, so available from the start of beam running; ND-GAr replaces the spectrometer in Phase II.**

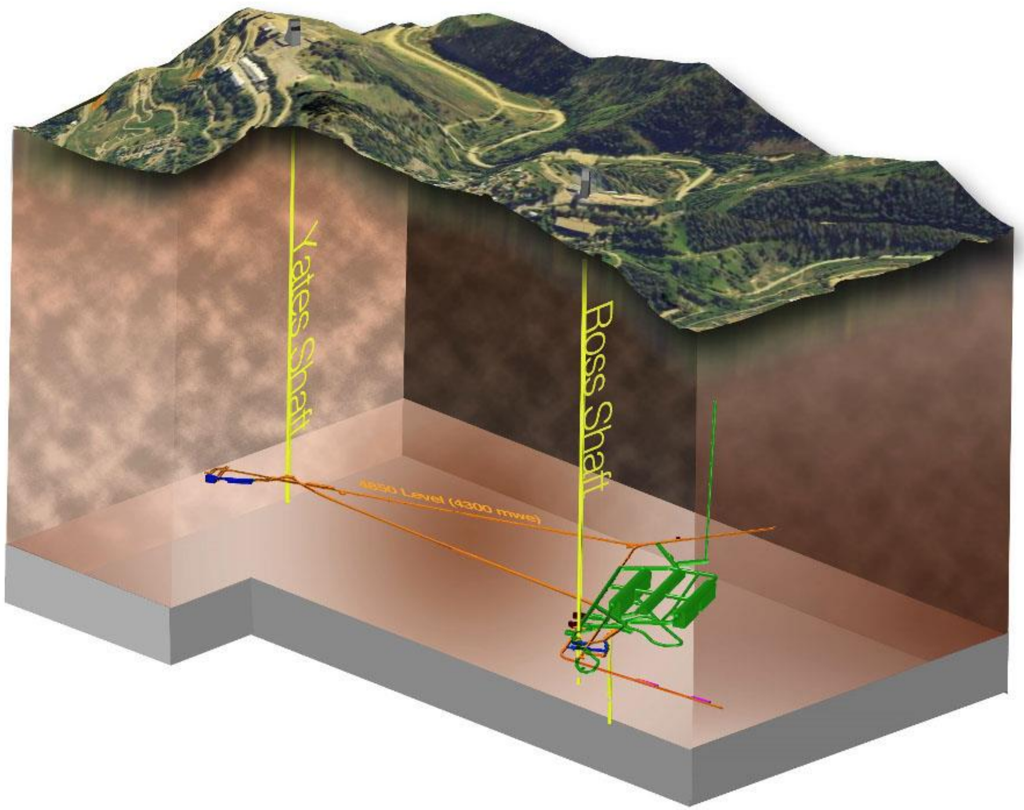

The DUNE caverns at SURF, 1478 m underground.

DUNE sensitivity

What the exposure buys — and what Phase II is for

70 kt·MW·yr
5σ mass ordering, any δCP

600 kt·MW·yr
5σ CP violation, 50% of δCP

1000 kt·MW·yr
3σ CP violation, 75% of δCP

8°–22°
1σ δCP resolution at 600 kt·MW·yr

- Phase II — modules 3 and 4, an upgraded near detector, a beam above 2 MW — gives 40 kt × 2 MW ≈ 80 kt·MW per year: 600 kt·MW·yr in six to seven years, 1000 in ten.
- Without Phase II — 20 kt at 1.2 MW, about 24 kt·MW per year — the same exposure takes 24–40 years.
- At 1000 kt·MW·yr: Δm^2_{32} to 0.4% and $\sin^2 2\theta_{13}$ to 5% — comparable to Hyper-K, from a completely different systematic direction.
- All of it assumes systematics at the few-per-cent level — the near detectors set the timescale.

