





AdCraft AI - Automating Ad Creation: Final Presentation

By: Shehbaz Virk, Kenneth Lam, Jun Park, Ethan Wang





Problem




70% of marketers are unhappy with their current marketing automation software **39% have no access to external data**

— [*Oracle Market Research*](#)



Manually **creating multiple ads for various locations** based on different weather conditions is **time-consuming and resource-intensive**

— [*Hunch Case Studies*](#)



37% of brand-side marketers report 'Data and Analytics' as the biggest skills gap faced by the marketing industry in 2024.

— [*Symmetrical Market Research*](#)




Domain Expert Insights



Parker is an experienced digital marketer specializing in creating and scaling online stores.




Focuses on creating online stores, running ads, and generating creatives (ad images/videos).



Workflow includes product discovery, ad testing, and scaling successful campaigns based on incoming data.

Domain Expert Insights



Manual and repetitive image editing tasks



Expensive to fine tune and personalize ads




Difficult to use AI models with frequent bugs




Return on Ad Spend fluctuations (ROAS)

Solution & Target Customer



Automate ad creation by using AI to generate images and captions




Online Businesses & Digital Marketing Teams



Impact



Turn aggregated customer reviews into actionable advertisement captions



Quickly develop high quality and creative advertisements based on starting images of your product

Initial Demo

Menu

- Home
- Instructions
- FAQ
- Settings and Privacy

Stable Diffusion Image Generator

Themes

Text Prompt:

Enter text prompt here



Upload Base Image (Optional):

Choose File No file chosen

Input for Caption Generation:

Input for caption generation

Generate Image






Image Generation AI

Caption Generation RAG



Image Generation Pipeline





User Image

**Object
Detection**

**Image
Masking**

**Prompt
Generation**

Stable Diffusion



**“Blender sitting
on a counter
with a white tile
splashboard
from a clean
kitchen”**



**Users Insert a
Starting Image
containing their
product**


**Classify the Focal
Subject through
Ultralytics’ YOLOv11
Segmentation Model**

**Generate a gaussian
feathered mask
around the detected
subjects from the
YOLO identified list**

**Use FoocusAI’s GPT2
prompt expansion
model with thematic
keyword guidance to
generate great
prompts**

**We pair the input
image, YOLOv11 mask,
and generated
prompts into a
SD-2-inpaint model
to create unique ads**

Image Generation




User Image

Object
Detection

Image
Masking

Prompt
Generation

Stable Diffusion



“Blender sitting
on a counter
with a white tile
splashboard
from a clean
kitchen”



Users Insert a
Starting Image
containing their
product


Classify the Focal
Subject through
Ultralytics' YOLOv11
Segmentation Model

Generate a gaussian
feathered mask
around the detected
subjects from the
YOLO identified list

Use FoocusAI's GPT2
prompt expansion
model with thematic
keyword guidance to
generate great
prompts

We pair the input
image, YOLOv11 mask,
and generated
prompts into a
SD-2-inpaint model
to create unique ads

Image Generation



User Image

**Object
Detection**

**Image
Masking**

**Prompt
Generation**

Stable Diffusion



**"Blender sitting
on a counter
with a white tile
splashboard
from a clean
kitchen"**



**Users Insert a
Starting Image
containing their
product**


**Classify the Focal
Subject through
Ultralytics' YOLOv11
Segmentation Model**

**Generate a gaussian
feathered mask
around the detected
subjects from the
YOLO identified list**

**Use FoocusAI's GPT2
prompt expansion
model with thematic
keyword guidance to
generate great
prompts**

**We pair the input
image, YOLOv11 mask,
and generated
prompts into a
SD-2-inpaint model
to create unique ads**

Image Generation



User Image



Object
Detection




Image
Masking



Prompt
Generation



Stable Diffusion



**"Blender sitting
on a counter
with a white tile
splashboard
from a clean
kitchen"**



Users Insert a
Starting Image
containing their
product



Classify the Focal
Subject through
Ultralytics' YOLOv11
Segmentation Model

Generate a gaussian
feathered mask
around the detected
subjects from the
YOLO identified list


