



Adobe

ReImagining Photoshop - The AI Editor of 2030

Team 65

Meet Doraemon

Time traveler from 2030

Generous !!! eager to help



Has access to advanced apps

Can Make anyone time travel

Meet Nobita

Aspiring content creator in 2025

Editing feels too complex



Joined editing contest

Doubts his future



Why is he sad !!!



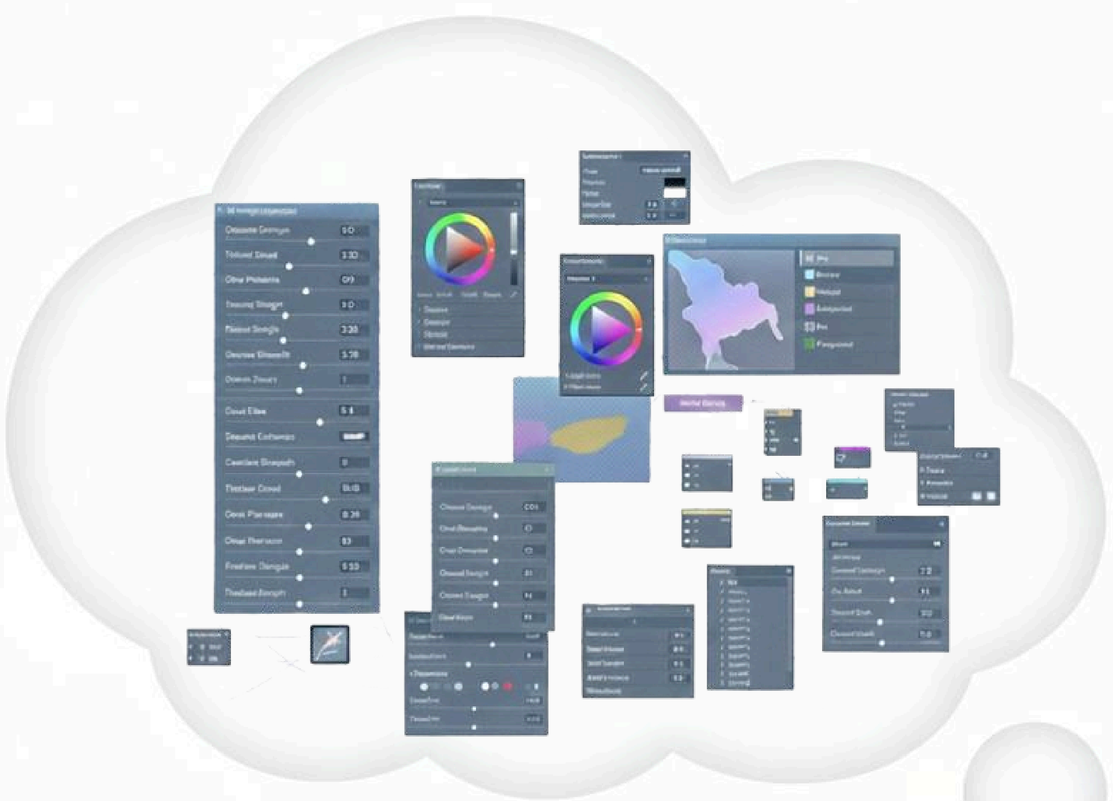
Sobbing.....

!?



Why sad buddy?





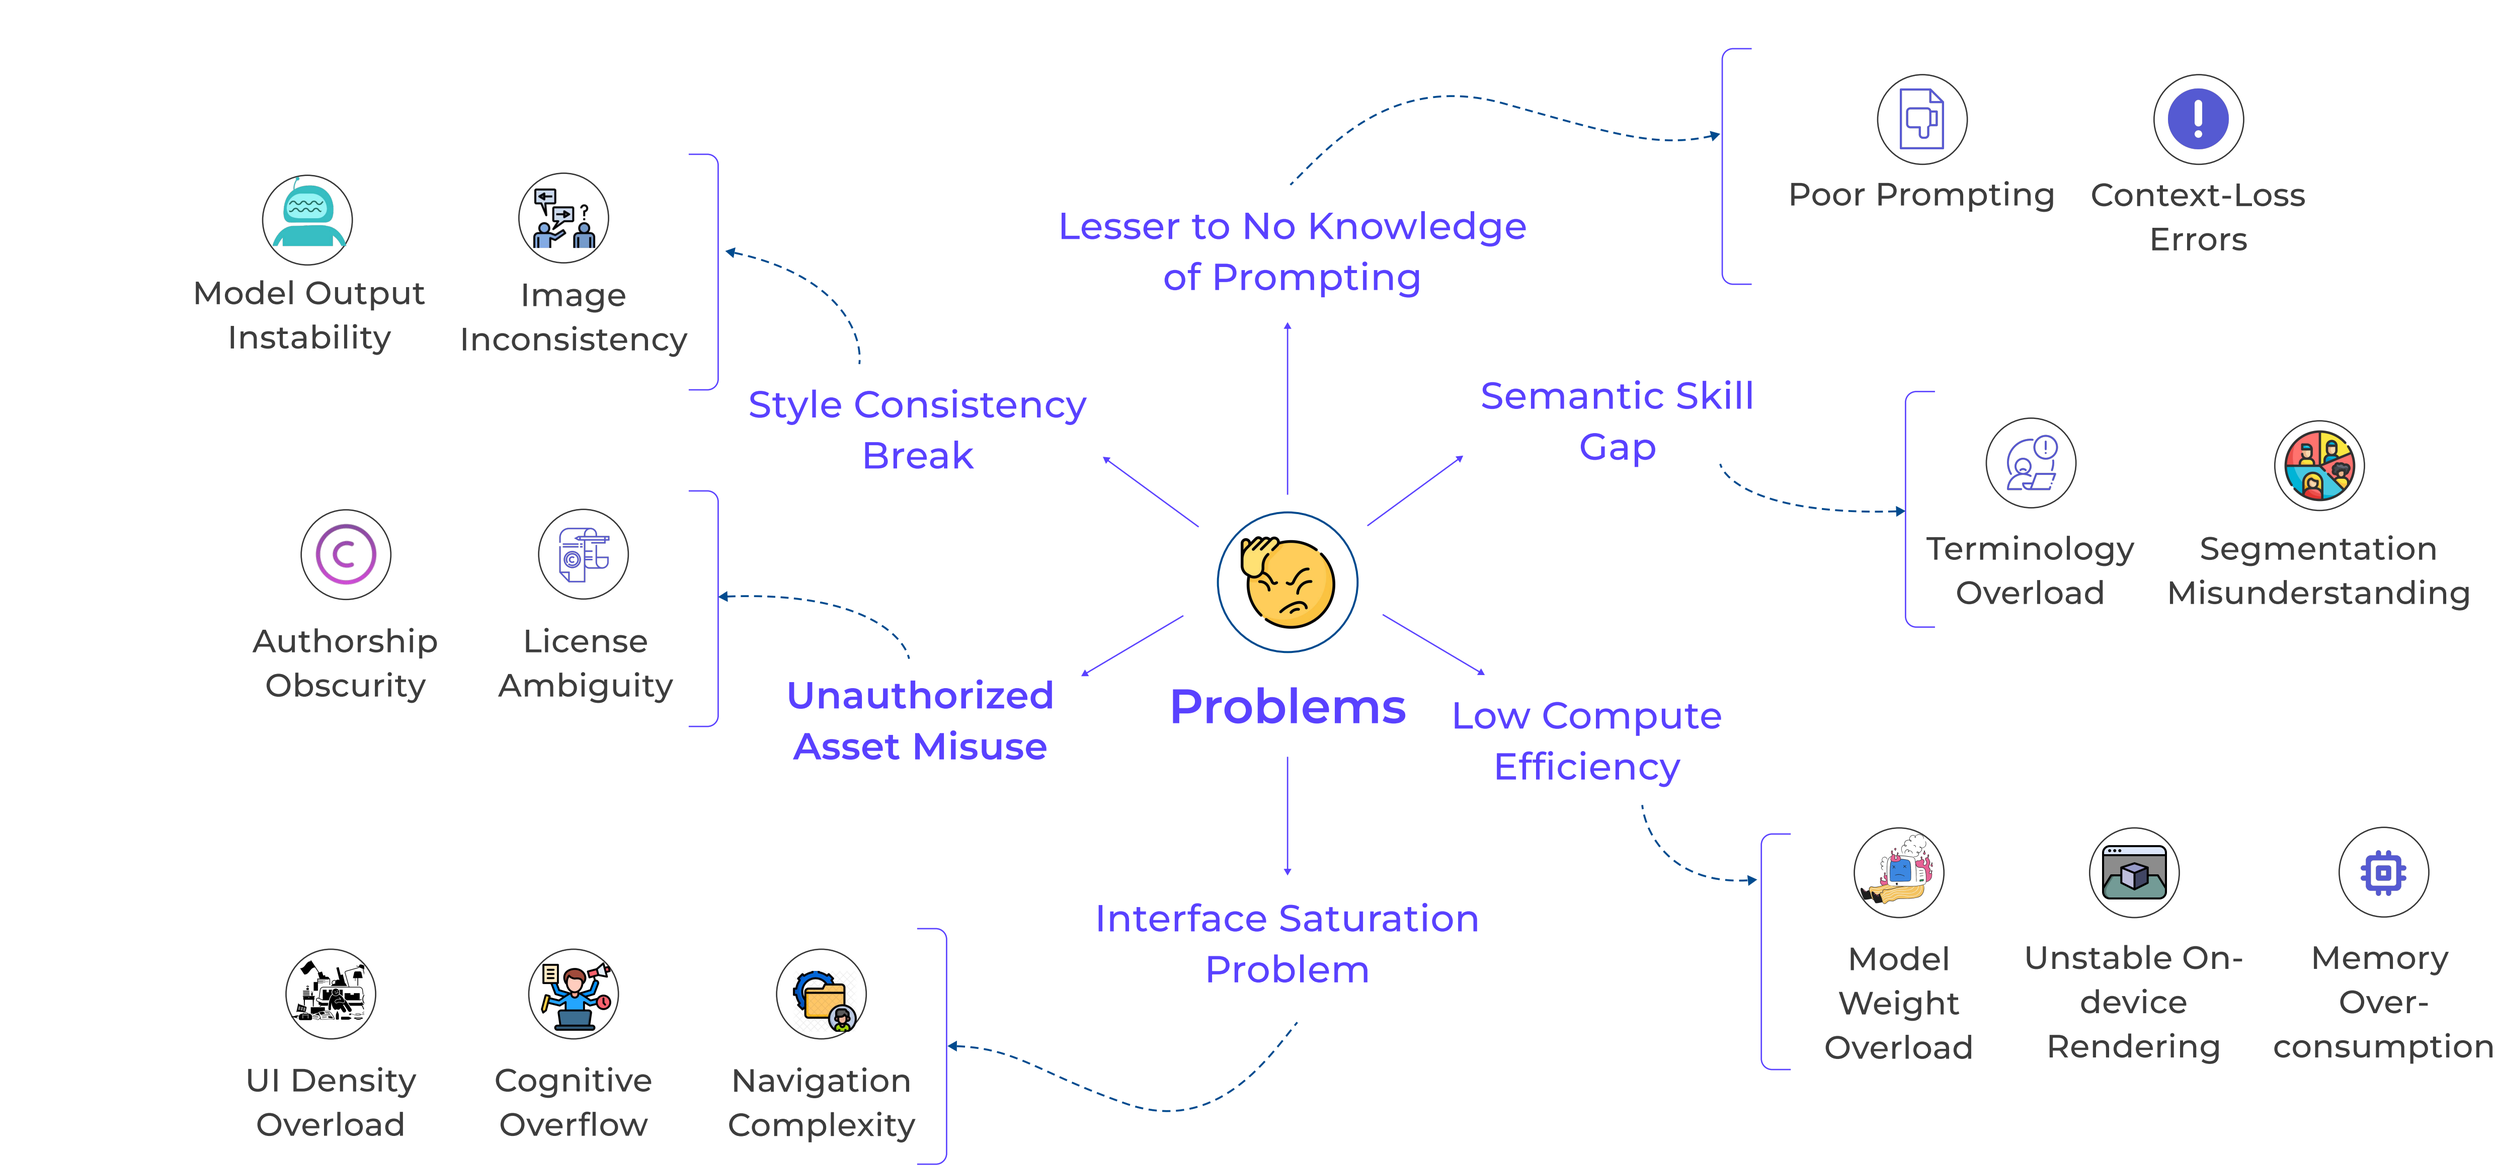
Editing is... too hard.
So many sliders...
So many terms...
I don't get anything.





Atleast you tried.
That's good.

Identifying the Current Problem



The worst part is.....

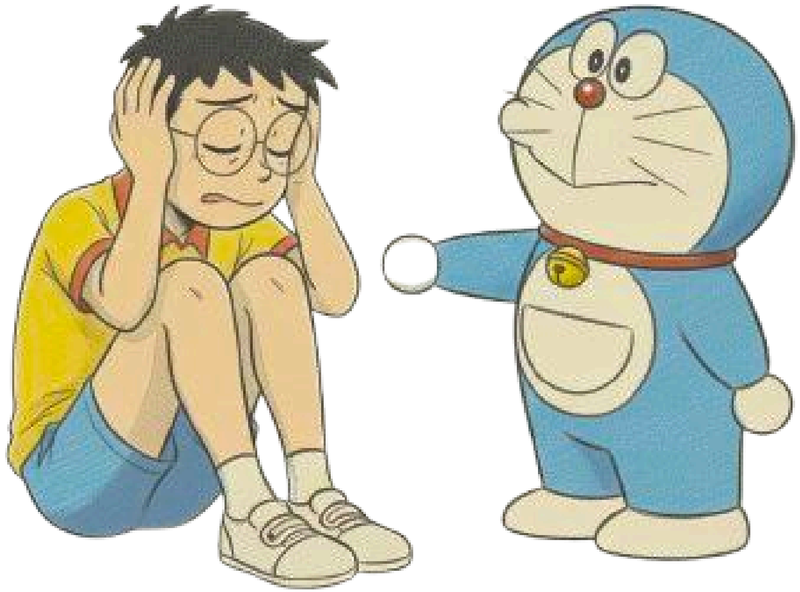
My image got pirated...

I was disqualified ! ! !

Oh no...

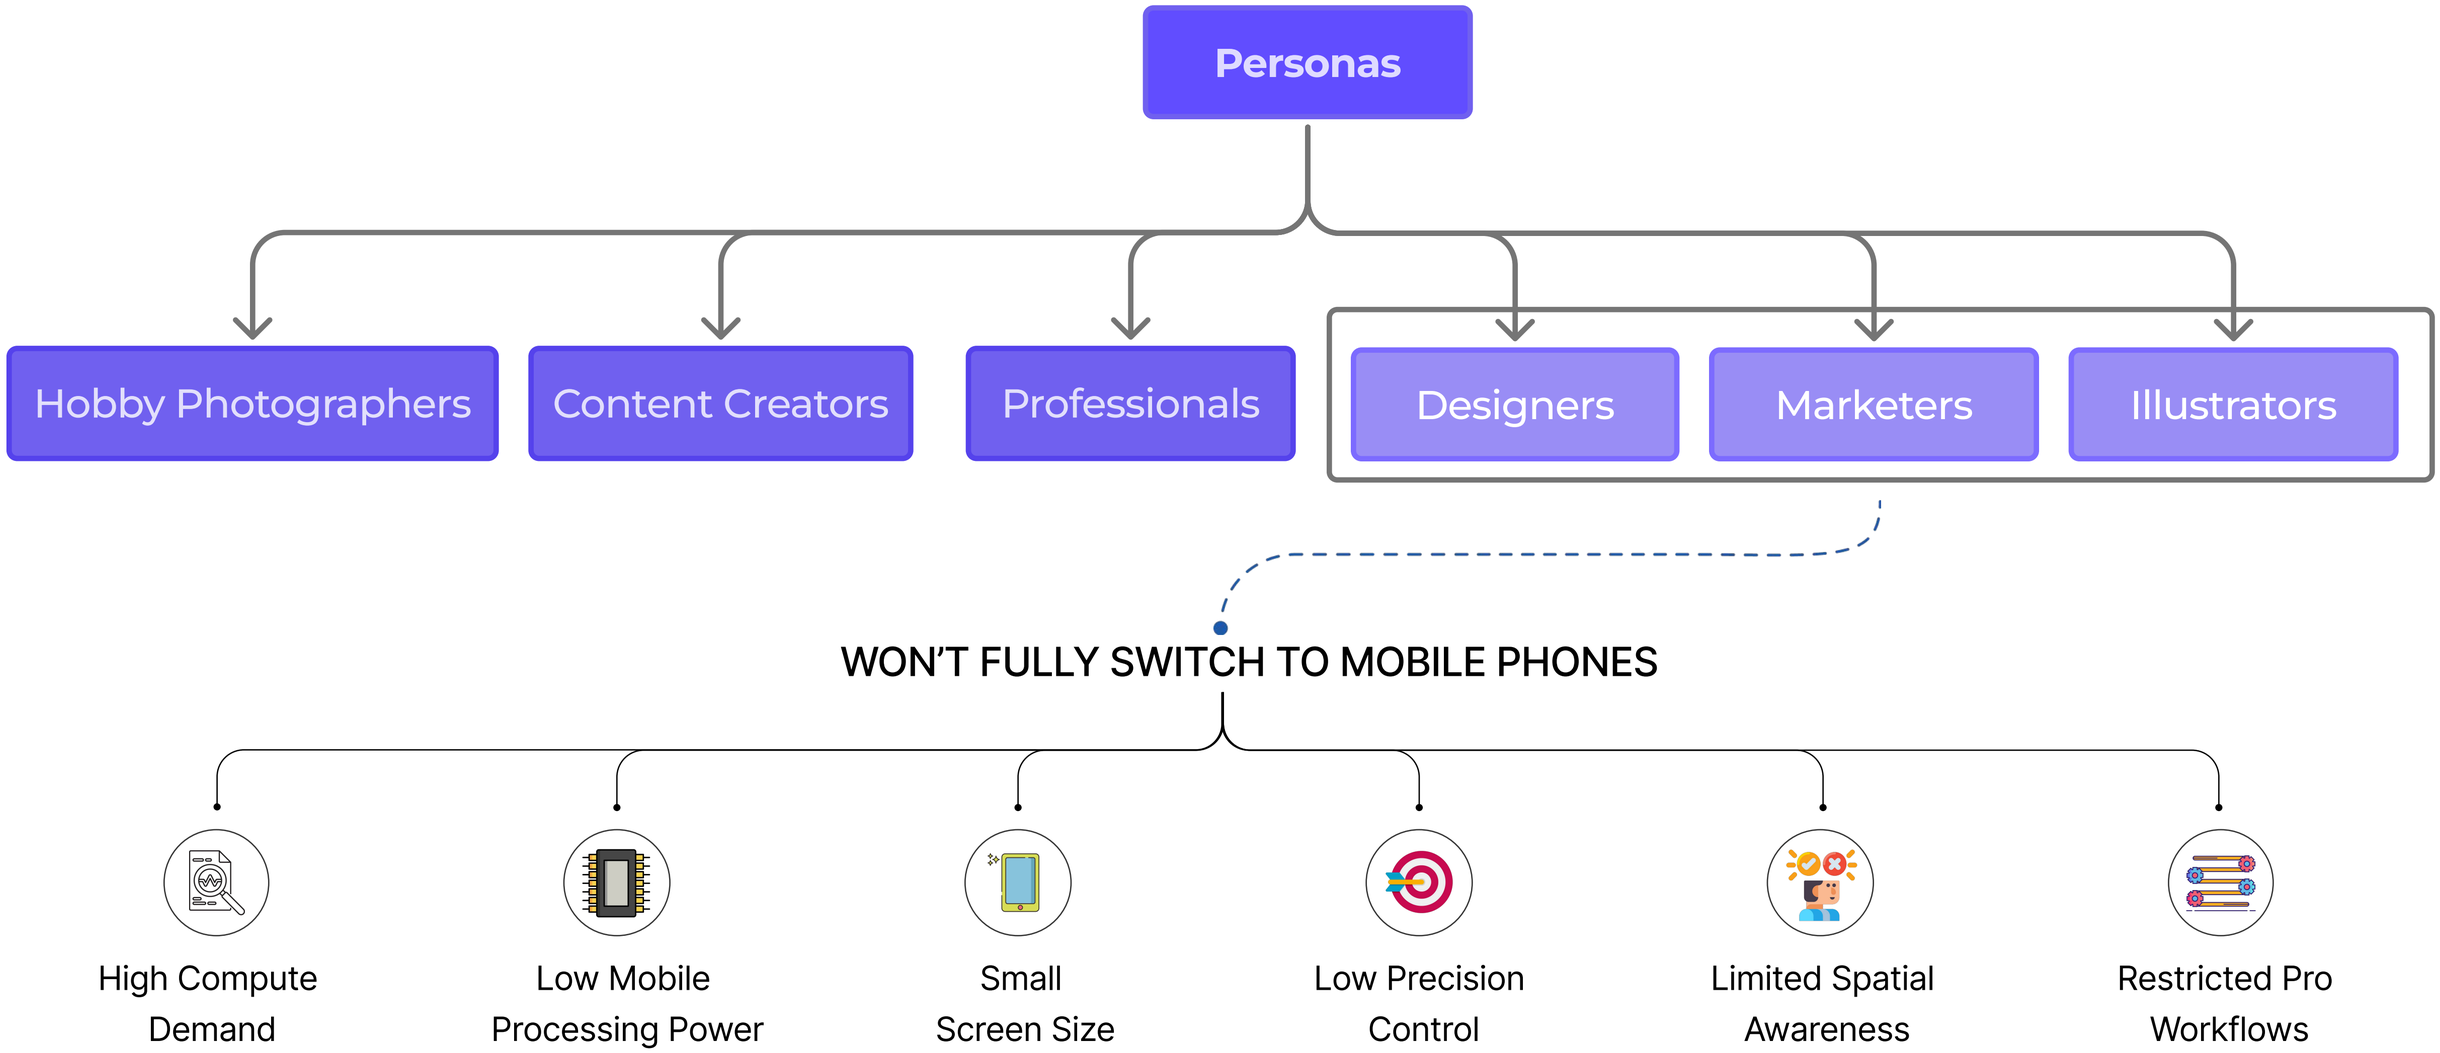
My friends struggled too....

