

MausamPanchayat

From a block forecast to a panchayat cell.

A working downscaler for Ministry of Earth Sciences / India Meteorological Department problem statement 26074. It turns a 0.25 degree (about 25 km) daily forecast into a 0.05 degree (about 5 km) field over India, serves that field from a cheap CPU machine, and refreshes it every day from the latest ECMWF IFS run. This brief states the problem, the method, the measured skill, what is feasible, and what the live prototype does and does not claim.

Organisation	Ministry of Earth Sciences, India Meteorological Department
Theme	Agriculture, FoodTech and rural development
Live prototype	https://mausam-panchayat.vercel.app
Model	MausamNet-R 1.2.0-corrdiff-r, 43,121,248 parameters
Where it runs	Google Cloud e2-standard-2, asia-south1 (Mumbai)
Date of this brief	28 September 2026

The weights are a tested checkpoint. They are not retrained overnight. Each morning the machine fetches the new 00Z forecast and runs the same network on it.

1. The problem statement

IMD issues forecasts and observations at block and district scale. A Gram Panchayat is smaller than that block. Farming decisions — irrigation, spraying, harvest, covering produce — are taken for a village, not for a 25 km square. Two panchayats inside one block can differ because of the Ghats, a valley, the coast, or simply because the rain cell missed one of them.

A block bulletin that says “light to moderate rain” can be a fair average and still be the wrong instruction for a dry village inside that block, or for a wetter village on the windward slope. Agro-advisories built only on the block inherit that ambiguity. Kisan call centres and extension officers are then asked “will it rain in my village?” and have no defensible 5 km answer.

PS 26074 asks for high-resolution information inferred from low-resolution weather, for agro-meteorological advisory. The functional core used here is FR-030 to FR-034: rainfall with an IMD category, maximum and minimum temperature, humidity, wind, and cloud, for days 1 to 5, each with a central estimate and a P10–P90 band, a stated method, a confidence tier, and provenance back to the model version and the input run.

What “high resolution” means in this build

The target grid is 0.05 degree, about 5 km, from 6.025°N to 37.475°N and 68.025°E to 97.475°E (630 by 590 cells). The source grid is 0.25 degree. That is the block-to-panchayat step the problem asks for. It is a 2D map, not a 3D atmospheric model. Terrain enters as flat covariate layers: elevation anomaly, slope, aspect, and land cover. A 3D dynamical core is a different product, and it is not what this checkpoint is.

2. What is actually running

The public page at <https://mausam-panchayat.vercel.app> does not draw a canned picture. A click fetches an 8 degree by 8 degree window (160 by 160 cells at 0.05 degree) around the chosen point and runs MausamNet on the Mumbai CPU. One lead takes a few seconds the first time that window is seen. Later leads, and a repeat of the same window, reuse the result.

- The visitor clicks the India map, types a district or city, or picks a city chip. Ocean cells return an error instead of a silent zero.
- The left map is the ECMWF IFS 0.25 degree forecast, drawn as blocks. The right map is the network’s 0.05 degree central estimate (the P50).
- A pin on either weather map reads that cell from the window already fetched. It does not pretend a new forward pass happened.
- Five lead days are the IFS 00Z issue. Rain is shown in millimetres with the IMD category. Wind is shown in km/h. Cloud is a humidity diagnostic, labelled as such.
- The page states the model version, the parameter count, whether the patch was cached, and the IFS run time.

Why the two maps can look so different

They are supposed to differ. The left map is the coarse forecast. The right map is the network’s estimate of the finer field, trained to resemble CHIRPS rain and ERA5-Land temperature, not to copy IFS millimetre for millimetre. On the 00Z run of 28 September 2026, the Pune cell at 18.525°N, 73.875°E had a 25 km block of 2.3 mm and a 5 km central estimate of 0.0 mm (wet-day probability about 0.06). In that same 8 degree window the wettest block was 45.6 mm, near the coast, and the wettest 5 km cell was 5.8 mm. Most 5 km land cells were under 0.1 mm.

If both maps share a linear colour scale up to that 46 mm block, the 5 km field looks blank. That is a display failure, not a failed call. The prototype now uses a square-root rain scale capped near the 98th percentile, prints the block value and the cell value side by side, and can move the pin to the wettest 5 km cell. A pale

map with a completed status line means the central estimate is dry.