Use FoocusAI's GPT2
prompt expansion
model with thematic
keyword guidance to
generate great
prompts

We pair the input
image, YOLOv11 mask,
and generated
prompts into a
SD-2-inpaint model
to create unique ads

Image Generation



“Blender sitting on a counter with a white tile splashboard from a clean kitchen”



Users Insert a Starting Image containing their product

Classify the Focal Subject through Ultralytics’ YOLOv11 Segmentation Model



Generate a gaussian feathered mask around the detected subjects from the YOLO identified list

Use FoocusAI’s GPT2 prompt expansion model with thematic keyword guidance to generate great prompts

We pair the input image, YOLOv11 mask, and generated prompts into a SD-2-inpaint model to create unique ads

Image Generation





Compare to Real Ads

AdCraft Generation



Base Image



Image Generation



Challenges and Iteration




Canny Edge Gaussian Blur



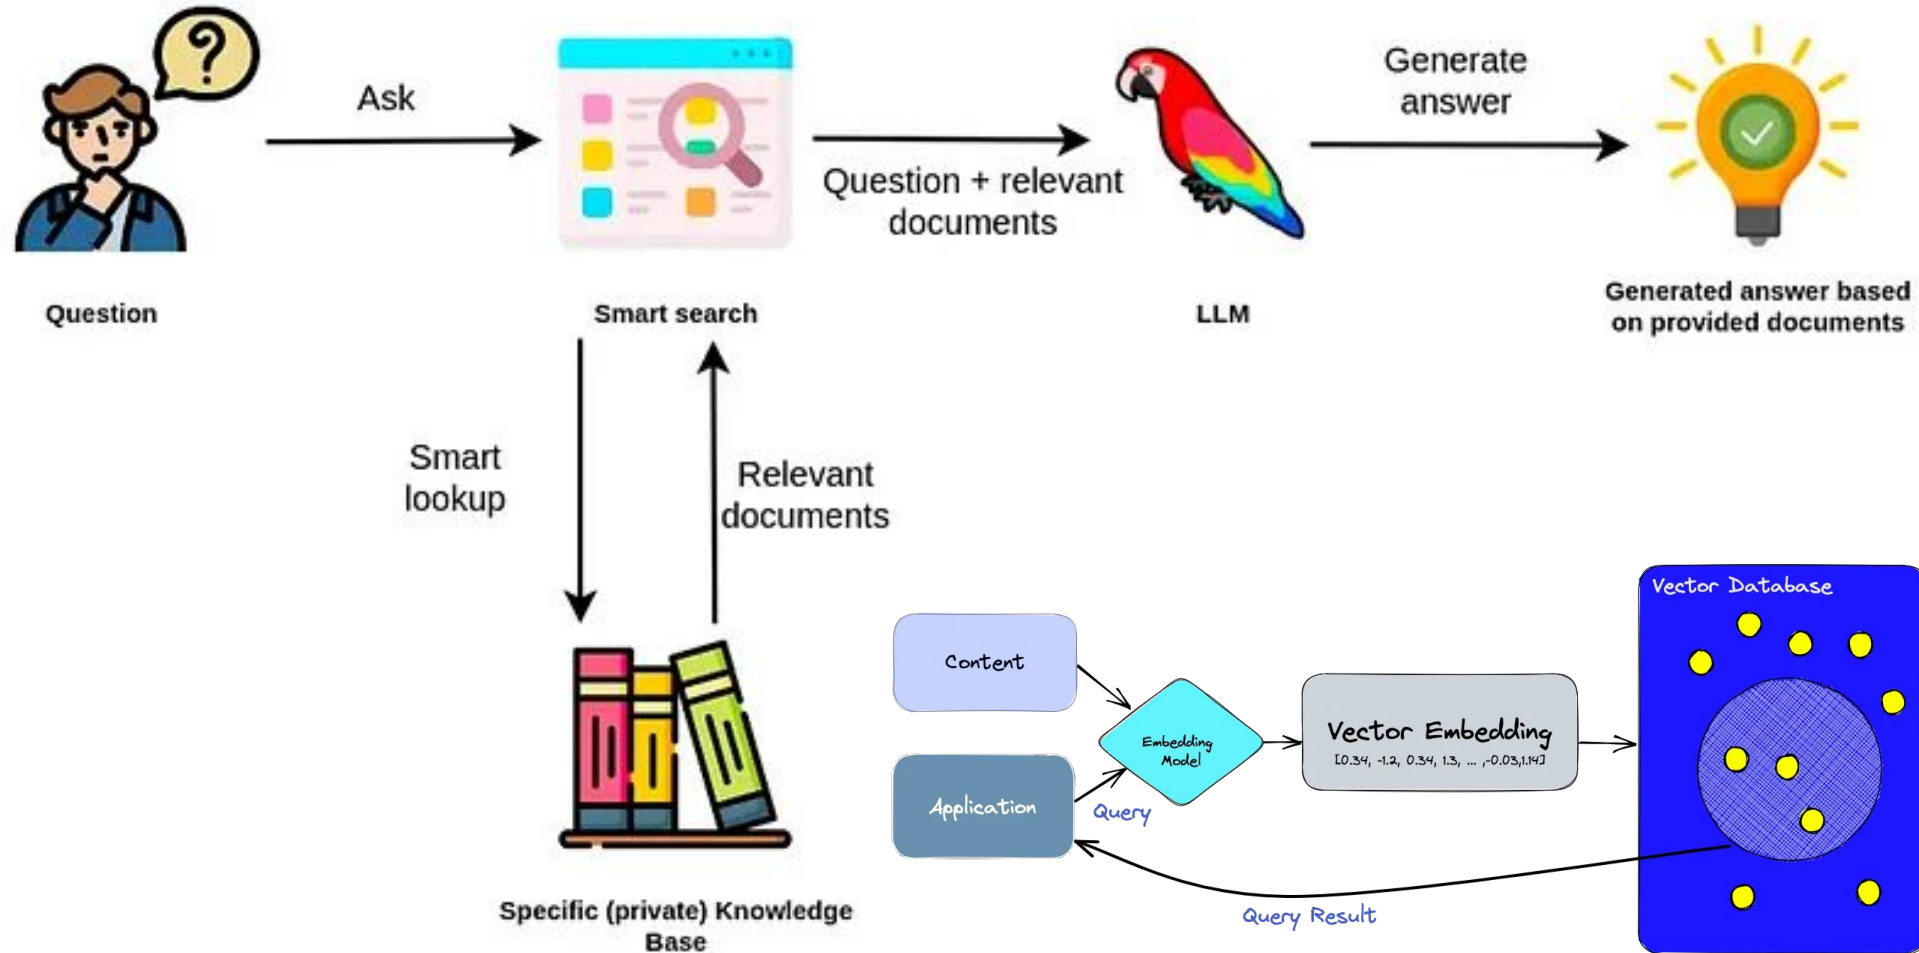
Inpaint Extrapolation

Image Generation


Caption Generation Pipeline



RAG Overview



Source: [What is RAG \(Retrieval-Augmented Generation\)?](#)



Product Choice

Subreddit: Car
Thread: BMW M5
Keyword: Weight

Reddit API



Search the product/idea parameters using Reddit API to get the 100 most recent posts