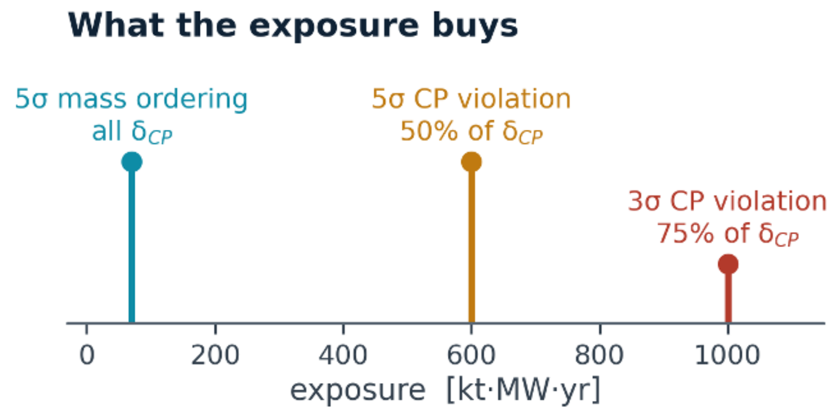

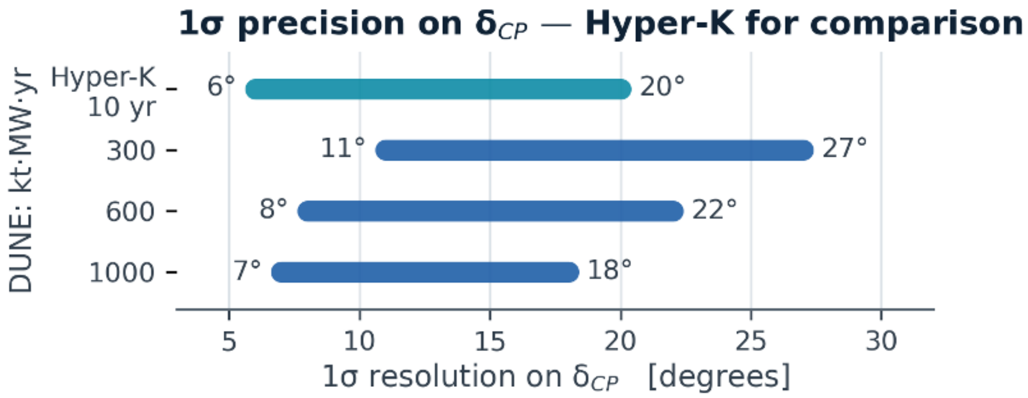

Best near $\delta_{CP} = 0$, worst near $-\pi/2$.
DUNE: ESPP 2026 input (arXiv:2503.23291).
Hyper-K: EPJC 86, 170 (2026), 10 years.

DUNE reach against exposure, with Hyper-K on the same δCP axis for comparison.

Hyper-K and DUNE side by side

Comparable reach, by very different routes

	Hyper-Kamiokande	DUNE
Technique	Water Cherenkov, 188 kt fiducial	Liquid argon TPC, >20 kt (Phase I)
Baseline · energy	295 km · 0.6 GeV, first maximum	1300 km · 2–3 GeV, first and second maxima
Beam	Narrow band, 2.5° off-axis, 1.3 MW	Wide band, on-axis, 1.2 → >2 MW
Matter effect	Few per cent — a clean δCP measurement	Tens of per cent — the ordering directly
Statistics driver	Fiducial mass	Beam power and module count
Systematics driver	Interaction model; IWCD is the answer	Argon interaction model; ND-LAr / ND-GAr
Ordering	Atmospherics, 3.8–6.2σ in ~6 yr	Beam, 5σ for all δCP at 70 kt·MW·yr
CP violation	5σ for >60% of δCP in 10 yr	5σ for 50% of δCP at 600 kt·MW·yr
First data	2028	2030 non-beam · 2031 beam

Different degeneracies, different systematics, different failure modes. The redundancy is the point — it is the only real protection against a systematic nobody has thought of.

Alongside, and afterwards

The rest of the programme that makes this work

1. **JUNO — the ordering from the reactor spectrum, and the world's best θ_{12} and Δm^2_{21}.**

2. **ESSnuSB — a proposed superbeam at the second maximum: larger CP asymmetry, far fewer events.**

3. **Atmospheric arrays** — IceCube Upgrade, KM3NeT-ORCA; and the Protvino-to-ORCA beam concept.

4. **Detector R&D** — THEIA, LLPDs, SuperFGD, high-angle TPCs, multi-PMTs, gadolinium, WbLS.

5. **The cross-section programme** — NA61/SHINE for hadroproduction and the flux; the MINERvA legacy dataset (finished 2019); ND280, SBND and IWCD taking data.

And the same detectors do more

Hyper-K and DUNE are also proton-decay detectors, supernova observatories and BSM search engines — from day one.

The field timeline

What arrives, and when

- ✓ **2025** JUNO begins data taking; first results within months
- ✓ **2026** ND280-upgrade physics data; Hyper-K detector installation; J-PARC towards 1.3 MW
- ★ **2028** Hyper-Kamiokande first physics data
- → **2030 DUNE first far-detector data — atmospherics and supernova**
- ★ **2031** DUNE first beam data
- → **early 2030s** Mass ordering at 3σ or better from several independent routes
- ★ **mid-2030s** 5σ CP violation over a large fraction of δCP
- ★ **>2030s:** Several proposed ideas for LBL experiments after HyperK and DUNE: ESSnuSB, P2O, Beta beams, Neutrino factory. ESSnuSB is currently the most active.