3. How the model is made

MausamNet-R is Stage A of a CorrDiff-style downscaler: a regression U-Net, not a diffusion sampler. The backbone is NVIDIA PhysicsNeMo's SongUNet. The checkpoint stores 43,121,248 parameters. Width is 96 channels, the multipliers are 1, 2, 2, 4, there are two residual blocks per stage, and attention sits on the 20 by 20 map. The patch is 160 by 160. Training used automatic mixed precision. The checkpoint step is 8,000. The algorithm id is corrdiff-v1 and the version string is 1.2.0-corrdiff-r.

Piece	Choice
Family	CorrDiff Stage A regression. Stage B diffusion was trained and not served.
Backbone	PhysicsNeMo SongUNet, 2D
Parameters	43,121,248, counted from the checkpoint
Input channels	28
Outputs	P10, P50, P90 for five variables, plus a wet-day logit at 2.5 mm
Rain transform	$\log(1 + \text{mm})$, inverted before scoring and display
Quantile head	Monotone: P50 = base + mid, gaps through a softplus
Loss	Pinball on the three quantiles, plus 0.2 times wet-day cross-entropy
Train / val / test	2016–2023 / 2024 / 2025
Hardware used to train	NVIDIA L40S

The 28 input channels

Ten coarse daily weather fields are lifted from 0.25 degree onto the 0.05 degree grid. Twelve static layers carry terrain and land cover. Then a land mask, the month's rain climatology, latitude, longitude, and the sine and cosine of the day of year. The network can use the static maps to move rain and temperature off the block average. It cannot invent a storm the coarse model did not contain.

Targets

Rain is CHIRPS on the 0.05 degree grid. Temperature, humidity, and wind targets are ERA5-Land. The live input is not those reanalyses. It is ECMWF IFS open data at 0.25 degree, the 00Z run, five lead days, fetched from the Google mirror of the open-data bucket. That domain shift is real: the network was scored against analyses and is now fed an operational forecast. The Pune example above is what that shift looks like on a day when IFS is wetter than the CHIRPS-like median the network learned.

4. Measured skill

The held-out year is 2025. Rain is scored against CHIRPS. The other variables are scored against ERA5-Land. The baseline is bilinear interpolation of the coarse field onto the fine grid. These are grid scores, not a station-by-station IMD verification. Station verification in the pilot pipeline did not have observation days to score.

Variable	MausamNet-R	Bilinear
Rain MAE	2.64 mm	3.86 mm
Rain RMSE	8.31 mm	10.14 mm
Rain correlation	0.636	—
Rain bias	−1.27 mm	—

Variable	MausamNet-R	Bilinear
CSI at 2.5 / 15.6 / 64.5 mm	0.505 / 0.343 / 0.059	—
Tmax MAE (RMSE)	0.57°C (0.75°C)	—
Tmin MAE	0.60°C	—
Relative humidity MAE	1.78 points	—
Wind MAE	0.18 m/s	—
P10–P90 coverage	about 0.80	—

Rain RMSE of 8.31 mm is the honest full-year number. It is better than bilinear interpolation and it is not below 6 mm, and it is nowhere near 0.5 mm. Daily rain at 5 km is a skewed, intermittent field. A single missed heavy cell dominates RMSE. The coarse field itself, placed on the fine grid, is already near that error: bilinear sits at 10.14 mm RMSE. Temperature is the strong result. Half a degree on daily maximum and minimum is inside the band a farm decision can use.

Critical success index falls as the threshold rises. The network is useful on the rain / no-rain question (CSI 0.505 at 2.5 mm) and weak on extremely heavy cells (CSI 0.059 at 64.5 mm). Heavy-rain alerts should not be read as a solved problem.

5. What was tried and left out of the served model

Several post-hoc ideas were measured on 2024 and scored on 2025. Only changes that improved the problem-statement scores on the held-out year were eligible to ship. These did not.

Idea	Result	Served?
Stage B diffusion, about 43.3M parameters, residual scale about 0.78	Quantile CRPS 1.988 against 1.796 for the regression. Failed the gate.	No
Forcing each 0.25 degree block mean back onto the source (rain within 10%, temperature within 0.5°C)	On a 61-day sample, rain MAE moved from 2.68 to 3.26 and CSI at the light-rain threshold fell. Temperature MAE moved from 0.56°C to 0.92°C.	No
Blending the network halfway back to the coarse field	CSI rose slightly and MAE got worse.	No
A Platt correction on the wet-day probability	2025 Brier score did not improve (0.083407 against 0.083403).	No
Cloud cover as a trained target	Not in the checkpoint. The API publishes $\text{clip}((\text{RH} - 40) / 55, 0, 1)$, method rh-diagnostic, band ± 0.15 .	Diagnostic only

The served field is the regression P50, with the P10–P90 band widened slightly by lead day for temperature, humidity, and wind. Rain quantiles are left as the network produced them. Hard block reconciliation looks responsible in a requirements list and made the scores worse, so it is not in the response.