Vectorize Data

post_id	score	post_text
3	0.453663	the new bmw m5 reportedly weighs 5368 pounds and we're worried
1	0.418749	the 2025 bmw m5 weighs 5390 pounds which is 780 pounds more than a toyota sienna full sized van and 925 pounds more than a ford f150 supercrew pickup truck
87	0.336094	carwow bmw m5 v amg e63 v rs6 v panamera turbo dyno weight
2	0.218442	the new bmw m5 is heavier than an x5 slower than the old m5 to 60

Vectorize Reddit posts to allow for similarity searches and run keyword search to retrieve top 5 posts

Image Description



"A sleek sedan cruising through breathtaking mountain scenery under a vibrant sky"

BLIP generated image description to ensure caption is contextualized to the ad image


Caption Generation

"No matter the weight, no matter the destination"

Combine the results of LLM summarization of top posts and the BLIP generated image description to create a final caption

Caption Generation





Product Choice



Reddit API



Search the product/idea parameters using Reddit API to get the 100 most recent posts



Vectorize Data

post_id	score	post_text
3	0.453663	the new bmw m5 reportedly weighs 5368 pounds and we're worried
1	0.418749	the 2025 bmw m5 weighs 5390 pounds which is 780 pounds more than a toyota sienna full sized van and 925 pounds more than a ford f150 supercrew pickup truck
87	0.336094	carwow bmw m5 v amg e63 v rs6 v panamera turbo dyno weight
2	0.218442	the new bmw m5 is heavier than an x5 slower than the old m5 to 60

Vectorize Reddit posts to allow for similarity searches and run keyword search to retrieve top 5 posts

Image Description



"A sleek sedan cruising through breathtaking mountain scenery under a vibrant sky"

BLIP generated image description to ensure caption is contextualized to the ad image




Caption Generation

"No matter the weight, no matter the destination"

Combine the results of LLM summarization of top posts and the BLIP generated image description to create a final caption

Caption Generation





Product Choice




Reddit API



Search the product/idea parameters using Reddit API to get the 100 most recent posts



Vectorize Data



Subreddit: Car
Thread: BMW M5
Keyword: Weight

post_id	score	post_text
3	0.453663	the new bmw m5 reportedly weighs 5368 pounds and we're worried
1	0.418749	the 2025 bmw m5 weighs 5390 pounds which is 780 pounds more than a toyota sienna full sized van and 925 pounds more than a ford f150 supercrew pickup truck
87	0.336094	carwow bmw m5 v amg e63 v rs6 v panamera turbo dyno weight
2	0.218442	the new bmw m5 is heavier than an x5 slower than the old m5 to 60

Vectorize Reddit posts to allow for similarity searches and run keyword search to retrieve top 5 posts

Image Description



Caption Generation



"A sleek sedan cruising through breathtaking mountain scenery under a vibrant sky"


BLIP generated image description to ensure caption is contextualized to the ad image

"No matter the weight, no matter the destination"

Combine the results of LLM summarization of top posts and the BLIP generated image description to create a final caption

Caption Generation





Product Choice



Reddit API



Search the product/idea parameters using Reddit API to get the 100 most recent posts



Vectorize Data



Image Description



Caption Generation



Subreddit: Car
Thread: BMW M5
Keyword: Weight



"A sleek sedan cruising through breathtaking mountain scenery under a vibrant sky"

BLIP generated image description to ensure caption is contextualized to the ad image


"No matter the weight, no matter the destination"

Combine the results of LLM summarization of top posts and the BLIP generated image description to create a final caption

Caption Generation

post_id	score	post_text
3	0.453663	the new bmw m5 reportedly weighs 5368 pounds and we're worried
1	0.418749	the 2025 bmw m5 weighs 5390 pounds which is 780 pounds more than a toyota sienna full sized van and 925 pounds more than a ford f150 supercrew pickup truck
87	0.336094	carwow bmw m5 v amg e63 v rs6 v panamera turbo dyno weight
2	0.218442	the new bmw m5 is heavier than an x5 slower than the old m5 to 60