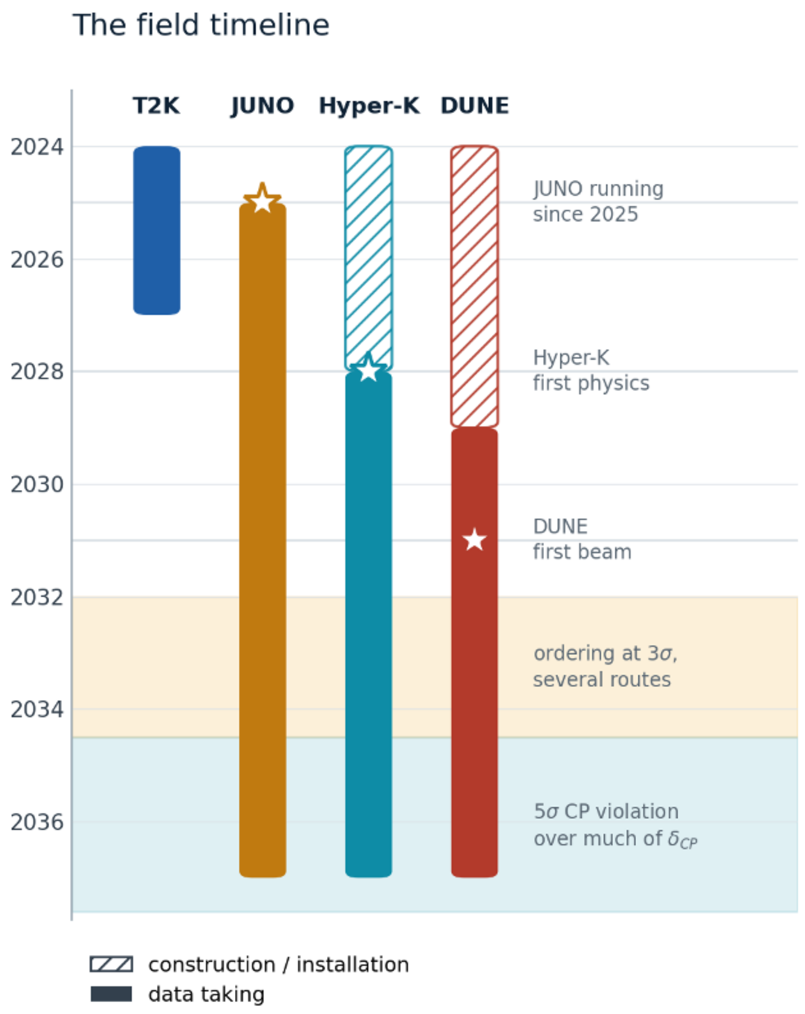

Shaded bands are projected sensitivity milestones, not scheduled dates.

Summary

Where the long-baseline programme stands

- **The current generation has taken us from "is θ_{13} non-zero" to a 95% C.L. exclusion of CP conservation and under 2% on $|\Delta m^2_{32}|$.**
- But the mass ordering is still open, and the T2K–NOvA joint fit shows how entangled it is with δCP. Neither question can be answered alone.
- We are systematics-limited as much as statistics-limited. The neutrino interaction model is the bottleneck, and the near-detector programme is the response.
- **Hyper-Kamiokande (2028) and DUNE (2031) attack the same questions with different baselines, energies and systematics; JUNO adds a third, matter-effect-free route to the ordering.**
- **The 2030s should deliver 5σ CP violation and the mass ordering and many more measurements!**

Thank you

Questions welcome.

Backup

Held in reserve for questions.

The three-flavour framework

Six parameters, and who measures which

Flavour states are superpositions of mass states through the PMNS matrix — three angles θ_{12}, θ_{13}, θ_{23}, one Dirac phase δCP, and two mass-squared splittings Δm^2_{21}, Δm^2_{32}.

If neutrinos are Majorana there are two further CP phases. They do not affect oscillations at all, which is why only six parameters appear here.

$$\begin{bmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{bmatrix} = \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix}$$

$$U = R_{23}(\theta_{23}) \cdot R_{13}(\theta_{13}, \delta_{CP}) \cdot R_{12}(\theta_{12})$$

atmospheric & beam

reactor & appearance

solar & KamLAND

for two flavours the whole thing collapses to one line:

$$P(\nu_\alpha \to \nu_\beta) = \sin^2 2\theta \sin^2\left(1.27 \frac{\Delta m^2\,[\mathrm{eV}^2]\, L\,[\mathrm{km}]}{E\,[\mathrm{GeV}]}\right)$$

Why δCP matters — and the honest caveat

Motivation, not derivation

- CP violation is one of Sakharov's three conditions for a matter–antimatter asymmetry.
- The CP violation seen in the quark sector is orders of magnitude too small to account for the baryon asymmetry.
- **But: the phase that drives leptogenesis is not in general the same phase we measure at low energy.**

$P(\nu_\mu \to \nu_e) \simeq$

$\sin^2\theta_{23}\sin^2 2\theta_{13}\sin^2\Delta_{31}$	the rate
$\mp\sin\delta_{CP}\cdot 8J_r\sin^3\Delta$	CP-odd: flips for $\bar{\nu}$
$+\cos\delta_{CP}\cdot(\mathrm{CP}-\mathrm{even})$	a nuisance
$+\cos^2\theta_{23}\sin^2 2\theta_{12}\sin^2\Delta_{21}$	solar, small

all modified by matter: $a = 2\sqrt{2}G_F n_e E_\nu$

Independently of cosmology

δCP is a fundamental parameter of nature. We have no idea why it takes the value it does — and until recently, no way to measure it.