The apps overwhelm everyone....

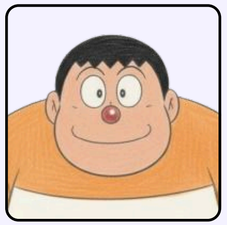


Who are going to be the priority users?

Out of the 6 selected user personas, we have dropped 3 segments.



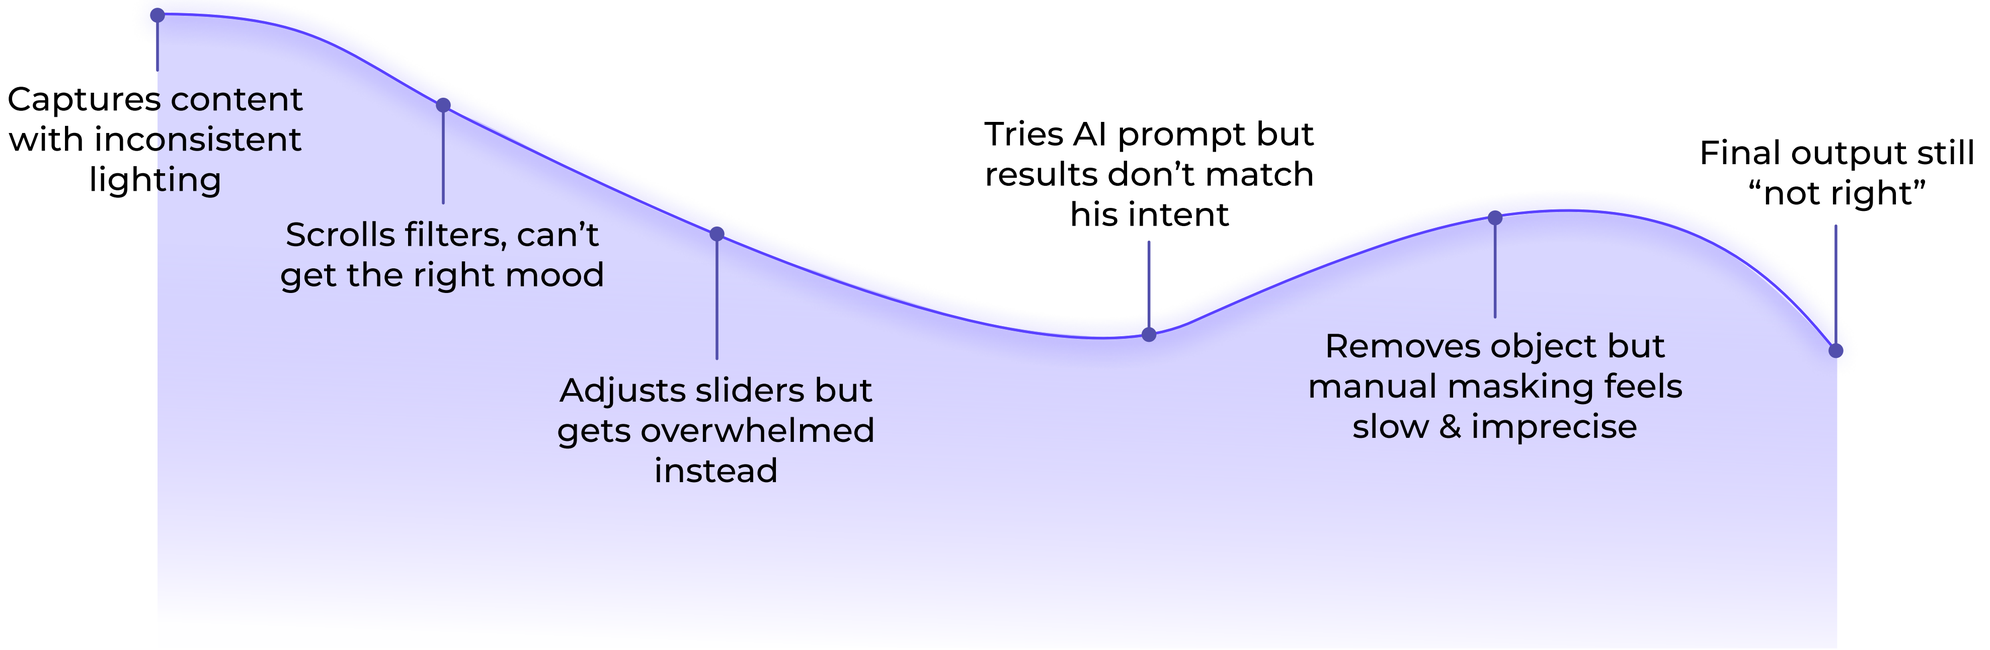
User Persona - Casual Content Creators



Giyan

17 Years, Male
Content Creator

"Editing feels like guesswork. I don't know which changes actually improve my photo."



Observations	Creators can't express themselves in technical controls	AI outputs vary widely, causing repeated trial-and-error	Selecting regions takes time & breaks creative flow	Users don't know which adjustments help or harm the final look
Analysis	Editing Needs to Adapt to Natural Language & Emotion	AI Should Guide Users Instead of Leaving Them to Experiment	Fast, intuitive area selection reducing friction	Creators need repeatable, predictable results aligning with their intended style

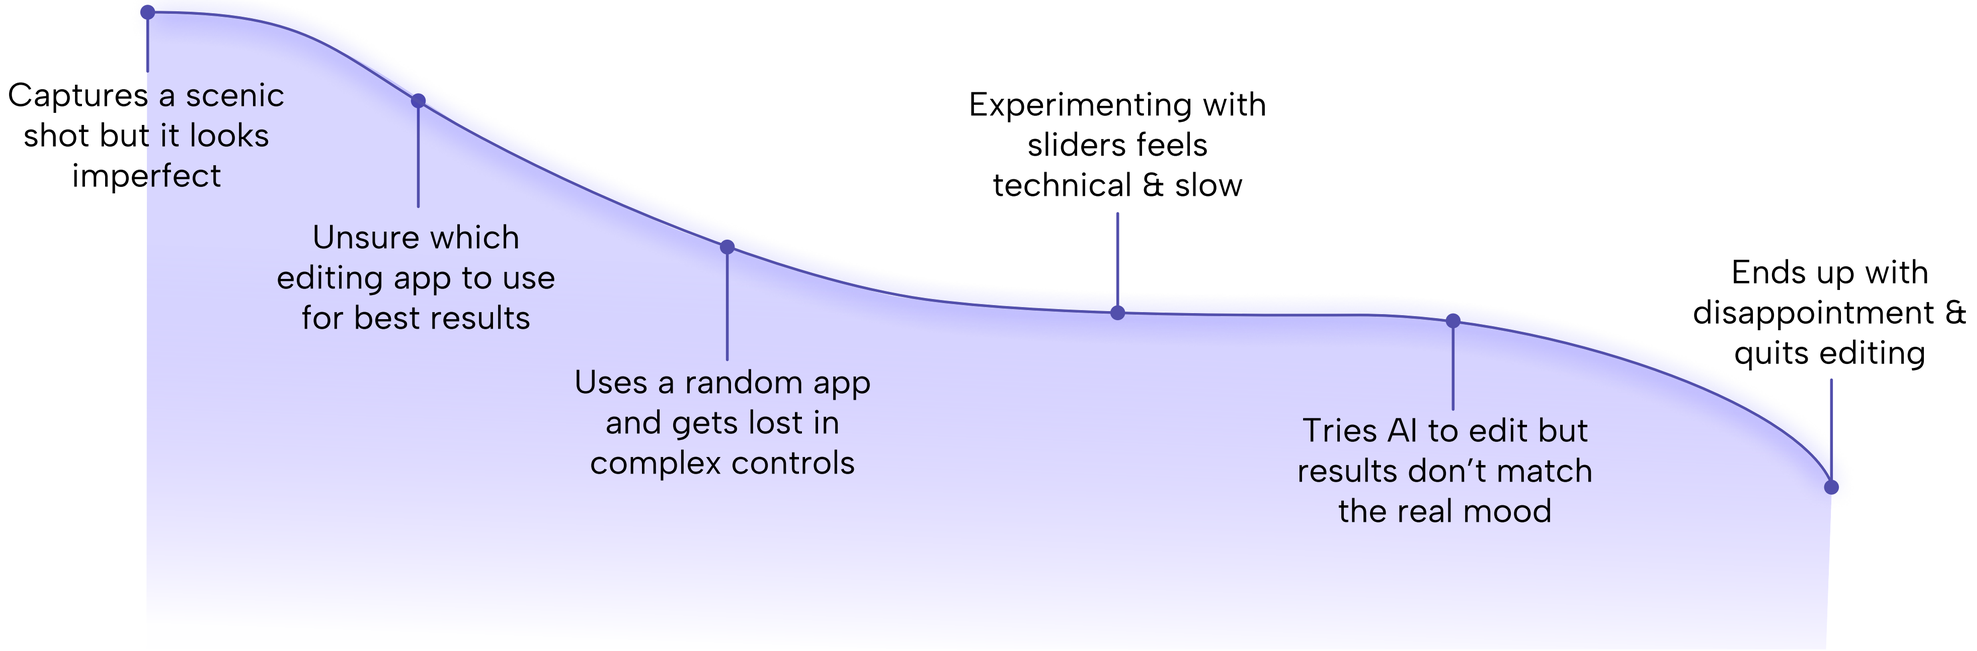
User Persona - Hobby Photographers



Dekisugi

16 Years, Male
Hobby Photographer

"I just want my edits to look natural and true to the moment, not artificial or overdone."



Observations	Can't recreate the original look while editing	Too many tools, none clearly tailored to their needs	Sliders and controls require knowledge they don't have	Outputs distort colours, mood, or depth they captured
Analysis	Hobby photographers want guidance, not complexity	Tools should interpret the scene's mood, not force technical steps	Edits should feel natural and intuitive, not technical	AI must enhance realism, not overwrite it by destroying the details

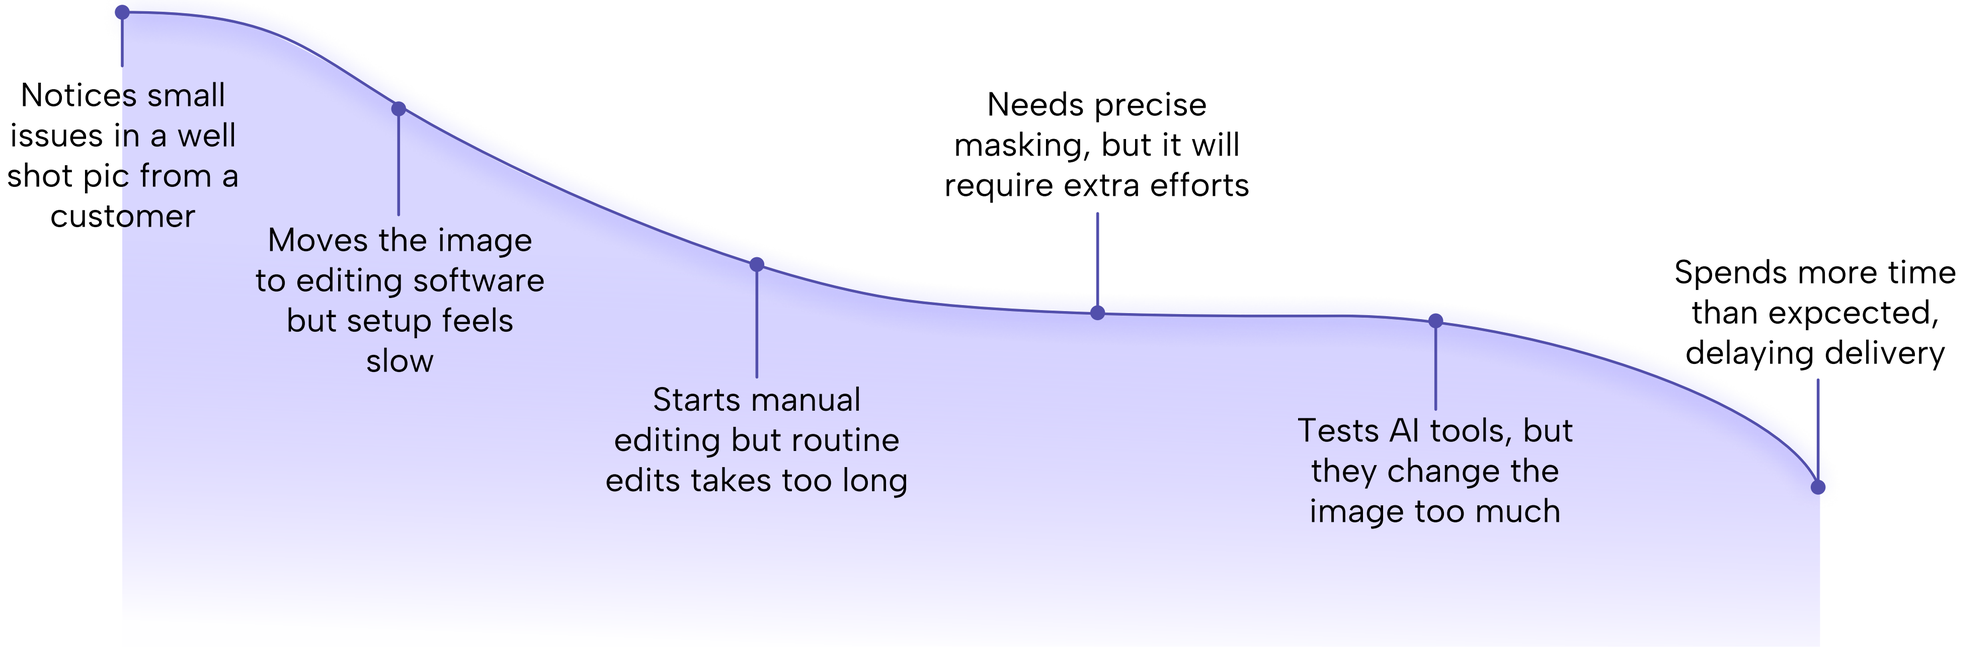
User Persona - Professionals



Jaiko

19 Years, Female
Professional Editor

"AI helps sometimes, but when it changes my style, fixing it takes even longer than manual editing."