Vectorize Reddit posts to allow for similarity searches and run keyword search to retrieve top 5 posts



Product Choice



Reddit API



Search the product/idea parameters using Reddit API to get the 100 most recent posts



Vectorize Data




Image Description



Caption Generation

"No matter the weight, no matter the destination"

Combine the results of LLM summarization of top posts and the BLIP generated image description to create a final caption



"A sleek sedan cruising through breathtaking mountain scenery under a vibrant sky"

BLIP generated image description to ensure caption is contextualized to the ad image

Caption Generation

Subreddit: Car
Thread: BMW M5
Keyword: Weight

post_id	score	post_text
3	0.453663	the new bmw m5 reportedly weighs 5368 pounds and we're worried
1	0.418749	the 2025 bmw m5 weighs 5390 pounds which is 780 pounds more than a toyota sienna full sized van and 925 pounds more than a ford f150 supercrew pickup truck
87	0.336094	carwow bmw m5 v amg e63 v rs6 v panamera turbo dyno weight
2	0.218442	the new bmw m5 is heavier than an x5 slower than the old m5 to 60

Vectorize Reddit posts to allow for similarity searches and run keyword search to retrieve top 5 posts

Caption Creation

Reddit Data Summarization

"BMW's new m5 has been unveiled and it's been met with a fair amount of criticism because of its weight"

Image Summarization (BLIP)

"A silver bmw sedan driving down a road with mountains in the background"

Finalized Caption

"No matter the weight, no matter the destination" (with BLIP)



vs.

"Power that outweighs the weight—BMW M5: Redefining performance, not limits" (w/o BLIP)

Caption Generation

Reddit Summary Evaluation

Model Type	Precision	Recall	F1 Score
BART	0.8325	0.8064	0.8586
T5	0.8211	0.7840	0.8514
Pegasus	0.8912	0.9046	0.8978
LLaMA 3.1	0.8695	0.9191	0.8936



BLIP Summary Tuning



a silver bmw sedan driving down a road with mountains in the background

a silver bmw car driving down a road with mountains in the background

a silver bmw sedan driving down a road in front of a mountain

a silver bmw sedan driving down a road in front of mountains

a silver bmw sedan driving down a road in front of a mountain range



Caption Tuning

Before Tuning	Llama Model Parameters
Isn't the M5 a bit...heavy?	Prompt: Generate a short creative marketing caption with these two inputs: (Reddit), (BLIP) Max_Tokens = 500
Unleash the Beast, Not the Extra Pounds. Our new (end of caption)	Prompt: Generate a creative marketing caption with these two inputs: (Reddit), (BLIP) Max_Tokens = 50
Unleash the Beast, Unleash Yourself. We know the critics said it, 'The new BMW M5 is too heavy.' But we say, 'The more you weigh, the more you can carry.' Meet the beast that's got the heart of a lion, the spirit of a wanderer, and the soul of a thrill-seeker. This isn't just a car, it's a key to unlocking the roads less traveled.	Prompt: Generate a creative marketing caption with these two inputs: (Reddit), (BLIP) Max_Tokens = 200

Caption Tuning

After Tuning	Llama Model Parameters
Unleash the beast: where power meets precision, and every curve is a mountain to conquer	Prompt: Generate <i>a very short and creative marketing caption, around 20 words</i> , using these two inputs: (Reddit), (BLIP) Max_Tokens = 1000
Experience the thrill of the open road with our powerful BMW M5 - built for the bold, not for the lightweight.	(Same as above)
Unleash power, not compromise. The new BMW M5, engineered for handling any type of road.	Prompt: Generate a very short and creative marketing caption, around 20 words, <i>to highlight the product positives</i> , using these two inputs: (Reddit), (BLIP) Max_Tokens = 1000

Demo on More Examples

Menu

- Home
- Instructions
- FAQ
- Settings and Privacy

Stable Diffusion Image Generator

Themes

Selected Theme: **Semi-realistic**

Text Prompt:

photorealistic bmw car driving in the mountains



2

Upload Base Image (Optional):