Where each experiment sits

Baseline against energy — and why almost everyone is on the first maximum

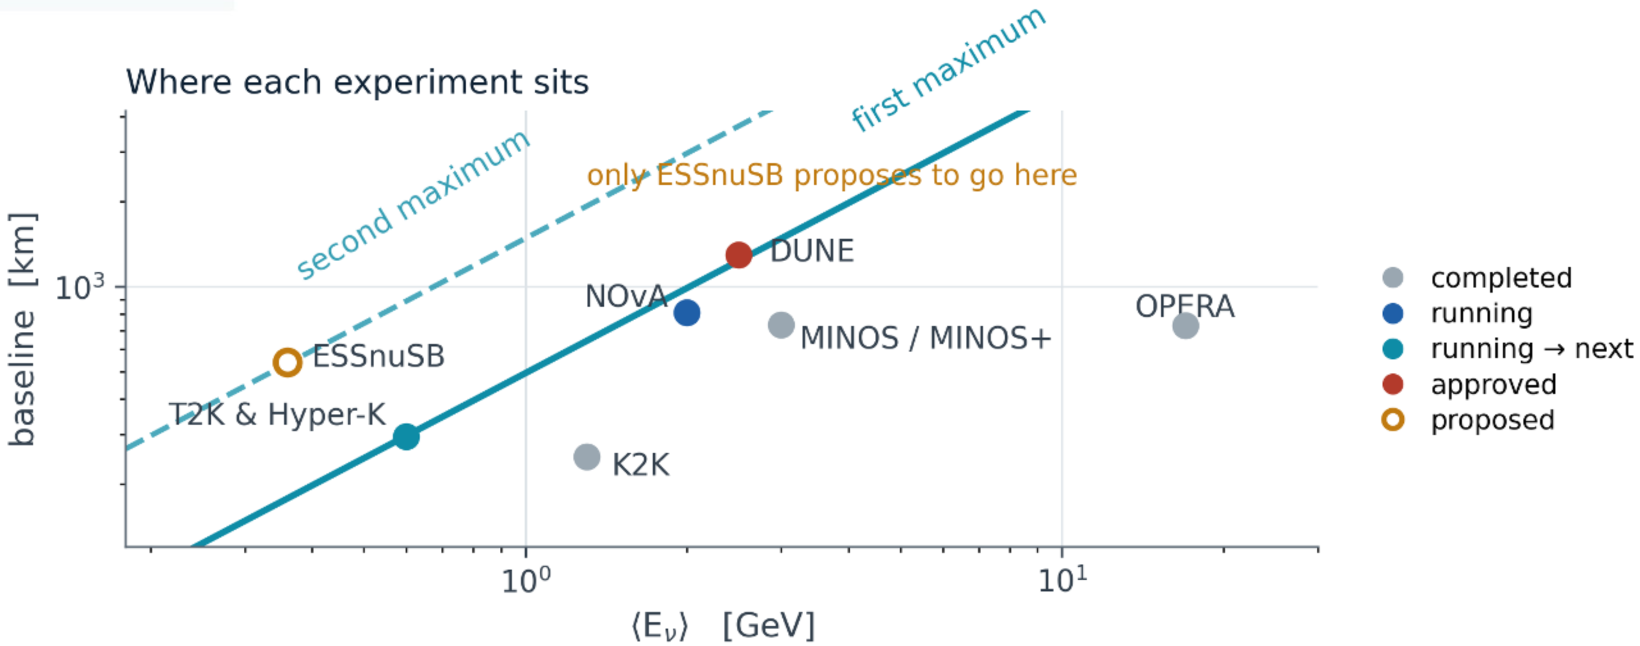

Backup — global fit parameters

NuFIT 6.1 plus the Neutrino 2026 update

Parameter	Normal ordering	Inverted ordering
$\sin^2\theta_{12}$	0.3088 (+0.0067 / −0.0066)	same
$\sin^2\theta_{13}$	0.02249 ± 0.00057	0.02261 ± 0.00056
$\sin^2\theta_{23}$	0.470 (+0.017 / −0.014)	0.555 (+0.013 / −0.016)
δCP	207° (+23 / −20)	283° (+24 / −28)
Δm^2_{21} / 10^{-5} eV²	7.537 (+0.094 / −0.10)	same
Δm^2_{31} / 10^{-3} eV²	+2.521 (+0.026 / −0.018)	−2.500 (+0.024 / −0.023)

Nuclear effects

Why the target is not a free nucleon

- At 0.6 GeV quasi-elastic, 2p2h and resonant pion production all contribute at once.
- **Eν is reconstructed from the lepton kinematics assuming quasi-elastic scattering off a nucleon at rest — anything else reconstructs at the wrong energy.**
- Final-state interactions can absorb the pion and move a resonant event into the CC0π sample. Mis-modelled absorption = shifted spectrum.
- A shifted spectrum is a biased Δm^2_{32}, and through ν vs ν̄ a biased δCP.
- None of this improves with more protons on target.