Observations	Want speed without loosing control	Precision still needs manual efforts	Higher workload increases editing fatigue	AI feels unpredictable, often changes their style
Analysis	Needs faster workflow with manual refinements	A feature allowing quick, accurate region selection could be provided	Tools that cut repetitive steps and make the overall workflow lighter	AI that respects their style, understand their intent and allows manual control

Prioritizing the Users

Prioritizing personas using a Multi-Metric Matrix

Metrics/Persona	Professionals	Casual Creators	Hobby Photog.	
Editing Frequency (15)	15	12	9	Behavior Metrics
Switching Need (15)	6	12	9	
Market Size (15)	9	9	12	Market Metrics
Market Growth (15)	3	15	12	
Workflow Speed (15)	15	9	3	Capability Metrics
AI Adaptability (15)	3	15	9	

Prioritizing the Users - Scoring



Metrics/Persona	Professionals	Casual Creators	Hobby Photographer
Behavior Metrics (35%)	7.35	8.4	6.3
Market Metrics (30%)	3.6	7.2	7.2
Capability Metrics (35%)	6.3	8.4	4.2
Total Score	17.25	**24**	17.7

Weightage Rationale

Behavior Metrics

Frequency + Need to switch

- Real user behavior is the strongest predictor of actual adoption
- Personas who edit often and feel more pain will shift to mobile faster

Market Metrics

Market Size + Growth

- Size & Growth matter, but they are secondary to actual intent and usability fit
- A large market is irrelevant if users aren't motivated or capable of adopting the workflow

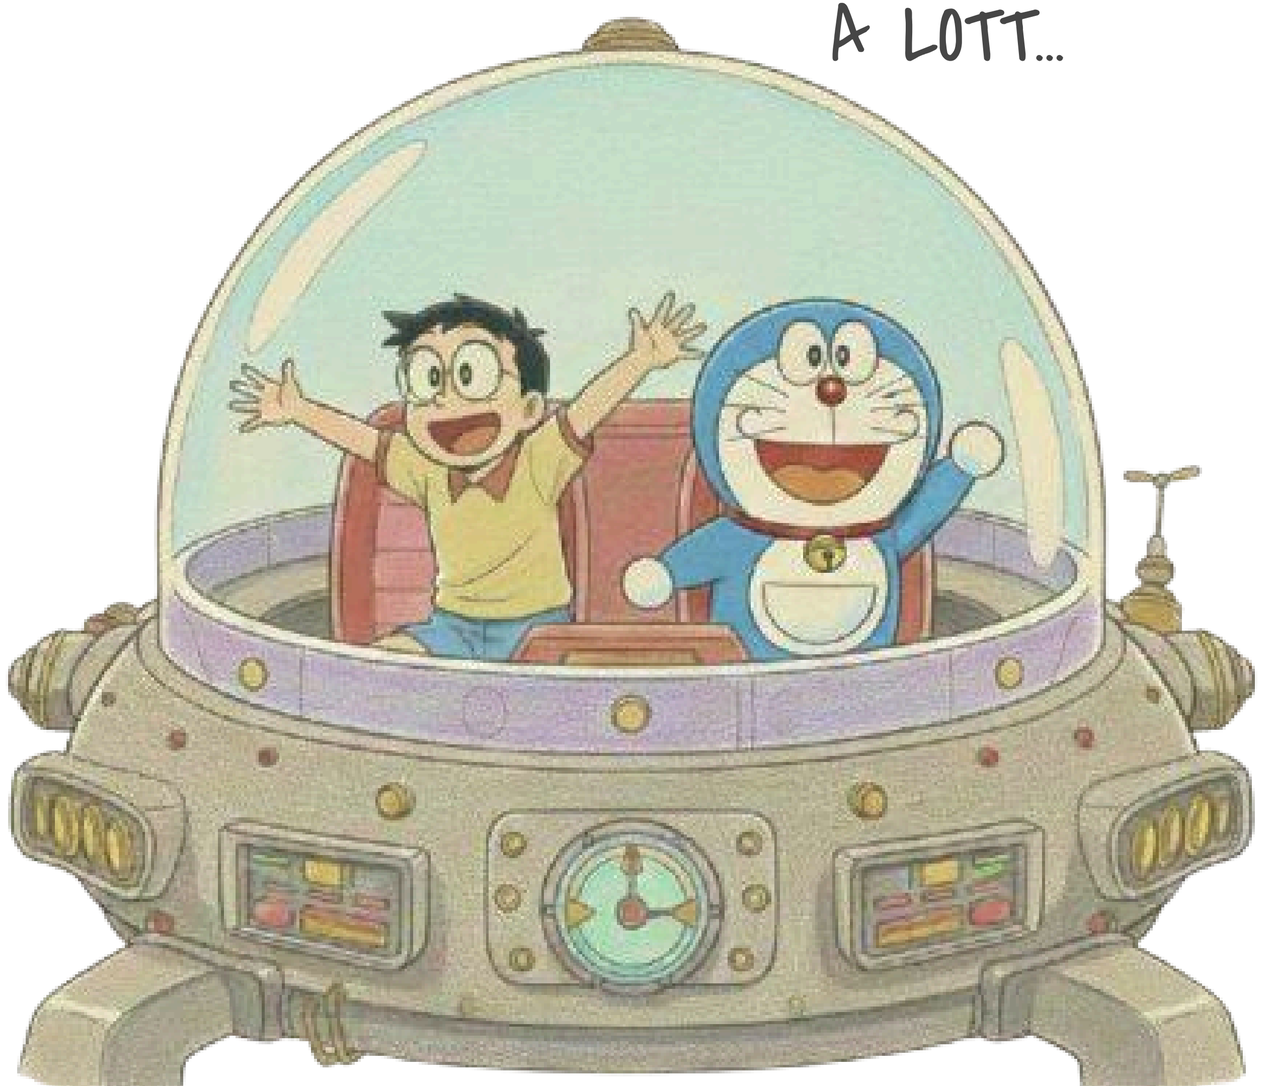
Capability Metrics

Speed + AI Readiness

- High capability determines how quickly users find value and stay retained
- Personas who benefit more from AI-led workflows deserve equal priority to behavioral need

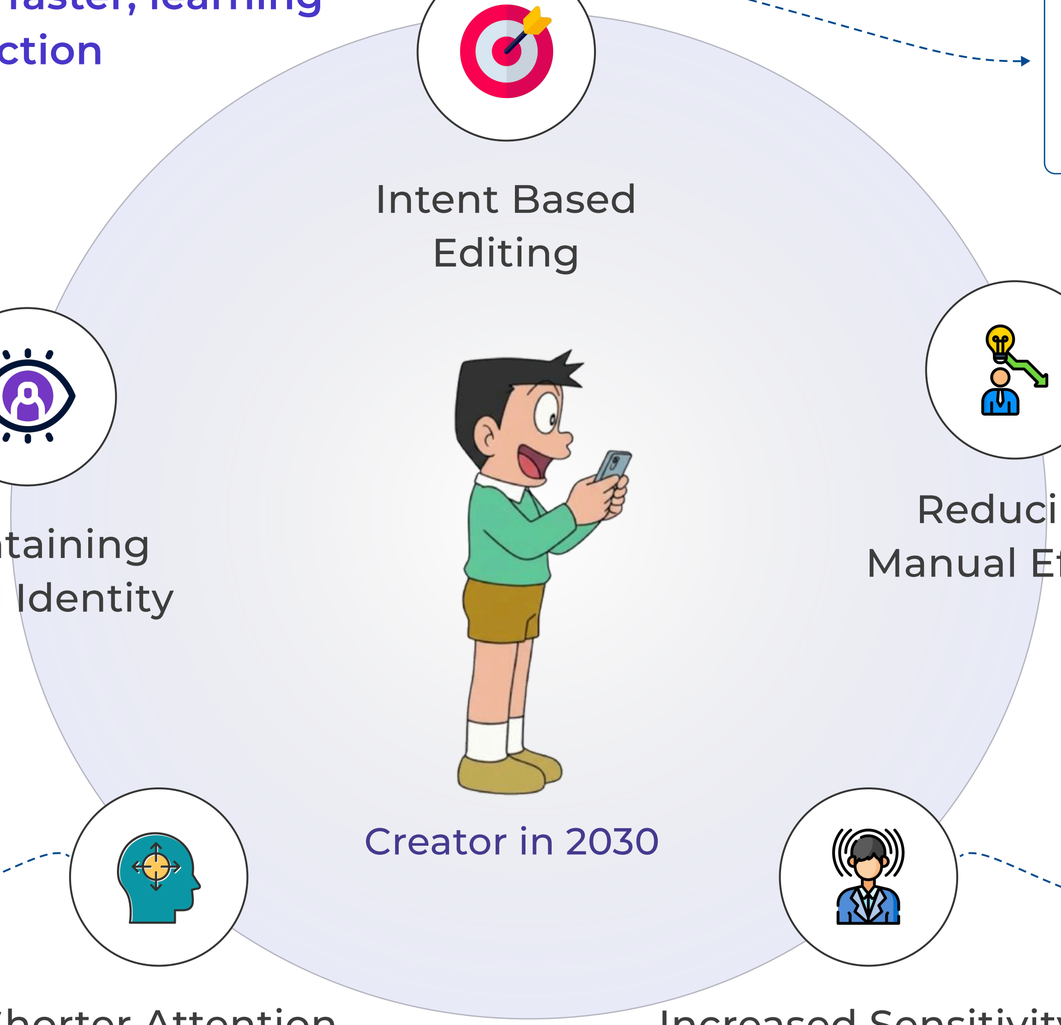
I wonder how things have changed
so much in just 5 years....

Yeahh...things have changed
Because user behavior have changed
A LOTT...



User Behavior in 2030

In 2030, creators win by evolving faster, learning smarter, and creating without friction



Creator in 2030

Intent Based Editing

- Wants a tool which can understand his Emotions

Maintaining Visual Identity

- Wants consistency across all edits

Reducing Manual Efforts

- Wants fewer steps for same precision

Shorter Attention Spans

- Needs faster workflows

Increased Sensitivity their about Data

- Worried about
  - Misuse
  - Plagiarism

What app do we use now ??



it's called.....KORTEX

A kotlin based android application





Workflow 1 : Creator Assist AI

Kortex understands user intent through voice, text, or reference images, translating emotions into precise, prioritised adjustments. It applies the values and surfaces the parameter changes, keeping creators fully informed about the edits.



Content Creator

- Voice based prompt
- Text based prompt
- Reference based prompt

“Make it Horror”





High Impact

Medium Impact

Low Impact

🔧 Supporting Adjustments:
- Sharpness: +0.20
- Highlights: -0.20
- Texture: +0.19

✨ Fine-Tuning:
- Brightness: -0.14
- Exposure: -0.13
- Vibrance: -0.10

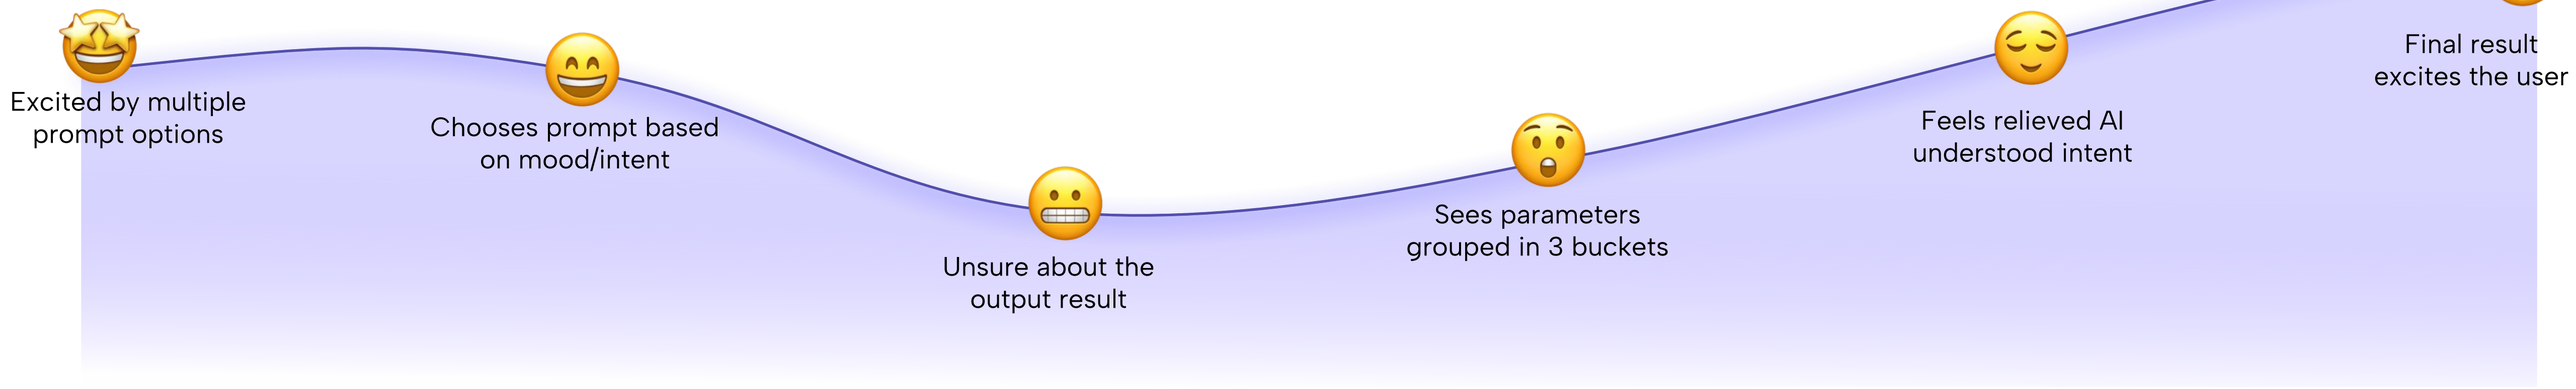
Pure Intent, No Technical Terms

Multiple Prompting Options

Emotion-Driven Input

Intelligent Parameter Prioritisation

Emotion Mapping



Workflow 1

Contextual Prompt



Input



“Make it Cheerful”



Result



Input



“Make it look like a Vintage Movie scene”