6. Daily operation

IFS 00Z is not reliably on the open-data mirrors at 08:00 IST. The jobs run at 14:00 IST and retry at 16:15 IST. Both the CPU prototype and the GPU panchayat pipeline use the same fetch.

- The fetch tries the Google mirror first, then Azure, then AWS. AWS has returned long 503 slow-down loops, so it is the last resort.
- The CPU job writes the forecast, then asks the local API to reload. Weights stay on disk. The health check then reports the new run time, for example 2026-09-28T00:00:00.

- The GPU job writes the same file, runs inference for the demo panchayat set, and publishes through the existing pipeline. If that date already has a success stamp, the retry does nothing.
- A known library clash between eccodes and HDF5 aborts the interpreter after a successful write. The fetch and the inference entry points exit immediately after the file is in place so cron does not treat a finished run as a failure.

Deployment

Piece	Where
Browser	https://mausam-panchayat.vercel.app (Next.js on Vercel)
Model API	FastAPI on the CPU VM, port 8080. The website calls it from the server, because a browser cannot call plain HTTP from an HTTPS page.
Machine	e2-standard-2, Debian 12, asia-south1-a, 2 vCPU, 8 GB, 30 GB disk
Threads	Two, matching the machine
Cost if left on	About \$0.07 an hour, on the order of \$50 a month
Checkpoint	regression_v2 best.pt, 172,588,187 bytes

One forward pass is one 8 degree tile and one lead. A cold Pune lead has been measured around 4 to 8 seconds on this machine. That fits an interactive page. It does not, by itself, downscale every cell in India for every request. A national pass is a batch of tiles, which is what the GPU pipeline is for.

7. Feasibility

Statistical downscaling to 5 km, conditioned on terrain, is feasible on the data that exists and is feasible on the machine that is running. The evidence is not a projection. The checkpoint exists, the 2025 scores exist, and the public page calls the checkpoint.

Feasible now

- Any land point in the India domain can be downscaled interactively, five days ahead, from the morning IFS run.
- Temperature, humidity, and wind are accurate enough, on the 2025 grid scores, to sit inside an advisory.
- Rain occurrence at 2.5 mm is better than copying the block, and the error is reported rather than hidden.
- The daily refresh does not need a GPU and does not need a retraining job.
- The serving cost of the interactive path is a small always-on CPU VM plus a static frontend.

Not feasible as stated, and not claimed

- A full-year 5 km rain RMSE under 6 mm, and any target near 0.5 mm, against CHIRPS. The coarse field's own error is already in that neighbourhood. Chasing it with post-hoc corrections made MAE and the rain/no-rain score worse.
- A reliable extremely-heavy-rain detector at 64.5 mm. CSI there is 0.059.
- Treating the live IFS field as if it were the ERA5/CHIRPS pair the model trained on. The maps will disagree, and the page now says by how much.
- A 3D nested weather model per panchayat. That is a different compute budget and a different science product.

8. Viability

The interactive product is viable at hackathon and pilot scale: one CPU VM, one checkpoint, one morning download. Stopping the VM outside working hours cuts the bill. Leaving it on is still a small line item next to a GPU.

A national product — every Gram Panchayat, every morning, with a stored forecast — is viable as a batch, not as a click. India at 0.05 degree is about 370,000 land-capable cells. The GPU pipeline already publishes

a demo set of 16 panchayats. Extending that loop is an operations task: more polygons, the same checkpoint, a success stamp per day. It does not require a new architecture.

What is not yet viable is a claim of official IMD skill. There is no sustained station verification in this deployment, cloud is not a trained field, the administrative Gram Panchayat polygon layer is not the thing being served (the grid cell under the point is), and the advisory sentence is a short rule on the cell, not a crop-stage GKMS bulletin. Those are product gaps, not reasons to distrust the grid that is being served.