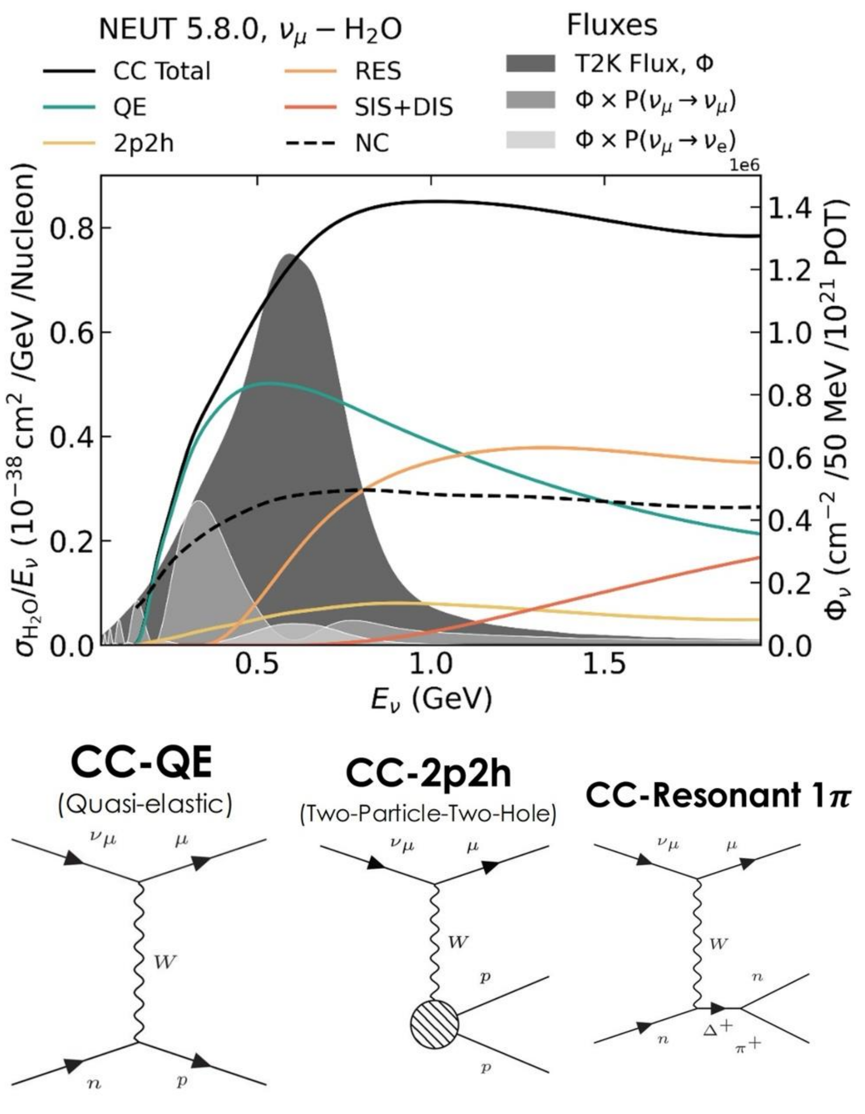

T2K flux against the channels that populate it.

What the near detector buys

Constraining the model, not removing it

- The ND280 fit constrains flux × cross-section together, cutting the far-detector rate uncertainty from 19.2% to 3.7% on the 1-ring μ sample and 20.3% to 5.7% on 1-ring e.
- **What it does not do: ND280 is hydrocarbon, Super-Kamiokande is water — different nucleus, different acceptance, different energy dependence.**
- The residual uncertainty lives in that extrapolation, and it does not shrink with statistics.
- The upgrade — SuperFGD, two high-angle TPCs, six time-of-flight planes — restores full 4π acceptance and a ~300 MeV/c proton threshold. It is not yet in the oscillation fit.