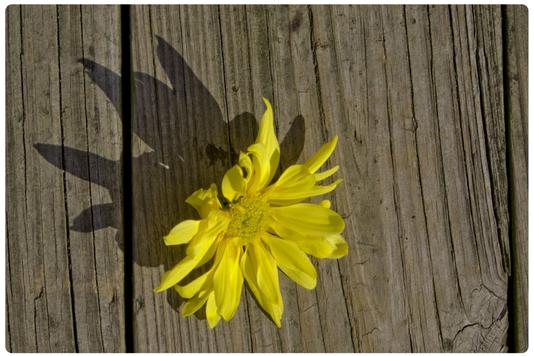
Result

Workflow 1

Reference Prompt



Input



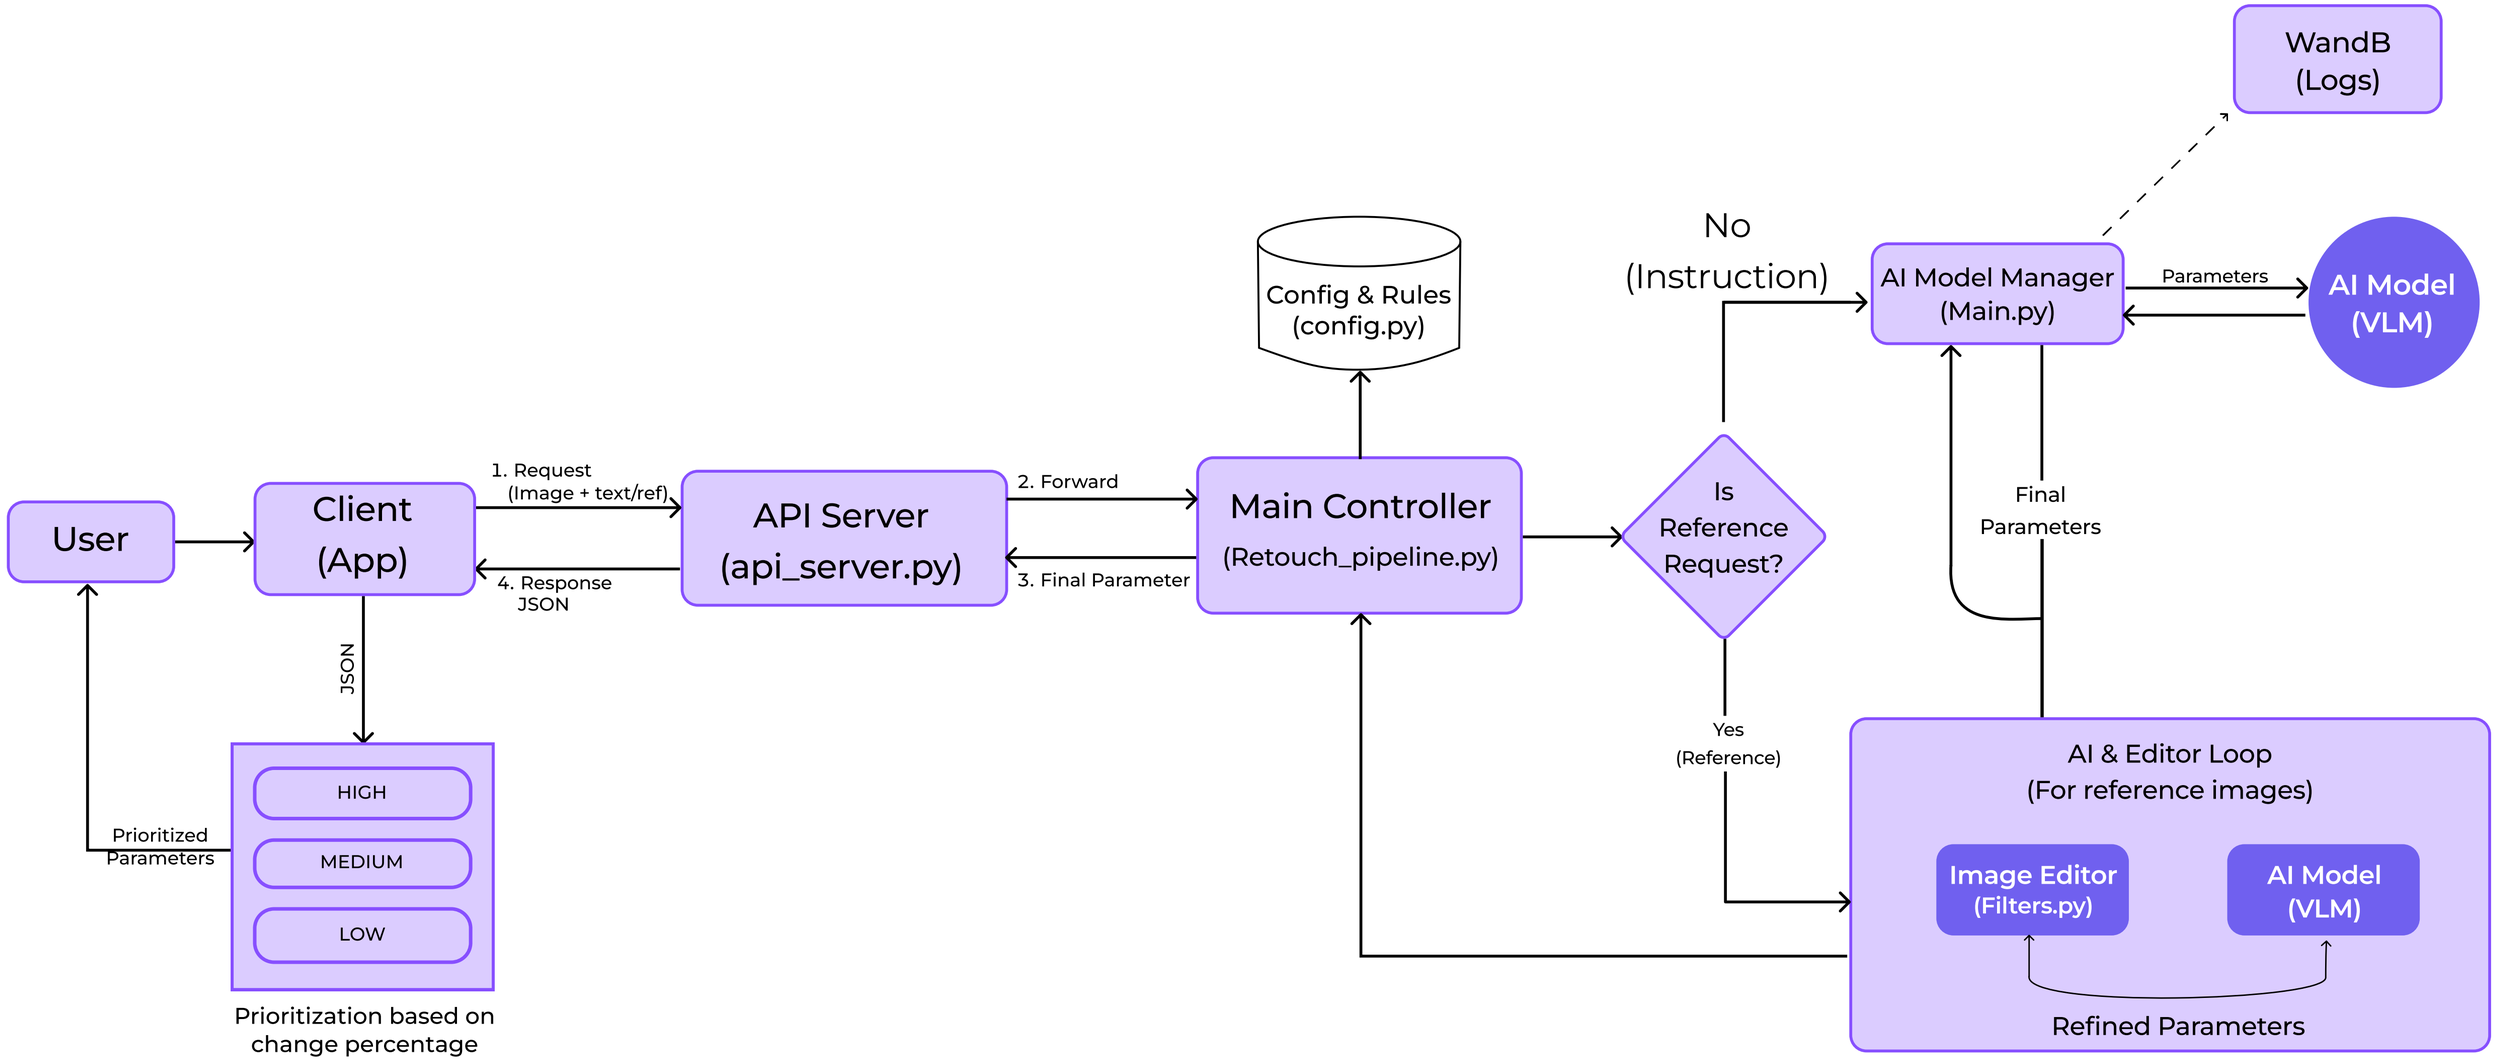
Reference



Result

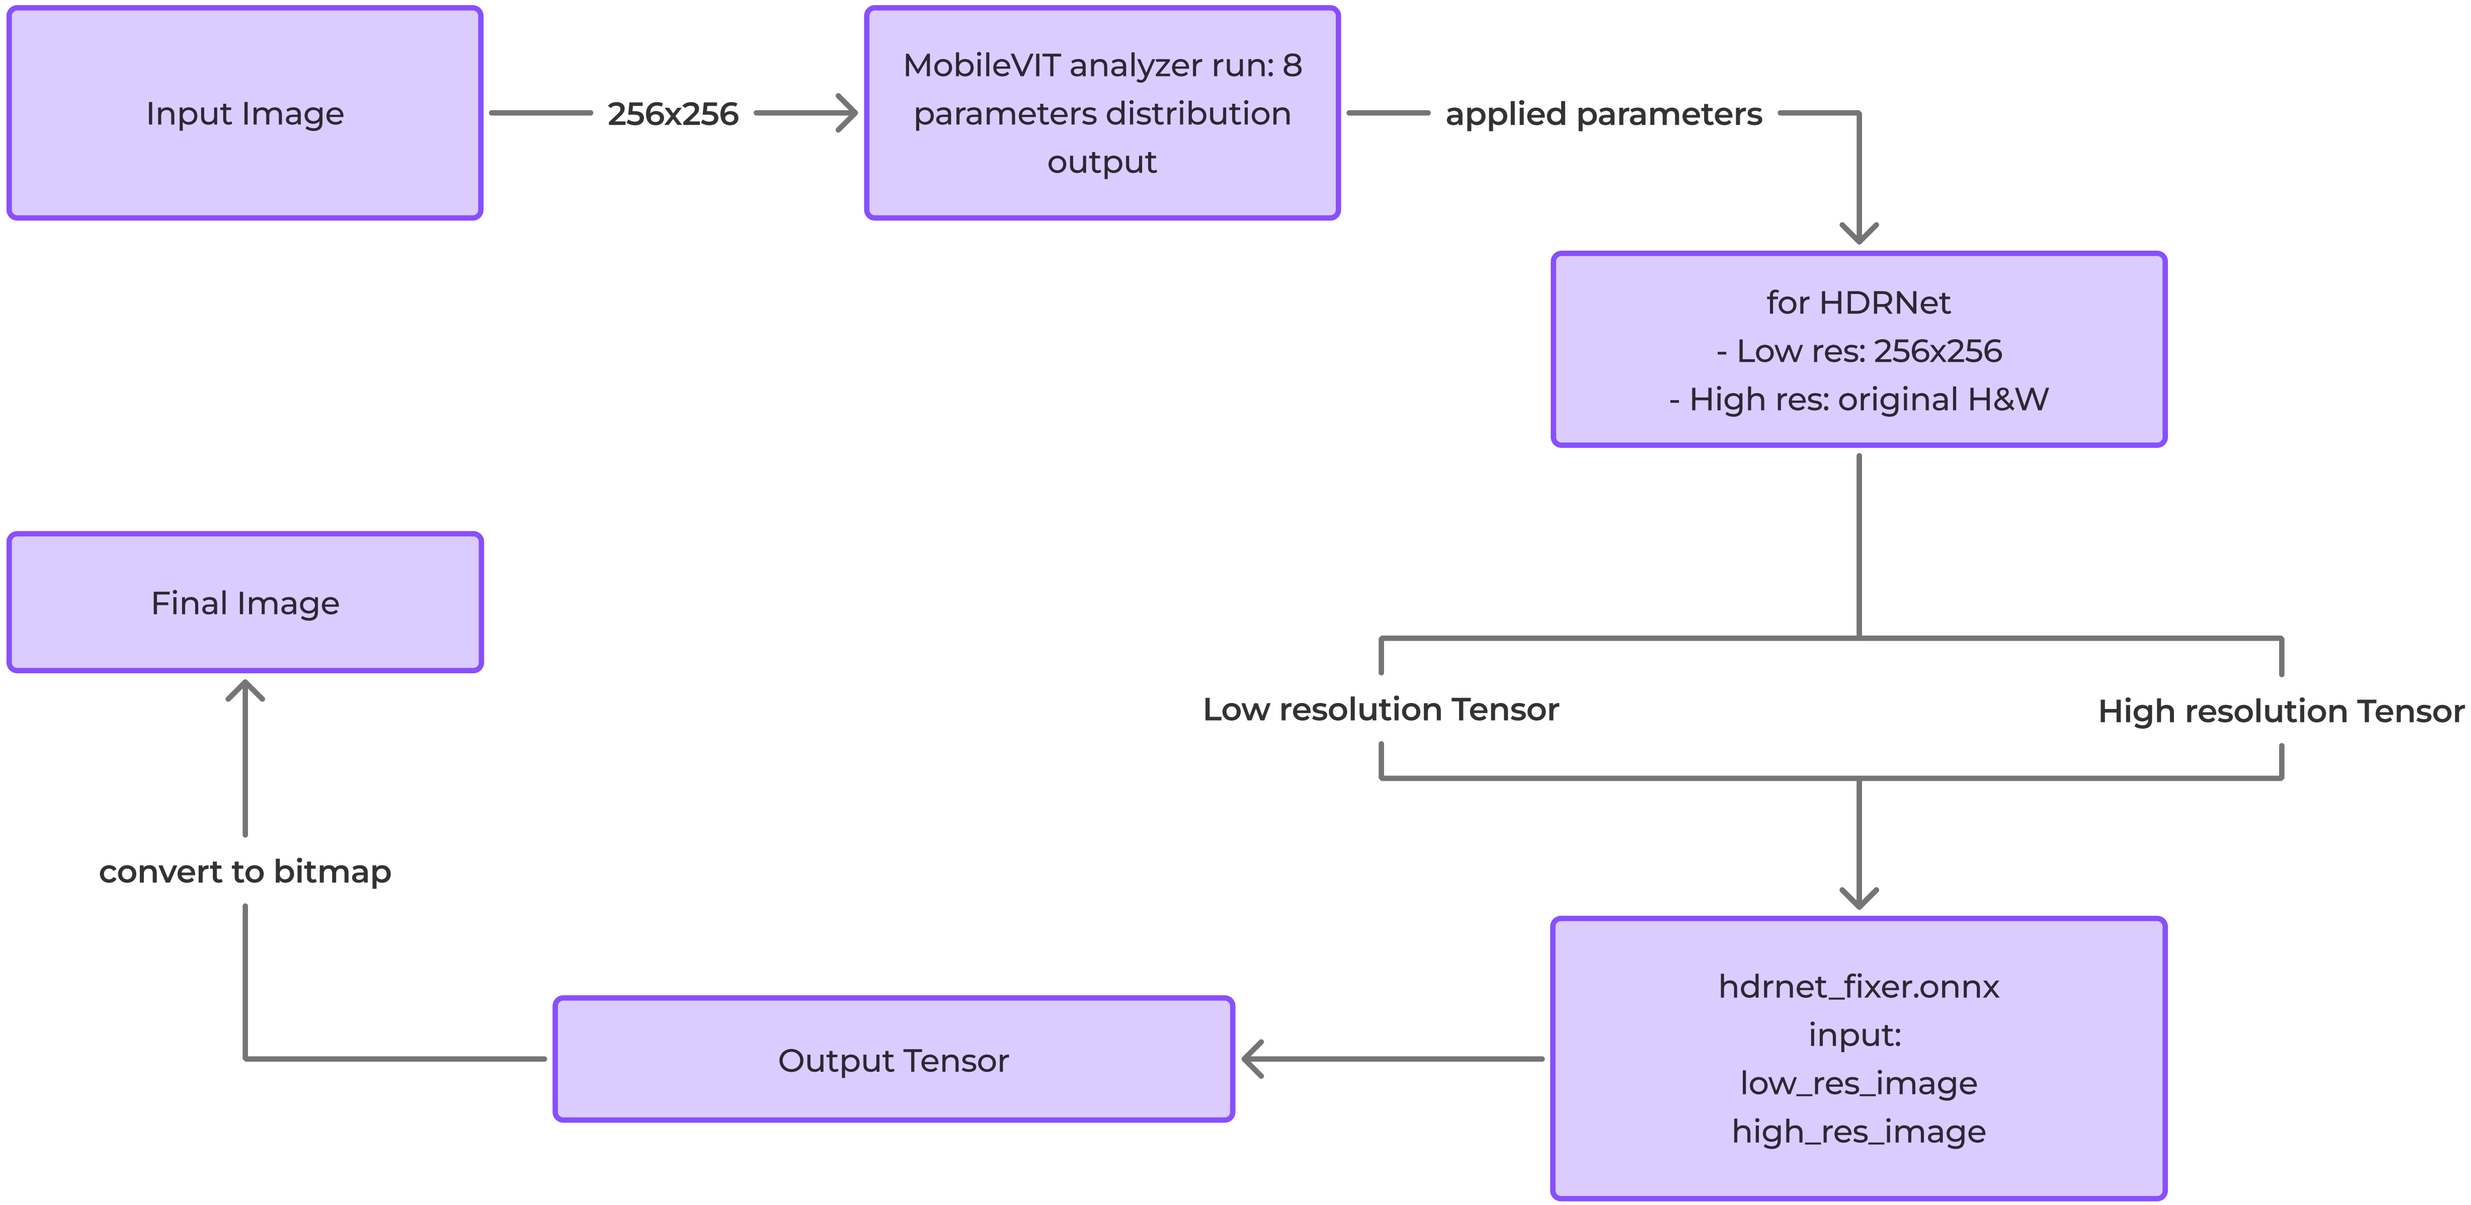
The Brain Behind

Retouch & Reference



The Brain Behind

Autoenhance



Compute Profile Summary

Hardware Environment: NVIDIA Tesla T4 (16GB VRAM) Testing Protocol: 30-minute sustained load test comparing Instruction Mode (Text-only prompt) vs. Reference Mode (Image-to-Image style transfer)

Metric Category	Metric	Instruction Mode (Prompt)	Reference Mode (Image)
System Resources	Peak System RAM	~7.7 GB	~6.5 GB
	CPU Utilization	< 10%	~30–40%
	Disk I/O (Read)	~58 MB	<5 MB
Performance	Inference Time	~12.5 s / step	~14.5 s / step
	Estimated TFLOPS	Spikes to ~0.18	Sustained ~0.17

Compute Profile Summary

Hardware Environment: NVIDIA Tesla T4 (16GB VRAM) Testing Protocol: 30-minute sustained load test comparing Instruction Mode (Text-only prompt) vs. Reference Mode (Image-to-Image style transfer)

Metric category	Metric	Instruction Mode (Prompt)	Reference Mode (Image)
GPU Resources	Peak VRAM	~11.5 GB	~12.8 GB
	Power Draw	Low / Idle	~70W (sustained)
	Temperature	Peaks ~43°C	Peaks ~47°C
	Utilization	Spiky (0–100%)	Consistent (~75–80%)

Workflow 2 : Section Edits

Our app enables precise sectional editing by letting creators select specific regions and apply targeted tools like object removal, background clean-up, AI fill, and local adjustments without affecting the rest of the image.