Cost shape

Mode	Machine	What you get
Interactive lookup	e2-standard-2, about \$0.07/hour	Any point, one 8 degree window, five days, a few seconds
Demo panchayat batch	GPU box, separate from this VM	The existing 16-panchayat pipeline, same morning file
National batch	Same GPU pattern, more tiles	Feasible later. Not what the public page does on each click.

9. Impact and benefits

The useful change is that a block number and a cell number can be different, and both are visible. On 28 September 2026 the Pune block was very light rain and the 5 km cell was dry. An advisory copied from the block would have mentioned rain. The cell advisory does not. That is the decision the problem statement is about, on a real morning file.

- An extension officer can point at a village instead of reading only the district or block line.
- The band and the confidence tier travel with the fetched point, so a sharp number is not shown as if it were certain. Coverage of the P10–P90 band on the 2025 test was about 80 percent, not 100.
- Temperature skill is the benefit that is already numerically strong: about 0.6°C mean absolute error on daily extremes.
- Rain benefit is mainly the light-rain question and a lower error than bilinear interpolation, not a guarantee of the heavy-rain total.
- The morning file updates without a scientist in the loop. The checkpoint does not silently change.

Benefits that this prototype does not deliver yet: Hindi and other languages, a PDF or SMS bulletin per panchayat, accounts, a full Local Government Directory polygon layer, soil moisture, evapotranspiration, and a crop-stage rule base covering the pilot crops. The advisory on the page is one sentence from the cell's rain, heat, and wind.

10. Technical approach, end to end

Step	What happens
1. Ingest	Download ECMWF IFS 00Z at 0.25 degree for five leads. Store one netCDF. Skip the download if that run is already stored.
2. Features	Lift the coarse fields onto the 0.05 degree grid. Stack terrain, land cover, climatology, coordinates, and day of year. 28 channels.
3. Infer	SongUNet on a 160 by 160 tile. Emit monotone quantiles and a wet logit. Mask the ocean.
4. Derived cloud	From relative humidity only. Not a fifth trained target.
5. Serve	FastAPI returns the point, the block value, and both maps for the window. The browser never holds the checkpoint.
6. Show	Shared square-root rain scale, a pin, and an explicit block-versus-cell table so a dry fine field is readable.
7. Advise	IMD rainfall category from the millimetres, plus a short farm sentence. Not an official warning.

Provenance on a response is the model version, the parameter count, the device, the IFS run timestamp, and the attribution line. Confidence for the trained variables uses the width of the P10–P90 band and

decays with lead day. The cloud diagnostic is capped so it cannot display as if it were a trained field.

11. Coverage of the problem statement

Requirement	Status in this prototype
FR-030 variables, D+1 to D+5, central value and P10–P90	Served for rain, tmax, tmin, humidity, and wind. Cloud is present as a diagnostic. IMD rain categories are the published thresholds, including very light from 0.1 mm.
FR-031 method stack	The served method is the terrain-conditioned U-Net. Historical orographic regressions, a tree-model residual, and block reconciliation were not the winners on the held-out scores. Reconciliation was measured and rejected.
FR-032 physical guards	Rain is held at or above zero, humidity is bounded, and minimum temperature is kept at or below maximum. The block-mean tolerance is implemented and tested, and it is not applied on the served path because it hurt MAE and CSI.
FR-033 confidence	A score and a high / medium / low tier are returned for the fetched point. The formula is the band width and the lead, not a rolling station-skill map.
FR-034 provenance	Model version, parameter count, source run, and attribution are on every response and on the page.
Low-resolution and high-resolution plots	Both are on the page, for the same window and the same colour scale, with the pin shared.
3D visualisation	Not built. The model is 2D. Terrain is a covariate, not a volume.
Every Gram Panchayat polygon	Not built. The page serves the 5 km cell under a point. Sixteen demo panchayats exist on the separate GPU pipeline.
Station verification	The module exists. This deployment has not been scoring a live station feed.
Hindi farmer dashboard, PDF bulletin, accounts	Not in this prototype. This PDF is the project brief, not a village bulletin.

The system to judge is the one that is on, not the one in the earlier design notes. It is a 43,121,248-parameter 2D regression, refreshed every morning from IFS, better than bilinear interpolation on the 2025 test, honest about an 8.31 mm rain RMSE, and visible as two maps and a pair of numbers at the pin.

Prototype: <https://mausam-panchayat.vercel.app>. Checkpoint version 1.2.0-corrdiff-r. Brief date 28 September 2026.