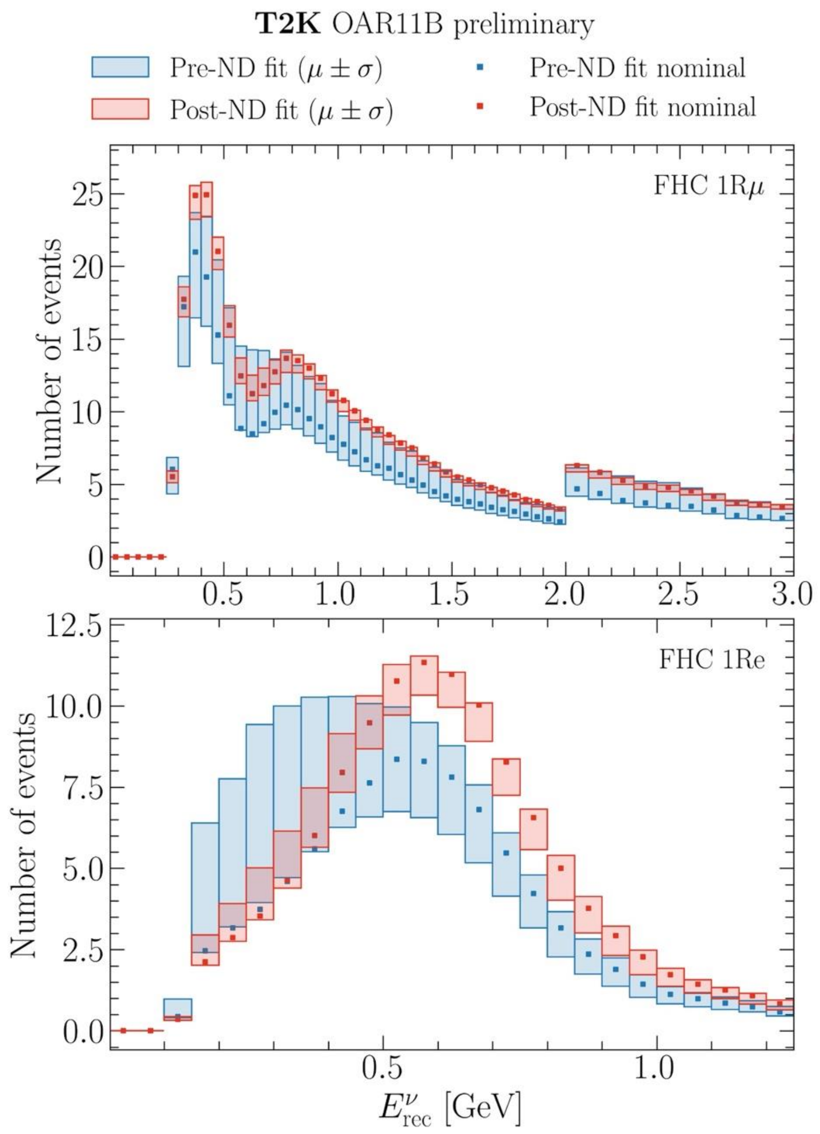

Super-K spectra before and after the ND280 constraint.

…and what it does not buy

The energy-reconstruction bias survives the constraint

- Even after the ND280 fit, the far-detector 1-ring μ sample keeps a long tail on the negative side of E_rec / E_true − 1: events reconstructed well below their true energy.
- **The constraint sharpens the peak and trims the tail. It does not remove it — the bias is built into the quasi-elastic assumption, not into the normalisation.**
- A residual bias in reconstructed energy is a residual bias on Δm^2_{32}, and it does not shrink with more protons on target.
- The fix is a near detector with the same nucleus and the same reconstruction as the far detector: IWCD for Hyper-K, ND-LAr with PRISM for DUNE.

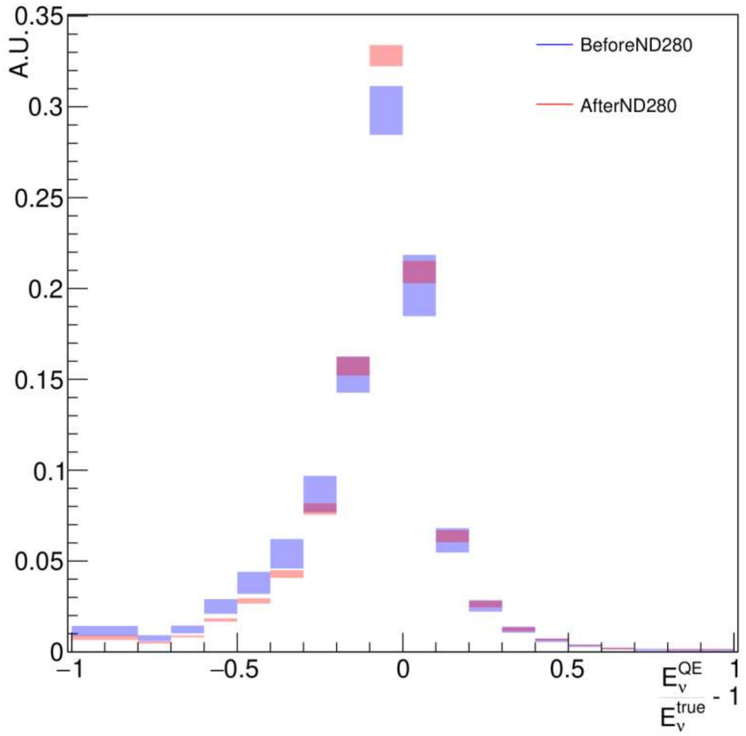

Bias in reconstructed neutrino energy, far-detector 1-ring μ sample, before and after the ND280 constraint.