Content Creator

AI Edits

- Background Removal
- Sectional Edits
- Object Removal
- Move Object

Select an object



Move, Resize or Rotate



Harmonize



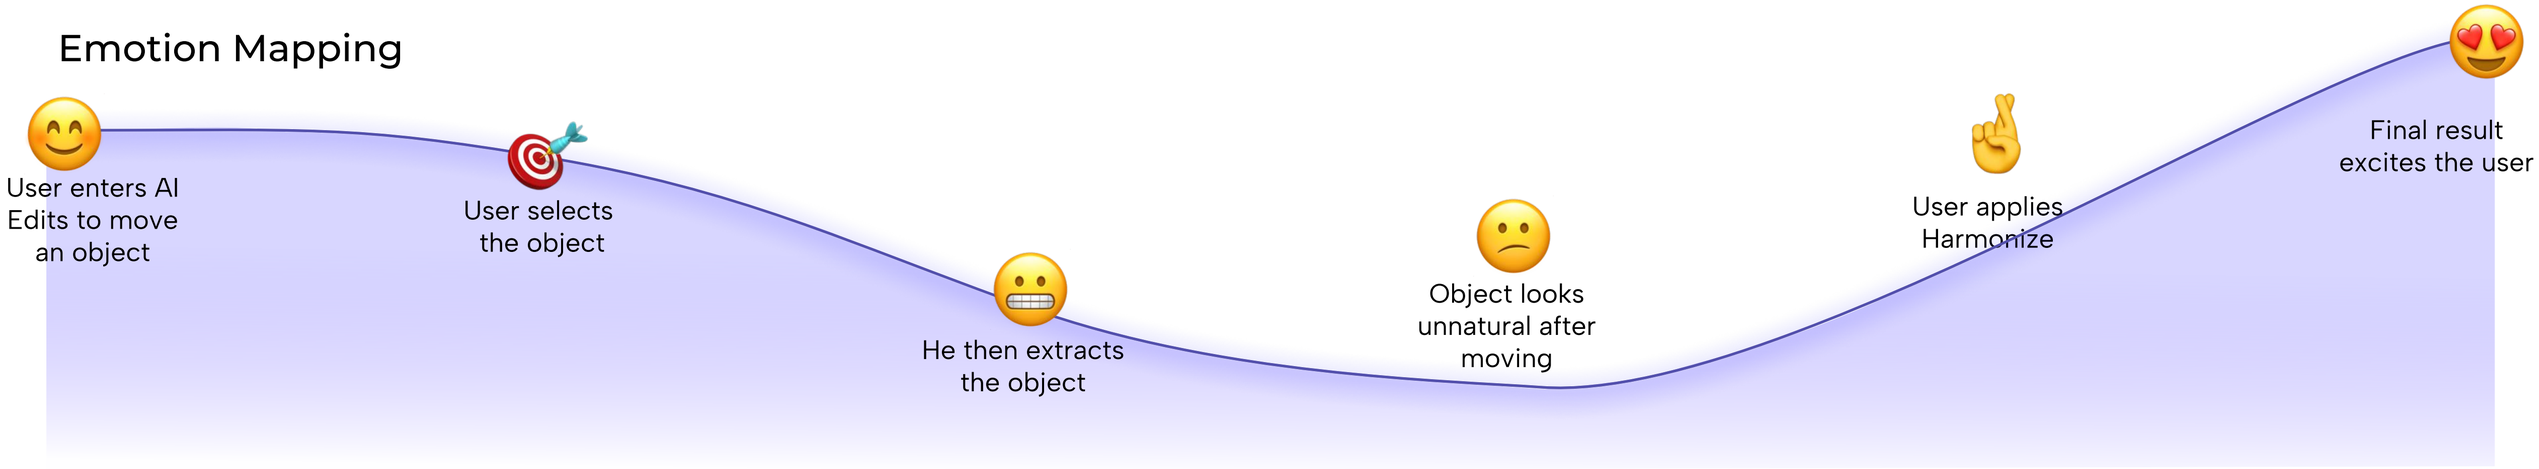
Real-Time Refinement & Preview

Multi-Tool Section Control

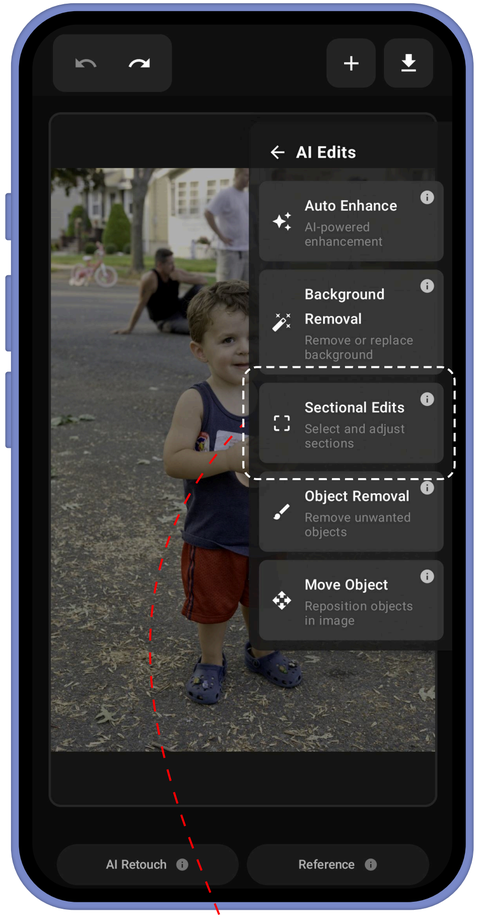
Selective Edits

Smart Harmonization for Realism

Emotion Mapping

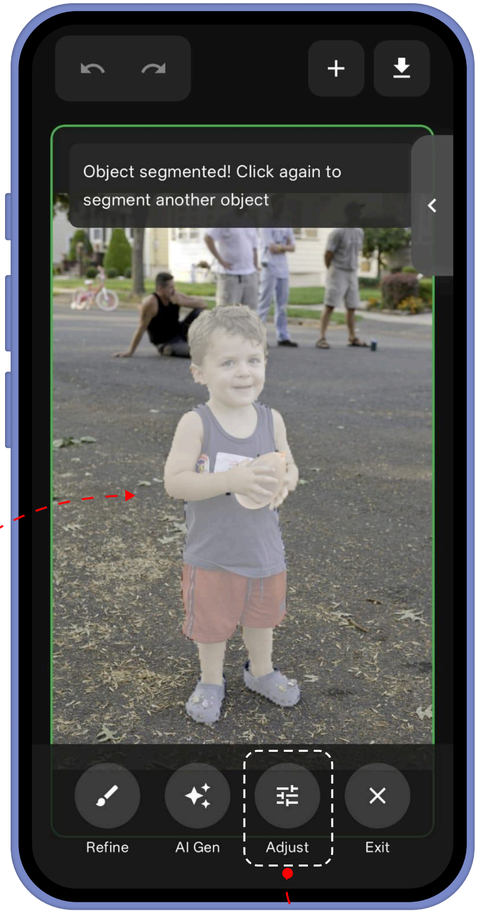


Sectional Edits - Object Adjustment

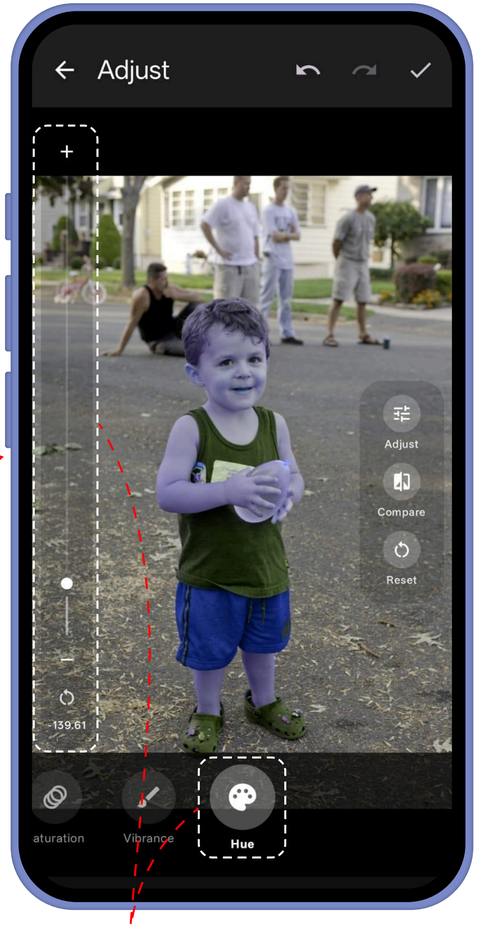


User taps on “Sectional Edits” to select objects

User taps on subject to select



User Taps on “Adjust”



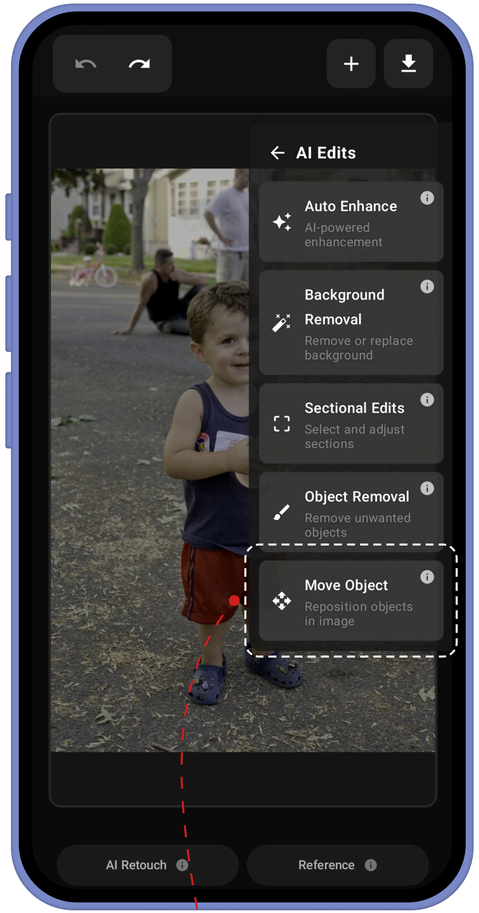
Object Adjustment

Precise Selection

Speed with Control

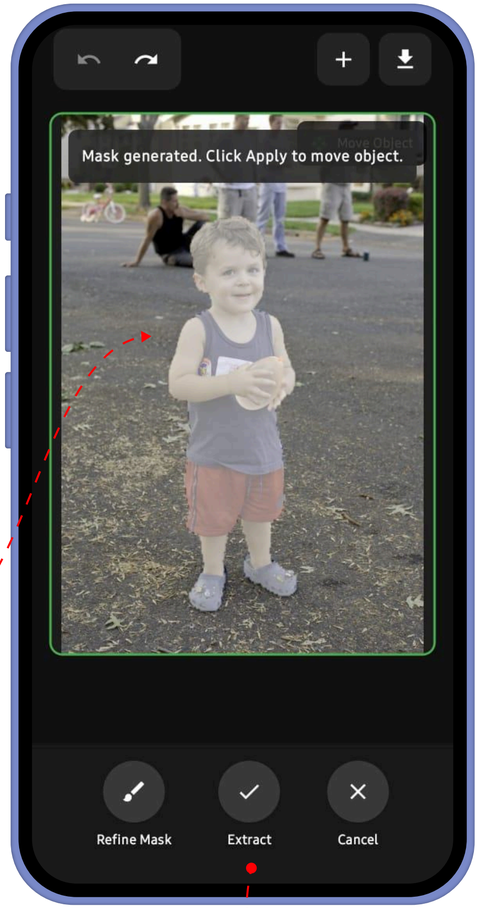
AI Assistance

Object Movement



User taps on “Move Object”

User taps on subject to select



User Taps on “Extract” (Subject gets extracted as an object)