Choose File

Input for Caption Generation:


caption for increased weight does not affect performance

Generate Image

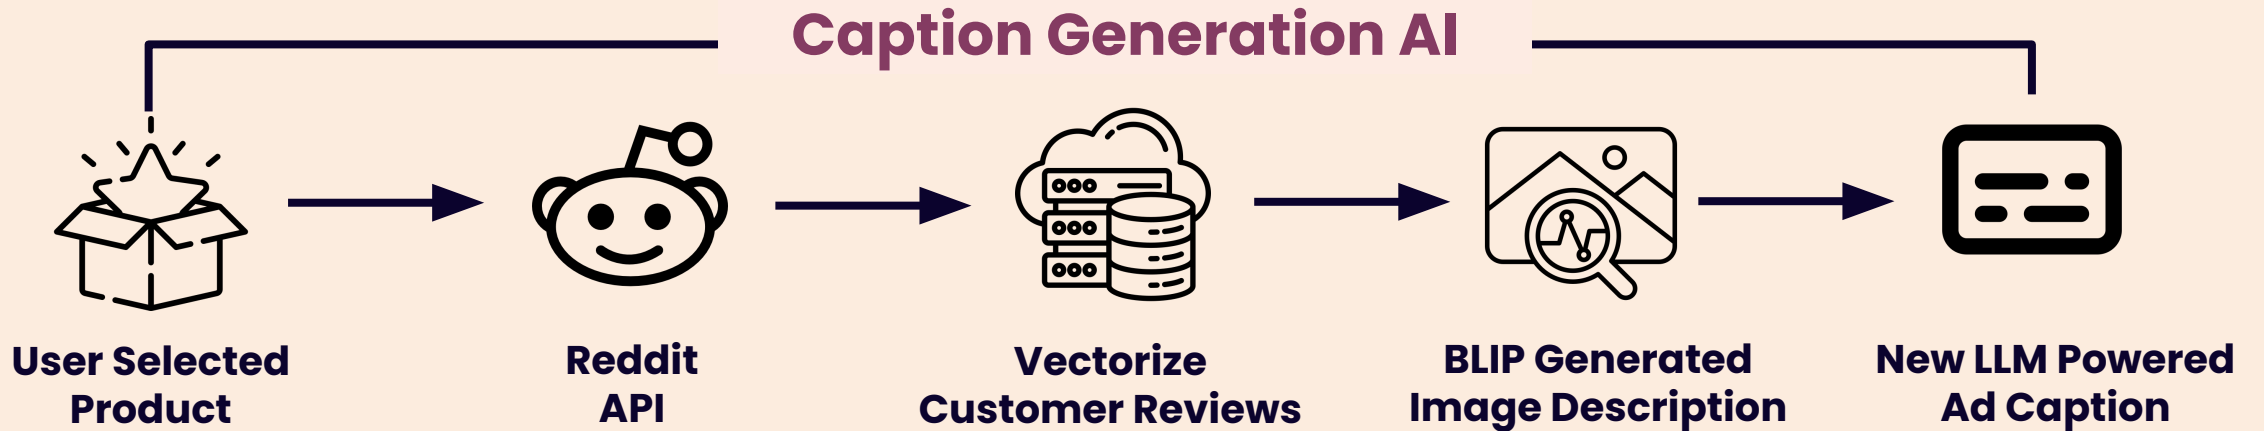
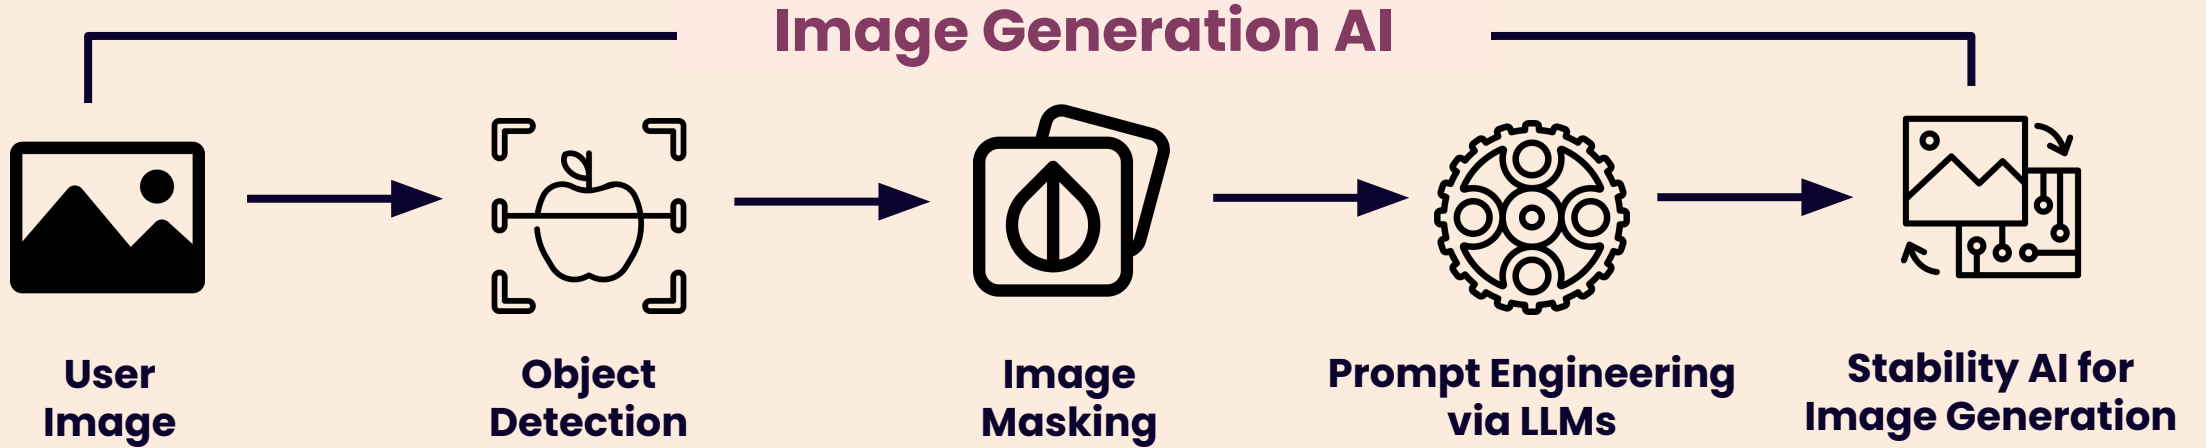
Generated Image:



Conclusion



Technical Summary



Ethics & Privacy Concerns



Using personal data such as Reddit reviews to generate advertisement captions



Training models on user uploaded or business specific data

Wrap-Up

Key Problems

- Difficult to properly use advanced AI models
- No External Data/Personalization
- Manual and Time-Consuming Content Creation

Our Solutions


- Easily leverage cutting edge AI models
- Reddit data as context for our RAG model
- GenAI – Stability AI powered ad generation

Potential Next Steps



Host the full platform so it's publicly accessible

- Currently Demo is up on our website. Web App, Image Generation, and Caption Generation model code accessible on GitHub




Create a multimodal RAG model that generates both images and captions together by training a model on specific data

- Currently separated into individual Image and Caption Generation Models

Q & A

Thanks!





Appendix

- <https://www.oracle.com/cx/marketing/automation/what-is-marketing-automation/statistics/>
- <https://www.semetrical.com/digital-marketing-skills-gap/>
- <https://github.com/Illyasviel/Fooocus> (Fooocus AI)