The cross-section programme

The rate-limiting step for everything above

- NA61/SHINE hadroproduction data fix the flux normalisation and shape — the input every accelerator beam experiment shares.
- Dedicated T2K measurements: ν_μCC0π on carbon and on water, ν_e CCπ$^+$, NC1π$^+$, and neutron tagging in NCQE — with transverse kinematic imbalance now isolating nuclear effects directly.
- SBND and the MINERvA legacy on argon and hydrocarbon; IWCD on water.
- **Also a near detector that removes the extrapolation: same nucleus, same technology, same reconstruction.**

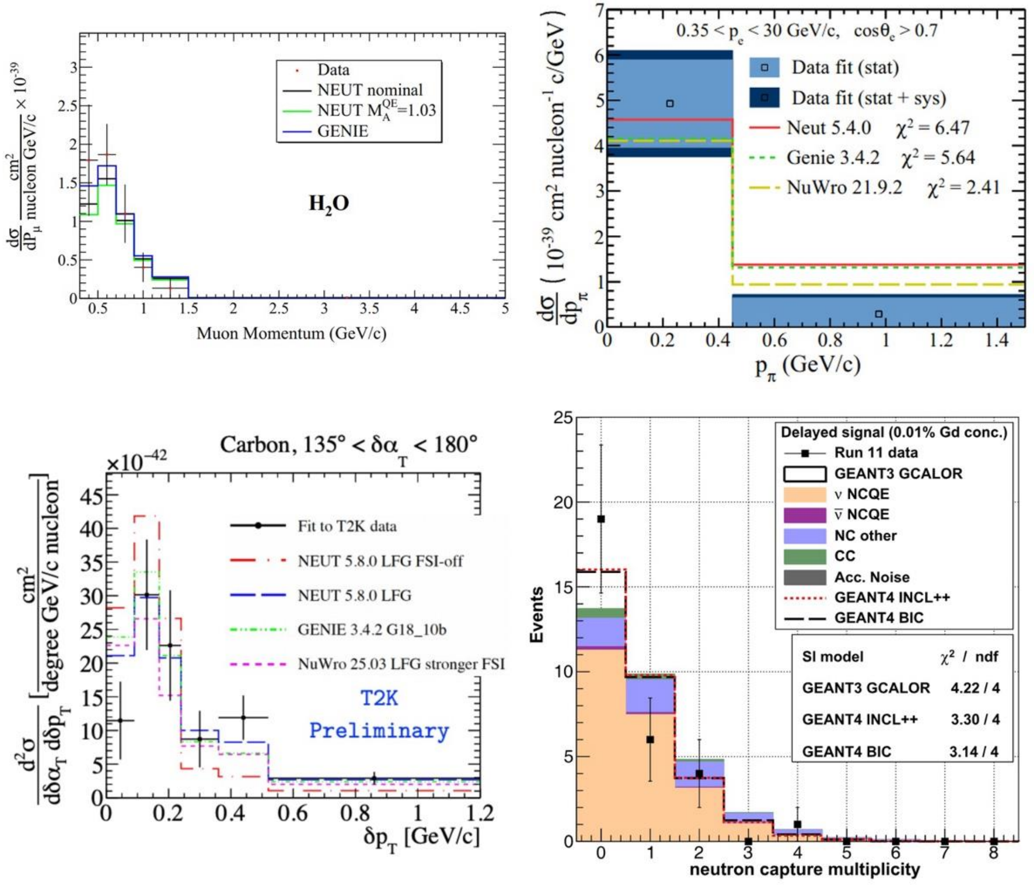

Recent T2K cross-section measurements against models.

Backup — the appearance probability

Why the off-axis angle sits where it does

- $P(\nu_\mu \to \nu_e) \approx$ a leading term in $\sin^2\theta_{23}\,\sin^2 2\theta_{13}$, a CP-odd term $\propto \sin\delta CP$, a CP-even term and a solar term — all modified by the matter potential.
- The CP-odd term peaks at the oscillation maximum and changes sign between ν and $\bar{\nu}$.
- At 295 km the matter term is a few per cent; at 1300 km it is tens of per cent. That single number is the difference between the Hyper-K and DUNE strategies.

$$P(\nu_\mu \to \nu_e) \simeq$$

$\sin^2\theta_{23}\sin^2 2\theta_{13}\sin^2\Delta_{31}$	leading term — sets the overall rate
$\mp \sin\delta_{CP} \cdot 8J_r \sin^3\Delta$	CP-odd — flips sign between ν and $\bar{\nu}$
$+ \cos\delta_{CP} \cdot (\mathrm{CP-even})$	same for ν and $\bar{\nu}$ — a nuisance here
$+ \cos^2\theta_{23}\sin^2 2\theta_{12}\sin^2\Delta_{21}$	solar term — small at 295 km

every term is modified by the matter potential $a = 2\sqrt{2}\,G_F n_e E_\nu$ — which is exactly why the ordering and δ_{CP} are hard to separate.