User moves the subject as per requirements

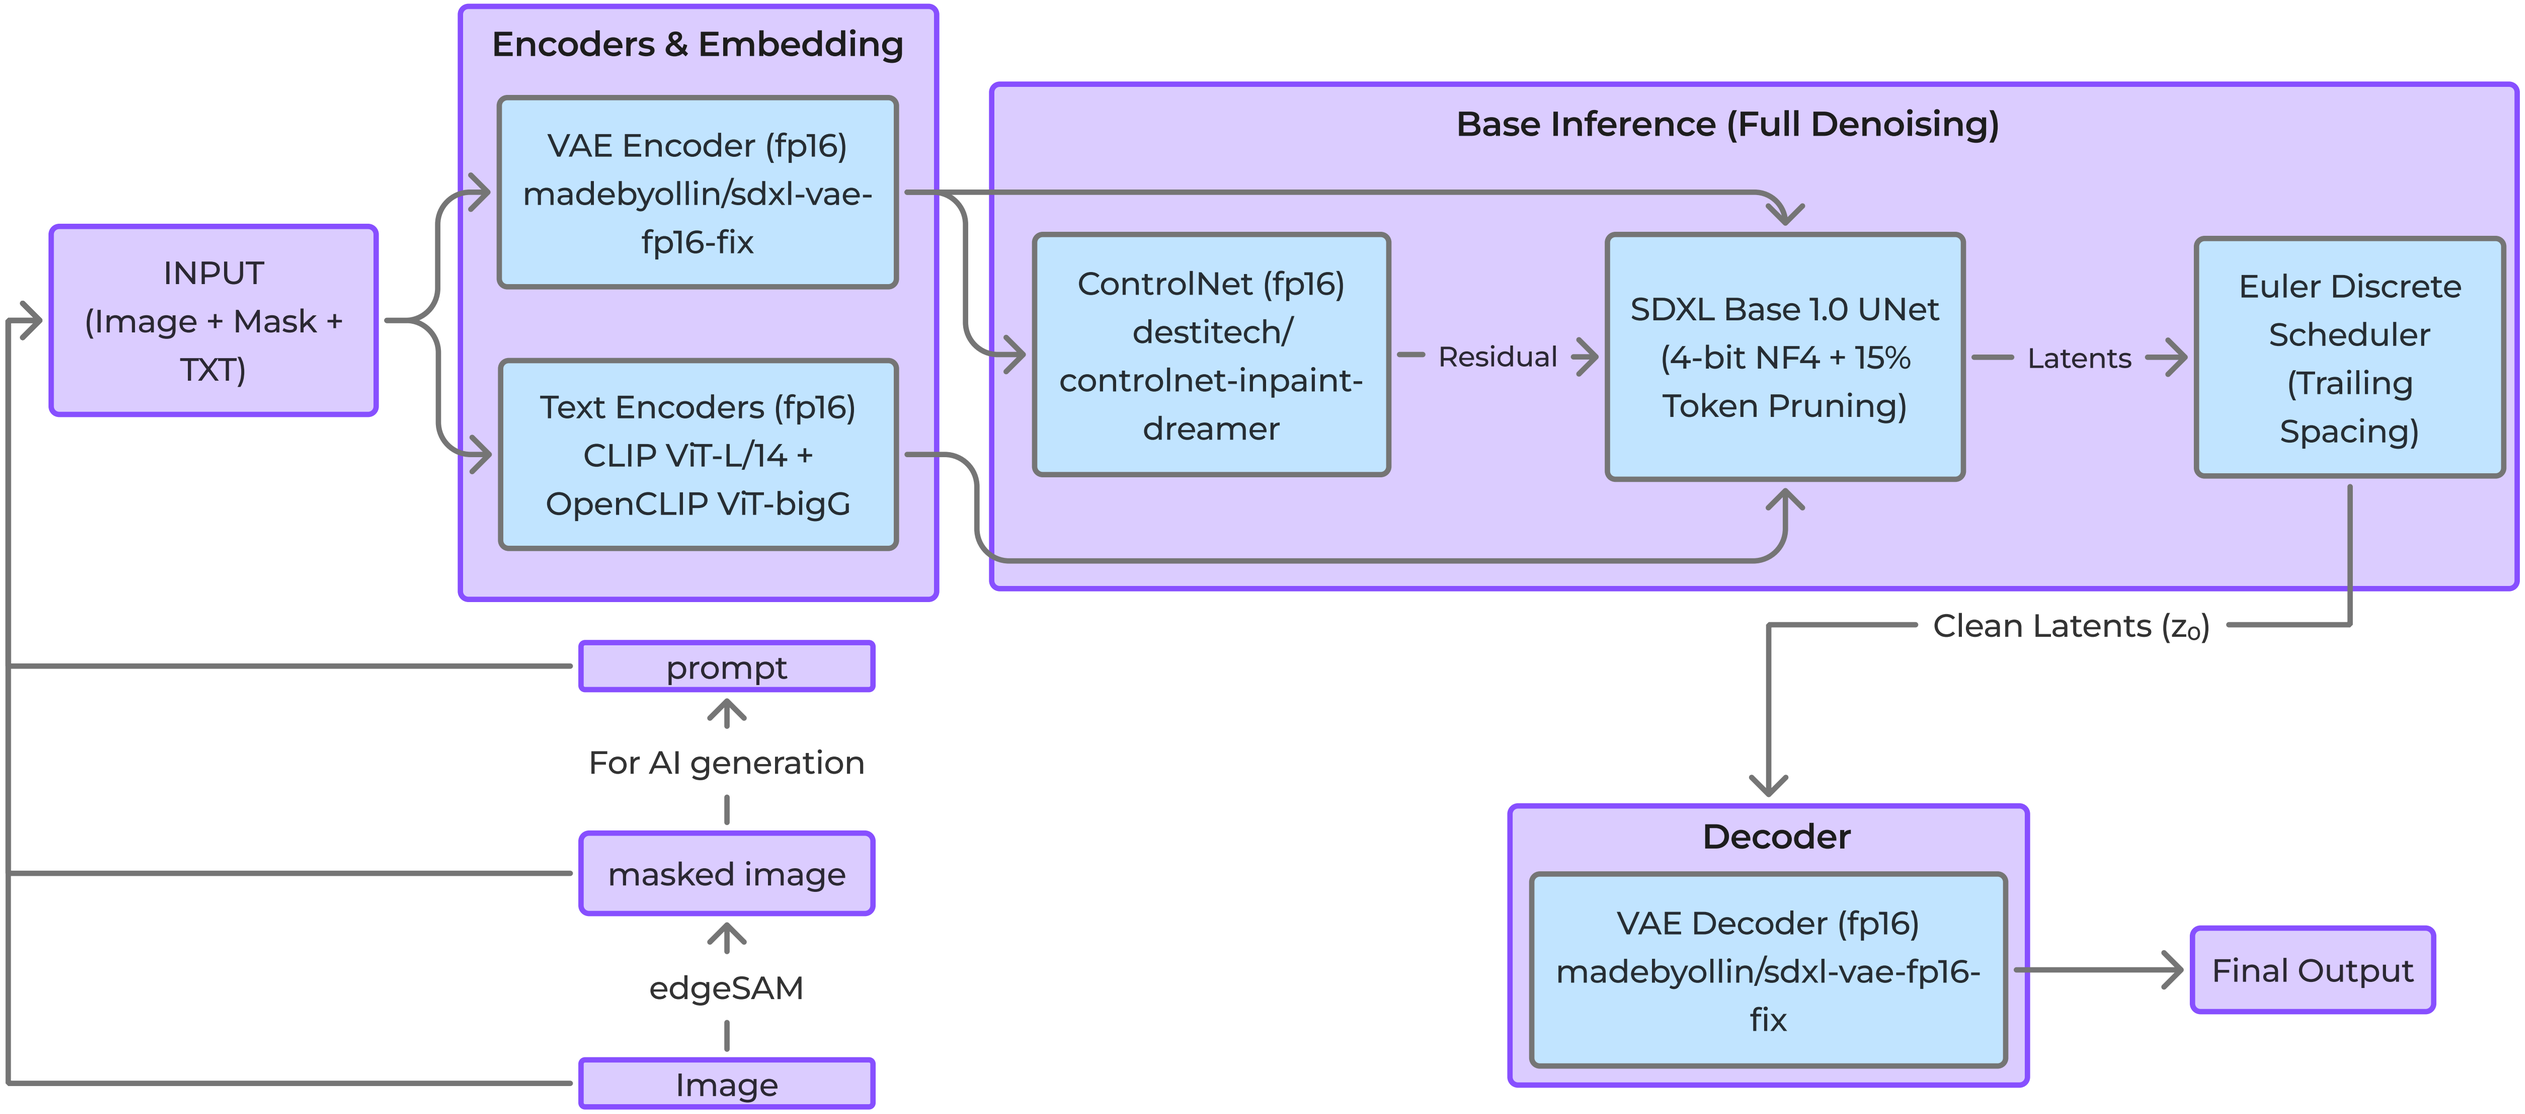


“Harmonize & Save” to get final result

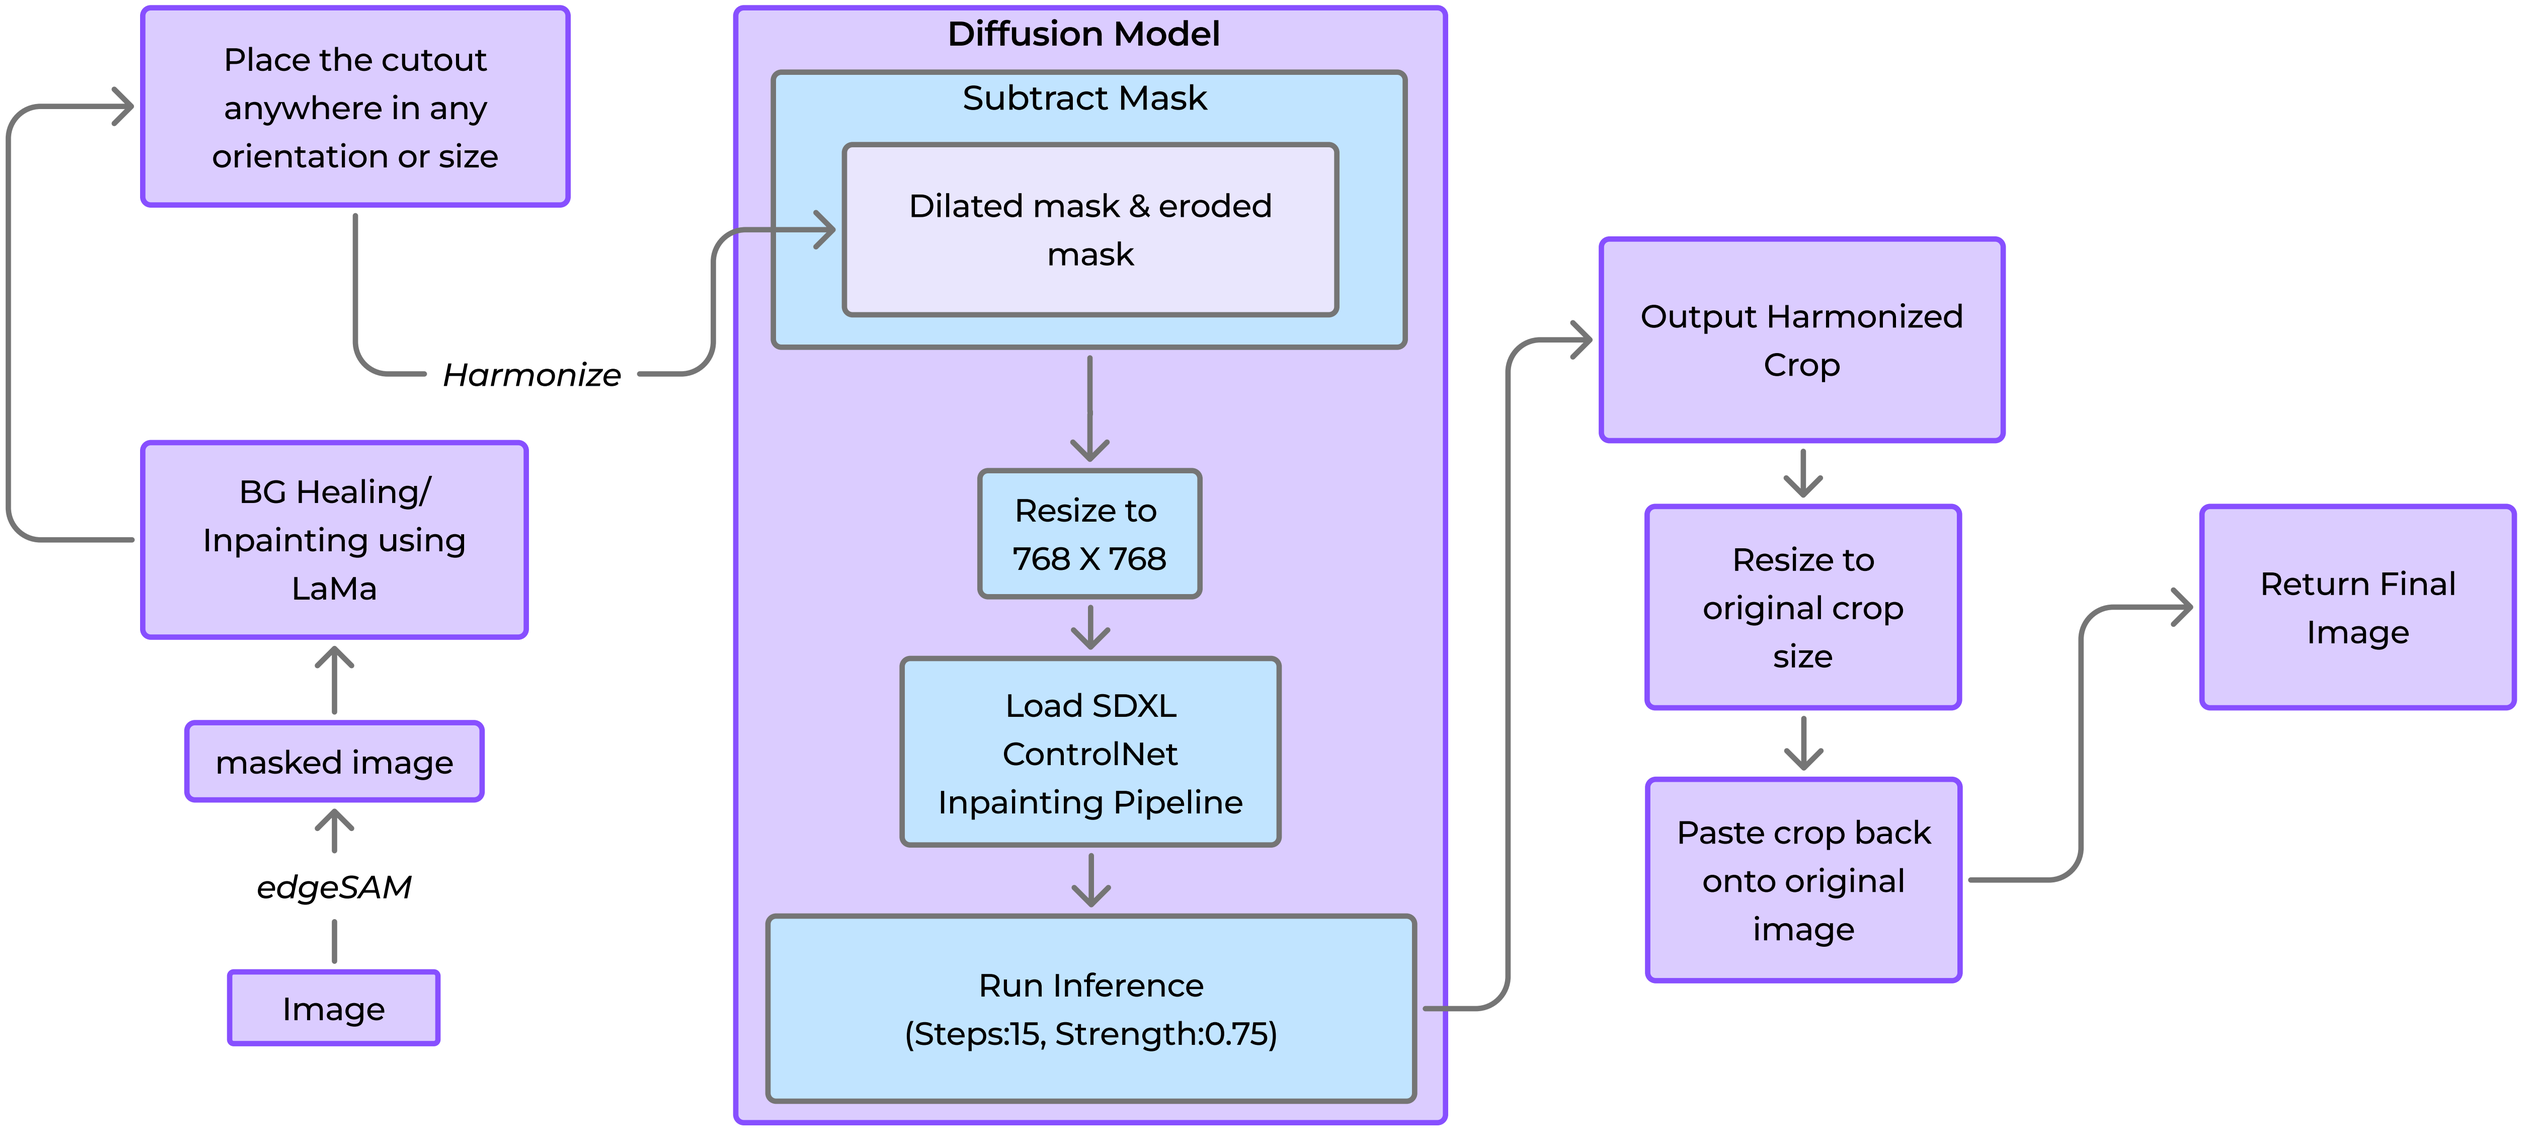
Precise Selection

Speed with Control

AI Assistance


Encoders & Embedding
VAE Encoder (fp16) madebyollin/sdxl-vae-fp16-fix
Text Encoders (fp16) CLIP ViT-L/14 + OpenCLIP ViT-bigG
INPUT (Image + Mask + TXT)
Base Inference (Full Denoising)
ControlNet (fp16) destitech/ controlnet-inpaint-dreamer
Residual
SDXL Base 1.0 UNet (4-bit NF4 + 15% Token Pruning)
Latents
Euler Discrete Scheduler (Trailing Spacing)
Clean Latents (z_0)
prompt
For AI generation
masked image
edgeSAM
Image
Decoder
VAE Decoder (fp16) madebyollin/sdxl-vae-fp16-fix
Final Output

The Brain Behind Harmonization



Compute Profile Summary

Metric Category	Metric Name	Harmonization	Generative Fill
VRAM	GPU Mem Allocated	7.4 GB (~45%)	9.0 GB (~56%)
Bandwidth	Time Accessing Mem	<10% avg (spiky)	40–45% constant
System RAM	Process Memory	Drops to ~2.8 GB	Steady at ~3.2 GB
Disk I/O	Disk Read (MB)	~250 MB	200 MB
Network	Network Traffic	High (~17 MB Recv)	Low (~9 MB Recv)
Compute	Total TFLOPS	~0.045	0.155

Technical Novelties and Optimization

EdgeSAM → Segmentation/Masking

Encoder runs once, caches the image embeddings, then decoder runs on user taps and segmentation is done in real time.

Kortex SDXL → Base Pipeline + ControlNet Inpaint Dreamer

Implemented NF4 quantization with ~15% Tokens merged

- SmartFill - from a given mask, controlnet guides diffusion model to adhere to the geometry and composition around the mask
- Edge-only Harmonization - It isolates only the boundary around the sticker using mask dilation/erosion, crops and upscales that region, then runs a lightweight inpainting pass on just the edges to blend it seamlessly back into the original image.

MiniCPM V-4.5 → AI Retouching

Utilized int4 quantized weights with fp16 activations for decreased model size while maintaining optimum model accuracy.
Unified 3d resampler to 16x compress tokens. Qwen 3B LLM backbone

Quantitative Benchmarking

Kortex SDXL achieves the best FID score, indicating more realistic and coherent generations than other SDXL-backbone models.

FID scores of Models with SDXL backbone (benchmarked on MagicBrush dataset)

Kortex SDXL (Ours)	InstruxtPix2Pix	Realvis_4	Juggernaut_XL	ThinkDiffusion_XL
115.8294	193.4221	146.1342	149.9384	148.3831

Metric	Kortex(NF4+ToMe)	SDXL(InstruxtPix2Pix)
CLIP Score	27.3156	25.5384
L1 Error	0.1971	0.2700
L2 Error	0.0616	0.1086

Our model **better captures user intent** from text prompts than instruction-based image editing models.

Technical Feasibility for 2025

72% of the models are deployed locally for speed and privacy, while **~28%** run on the cloud simulating 2030 on-device compute.

Model	Features	Environment
LaMa	Background healing	On-Device
EdgeSAM	Segmentation/Masking	On-Device
MobileViT + HDRNet	Auto-enhancing	On-Device
Vosk Speech Model	speech to text	On-Device
Kortex SDXL	Base Pipeline + Inpaint Dreamer	Hybrid
MiniCPM-V	AI Retouch+Reference	Hybrid

Will it be feasible in 2030? - The 2030 Compute (1)

The Growth Model - Explanation

To validate the decision of using the NVIDIA Tesla T4 as a proxy for 2030 mobile hardware, we employed a quantitative forecasting model based on historical Mobile NPU scaling trends (2017-2024)

We model the Neural Processing Unit (NPU) performance (P) at year y using an exponential growth function starting from a hard baseline in 2024

$$P_y = P_{base} \times (1 + r)^{(y-2024)}$$

- The baseline performance in 2024, established at 40 TOPS (Trillions of Operations Per Second)
- This represents the industry midpoint between the Apple A17 Pro (35 TOPS) and Qualcomm Snapdragon 8 Gen 3 (45 TOPS)

The target projection year (2030)

The Compound Annual Growth Rate (CAGR) of NPU efficiency

Sources:

1. HAI Index 2025 2. HAI Index 2024

Will it be feasible in 2030? - The 2030 Compute (2)

Scaling Scenarios

We projected three scaling scenarios for the 6-year horizon (2024–2030):

Scenario	Growth Rate (r)	Rationale	2030 Projection (P_2030)
Conservative	10%	Constraints on thermal envelope and power draw limit scaling	~71 TOPS
Realistic	**20%**	Continued architectural improvements (N2/A14 process nodes)	**~120 TOPS**
Optimistic	30%	Aggressive low-precision marketing (INT2/INT4) counting	~193 TOPS

Hardware Mapping

Our project targets the Realistic Scenario, projecting that flagship mobile processors in 2030 will deliver approximately 100-120 TOPS of INT8 inference performance

The NVIDIA Tesla T4, widely used in this project's server simulation, offers:

- ~130 TOPS (INT8)
- ~65 TFLOPS (Mixed)

By optimizing our SDXL pipeline (via 4-bit NF4 quantization and Token Merging) to run efficiently on the T4, we effectively demonstrate a workload that will fit comfortably within the thermal and compute budget of a high-end mobile device in 2030

Will it be feasible in 2030? - The 2030 Compute (3)

2025 Baseline: Current Smartphone AI Hardware Limits

Avg smartphone DRAM ~7.6-9 GB	**Flagship DRAM** 12-16 GB; extreme 24 GB	**NPU TOPS** ~35-45 TOPS
DRAM bandwidth ~85 GB/s	**Effective bits/weight** 4-bit, ~1.0 B/param	**Practical on-device LLM** 7-8B params (phone)

Based on current trends and uncertainty in AI growth, we have classified future outcomes into three scenarios.

Modeling Approach

Conservative

Assumed slow AI growth limited by power, cost, and hardware constraints. On-device AI remains focused on basic use cases.

Realistic

Assumed steady, roadmap-led growth in hardware and on-device AI. This represents the most likely industry outcome.

Optimistic

Assumed rapid AI acceleration driven by major breakthroughs. Premium smartphones evolve into personal AI machines.

Will it be feasible in 2030? - The 2030 Compute (4)

How we built this Growth Model?

Choose a baseline year and fix a hard reference value

Pick a year where the metric is well measured. Call this value the baseline.

Identify historical scaling behavior for the metric

Look at **past data** to understand how the metric typically changes each year. Examples of patterns: steady percentage growth, slowing growth, or step jumps.

Purpose: Determine whether exponential form is suitable.

Select a compound annual growth rate ‘r’

Based on the observed behavior, choose r. This may be a single value or a set of scenarios such as conservative, realistic or optimistic.

Purpose: Express the yearly improvement factor numerically.

Define the target projection year

Represent the year you want to forecast as y. Compute the number of years since the baseline using (y-baseline year).

Purpose: Set up the time variable that drives the model.

Construct the exponential projection function

P_y = P_base × (1 + r)^(y - baseline year). This mirrors the compounding mechanism observed historically.

Purpose: Produce a forward looking model consistent with prior data.

Use scenarios if needed

If the metric may vary under different constraints, build three versions one using a low r, one using a mid r, one using a high r. This keeps the model flexible while preserving the same structure.

Purpose: Express uncertainty without changing the core equation.

Instances that validate how this model was developed

Global average DRAM per smartphone rises from 4.3 GB (2020) → 5.1 GB (2021) → 6.2 GB (2022) → 6.9 GB (2023) → 7.6 GB (2024F) → ~9 GB (2025F), implying ~15–16% CAGR.

Sources

Effective bandwidth of a typical flagship memory subsystem (64-bit bus) goes from ≈51 GB/s (LPDDR5-6400, ~2020) to ≈86 GB/s (LPDDR5X-10700, Samsung 2nd-gen, 2024) – about 11% compound annual growth.

Sources

AI Ethics

Why?

It gives us a policy baseline by country so we can design features, safeguards, and messaging where they matter most.

	Content Authenticity	Deepfake & Harmful Use	Access, Data Boundaries	Commercial & Training Rules
US	3	4	1	2
Brazil	2	1	3	4
W. Europe	4	3	1	2
India	2	1	3	4
Japan & SK	3	1	2	4

Across our four AI-governance buckets, our priority markets consistently put the most weight on:

Content Authenticity

and

Commercial & Training Rules

AI Ethics

Content Authenticity

- **Source and edit history**
- Tools to inspect provenance
- **Tamper detection and integrity**
- **AI-use disclosure and labeling**
- Persistent metadata and signatures

Deepfake & Harmful Use

- Limits on political deepfakes
- Filters for harmful or hateful content
- Safeguards for non-consensual imagery
- Human review of high-risk outputs
- Prohibitions on deceptive impersonation

Access, Data Boundaries

- Separation of user and training data
- Regional data-storage controls
- Role-based access and audit logs
- Limits on cross-border transfers
- User rights for export and deletion

Commercial & Training Rules

- **Rights over user-generated content**
- **Obligations to respect third-party rights**
- Permissions for commercial exploitation
- Restrictions on training new models
- Clarification of liability and indemnity

But... what's the point?
Images get pirated so easily...



