

Seed · Seoul, South Korea

Tinyfold

Compiles and compresses AI models to run on the neural chips inside appliances, cars and cameras



Tinyfold

What's broken today


Consumer electronics makers want speech, vision and small language models running on device, but each product line ships with a different neural chip and every vendor toolchain behaves differently. One appliance maker told us that porting a wake-word model and a vision model to three chip families took six engineers four months, and accuracy still dropped after quantisation.



Our solution

Tinyfold Bench

Tinyfold takes a trained model, searches quantisation and pruning settings layer by layer, and compiles it for the target chip, checking accuracy against the customer's own test set at each step. Engineers see size, latency, power and accuracy side by side for 14 supported chips and export a ready-to-flash package in hours instead of months.



Layer sensitivity

 Tinyfold Bench

Projects

Models

Targets

Benchmarks

Exports

Layers

142

Sensitive layers

9

Search runs

64

Accuracy budget

1.0 pts

Search models...

Apply suggested precision

Layer sensitivity

fridge-vision-v3 on NPU-A, int8 baseline

Most sensitive layers

Layer	Type	Precision	Accuracy hit
head.cls	Linear	int8	-0.31 pts
stage4.attn.2	Attention	int8	-0.22 pts
stage4.conv.5	Conv 3x3	int8	-0.14 pts
stage3.attn.1	Attention	int8	-0.09 pts
stem.conv.0	Conv 7x7	int8	-0.07 pts

Showing 5 of 73

View all →

Insight

Keeping head.cls and stage4.attn.2 in fp16 recovers 0.5 pts for 0.8 MB of extra size.

 Tinyfold

HS

Han Seo-yeon

Why now

The timing

Every new appliance and car platform now ships with a neural chip, but vendor toolchains are fragmented and makers cannot afford a porting team per chip.

Market

Device makers in Korea, Japan, Taiwan and China shipping over 1 billion products a year with neural chips; on-device AI tooling spend is projected to approach US\$2B by 2029.



Tinyfold

How we make money

Annual licence for engineering teams starting at US\$60K, plus a per-device royalty of a few cents once a compiled model ships in volume.

Traction

Where we are

6

Paying device makers

US\$41K

Monthly revenue

14

Chip targets supported

4.1x

Median size reduction

- 6 paying device makers, US\$41K monthly revenue
- First compiled model shipping in a mid-range robot vacuum line
- 14 chip targets supported, up from 4 a year ago



Tinyfold

Team

Who's building it

HS

Han Seo-yeon

CEO and Co-founder

PJ

Park Ji-hoon

CTO and Co-founder

CM

Choi Min-jae

Compiler Lead



Tinyfold

The ask

What we're raising

Raising US\$4M seed to reach 20 device makers and first royalties by 2027

- Support 10 more chip families, including automotive
- Hire four compiler engineers in Seoul
- Open a field engineering office in Taipei



Tinyfold