Bi-probability: the ordering–δ_{CP} degeneracy, and how baseline breaks it

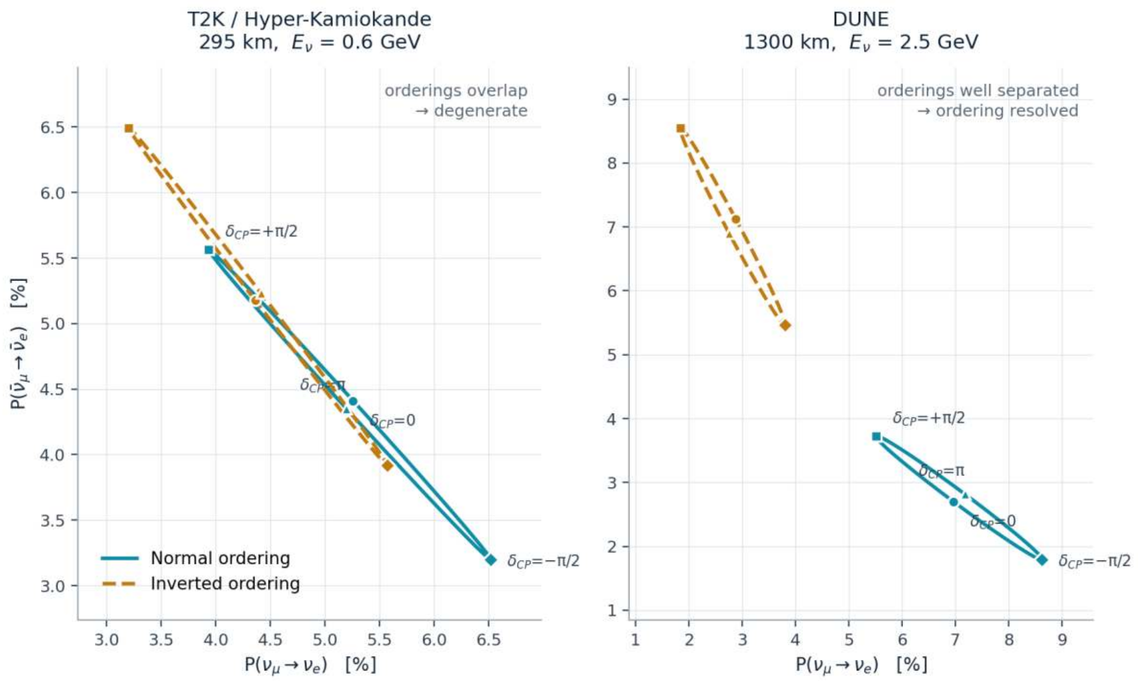

Ellipses traced by δ_{CP} over $[0, 2\pi)$. At 295 km the two orderings overlap; at 1300 km the matter effect separates them.

Backup — how these fits are done

What each result assumes

Result	Method	θ_{13} input	Priors / notes
T2K Run 1–11 v2	Frequentist (pseudo-experiments) for the CP exclusion; Bayesian for posteriors	Reactor constraint applied; also quoted without	Flat in δCP and in $\sin \delta CP$ — both shown
NOvA 10 years	Bayesian MCMC	Daya Bay 1D; the stronger ordering number uses the 2D Δm^2_{32}–$\sin^2 2\theta_{13}$	Flat in $\sin^2\theta_{23}$ and δCP; Gaussian systematics
T2K + SK atmospheric	Bayesian	Reactor constraint	Common interaction model, correlated detector systematics
T2K + NOvA	Bayesian	Reactor constraint; results also given without	Bayes factors quoted both ways
NuFIT 6.1	Frequentist global χ^2	Short-baseline reactors included	Predates JUNO; SK-atm shown separately

Frequentist: how often would data like this arise if the hypothesis were true. Bayesian: given a prior, how probable is the hypothesis. Both are quoted because the mass ordering is a discrete choice, where the frequentist σ is awkward to define.

ESSnuSB — the second-maximum option

A design study for what comes after Hyper-K and DUNE

1. **The idea — sit on the SECOND maximum: the CP asymmetry is ~3× larger and matter effects are negligible.**
2. **The beam — the 5 MW ESS linac at Lund, upgraded to 2.5 GeV and 28 Hz, with an accumulator ring to compress the pulse.**
3. **The detectors — 2 × 270 kt water Cherenkov (≈20 × Super-K) at 360 km, 1000 m deep, plus a three-part near complex.**
4. **The prize — δCP precision that separates flavour-symmetry and leptogenesis models, not just CP violation itself.**
5. **The catch — a design study, not a project: CDR 2022, pTDR ~2030, physics not before the late 2030s.**

Why it is in this talk

The only proposal attacking δCP from the second maximum — the lever if Hyper-K and DUNE leave an ambiguity.