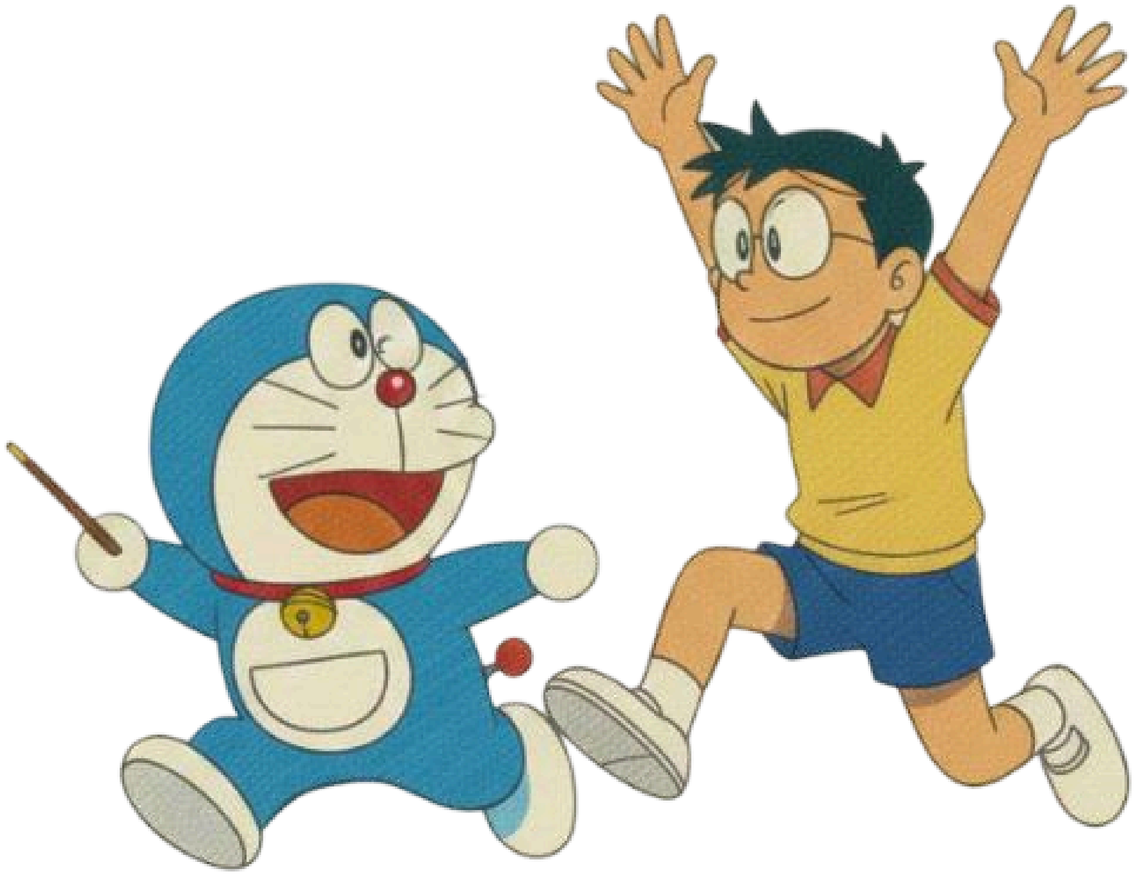
I know....
That's why Kortex has features like LSB Watermark to stop piracy.



Whoa!"

That's crazy

I'll tell all my friends!



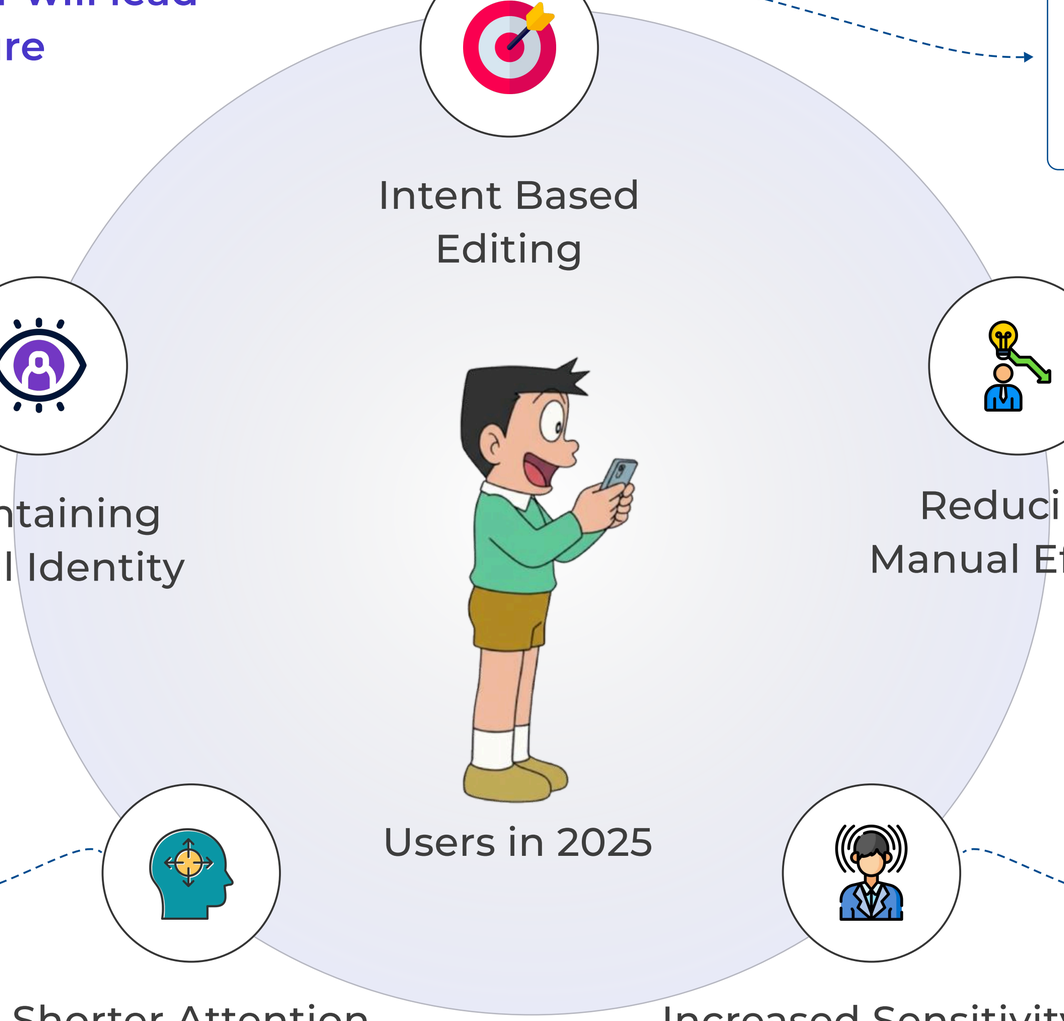
We Are Now Open To Questions



Appendix

Behavior Validation

Changes in current habits of user will lead to the change in behavior in future



Users in 2025

Intent Based Editing

Around 50–60% of active photo editors now use generative AI tools in their workflows, Source

Maintaining Visual Identity

70–80% of casual content creators say it is important to keep a consistent visual identity Source

Reducing Manual Efforts

80% adopt AI mainly to reduce repetitive edits and save time Source

Shorter Attention Spans

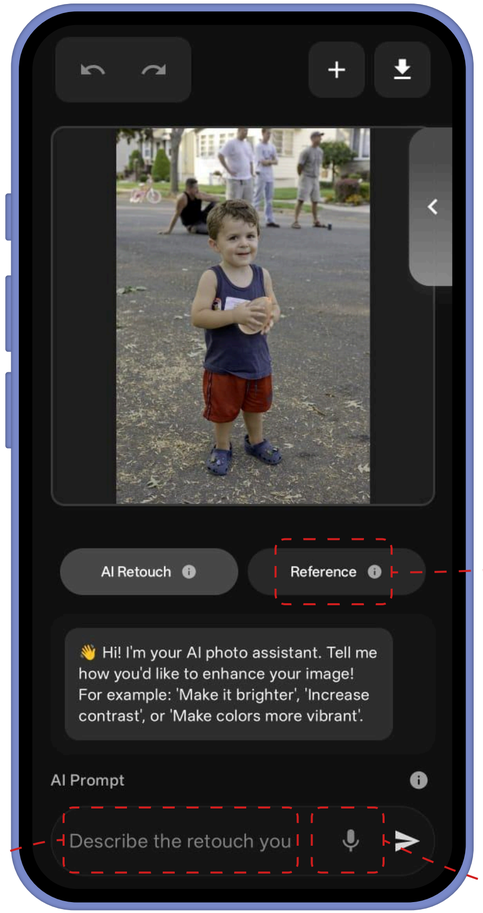
Attention span dropped from 2 min to 40 sec Source

Increased Sensitivity their about Data

71% of adults are concerned about their privacy Source

Workflow 1

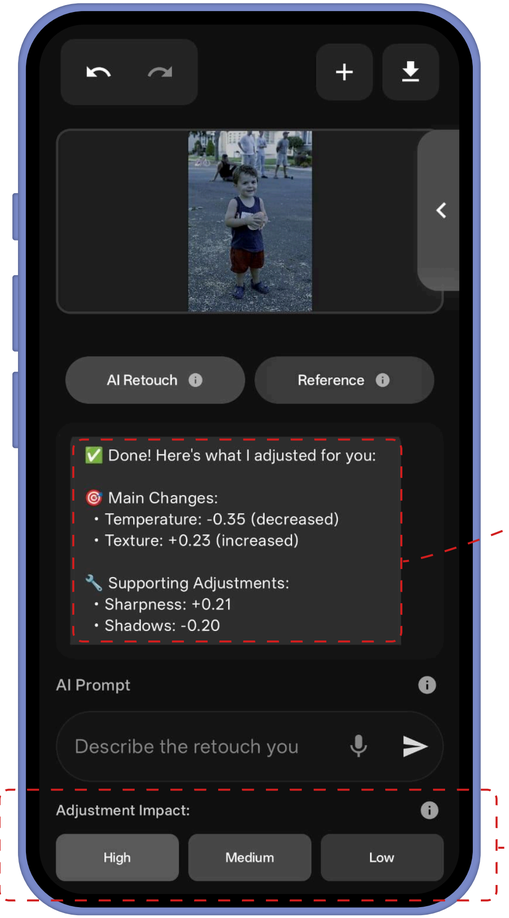
Contextual Prompt



Reference Based Edits

Text Prompt

Voice Based Prompt



Transparent Adjustments

Intelligent Parameter Prioritisation

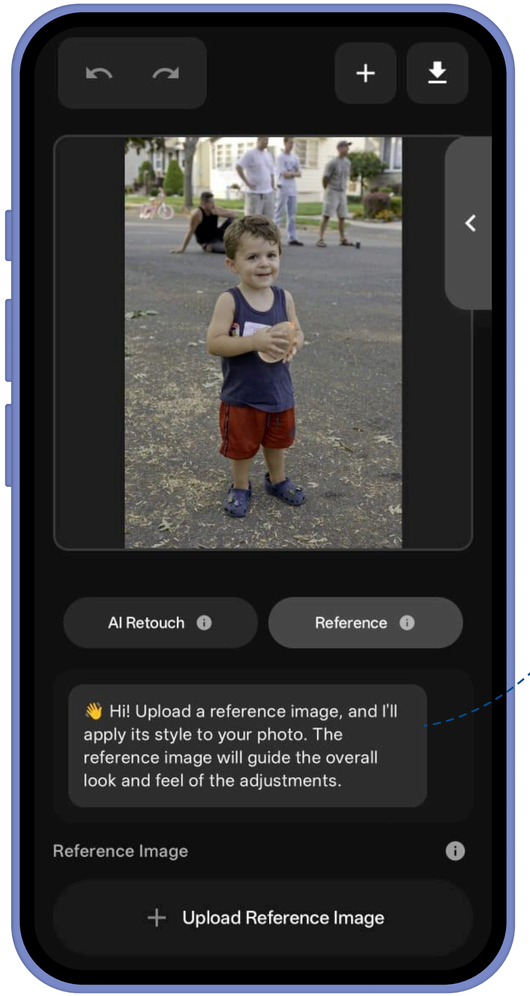
Multiple Prompting Options

Pure Intent, No Technical Terms

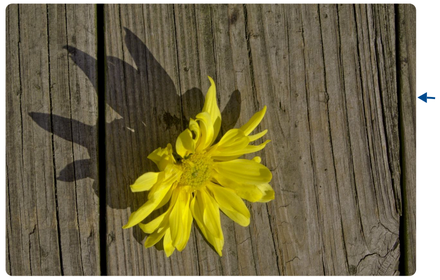
Intelligent Prioritisation

Workflow 1

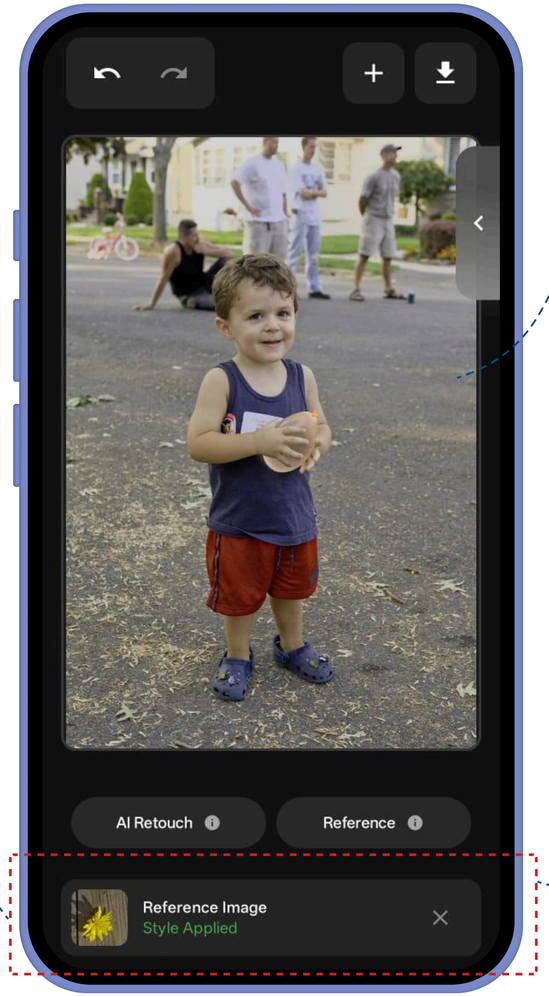
Reference Prompt



User Guide



Used Reference Image



Enhanced
Enlarged Image

AI Style Transfer

Multiple Prompting Options

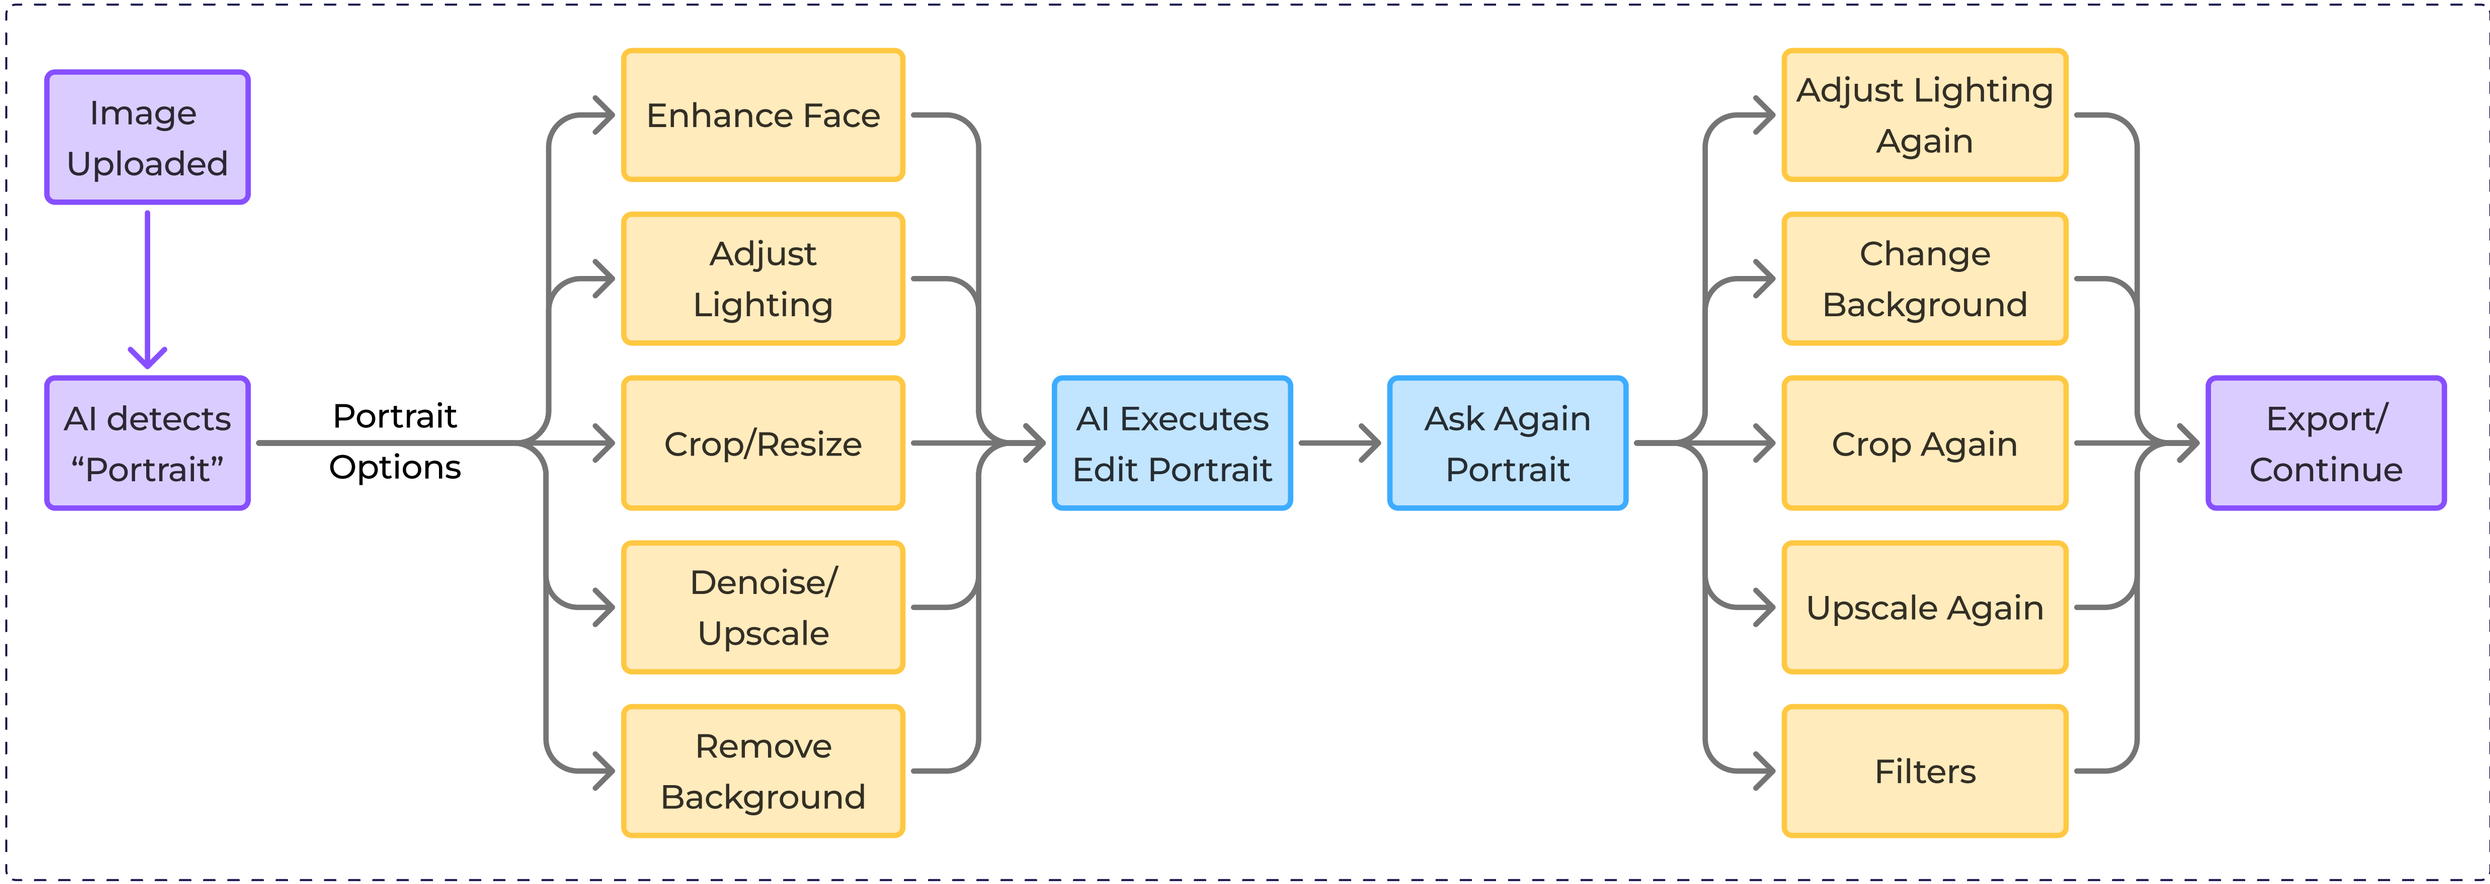
Human-Guided Refinement

Reference-Aligned Output

Future Scope : Context Aware Suggestive Edits

AI will analyze every image the user uploads and suggest the most suitable edits based on its content. After the user applies an edit, the AI will re-evaluate the updated image and recommend the next set of improvements. This iterative loop continues until the user decides to stop.

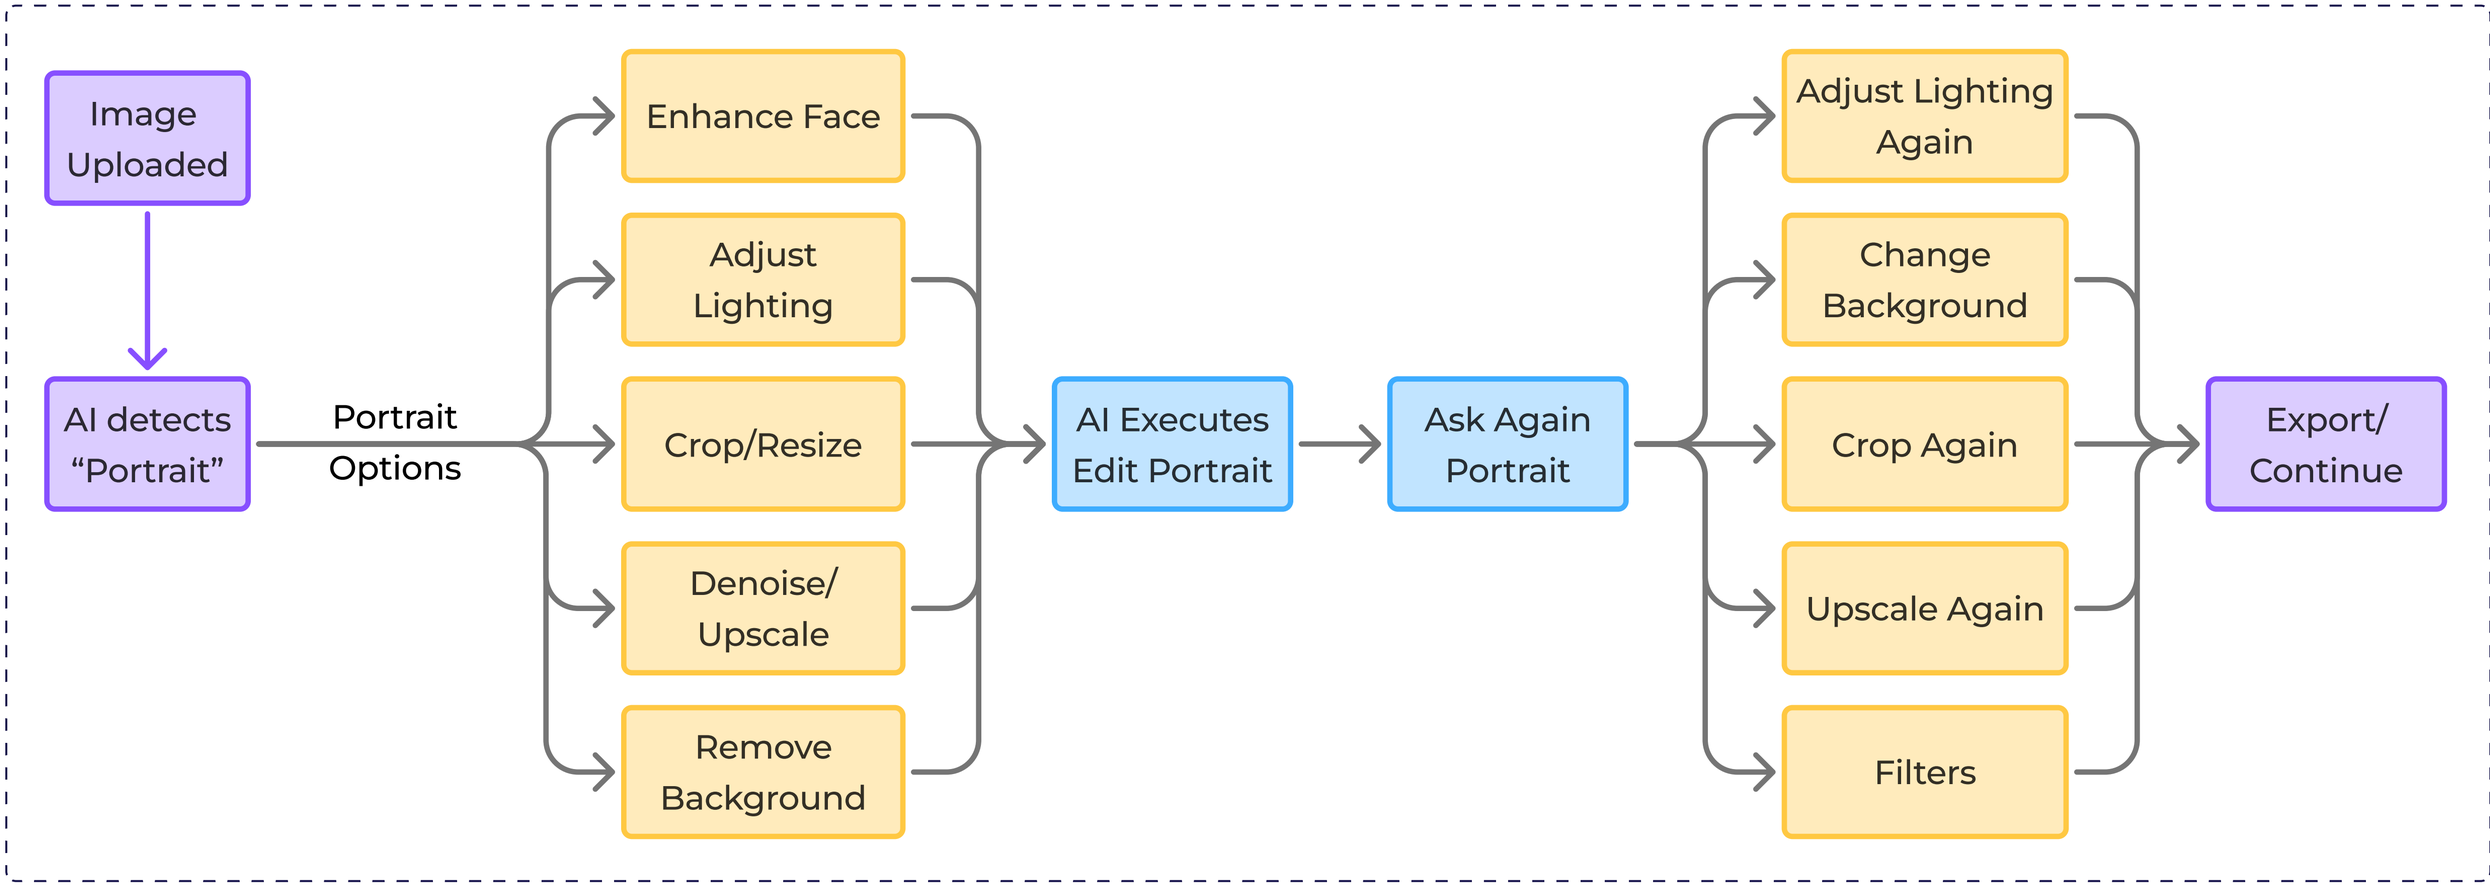
Process Flow



Future Scope : Context Aware Suggestive Edits

AI will analyze every image the user uploads and suggest the most suitable edits based on its content. After the user applies an edit, the AI will re-evaluate the updated image and recommend the next set of improvements. This iterative loop continues until the user decides to stop.

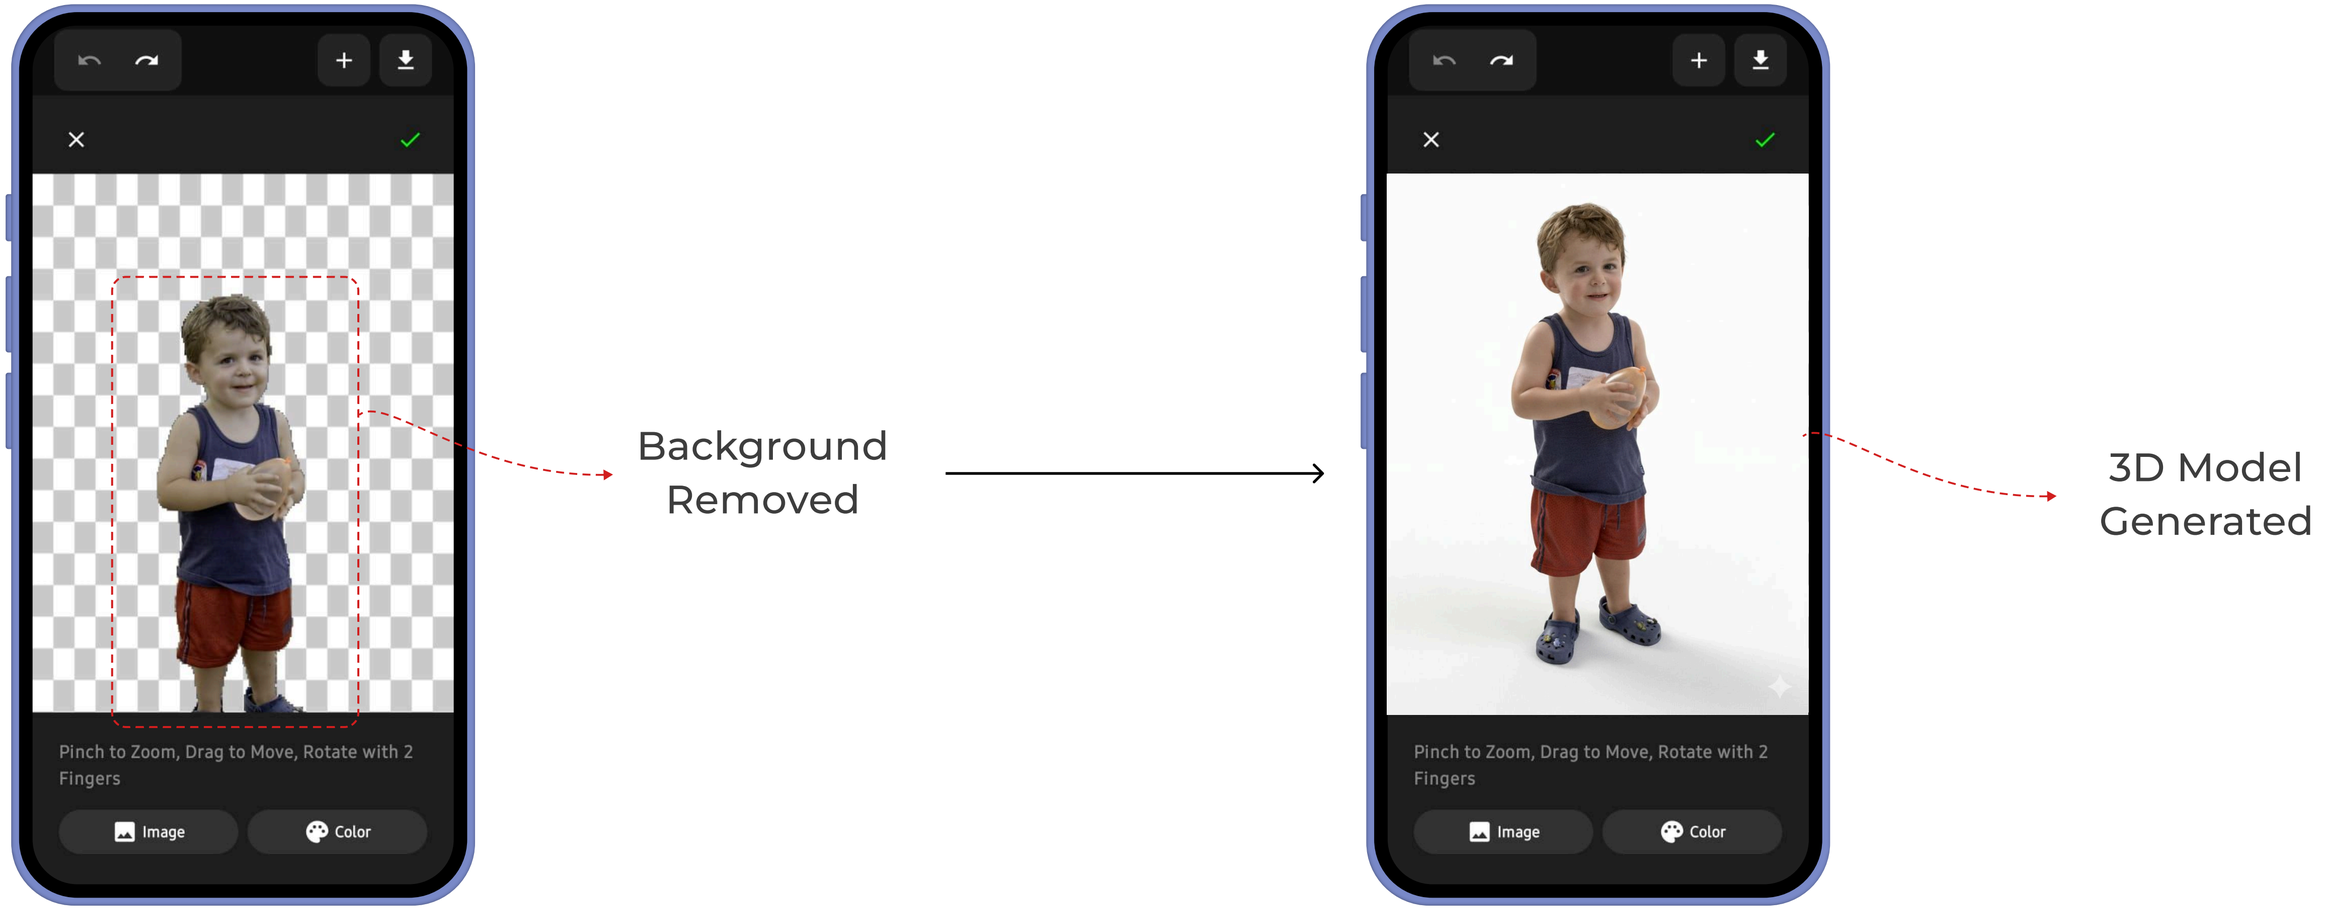
Process Flow



Future Scope

Object Removal Extended to Shadow & Reflection Intelligence

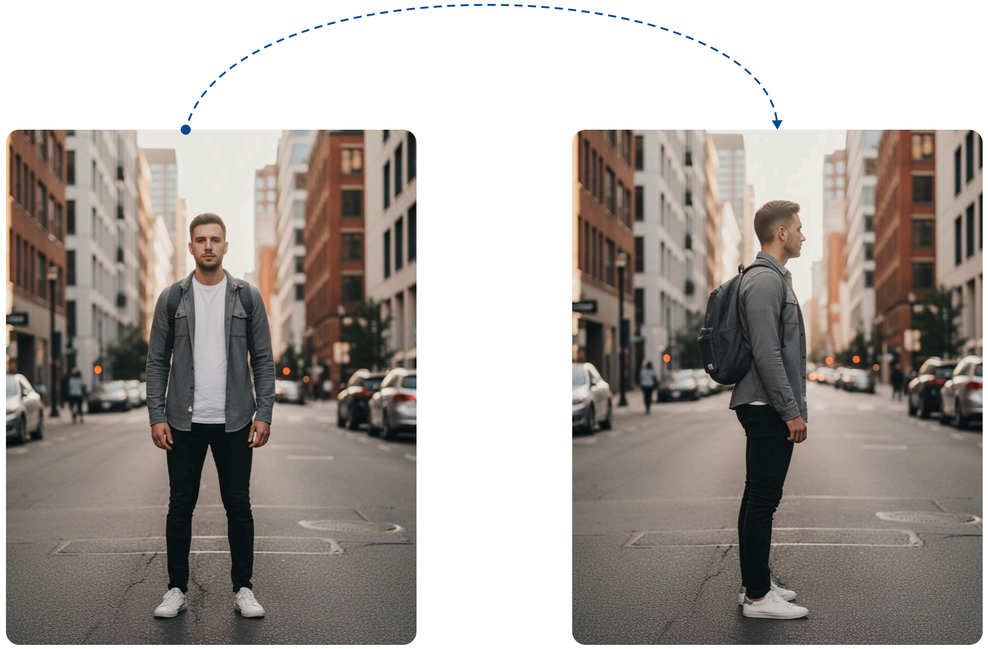
Background removal is common today, but in the future it could evolve to generating full 3D models directly from simple 2D images.



Future Scope

Reimagining Reality: A Next Frontier of Creative Control

Camera Angle Shifting



Lighting Re-Construction

